The Network-Centric Incident Response and Forensics Imperative v1.0

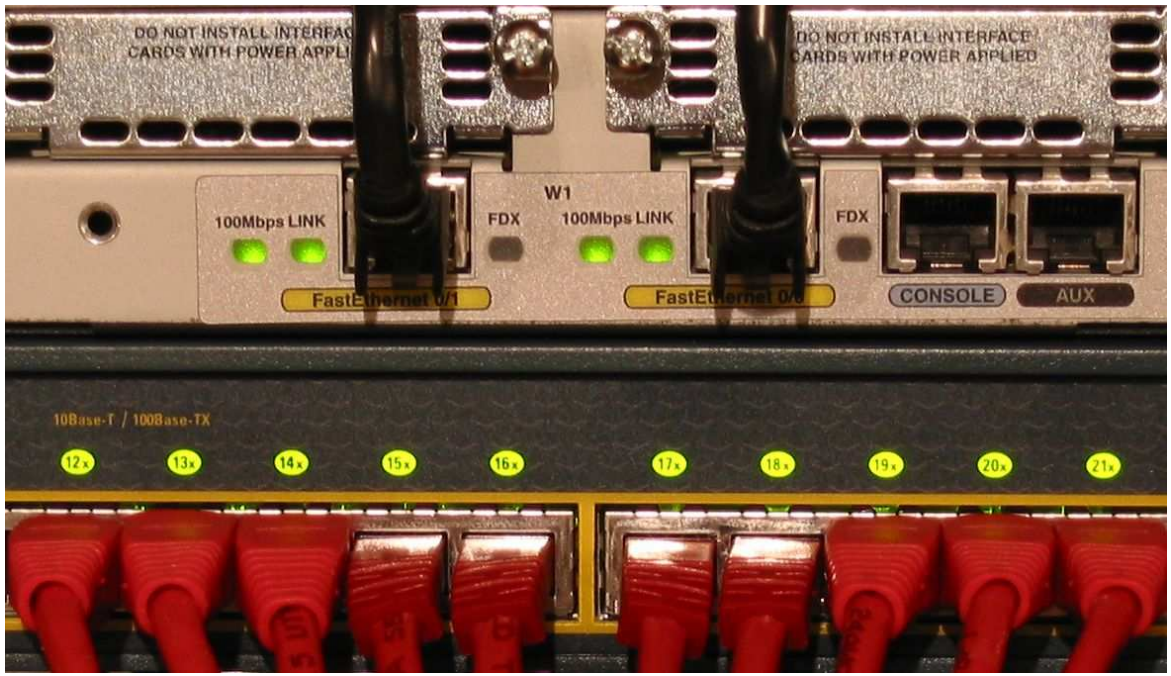

TAOSECURITY
THE WAY OF DIGITAL SECURITY

Richard Bejtlich
richard@taosecurity.com
www.taosecurity.com / taosecurity.blogspot.com

Copyright 2006 Richard Bejtlich

Introduction

- Bejtlich ("bate-lik") biography
  - TaoSecurity (05-present)
    - ManTech (04-05)
    - Foundstone (02-04)
    - Ball Aerospace (01-02)
    - Captain at US Air Force CERT (98-01)
    - Lt at Air Intelligence Agency (97-98)
  - Author
    - Tao of Network Security Monitoring: Beyond Intrusion Detection (solo, Addison-Wesley, Jul 04)
    - Extrusion Detection: Security Monitoring for Internal Intrusions (solo, Addison-Wesley, Nov 05)
    - Real Digital Forensics (co-author, Addison-Wesley, Sep 05)
    - Contributed to Incident Response, 2nd Ed and Hacking Exposed, 4th Ed

