

572.4

Commercial Tools, Wireless, and Full-Packet Hunting

SANS

572.4

Commercial Tools, Wireless, and Full-Packet Hunting

SANS

Copyright © 2017, The SANS Institute. All rights reserved. The entire contents of this publication are the property of the SANS Institute.

PLEASE READ THE TERMS AND CONDITIONS OF THIS COURSEWARE LICENSE AGREEMENT ("CLA") CAREFULLY BEFORE USING ANY OF THE COURSEWARE ASSOCIATED WITH THE SANS COURSE. THIS IS A LEGAL AND ENFORCEABLE CONTRACT BETWEEN YOU (THE "USER") AND THE SANS INSTITUTE FOR THE COURSEWARE. YOU AGREE THAT THIS AGREEMENT IS ENFORCEABLE LIKE ANY WRITTEN NEGOTIATED AGREEMENT SIGNED BY YOU.

With the CLA, the SANS Institute hereby grants User a personal, non-exclusive license to use the Courseware subject to the terms of this agreement. Courseware includes all printed materials, including course books and lab workbooks, as well as any digital or other media, virtual machines, and/or data sets distributed by the SANS Institute to the User for use in the SANS class associated with the Courseware. User agrees that the CLA is the complete and exclusive statement of agreement between The SANS Institute and you and that this CLA supersedes any oral or written proposal, agreement or other communication relating to the subject matter of this CLA.

BY ACCEPTING THIS COURSEWARE YOU AGREE TO BE BOUND BY THE TERMS OF THIS CLA. BY ACCEPTING THIS SOFTWARE, YOU AGREE THAT ANY BREACH OF THE TERMS OF THIS CLA MAY CAUSE IRREPARABLE HARM AND SIGNIFICANT INJURY TO THE SANS INSTITUTE, AND THAT THE SANS INSTITUTE MAY ENFORCE THESE PROVISIONS BY INJUNCTION (WITHOUT THE NECESSITY OF POSTING BOND), SPECIFIC PERFORMANCE, OR OTHER EQUITABLE RELIEF.

If you do not agree, you may return the Courseware to the SANS Institute for a full refund, if applicable.

User may not copy, reproduce, re-publish, distribute, display, modify or create derivative works based upon all or any portion of the Courseware, in any medium whether printed, electronic or otherwise, for any purpose, without the express prior written consent of the SANS Institute. Additionally, User may not sell, rent, lease, trade, or otherwise transfer the Courseware in any way, shape, or form without the express written consent of the SANS Institute.

If any provision of this CLA is declared unenforceable in any jurisdiction, then such provision shall be deemed to be severable from this CLA and shall not affect the remainder thereof. An amendment or addendum to this CLA may accompany this courseware.

SANS acknowledges that any and all software and/or tools, graphics, images, tables, charts or graphs presented in this courseware are the sole property of their respective trademark/registered/copyright owners, including:

AirDrop, AirPort, AirPort Time Capsule, Apple, Apple Remote Desktop, Apple TV, App Nap, Back to My Mac, Boot Camp, Cocoa, FaceTime, FileVault, Finder, FireWire, FireWire logo, iCal, iChat, iLife, iMac, iMessage, iPad, iPad Air, iPad Mini, iPhone, iPhoto, iPod, iPod classic, iPod shuffle, iPod nano, iPod touch, iTunes, iTunes logo, iWork, Keychain, Keynote, Mac, Mac Logo, MacBook, MacBook Air, MacBook Pro, Macintosh, Mac OS, Mac Pro, Numbers, OS X, Pages, Passbook, Retina, Safari, Siri, Spaces, Spotlight, There's an app for that, Time Capsule, Time Machine, Touch ID, Xcode, Xserve, App Store, and iCloud are registered trademarks of Apple Inc.

Governing Law: This Agreement shall be governed by the laws of the State of Maryland, USA.



Commercial Tools, Wireless, and Full- Packet Hunting

© 2017 Lewes Technology Consulting, LLC and Mat Oldham | All Rights Reserved | Version # FOR572_C01_01

Authors:

Phil Hagen, Lewes Technology Consulting, LLC

phil@lewestech.com | @philhagen

Mat Oldham

mat.oldham@gmail.com | @roujisecurity

<http://twitter.com/sansforensics>

SANS DFIR

DIGITAL FORENSICS & INCIDENT RESPONSE


FOR408
Windows Forensics

FOR518
Mac Forensics

FOR526
Memory Forensics In-Depth

FOR585
Advanced Smartphone Forensics

OPERATING SYSTEM & DEVICE IN-DEPTH



INCIDENT RESPONSE & ADVERSARY HUNTING

FOR508
Advanced Incident Response


FOR572
Advanced Network Forensics and Analysis


FOR578
Cyber Threat Intelligence


FOR610
REM: Malware Analysis


SEC504
Hacker Tools, Techniques, Exploits, and Incident Handling


MGT535
Incident Response Team Management


[@sansforensics](https://twitter.com/sansforensics)


[sansforensics](https://www.facebook.com/sansforensics)


[dfir.to/DFIRLinkedInCommunity](https://www.linkedin.com/company/dfir/)


[dfir.to/gplus-sansforensics](https://plus.google.com/dfir.to/gplus-sansforensics)


dfir.to/MAIL-LIST

This page intentionally left blank.

SANS DFIR

DIGITAL FORENSICS & INCIDENT RESPONSE



FOR408
Windows Forensics



FOR518
Mac Forensics



FOR526
Memory Forensics
In-Depth



FOR585
Advanced Smartphone
Forensics

OPERATING
SYSTEM &
DEVICE
IN-DEPTH

INCIDENT
RESPONSE &
ADVERSARY
HUNTING



FOR508
Advanced Incident Response



FOR572
Advanced Network Forensics
and Analysis



FOR578
Cyber Threat Intelligence



FOR610
REM: Malware Analysis



SEC504
Hacker Tools, Techniques,
Exploits, and Incident Handling



MGT535
Incident Response
Team Management



@sansforensics



sansforensics



dfir.to/DFIRLinkedInCommunity

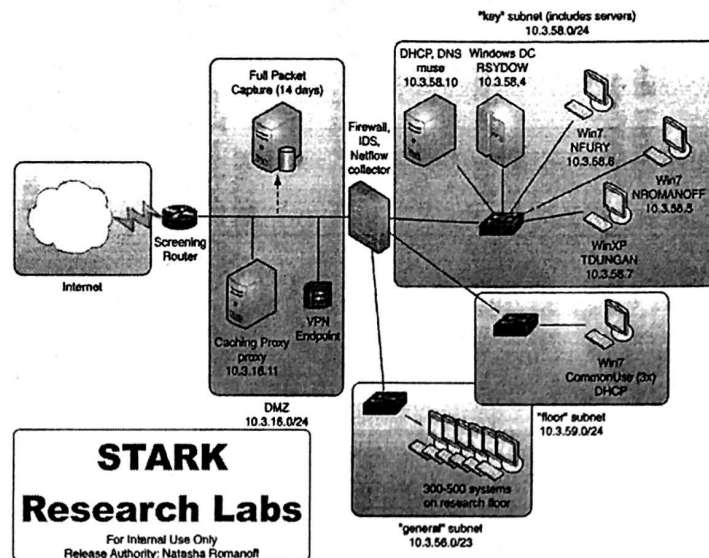


dfir.to/gplus-sansforensics



dfir.to/MAIL-LIST

Victim Network




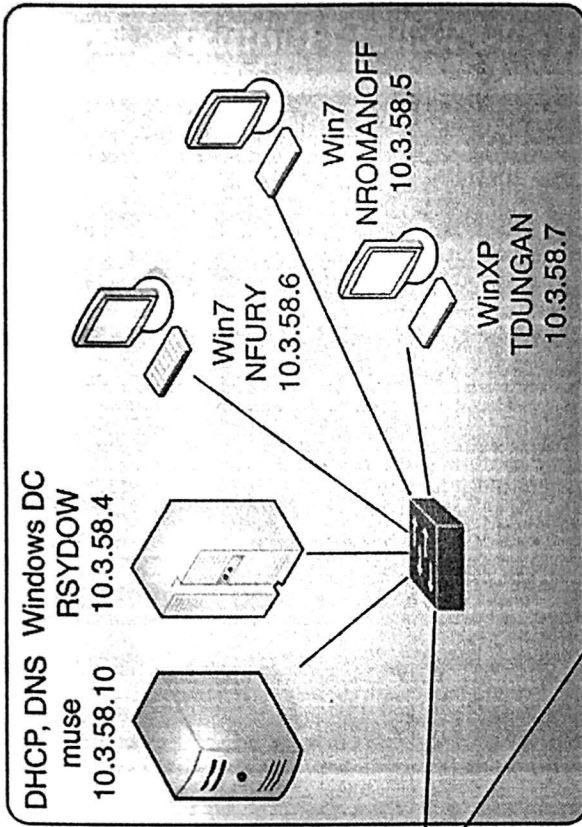
STARK
Research Labs
 For Internal Use Only
 Release Authority: Natasha Romanoff

The SRL network is representative of many medium-sized enterprises, with several major internal segments separated by a single firewall, but very little network protection between hosts once on the inside of that perimeter.


Key systems we are aware of (not an exhaustive list):

Host	IP Address	Role
Router	10.3.16.1	Inside interface of Internet screening router
Firewall	10.3.56.1/23 10.3.58.1/24 10.3.59.1/24 10.3.16.99/24	General administration and users Key users & SHIELDBASE servers subnet Production floor systems External (DMZ) interface
VPN	10.3.16.5/24	VPN Concentrator external interface
Proxy	10.3.16.11/24	Caching proxy for outgoing HTTP
Sniffer	10.3.56.254/24	Tcpdump rotating buffer. Approx 14 days of storage.
Webserver	10.3.16.3/24	External presence of SRL
rsydow	10.3.58.4	SHIELDBASE Domain Controller – Windows 2008r2
tdungan	10.3.58.7	Tim Dungan's R&D workstation
nromanoff	10.3.58.5	Natasha Romanoff's workstation
nfury	10.3.58.6	Nick Fury's workstation


"key" subnet (includes servers)
10.3.58.0/24




"floor" subnet
10.3.59.0/24



"general" subnet
10.3.56.0/23



DMZ
10.3.16.0/24



STARK

Research Labs

For Internal Use Only
Release Authority: Natasha Romanoff

Commercial Tools, Wireless, and Full-Packet Hunting

Simple Mail Transfer Protocol (SMTP)

Commercial Network Forensic Tools

Lab 09: NetworkMiner

Wireless Network Forensics

Automated Tools and Libraries

Lab 10: Using Command-Line Tools for Analysis

Full-Packet Hunting with Moloch

Lab 11: Network Forensic Analysis Using Moloch

This page intentionally left blank.

Simple Mail Transfer Protocol (SMTP)

This page intentionally left blank.

SMTP: An Old Standby

- First defined in 1982 by RFC 821
 - Revised and updated many times since
 - Still does most heavy lifting in mail exchanges today
- The “sending” and “relay” protocol
 - Most SMTP “servers” also act as clients
- Messages traverse multiple hops between sender and recipient
 - Most of these hops use SMTP

E-mail is one of the oldest and most lasting, ubiquitous protocols on the Internet. What started in the early days of ARPAnet remains a fundamental protocol used to send messages between humans and systems around the world. Simple Mail Transfer Protocol, or SMTP, is the “sending” and “relay” protocol that ensures mail messages arrive at their intended destinations.

Just as a piece of paper mail or a package crosses multiple mail stations, delivery trucks, or other physical points on its destination, e-mail traverses a series of hops between the originator clicking “Send” and the recipient seeing it arrive in their Inbox. There are a number of roles that different servers and server processes fill during this end-to-end process. As we’ll see next, SMTP takes care of most of them. Even in a webmail model, the HTTP-based webmail application running on the server conducts SMTP transactions to put an e-mail message into play.

In this “relay” role, it’s important to recognize that an SMTP server then becomes an SMTP client when passing the e-mail message on to the next system.

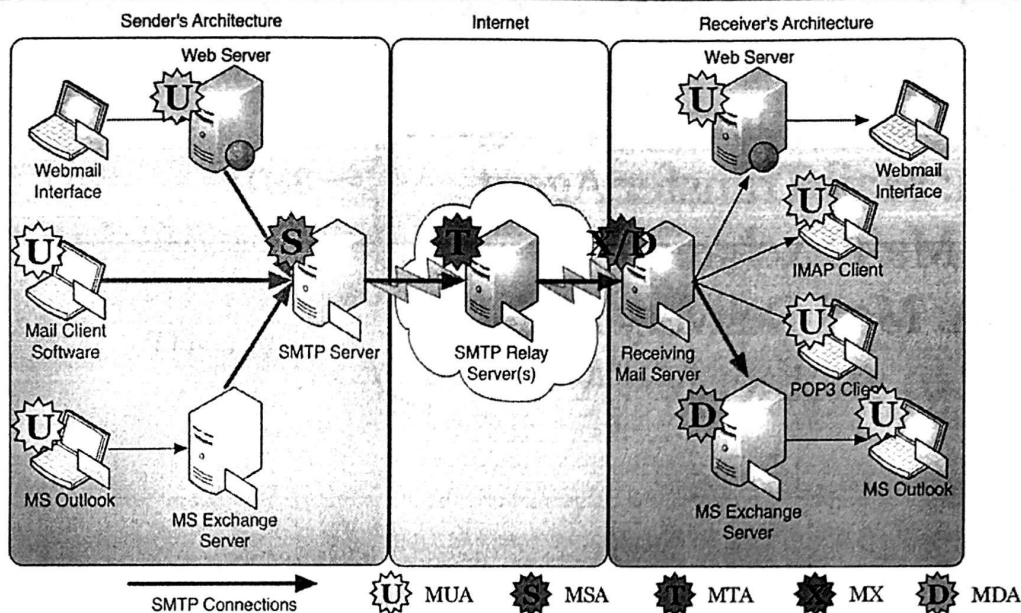
M-Acronym Soup

- MUA: Mail User-Agent
 - MSA: Mail Submission Agent
 - MTA: Mail Transfer Agent
 - MX: Mail Exchanger
 - MDA: Mail Delivery Agent
 - Back to another MUA
- 
- Often SMTP
- SMTP
- Non-SMTP

Each step in the overall end-to-end SMTP process is performed by a specific piece of client or server software, identified by the following acronyms:

- **Mail User-Agent (MUA):** The client software an individual uses to send mail. For example, Outlook, macOS's Mail.app, etc.
- **Mail Submission Agent (MSA):** The server software that receives a message after the user sends it in the MUA. The sending user's MUA connects directly to the MSA and submits the message into the overall SMTP process.
- **Mail Transfer Agent (MTA):** The server software that passes the message along to other servers (MTAs) in the series of hops between sender and receiver. There can be (and often are) multiple MTAs involved in the transfer of each mail message. MTAs use SMTP in each step of relaying the message along the delivery path. An MTA is both an SMTP client and server in these steps.
- **Mail Exchanger (MX):** The system identified as a responsible receiver for mail sent to a given hostname or domain. This system is most often designated in a DNS record with the "MX" record type. Each destination host or domain can have multiple MXes for redundancy and load balancing purposes. The MX serves as a "last-step" MTA in this role, using an SMTP connection.
- **Mail Delivery Agency (MDA):** The server software that provides mail messages to a user after successful authentication. This is not an SMTP step, but rather uses other protocols such as Post Office Protocol (POP3), Internet Message Access Protocol (IMAP), or Messaging Application Programming Interface (MAPI). The receiving user's own MUA makes a connection to the MDA for retrieval of mail data.

E-Mail Message Flow



SANS DFIR

FOR572.4 | Advanced Network Forensics and Analysis 10

Here you can see examples of several common e-mail paths. Although there are too many possible configurations to detail here, these cover those you'd expect to see in most environments.

Note that although some clients (MUAs) such as the webmail and Outlook systems don't natively use SMTP to send e-mail messages, their immediate upstream systems (MSA) pass sent messages via SMTP-speaking paths very soon into the message's lifecycle. As soon as the message is routed via the Internet, SMTP is the name of the game. Usually one or more relays (MTAs) will receive a message, and then pass it along based on the preconfigured routing rules. Finally the designated recipient server (MX) receives the message and places it into the receiving user's mailbox (MDA). At the other end of the process, the user's MUA receives and displays the message. (Unless they ignore or delete it...)

What's interesting to note is that at each stage of the transmission, servers generally add a message header indicating the server's hostname, date, and time the message was processed, and some other details that can be valuable in determining the path a message took on its journey from sender to recipient. Of course, these headers are only "required" by RFC and not law—so some servers may alter the existing or expected header values.

An example of a message a Gmail user sent to the SANS DFIR mailing list is on the next page. The headers are typically in reverse chronological order. Therefore, the last header indicates that the originator used the web interface to send the message, with each subsequent node adding its identity as well as the previous node. This establishes a clear path that the message took, including both hostnames and corresponding IP addresses (public and private).

From: Sending User <sendinguser@gmail.com>
To: dfir@lists.sans.org

Received: from lists21a.den.sans.org (lists.sans.org [66.35.59.40]) by mail.example.com (8.13.8/8.13.8) with ESMTP id r79HxUbk021274 for <listssubs@receiver.com>; Fri, 9 Aug 2013 13:26:21 -0400

Received: from lists21a.den.sans.org (localhost.localdomain [127.0.0.1]) by lists21a.den.sans.org (Postfix) with ESMTP id 64A29102E1 for <listssubs@receiver.com>; Fri, 9 Aug 2013 17:25:30 +0000 (UTC)

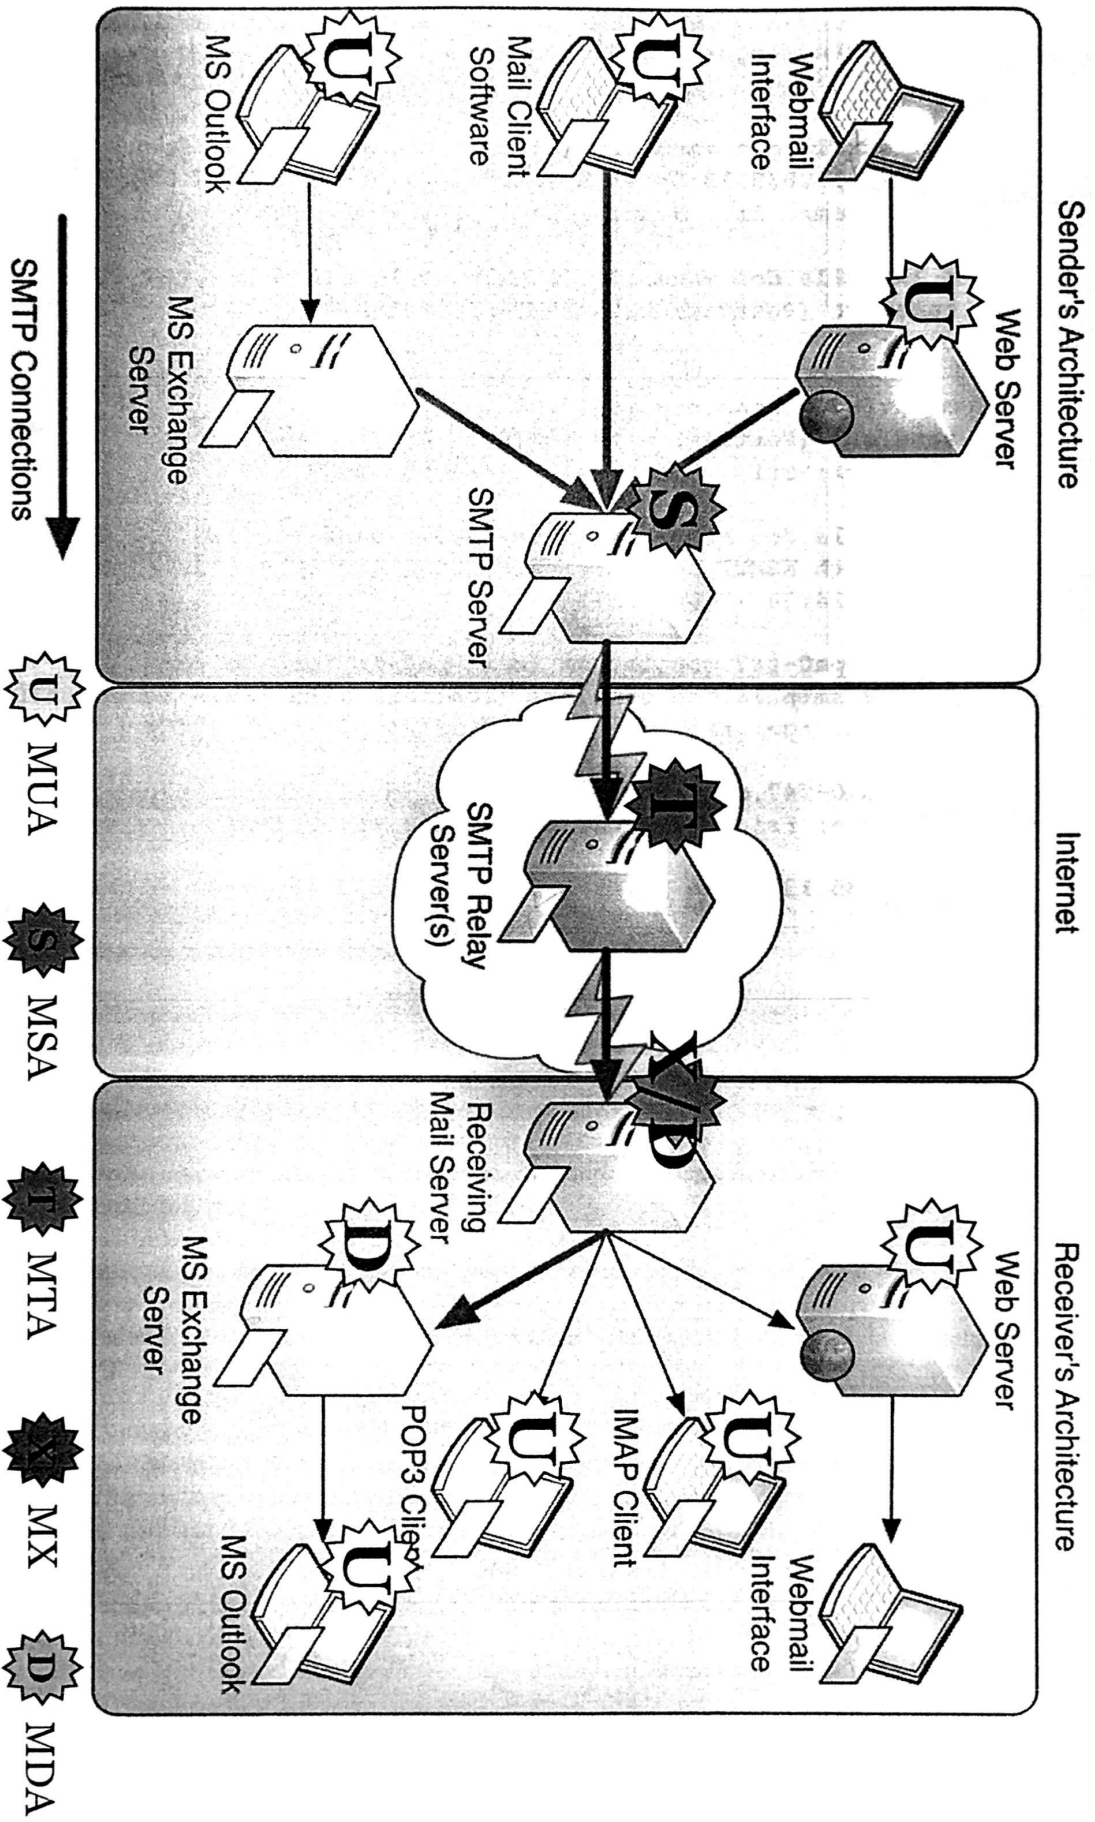
Received: from savfw21a.sans.org (savfw21a.den.sans.org [10.2.2.25]) by lists21a.den.sans.org (Postfix) with ESMTP id 4117A401ED for <dfir@lists.sans.org>; Fri, 9 Aug 2013 17:24:56 +0000 (UTC)

Received: from smtp21a.den.sans.org (smtp21a.den.sans.org [10.2.2.12]) by savfw21a.sans.org with ESMTP id 1dB14gc5lLQxZCb5 for <dfir@lists.sans.org>; Fri, 09 Aug 2013 17:24:54 +0000 (GMT)

Received: from mail-pa0-f47.google.com (mail-pa0-f47.google.com [209.85.220.47]) by smtp21a.den.sans.org (Postfix) with ESMTP id 9C86410188 for <dfir@lists.sans.org>; Fri, 9 Aug 2013 17:24:53 +0000 (UTC)

Received: by mail-pa0-f47.google.com with SMTP id k113sh5142167pab.20 for <dfir@lists.sans.org>; Fri, 09 Aug 2013 10:24:53 -0700 (PDT)

Received: by 10.70.35.131 with HTTP; Fri, 9 Aug 2013 10:24:53 -0700 (PDT)



SMTP Transmission Characteristics

- Uses TCP/25, TCP/587 (legacy SSL: TCP/465)
- Created in simpler times
 - English ASCII was sufficient, SPAM didn't exist, and trust ruled the Internet
 - Protocol extensions and updates address these and other developments

Primarily, SMTP is associated with TCP port 25 but has more recently expanded to include port 587 as well. Typically, SMTP usage over port 587 requires authentication, which we'll discuss shortly. For a while, TCP/465 was also used for SMTP wrapped in an SSL connection. However, this latter use has been depreciated in favor of the "STARTTLS" option on ports 25 or 587. STARTTLS also provides encryption, as we'll see.

It's important to recognize that many legacy protocols—including SMTP—were created when the Internet was a very different entity than it is now. Today's globally ubiquitous and often hostile environment demands a very different mindset for protocol design. For example, the original SMTP specification had no considerations for authentication/authorization, non-English text, SPAM abatement, or even attachments. These features, which we consider integral to the modern e-mail landscape, were all established after SMTP was already in widespread use. We'll look at a few of the protocol extensions that brought such features into being and how they affect our analysis of the protocol.

