



The In's and Out's of SIEM Technology

Dr. Eugene Schultz, CISSP, CISM

Chief Technology Officer

Emagined Security

EugeneSchultz@emagined.com

- IX National Computer and Information Security Conference
- Bogota, Columbia
- June 18, 2009





**IX JORNADA
de SEGURIDAD
INFORMÁTICA**
Monitoreo y Evolución de
La Inseguridad Informática
Junio 17, 18 y 19 de 2009

Outline

Introduction

SIEM functionality

How SIEM tools generally work

How to select the right SIEM tool for you

Conclusion



Radical changes on the attack front

- Commonly occurring attack methods have changed substantially over the last five to ten years
 - “Frontal assault hacks” used to be common
 - With motivation for cyberattacks having changed so substantially, attack methods have changed accordingly
- Today’s attacks
 - Are much more subtle
 - Target applications (Web, Microsoft Office, Adobe Acrobat, and more)
 - Often involve sending many small pieces of content that must be reassembled by the destination host



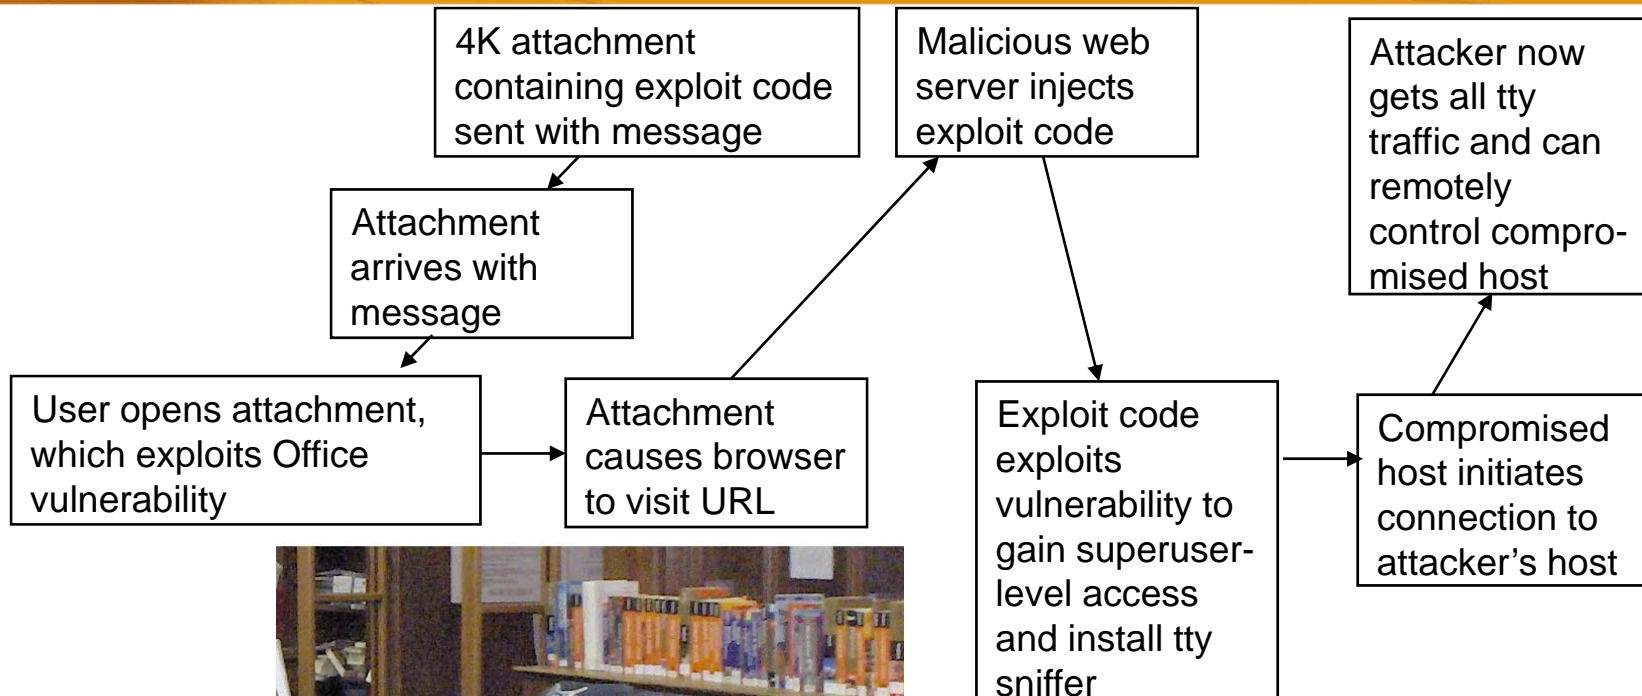
Radical changes on the attack front

- Today's more subtle attacks have substantially changed the nature of intrusion detection
 - Less reliance on conventional IDSs (and IPSs) per se
 - Indications of attacks are now usually very small and subtle, something that has made event correlation necessary (more on this soon)
 - More reliance on data mining based on large databases containing system log information over a long period of time (e.g., up to one full year)