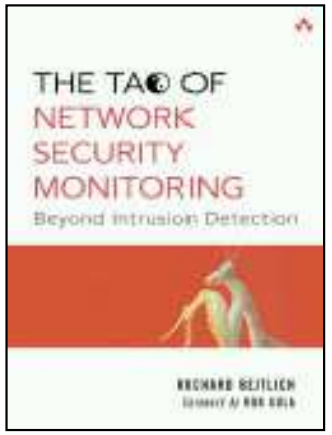

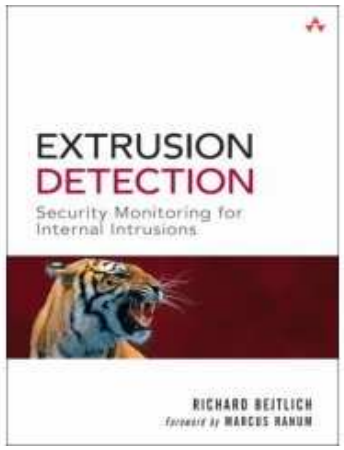

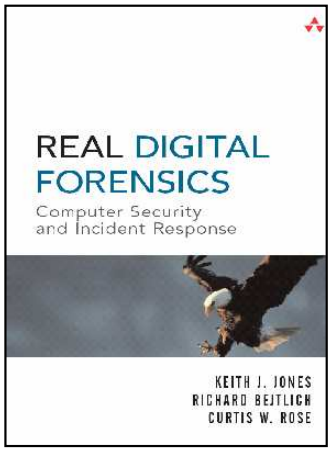

Copyright 2006 Richard Bejtlich

Traditional Host-Centric IR and Forensics

- Standard host-centric incident response and forensics scenario assumes:
  - Investigators know what systems are suspected of being compromised
  - Live response will yield reliable results that can be interpreted
  - Forensic duplication of a hard drive will show evidence of compromise

Copyright 2006 Richard Bejtlich

Reality: What Is Compromised?

- In reality, investigators may not know what systems are affected

Is it realistic to perform host-centric IR and forensics on all of these systems? Or even a subset?