Basic SMTP Transaction

```

S: 220 mail.identityvector.com ESMTP Sendmail 8.13.8/8.13.8: Tue, 4 Nov 2014 11:38:44 -0500
C: EHLO metrowg.pronaceouses.com
S: 250-mail.identityvector.com Hello metrowg.pronaceouses.com [138.128.6.57],
  pleased to meet you
S: 250-ENHANCEDSTATUSCODES
S: 250-PIPELINING
S: 250-8BITMIME
S: 250-SIZE
S: 250-DSN
S: 250-ETRN
S: 250-STARTTLS
S: 250-DELIVERBY
S: 250 HELP
C: MAIL FROM:<BraylonHuff@metrowg.pronaceouses.com>
S: 250 2.1.0 <BraylonHuff@metrowg.pronaceouses.com>... Sender ok
C: RCPT TO:<recip@identityvector.com>
S: 250 2.1.5 <recip@identityvector.com>... Recipient ok
C: DATA
S: 354 Enter mail, end with "." on a line by itself
C: Message-ID: <92861425.aLg7b0EEmFjsWCZ20141104083084444@ml1.metrowg.pronaceouses.com>
C: Date: Tue, 04 Nov 2014 08:38:44 -0800
C: From: Laxy Weight Loos <Braylon@pronaceouses.com>
C: To: <recip@identityvector.com>
C: Reply-to: <Braylon@pronaceouses.com>
C: Subject: some believable spam subject line
C:
C: _____ Header/body separator
C: This is some content that is supposed to look legit...
C: Completely Legit.
C: Click the link, because you do what you're supposed to do.
C: Just like a good little human does.
C: Obey your email master.
C: http://hfdk1ashgfjkd1shgfdlsjbg1ksd.cs.oc/pwnme/please
C: .
S: 250 2.0.0 sM4Gci0h008308 Message accepted for delivery
C: QUIT
S: 221 2.0.0 mail.identityvector.com closing connection
  
```

Negotiation & Envelope

SMTP Headers

Header/body separator

SMTP Body

Body

terminator

Teardown

This is perhaps the most rudimentary example of an SMTP submission, as seen from a network vantage point. In this case, the “metrowg.pronaceouses.com” server is relaying a message to the “mail.identityvector.com” server. For the purposes of this illustration, we’ll state that the “mail.identityvector.com” server has been designated as an MX for the “identityvector.com” domain, so this transaction represents a transaction between an MTA and an MX.

First, the exchange contains a series of greetings between the client and server that establish the mutually supported protocols and protocol extensions. This starts with an “EHLO”, or “extended hello” command, but some exchanges may include an older “HELO” instead.

After the SMTP negotiation, the sender sends SMTP header data. Long headers are wrapped and indented on subsequent lines. The server knows to treat indented lines as continuations of their immediate predecessor.

The separator between headers and body content is a “double-CRLF”, consisting of a carriage return and newline followed by another carriage return and newline. In hex, these bytes are 0x0D0A0D0A. This sequence is a common header/body separator in other ASCII protocols such as HTTP as well.

The message body here contains HTML content, and the client indicates it is complete with the message by sending a single period character on a line by itself.

After the body—and therefore the entire message—is complete, the server provides the queued message identifier (which is generally available in the mail server’s logs) and the client initiates a clean teardown procedure.

```

S: 220 mail.identityvector.com ESMTP Sendmail 8.13.8/8.13.8; Tue, 4 Nov 2014 11:38:44 -0500
C: EHLO metrowg.pronaceouses.com
S: 250-mail.identityvector.com Hello metrowg.pronaceouses.com [138.128.6.57],
  pleased to meet you
S: 250-ENHANCEDSTATUSCODES
S: 250-PIPELINING
S: 250-8BITMIME
S: 250-SIZE
S: 250-DSN
S: 250-ETRN
S: 250-STARTTLS
S: 250-DELIVERBY
S: 250 HELP
C: MAIL FROM:<BraylonHuff@metrowg.pronaceouses.com>
S: 250 2.1.0 <BraylonHuff@metrowg.pronaceouses.com>... Sender ok
C: RCPT TO:<recip@identityvector.com>
S: 250 2.1.5 <recip@identityvector.com>... Recipient ok
C: DATA
S: 354 Enter mail, end with "." on a line by itself
C: Message-ID: <92861425.aLg7b0EBmFjmWCZ2014110408308444@mx1.metrowg.pronaceouses.com>
C: Date: Tue, 04 Nov 2014 08:38:44 -0800
C: From: Lazy Weight Loss <Braylon@pronaceouses.com>
C: To: <recip@identityvector.com>
C: Reply-to: <Braylon@pronaceouses.com>
C: Subject: some believable spam subject line
C:
C: This is some content that is supposed to look legit...
C: Completely legit.
C: Click the link, because you do what you're supposed to do.
C: Just like a good little human does.
C: Obey your email master.
C: http://hfdkiasghfjkdhsingfdhsjngtksd.cz.cc/pwmmw/pleaseBody
C:
S: 250 2.0.0 sA4Gcioh008308 Message accepted for delivery
C: QUIT
S: 221 2.0.0 mail.identityvector.com closing connection

```

Negotiation
& Envelope

SMTP
Headers

SMTP Body

Teardown

Header/body
separator

Body

terminator

The Results

```
Return-Path: <BraylonHuff@metrowy.pronaceouses.com>
Received: from metrowy.pronaceouses.com (metrowy.pronaceouses.com [138.128.6.57])
  by mail.identityvector.com (8.13.8/8.13.8) with ESMTP id sA4Gciob008308
  for <recip@identityvector.com>; Tue, 4 Nov 2014 11:38:45 -0500
Message-ID: <92861425.aLg7b0EBmFjmwKCE20141104083084444@mail.metrowy.pronaceouses.com>
Date: Tue, 04 Nov 2014 08:38:44 -0800
From: Lazy Weight Loss <Braylon@pronaceouses.com>
To: <recip@identityvector.com>
Reply-to: <Braylon@pronaceouses.com>
Subject: Too Self-Conscious about your body?
Content-Type: text/html; charset="us-ascii"
Content-Transfer-Encoding: 7bit
MIME-Version: 1.0
X-Spam-Checker-Version: SpamAssassin 3.3.1 (2010-03-16) on
  quaff.identityvector.com
X-Spam-Level:
X-Spam-Status: No, score=-1.8 required=5.0 tests=BAYES_00,HTML_MESSAGE,
  MIME_HTML_ONLY,RP_MATCHES_RCVD,SPF_HELO_PASS,SPF_PASS autolearn=no
  version=3.3.1
X-Virus-Scanned: clamav-milter 0.98.4 at mail.identityvector.com
X-Virus-Status: Clean
X-Graylist: Sender passed SPF test, not delayed by milter-greylist-4.4.3
  (mail.identityvector.com [205.186.148.46]);
  Tue, 04 Nov 2014 11:38:46 -0500 (EST)
```

New "Received:" header

Arbitrary "X-" headers
added by MX server

Header/body separator

```
This is some content that is supposed to look legit...
Completely Legit.
Click the link, because you do what you're supposed to do.
Just like a good little human does.
Obey your email master.
http://hfdklashgfjkdlshgfdlsjhgklsd.cx.cc/pwmma/please
```

The resulting message, as delivered to the recipient's Inbox (and possibly stored on disk, depending on the MDA software), is shown here. You can see that the headers included by the sending SMTP server remain intact during the MTA-to-MX transaction, but that the MX has added some headers of its own. Most important to note is that an additional "Received:" header is added at each hop in the transaction. It's also useful to recognize that any SMTP server in the path can also add, remove, or alter whatever SMTP headers or body data it's configured to mangle. Although this is uncommon, it does bear consideration when thinking from a forensic point of view.

Despite the few header changes from the MX's own delivery process, it's easy to see that the majority of the network transaction from the previous slide has remained intact.

Any MTA or MX along the way can optionally add whatever "X-" headers, indicating free-form extension data. Often this includes antivirus, anti-spam, and other features, but can truly be anything the server is configured to add. Most webmail providers now include a header containing the encoded or encrypted IP address of the message originator. These can aid in tracking down abuse and other malicious activity.

New "Received:" header

Return-Path: <BraylonHuff@metrowg.pronaceouses.com>
Received: from metrowg.pronaceouses.com (metrowg.pronaceouses.com [138.128.6.57])
by mail.identityvector.com (8.13.8/8.13.8) with ESMTP id sA4Gci0h008308
for <recip@identityvector.com>; Tue, 4 Nov 2014 11:38:45 -0500
Message-ID: <92861425.aLg7b0EBmFjmWCZ20141104083084444@mx1.metrowg.pronaceouses.com>
Date: Tue, 04 Nov 2014 08:38:44 -0800
From: Lazy Weight Loss <Braylon@pronaceouses.com>
To: <recip@identityvector.com>
Reply-to: <Braylon@pronaceouses.com>
Subject: Too Self-Conscious about your body?
Content-Type: text/html; charset="us-ascii"
Content-Transfer-Encoding: 7bit
MIME-Version: 1.0
X-Spam-Checker-Version: SpamAssassin 3.3.1 (2010-03-16) on
quaff.identityvector.com
X-Spam-Level:
X-Spam-Status: No, score=-1.8 required=5.0 tests=BAYES_00,HTML_MESSAGE,
MIME_HTML_ONLY,RP_MATCHES_RCVD,SPF_HELO_PASS,SPF_PASS,autolearn=no
version=3.3.1
X-Virus-Scanned: clamav-milter 0.98.4 at mail.identityvector.com
X-Virus-Status: Clean
X-Greylist: Sender passed SPF test, not delayed by milter-greylist-4.4.3
(mail.identityvector.com [205.186.148.46]);
Tue, 04 Nov 2014 11:38:46 -0500 (EST)

Arbitrary "X-" headers added by MX server

Header/body separator

This is some content that is supposed to look legit...
Completely legit.
Click the link, because you do what you're supposed to do.
Just like a good little human does.
Obey your email master.
<http://hfdkklashgfjkdishgfdlsjhglksd.cz.cc/pwnme/please>

Adapt and Overcome

- Evolution of the Internet required changes
 - Attachments → MIME/base64 encoding
 - SPAM → Authentication
 - Privacy data → Encryption
 - International Character Sets → Unicode

As we mentioned previously, SMTP has evolved significantly since its early incarnations. We'll take a look at each of these more significant modifications in turn.

MIME Parts and base64 Encoding

- ASCII-based protocol couldn't handle binary
 - Attachments, internationalized characters, etc.
 - Base64 handles binary, but not human-readable
- MIME uses separator line to establish “parts”
 - Look for “Boundary” in headers
 - Encoding specified for each MIME part
 - Allows carving individual attachments
 - Some parts contain additional metadata

Although it was great for (English) text, SMTP was never designed to handle that 175-MB PowerPoint file, endless JPEGs of kitty cats, or even non-Western languages. To bring these now-fundamental features into existence, base64 encoding was selected. This encoding method represents arbitrary binary bytes in an all-ASCII character set consisting of 64 possible characters. However, dropping massive blocks of base64 text into the body of an email was not an option that most humans would be able to deal with.

To address this, the Multipart Internet Mail Extension, or MIME, standard was adopted for use in SMTP. E-mail messages that include MIME data contain a “boundary” indicator in the headers. This string designates the start of each new MIME part in the overall mail message.

Each part can also include its own headers, which often include useful metadata such as the attachment's filename, encoding method, etc.

By tracking these boundaries and referencing each part's header metadata, we can easily discern each subsequent message part, making for rather simple attachment carving.

Content-Type: multipart/mixed; boundary=Apple-Mail-E39F65D3-710E-43DD-AEE1-62EFF51AB58

Content-Transfer-Encoding: 7bit

Mime-Version: 1.0 (1.0)

...

--Apple-Mail-E39F65D3-710E-43DD-AEE1-62EFF51AB58

Content-Type: text/plain; charset=us-ascii

Content-Transfer-Encoding: quoted-printable

Here is that photo from the fundraiser.=20

--Apple-Mail-E39F65D3-710E-43DD-AEE1-62EFF51AB58

Content-Type: image/jpeg; name=photo.JPG

Content-Disposition: inline; filename=photo.JPG

Content-Transfer-Encoding: base64

/9j/4Q/+RXhpZgAATU0AKGAAAgACWEPAAIAAAAGAAAkGEQAAIAAAAJAAAAmAESAAAMAAAAABAAYAAEAAUAAAAABAAAAogEbaAUAAAAABAAAAqgEoAAAMAAAAABAAIAAAEAAIAAAAEENI4xAAEYAAIAAAAU

.....

Hr1rrG0a00XS2im10a6lQ7mUFS0gzxx7A/zr+5uiuef0IpOV/wC1NP8Arz/91NI/sc0v+ai/8tH/

AL5P/9k=

--Apple-Mail-E39F65D3-710E-43DD-AEE1-62EFF51AB58

Content-Type: text/plain; charset=us-ascii

Content-Transfer-Encoding: 7bit

Thanks again for your help!

--Apple-Mail-E39F65D3-710E-43DD-AEE1-62EFF51AB58--

From SMTP Headers

(MIME Part Separators)

MIME Part 1

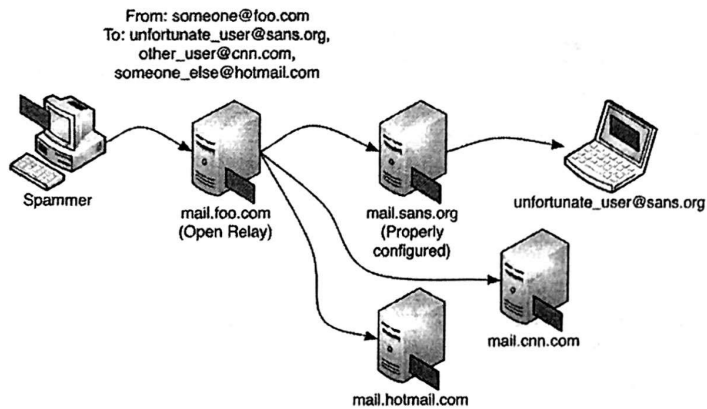
MIME Part 2

MIME Part 3

Terminator

Authentication

- Authentication prevents wanton SPAMming
- Multiple authentication methods



- Plain text (should require encrypted SMTP)
- LOGIN, PLAIN – base64-encoded strings
- Encrypted (generally considered “secure”)
- CRAM-MD5, DIGEST-MD5, GSSAPI, NTLM, etc.

Before spam was such a pervasive problem, there was no need to authenticate users who were requesting to send e-mail. Server operators had a reasonable expectation that most or all SMTP transactions were legitimate and MTAs unquestioningly passed e-mail messages on toward their final recipients.

In the example illustrated on this slide, a spammer is using the improperly configured SMTP server at “mail.foo.com” to relay a forged message from “someone@foo.com”. Although this depicts just one message, even a small spamming operation can send many hundreds of thousands of messages per day using this methodology.

Obviously, that model was soon exploited by spammers, so RFC 2554 was created in 1999, establishing the SMTP-AUTH standard. (This RFC has since been replaced by RFC 4954.) SMTP-AUTH establishes multiple authentication methods, some of which are considered more secure than others.

At the low end of the scale are the LOGIN and PLAIN methods. As you'll see in a moment, these are trivially simple, consisting merely of base64-encoded username and password credentials. Remember: Encoded is not encrypted! These methods should never be used over unencrypted connections, and many SMTP servers will not advertise their availability unless SSL or TLS is in use. Of course, this is not an absolute, and many servers do permit these methods over unencrypted connections.

Because of this major security shortfall, the RFCs also created a number of authentication methods that provide a greater level of security. These are generally considered acceptable even over non-SSL/TLS connections because they provide their own cryptographic protections. Although we won't examine all of these methods in detail, it's important to note that there are multiple authentication methods and the two encoded methods can provide valuable credential evidence for the investigator.

Authentication Example

```
S: 220 esmtp.starklabs.com ESMTP
C: EHLO client.marvelmail.com
S: 250-esmtp.starklabs.com
S: 250-PIPELINING
S: 250-8BITMIME
S: 250-SIZE 255555555
S: 250 AUTH LOGIN PLAIN CRAM-MD5
C: AUTH LOGIN
S: 334 VXN1cm5hbWU6
    "Username:"
C: dG9ueWR1bmdhbgo=
    "tonystark"
S: 334 dG9ueXN0YXJr
    "Password:"
C: U3RAcmtMYWJzUEAkJAo=
    "St@rkLabsP@$@"
S: 535 authentication failed (#5.7.1)

S: 220 esmtp.starklabs.com ESMTP
C: EHLO client.marvelmail.com
S: 250-esmtp.starklabs.com
S: 250-PIPELINING
S: 250-8BITMIME
S: 250-SIZE 255555555
S: 250 AUTH LOGIN PLAIN CRAM-MD5
C: AUTH PLAIN
dG9ueXN0YXJrQHNOYXJrbGFicy5jb20AdG9ueXN0YXJrAE15U
EBTczIwcmRUb2RheQA=
    "tonystark@starklabs.com\x00tonystark\x00MyP@Ss20rdToday\x00"
S: 235 ok, go ahead (#2.0.0)
```

Shown here are examples of both “insecure” authentication methods we discussed on the previous slide—LOGIN and PLAIN. As shown in the server's capability listing, the esmtp.starklabs.com server can perform authentication using the LOGIN, PLAIN, and CRAM-MD5 methods.

After the client requests the LOGIN method, the server responds with the base64-encoded string “Username:”. The client responds in kind with the encoded username. The parties then perform a similar exchange for the password. In this case, the authentication attempt was unsuccessful.

In the second example, the client requests the PLAIN method. For this method, the client can immediately send an encoded string with either two or three components, each terminated with a null byte. These components are:

- Authorization identity (optional): The entity as which the client is attempting to authenticate.
- Authentication identity: The username
- Password

In the second case, the authentication request was successful, so the client would then initiate with the e-mail message submission process.

References:

<http://for572.com/37zjq>

Encryption

- Encrypted SMTP helps prevent observing and spoofing content
- SMTP options may not be available in plaintext
 - Plaintext/reversible authentication methods
- Two methods:
 - Native SSL: Starts SSL conversation before SMTP
 - STARTTLS: “Go secure” on plaintext connection

As you know, encryption can provide both confidentiality and integrity depending on how it is implemented. We'll discuss encryption at great length later, but let's briefly touch on its application within the context of SMTP.

Considering the inadequacy of the LOGIN and PLAIN authentication methods to protect the user's login credentials, encryption becomes very important in most SMTP transactions. Most SMTP servers will not even advertise these two methods unless the connection has been secured.

There are two primary mechanisms of securing the SMTP connection: SSL and STARTTLS. Although not as common today as it once was, the SSL mechanism starts the new TCP connection with an immediate exchange to encrypt the connection. After it is secured, the SMTP transaction starts.

On the other hand, with the STARTTLS mechanism, the connection starts as a standard, unencrypted one. The client requests that the server “goes secure”, at which point, an SSL negotiation occurs over the same connection. The benefit in using STARTTLS is that a separate port doesn't need to be allocated and managed across an architecture. In this model, an SMTP server may present different capabilities on the same connection, depending on whether a STARTTLS command has been issued and the negotiation completed.

STARTTLS Example

```
S: 220 esmtp.starklabs.com ESMTP
C: EHLO client.marvelmail.com
S: 250-esmtp.starklabs.com
S: 250-8BITMIME
S: 250-STARTTLS
S: 250-SIZE 255555555
S: 250 AUTH CRAM-MD5
C: AUTH LOGIN
S: 504 5.3.3 AUTH mechanism LOGIN not available
C: STARTTLS
```

Plaintext

S&C: <TLS negotiation and establishment>

```
C: EHLO client.marvelmail.com
S: 250-8BITMIME
S: 250-SIZE 255555555
S: 250 AUTH LOGIN PLAIN CRAM-MD5
C: AUTH LOGIN
S: 334 VXN1cm5hbWU6
C: dG9ueWR1bmdhbgo=
S: 334 UGFzc3dvcmQ6
C: TXlQQFNzMjByZFRvZGF5Cg==
S: 235 ok, go ahead (#2.0.0)
```

Encrypted

Single TCP
Connection

Here, you can see an example of the STARTTLS process. The client system connects to the esmtp.starklabs.com server, which advertises that it can use the STARTTLS mechanism as well as authenticate using the CRAM-MD5 method.

The client initiates the negotiation, and everything shown in the shaded region represents the conversation that occurs over the encrypted channel. The key point here is that the same connection is used throughout.

During the encrypted phase of the conversation, the server advertises that it will accommodate the LOGIN and PLAIN methods in addition to the CRAM-MD5 method. The client then selects the LOGIN method and proceeds with the SMTP transaction.

References:

<http://for572.com/lunt8>

Investigative Relevance

- Bad people still gotta talk!
- Data theft via e-mail (PCI, keylog, etc.)
 - The Perfect Keylogger, many others
- Network-based monitoring of spearphishing

Of course, we're most interested in knowing how SMTP is relevant during an incident response or investigation. After all SMTP is a rather old protocol, and probably not one that most malicious actors use... or do they?

One of the key things to remember is that bad guys still have to communicate somehow! In many cases, they'll use what's available—such as e-mail! Although it's unlikely that a nation-state actor would use e-mail within a victim's network to coordinate operations, a rogue employee might certainly use e-mail as a communication mechanism, or possibly a means of sending privileged information outside of the network.

In other pockets of the pure criminal world, software keyloggers are still used quite efficiently. In fact, many credit card readers are no more than keyboard devices from the computer's perspective, so a keylogger on a point-of-sale or cash register system is a common tactic for PCI data theft. Keylogger software can be configured to send captured data to an external e-mail account, providing attackers with a very easy data theft collection mechanism.

Another key presence of SMTP in the modern threat environment is the pervasive use of spearphishing. There is no discounting the fact that this attack method is both widely and efficiently used by strategic attackers to gain footholds within their targets' networks. Of course, the penultimate leg of every such message is across SMTP—right before the MDA puts it into the Inbox of their target's client software. The ability to examine SMTP exchanges can provide useful insight to the flow of spearphishing messages and their payloads. Understanding the nature of SMTP headers can show the hops a message took in getting to its destination, and knowing how they can be forged or altered can help to keep “red herring” findings in check. Using the MIME standard to effectively and efficiently carve attachments can provide a reverse engineering team with the initial dropper binary that an attacker used to gain access to the target's environment. Of course, analysis of this dropper may very well lead to additional network indicators that can be used to scope an incident and start to plan toward remediation.



Commercial Network Forensic Tools

This page intentionally left blank.

Value of Commercial Products

- When an organization procures software, there are many unseen background activities
- Before selecting a product consider:
 - Legal requirements
 - Support structure
 - Development roadmap
 - Deployment scalability

Commercial products are not “just like free products but they cost money”; there are a number of legal and support aspects that are important to organizations and businesses.

Although an open source product or tool may suit your needs today, businesses may ask many key questions surrounding the product, the development team, the environment, and their past and future plans.

Commercial contracts will necessitate appropriate pre-contract checks to ensure that the vendor is not bankrupt. After all, we certainly want the product we ordered to be delivered! Additionally, the contracts department will often check for conflicts of interest that may prevent the contracting with the vendor. Because the supply chain for many commercial products can vary greatly, there is also an opportunity for the buyer to validate the company and products they are about to purchase to some extent. As has been reported widely,^[1] the convoluted supply chain can sometimes thwart even the most formal of organizations.

The procurement process should ensure that the correct elements of the vendor’s product are specified in the contract (that what the organization wants is what is actually ordered) and that the contract is correctly fulfilled.

The process of ordering a product tends to add legitimacy to a project and can galvanize staff to help support and engage with those trying to deliver the new capability.

In the pricing and planning stage of the project, the buyer’s personnel should engage with the vendor’s pre-sales to ensure that they understand the costs associated with the full product lifecycle.

The legal input to the procurement process is generally the review of the contract from the purchasing department (if they are involved). However, it can also help in the review of contract terms.

The biggest difference between open source and the commercial products is that the latter don't often include total disclaimers that remove all liability from the authors for any damage or disruption caused by the installation and use of the software. This is an important delineation and demonstrates a level of responsibility for the code they developed and charge for. Importantly, if the tool fails to work correctly or fails to meet the contracted standard, then there should be some legal recourse to getting the vendor to fix the issue, send staff to configure/tune the product, and even provide compensation (or service credits). We live in a heavily litigious society, so when a large organization is looking to implement some new capability/technology or barrier, the lawyers like to draft a contract addressing legal recovery of damages if the product fails to meet expectations.

Additionally, if there is a compromise and the vendor's product did not function as required, there is some legal responsibility that the victim organization would certainly like to shift to the vendor. This is common if an attack or compromise was not detected by a product despite being correctly installed and managed. In the event of regulatory fines, this legal link and responsibility could prove to be a career lifeline.

Finally, there is the matter of the receipt of complete and packaged goods. The advantage of deploying a firewall from a trusted vendor is that the vendor's supply chain has hopefully been validated. For example, most Common Criteria evaluated products^[2] have their supply chain included in the evaluation, meaning that there is a process for tracking the device from manufacture to delivery. (This is usually in the Common Criteria Security Target.) If the organization has followed this procedure and extended it internally to the installation, patching and updating of the device in accordance with the product and evaluation guidelines, then the organization can demonstrate that they have followed the secure deployment process for the product. No such deployment guides or processes exist for open source software.

Support is vital from the outset. There are many considerations when looking at network forensics starting in the early requirements capture phase, matters such as network design, hardware location, and processing capacity will all need to be considered.

In the planning stages the following are critical to developing a good solution:

- The network architecture (including geographical span of deployment)
- The bandwidth and average usage to the Internet and between sites
- The assets considered most valuable/critical to the organization
- What they are (PCI, PII, IP, or other internal confidential data)
- Where they are located
- How they are protected
- How they would most likely be attacked
- What regulatory requirements must or should be considered in the solution
- The location and number of users of the admin/analysis console

Once implemented, good support will assist in tuning the equipment and training the buyer's staff. Generally, this is something that neither community nor forum-based support can match.

Finally, good solution vendors can often provide experienced consultants and equipment to help in the event of a major breach that requires extra staff, hardware, and processing power. This support is regularly covered by standing NDAs and ad-hoc consulting frameworks.

Many commercial organizations are underpinned by strategies and plans for improvement, development, and expansion. Because of these, it is in their interest to develop products and clients to match their growth plans. More importantly, they will have plans to ensure their offering keeps pace with operating system, hardware, and networking speed improvements. For example, a 1Gbps capture capability will not cope on a 10Gbps link, so a company offering only 1Gbps taps would not fare well for a good deal of the modern market.

Probably the most critical factor when evaluating security products is their approach and history of security patches and updates. Security monitoring tools that are vulnerable to either DoS or remote exploit attacks are extremely dangerous as an attacker could

- Gain control of the security tool
- Disable the tool—preventing detection of other attacks
- Destroy evidence of the attacker's actions
- Hide the exfiltration of data

In addition to security updates, host platform service packs and kernel releases patches can result in fundamental changes bringing additional capabilities and performance attributes, e.g., TCP/IP Offload Engine (TOE), IPv6, or direct SAN integration.

Scalability is probably the biggest challenge facing the network forensics analyst. Although a malware analyst can quickly identify the storage requirements for their tasks (a rough rule of thumb is for each machine: (target system total HDD size x 3 + memory size x 2.5) x 2 for redundancy/backup), the same cannot be said for the network analyst.

Full packet captures on fast links can quickly produce tens of terabytes of data, all of which needs to be parsed, indexed, and potentially linked to other packets, connections, hosts, or entities on the network. The products and tools should have multiple options for remote analysis, tunable QoS, and bandwidth-efficient methods for centralizing the data for analysis.