**IX JORNADA
de SEGURIDAD
INFORMATICA**
Monitoreo y Evolución de
La Inseguridad Informática
Junio 17, 18 y 19 de 2009

Today's "hacks"

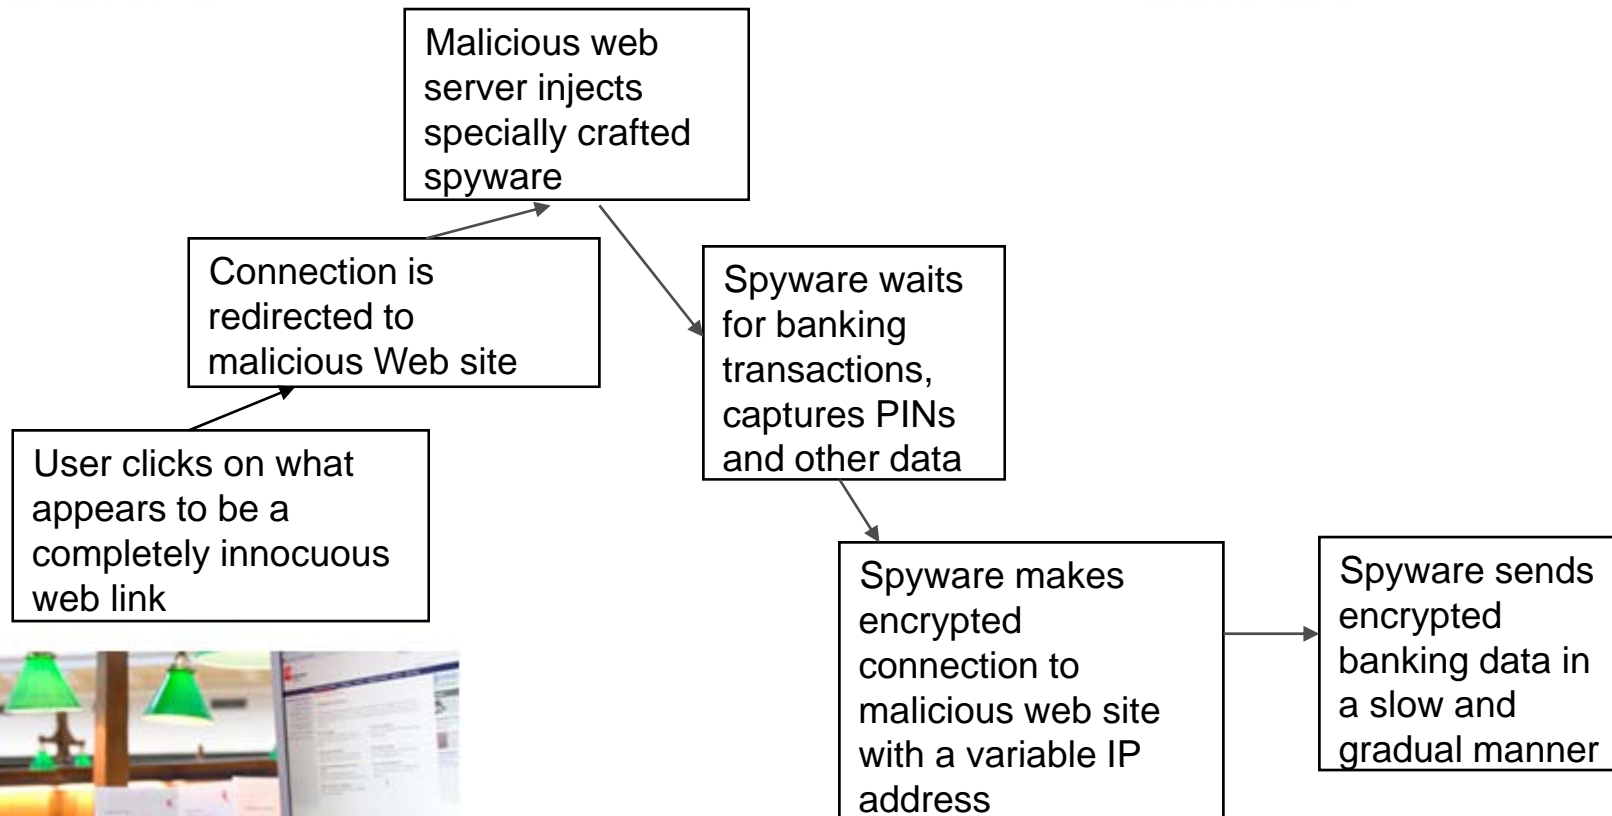



Continued on next slide



**IX JORNADA
de SEGURIDAD
de INFORMÁTICA**
Monitoreo y Evolución de
La Inseguridad Informática
Junio 17, 18 y 19 de 2009

Today's "hacks"