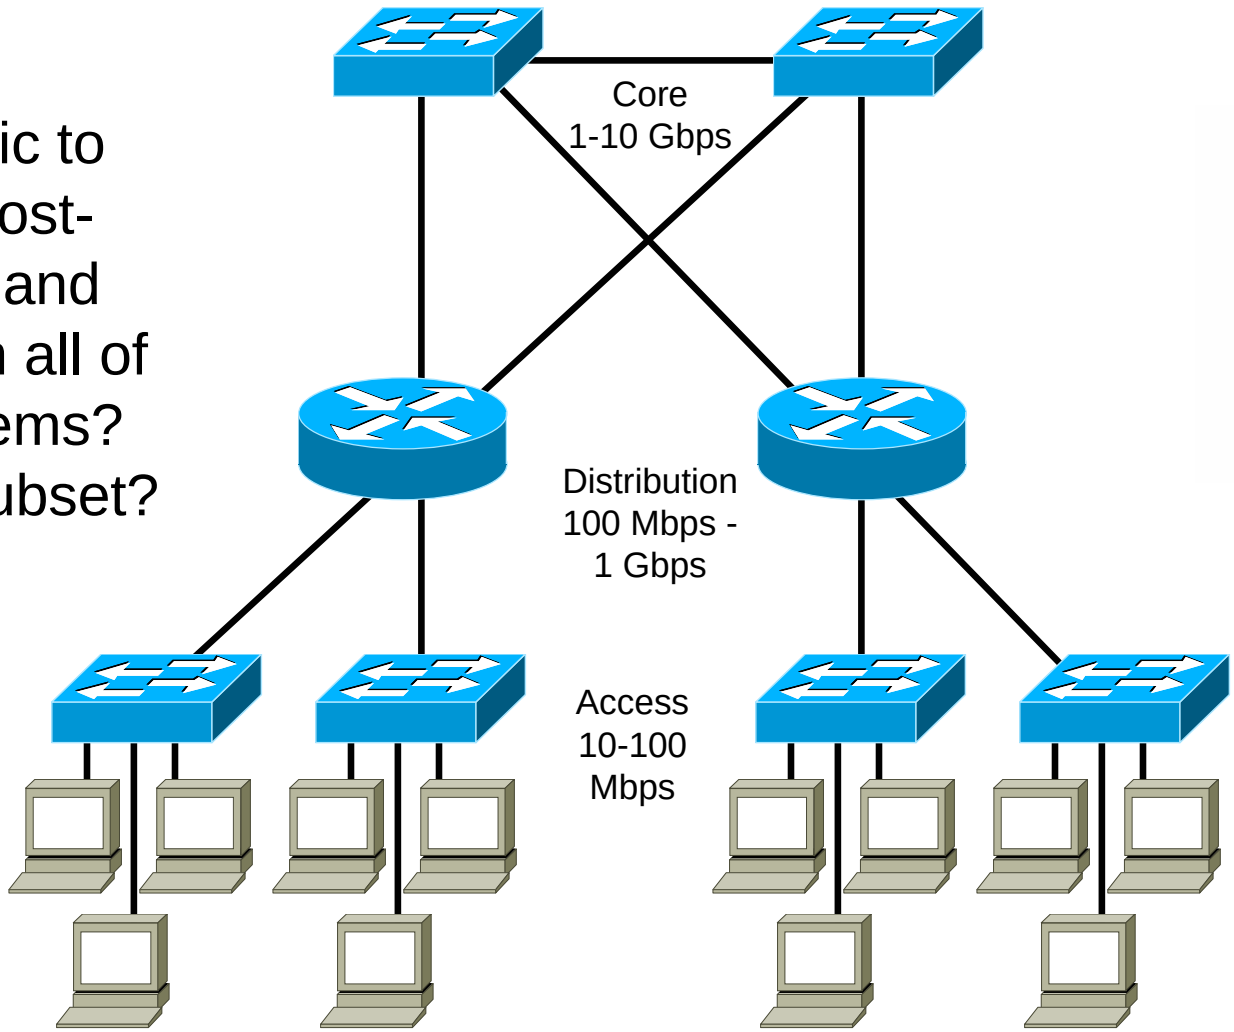

Copyright 2006 Richard Bejtlich

Reality: Is Data from Live Victims Trustworthy?

- Live response can be easily subverted by rootkits

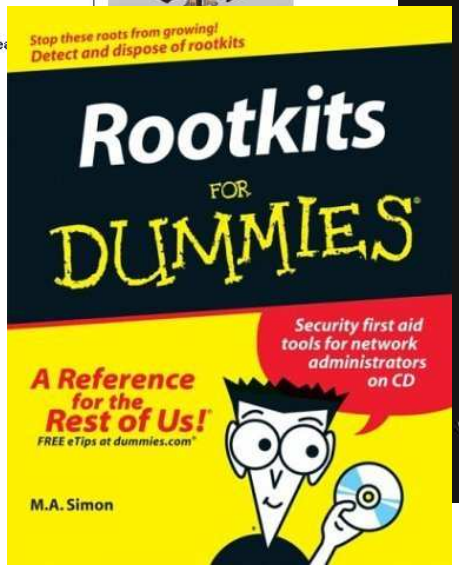

Copyright 2006 Richard Bejtlich

Reality: Will Hard Drive Duplication Be Useful?

- Some malware (and intruders) leave little or no evidence on the victim hard drive, frustrating forensic duplication and analysis
- For example, SQL Slammer did not copy any files to disk -- it was entirely memory resident
- High-end intruders take steps to leave as small a footprint as possible on hard drive
- Hardware-based disk encryption (Seagate Momentous FDE laptop drive) will hamper forensic investigations

TECHNICAL DETAILS:

W32.Slammer is a memory resident worm that propagates via UDP Port 1434 and exploits a vulnerability in SQL Server 2000 systems and systems with MSDE 2000 that have not applied the patch released by Microsoft Security Bulletin MS02-039. This bulletin was first available on July 24, 2002.

Copyright 2006 Richard Bejtlich

Reality: Three Principles

- Some intruders are smarter than you
  - It's too easy to fool every intrusion detection system ever built or to be built
  - The extreme case involving a rogue trusted insider must still be addressed
- Many intruders are unpredictable
  - Sophisticated intruders, in the aggregate, are always ahead of defenders
- Prevention eventually fails
  - Enterprise is too complex, staffed by overworked, under-resourced administrators meeting "business requirements"
  - Every enterprise will eventually be compromised

This intruder is probably not smarter than you, but may be unpredictable...