Additionally, the challenge facing the network analyst is the pace at which networks are increasing their speed at every point in the environment. Connectivity to the desktop, server networks, on virtual switches inside virtualized environments, on enterprise core switches and even WAN links to the Internet is expanding at an increasingly rapid rate. This means that faster capture and storage solutions are required to ensure that all desired data is actually captured and that packets are not lost. This problem is also being faced by the IDS/IPS community, which is struggling to provide the necessary sustained throughput for adequate detection and protection of assets.

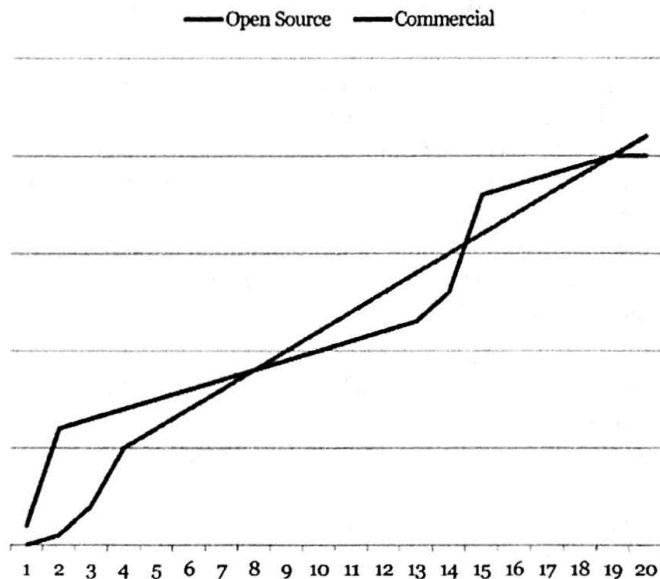
Commercial products tend to provide off-the-shelf solutions for this as it helps them to land large customers. Open source projects often support large enterprise deployments; however, such deployments tend to require a significant amount of planning, engineering, administration, and maintenance to keep in top shape.

References

- [1] <http://for572.com/ybgx7>
- [2] <http://for572.com/60onh>

Generic Cost Models: Open Source vs. Commercial

- Open source != free
- All solutions cost!
- Costs to buy and maintain are loaded at different times
- Consider lifecycle costs before selecting “the cheapest”



Before committing to a solution, consider both open source and commercial models as all solutions cost; some cost more in time and effort, others in money.

By means of an example, we have a hypothetical model represented in the cost chart on this slide. Initially the commercial tool installation costs time as the Request for Proposal (RFP) needs to be written and proposals assessed before being purchased and installed and tuned by skilled consultants. Meanwhile the open source solution needs to be fully researched, tested, and tuned, and the deployment planned and then conducted—all of which take time.

The initial open source solution costs less as it uses reclaimed equipment, or bare metal servers that don't attract the same [appliance] profit margins for the vendor.

Once installed, however, the refined update and management interface of the commercial tool is easier to use resulting in a lower support time and thus lower ongoing cost. Meanwhile the forum-based support for the open source tool has a higher cost due to the increased research and patching/updating activities necessary. (Open source tool authors tend to release individual patches as soon as they are complete rather than as rollup patches that commercial vendors tend to use.)

Remember, if a support technician who costs \$60k per year, or \$37 per hour (200 days at 8 hours per day), needs to spend 5 hours a week supporting an open source device, but only 1 hour a week to support a commercial tool, the increased cost of the open source tool is \$37 per hour x 4 extra hours x 52 weeks per year or an impressive \$7,800 per year.

Thus, the higher ongoing costs of supporting and managing the open source solution pushes the open source cost line above that of the commercial tool. However, just before the 4-year point (16 quarters) the commercial vendor knocks on the door to discuss a “hardware refresh,” whereupon another capital purchase is made to replace functioning, but no longer supported, equipment.

Commercial products have well-defined upfront costs, but they are generally high. Commercial products front load their costs, because many take the approach of an appliance with its own dedicated, supported hardware. This makes them appear much more expensive than open source solutions, which often omit the cost of hardware, favoring “extra” gear. However, once a vendor has been selected, pre-sales engineers from the vendor often provide support as they help to scope requirements and the operating environment that the solution requires.

A downside of commercial products is the lack of access to their internals. With an open source solution, a test deployment can be customized to meet the organization’s unique requirements. Commercial vendors tend to prefer webinars, videos, and demonstrations to highlight their value to potential buyers. These will certainly show the product in the best possible light and may not be relevant to the client’s intended use of the product. Therefore, it is recommended that test deployment systems be obtained from vendors before reaching a final decision. Most vendors commonly grant 30-day trial periods for serious potential buyers. Putting hands on an actual product in the target environment will provide excellent insight to the installation, configuration, and operational issues that may affect the full scale project.

Additionally, vendors usually provide incremental training on their products, reducing the time for staff to become proficient on the new equipment. This training is often conducted in the customer’s own environment, resulting in customized training and an element of tuning the new solution.

Open source products have many advantages in terms of a low initial monetary cost. Security deployments using open source software tend to start by using reclaimed or spare hardware. There are a number of hidden costs that many engineers and management overlook—research, testing, engineering, deployment, and ongoing care and feeding all require experienced personnel with potentially high salaries. These costs should be factored into the evaluation of any technical solution, but more often tend to be overlooked for open source solutions.

Before reaching the moment of software installation, the organization must undergo several steps:

- A basic requirements study—they know what they want to accomplish with the project
- Research into options—they have selected what they consider the best solution
- Identification of installation requirements—as they get the right hardware together
- Configuration review—to identify how to implement the features they want
- Troubleshooting—because new installations rarely go to plan

Furthermore, moving from the initial trial to a solution that will meet the needs of the enterprise will take time as the value of the trial needs to be assessed. After this and further research, a comprehensive solution must be planned, proposed, and purchased. If the open source solution is a market leader, such as the Snort IDS, trained and experienced consultants can be found from the open market. However, if the solution is more bespoke or the open source project new, external support may be limited.

The reason for the concern is that organizations that implement enterprise solutions are investing time and money in that solution. If the open source project dies, the developers move on, and updates stop, the organization will be left with a legacy tool that may start to fail in terms of reliability, sustainability, and flexibility.

Network Forensics Tool Types

- Tools/products/appliances fall generally into:
 - Full packet capture, collation, and storage
 - RSA NetWitness Investigator
 - NetFlow collection, collation, and storage
 - QRadar (Q1 Labs, acquired by IBM)
 - Analyst supporting tools
 - NetworkMiner (NETRESEC)
 - SteelCentral Packet Analyzer Personal Edition (Riverbed)
- Many product lines starting to converge across these boundaries

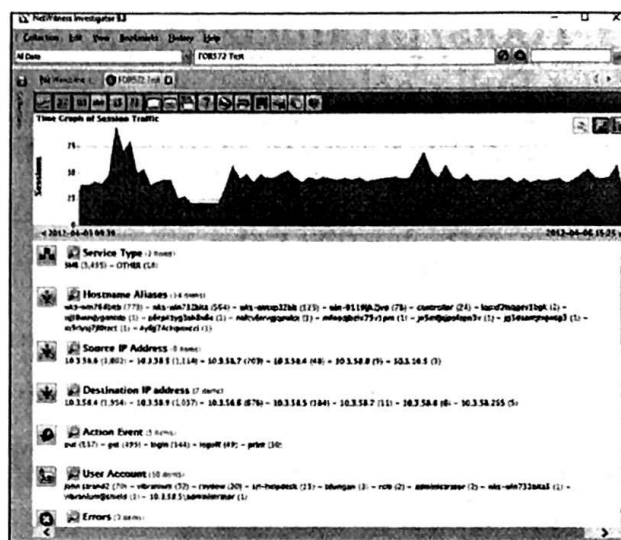
Commercial (and open source) network forensic tools tend to fall into one of three groups:

- **Full packet capture:** The most data-hungry of the lot. Commercial offerings in this group will generally bring the most expensive implementations, because large, fast data storage is still not cheap.
- **Summary capture:** This group allows for greatly reduced storage requirements associated with packet metadata or IPFIX data. This is also known as NetFlow, Sflow, or JFlow data.
- **Analyst supporting tools:** Tools that support the analyst's processes often simplify the display or presentation of large batches of network data, distilling them to a human-readable form.

We will conduct a brief review of some of the useful aspects of the listed products' capabilities.

RSA NetWitness

- Part of Security Analytics
- Distributed architecture
 - Decoders (Collectors)
 - Concentrator
 - Brokers
- Applications
 - Investigator
 - Informer/Visualize
- Uses “drills” to pivot through data



EMC/RSA has integrated most of the legacy NetWitness functionality into their Security Analytics platform. Security Analytics collects, indexes, and stores full packet data for as long as the user can afford the storage.

The core product includes decoders that collect the data from network choke points and perform various proprietary initial analysis activities such as extracting various bits of metadata. The decoders then compress and send the data to concentrators for consolidation and storage. Brokers query the concentrators' data on behalf of the human analyst.

At this point, we have only the data in a centralized location. The Security Analytics product portfolio allows the customer to purchase additional modules, which drill down into or mine the data. Investigator is a separate (and free) software package that provides the primary tools to allow an analyst to drill down into data chunks. Investigator can also attach to a fleet-style deployment of the broader Security Analytics suite, though this is not in the free tier of service.

The bright blue text items in this screenshot are links to pivot the data based upon the selection. This pivoting process gives NetWitness its power, because the analyst can drill into areas of interest based upon the information they have and deduce. Additionally, the session activity time chart (the top part of the screenshot) can also be used to zoom in on a specific time period, and upon applying the filter, the pivot links are adjusted to reflect the reduced time window. Some of the data points an analyst can use as pivot points include MAC addresses, source and destination IP addresses, services, and http actions (commands), e.g., `get/put/login`.

Scale is one of the strong points for Security Analytics and NetWitness's legacy before the acquisition. Although the solution is not cheap, it has been designed to meet large networks' requirements, including fast packet capture and high-volume storage.

NetWitness Investigator 9.8

Collection Edit View Bookmarks History Help

All Data FOR572 Test

FOR572 Test

Time Graph of Session Traffic

Sessions

< 2012-04-03 09:39 2012-04-06 15:25 >

- Service Type** (2 items)
 - SMB (3,455) - OTHER (18)
- Hostname Aliases** (14 items)
 - wks-win764bitb (773) - wks-win732bita (564) - wks-winxp32bit (125) - win-9119jkk2jvp (78) - controller (24) - lqpxf2lsqgev1bgk (2) - ujil8wuxiygdntdp (1) - p6epil1yg3ah8x6e (1) - nafcv6rvrjgqowkx (1) - mfwzjlbzix75v1pm (1) - jn5mfajpsfapn3v (1) - jg1dsantqjwsg3 (1) - iq9rlyvj7j0tact (1) - 4y6ij74ckqinxcci (1)
- Source IP Address** (6 items)
 - 10.3.58.6 (1,602) - 10.3.58.5 (1,114) - 10.3.58.7 (703) - 10.3.58.4 (48) - 10.3.58.8 (3) - 10.3.16.5 (3)
- Destination IP address** (7 items)
 - 10.3.58.4 (1,554) - 10.3.58.9 (1,037) - 10.3.58.6 (676) - 10.3.58.5 (184) - 10.3.58.7 (11) - 10.3.58.255 (5)
- Action Event** (5 items)
 - put (537) - get (495) - login (144) - logoff (49) - print (30)
- User Account** (10 items)
 - john strand2 (70) - rsydow (32) - vibranium (32) - srl-helpdesk (13) - tdungan (3) - rob (2) - administrator (2) - wks-win732bita\$ (1) - vibranium@shield (1) - 10.3.58.5\administrator (1)
- Errors** (3 items)

QRadar (QI Labs, Now IBM)

- SIEM with investigative features
 - IPFIX integration, full-packet capture
- Links analysis of flows to other log/event activity



SANSDFIR

FOR572.4 | Advanced Network Forensics and Analysis

36

QRadar is designed as a SIEM so it primarily digests logs (Windows Event Logs, security devices like firewalls, UNIX-like syslog data, etc.) and various IPFIX/NetFlow/JFlow/SFlow messages from routers and core switches. It then collates this information and raises contextual alerts and events.

QRadar provides a hybrid capability. The power of the product comes from the blending of the network connection data with the system alerts and logs. The QRadar device is configurable so that both network and log events can be linked to alert rules, allowing for faster response to malicious network traffic. This speed of reaction can be the difference between event containment and an environment-wide catastrophe.

The correlative value from a platform like QRadar can enhance an organization's investigative abilities significantly. This trend has seen a number of platforms with historical roots primarily in the log collection, network visibility, or packet capture markets merge into unified platforms aiming to meet as many customer requirements as possible in one "convenient" purchase.

IBM Security QRadar SIEM

Dashboard Overview: Vulnerability Management, Remote Dashboard, Delta Dashboard, Add Item...

Security News: Last updated Tue May 21 17:17:26 GMT 2013. Includes headlines like "Threat of Cyber Attacks Comes From China" and "Cisco to Use Intel-based Software Center for \$475 million".

Security Advisories: Last updated Tue May 21 17:17:26 GMT 2013. Includes "CVE-2013-0283 - Microsoft Exchange Server Remote Code Execution" and "CVE-2013-0284 - Microsoft Exchange Server Remote Code Execution".

Network All: Vulnerability Count table.

Vulnerability	Vulnerability Count
ICMP: Traceroute Request	86
Trace Route Information	84
Web Services: Denial of Service	58
2012-0814 - OpenSSH - Information Disclosure	41
2011-5000 - OpenSSH - Denial of Service Issue	41
OpenSSH: Public Parameter Validation	34
Shared Secret Authentication Attacks	28
SSL - Self-Signed Certificate	32
Information Leak - Network Information Disclosure	28
TRACE - Possible Attacker's IP Via Method	21
TRACE:TRACE - Cross-site tracking attack via HTTP	21

Vulnerability Count / Open Service: Bar chart showing counts for various protocols like http, ssh, https, etc.

Vulnerability Count / Asset: Bar chart showing counts for various IP address ranges.

PO Failures: Bar chart showing counts for various protocols.

Scan in Progress: Last updated Tue May 21 17:17:26 GMT 2013. Lists scans like "Nessus test scan - 2013-05-20 16:53:16".

Scans Completed: Last updated Tue May 21 17:17:26 GMT 2013. Lists completed scans like "Nessus test scan - 2013-05-20 16:53:16".

Impact AP: Vulnerability Count / Impact pie chart showing distribution of impact types like Disclosure, Downtime, etc.

Latest Published Vulnerabilities: Last updated Tue May 21 17:17:26 GMT 2013. Lists vulnerabilities like "CVE-2013-0283 - Microsoft Exchange Server Remote Code Execution".

IBM QRadar SECURITY INTELLIGENCE - INCIDENT OVERVIEW

Recent Incidents: Bubble chart showing incident counts and categories like "31 IP Connections", "19 SMTP Mail", "25 IP Connections (country: US, IP)", "10 IP Connections", "74", "28", "22", "8", "20", "21", "32", "33", "5", "82 Local POP", "54", "9", "106", "12", "68", "30", "70", "94", "10", "11", "25", "36", "27", "97", "34", "38", "13", "14", "29", "7", "16", "42", "1", "114", "24", "1", "107", "28", "66".

Incidents by Magnitude: 100

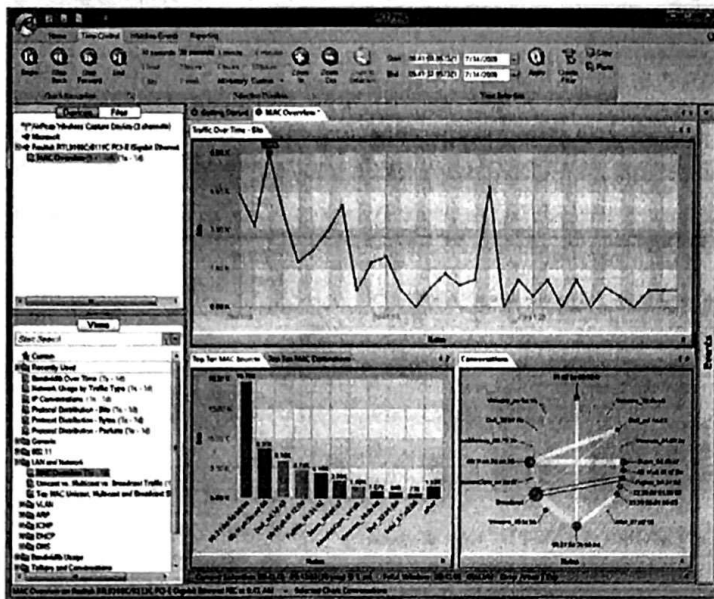
Average Incident Duration: 1 Minutes

Top Categories: Bar chart showing counts for various categories.

Top Incidents by Event Count: Pie chart showing distribution of event counts.

Security Intelligence News | Innovative New Solutions for Securing the Internet of Things

Riverbed SteelCentral Packet Analyzer Personal Edition



- Third-party Wireshark GUI
- Very fast even with large captures
 - Good to broadly search for activity
- Designed for report-ready visualizations

SAASDFIR

FOR5724 | Advanced Network Forensics and Analysis

38

Riverbed's SteelCentral Packet Analyzer Personal Edition is an excellent GUI support interface for much of Wireshark's deep-dig capabilities. Once a pcap file has been imported or acquired live with a AirPcap/TurboCap device, various analyses can be conducted on the file. The filters produce graphical output suitable for human visual analysis. The operator can apply several filters to narrow down their search. At any point, the selected data can be exported to Wireshark for other traditional work.

The tool works with large files mainly because it generates metadata for the capture file, and then allows the analyst to drill down based on the high-level rendering.


As an example of its performance, using just a standard Core i5 laptop (8GB RAM), it took just 2 minutes and 12 seconds to search and index a 5.5GB pcap file.

This type of product allows the operator to focus on doing the analysis regardless of the display limitations of Wireshark and the increasingly large files generated by fast networks.

The developer has also invested in the visualization and reporting aspects of the tool, with a goal of streamlining all phases of a typical case or investigation, while also supporting short-notice tactical examinations that are often involved with threat hunting operations.

Home | Time Control | Watches/Events | Reporting
 Begin | Step Back | Step Forward | End
 10 seconds | 30 seconds | 1 minute | 10 minutes
 1 hour | 3 hours | 6 hours | 12 hours
 1 day | 1 week | All History | Custom
 Selection Duration

Copy | Paste | Create Filter | Apply
 Start 09:41:03:957321 7/14/2009
 End 09:41:33:957321 7/14/2009
 Zoom In | Zoom Out | Zoom to Selection



MAC Overview
 Top Ten MAC Sources | Top Ten MAC Destinations | Conversations

Top Ten MAC Sources

MAC Address	Bytes
00:1f:c8:81:2d:a6	19.78K
00:1f:c8:81:2d:a6	8.30K
00:1f:c8:81:2d:a6	6.98K
00:1f:c8:81:2d:a6	4.74K
00:1f:c8:81:2d:a6	4.18K
00:1f:c8:81:2d:a6	2.88K
00:1f:c8:81:2d:a6	1.92K
00:1f:c8:81:2d:a6	1.07K
00:1f:c8:81:2d:a6	860
00:1f:c8:81:2d:a6	736
00:1f:c8:81:2d:a6	1.92K

Top Ten MAC Destinations

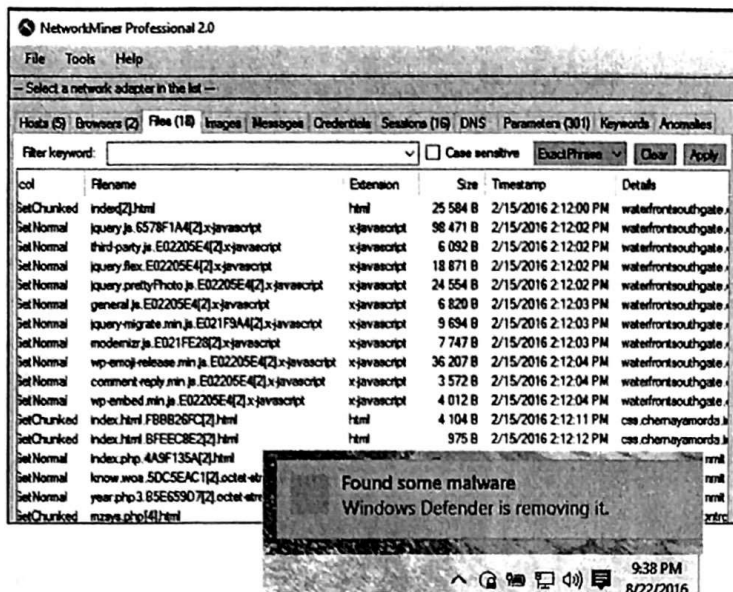
MAC Address	Bytes
00:21:64:3c:b5:ba	1.92K
00:21:64:3c:b5:ba	1.07K
00:21:64:3c:b5:ba	860
00:21:64:3c:b5:ba	736
00:21:64:3c:b5:ba	1.92K
00:21:64:3c:b5:ba	1.07K
00:21:64:3c:b5:ba	860
00:21:64:3c:b5:ba	736
00:21:64:3c:b5:ba	1.92K
00:21:64:3c:b5:ba	1.07K
00:21:64:3c:b5:ba	860
00:21:64:3c:b5:ba	736

Conversations

Start Search
 Custom
 Recently Used
 Bandwidth Over Time (1s - 1d)
 Network Usage by Traffic Type (1s - 1d)
 IP Conversations (1s - 1d)
 Protocol Distribution - Bits (1s - 1d)
 Protocol Distribution - Bytes (1s - 1d)
 Protocol Distribution - Packets (1s - 1d)
 Generic
 802.11
 LAN and Network
 MAC Overview (1s - 1d)
 Unicast vs Multicast vs Broadcast Traffic (1s - 1d)
 Top MAC Unicast, Multicast and Broadcast S
 VLAN
 ARP
 ICMP
 DHCP
 DNS
 Bandwidth Usage
 Talkers and Conversations

Current Selection: 09:41:03 - 09:41:33 (30 secs) @ 1 sec - Total Window: 09:41:03 - 09:43:44 - Drop After: 1 Day
 Selected Chart: Conversations

NetworkMiner (NETRESEC) (1)



- Commercial and free editions
- Deep packet digger
- Live or pcap
- Extracts:
 - Files & images
 - Mail messages
 - Credentials
- Exports to CSV

SANSDFIR

FOR572.4 | Advanced Network Forensics and Analysis

40

NetworkMiner is an exceptionally useful dual-function tool that can both passively sniff traffic from the local interface or read previously captured data in the pcap format.

When used in live mode, the interface allows for analysis while the tool is still capturing data. The decoding and session identification features run in real-time. This allows the tool to reassemble session contents, extract files and images, and decode various client-to-server messages including cookies, URL parameters, and credentials for plaintext protocols. These objects are written to disk in real time, so it's important to run NetworkMiner on a system that doesn't have antivirus or anti-malware software, as those files will be removed or quarantined as soon as they are written.

The simple GUI allows an analyst to monitor and investigate any aspect of the session while it continues to capture data. Furthermore the ability to stop recording but retain the session data and then restart capturing again is an advantage over Wireshark (where the user has to save and close each capture session before starting a new capture). This feature is excellent when analyzing particular system activity that occurs at known or regular times, such as malware beaconing or connections to C&C servers.

The free version lacks a few features from its commercial sibling, but the free license does permit using NetworkMiner on real-world cases. Often, such licenses restrict use to demonstration or educational use only.

-- Select a network adapter in the list --

Hosts (5) Browsers (2) Files (18) Images Messages Credentials Sessions (16) DNS Parameters (301) Keywords Anomalies

Filter keyword: Case sensitive ExactPhrase

col	Filename	Extension	Size	Timestamp	Details
SetChunked	index[2].html	html	25 584 B	2/15/2016 2:12:00 PM	waterfrontsouthgate
Set Normal	query.js.6578F1A4[2].x-javascript	x-javascript	98 471 B	2/15/2016 2:12:02 PM	waterfrontsouthgate
Set Normal	third-party.js.E02205E4[2].x-javascript	x-javascript	6 092 B	2/15/2016 2:12:02 PM	waterfrontsouthgate
Set Normal	query.flex.E02205E4[2].x-javascript	x-javascript	18 871 B	2/15/2016 2:12:02 PM	waterfrontsouthgate
Set Normal	query.prettyPhoto.js.E02205E4[2].x-javascript	x-javascript	24 554 B	2/15/2016 2:12:02 PM	waterfrontsouthgate
Set Normal	general.js.E02205E4[2].x-javascript	x-javascript	6 820 B	2/15/2016 2:12:03 PM	waterfrontsouthgate
Set Normal	query-migrate.min.js.E021F9A4[2].x-javascript	x-javascript	9 694 B	2/15/2016 2:12:03 PM	waterfrontsouthgate
Set Normal	modernizr.js.E021FE28[2].x-javascript	x-javascript	7 747 B	2/15/2016 2:12:03 PM	waterfrontsouthgate
Set Normal	wp-emoji-release.min.js.E02205E4[2].x-javascript	x-javascript	36 207 B	2/15/2016 2:12:04 PM	waterfrontsouthgate
Set Normal	comment-reply.min.js.E02205E4[2].x-javascript	x-javascript	3 572 B	2/15/2016 2:12:04 PM	waterfrontsouthgate
Set Normal	wp-embed.min.js.E02205E4[2].x-javascript	x-javascript	4 012 B	2/15/2016 2:12:04 PM	waterfrontsouthgate
SetChunked	index.html.F8BB26FC[2].html	html	4 104 B	2/15/2016 2:12:11 PM	css.chemayamorda.ji
SetChunked	index.html.BFECC8E2[2].html	html	975 B	2/15/2016 2:12:12 PM	css.chemayamorda.ji
Set Normal	index.php.4A9F135A[2].html	html	96 836 B	2/15/2016 2:12:14 PM	and.theleadersummit
Set Normal	know.woa.5DC5EAC1[2].octet-stream	octet-stream	428 052 B	2/15/2016 2:12:20 PM	and.theleadersummit
Set Normal	year.php3.B5E659D7[2].octet-stream	octet-stream	428 052 B	2/15/2016 2:12:31 PM	and.theleadersummit
SetChunked	mzsys.php[4].html	html	20 B	2/15/2016 2:12:31 PM	lect.ac.in/webcontn
SetChunked	mzsys.php[5].html	html	20 B	2/15/2016 2:13:15 PM	lect.ac.in/webcontn

Case Panel
 Filename MD5
 2016-02-15-Admedia... d731ab8a6f17215fe0bb825a4d27

Reload Case Files

NetworkMiner (NETRESEC) (2)

- Primarily a very focused analysis tool
- The decoding and collection of extracted data speeds up analysts' work
- Ideal tool for investigating malicious staff or rogue internal activity
- Pre-loaded keywords aid investigations
- Less useful at APT style detection activities
- Does not scale well—requires CapLoader

Given the extensive object extraction capabilities, NetworkMiner does suffer from somewhat slow processing speed when consuming very large packet captures. For this reason, it is best suited to use after data reduction preprocessing with `tcpdump` or another BPF-capable tool.