- We need a better way to obtain a wide range of detailed security-related information, analyze it down to the level of minute details, and respond if necessary
- SIEM technology provides such a way



**IX JORNADA
de SEGURIDAD
INFORMÁTICA**
Monitoreo y Evolución de
la Inseguridad Informática
Junio 17, 18 y 19 de 2009

Outline

Introduction

SIEM functionality

How SIEM tools generally work

How to select the right SIEM tool for you

Conclusion



What does a SIEM tool do? A high-level view...

- Aggregates and securely stores security event log and other data from hosts and devices throughout the network
- Makes log and other information easy to obtain through built-in reporting functionality
- Uses event correlation to determine whether attacks and security breaches have occurred (more on this shortly)
- Facilitates incident response
- Helps in achieving compliance



**IX JORNADA
de SEGURIDAD
INFORMÁTICA**
Monitoreo y Evolución de
la Inseguridad Informática
Junio 17, 18 y 19 de 2009

SIEM = SIM + SEM

- SIEM technology is actually a combination of SIM (Security Information Management) and SEM (Security Event Management) technology
 - SIM = log archival, management, reporting and compliance more than anything else
 - SEM = intrusion detection, event correlation, and incident response support



**IX JORNADA
de SEGURIDAD
INFORMÁTICA**
Monitoreo y Evolución de
La Inseguridad Informática
Junio 17, 18 y 19 de 2009

Reasons to use a SIEM tool

- Log management and archival
- Greater efficiency and convenience in accessing and analyzing security-related data
- Greater ability to separate “the wheat from the chaff” in discovering incidents
- Optimizes use of analyst time
- Greatly expands the scope of threat analysis
- Boosts awareness of security-related threats
- Provides ability to work around weaknesses in intrusion detection technology




Reasons to use a SIEM tool

- Elevates the functional level of intrusion detection expertise
- Provides information needed in incident response and forensics efforts
- Improves ability to keep networks healthy
- Saves time and money
- Helps with compliance by providing compliance reporting
 - ISO/IEC 27001/27002
 - PCI Data Security Standard (PCI-DSS)
 - Health Insurance Portability HIPAA
 - Sarbanes Oxley (SOX)
 - Gramm-Leach-Bliley Act (GLBA)
 - FISMA
 - More...



Where does input to SIEM tools come from?





**IX JORNADA
de SEGURIDAD
de INFORMÁTICA**
Monitoreo y Evolución de
La Inseguridad Informática
Junio 17, 18 y 19 de 2009

What types of output can SIEM tools process?

- Syslog messages*
- Intrusion detection system (IDS) output
- Firewall logs
- Windows Event Logs
- Vulnerability scan tool output
- SQL queries
- HTTP-S GETs
- LEA—Log Export Application Program Interface
- More



A syslog message (from Snort)

```
[**] [1:1913:8] RPC STATD UDP stat mon_name format string  
exploit attempt
```

```
[**]
```

```
[Classification: Attempted Administrator Privilege Gain] [Priority:  
1]
```

```
11/04-04:27:16.655166 192.168.1.1:807 -> 10.1.1.1:956
```

```
UDP TTL:3 TOS:0x0 ID:0 IpLen:20 DgmLen:1104 DF
```

```
Len: 1076
```