Copyright 2006 Richard Bejtlich

Introducing Trust

- When data from victim systems cannot be trusted, investigators may be forced to turn to other data sources
- The less contact users (and intruders) have with a system, the more trusted it is
- Network infrastructure, and especially data collected passively on the network, can be trusted to a higher degree than some host-based data
- Network data may still be blinded by encryption, degraded by high bandwidth, or unavailable due to lack of visibility -- but simply knowing traffic patterns can often solve a case
- Is there a network-centric approach to IR and forensics? Yes -- but most people think it's IDS or IPS

Copyright 2006 Richard Bejtlich

The Problem with IDS/IPS

- Most people install an IDS or IPS in monitoring mode and wait for alerts

Event	Example	IDS Action
Event 1	Ping Web site IP address	**Ignore**
Event 2	Visit Web site	**Ignore**
Event 3	Exploit Web site flaw	**Alert**
Event 4	FTP to retrieve tools	**Ignore**
Event 5	Install back door	**Ignore**
Event 6	Communicate with back door	**Ignore**
Event 7	Connect via SSH to another site	**Ignore**
Event 8	Transfer local exploit via SCP	**Ignore**
Event ...	And so on...	**?**

- Thanks to the IDS alert, an analyst is aware of a Web site problem -- but what about activity before or after the alert?

Copyright 2006 Richard Bejtlich

The Problem with IDS/IPS

- The previous slide presented a best-case scenario -- at least the attack was detected by the IDS! But what do you get with that alert?
  - Cryptic message about an attack
  - Maybe a packet that specifically triggered an alert
  - A reference to visit the vendor's Web site for more generic info
- Factors compounding the problem
  - Attack over HTTPS using SSL
  - Attack using insertion and evasion methods
  - Attack using a zero-day exploit undetected by any IDS
- Scarier scenarios
  - Use stolen credentials and connect via SSH
  - Compromise a customer or employee and ride their VPN
  - Go rogue and steal from your own company

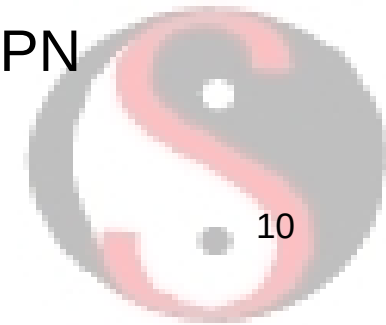

Copyright 2006 Richard Bejtlich

The Problem with IDS/IPS Vendors

- Many security developers and vendors believe one or more of the following
  - Attacks can be understood prior to execution
  - Methods to detect or prevent attacks can be encapsulated in programming logic
  - Customers will purchase, properly configure, and effectively deploy products offering sufficient defensive logic
  - The customer's environment will behave as anticipated by the developers and vendors
- Accordingly, developers and vendors field *alert-centric* products which act on those beliefs
- All of these beliefs must hold true in order to counter sophisticated threats, but few do

Copyright 2006 Richard Bejtlich

The Problem with IDS/IPS Operations

- Investigations with alert-centric systems quickly end, often without resolving the incident

Analyst sees original alert

ALERT

Queries database for alerts

Database returns single alert

ALERT

Investigation ends

- Analysts stuck with only alert data to inspect cannot make validation and escalation decisions
  - MSSPs call customers to ask if they have been compromised
  - Security personnel ignore alerts because they have no other data

Copyright 2006 Richard Bejtlich

NSM Theory

- Network security monitoring is the collection, analysis and escalation of indications and warning to detect and respond to intrusions