Although excellent for internal investigations, live investigations, and malware beaconing analysis, the tool does not scale to hunt potential APT breaches, as it cannot handle the volume by itself.

CapLoader (NETRESEC)

- Reads and indexes large pcap files
- Host/service discovery
- Export flows to NetworkMiner, Wireshark, etc.

The screenshot displays the CapLoader application window. At the top, there are menu options like 'File', 'View', 'Tools', and 'Help'. Below that, there's a 'Hosts' section with a search bar and a list of hosts. The main area is a table of network flows. The table has columns for 'Flow ID', 'Host', 'Port', 'Service', 'Flow Type', 'Status', 'Start', 'End', 'Duration', 'Bytes', 'Packets', 'Bytes/s', 'Packets/s', 'Flow ID', and 'Host'. The table contains several rows of data, including flow IDs like 1000000000, 1000000001, 1000000002, etc., and host names like 10.0.0.1, 10.0.0.2, etc.

The NETRESEC Company created a tool designed to handle large captures called CapLoader.

After being loaded, all network traffic is indexed into flows and presented to the user. Individual flows, or all traffic to/from a host, or using a specified protocol can then be exported to NetworkMiner for object extraction or to Wireshark for deep-packet exploration. This makes the tool a “GUI for data reduction” of sorts.

The tool is separate from NetworkMiner and is not a security tool. However, as a highly efficient data management tool, CapLoader can help to significantly speed up investigative and hunting workflows.

File Edit View Tools Help
 Identify protocols
 Show empty flows
 Parse DNS
 Frames limit: 10

Rows (443) Services (42) Hosts (7)
 Hide Selected Flows | Invert Hiding | Show All Rows | Selected Flows: 41
 Display Filter (BP)
 Keyword Filter

Clear Apply
 - Exact Phrase - Clear Apply

Flow ID	Flow Name	Src (bytes)	MD5	Detail	Endname	Full Path
1	stsk-20120403-f4f-emb-emb2.pcap	289 409 392	c7d129f6dc2...	ETHERNET	Little-endian	Z:\vdh\home\lmp\ludb_08_evidence\stsk-201204...

Auto-extract flows on
 Extracted Flows
 Rows: 41
 Filename: stsk-20120403-115f27194.pcap
 Src: 385 518 B

Flow ID	Client IP	Client Port	Server IP	Server Port	Transport	Hostname	Area Domain	Start	End	Duration	Frames	Payload Bytes	PCR	TCP Flags	Flow ID	Gen. Chart
1001	10.3.58.6	5575	10.3.58.9	139	TCP			4/4/2012 2:09 PM	4/4/2012 2:09 PM	00:00:14.5572770	16	3215	-0.16	AP	1	
1002	10.3.58.6	5575	10.3.58.9	445	TCP			4/4/2012 2:09 PM	4/4/2012 2:09 PM	00:00:14.5590700	16	4747	0.82	AP	1	
1003	10.3.58.5	5582	10.3.58.9	445	TCP			4/4/2012 2:11 PM	4/4/2012 2:12 PM	00:00:00.4118540	28	7255	0.17	AP	1	
1004	10.3.58.6	5583	10.3.58.9	445	TCP			4/4/2012 2:12 PM	4/4/2012 2:12 PM	00:00:00.2002167	14	4593	0.59	AP	1	
1005	10.3.58.5	43328	10.3.58.6	139	TCP			4/4/2012 2:12 PM	4/4/2012 2:13 PM	00:00:00.01658530	14	2191	-0.13	AP	1	
1006	10.3.58.5	43321	10.3.58.6	139	TCP			4/4/2012 2:12 PM	4/4/2012 2:13 PM	00:00:15.9115457	16	2284	-0.18	AP	1	
1007	10.3.58.7	4372	10.3.58.9	445	TCP			4/4/2012 2:16 PM	4/4/2012 2:18 PM	00:00:11.3023202	18	4358	0.43	AP	1	
1008	10.3.58.5	4372	10.3.58.9	445	TCP			4/4/2012 2:18 PM	4/4/2012 2:18 PM	00:00:10.5858670	16	5166	0.53	AP	1	
1009	10.3.58.6	5584	10.3.58.9	445	TCP			4/4/2012 2:21 PM	4/4/2012 2:21 PM	00:00:10.4358730	12	1756	-0.10	AP	1	
1040	10.3.58.6	5586	10.3.58.9	139	TCP			4/4/2012 2:21 PM	4/4/2012 2:21 PM	00:00:10.4246580	16	2216	-0.16	AP	1	
1041	10.3.58.6	5586	10.3.58.9	445	TCP			4/4/2012 2:21 PM	4/4/2012 2:21 PM	00:00:10.4185080	15	4747	0.82	AP	1	
1042	10.3.58.5	49323	10.3.58.6	139	TCP			4/4/2012 2:24 PM	4/4/2012 2:25 PM	00:00:15.8479540	14	2191	-0.13	AP	1	
1043	10.3.58.5	49324	10.3.58.9	139	TCP			4/4/2012 2:24 PM	4/4/2012 2:25 PM	00:00:15.8201790	16	2284	-0.18	AP	1	
1044	10.3.58.6	5587	10.3.58.9	445	TCP			4/4/2012 2:27 PM	4/4/2012 2:27 PM	00:00:14.1907730	16	5110	0.53	AP	1	
1045	10.3.58.7	2001	10.3.58.9	445	TCP			4/4/2012 2:31 PM	4/4/2012 12:23 PM	21:52:26.5331930	36	9716	0.48	AP	1	
1046	10.3.58.7	2021	10.3.58.5	445	TCP			4/4/2012 2:33 PM	4/4/2012 2:34 PM	00:01:06.5344450	12	1934	-0.06	AP	1	
1047	10.3.58.6	5589	10.3.58.9	445	TCP			4/4/2012 2:33 PM	4/4/2012 2:33 PM	00:00:13.2895400	12	1756	-0.10	AP	1	
1048	10.3.58.6	5590	10.3.58.9	139	TCP			4/4/2012 2:33 PM	4/4/2012 2:33 PM	00:00:13.2893960	16	4747	-0.16	AP	1	
1049	10.3.58.6	5591	10.3.58.9	445	TCP			4/4/2012 2:33 PM	4/4/2012 2:33 PM	00:00:13.2895400	16	4747	0.82	AP	1	
1050	10.3.58.5	49327	10.3.58.9	445	TCP			4/4/2012 2:33 PM	4/4/2012 2:33 PM	00:00:10.4735920	16	5108	0.53	AP	1	
1051	10.3.58.7	2053	10.3.58.5	445	TCP			4/4/2012 2:34 PM	4/4/2012 2:34 PM	00:00:00.0950000	335	237677	0.93	AP	1	
1052	10.3.58.7	2076	10.3.58.9	445	TCP			4/4/2012 2:36 PM	4/4/2012 2:36 PM	00:00:11.3844420	43	10190	0.54	AP	1	

Lab 09

NetworkMiner

This page intentionally left blank.

Lab 09 Objectives: NetworkMiner

- Familiarize yourself with basic operations of commercial network forensics tools
- Load pcap into NetworkMiner, conduct analysis and mining of the data
- Consider scalability of network forensics tools in your analytic workflow

This page intentionally left blank.

Lab 09 Takeaways: NetworkMiner

- Commercial tools often provide a polished GUI, pre/post sales advice, and support
 - Exposed features limited by developer's preference
- NetworkMiner quickly extracts many different objects from pcap files
 - Comprehensive capabilities come with speed costs
- Knowing strengths and weaknesses of many tools can free up analysts' time and capacity

This page intentionally left blank.

Wireless Network Forensics

This page intentionally left blank.

Limitations

- This is not a wireless security course!
- We focus on common 802.11 a/b/g/n network features/functions
- We will not cover Zigbee, Z-Wave, GSM, or Bluetooth...
 - But with the right hardware, Wireshark parses these, too!

With limited classroom time, we could never cover the full depth of this vast and complex subject without limiting ourselves to a reasonable set of protocols, medium, and frequency.

If you want to study this area further, we recommend you consider taking SANS SEC617, “Wireless Ethical Hacking, Penetration Testing, and Defenses,”^[1] where you’ll spend considerable time using the hardware and software to uncover the security issues of the likes of Zigbee, GSM, and Bluetooth.

References:

[1] <http://for572.com/g3ytb>

Why Is Wireless Such a Target?

Switched Network

- Logical circuit created between hosts for "privacy"
- Sniffing data requires active changes/attack
- Spoofing causes broad impact to users
- Defined boundary

Wireless Network

- Broadcast medium, requires host to be in range of transmission
- Sniffing requires only a wireless network card
- Spoofing is trivial
- No RF boundary

Wireless remains both a risk and exposed attack surface due to the nature of the medium through which it connects. Although wired communication has seen the introduction of switched communications as a method for reducing the numbers of hosts that can "see" the signal, wireless communication by its very nature is broadcast.

Although sniffing on a switch is made easier with tools like Cain and Abel, Bettercap, or Subterfuge, these attacks involve manipulating the environment in ways that have become readily detectable. Wireless sniffing, however, is completely passive and can often be conducted with the native OS and natively installed wireless card. If specialized long range or sensitive equipment is required, it is openly available from online retailers for commodity prices. Similarly, injecting spoofed traffic to a wired/switched network generally requires a significant technology footprint and often causes a noticeable impact to the users in the victim's environment. Wireless attacks often cause an observable negative impact as well, but depending on the attacker's goal, it's quite feasible for them to achieve their goals in a short period of time—before users start to complain.

Finally, while wired attacks require the individual to be physically connected to the infrastructure (which generally requires them to be inside the target's physical spaces), wireless networks do not impart that limitation. Indeed, with inexpensive but tuned equipment, an attacker could be several miles away from the target network and still collect the data being transmitted.

Actively injecting packets to the targeted wireless network would require the attacker to be in closer proximity to the victim's environment. However, in more densely populated areas like cities, a radius of 100 meters would include many businesses, apartments, and coffee shops—any of which could harbor a potential hidden attacker. Even in a more isolated environment, an attacker could easily conceal a miniaturized hardware platform and interact with it from a safe distance.


Put simply: Attackers focus on the weak links within their targets' environments. For all the preceding reasons, wireless is a decidedly weak link—and therefore a very attractive attack vector.

Modes: Master and Managed



Master

- Access Point (AP)
- BSSID is the AP's MAC
- Handles access and traffic controls



Managed

- Mobile clients/stations
- Have MAC address
- Connect to "SSID"


Master Mode or Infrastructure Architecture Mode

When a centralized device is given control of part of the RF Spectrum (the channel and SSID), it is deemed to be in infrastructure architecture mode. This device will be in **Master Mode** and for that BSSID, SSID, and frequency, it will determine who can communicate with it, associate to it, and following a successful challenge response (i.e., authentication), who can transmit and how much time they can transmit for.

Managed Mode

Stations that connect to access points (APs) and submit to their control are deemed to be in **Managed Mode**. In this mode, they will act upon the instructions of the Access Point, including the encryption and authentication to be used. The AP can also instruct the station to disassociate from the Access Point. Managed Mode is the default operating mode of laptops and mobile devices that regular users operate. When in this mode, any network sniffing conducted by the client will reveal typical Ethernet type frames, packets, and data. In this mode, the client is not able to see any of the 802.11 management communications from the AP to the client's RF radio on their wireless card.

Modes: Ad-Hoc and Monitor



Ad-Hoc (IBSS)

- Short-lived peer-to-peer
- LAN or PAN use
- File/print/setup services

Monitor (RFMON)

- Passive monitoring and collection
- Electronically undetectable
- Only one channel (frequency) at a time

SANSDFIR

FOR572.4 | Advanced Network Forensics and Analysis

52

Ad Hoc or Independent Basic Service Set (IBSS) Networks


When clients communicate directly with each other and don't rely on an AP to manage the wireless environment, they engage in "ad-hoc" networking. This is designed for point-to-point communication between devices to share small amounts of data. Common implementations are tethered phones, wireless printers, Apple's AirDrop file exchange protocol, and "setup" networks for embedded/IoT devices. Devices trust each other and try to manage communication between themselves by means of back off algorithms to ensure that both get a fair share of the airwaves.

Monitor or RFMON mode

In the fourth mode of operation, the wireless card no longer transmits but instead listens to all 802.11 traffic on a set frequency. An RFMON interface itself is not dangerous, however, if a packet capture utility is connected to the monitor interface, then *all* 802.11 traffic observed by the network card's radio on the frequency it is listening to will be captured. This includes the other network and RF management and control frames transmitted by the target AP and all other 802.11-enabled devices on that frequency within the range of the sniffing station.

RFMON is completely passive and cannot be electronically detected by others in the vicinity. The downside, from the attacker's point of view, is that they can listen only to one channel/frequency at a time. To address this, RFMON-focused tools may hop through multiple channels or easily permit the simultaneous use of multiple network interfaces to cover as wide a swath of the radio frequency spectrum as possible.


Basic Service Set (BSS)



In IEEE terms, there are several other important variations of this implementation.


The Basic Service Set (BSS) is the core of the 802.11 wireless networking concepts. It is a collection of one or more stations that are associated with an access point. The BSS can be identified by the BSS Identifier or BSSID, which is the MAC address of the Access Point to which all the stations are connected.

Extended Service Set (ESS)




A collection of access points with the same SSID that provides access to a common infrastructure and the stations that are associated with the APs is known as an Extended Service Set (ESS).

These are commonly deployed by organizations (and many homes) as users seek a seamless roaming experience. An ESS allows stations to move about the building connecting and reconnecting to different APs while still accessing internal resources.



BSS = Footprint of AP
 BSSID = 58:6d:af:a1:d6:d6
 SSID = FOR572




ESS = Total footprint of all APs
 BSSID = Various
 SSID = FOR572

Wireless Distribution System (WDS)

BSSID = 58:6d:af:a1:c7:f1
SSID = FOR572
Station MAC = 00:15:17:6d:a8:86

- “Extends” wireless network
- APs communicate with each other to relay packets to correct station



Relayed frames
between access
points populate all
four MAC
addresses in
802.11 header

BSSID = 58:6d:af:a1:d6:d6
SSID = FOR572
Station MAC = 14:10:9f:a3:dc:4b

SANSDFIR

FOR572.4 | Advanced Network Forensics and Analysis

56

A specialized configuration is a Wireless Distribution System (WDS). This is where an AP relays traffic from the internal network over a station-less link to another AP, for example, within a large building that uses wireless “extenders” or in a campus environment where access points have RF visibility to each other. There can be two or more APs in a WDS deployment to widen the physical footprint the network provides. It is interesting that in this mode, packets relayed between access points use all four MAC address fields in the 802.11 frame header, as we will see later in this material.

WARNING!

- What we discuss in this section may not be legal in your country!
- Always check with your legal staff before collecting or even monitoring Wi-Fi networks
- **BE CAREFUL!**

Note, in some parts of the world, even having these tools and hardware could get you arrested and potentially subject to fines or other punishments.

Always understand the laws of the country you are in, and those you are travelling to. If you do not require these tools (e.g., on vacation), consider leaving them at home.

If you require them for work purposes, consider having a letter of introduction that specifies your need for having the tools, what the tools are, and why you have them in that country/location.

Tools Required

- **Operating System:**
 - Linux drivers tend to permit Monitor Mode more readily than a card's MS Windows driver
- **Wireless card:**
 - MiniPCI card installed inside the laptop:
More covert, but lower power
 - External USB wireless adapter:
Less covert but more power, allow external antennas

Most Linux OS distributions have good wireless manipulation tools built in, but more importantly, they provide better control of the wireless adapter through better driver support and the inclusion of libpcap by default.

Windows can be augmented with drivers for certain wireless cards; for example, the Riverbed AirPcap USB wireless card^[1] is shipped with Windows drivers that will allow for promiscuous mode sniffing of 802.11 traffic.

When considering a wireless card choice, although internal cards are more covert, you are limited in that many laptops have only one internal MiniPCI slot. You may be fortunate that your laptop manufacturer has included an Atheros chipset card, which is suitable for most wireless capture activities (in a Linux OS).

With USB wireless adapters, you are not limited by slots or space, only by USB ports. Although some laptops can run only one device per port, others can power USB hubs with multiple USB wireless adapters. This can provide monitor mode visibility on multiple channels simultaneously.

Although entry-level analysis can be conducted with just one capable card, if you are performing wireless forensic examinations or collecting wireless data on a strategic basis, you'll likely want to invest in a good collection of USB wireless cards.

References:

[1] <http://for572.com/10oym>

Favorite USB Adapter – Alfa

- Alfa AWUS051NH USB 802.11 a/b/g/n card
 - 2.4 and 5GHz support
- Usually supplied with RP-SMA type antenna or 6dB flat panel and mount
- Ralink RT2770/2750 Chipset
 - Uses Linux rt2800usb driver



Perhaps the most popular hardware kits used in both strategic and tactical wireless network collection and monitoring operations are from the Alfa company. Although international availability is occasionally limited, the comparatively low price and the high power/high gain feature set cause investigators and penetration testers worldwide to seek out its gear.

These devices can use a variety of different antennas to best suit various use cases, whether covert or overt. The driver support is quite good on UNIX-like operating systems (including Apple's macOS), and drivers are available for some of the Microsoft Windows family of operating systems.

Software Wi-Fi Tools

- Wireshark
- Kismet
- NetSpot
- Aircrack-ng suite



There are a variety of software tools useful for analyzing 802.11 packets and traffic, including:

- **Wireshark**^[1]: This staple tool can also parse many fields from 802.11 wireless traffic.
- **Kismet**^[2]: Kismet is a passive wireless monitoring tool, potentially allowing users to conduct surveys in areas where they are not licensed or authorized to transmit but still legally permitted to collect. Kismet supports GPS input for mapping purposes.
- **NetSpot**^[3]: This is a mapping utility for Windows and macOS that can use an image as a background for mapping wireless coverage and available networks. It's quite ideal for locating rogue access points and dead spots alike.
- **Aircrack-ng**^[4]: Aircrack-ng is a fork of the original Aircrack tools that were developed to assist in the auditing of WEP and WPA implementations. The tools have been expanded over the years and now include the code from tools like coWPAtty and AirGraph-ng.
- **tcpdump**^[5]: Even tcpdump can now handle 802.11 traffic. Recent versions of this staple to every network forensicator's toolbox can handle acquiring wireless traffic when the network interface is in monitor mode.

In this course, we will be focusing on Wireshark, because the other tools are best used to capture live traffic for analysis. However, understanding the other tools and how they can be useful in identifying malicious activity is important for many wireless network forensic use cases.

References:

- [1] <http://for572.com/7itb1>
- [2] <http://for572.com/7d1z9>
- [3] <http://for572.com/mcqvz>
- [4] <http://for572.com/hk7o0>
- [5] <http://for572.com/7itb1>

NetSpot - New survey (E3-2012.netspd)

DISCOVER SURVEY EXPORT USER GUIDE ASK A QUESTION

Zone surveys: 31 Jun 5, 2012

South Hall

The signal-to-interference ratio is similar to the signal-to-noise ratio, but in this case the interference is specific to co-channel interference from other radio transmitters.

RESUME SCAN

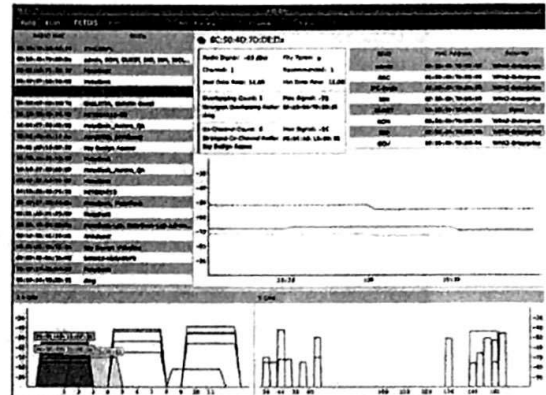
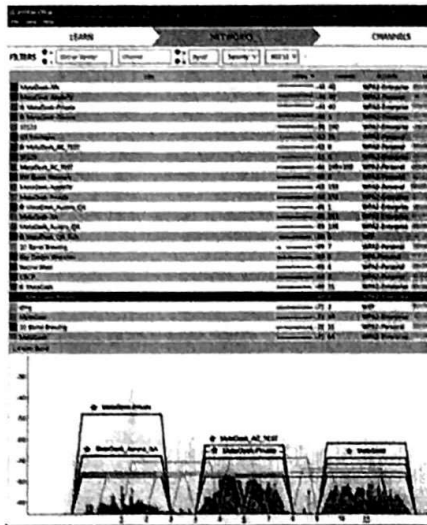
South Hall...re Show

Signal-to-interference ratio: 0dB, 43dB, 36dB

Device	MAC Address	Signal Strength (dBm)
ASUS_5G	10:BF:48:53:A4:34 / ASUS...	-45
ASUS...	10:BF:48:53:A6:8C / ASUS...	-45
ASUS...	C8:60:00:94:AB:A5 / ASUS...	-45
Audio 17	20:AA:4B:0D:8D:A9 / Cisco...	-45
Audio 23	20:AA:4B:0D:8D:AA / Cisco...	-45
BGI Tech Wifi	60:33:4B:E2:D2:8F / Apple, 5...	-45
Cisco42971	60:33:4B:E2:D2:90 / Apple...	-45
Disney	58:6D:8F:56:4B:9A / Cisco...	-45
E3 2012	58:6D:8F:56:4B:9B / Cisco...	-45
Xirrus...	00:0F:7D:9A:B2:90 / Xirrus...	-45
Xirrus...	00:0F:7D:9A:B2:C0 / Xirrus...	-45
Xirrus...	00:0F:7D:9A:B2:F0 / Xirrus...	-45
Xirrus...	00:0F:7D:9A:B3:70 / Xirrus...	-45

InSSIDER for Windows and macOS

- InSSIDER for Windows
- InSSIDER for macOS
- Supports ALL 802.11 frequencies/channels!



SANSDFIR

FOR572.4 | Advanced Network Forensics and Analysis

62

The following are some of the tools available for MS Windows:

- InSSIDer^[1]
- Eye P.A.^[2]

Metageek has also maintained updates for their InSSIDer software^[3] for Windows and macOS. A useful, low-cost addition to the investigator's tool set, InSSIDer for Mac supports all the channels and frequencies that the host device supports. This is a good feature as Apple hardware generally has excellent support for the latest 802.11 Wi-Fi standards. The software is also available for Android on the Google Play store.

Commercial tools for detecting rogue Access points include:

- VisiWave^[4]
- Ekahau Site Survey^[5]
- NetSpot^[6]

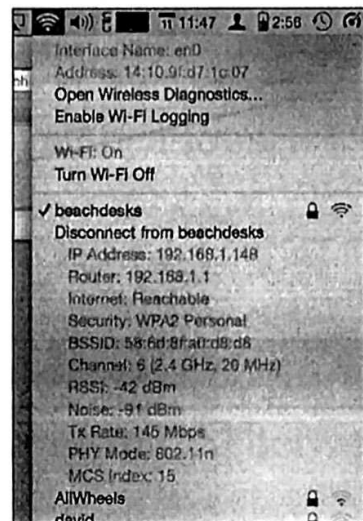

Although these tools have great commercial grade support, the pricing may be high for organizations that don't have a continuous need for wireless visibility.

References:

- [1] <http://for572.com/ong5->
- [2] <http://for572.com/4t-yd>
- [3] <http://for572.com/4ui-6>
- [4] <http://for572.com/d6tc5>
- [5] <http://for572.com/r-mx5>
- [6] <http://for572.com/mcqv6>

Built-in macOS Wi-Fi Tools

- Apple macOS has built-in monitor mode and easily accessible diagnostics



Apple's macOS provides a handy sniffer built into the `airport` tool (which is buried deep in the file system). The binary is located in the following location (for macOS Sierra – other versions may vary):

```
/System/Library/PrivateFrameworks/Apple80211.framework/Versions/Current/Resources/airport
```

The syntax to use the sniffer is simply `'airport <device> sniff <channel>'`.

```
$ sudo airport en0 sniff 6
```



You will notice the regular wireless icon changes (see the left image on this slide) to what some refer to as “The Eye of Sauron” (in deference to the *Lord of the Rings* trilogy's all-seeing eye).

The output is only written to a file when the process exits. The output files are written to the `/tmp/` directory and have the following format:

```
/tmp/airportSniffXXXXXX.cap
```

One additional handy tip for Mac users: When connected to an Access Point, if you hold the “Option” key before clicking the Wi-Fi icon, you will get the technical details of the device to which you are connected. This “advanced settings” view is in the screenshot on the right.

The 802.11 Frame



Note: When you see the icon shown at the top right corner of this slide, refer to the SANS 802.11 Pocket Reference Guide that was provided with your course materials. It can also be downloaded from Josh Wright's website.^[1]

The 802.11 frame is like most OSI-based protocols in that it is layered and contains nested data. Of particular note is that 802.11 is a Layer 2 framing protocol, meaning it replaces the Ethernet header. This framing is performed at the wireless driver, so higher-layer tools may not have the knowledge that the packet was/will be sent across a wireless connection. This is important because it may change the analyst's understanding of the traffic depending on what utility was used to perform the traffic capture.


The majority of the control information is contained in the first section called "Frame Control". In this section, the type and sub-type of frame are defined allowing the receiving station to understand what it should do with the contents.

Over the next few slides, we will cover the key parts of 802.11 frames so that you as an analyst can move seamlessly from the wired to the wireless environment.

References:

[1] <http://for572.com/6sto->

802.11 Frame Type



Protocol Version	Frame Type	Frame Subtype	ToDS	FromDS	More Fragments	Retry	Power Management	More Data	Security	Reserved
------------------	------------	---------------	------	--------	----------------	-------	------------------	-----------	----------	----------

- Type 0 – Management frames
 - These include probes, beacons, authentications, associations, and others
- Type 1 – Control frames
 - Request-to-send, clear-to-send, acknowledgments
- Type 2 – Data frames
 - The actual data (in IP form) to be transmitted between the devices

There are three types of 802.11 frames:

- Management frames
- Control Frames
- Data frames


Unless you have a card that is able to go into Monitor mode, you will never see the Management or Control frames, nor will you see the 802.11 headers of the data frames.

- **Management frames** are used to manage the BSSID's service area (i.e., its BSS). Management frames control who can connect to an AP and who can stay connected to the AP, thus they are limited to the associated (or attempting to associate) stations.
- **Control frames** are used to manage the medium (RF in this case). Control frames are transmitted to everyone regardless of the BSSID they are associated to as control frames manage who can communicate on the shared medium. Although the inherent forensic value of these frames is low, they often provide valuable insight to other activity on the wireless network, such as congestion and certain denial of service methodologies.
- **Data frames** are the actual data that the whole network is there to actually move from host A to host B. Data frames are addressed to individual nodes on the wireless network and have upper TCP/IP protocol data (i.e., the actual IP packets) in the payload.