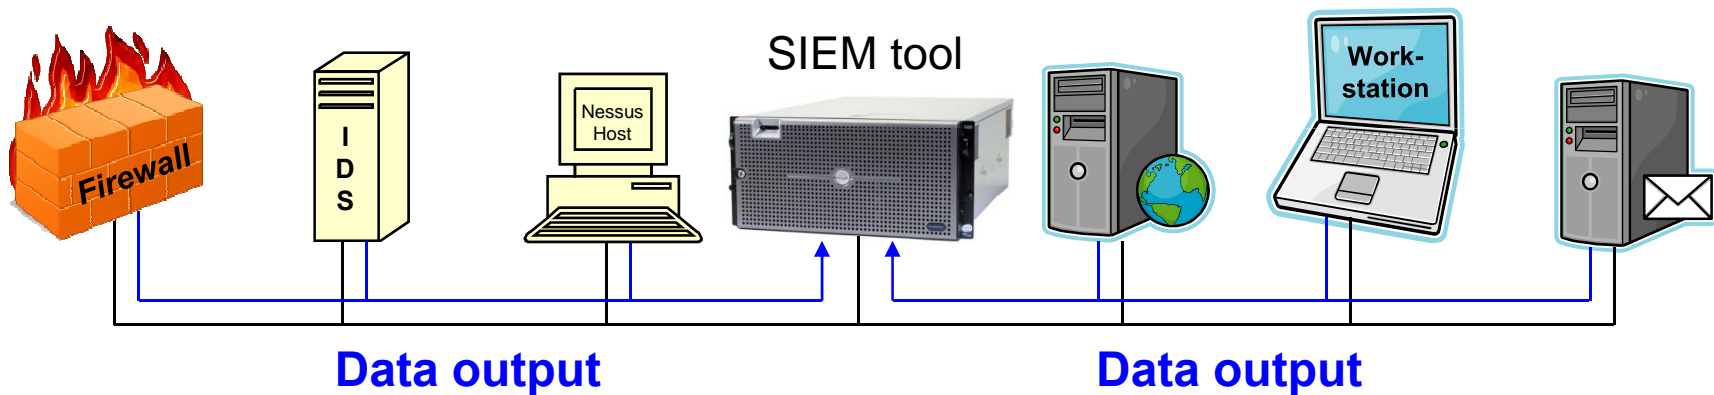
```
[Xref => http://www.securityfocus.com/bid/1480]
```

```
[Xref => http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2000-0666]
```



Data aggregation

- Data from various sources are collected
- Makes data available at a single location





**IX JORNADA
de SEGURIDAD
INFORMÁTICA**
Monitoreo y Evolución de
la Inseguridad Informática
Junio 17, 18 y 19 de 2009

Event correlation

- Much more powerful than data aggregation
- Takes multiple isolated events, combining them into a single relevant security incident
- Requires comparative observations based on multiple parameters such as
 - Source/destination IP addresses
 - Identifiable network routes
 - Type of attack
 - Type of malware installed on compromised systems
 - The time the activity began or ended
 - Many more

Continued on next slide





**IX JORNADA
de SEGURIDAD
INFORMÁTICA**
Monitoreo y Evolución de
la Inseguridad Informática
Junio 17, 18 y 19 de 2009

Event correlation

- Must include *all* traffic, regardless of whether it is
 - Inbound
 - Outbound
 - Internal



**IX JORNADA
de SEGURIDAD
INFORMÁTICA**
Monitoreo y Evolución de
la Inseguridad Informática
Junio 17, 18 y 19 de 2009

Outline

Introduction

SIEM functionality

How SIEM tools generally work

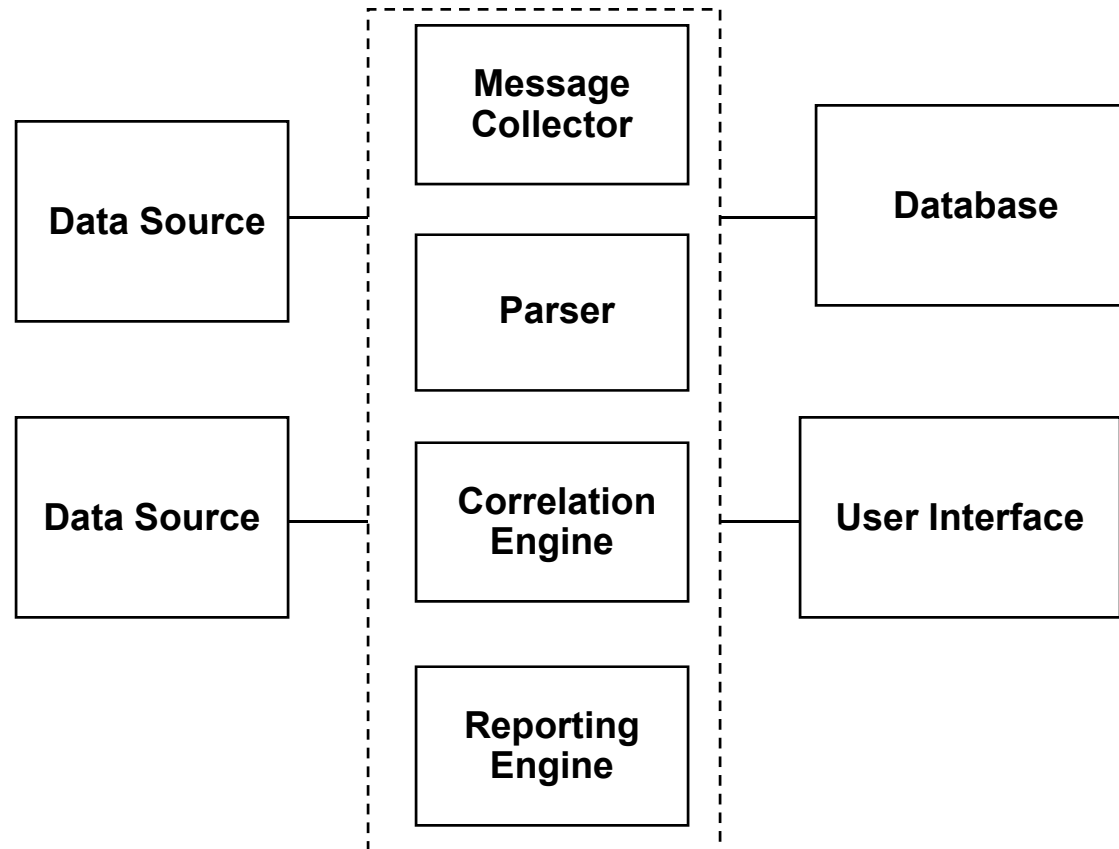
How to select the right SIEM tool for you