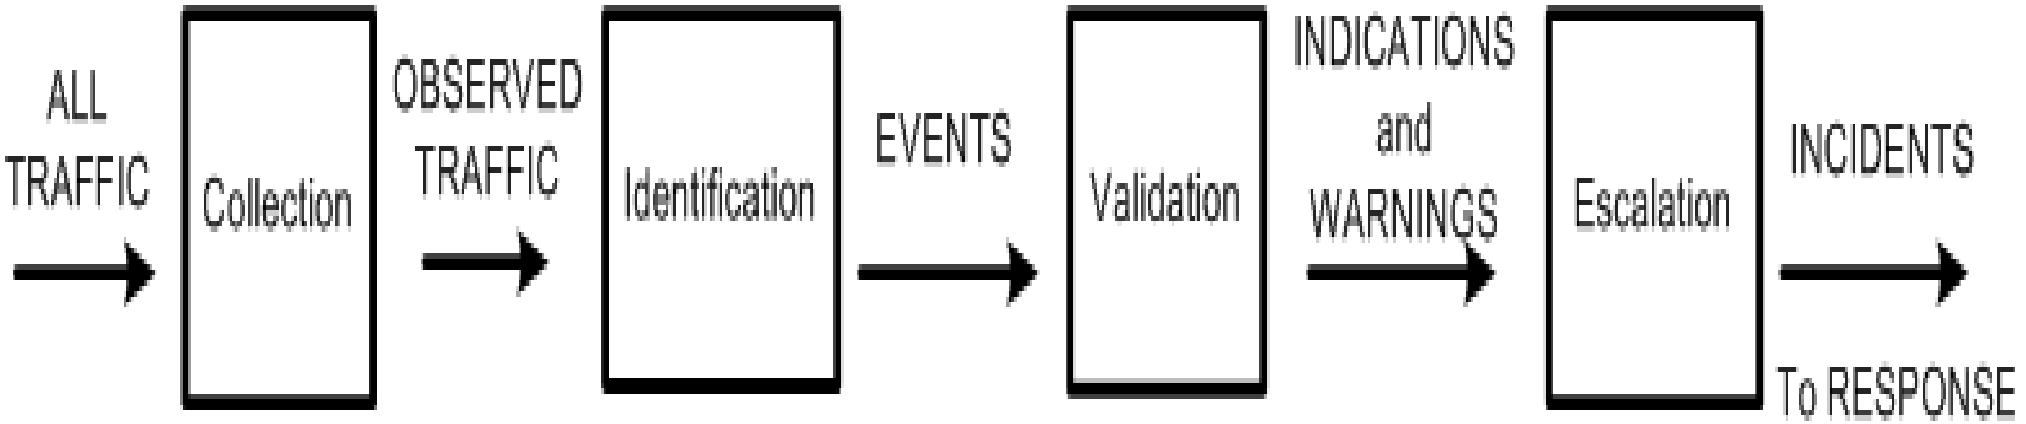

- NSM gives analysts the data they need to make decisions
- NSM treats all data as indicators, not "false positives" or "false negatives"

Copyright 2006 Richard Bejtlich

NSM Data Types

- NSM relies upon four forms of traffic-centric data
  - Statistical data (Capinfos, Tcpdstat, Trafshow)
    - Descriptive, high-level view of aggregated events
  - Session data (Argus, SANCP, NetFlow)
    - Summaries of conversations between systems
    - Content-neutral, compact; encryption no problem
  - Full content data (Tcpdump, Tethereal, Snort as packet logger)
    - All packet details, including application layer
    - Expensive to save, but always most granular analysis
  - Alert data (Snort, Bro, other IDSs)
    - Traditional IDS alerts or judgments (“RPC call!”)
    - Context-sensitive, either by signature or anomaly
- Sguil (www.sguil.net) is an interface to much of this in a single open source suite

Copyright 2006 Richard Bejtlich

NSM Data Types

Alert data

2004-12-06 18:25:47	10.200.211.32	49425	10.200.211.99	1023	6	RPC mountd TCP export request
2004-12-06 18:25:52	10.200.211.32	951	10.200.211.99	111	17	RPC portmap NFS request UDP
2004-12-06 18:25:52	10.200.211.32	628	10.200.211.99	1022	17	RPC mountd UDP mount request

Session data

2004-12-06 18:25:52	2004-12-06 18:25:52	10.200.211.32	628	10.200.211.99	1022	17	1	108
2004-12-06 18:25:52	2004-12-06 18:25:52	10.200.211.32	796	10.200.211.99	2049	17	1	108
2004-12-06 18:33:41	2004-12-06 18:33:41	10.200.211.32	49579	66.93.110.10	80	6	1	0