References:

<http://for572.com/f40a9>

802.11 Management Frame Subtype



Protocol Version	Frame Type	Frame Subtype	ToDS	FromDS	More Fragments	Retry	Power Management	More Data	Security	Reserved
------------------	------------	---------------	------	--------	----------------	-------	------------------	-----------	----------	----------

- Association Req. (0x00)
- Association Resp. (0x01)
 - Successful code: (0x0000)
- Reassoc. Req. (0x02)
- Reassoc. Resp. (0x03)
- Probe Request (0x04)
- Probe Response (0x05)
- Beacon frame (0x08)
- ATIM (0x09)
- Disassociation (0x0a)
- Authentication (0x0b)
- Deauthentication (0x0c)

Management frames, as their name suggests, are used to manage the communications on any WLAN, whether it's an AP beaconing out its supported speeds and SSID or a station sending out a probe request looking for an AP to which it was previously connected.


Management frames are important from a forensic point of view for several reasons:

- Management (and Control) frames are not encrypted
- They allow clients and APs to be located by triangulating their power or by manually plotting locations and relative power settings
- Attackers will look in these frames for information about the BSS's and ESS's capabilities
- Attackers manipulate APs and their clients by injecting or corrupting these frames

Therefore, knowing how to capture these frames and then decode and understand their contents are important skills when handling a wireless-based incident or investigation.

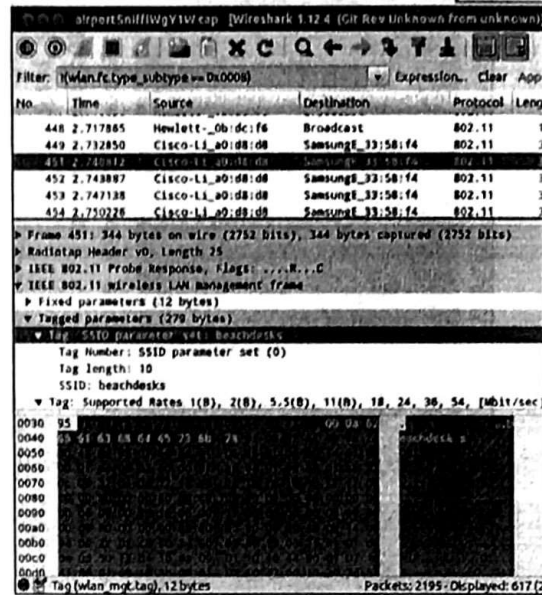
802.11 Fields: SSID and Other Tagged Parameters

• SSID: Service Set Identifier



Common Management Tag Values

0	SSID	1	Supported data rates
2	Frequency Hopping Channel Set	3	Direct Sequence Channel Set
4	Contention Free period	5	Traffic Indication Map
6	IBSS (Ad-hoc) parameter set	7	Country Information
0x30	RSN Information Element	0x85	Cisco CCX Extensions 1
0x88	Cisco CCX Extensions 2	0x95	Cisco CCX Extensions 3
0x2D	High Throughput (.11n) capability	0x34	AP Neighbor Report
0x3d	High Throughput (.11n) Information	0x2E	QoS Capability
0x22	Transmit Power Control Request	0x23	Transmit Power Control Response
0x24	Supported Channels	0x32	Extended supported data rates



The Service Set Identifier (SSID) is a textual tag that identifies a wireless network to the users. It is transmitted in the AP's beacons that are a type of Management frame. Beacons are commonly transmitted between 10 and 20 times a second, making them one of the most common management frame subtypes in 802.11 traffic.

In the past, it was suggested that restricting beacon frames from being broadcast would increase the security of a network by making it a "Hidden Network" or "Cloaked Network." However, this is a technical fallacy because the SSID is present—in plaintext format—in other management frames including those sent during the wireless network association process. If an attacker wanted to learn the SSID of a targeted network, he could simply spoof a de-authentication packet to a station, and then monitor the re-association process, which includes the SSID.

The SSID, along with numerous other values in the 802.11 frame header, is contained in a field of "tagged values". A list of common tags are on your handout and shown above. Each tag consists of a single byte that tells what field follows, then a length value, and finally the tag itself. It is important to recognize that these tags often include critical information to the normal operation of the wireless network, including how it is protected, what authentication methods are allowed, etc.

Filter: `(wlan.fc.type_subtype == 0x0008)` Expression... Clear Apply Save

No.	Time	Source	Destination	Protocol	Length	Info
448	2.717865	Hewlett-Ob:dc:f6	Broadcast	802.11	127	Data, SN=5, FN=0, Flags=.P...F.C
449	2.732850	Cisco-Li_a0:d8:d8	SamsungE_33:58:f4	802.11	344	Probe Response, SN=2395, FN=0, Flags=.....
451	2.740812	Cisco-Li_a0:d8:d8	SamsungE_33:58:f4	802.11	344	Probe Response, SN=2395, FN=0, Flags=.....R
452	2.743887	Cisco-Li_a0:d8:d8	SamsungE_33:58:f4	802.11	344	Probe Response, SN=2395, FN=0, Flags=.....R
453	2.747138	Cisco-Li_a0:d8:d8	SamsungE_33:58:f4	802.11	344	Probe Response, SN=2395, FN=0, Flags=.....R
454	2.750226	Cisco-Li_a0:d8:d8	SamsungE_33:58:f4	802.11	344	Probe Response, SN=2395, FN=0, Flags=.....R

Frame 451: 344 bytes on wire (2752 bits), 344 bytes captured (2752 bits)

Radiotap Header v0, Length 25

IEEE 802.11 Probe Response, Flags:R...C

IEEE 802.11 wireless LAN management frame

Fixed parameters (12 bytes)

Tagged parameters (279 bytes)

Tag: SSID parameter set: beachdesks

Tag Number: SSID parameter set (0)

Tag length: 10


SSID: beachdesks

Tag: Supported Rates 1(B), 2(B), 5.5(B), 11(B), 18, 24, 36, 54, [Mbit/sec]

0030	95	98 63 f3 c8 7b 00 00 00 64 00 11 04 00 0a 62	eachdesk s...
0040	65	61 63 68 64 65 73 6b 73 01 08 82 84 8b 96 24	eachdesk s...
0050	80	48 6c 03 01 06 2a 01 06 2f 61 06 30 14 01 00	eachdesk s...
0060	00	0f 9c 04 01 00 00 0f 8c 04 01 00 00 0f ac 02	eachdesk s...
0070	0c	00 32 04 0c 12 18 60 2d 1a fc 18 1b ff ff 00	eachdesk s...
0080	09	00 00 00 00 00 00 00 00 00 00 00 00 00 00 00	eachdesk s...
0090	00	00 00 00 3d 16 08 00 17 00 00 00 00 00 00 00	eachdesk s...
00a0	00	00 00 00 00 00 00 00 00 00 00 00 00 4a 0e 14 00	eachdesk s...
00b0	0a	00 2c 01 c8 00 14 00 05 00 19 00 7f 01 01 d6	eachdesk s...
00c0	6e	00 50 f2 04 19 4a 00 01 10 19 4a 00 01 02 10	eachdesk s...
00d0	af	00 01 00 10 3b 00 01 02 10 47 00 10 3d 7d 54	eachdesk s...

Tag(wlan.mgmt.tag), 12 bytes Packets: 2195 · Displayed: 617 (28.1%) · Load time: 0:0... Profile: Default

802.11 Control Frame Subtype



Protocol Version	Frame Type	Frame Subtype	ToDS	FromDS	More Fragments	Retry	Power Management	More Data	Security	Reserved
------------------	------------	---------------	------	--------	----------------	-------	------------------	-----------	----------	----------

- There are three control frame subtypes:
 - Request-to-send (0x1b)
 - Clear-to-send (0x1c)
 - Acknowledgment (0x1d)

As we identified earlier, the RF spectrum is a hub/broadcast medium and as a result, there are transmission collisions.

With wired Ethernet, all stations access the same medium. Therefore, the standard “Carrier Sense, Multiple Access with Collision Detection” (CSMA/CD) is used. Because stations in this model can transmit and receive data at the same time, they can detect if they attempt to transmit at the same time another station is doing so, and delay a miniscule amount of time to avoid crosstalk.

However, in a wireless environment, clients are not always in range of each other and therefore cannot determine if another station in the BSS is utilizing the shared medium of the RF spectrum. This is referred to as the “hidden node” problem.^[1] Further, once a station starts transmitting a data stream, it can no longer receive traffic on that frequency—including any control frames that would indicate another station needs to send traffic. Therefore, in the wireless realm, Carrier Sense, Multiple Access with Collision Avoidance (CSMA/CA) is used.

In CSMA/CA, control frames are used with specific subtypes to manage access to the frequency, which prevents rather than detects collisions. There are three Control Frame subtypes: “Request to Send”, “Clear to Send”, and “Acknowledgment”. These frames are used by the in-RF-range stations and access points to manage the use of the RF spectrum.

These frames are also unencrypted, and are also unique among 802.11 participants in that these frames are “per channel” and not tied to any specific SSID or BSSID. Rather, all stations and access points on a given frequency will exchange these messages to share that medium they all use.

References:

[1] <http://for572.com/-wva2>

802.11 Data Frames



Protocol Version	Frame Type	Frame Subtype	ToDS	FromDS	More Fragments	Retry	Power Management	More Data	Security	Reserved
------------------	------------	---------------	------	--------	----------------	-------	------------------	-----------	----------	----------


- There are lots of different data frame subtypes, including the Null function (0x24), indicating no data
- However, the most interesting is likely to be those where the “Logical-Link Control Type” is IP (0x0800)

Data frames carry the upper layer protocols (TCP/IP) within their payload section of the frame, thus 802.11 Data frames are the Layer 2 protocol within which IP is encapsulated.

References:

<http://for572.com/-95md> (section 7.2.2)

802.11 Data Frame Subtype



Protocol Version	Frame Type	Frame Subtype	ToDS	FromDS	More Fragments	Retry	Power Management	More Data	Security	Reserved
------------------	------------	---------------	------	--------	----------------	-------	------------------	-----------	----------	----------

- Many data frame subtypes, but most forensically relevant is “IP” traffic (0x0800)
- This is the only frame type that can be encrypted, indicated by the single “Security” bit
 - No differentiation between WEP/WPA/WPA2/etc.
 - Old tools assume all encrypted traffic is WEP

Data frames carry the upper layer protocols (TCP/IP) within their payload section of the frame, thus 802.11 data frames are the Layer 2 protocol within which IP is encapsulated. This is where most network forensic processes will focus, as we are typically looking at data payloads.

Data frames are the only 802.11 frame type that can be encrypted, but this is indicated in the frame header by just one single bit. Therefore, it is not possible to further differentiate WEP from WPA from WPA2 or any other encryption parameters. This differentiation is made using the tagged parameters discussed previously.

Old tools (such as the venerable but woefully outdated NetStumbler) assume that all frames with the security bit set are WEP because when they were written, WEP was the only security that was available at the 802.11 frame level. These tools should be replaced with something newer.

References:

<http://for572.com/-95md> (section 7.2.2)

802.11 Frame Control Fields: ToDS and FromDS (I)





Protocol Version	Frame Type	Frame Subtype	ToDS	FromDS	More Fragments	Retry	Power Management	More Data	Security	Reserved
------------------	------------	---------------	------	--------	----------------	-------	------------------	-----------	----------	----------

- The “DS” is the Distribution System
 - Backend infrastructure that “feeds” the 802.11 segments
- Analysts can determine the flow of traffic by examining these bit flags:

	To DS	From DS
Wired to Wireless	0	1
Wireless to Wired	1	0
Ad Hoc	0	0
WDS Network	1	1

The ToDS and FromDS bit flags allow the direction of the frame to be identified. When a frame is being sent from one station to another, it is actually “reflected” by the access point, which also rewrites the header, effectively re-framing the transmission. The directionality bits will indicate the leg of the transmission the frame was on at the time it was captured. This also means that typically, a passive capture including station-to-station traffic will have duplicated content.

802.11 Frame Control Fields: ToDS and FromDS (2)



Address Order	
From DS Set, To DS Clear:	From DS Clear, To DS Set:
Address 1: Destination	Address 1: BSSID
Address 2: BSSID	Address 2: Source
Address 3: Source	Address 3: Destination
From DS Clear, To DS Clear:	From DS Set, To DS Set:
Address 1: Destination	Address 1: Receiver
Address 2: Source	Address 2: Transmitter
Address 3: BSSID	Address 3: Destination
	Address 4: Source

Frame Control Sub-Field Data	
Protocol: 0, only supported protocol identifier	More Frag: Set, more fragments remaining
Type:	Retry: Set, packet is being retransmitted
0 Management Frame	Power Management: Set, STA is entering power conservation state
1 Control Frame	More Data: Set, AP has more buffered frames for STA
2 Data Frame	WEP/Privacy Bit: Set, data frame is encrypted using WEP, TKIP or CCMP
Subtype: Function of the frame based on frame type	Strict: Set, station requires frames to be delivered in order
From DS set, To DS Clear: From Wired to Wireless	
From DS clear, To DS Set: From Wireless to Wired	
From DS clear, To DS Clear: Ad-hoc is type is data	
From DS Set, To DS Set: WDS network	

Here Host A sends a packet, say an ICMP Echo Request, to Host B. However, the stations are in a BSS and as such they are only communicating with the access point, not directly with each other. Therefore, the ICMP message is transmitted to the AP (not directly Host B) with the ToDS bit-flag set and the FromDS bit-flag is clear.

When the access point retransmits the packet containing the ICMP Echo Request to the Host B, the FromDS bit-flag will be set and the ToDS bit-flag will now be cleared.

If sniffing this in Managed mode, the ICMP Echo request will be observed normally, and from the pcap, the user may not even notice that Host B is wireless client. However, if seen by a system in Monitor Mode, both legs will be seen, appearing to be a duplicated packet unless carefully inspected to identify the directionality of each within the BSS.

The re-framing process also rewrites the MAC addresses in the 802.11 frame header. Just like on wired networks all adapters are required to have a unique 6-byte hexadecimal Media Access Control (MAC) addresses. These values include both the BSSID and the MAC address of the station(s), depending on the leg on which the frame is currently being sent.

In the 802.11 header, there are always at least three address fields, and four in a WDS. The content of these fields is context-dependent. It is the ToDS and FromDS bits that define the context and thus the values that are placed in the fields.

Although the packet reference is concise, it does not mention that the same addresses can be in multiple fields within the same frame. If the Access Point is sending a frame to a Wi-Fi host (say the admin), the MAC address of the AP (also the BSSID) will be in Address location 2 and 3; similarly, the reply will go from the Wi-Fi host to the AP and the AP's MAC will be in field 1 and 3. This can cause confusion for new analysts as they are not used to seeing a MAC address in two places in a raw networking capture.

Managed Mode Sniffing

The image displays two screenshots of the Wireshark network protocol analyzer interface, showing traffic captured in Managed Mode. Both screenshots have a filter set to 'icmp'.

Left Screenshot:

- Packet List:**

No.	Time	Source	Destination	Protocol	Length	Info
752	2.756458	192.168.1.148	74.125.226.35	ICMP	68	
753	2.776718	74.125.226.35	192.168.1.148	ICMP	68	
754	3.759010	192.168.1.148	74.125.226.35	ICMP	98	
762	4.763259	192.168.1.148	74.125.226.35	ICMP	98	
764	4.786785	74.125.226.35	192.168.1.148	ICMP	98	
765	5.767335	192.168.1.148	74.125.226.35	ICMP	98	
- Packet Details:**
 - Ethernet II, Src: Apple_d7:1c:07 (14:10:9f:d7:1c:07), Dst: Cisco-L1_a0:d8:d6 (58:6d:8f:a0:d8:d6)
 - Destination: Cisco-L1_a0:d8:d6 (58:6d:8f:a0:d8:d6)
 - Source: Apple_d7:1c:07 (14:10:9f:d7:1c:07)
 - Type: IP (0x0800)
 - Internet Protocol Version 4, Src: 192.168.1.148 (192.168.1.148), Dst: 74.125.226.35 (74.125.226.35)
 - Internet Control Message Protocol
 - Type: 0 (Echo (ping) request)
 - Code: 0
 - Checksum: 0xc9b9 [correct]
 - Identifier (BE): 54429 (0xd49d)
 - Identifier (LE): 40A04 (0xbdd4)
- Packet Bytes:** 0000 58 6d 8f a0 d8 d6 14 10 9f d7 1c 07 08 00 45 00 70E.
0010 00 54 4a 88 d0 d0 40 01 a1 46 c0 08 01 94 4a 7d .T...@.AP...J.
0020 e2 23 ..
0030 [REDACTED]
0040 [REDACTED]
0050 [REDACTED]
0060 [REDACTED]

Right Screenshot:

- Packet List:**

No.	Time	Source	Destination	Protocol	Length	Info
752	2.756458	192.168.1.148	74.125.226.35	ICMP	68	Echo (ping) request Id=0xd49d
753	2.776718	74.125.226.35	192.168.1.148	ICMP	68	Echo (ping) reply Id=0xd49d
754	3.759010	192.168.1.148	74.125.226.35	ICMP	98	Echo (ping) request Id=0xd49d
762	4.763259	192.168.1.148	74.125.226.35	ICMP	98	Echo (ping) request Id=0xd49d
764	4.786785	74.125.226.35	192.168.1.148	ICMP	98	Echo (ping) reply Id=0xd49d
765	5.767335	192.168.1.148	74.125.226.35	ICMP	98	Echo (ping) request Id=0xd49d
- Packet Details:**
 - Frame 753: 98 bytes on wire (784 bits), 98 bytes captured (784 bits)
 - Ethernet II, Src: Cisco-L1_a0:d8:d6 (58:6d:8f:a0:d8:d6), Dst: Apple_d7:1c:07 (14:10:9f:d7:1c:07)
 - Destination: Apple_d7:1c:07 (14:10:9f:d7:1c:07)
 - Source: Cisco-L1_a0:d8:d6 (58:6d:8f:a0:d8:d6)
 - Type: IP (0x0800)
 - Internet Protocol Version 4, Src: 74.125.226.35 (74.125.226.35), Dst: 192.168.1.148 (192.168.1.148)
 - Internet Control Message Protocol
 - Type: 0 (Echo (ping) reply)
 - Code: 0
 - Checksum: 0xd39f [correct]
 - Identifier (BE): 54429 (0xd49d)
- Packet Bytes:** 0000 14 10 9f d7 1c 07 58 6d 8f a0 d8 d6 08 00 45 20X.....E
0010 00 54 b4 59 d0 d0 35 01 e2 52 4a 7d e2 23 c0 a0 .T.Y..S..RJ)..
0020 01 94 ..
0030 [REDACTED]
0040 [REDACTED]
0050 [REDACTED]
0060 [REDACTED]

In Managed Mode, the user's view of the captured traffic is akin to that of traditional wired sniffing, with none of the 802.11 underlying information being visible.

For this to be conducted, the user must be connected to an AP (including having the necessary PSK, WEP, or other credentials). As the user is connected to the AP, they will see the data free of any 802.11-based encryption.

managed_sniffing.pcap [Wireshark 1.12.4 (Git Rev Unknown from unknown)]

Filter: icmp Expression... Clear Apply Save

No.	Time	Source	Destination	Protocol	Length	Info
752	2.756458	192.168.1.148	74.125.226.35	ICMP	98	Echo (ping) request id=0xd49d,
753	2.776718	74.125.226.35	192.168.1.148	ICMP	98	Echo (ping) reply id=0xd49d,
758	3.759010	192.168.1.148	74.125.226.35	ICMP	98	Echo (ping) request id=0xd49d,
763	4.763259	192.168.1.148	74.125.226.35	ICMP	98	Echo (ping) request id=0xd49d,
764	4.786765	74.125.226.35	192.168.1.148	ICMP	98	Echo (ping) reply id=0xd49d,
765	5.767535	192.168.1.148	74.125.226.35	ICMP	98	Echo (ping) request id=0xd49d,

▼ Ethernet II, Src: Apple_d7:1c:07 (14:10:9f:d7:1c:07), Dst: Cisco-Li_a0:d8:d6 (58:6d:8f:a0:d8:d6)

- Destination: Cisco-Li_a0:d8:d6 (58:6d:8f:a0:d8:d6)
- Source: Apple_d7:1c:07 (14:10:9f:d7:1c:07)
- Type: IP (0x0800)

► Internet Protocol Version 4, Src: 192.168.1.148 (192.168.1.148), Dst: 74.125.226.35 (74.125.226.35)

▼ Internet Control Message Protocol

Type: 8 (Echo (ping) request)

Code: 0

Checksum: 0xcb9f [correct]

Identifier (BE): 54429 (0xd49d)

Identifier (LE): 40404 (0x9dd4)

```

0000 58 6d 8f a0 d8 d6 14 10 9f d7 1c 07 08 00 45 00  Xm.....E.
0010 00 54 4a 86 00 00 40 01 41 46 c0 a8 01 94 4a 7d  .TJ...e. AF....J}
0020 e2 23 08 00 0b 9f d4 9d 00 00 55 50 ee c4 00 0f  ..#.....UP...
0030 28 a3 08 09 0a 0b 0c 0d 0e 0f 10 11 12 13 14 15  .....
0040 16 17 18 19 1a 1b 1c 1d 1e 1f 20 21 22 23 24 25  .....
0050 26 27 28 29 2a 2b 2c 2d 2e 2f 30 31 32 33 34 35  .....
0060 36 37

```

managed_sniffing.pcap [Wireshark 1.12.4 (Git Rev Unknown from unknown)]

Filter: icmp Expression... Clear Apply Save

No.	Time	Source	Destination	Protocol	Length	Info
752	2.756458	192.168.1.148	74.125.226.35	ICMP	98	Echo (ping) request id=0xd49d,
753	2.776718	74.125.226.35	192.168.1.148	ICMP	98	Echo (ping) reply id=0xd49d,
758	3.759010	192.168.1.148	74.125.226.35	ICMP	98	Echo (ping) request id=0xd49d,
763	4.763259	192.168.1.148	74.125.226.35	ICMP	98	Echo (ping) request id=0xd49d,
764	4.786765	74.125.226.35	192.168.1.148	ICMP	98	Echo (ping) reply id=0xd49d,
765	5.767535	192.168.1.148	74.125.226.35	ICMP	98	Echo (ping) request id=0xd49d,

► Frame 753: 98 bytes on wire (784 bits), 98 bytes captured (784 bits)

▼ Ethernet II, Src: Cisco-Li_a0:d8:d6 (58:6d:8f:a0:d8:d6), Dst: Apple_d7:1c:07 (14:10:9f:d7:1c:07)

- Destination: Apple_d7:1c:07 (14:10:9f:d7:1c:07)
- Source: Cisco-Li_a0:d8:d6 (58:6d:8f:a0:d8:d6)
- Type: IP (0x0800)

► Internet Protocol Version 4, Src: 74.125.226.35 (74.125.226.35), Dst: 192.168.1.148 (192.168.1.148)

▼ Internet Control Message Protocol

Type: 0 (Echo (ping) reply)

Code: 0

Checksum: 0xd39f [correct]

Identifier (BE): 54429 (0xd49d)

```

0000 14 10 9f d7 1c 07 58 6d 8f a0 d8 d6 08 00 45 20  .....Xm.....E
0010 00 54 b4 59 00 00 35 01 e2 52 4a 7d e2 23 c0 a8  .T.Y..5. .RJ}.#..
0020 01 94 00 00 d3 9f d4 9d 00 00 55 50 ee c4 00 0f  ..#.....UP...
0030 28 a3 08 09 0a 0b 0c 0d 0e 0f 10 11 12 13 14 15  .....
0040 16 17 18 19 1a 1b 1c 1d 1e 1f 20 21 22 23 24 25  .....
0050 26 27 28 29 2a 2b 2c 2d 2e 2f 30 31 32 33 34 35  .....
0060 36 37

```

RFMON Sniffing – WPA2

- WPA2-protected data frames effectively opaque when captured in monitor mode
- Later acquisition of key material allows decryption

No.	Time	Source	Destination	Protocol	Length	Info
1495	2016-08-23 19:39:31.9149..	Apple_9e:4a:57	Broadcast	802.11	105	Data, SN=3675, FN=0, Flags=.pm...F.C
1496	2016-08-23 19:39:31.9150..	Apple_9e:4a:57	Broadcast	802.11	105	Data, SN=3676, FN=0, Flags=.pm...F.C
1497	2016-08-23 19:39:31.9151..	Apple_9e:4a:57	Broadcast	802.11	105	Data, SN=3677, FN=0, Flags=.pm...F.C
1498	2016-08-23 19:39:31.9155..	Apple_9e:4a:57	Broadcast	802.11	105	Data, SN=3678, FN=0, Flags=.pm...F.C
1499	2016-08-23 19:39:31.9156..	Apple_9e:4a:57	Broadcast	802.11	105	Data, SN=3679, FN=0, Flags=.pm...F.C
1500	2016-08-23 19:39:31.9157..	Apple_9e:4a:57	Broadcast	802.11	105	Data, SN=3680, FN=0, Flags=.pm...F.C
1501	2016-08-23 19:39:31.9162..	Apple_9e:4a:57	Broadcast	802.11	105	Data, SN=3681, FN=0, Flags=.pm...F.C
1502	2016-08-23 19:39:31.9163..	Apple_9e:4a:57	Broadcast	802.11	105	Data, SN=3682, FN=0, Flags=.pm...F.C
1503	2016-08-23 19:39:31.9164..	Ubiquiti_9b:48:d5	Broadcast	802.11	123	Data, SN=3683, FN=0, Flags=.pm...F.C

▶ Frame 1496: 105 bytes on wire (840 bits), 105 bytes captured (840 bits)
▶ Radiotap Header v0, Length 25
▶ 802.11 radio information
▶ IEEE 802.11 Data, Flags: .pm...F.C
▶ Data (44 bytes)
Data: 3387280abaa1eb36522f082b7529afb988884aefb64ed2b8...
[Length: 44]

If the user is not connected to the AP, then the only sniffing that can be conducted is using RFMON and locking the radio in the Wi-Fi card to the frequency or channel of the targeted BSSID and SSID. However, without the Pre-Shared Key (PSK), any encrypted data frames will not be readable.

This screenshot shows that Wireshark reflects data frames (created by ARP scanning the wireless subnet) as arbitrary "Data" because the content is not accessible without the encryption key material. Although the attacker can sniff this data, no further decoding is possible without the PSK. The only information that can be used would be the Layer 2 addresses, timing, and volume of the traffic. In this case, the ARP scan activity stands out because of its size.