Conclusion



**IX JORNADA
de SEGURIDAD
de INFORMÁTICA**
Monitoreo y Evolución de
la Inseguridad Informática
Junio 17, 18 y 19 de 2009

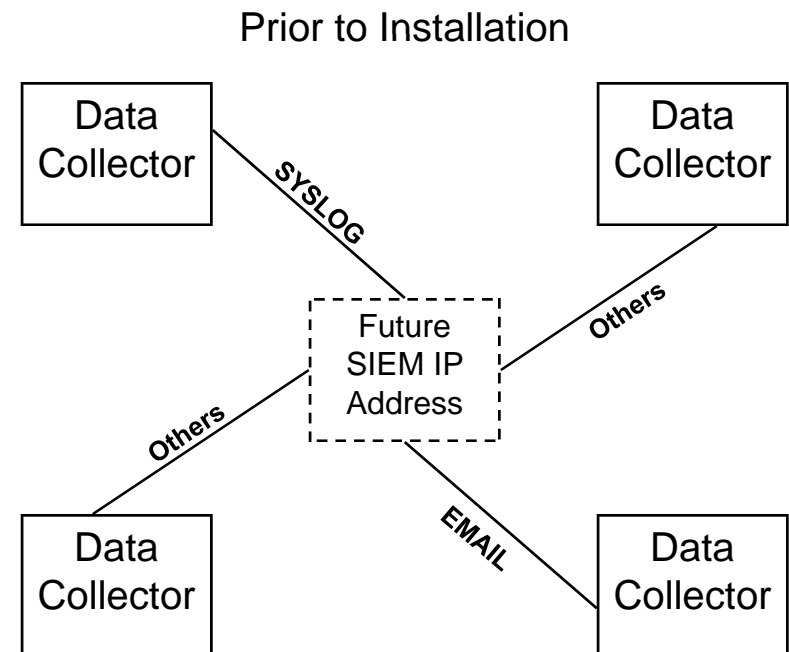
A typical SIEM architecture





Initial set-up

- Three initial procedures are necessary
 - Assign an IP address to the system on which your SIEM tool runs
 - Determine which hosts will provide data to the SIEM tool
 - Point all your data collectors' output to the IP address you have chosen





**IX JORNADA
de SEGURIDAD
de INFORMÁTICA**
Monitoreo y Evolución de
La Inseguridad Informática
Junio 17, 18 y 19 de 2009

How SIEM tools get log data ("push model")

1. syslog daemon enabled
2. `/etc/syslog.conf` configured
3. syslog daemon restarted



Data collector



`%PIX|ASA-6-308001: console enable password
incorrect for number tries (from 10.1.1.15)`

udp port
514



SIEM tool



Configuring syslog

- syslog must run on each host/network device that sends syslog output to SIEM tool
- For optimal analysis, you must usually send *all* syslog output to SIEM tool
- Example: in Unix and Linux hosts configure `/etc/syslog.conf` as follows and then *restart syslog*:

```
*.* @<IP_address_of_SEM_tool>
```