Full content data

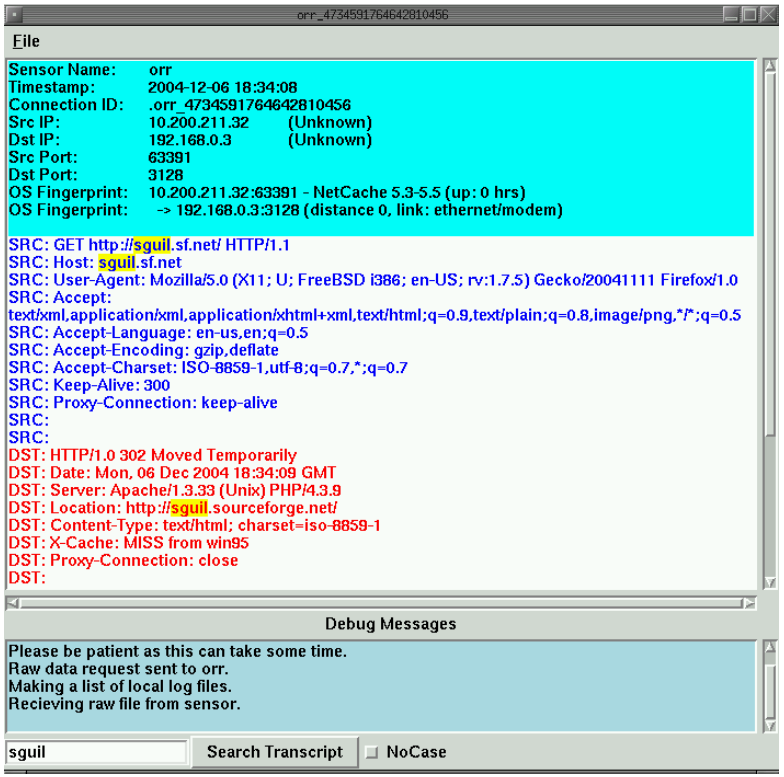

Statistical data

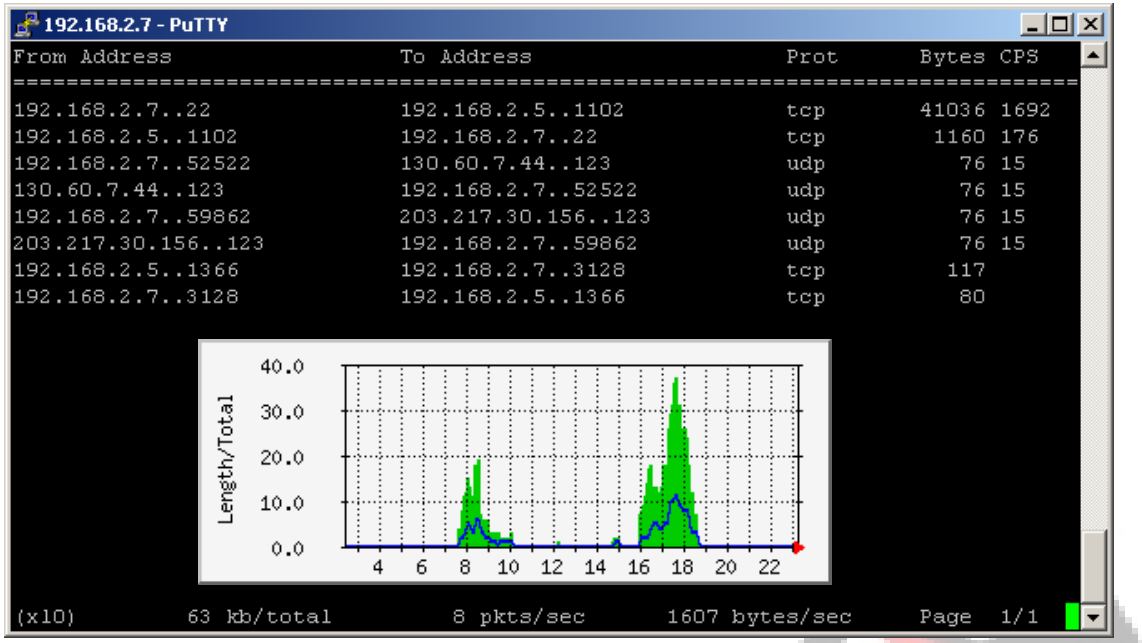

Copyright 2006 Richard Bejtlich

NSM Principles

- NSM does not try to anticipate attacks
- NSM uses a "dumb is smart" approach
  - NSM does not rely on fancy systems to pass judgements on network traffic, to the exclusion of all other collection mechanisms
  - NSM does leverage smart systems (IDS, network anomaly detection, etc.) for initial clues
- NSM session and full content collection is completely content neutral
  - Session and full content data are collected whether or not any other system thinks they are interesting
- NSM is not SIM/SEM: a SIM/SEM collects and correlates log sources which may or may not have any value