However, if the attacker retains this pcap file and later obtains the wireless network's authentication material (PSK, WEP key, or PairWise Master Key (PMK)^[1] for the session), they will be able to decrypt the traffic. The encryption material could be gained through brute-forcing, social engineering, or other means.

References:

[1] The PMK is what is generated during the 4-way handshake between an AP and a station. From the PMK, the session keys are derived. Therefore, knowledge of the PMK (by knowing the PSK and seeing the 4-way handshake) will render that session's encryption compromised.

RFMON Sniffing – No WPA/WEP

```
▼ Radiotap Header v0, Length 25
Header revision: 0
Header pad: 0
Header length: 25
▶ Present flags
MAC timestamp: 27750717
▶ Flags: 0x10
Data Rate: 1.0 Mb/s
Channel frequency: 2437 [BG 6]
▶ Channel type: 802.11g (0x0480)
SSI Signal: -64 dbm
SSI Noise: -95 dbm
Antenna: 0
▼ IEEE 802.11 Data, Flags: .....F.C
Type/Subtype: Data (0x0020)
▶ Frame Control Field: 0x0802
.000 0000 0000 0000 ▶ Duration: 0 microseconds
Receiver address: Broadcast (ff:ff:ff:ff:ff:ff)
Destination address: Broadcast (ff:ff:ff:ff:ff:ff)
Transmitter address: Cisco-Li_a8:64:a7 (00:23:69:a8:64:a7)
BSS Id: Cisco-Li_a8:64:a7 (00:23:69:a8:64:a7)
Source address: Apple_d7:1c:07 (14:10:9f:d7:1c:07)
Fragment number: 0
Sequence number: 930
▶ Frame check sequence: 0x56ff4649 [correct]
▼ Logical-Link Control
▶ Address Resolution Protocol (request)
Hardware type: Ethernet (1)
Protocol type: IP (0x0800)
Hardware size: 6
Protocol size: 4
Opcode: request (1)
Sender MAC address: Apple_d7:1c:07 (14:10:9f:d7:1c:07)
Sender IP address: 192.168.75.27 (192.168.75.27)
```

```
▼ Radiotap Header v0, Length 25
Header revision: 0
Header pad: 0
Header length: 25
▶ Present flags
MAC timestamp: 27751405
▶ Flags: 0x12
Data Rate: 54.0 Mb/s
Channel frequency: 2437 [BG 6]
▶ Channel type: 802.11g (0x0480)
SSI Signal: -63 dbm
SSI Noise: -95 dbm
Antenna: 0
▼ IEEE 802.11 Data, Flags: .....F.C
Type/Subtype: Data (0x0020)
▶ Frame Control Field: 0x0802
.000 0000 0010 1100 ▶ Duration: 44 microseconds
Receiver address: Apple_d7:1c:07 (14:10:9f:d7:1c:07)
Destination address: Apple_d7:1c:07 (14:10:9f:d7:1c:07)
Transmitter address: Cisco-Li_a8:64:a7 (00:23:69:a8:64:a7)
BSS Id: Cisco-Li_a8:64:a7 (00:23:69:a8:64:a7)
Source address: Cisco-Li_a8:64:a5 (00:23:69:a8:64:a5)
Fragment number: 0
Sequence number: 931
▶ Frame check sequence: 0xb45f720e [correct]
▼ Logical-Link Control
▶ Address Resolution Protocol (reply)
Hardware type: Ethernet (1)
Protocol type: IP (0x0800)
Hardware size: 6
Protocol size: 4
Opcode: reply (2)
Sender MAC address: Cisco-Li_a8:64:a5 (00:23:69:a8:64:a5)
Sender IP address: 192.168.75.1 (192.168.75.1)
```

Here the sniffing was again with RFMON, however, the network is not encrypted at the 802.11 level (no WEP/WPA/WPA2).

This provides the same endpoint, timing, and volumetric analysis, as well as content examination. Wireshark's encapsulation view model will provide visibility to both the framing and the content, with various field names at Layer 2 and up available for use.

No.	Time	Source	Destination	Protocol	Length	Info
1495	2016-08-23 19:39:31.9149...	Apple_9e:4a:57	Broadcast	802.11	105 Data, SN=3675, FN=0, Flags=.pm..F.C	
1496	2016-08-23 19:39:31.9150...	Apple_9e:4a:57	Broadcast	802.11	105 Data, SN=3676, FN=0, Flags=.pm..F.C	
1497	2016-08-23 19:39:31.9151...	Apple_9e:4a:57	Broadcast	802.11	105 Data, SN=3677, FN=0, Flags=.pm..F.C	
1498	2016-08-23 19:39:31.9155...	Apple_9e:4a:57	Broadcast	802.11	105 Data, SN=3678, FN=0, Flags=.pm..F.C	
1499	2016-08-23 19:39:31.9156...	Apple_9e:4a:57	Broadcast	802.11	105 Data, SN=3679, FN=0, Flags=.pm..F.C	
1500	2016-08-23 19:39:31.9157...	Apple_9e:4a:57	Broadcast	802.11	105 Data, SN=3680, FN=0, Flags=.pm..F.C	
1501	2016-08-23 19:39:31.9162...	Apple_9e:4a:57	Broadcast	802.11	105 Data, SN=3681, FN=0, Flags=.pm..F.C	
1502	2016-08-23 19:39:31.9163...	Apple_9e:4a:57	Broadcast	802.11	105 Data, SN=3682, FN=0, Flags=.pm..F.C	
1503	2016-08-23 19:39:31.9164...	Ubiquiti_9b:48:d5	Broadcast	802.11	123 Data, SN=3683, FN=0, Flags=.pm..F.C	

▶ Frame 1496: 105 bytes on wire (840 bits), 105 bytes captured (840 bits) on RadioTap
 ▶ RadioTap Header v0, Length 25
 ▶ 802.11 radio information
 ▶ IEEE 802.11 Data, Flags: .pm..F.C
 ▶ Data (44 bytes)
 Data: 3387288aba1eb36522f082b7529af9880884aefb64ed2b8...
 [Length: 44]

```

RadioTap Header v0, Length 25
Header revision: 0
Header pad: 0
Header Length: 25
Present flags
MAC timestamp: 27750717
Flags: 0x10
Data Rate: 1.0 Mb/s
Channel frequency: 2437 [BG 6]
Channel type: 802.11g (0x0480)
SSI Signal: -64 dbm
SSI Noise: -95 dbm
Antenna: 0

IEEE 802.11 Data, Flags: .....F.C
Type/Subtype: Data (0x0020)
Frame Control Field: 0x0802
Duration: 0 microseconds
Receiver address: Broadcast (ff:ff:ff:ff:ff:ff)
Destination address: Broadcast (ff:ff:ff:ff:ff:ff)
Transmitter address: Cisco-L1_a8:64:a7 (00:23:69:a8:64:a7)
BSS Id: Cisco-L1_a8:64:a7 (00:23:69:a8:64:a7)
Source address: Apple_d7:1c:07 (14:10:9f:d7:1c:07)
Fragment number: 0
Sequence number: 930
Frame check sequence: 0x56ff4649 [correct]

Logical-Link Control
Address Resolution Protocol (Request)
Hardware type: Ethernet (1)
Protocol type: IP (0x0800)
Hardware size: 6
Protocol size: 4
Opcode: request (1)
Sender MAC address: Apple_d7:1c:07 (14:10:9f:d7:1c:07)
Sender IP address: 192.168.75.27 (192.168.75.27)
  
```

```

RadioTap Header v0, Length 25
Header revision: 0
Header pad: 0
Header Length: 25
Present flags
MAC timestamp: 27751405
Flags: 0x12
Data Rate: 54.0 Mb/s
Channel frequency: 2437 [BG 6]
Channel type: 802.11g (0x0480)
SSI Signal: -63 dbm
SSI Noise: -95 dbm
Antenna: 0

IEEE 802.11 Data, Flags: .....F.C
Type/Subtype: Data (0x0020)
Frame Control Field: 0x0802
Duration: 44 microseconds
Receiver address: Apple_d7:1c:07 (14:10:9f:d7:1c:07)
Destination address: Apple_d7:1c:07 (14:10:9f:d7:1c:07)
Transmitter address: Cisco-L1_a8:64:a7 (00:23:69:a8:64:a7)
BSS Id: Cisco-L1_a8:64:a7 (00:23:69:a8:64:a7)
Source address: Cisco-L1_a8:64:a7 (00:23:69:a8:64:a7)
Fragment number: 0
Sequence number: 931
Frame check sequence: 0xb45f720e [correct]

Logical-Link Control
Address Resolution Protocol (Reply)
Hardware type: Ethernet (1)
Protocol type: IP (0x0800)
Hardware size: 6
Protocol size: 4
Opcode: reply (2)
Sender MAC address: Cisco-L1_a8:64:a5 (00:23:69:a8:64:a5)
Sender IP address: 192.168.75.1 (192.168.75.1)
  
```

Key Difficulties in Wi-Fi Attacks

- Difficult to detect at low volume
- Difficult to validate
- Used to manipulate APs and clients
- Used to exploit weaknesses in 802.11 standard



One significant problem with wireless networks is the fact that anyone can generate packets and place them in the medium. Additionally, an attacker doesn't need to be in the building/campus or in the organization's property—they can attack from the coffee shop across the road.

Therefore, data can easily be injected to manipulate both access points and client stations to behave in ways that benefit attackers by causing them to send sensitive/useful data back into the wireless environment, or by exploiting any weaknesses in the 802.11 standards themselves. This is compounded by the fact that 802.11 management frames are necessarily unencrypted, making it quite trivial for someone intent on causing harm to forge these packets.

With only basic wireless equipment, the analyst will struggle to identify injected packets and frames. Identification is possible, but without tools like Wireless IDS (WIDS), the manual checking for injected data does not scale much past the classroom.

The Main Attacks

- WPA/WPA2 – Offline: PSK dictionary attack,
Online: Forged de-authentication
 - Note: “WPA2” formally known as “RSN”
- DoS – Online: RF or protocol attacks
- Evil Twin – Online: Spoofed access point

We'll briefly consider some of the types of attacks against networks protected with typical measures. In any case, consider that these attacks can be either online or offline. An “online” attack is one in which the attacker interacts with the environment, whereas an “offline” attack requires only RFMON collection.

- Wireless Protected Access (WPA/WPA2)
 - Offline attack: Brute-forcing Pre-shared Key
 - Offline attack: Force stations to leave the network, divulging partial authentication data upon automatic/immediate re-association
- Denial of Service (DoS)
 - Online attack: Can use inherent nature of RF to kill signal-to-noise ratio or attack protocol weaknesses
- Evil Twin
 - Rogue AP advertises known or expected SSIDs, exploiting stations' typical “automatic association to known SSID” behavior

One note to keep in mind while examining encrypted wireless traffic is that the standard commonly known as “WPA2” is actually the “Robust Security Network” or “RSN” standard. Wireshark uses the RSN designation in its parsing functions.

WPA/WPA2 PSK Attacks

- WPA and WPA2 deemed “stronger” than WEP
 - WPA superseded by WPA2 in 2006
- Main attacks focus on pre-shared keys
 - Requires capture of 4-way handshake during association
 - Session keys derived from: SSID, Station/AP MAC addresses, Station/AP nonce, and pre-shared key
- Offline: Brute-force captured 4-way handshake(s)
 - Detection: Physical detection of traffic capture
- Online: Spoof de-auth to cause more handshakes
 - Detection: Wireless IDS

The WPA and WPA2 cryptographic implementations are currently deemed secure, so attackers have focused on the values used to generate session keys. These keys are part of the Pairwise Master key (PMK) that is derived from various values that are transmitted in the clear.

This attack requires the interception of the following values:

- Network SSID
- MAC address of the Station
- MAC address of the Access Point
- A NONCE transmitted by the Station during the association process
- A NONCE transmitted by the Access Point during the association process

These are all transmitted when every client connects to the Wireless network through what is called a 4-way handshake. The attacker lacks only the PSK (the encryption password) that is chosen by the admin establishing the network—the other five values are passed across the connection in the clear in management frames. The Aircrack Wiki has a detailed discussion of the entire authentication process including packet captures for successful and failed attempts.^[1]

A completely passive attack (as accomplished by Josh Wright's **coWPAtty** tool^[2]), requires a capture containing a 4-way handshake and a dictionary of words that might be the password.

However, an impatient attacker, or one attacking a relatively quiet environment, may not have time to wait for a “natural” 4-way handshake to occur. Instead, they could simply craft a packet that appears to come from the access point that directs a client station to de-associate itself from the network. The wireless driver on the client will seek to re-establish the connection as soon as possible for usability purposes. This re-association process requires a new 4-way handshake—of course, the attacker will be listening for this exchange, capturing it for offline brute-forcing.

Detecting a passive WPA/WPA2 attack can be quite difficult—as with any passive attack. This would require spotting the attacker while collecting the traffic, or the so-called “shady character with a big antenna parked in the parking lot” method. Obviously, this is not ideal, but it is a harsh reality in the realm of wireless investigations.

An online attack may be more easily detected with the use of a Wireless IDS (WIDS). These platforms can spot potentially spoofed traffic by keeping a baseline of typical behavior in the environment, and then alerting when traffic deviates from those baselines. In this case, knowing the normal rate of “de-authentication” packets would be a useful metric to track.



References:

[1] <http://for572.com/ivujc>

[2] <http://for572.com/dr0b3>

DoS Attack: RF Overload

- With a shared broadcast medium, DoS is easy
 - RF spectrum can be polluted with a simple jammer
 - Unintentional DoS-by-overload



SANSDFIR

FOR5724 | Advanced Network Forensics and Analysis

83

Before looking at the 802.11 protocol, the fact that the RF medium is also vulnerable to DoS attacks should also be considered. The items to the left on the slide are the hardware used in an example online guide^[1] on how to make an RF Jammer.

By overloading the RF environment on a set frequency (say 2.4Ghz to 2.4835Ghz), all 802.11 traffic would effectively be blocked. NOTE: Transmitting on a licensed frequency is illegal unless you own the license and conducting jamming activities is also quite likely illegal. SANS does not recommend, advise, or condone RF Jamming, illegal transmission on licensed, or unlicensed frequencies. The device shown is absolutely illegal in most countries.


However, an RF DoS does not have to be intentional. The screenshot to the right of the slide was taken during a security audit of a video game authoring studio. The staff complained that it struggled to get traditional laptop wireless working in its building. Upon checking it was discovered that the 300+ Microsoft Xbox consoles and 250+ Sony PlayStation consoles were generating sufficient 2.4Ghz RF to overload all 802.11 traffic operating in the same range.

References:

[1] <http://for572.com/1ui9m>

DoS Attack: 802.11 Management and Control Frames

- Spoofed de-authentication attack: Everybody off!
 - Uses forged management frames
 - Sends clients into de-auth/re-auth loop
 - Sent to client station: Forge AP-to-client de-auth
 - Sent to AP: Force AP-to-broadcast de-auth
- RTS/CTS control frame attack
 - If everyone is waiting for forged CTS to expire, no one can transmit



De-authentication Attacks^[1]

A de-authentication attack is conducted by an attacker spoofing de-authentication packets to control the behavior of the stations within their broadcast range. These use management frames, meaning they are unencrypted and unauthenticated. The attacker can target a single client station or the entire BSS of an access point. To target a single host, the spoofed packet is created with the BSSID as the source and the target as the destination. The wireless driver on the victim's system will interpret this packet directly, and dutifully remove itself from the network. Of course, the system still wants to communicate, so it needs to re-associate to the access point. This happens quickly, but if the attacker continues to flood the channel with these, the victim will never remain connected for any meaningful amount of time. To target the entire BSS, the attacker forges the same packet, but with a destination of the broadcast address—so all client stations in range will receive and process the de-authentication packet.

De-authentication attacks can be devastating to the target and are surprisingly simple to implement by the attacker. Although the 802.11 standard does not provide a means to authenticate these packets, a WIDS would see the unusual spike in such packets and alert the administrator to the odd behavior.

RTS/CTS Attacks

802.11 uses the airwaves for Layer 1, which are a shared medium much like that of a hub-built wired environment. However, one critical difference is that each station in the wireless environment may not have full visibility to all other hosts in the same environment. Therefore, at the 802.11 level, traffic collisions are avoided through the use of the CSMA/CA feature, as discussed previously.

When a client wishes to send traffic, it first sends a control frame called a "Request to Send" or "RTS". This includes a period of time for which the client station wishes to have exclusive access to the channel. When the client receives a "Clear to Send" (CTS) message, it can transmit for that duration of time. When any other client station sees a CTS destined for another party, it will go into a wait state and does not transmit. However, if those

other stations also saw the original RTS, they can delay for the specified interval. If the original RTS was not seen, the other clients assume the original sender was of RF range and use an internally calculated random delay instead. This is absolutely necessary because client stations are frequently out of range without client stations in the same BSS, but it also opens client stations to a denial of service.

If an attacker forges CTS packets and sends them into the wireless environment, all clients who see the forged packets will assume they were not in range of the original RTS—an RTS that never existed in the first place. But the unknowing clients will enter the wait state until the CTS expires before attempting to communicate with a new RTS of their own. However, if the attacker floods the environment with these forged CTS packets, no amount of waiting would lead to the CTS expiring—the attacker would simply command the airwaves indefinitely. This CTS Attack is explained in considerable depth in the paper “Smart Attacks based on Control Packets Vulnerabilities with IEEE 802.11 MAC.”^[2]

References:

- [1] <http://for572.com/2qnwz>
- [2] <http://for572.com/ey46p>

Spotting 802.11 DoS Attacks

- Users report poor performance
 - No one can connect/send data
 - Unusually low throughput
- Technological countermeasures
 - Frequent channel changes
- Wireless IDS identifies anomalous traffic patterns
 - Alerts often include access point(s) where problem exists
 - WIDS generally requires RFMON = legal involvement


A DoS attack is usually easy to spot when it's successful. Generally, nobody has access to the network and user complaints are not too far behind. This often leads to a geographic area that is affected more than others, helping to localize the investigation.

However, most wireless platforms are built to mitigate such attacks by automatically changing channels when the selected one loses its quality or reliability. This change in channel is often logged, so a rapid spike in channel changes or if this occurs more frequently than expected, it may signify that the architecture is under attack but handling the situation so far. A proactive monitoring team should see these log entries and consider the possibility of a DoS attack.

From a protocol perspective, a quality WIDS platform will detect malformed traffic or a high volume of atypical traffic. This would likely include the access point(s) that are affected by the anomalies, which would also be a helpful detail for investigators, because it would geolocate their response actions.

Evil Twin Attack

- Attacker hosts an access point with expected SSID
- Client stations attach to SSID regardless of BSSID
 - Highest priority on list at highest power usually wins
- Look for unexpected BSSIDs on known SSID
 - Also look for unusual spike in SSID count



In an Evil Twin attack, the victim is tricked into connecting to an SSID he believes to be safe and trusted. However, the victim does not realize that the BSSID providing the familiar network name is under the control of an attacker.

The attacker has used the expected SSID (and likely cloned any captive portal in use). To maximize the likelihood that a victim client station will associate to the Evil Twin and not the real BSSID, the attacker uses a card with a very strong signal. The attack can be conducted in or close to the organization's offices if powerful wireless cards and amplifying antennae are used. (For example, the Alfa adapter discussed earlier can be obtained in 2000mW versions, whereas a typical AP provides only 600-800mW. The legality of these amplified APs is dubious and varies greatly by country—you have been warned!)

When the victim connects to the attacker's Evil Twin BSSID, he uses the connection as normal—even providing credentials if requested. With those credentials, the attacker will be able to impersonate the user, intercept and manipulate traffic, and more. At this point, the attacker simply becomes the victim's upstream infrastructure.

Preventing Evil Twin access points is close to impossible, because anyone can set up a device that acts like an access point, and the management frames are unencrypted and unauthenticated. Instead, incident responders should seek to keep a reliable inventory of trusted BSSIDs that provide service on SSIDs for which the organization is responsible. Whenever a new BSSID advertises that network, there is a chance a rogue access point has been activated. Another good metric to keep updated is the list of SSIDs typical for a given location. If a large number of new network names suddenly appears, this may be a special category of super-efficient Evil Twin attack, where a device dynamically impersonates dozens of SSIDs on demand.

Evil Twin: Social Engineering Toolkit, Pineapple

- Social-Engineer Toolkit
 - Automates Evil Twin
- Wi-Fi Pineapple Router
 - Spoofs SSIDs at scale

```
attacker@attacker: ~
File Edit View Search Terminal Help

Join us on irc.freenode.net in channel #setoolkit

The Social-Engineer Toolkit is a product of TrustedSec.

Visit: https://www.trustedsec.com

Select from the menu:

1) Spear-Phishing Attack Vectors
2) Website Attack Vectors
3) Infectious Media Generator
4) Create a Payload and Listener
5) Mass Mailer Attack
6) Arduino-Based Attack Vector
7) SMS Spoofing Attack Vector
8) Wireless Access Point Attack Vector
9) QRCode Generator Attack Vector
10) Powershell Attack Vectors
11) Third Party Modules

99) Return back to the main menu.

set>
```



The Social-Engineer Toolkit^[1] is a software tool that allows a malicious user or attacker to quickly set up a soft access point—a turnkey Evil Twin, if using an existing SSID. Any victim systems that attach to the SET-based access point instead of the real infrastructure would then be subject to any traffic blocking, manipulation, or interception that the attacker wants to perform. As a robust toolkit intended for penetration testers, it provides added functionality such as automated credential theft, browser exploits, and more.

A hardware approach that can perform a similar function at scale is the Wi-Fi Pineapple Router^[2] from Hak5. This hardware device listens for Wi-Fi probe requests, and then dynamically creates wireless networks with matching SSIDs. Victim systems then attach to these “trusted” network names, allowing the attacker to perform the same Evil Twin functions against a wider base of victims.

Whether an ordinary computer running specialized software or dedicated hardware designed to support penetration testers, the miniaturization of hardware and growing availability of platforms that can provide a malicious access point means detecting rogue actors in the wireless environment will become an increasingly challenging aspect for incident responders for quite a long time.

References:

- [1] <http://for572.com/n83-4>
- [2] <http://for572.com/sxg5u>

Staged Approach to Wi-Fi Hacking

- Attackers tend to use the following modes:
 - Monitor (RFMON) to detect APs and Stations
 - Potential injection of packets
 - Master Mode to pretend to be an AP
 - Managed Mode when they have stolen credentials
- Consider if wireless is the path or the target
 - If just the path, nothing changes to investigate

Although there is no single attack pattern in any environment—wired or wireless—the unique attack surface that a wireless network provides often means attackers tend to follow a known pattern of behavior.

First, the attacker will use monitor mode to gain intelligence about the victim's environment—SSIDs, protection mechanisms, volume of client activity, etc. This is undetectable and may occur over days or weeks, ensuring the attacker has enough data collected to continue the attack.

If the attacker sees the victim's wireless network is protected by WPA or WPA2, they can conduct offline attacks against the data collected previously, or prod network members to re-associate so they divulge the 4-way handshake, giving the attacker sufficient source material for an offline attack. In either case, the attacker seeks to derive the PSK used to associate with the network.

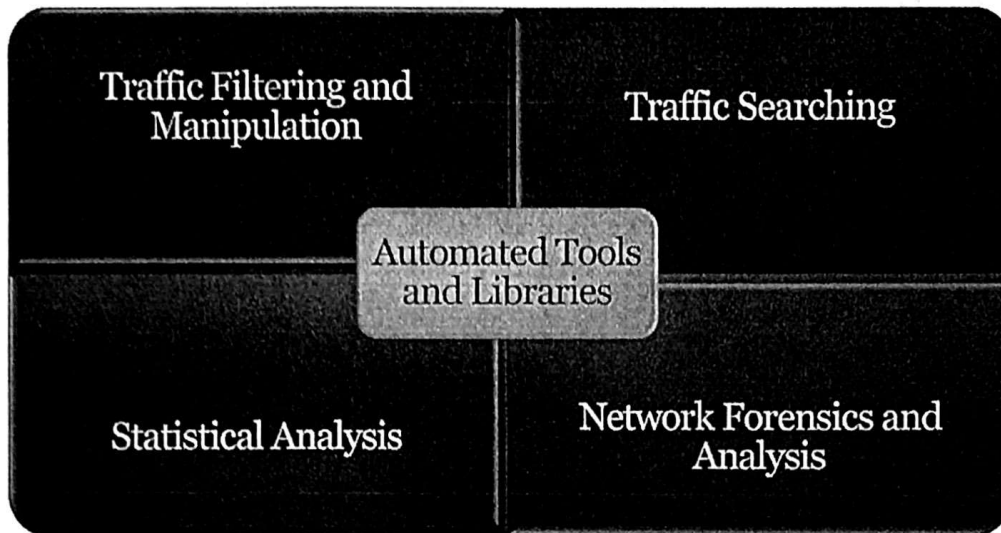
Depending on their goals, this may be when the attacker simply uses the PSK to access the victim's network and start conducting an attack from the inside. However, if they wish to conduct further reconnaissance or damage to the victim's users, they may instead take the Evil Twin route, setting up a rogue access point and intercepting the traffic of the unwitting users that attach to the Evil Twin instead of the legitimate infrastructure. Of course at this point, the attacker has virtually limitless options to wreak havoc at their disposal.

The means to investigate this sequence of events center largely around log data and historic norms of the victim organization's traffic patterns. However, this notional sequence of events centers around the idea of a wireless attack, which should not be confused with an attack that simply uses wireless as its delivery mechanism. In the latter case, our methods do not significantly change from those used in any other network investigation. It is only when the attack leverages inherent weaknesses in wireless protocols that our investigative methodologies significantly change.

Automated Tools and Libraries

This page intentionally left blank.

Overview



As with most technical tasks, there are a number of tools that an individual may use to accomplish a specific task. Each tool provides the user a specific set of capabilities and leaves the user to determine whether it may or may not be the best tool for the job. Network forensics and analysis are no different. Many of the early open source tools were developed to provide solutions to specific problems (such as `tcpdump` and `tcpflow`), while more recent tools take ideas from the early applications and extend them (`flowgrep` and Wireshark). For our purposes, we will use the following four broad categories to describe the tools' functionalities:

- Traffic Filtering and Manipulation These tools provide users the basic ability to slice and dice network traffic.
- Traffic Searching These tools provide users the ability to quickly search through traffic based on unique strings or byte sequences to determine whether the traffic of interest contains what the user is looking for.
- Statistical Analysis These tools provide users the ability to obtain information to detect anomalies or provide some ability to quantify metrics associated with network traffic.
- Network Forensics and Analysis Tools These tools typically provide users the ability to extract additional information from the application layer of network protocols to include images, files, and other items beyond just metadata.

libpcap/winpcap

- Provides an API for capture and display of network traffic

libnids

- Performs IP defragmentation, TCP stream reassembly, and TCP port scan detection

There are several libraries that abstract the lower level mechanics of manipulating raw traffic collected from a network interface to perform some action. Two of the most popular libraries used for development of applications include `libpcap` and `libnids`.