Use TAB, not spaces here



**IX JORNADA
de SEGURIDAD
INFORMATICA**
Monitoreo y Evolución de
La Inseguridad Informática
Junio 17, 18 y 19 de 2009

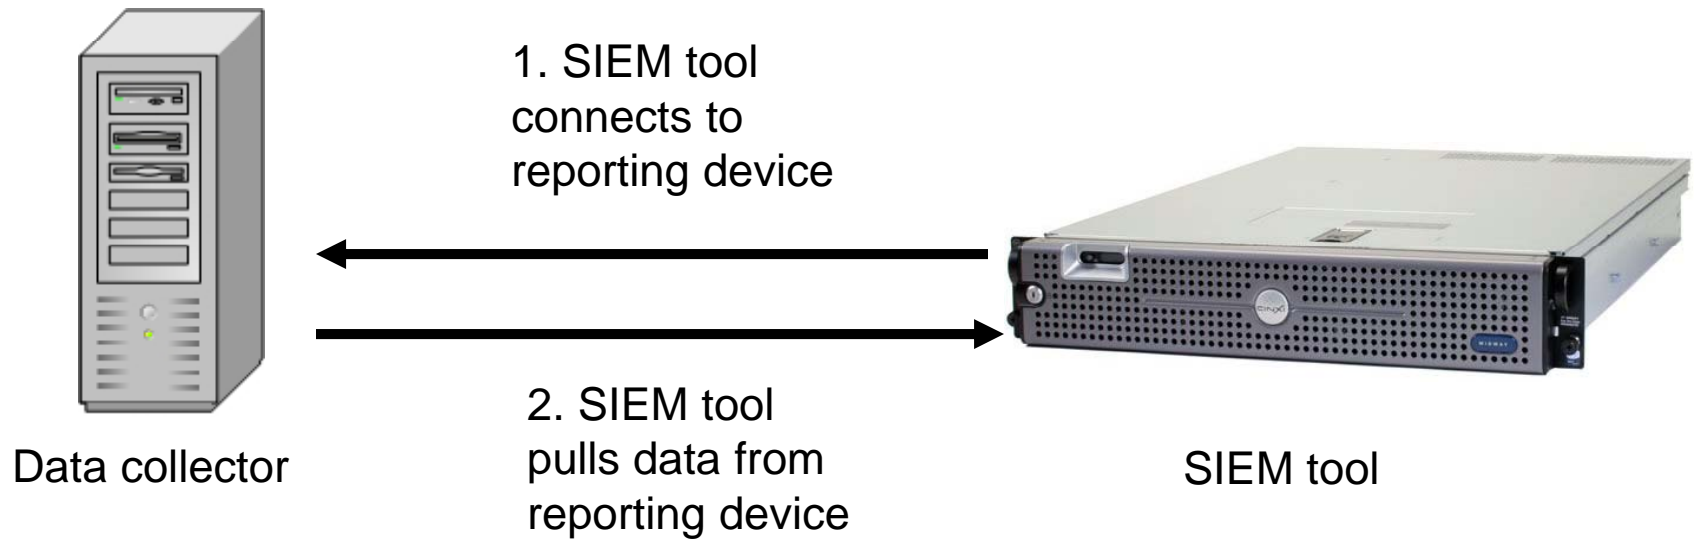
About agents

- An agent is a piece of software that acts on behalf a user or other process or program
- Used in SIEM tools to enable data collectors to send log and other data to the SIEM tool (server) when data collectors do not have syslog functionality
 - One agent must be installed on each data collector
- People's general feeling about agents in the SIEM arena is that they have numerous limitations and liabilities
 - Negative performance impact on hosts on which they reside
 - Add complexity to host environment
 - May introduce security vulnerabilities



**IX JORNADA
de SEGURIDAD
de INFORMÁTICA**
Monitoreo y Evolución de
la Inseguridad Informática
Junio 17, 18 y 19 de 2009

How SIEM tools get data ("pull model")





Down sides to the “pull model”

- Each SIEM tool must directly access each data collector to which the pull model applies—opens up vulnerabilities in the host on which the data collector runs
- Certain kinds of “pull model” access methods are
 - Poorly authenticated (if at all)
 - Unencrypted
- If connections in the “pull model” are encrypted, key management can be very time-consuming and difficult



Typical firewall port access requirements

<u>Protocol</u>	<u>Port</u>	<u>Direction</u>	<u>Why</u>
TCP	20	Inbound	ftp delivery of Blue Coat logs to SIEM tool
TCP	20	Outbound	Archiving logs via ftp
TCP	21	Inbound	ftp delivery of Blue Coat logs to SIEM tool
TCP	21	Outbound	Archiving logs via ftp
TCP	22	Inbound	Remote configuration of SIEM tool
TCP	25	Inbound	Receipt of Nessus vulnerability scans
TCP	25	Outbound	Sending incident and case alerts
UDP	53	Outbound	DNS resolution requests
TCP	80	Outbound	Cisco IDS polling
TCP/UDP	111	Outbound	NFS mounts of backups and reports
UDP	123	Inbound	NTP requests and responses if SIEM tool is configured as an NTP server
UDP	123	Outbound	NTP requests and responses
UDP	137	Outbound	CIFS/SMB mounting of backups and reports
UDP	138	Outbound	CIFS/SMB mounting of backups and reports
TCP/UDP	139	Outbound	CIFS/SMB mounting of backups and reports
UDP	161	Inbound	SNMP requests
UDP	162	Outbound	SNMP SEM notification traps