Copyright 2006 Richard Bejtlich

Detection with NSM

- Revisit intrusion scenario when NSM data is available

Event	Example	IDS Action	Helpful NSM Collection
Event 1	Ping Web site IP address	**Ignore**	**Session**
Event 2	Visit Web site	**Ignore**	**Session, Full Content**
Event 3	Exploit Web site flaw	**Alert**	**Alert, Session**
Event 4	FTP to retrieve tools	**Ignore**	**Session, Full Content**
Event 5	Install back door	**Ignore**	**Session, Full Content**
Event 6	Communicate with back door	**Ignore**	**Session, Full Content ***
Event 7	Connect via SSH to another site	**Ignore**	**Session**
Event 8	Transfer local exploit via SCP	**Ignore**	**Session**
Event ...	And so on...	?	

- Analysts have much more data to review

* if unencrypted (more common than you might think)

Copyright 2006 Richard Bejtlich

Detection with NSM

- Investigations with NSM present many more options

Analyst sees original alert

ALERT

Queries database for alerts

Database returns single alert

ALERT

Queries database for sessions

Analyst sees FTP to retrieve tools

SESSIONS

Reconstructs FTP control and data channels

FULL CONTENT

FTP data channel allows analysis of intruder back door

Queries database for sessions

Analyst sees connections to other IPs

SESSIONS

and so on...

Copyright 2006 Richard Bejtlich

- Identify the scope of the intrusion
  1. Deploy sensor with Sguil, or in a pinch, Tethereal and Argus
  2. Conduct traffic threat assessment using session data to discover anomalous connections
  3. Validate anomalous connections using full content data
  4. Evidence of suspicious or malicious connections points me to potential victims
  5. Review all connections to or from potential victims

Copyright 2006 Richard Bejtlich

Network-Centric IR and Forensics

6. Resulting systems receive host-based live response
7. Any time evidence reveals high probability of compromise, I recommend disabling access to victim until further course of action decided upon by client
8. Recommend hard drive forensics
9. As intruder's modus operandi is learned, deploy custom Snort signatures to provide alert data
10. Cage high-value targets if necessary

Copyright 2006 Richard Bejtlich

Network-Centric IR and Forensics

- A cage is usually a device in bridging mode that controls ingress and egress traffic to a target
- Deployed to let an intruder return, perhaps revealing his motive and/or the scope of the intrusion
- I encourage clients to not cage systems unless they are willing to watch them very closely and accept the consequences of misconfiguring the cage

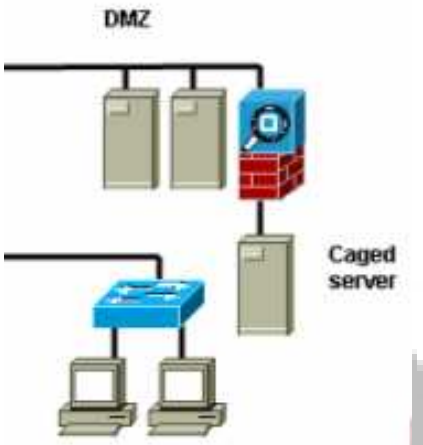

Copyright 2006 Richard Bejtlich

The Network-Centric IR and Forensics Imperative

- In a potential or actual compromise situation, demarcations between trusted and untrusted data sources and resources must be made
- A properly built and deployed NSM sensor can
  - Provide trusted data and remain invisible to an intruder
  - Validate or disprove host-based findings
  - Help scope an incident to conserve and direct host-centric response and forensics
  - Survive adversarial scrutiny due to its independence and reliability
  - Not interfere with hosts, as might be the case with host-based agents
- Where can you introduce a network incident monitoring and forensics appliance?

Copyright 2006 Richard Bejtlich