The `libpcap` libraries were originally developed by the Network Research Group at Lawrence Berkeley Laboratory (LBL). In the late 1980s, three researchers from LBL developed the `tcpdump` application to allow them to capture and display network traffic from a system connected to the LBL network. The `libpcap` library came about through stripping the packet capture code from the `tcpdump` application into a consolidated library and API that could then be used with other applications. Today, many open source and commercial applications use the `libpcap` library as a foundation to handle the traffic capture and filtering portions of the application.

The `libnids` library is a popular library that was developed in response to a paper published in January 1998 by Thomas H. Ptacek and Timothy N. Newsham called "Eluding Network Intrusion Detection."^[1] As part of the paper, Ptacek and Newsham explained a number of ways to bypass network intrusion detection systems (NIDS). `libnids` was developed in response to the paper as a way to mitigate ways to bypass NIDS devices, which included implementing the ability to perform IP defragmentation, TCP stream reassembly, and TCP port scan detection.

Both `libpcap` and `libnids` have a number of wrappers for use with various programming languages, including C/C++, Python, Perl, and Java, and they are available on both UNIX- and Windows-based systems.

References:

[1] <http://for572.com/5uf9t>

- Captures or processes traffic and outputs individual streams to a given directory

```
$ tcpflow -r ftp-example.pcap
$ ls
149.020.020.135.00021-192.168.075.029.37028
149.020.020.135.30321-192.168.075.029.49897
...
149.020.020.135.30893-192.168.075.029.38410
149.020.020.135.30904-192.168.075.029.60692
192.168.075.029.37028-149.020.020.135.00021
ftp-example.pcap
$ file 149.020.020.135.30893-192.168.075.029.38410
149.020.020.135.30893-192.168.075.029.38410: RPM v3.0 bin i386/x86_64
$ md5sum 149.020.020.135.30893-192.168.075.029.38410
70710806ef778f84a709b1a2318fd14b 149.020.020.135.30893-192.168.075.029.38410
$ rpm -K 149.020.020.135.30893-192.168.075.029.38410
149.020.020.135.30893-192.168.075.029.38410: rsa sha1 (md5) pgp md5 OK
```

Similar to `tcpdump`, `tcpflow` is a command line, `libpcap`-based application that captures data transmitted across network interfaces or processes data from network captures. It can filter network traffic using the Berkeley Packet Filter (BPF) syntax. However, unlike `tcpdump`, `tcpflow` is knowledgeable of the sequence numbers associated with TCP connections, allowing it to perform reconstruction of the data streams and stores each flow in its own file for later analysis. The output files contain the continuous data portion of the source, not new `pcap` files. `tcpflow` is commonly used for understanding network packet flows and performing network forensics, for example, reassembling the contents of HTTP sessions. Using `tcpflow`, a user can reconstruct web pages downloaded via HTTP or even extract malware or files transferred via HTTP or FTP. `tcpflow` is also commonly used during the analysis of undocumented network protocols aiding in reverse engineering of the protocols.

The preceding example shows `tcpflow` processing the `ftp-example.pcap` file from the `/sample_pcaps/` directory on your FOR572 USB. The `149.020.020.135.30893-192.168.075.029.38410` stream corresponds to a particular RPM file, which the subsequent commands validate.

References:

<http://for572.com/odbi7>

<http://for572.com/wujhc>

- Low-level Python framework to parse, decode, and create (and forge) packets
 - Build your own sniffer and/or pcap reader
 - Custom Pythonic parsing engines
- Intended purpose is to create and inject packets to a network segment
 - “Attack” methods not usually of forensic interest
 - Can be used to examine undocumented protocols
 - Interact with client/server software in sandbox

The `scapy`^[1] framework is a Python tool commonly used by penetration testers and attackers alike. Although this is not a primary focus for the forensically focused investigator, it can still be a good tool to understand. If faced with an undocumented protocol, for example, a researcher could use `scapy` to build a custom parser for command and control traffic, use it to extract certain parameters from an HTTP stream, reconstruct files, etc. A reverse engineer could use it to interface with malware executing in a controlled sandbox environment.

`scapy` can also be used to sniff live packets from a network interface as well as read packet data from a pcap file. For example, a DNS sniffer can be created with just a handful of lines of Python code.^[2] This type of functionality can be extended for new and existing protocols, and could be a great benefit for tactical network monitoring during an incident response.

References:

[1] <http://for572.com/6bvmx>

[2] <http://for572.com/ugwon>

- Python-based protocol parsing framework
 - Open source project from U.S. Army Research Lab
- TCP stream reassembly (IPv4 and IPv6)
- GeoIP integration for geo and ASN queries
- Plugins to extend parsing as needed
 - Show relevant data on screen, perform file extraction, etc.

Another promising entrant into the field of traffic parsing utilities is Dshell,^[1] a Python-based protocol parsing framework. It is fully free and open source, and anyone with a working knowledge of Python can write general or incident-specific parsers and analytic helper modules to aid in their work.

In the examples on the next page, you can see the results of running Dshell against the “ftp-example.pcap” file on your FOR572 USB drive, including a general NetFlow-like view showing the single command channel connection and multiple data channels, as well as a full object extractor, reconstructing the files retrieved via the data channels.

References:

[1] <http://for572.com/cmdlh>

```
$ cd /mnt/hgfs/sample_pcaps/
$ dshell
```

```
Dshell> decode -d netflow ftp-example.pcap
2013-11-22 16:38:30.883332 192.168.75.29 -> 149.20.20.135 (---> US) TCP 60090 30351 0 2 2896 0.1906s
2013-11-22 16:38:33.889745 192.168.75.29 -> 149.20.20.135 (---> US) TCP 49897 30321 0 0 0 0.2079s
2013-11-22 16:38:38.452075 192.168.75.29 -> 149.20.20.135 (---> US) TCP 37110 30487 0 0 0 0.1772s
2013-11-22 16:38:42.970871 192.168.75.29 -> 149.20.20.135 (---> US) TCP 37583 30669 0 0 0 0.1780s
2013-11-22 16:38:49.561496 192.168.75.29 -> 149.20.20.135 (---> US) TCP 49943 30760 0 424 0 612916 2.2349s
2013-11-22 16:39:00.453416 192.168.75.29 -> 149.20.20.135 (---> US) TCP 38410 30893 0 705 0 1019540 1.8620s
2013-11-22 16:39:11.299795 192.168.75.29 -> 149.20.20.135 (---> US) TCP 60692 30904 0 654 0 945952 1.9983s
2013-11-22 16:39:18.431439 192.168.75.29 -> 149.20.20.135 (---> US) TCP 60621 30444 0 1912 0 2767424 3.3674s
2013-11-22 16:39:37.628648 192.168.75.29 -> 149.20.20.135 (---> US) TCP 51851 30472 0 19260 0 27888036 12.5452s
2013-11-22 16:38:14.266541 192.168.75.29 -> 149.20.20.135 (---> US) TCP 37028 21 28 43 359 3660 99.3606s
```

```
Dshell> decode -d ftp --ftp_dump ftp-example.pcap
```

```
WARNING:ftp:Using output directory: /home/sansforensics/ftpout
```

```
ftp 2013-11-22 16:38:14 192.168.75.29:37028 -> 149.20.20.135:21 ** User: ftp, Pass: for572@lewestech.com, RETR File: centos/LIST
(2896 bytes written to RETR_centos_LIST) **
ftp 2013-11-22 16:38:14 192.168.75.29:37028 -> 149.20.20.135:21 ** User: ftp, Pass: for572@lewestech.com, RETR File: centos/6.4/LIST
**
ftp 2013-11-22 16:38:14 192.168.75.29:37028 -> 149.20.20.135:21 ** User: ftp, Pass: for572@lewestech.com, RETR File: centos/6.4/os/LIST **
ftp 2013-11-22 16:38:14 192.168.75.29:37028 -> 149.20.20.135:21 ** User: ftp, Pass: for572@lewestech.com, RETR File: centos/6.4/os/x86_64/LIST **
ftp 2013-11-22 16:38:14 192.168.75.29:37028 -> 149.20.20.135:21 ** User: ftp, Pass: for572@lewestech.com, RETR File: centos/6.4/os/x86_64/Packages/LIST (611468 bytes written to RETR_centos_6.4_os_x86_64_Packages_LIST) **
ftp 2013-11-22 16:38:14 192.168.75.29:37028 -> 149.20.20.135:21 ** User: ftp, Pass: for572@lewestech.com, RETR File: centos/6.4/os/x86_64/Packages/yum-3.2.29-40.el6.centos.noarch.rpm (1019540 bytes written to RETR_centos_6.4_os_x86_64_Packages_yum-3.2.29-40.el6.centos.noarch.rpm) **
ftp 2013-11-22 16:38:14 192.168.75.29:37028 -> 149.20.20.135:21 ** User: ftp, Pass: for572@lewestech.com, RETR File: centos/6.4/os/x86_64/Packages/xchat-2.8.8-1.el6.x86_64.rpm (945952 bytes written to RETR_centos_6.4_os_x86_64_Packages_xchat-2.8.8-1.el6.x86_64.rpm) **
ftp 2013-11-22 16:38:14 192.168.75.29:37028 -> 149.20.20.135:21 ** User: ftp, Pass: for572@lewestech.com, RETR File: centos/6.4/os/x86_64/Packages/zenity-2.28.0-1.el6.x86_64.rpm (2767424 bytes written to RETR_centos_6.4_os_x86_64_Packages_zenity-2.28.0-1.el6.x86_64.rpm) **
ftp 2013-11-22 16:38:14 192.168.75.29:37028 -> 149.20.20.135:21 ** User: ftp, Pass: for572@lewestech.com, RETR File: centos/6.4/os/x86_64/Packages/scenery-backgrounds-6.0.0-1.el6.noarch.rpm (27888036 bytes written to RETR_centos_6.4_os_x86_64_Packages_scenery-backgrounds-6.0.0-1.el6.noarch.rpm) **
```

- Extracts portions of packet traces based on start and end times
- Merges several packet captures together into a single pcap file

```
$ cd /mnt/hgfs/sample_pcaps/  
$ tcpslice nitroba.pcap -r  
nitroba.pcap Tue Jul 22 01:51:07 2008 Tue Jul 22 06:13:47 2008  
$ tcpslice 2008y7m22d02h00m00s 2008y7m22d03h00m00s nitroba.pcap -w nitroba_0100-0200.pcap  
$ tcpslice -r nitroba_0100-0200.pcap  
nitroba_0100-0200.pcap Tue Jul 22 02:00:11 2008 Tue Jul 22 02:30:21 2008
```

`tcpslice` is a command-line, `libpcap`-based application that extracts portions of pcap files based on the user's date and time specifications. In addition to extracting files, `tcpslice` can also be used to merge several files together. `tcpslice` is most commonly used to extract a time period of interest from a larger pcap file, making lower-level analysis less cumbersome.

References:

<http://for572.com/emrin>

- Search traffic for strings or regular expressions
- Can write matching packets to disk
- Suitable for plaintext protocols (HTTP, SMTP, FTP)
- Works only within each individual packet!
- For IPv6, use pngrep.pl

```
$ ngrep -q -I ftp-example.pcap 'RETR'
input: ftp-example.pcap
match: RETR

T 192.168.75.29:37028 -> 149.20.20.135:21 [AP] RETR yum-3.2.29-40.el6.centos.noarch.rpm..
T 192.168.75.29:37028 -> 149.20.20.135:21 [AP] RETR xchat-2.8.8-1.el6.x86_64.rpm..
T 192.168.75.29:37028 -> 149.20.20.135:21 [AP] RETR zenity-2.28.0-1.el6.x86_64.rpm..
T 192.168.75.29:37028 -> 149.20.20.135:21 [AP] RETR scenery-backgrounds-6.0.0-
1.el6.noarch.rpm..
```

Ngrep^[1] (or network grep) is a libpcap-based application that allows users to specify extended regular or hexadecimal expressions to match against data payloads of packets. Similar to most other libpcap-based tools, it can filter network traffic using the Berkeley Packet Filter (BPF) syntax. ngrep is commonly used to quickly obtain information from plaintext protocols such as HTTP, SMTP, FTP, etc. Like tcpdump, ngrep can also read from a live network interface or capture file and can also capture data to a pcap file. However, because ngrep is packet-based, if the regular expression's target pattern crosses packet boundaries, ngrep will not match the pattern.

One notable shortfall of ngrep is that the primary codebase has not yet been updated to address IPv6 traffic. For this, check out Jim Clausing's pngrep.pl,^[2] a pure-perl version that overcomes the IPv4 limitations.

References:

[1] <http://for572.com/1jey->

[2] <http://for572.com/c-rj8>

- Combines functions of three different utilities
 - tcpflow – Traffic capture and collection
 - ngrep – Searching
 - tcpkill – Killing TCP connections
- Provides minimal functionality needed for an IDS/IPS device

At its core, flowgrep is the combination of ideas from several programs, including tcpflow, ngrep, and tcpkill. Essentially, this makes flowgrep a very basic intrusion detection/intrusion prevention system using the libnids library. Because flowgrep uses libnids, it mitigates several shortcomings of previously mentioned applications. flowgrep can perform IP and UDP defragmentation, reassemble TCP stream, and search for regular expressions across packet boundaries. Similar to tcpflow, flowgrep can log network streams that match the specified regular expression to a pcap file for later use. flowgrep also provides users the ability to kill network connections that match the specified regular expression.

References:

<http://for572.com/qbhn1>


- Provides 15 different observed and calculated traffic statistics
 - Observed: Source/dest IP addresses and ports
 - Calculated: Packet count, average packet size, standard deviation, bandwidth
- Equivalent of vmstat for network traffic

tcpstat is a libpcap-based application that provides network interface statistics. tcpstat acquires this information either by monitoring a specific interface or by reading an existing pcap file. It provides more than 15 different types of statistics including the number of packets passed through the interface, the average size of each packet, the standard deviation of the packet size, and the bandwidth in bits per second. tcpstat was designed to provide a network-focused cousin to the popular vmstat program that monitors host-based performance such as CPU, memory, and I/O performance.

A similar program, tcpdstat, was developed by Sony of Japan and also provides several output options including the number of packets, average rate and standard deviation, the number of unique source and destination address pairs, and information based on the breakdown of protocols. The output can easily be converted to HTTP format for graphical review.

References:

<http://for572.com/esydj>
<http://for572.com/p6zwg>
<http://for572.com/gql7p>



The graph provided is a graph depicting the output of tcpstat data when plotted by gnuplot. Contained within the graph are examples of network traffic breakdown by protocol. Through the combination of tcpstat with a gnuplot script, an analyst can produce interesting graphs to help identify spikes in activity that are indicative of potential data theft or file transfers.

```
$ cd /mnt/hgfs/sample_pcaps/
$ tcpstat -r ftp-example.pcap -o "%R\t%A\n" 1 > arp.data
$ tcpstat -r ftp-example.pcap -o "%R\t%C\n" 1 > icmp.data
$ tcpstat -r ftp-example.pcap -o "%R\t%T\n" 1 > tcp.data
$ tcpstat -r ftp-example.pcap -o "%R\t%U\n" 1 > udp.data
```

Where the `-o` option indicates the output format for traffic over a given time interval (in this case, every 60 seconds). The output is then redirected into a specified data file. gnuplot uses a scripting language to generate the graphs. The following gnuplot script is on your USB as `"/sample_pcaps/ftp-example.gnuplot"`:

```
set output "proto_breakdown_over_time.png";
set term png size 1024, 768;
set grid;
set yrange [ -10 : 3000 ];
set title "Protocol Breakdown";
set xlabel "seconds";
set ylabel "packets/s";
plot "arp.data" using 1:($2) lw 3 smooth csplines title "ARP", \
     "icmp.data" using 1:($2) lw 3 smooth csplines title "ICMP", \
     "tcp.data" using 1:($2) lw 3 smooth csplines title "TCP", \
     "udp.data" using 1:($2) lw 3 smooth csplines title "UDP";
```

The following commands then generate and view the actual graph image file:

```
$ cd /mnt/hgfs/sample-pcaps/  
$ gnuplot ftp-example.gnuplot  
$ display proto_breakdown_over_time.png
```

- Provides network-based stats associated with observed traffic
 - Top talkers (most active IPs, ports, and protocols)
 - Traffic amount and byte counts
 - Overall throughput
- HTTP web server provides detailed output, searching, and filtering
 - Includes basic intelligence such as IP address geolocation

Similar to the venerable UNIX “top” command, ntopng provides real-time usage information associated with network usage. ntopng is a libpcap-based application that operates by observing active network interfaces in real time to provide statistical information associated with the observed traffic. ntopng provides an interactive mode where information is displayed to the command-line interface and a web-based mode, which provides the statistical information via a web browser. It can also receive NetFlow data and perform targeted full-packet capture on a per-host basis.

However, it does not consume existing data, so it is limited to real-time usage and pre-incident deployment. In a tactical collection platform, however, ntopng provides solid performance in a freemium licensing model.

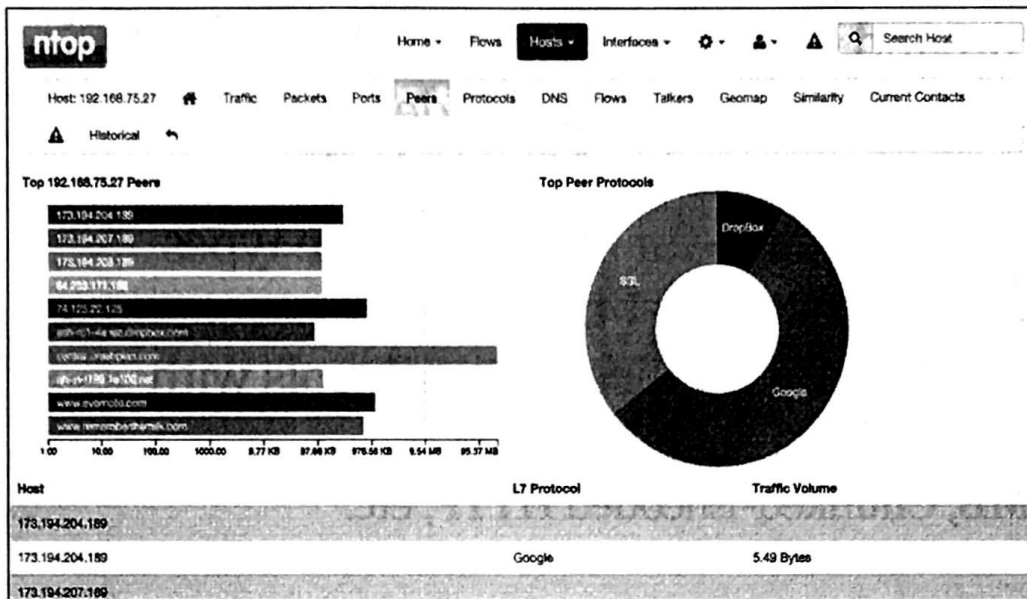
References:

<http://for572.com/iysd4>

The screenshot shows the ntopng web interface. At the top, there are navigation tabs: Home, Flows (selected), Hosts, and Interfaces. A search bar is present on the right. The main content area is titled 'Active Flows' and displays a table of network flows. The table has columns for 'Info', 'Application', 'L4 Proto', 'Client', 'Server', 'Duration', 'Breakdown', 'Actual Thpt', and 'Total Bytes'. There are 8 rows of data, each representing a different network flow with various protocols like SSL, SIP, ICMP, HTTP, and SSDP.

Info	Application	L4 Proto	Client	Server	Duration	Breakdown	Actual Thpt	Total Bytes
Info	SSL	TCP	central.craahplan.co...:443	Phils-MBP-2.Identity...:57770	57 min, 25 sec	Client	0 bps	210.62 MB
Info	SIP	UDP	192.168.75.129:10000	69.68.252.114:10000	4 days, 21 h, 36 min, 56 sec	Client Server	0 bps	27.55 MB
Info	ICMP	ICMP	192.168.75.1	192.168.75.130	1 day, 3 h, 30 min, 11 sec	Client Server	783.08 bps	9.08 MB
Info	HTTP	TCP	Marisa-MBP-2.Identity...:55845	www.trbimg.com:80	3 min, 5 sec	Server	0 bps	8.14 MB
Info	HTTP	TCP	Marisa-MBP-2.Identity...:55820	www.trbimg.com:80	2 min, 57 sec	Server	0 bps	5.6 MB
Info	Facebook	TCP	Marisa-MBP-2.Identity...:55997	ix-fbodn-shv-01-lad3...:443	45 sec	Server	0 bps	4.27 MB
Info	SSDP	UDP	192.168.75.48:1900	239.258.255.250:1900	1 day, 6 min, 49 sec	Client	0 bps	2.98 MB
Info	Unknown	UDP	192.168.75.77:54321	213.33.117.249:49049	2 days, 11 h, 47 min, 56 sec	Client Server	0 bps	2.26 MB

From a user-interface perspective, ntopng provides a graphical interface including metadata associated with each flow such as the client/server ports as determined by the TCP three-way handshake and the amount of traffic transferred between the client server before the two systems over a given time period. You can also dive into individual flows to obtain more information or filter by client/server or dive deeper into traffic associated with a specific client or server address.



This screenshot indicating the type of information a user can obtain from ntopng when deep-diving into a specific host; however, similar information can be obtained from other metadata, such as ports. The upper left portion of the screen contains information about what other peers (IPs or hostnames) the given IP address connected to and the amount of data transferred between the two peers. The upper right graph contains a breakdown in the protocols used by the IP address. The bottom chart contains information on the individual flows including the peer the original IP address queried connected to, the associated protocol, and the given amount of traffic. The graphic visualization aids users in making comparisons between different types of traffic that may help users also find anomalies.

- Command line tool to extract files from traffic
- Extraction based on file header and footer information contained in config file
- Easily extensible for additional file types
- Works across packet boundaries
- Unsuitable for application protocols with per-packet overhead
 - SMB, Chunked-Encoded HTTP, etc.

tcpextract^[1] is a command-line, libpcap-based application that extracts files from a live network stream or pcap file based on the “magic” file type headers and footers. If you’re familiar with file carving from static media using tools like foremost^[2] or scalpel,^[3] tcpextract does this against network data. It performs TCP stream reassembly and uses a configuration file to identify headers and footers of various file types. It easily extracts files such as images and office documents, and it conveniently handles data spread out across packet boundaries. Upon identifying a file, tcpextract writes the file to disk. Note, however, that the filenames are not recovered. Instead, tcpextract uses the frame number as the name for the carved file with the file extension defined in the configuration file.

Although tcpextract is a very handy tool to have, it is not aware of higher-layer protocols, such as SMB. Therefore, it is limited to protocols where file content is passed contiguously in a stream, without any Layer 5-7 headers in each packet.

References:

[1] <http://for572.com/uqlok>

[2] <http://for572.com/5v4a0>

[3] <http://for572.com/yjeaw>

```
$ cd /mnt/hgfs/sample_pcaps/
$ cat rpm-tcpextract.conf
# RPM files
# http://www.rpm.org/max-rpm/sl-rpm-file-format-rpm-file-format.html
rpm(400000000, \xed\xab\xee\xdb);

$ tcpextract -c rpm-tcpextract.conf -f ftp-example.pcap -o /tmp/
Found file of type "rpm" in session [149.20.20.135:44408 -> 192.168.75.29:2710], exporting
to /tmp/00000000.rpm
Found file of type "rpm" in session [149.20.20.135:47224 -> 192.168.75.29:5357], exporting
to /tmp/00000001.rpm
Found file of type "rpm" in session [149.20.20.135:60534 -> 192.168.75.29:52716],
exporting to /tmp/00000002.rpm
Found file of type "rpm" in session [149.20.20.135:2167 -> 192.168.75.29:35786], exporting
to /tmp/00000003.rpm

$ file /tmp/0000000*.rpm
/tmp/00000000.rpm: RPM v3.0 bin i386/x86_64
/tmp/00000001.rpm: RPM v3.0 bin i386/x86_64
/tmp/00000002.rpm: RPM v3.0 bin i386/x86_64
/tmp/00000003.rpm: RPM v3.0 bin i386/x86_64
```

The “rpm-tcpextract.conf” file is provided on your course USB drive as “/sample_pcaps/rpm-tcpextract.conf”. It contains a single carving signature, which uses the magic values for the RPM file as defined by the file specification.^[1] Running tcpextract against the “ftp-example.pcap” file with this configuration results in four extracted files, which can be validated based on their contents. FTP is an ideal protocol to use to demonstrate this functionality because the file contents are transmitted in a contiguous byte stream. RPM files are also ideal to use while testing file reconstruction tools because their content can be cryptographically verified against what is expected.

References:

[1] <http://for572.com/20xat>



Lab 10

Using Command-Line Tools for Analysis

This page intentionally left blank.

Lab 10 Objectives: Using Command-Line Tools for Analysis

- Become familiar with command line tools to perform common network forensic tasks
- Reconstruct transferred files at the command line
- Identify FTP transfers using `ngrep`
- Recover files using both `tcpflow` and `tcpextract` based on network and file metadata and file content signatures

This page intentionally left blank.

Lab 10 Takeaways: Using Command-Line Tools for Analysis

- Graphical tools often make analysis easier; they are sometimes infeasible
 - Not often suitable for remote analysis
 - Don't lend themselves to repeatable, scalable analysis
- Command-line tools tend to be built for a single purpose and are not all-in-one solutions
 - Analysts must often identify the best tool for each job
 - Combining such single-purpose tools can provide a very robust and efficient toolchain for the broader task
 - Scripting with these tools creates repeatable processes

This page intentionally left blank.

Full-Packet Hunting with Moloch

This page intentionally left blank.

Moloch at a Glance

- Network forensics and analysis tool
 - Full packet capture
 - Protocol parsing/indexing
 - pcap reduction/retrieval
- Open-source project started at AOL
- Not an IDS



Of course command-line tools can provide scalability in terms of processing large amounts of data. However, maintaining the custom scripts, linking tools and processes together, and reviewing the data can all be a bit cumbersome. If we want an easier way to search and review data, a tool like Wireshark may quickly come to mind, but it has scaling problems of its own, especially based on the volume of data being processed. We have discussed a number of commercial solutions that can perform full content capture and indexing of network traffic, but these are often quite expensive. Fortunately, there is an excellent, scalable, free, open-source tool that rivals some commercial platforms in terms of features and usability.

Moloch is a large-scale packet-capturing, indexing, and database system that provides a simple web interface for browsing, searching, and exporting pcap files. Originally started by employees at AOL, Moloch is not meant to replace existing intrusion detection systems, but instead, it works alongside them to store and index all the network traffic in a standard pcap format while providing fast access.