Continued on next slide



**IX JORNADA
de SEGURIDAD
INFORMÁTICA**
Monitoreo y Evolución de
La Inseguridad Informática
Junio 17, 18 y 19 de 2009


Typical firewall port access requirements

<u>Protocol</u>	<u>Port</u>	<u>Direction</u>	<u>Why</u>
TCP	443	Inbound	Portal and console connections to SIEM tool
TCP	443	Outbound	Cisco IDS polling
TCP/UDP	445	Outbound	CIFS/SMB mounting of backups and reports
UDP	500	Inbound	ISAKMP traffic for Check Point
UDP	500	Outbound	ISAKMP traffic for Check Point
UDP	514	Inbound	Syslog messages
UDP	514	Outbound	Notifications by the SIEM tool
TCP/UDP	1433	Outbound	MS-SQL traffic for ISS IDS polling
TCP/UDP	1434	Outbound	MS-SQL traffic for ISS IDS polling
TCP/UDP	2049	Outbound	NFS mounting of backups and reports
UDP	2055	Inbound	Netflow data
UDP	6343	Inbound	sFlow data
TCP	40000	Inbound	Console/appliance traffic; SSL encryption
TCP	40001	Inbound	Console/appliance traffic; SSL encryption
TCP	41000	Inbound	Console/appliance traffic; SSL encryption
TCP	46645	Inbound	Master console traffic
TCP	46646	Outbound	Master console traffic
ICMP	N/A	Inbound	
ICMP	N/A	Outbound	



How most intrusion detection systems (IDSs) work

Attack "signature"



ALARM
"site exec"
attack!

```
quote site exec exec echo toor::0:0:>:::/:/bin/sh >> /etc/passwd
```



**Intende
d Victim**



About rule-based intrusion detection

- Rules are logic conditions based on knowledge concerning how real-life attacks occur
 - Based on combinations of possible indicators of attacks, combining them to see if a rule condition has been fulfilled
 - Example of logic: Condition A leads to Condition B (which might also lead to Condition C)
 - Intrusion detection system (IDS) signatures can comprise one or more conditions
 - *Sequence and timing* of events are critical
 - For example, a remote FTP login followed by an attempt to obtain a copy of a Unix host's password file usually means that an attack has occurred
- Evaluation
 - Higher correct detection rates and lower false alarm rates
 - Provides a much more comprehensive analysis—IDS rules are based only on the events that any one IDS can detect



About event correlation

- Event correlation goes well beyond typical rules in that they incorporate logic based on series of related events known to occur in connection with attacks
- By default, they do not rely on any one
 - IDS signature
 - Device
 - Rule
- Involve correlation of multiple threads associated with potential incidents
- Provide a very powerful type of attack pattern analysis based on independent inputs



How event correlation usually works

- One event correlation rule must be created for each attack pattern
- Each event correlation rule extracts information in messages and other information received from collection devices
- Causes threads that listen for additional event-related information to be created
- If the logic of an event correlation rule is met, the rule triggers an incident
- If insufficient information needed to fulfill the logic of the potential attack related to a thread is not received within a criterion period of time, the thread dies



About event correlation rules

- For example, an event correlation rule might specify that Event A followed by Event B followed by Event C occurs whenever a certain attack occurs
 - Event A might be a vulnerability scan from a particular IP address
 - Event B might be a successful connection from the same source IP address to an internal host afterwards
 - Event C might be a successful connection from the same internal host to the same external IP address afterwards
 - Each of these events would have to be detected by multiple reporting devices
 - This chain of events is a very strong indicator that there has been a successful attack against an internal host

A
↓

B

↓

C

=


Attack



**IX JORNADA
de SEGURIDAD
INFORMÁTICA**
Monitoreo y Evolución de
La Inseguridad Informática
Junio 17, 18 y 19 de 2009

Major advantage of event correlation— lowering false alarms

- Signature-based IDSs cannot weed out false alarms
 - Suppose that event B is a false alarm produced by an IDS
 - Low-level rule would still show three separate events
- Event correlation significantly reduces false alarms
 - If A is true, but B is not true, the logic condition (A \rightarrow B and B \rightarrow C) will *not* be met
- Benefit—reduction in the high toll that false alerts inflict
 - Fewer non-working hour phone calls
 - Less frustration within IT
 - Increased productivity





Another advantage: Zero-day attack discovery

- Attackers continually change their tactics
- Good event correlation logic is not dependent on any particular attack method
 - Good event correlation rules consider many patterns of activity instead
- Events that are analyzed are normally (but not always) within a specified time interval, however
- Benefit—provides a very powerful way to
 - Detect zero-day attacks
 - Improve security