References:

<http://for572.com/ftea9>

<http://for572.com/8j15n>

Platform Components

- **Capture**
 - Captures and parses live or existing network traffic Creates Session Profile Information (SPI-Data)
- **Elasticsearch**
 - Stores, indexes, and searches SPI-data generated by the capture component
- **Viewer**
 - Web-based GUI interface that provides browsing and querying against indexed data and pcap file retrieval

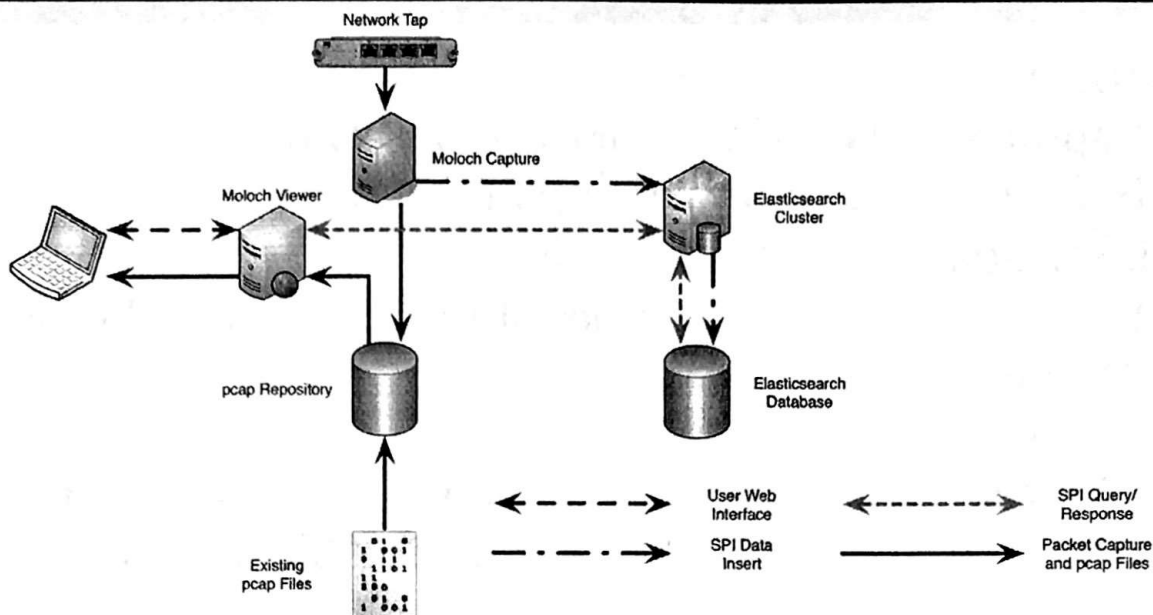
Moloch consists of three main components, which work in concert with each other to provide the overall Moloch feature set.

The capture mechanism can passively listen on a network interface or read existing packet capture data. It parses the traffic, tracks Session Profile Information (SPI-Data), and writes the raw packets to disk for later use and extraction.

Moloch uses the Elasticsearch search and analytic engine to store SPI-Data records and expose them to the user through the viewer component. Because Elasticsearch is such a widely used and well-documented part of the overall platform, its use also allows third-party access and integration to the broader scope of Elasticsearch analytic capabilities.

The viewer component is Moloch's integrated web-based GUI that provides user-level query capabilities, session-level pcap data reduction and extraction, as well as an API that allows for programmatic and repeatable access to the underlying data stored in Elasticsearch.

Components and Architecture



SANS DFIR

FOR572.4 | Advanced Network Forensics and Analysis

114

Architecturally, each component of Moloch can scale independently from the others. For smaller-scale live deployments and small-to-moderately sized pcap data loads, all components can run simultaneously on a single system. This is the model used in the FOR572 Moloch VM that you will use in the lab.

However, for larger-scale usage, Moloch can operate under a “fleet” model, with multiple capture platforms deployed throughout the environment, each storing full content and indexing network traffic under its own purview.

The resulting SPI-Data from this fleet of capture platforms would then be fed back to the Elasticsearch node or cluster for unified storage. This cluster could range from a single system to hundreds in a cluster if needed, given the horizontal scaling inherent in Elasticsearch.

Because the Moloch viewer process interacts directly with the Elasticsearch, the user can query all available SPI-Data from the entire capture platform fleet in one interface. This allows for more comprehensive and streamlined analysis, to include pcap session extraction for deeper examination with additional tools beyond Moloch itself.

References:

<http://for572.com/8xc62>

Moloch: Limitations

- Technical
 - Capture speed > parsing capabilities
 - Protocol support less robust than Wireshark
 - Handling corrupt input data not sufficiently resilient
- Platform complexity for large environments
 - Capture device tuning needed for higher speed networks
 - Database/Elasticsearch tuning and administration
 - No commercial support for bugs or performance problems

Although Moloch fits a need between expensive, out-of-the-box commercial solutions and more limited command-line tools or deep-dive utilities such as Wireshark, it is not without its shortcomings. When collecting live traffic, it is not difficult to overwhelm the protocol parsers, resulting in lost data unless the architecture is scaled to accommodate high data rates. From the analyst's perspective, Moloch protocol support is notably more sparse than that of Wireshark or many common commercial platforms. Obscure protocols are not natively supported, and adding parsing support for them requires development experience.

In terms of operations, the complexity a Moloch deployment increases significantly in larger and more diverse environments. Using Moloch in such an organization requires knowledge on how to tune both the capture application and the underlying Elasticsearch database. Although there are a number of resources available online regarding best practices for tuning Elasticsearch and troubleshooting common issues, this often involves a significant amount of research. Also, it's necessary to note that there is no commercial support available for Moloch at this time.

However, the value proposition for this platform is certainly noteworthy, and Moloch is a great addition to any network forensicator's toolkit.

Session Viewer

The screenshot displays the Session Viewer interface. At the top, there is a table with columns for 'Time', 'Source IP', 'Destination IP', 'Port', and 'Protocol'. Below the table, a detailed view of a selected session is shown, including protocols, start/stop times, source/destination IP and MAC addresses, and request/response headers. The 'Request Headers' section lists various fields like 'Accept', 'User-Agent', and 'X-Forwarded-For'. The 'Response Headers' section lists 'Age', 'Cache-Control', and 'Content-Type'. Below the headers, the body content is displayed in a hex dump format, with a 'Source' and 'Destination' column.

The main data Moloch indexes to Elasticsearch is called Session Profile Information or SPI-Data. As the Moloch capture process reads and parses network traffic or pcap files, it extracts interesting fields depending on the underlying protocol.

For example, a simple HTTP request would not only include typical network-level fields such as the timestamp of the packet, the source and destination IP addresses, source and destination ports, etc., but it would also contain multiple application-level fields from the HTTP session itself such as the Request URI, User-Agent, and Cookie strings among other things.

Note that in the screenshot above, the HTTP values displayed in the unrolled session metadata include some of these values, but not others (the URI is not displayed, for example). Although Moloch parses out such values, only those that the user is most likely to use as a pivot point are shown. Each value can be clicked to augment the existing query string, as is familiar with numerous other investigative tools and platforms.

Also take note that Moloch is a session-centric platform. The SPI-Data shown above includes values from both the request and the response. In this case, because this session consists of two HTTP request/response pairs in a keep-alive connection, the session's SPI-Data values include parameters from both in the various lists presented.

Below the SPI-Data fields, Moloch presents the content from the session in the standard red/blue coloration to indicate client-to-server and server-to-client communications, respectively. The interface allows the user to show these communications in native, hexdump, or several other formats, as well as to use several different decoding algorithms such as compression and XOR, and to display images in their native form.

Custom port == 80 Search Actions

1000 Showing 1 to 1,000 of 1,137 entries (filtered from 5,619 total entries)

2013/06/08 16:22:19	2013/06/08 16:22:37	10.3.16.11	39104	212.58.246.93	80	118	90,450 / 96,254	for:572-moloch	Column visibility	First	Previous	1	2	Next	Last
tcp	2013/06/08 16:22:19	10.3.16.11	32770	74.125.226.198	80	26	11,020 / 12,752	for:572-moloch	www.bbc.co.uk/news/world-europe-22825640	www.bbc.co.uk/news/world-europe-22825640					
tcp	2013/06/08 16:24:18	10.3.16.11		USA					if.yimg.com/ih/ty/6SFNSP8/default.jpg	if.yimg.com/ih/ty/6SFNSP8/default.jpg					

Download Pcap	Source Raw	Destination Raw	Permalink	Actions
	id: 130608-40bXbz2Sq4pDwJuwj6Fvbiq4 Protocols: http, tcp Start: 2013/06/08 16:22:19 Stop: 2013/06/08 16:24:18 Node: for:572-moloch IP Protocol: tcp Src: Packets: 14 Bytes: 1790 Databytes: 858 Dst: Packets: 12 Bytes: 10962 Databytes: 10162 Ethernet: Src Mac: 08:00:27:55:b6:1a Dst Mac: 08:00:27:5f:27:d5 Src IP/Port: 10.3.16.11:32770 Dst IP/Port: 74.125.226.198:80 (USA) [AS15169 Google Inc.] [ARIN] Payload8: Src: 47455420276692f (GET /) Dst: 485454502612630 (HTTP/1.0) Tags: <input type="checkbox"/> Files: /data/moloch/raw/lab-05.pcap			

HTTP

```

Method: GET
Status code: 200
Hosts: i.yimg.com
User Agents: Mozilla/5.0 (Windows NT 6.1; WOW64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/27.0.1453.110 Safari/537.36
XFFs: 10.3.59.53 (--) [-]
Request Headers: accept, accept-encoding, accept-language, cache-control, connection, host, referer, user-agent, via, x-forwarded-for
Request Versions: 1.0
Resp Headers: age, cache-control, connection, content-length, content-type, date, expires, server, x-content-type-options, x-ssr-protection
Resp Versions: 1.0
Body MD5s: f7421748b64e1ab0deebb463cd8f8f8f, aa302be47bc0bd68cca5cdd6f7293193
libfile content type: image/jpeg
  
```

natural ascii utf8 hex Line Numbers Uncompress Show Images & Files Show Timestamps Other Decodings

Source

```

GET /vi/hdy6SFNSP8/default.jpg HTTP/1.0
Host: i.yimg.com
User-Agent: Mozilla/5.0 (Windows NT 6.1; WOW64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/27.0.1453.110 Safari/537.36
Accept: */*
Referer: http://news.google.com/
Accept-Encoding: gzip, deflate, sdch
Accept-Language: en-US,en;q=0.8
  
```

Destination

Moloch: Filtering

- SPI-Data is viewed through the Moloch viewer
- Implements a simple query language similar to Wireshark to build complex expressions
- Search capabilities depend on field types
 - Search operators: &&, ||, ==, !=, EXISTS!, ()
 - String fields: Tokens, lists, wildcards, regular expressions
 - IP fields: Dotted quads, partial IPs, CIDR blocks, lists
 - Numeric fields: ==, !=, >, <, lists
 - Date fields: Absolute and relative timestamps, lists

Users can query the indexed SPI-Data using the Moloch viewer. Similar to Wireshark, Moloch has implemented a simple query language for building expressions. It supports logical combinations of individual queries using the “&&” and “||” operators for AND and OR, respectively. Equality is tested with the “==” operator and inequality with “!=”. A unique case using the special pseudo-value of “EXISTS!” (including the exclamation point) allows the user to query records where a specific field has been identified and indexed. (For example, “cert.issuer.cn == EXISTS!”.) All of these can be grouped with parentheses to enforce order of operation.

There are several other query types available, depending on the type of data in a given field.

String Fields

String fields can be searched in several ways, and their case may be converted to lowercase in the index for some protocols.

- **Tokens:** It is important to recognize that Moloch often “tokenizes” strings during the indexing process. This means they are broken up based on separators in the source. Characters such as “_”, “/”, “.”, and others are separators, meaning the source string “www.sans.org” will match a query of “www”, “sans”, or “org”, as well as the original string.
- **Lists:** This is a shorthand method for doing multiple OR queries within the same field. For example, the query string “http.uri == [www, moloch]” is equivalent to “http.uri == www || http.uri == moloch”.
- **Wildcards:** A “*” character within a query expression will match any number of arbitrary characters and a “?” character will match any arbitrary single character. For example, the query string “http.uri == “www.sans.*” will match records with the values “www.sans.org” and “www.sans.edu”, whereas “http.uri == “www.sans.?” will match records with the values “www.sans.a” and “www.sans.9” but not “www.sans.org”.
- **Regular expressions:** Regex queries can be run against text strings and must be enclosed with forward slashes (such as “/mx[1-4]\.example\.com/”). Note that Moloch uses the underlying Elasticsearch regex engine, which is not as fully-featured as the PCRE engine.

IP Address Fields

Queries against IP address fields can be performed using a full or partial IP address or CIDR block notation. For fields that include ports, you can use both “1.2.3.4:80” and “10.3.58/24:53” formats. Similar to strings, lists are also accepted in IP address fields.

Numeric Fields

In addition to the standard “==” and “!=” operations, Moloch also allows simple range comparisons using the “>”, “<”, “>=”, and “<=” operators. Lists are also supported.

Date Fields

Dates can be queried with simple equality and inequality and range operators, as well as Splunk-like relative timestamps.^[1]

References:

[1] <http://for572.com/xogaz>

Moloch: External Data Lookups

- WISE: “With Intelligence See Everything”
- Intelligence framework to integrate data feeds
- Allows Moloch to query third parties on data points
 - IP addresses
 - Email addresses
 - Domain names
 - MD5s of extracted files

The Moloch developers understand the need to integrate external intelligence feeds and other data lookups into the tool, because use cases and available data sources vary between organizations and teams. The framework for doing this in Moloch is called “WISE”, which stands for “With Intelligence, See Everything”.

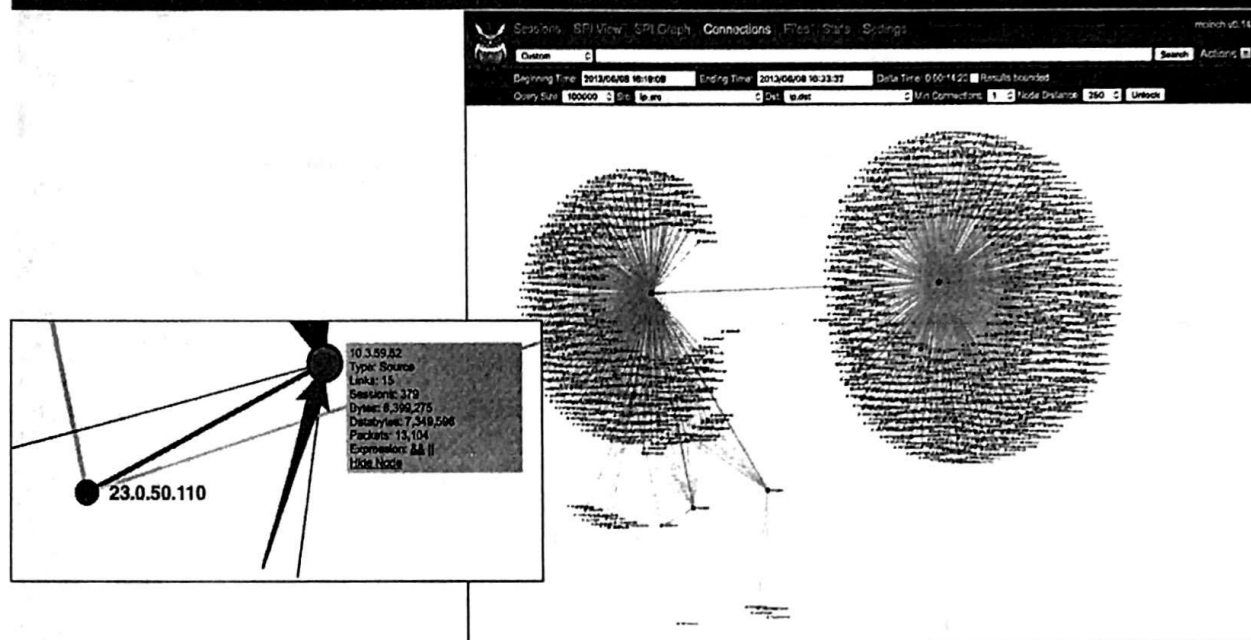
While capturing live data or parsing existing pcap files, Moloch queries the WISE subsystem to obtain additional information about values it observes. This may include lookups against IP addresses, email addresses, domain and hostnames, MD5 hashes of extracted files, and more. WISE lookups can use both free and commercial data sources, and the results are stored in the Elasticsearch database document for the corresponding session.

The latest version of available sources is maintained on the Moloch Wiki.^[1]

References:

[1] <http://for572.com/-021a>

Connection Exploration



SANSDFIR

FOR5724 | Advanced Network Forensics and Analysis 121

Another very useful feature in the Moloch viewer interface is the “Connections” tab. This tab allows the user to visualize values from two series of data points, with any parsed field available for each of those series. In this manner, the user can explore centers of gravity between different items within the source data.

In the example above, the Connections tab reflects the association between source and destination IP addresses, which would be useful for the analyst in seeking patterns of common or outlier communications.

This view is both interactive and analytic. The user can interact with the visualized connections by hovering the mouse over an individual node to get more information about it, including link, session, and packet counts and the total databytes involving the node. The user can drag nodes around the canvas to more comprehensively explore the presented data.

Moloch query strings for each node’s values can be built using the “&” and “|” links on a popup that is displayed by clicking on it.

Remember, every available field can be used for the data series, so a wide range of useful associations can be explored, including but not by any means limited to:

- Source IP address and HTTP hostname (`ip.src, host.http`)
- Source IP address and destination ASN (`ip.src, asn.dns`)
- SMB username and SMB filename (`smb.user, smb.fn`)



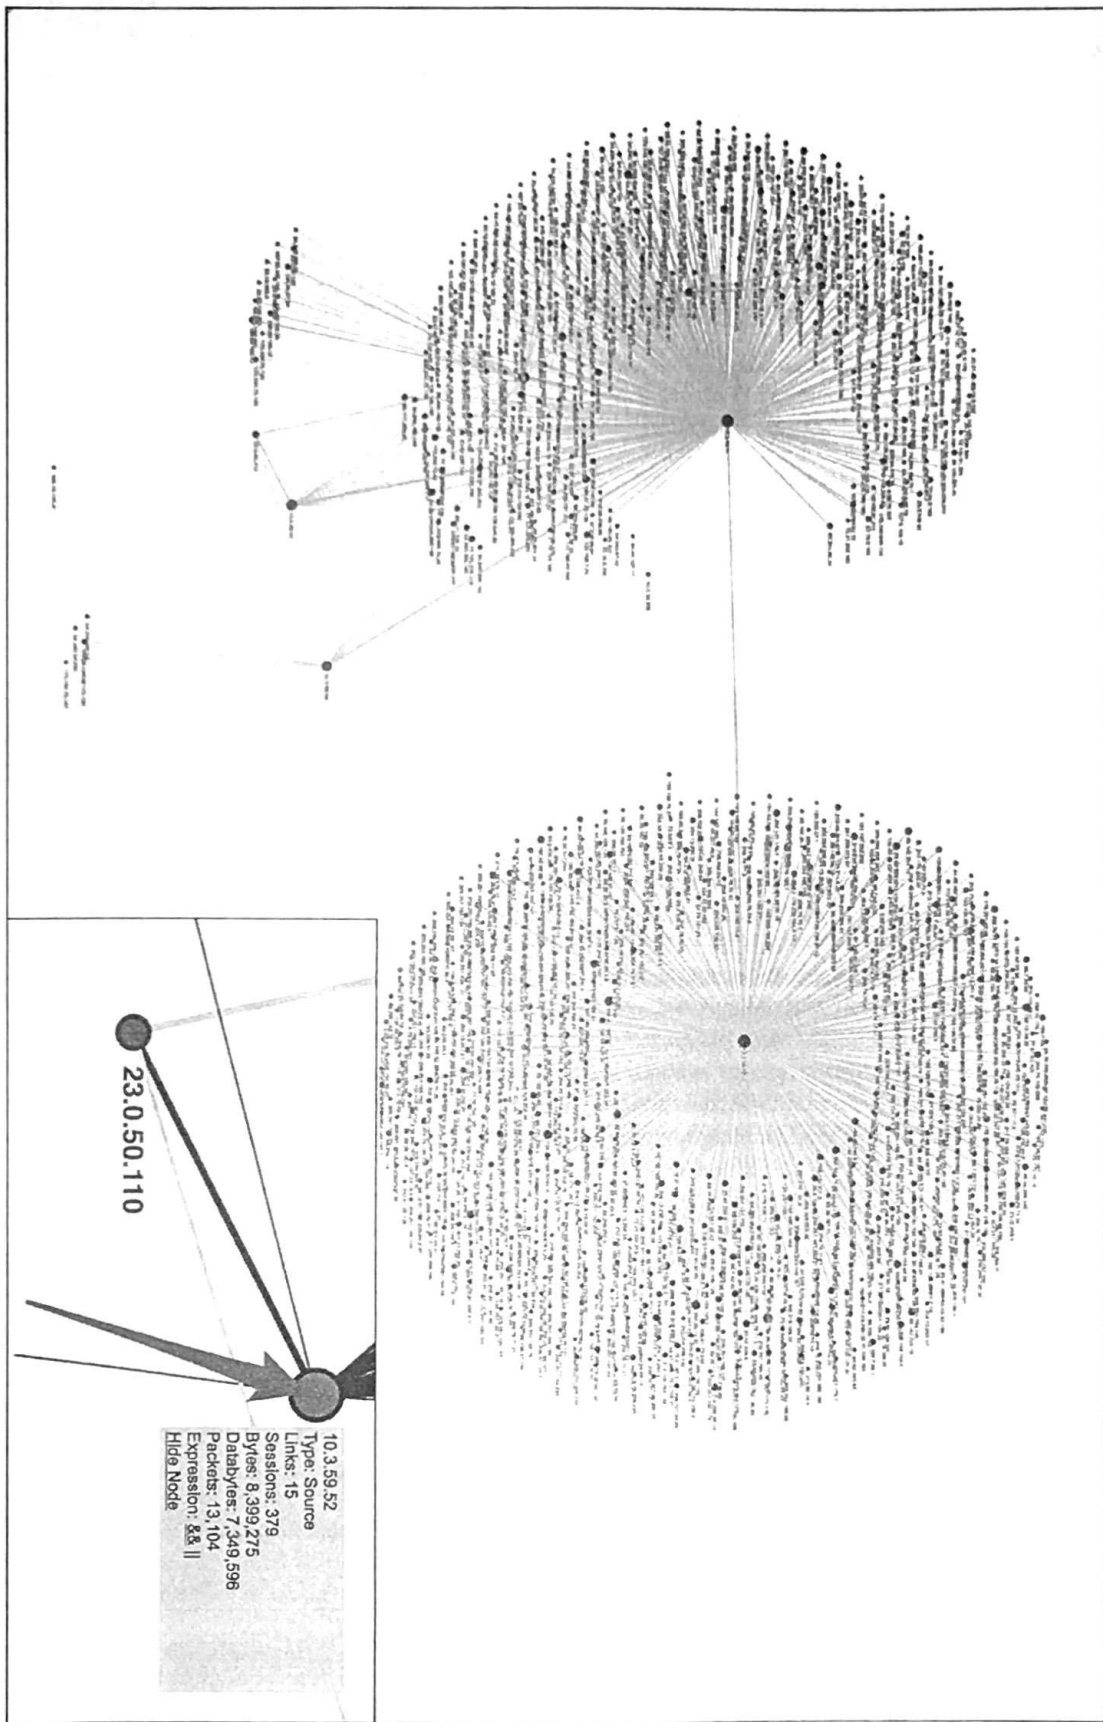
Sessions | SPI View | SPI Graph | Connections | Files | Stats | Settings

meleoch v0.14.2

Custom

Search Actions

Beginning Time: 2013/06/08 16:19:09 Ending Time: 2013/06/08 16:33:32 Delta Time: 0 00:14:23 Results bounded
 Query Size: 100000 Src: ip.src Dst: ip.dst Min Connections: 1 Nddle Distance: 250 Unlock



Lab II

Network Forensic Analysis Using Moloch

This page intentionally left blank.

Lab 11 Objectives: Network Forensic Analysis Using Moloch

- Become familiar with the Moloch full-packet capture and analysis platform
- Load captured traffic from existing pcap files into the Moloch platform
- Build search filters for traffic that has been indexed into Moloch
- Export traffic within the Moloch interface for external file extraction

This page intentionally left blank.

Lab 11 Takeaways: Network Forensic Analysis Using Moloch

- Moloch is a very useful tool for analysis of both existing PCAP files and in collection mode
- Enables a good forensic analysis workflow that equally supports both post-incident analysis and live IR actions
- Search filters can quickly identify sessions of interest
- Exploring session associations based on various fields can bring to light correlations that would have otherwise gone unidentified

This page intentionally left blank.



Commercial Tools, Wireless, and Full- Packet Hunting

© 2017 Lewes Technology Consulting, LLC and Mat Oldham | All Rights Reserved | Version # FOR572_C01_01

Authors:

Phil Hagen, Lewes Technology Consulting, LLC

phil@lewestech.com | @philhagen

Mat Oldham

mat.oldham@gmail.com | @roujisecurity

COURSE RESOURCES AND CONTACT INFORMATION



AUTHOR CONTACT

Phil Hagen/Lewes Technology Consulting, LLC
phil@lewestech.com | @PhilHagen

Mat Oldham
mat.oldham@gmail.com | @roujisecurity



SANS INSTITUTE

8120 Woodmont Ave., Suite 310
Bethesda, MD 20814
301.654.SANS(7267)



DFIR RESOURCES

digital-forensics.sans.org
Twitter: @sansforensics



SANS EMAIL

GENERAL INQUIRIES: info@sans.org
REGISTRATION: registration@sans.org
TUITION: tuition@sans.org
PRESS/PR: press@sans.org

This page intentionally left blank.

"As usual, SANS courses pay for themselves by Day 2. By Day 3, you are itching to get back to the office to use what you've learned."

Ken Evans, Hewlett Packard Enterprise - Digital Investigation Services

SANS Programs
sans.org/programs

GIAC Certifications
Graduate Degree Programs
NetWars & CyberCity Ranges
Cyber Guardian
Security Awareness Training
CyberTalent Management
Group/Enterprise Purchase Arrangements
DoDD 8140
Community of Interest for NetSec
Cyber Security Innovation Awards

SANS Free Resources
sans.org/security-resources

- E-Newsletters
 - NewsBites: Bi-weekly digest of top news
 - OUCH!: Monthly security awareness newsletter
 - @RISK: Weekly summary of threats & mitigations
- Internet Storm Center
- CIS Critical Security Controls
- Blogs
- Security Posters
- Webcasts
- InfoSec Reading Room
- Top 25 Software Errors
- Security Policies
- Intrusion Detection FAQ
- Tip of the Day
- 20 Coolest Careers
- Security Glossary



Search SANSInstitute

SANS Institute

8120 Woodmont Avenue | Suite 310
Bethesda, MD 20814
301.654.SANS(7267)
info@sans.org