Another benefit: Determining whether attack was successful

- The availability of vulnerability scan data on a SIEM tool can lead to determining whether or not an attack was successful
 - Event correlation logic can be expanded to include an additional step, D (A -> B -> C -> D), where D represents the fact that the host was vulnerable to the particular attack in question
 - If condition D is true, and A, B, and C are also true, a successful attack has occurred
 - If condition D is false, but A, B and C are true, the attack has been unsuccessful
- Benefit—this further reduces false alarms and thus reduces the number of incorrect alerts



Example: telnet from external to internal host

```
ns5xt: NetScreen device_id=ns5xt system-  
notification-00257(traffic):  
start_time="2005-09-15 09:41:44"  
duration=5 policy_id=0 service=tcp/port:  
proto=6 src zone=Trust dst zone=Untrust  
action=Permit sent=1034 rcvd=19829  
src=206.3.3.2 dst=10.1.1.167  
src_port=1059 dst_port=23 translated  
ip=10.1.1.162 port=23
```

A telnet connection like this almost never causes an IDS, intrusion prevention system (IPS), firewall, or any other device to identify this event as malicious



Example: ssh connection from internal to external host afterwards

```
ns5xt: NetScreen device_id=ns5xt system-  
notification-00257(traffic):  
start_time="2005-09-15 09:41:44"  
duration=5 policy_id=0  
service=tcp/port:1214 proto=6 src  
zone=Trust dst zone=Untrust  
action=Permit sent=1034 rcvd=19829  
src=10.1.1.167 dst=206.3.3.2  
src_port=1059 dst_port=22 translated  
ip=10.1.1.162 port=22
```

Once again, virtually no IDS, IPS, firewall, and so on would identify this as a malicious event



Inbound telnet followed by outbound ssh to the same host often indicates malicious activity

- A good event correlation algorithm is capable to identifying and reporting this *pattern* of events

BUT

- IDSs, IPSs, firewalls, and even the targeted host itself will almost without exception fail to identify and report this pattern
- Shows just how valuable event correlation can be!



The problem with low-level event correlation rules

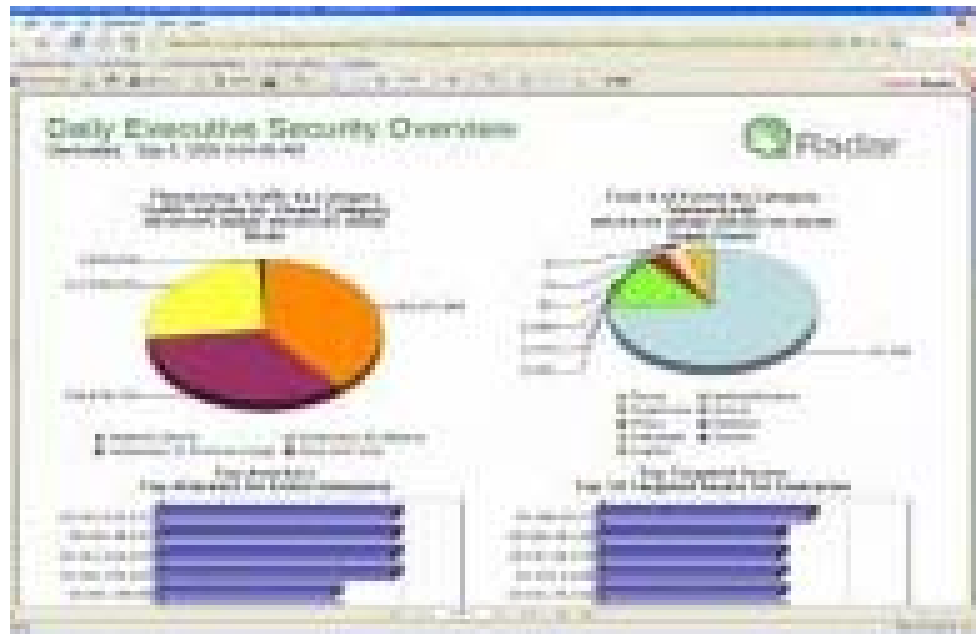
- Some SIEM tools perform event correlation at a very low level (e.g., by simply confirming that an event occurred because of more than one indication of the event)
 - These tools require many hundreds (if not thousands) of low-level rules
- A SIEM tool that embodies hundreds of low-level rules is
 - Inefficient
 - Likely to miss patterns of malicious network activity
 - Likely to yield false alarms that tools based higher level event correlation would have caught and eliminated
- Having fewer higher level event correlation rules provides a much more efficient and powerful way to correlate the output of devices that provide intrusion detection and other critical threat management data



IX JORNADA de SEGURIDAD de INFORMÁTICA


Monitoreo y Evolución de la Inseguridad Informática
Junio 17, 18 y 19 de 2009

Reporting output





Gaining access to a SIEM server



Workstation
with Web
browser



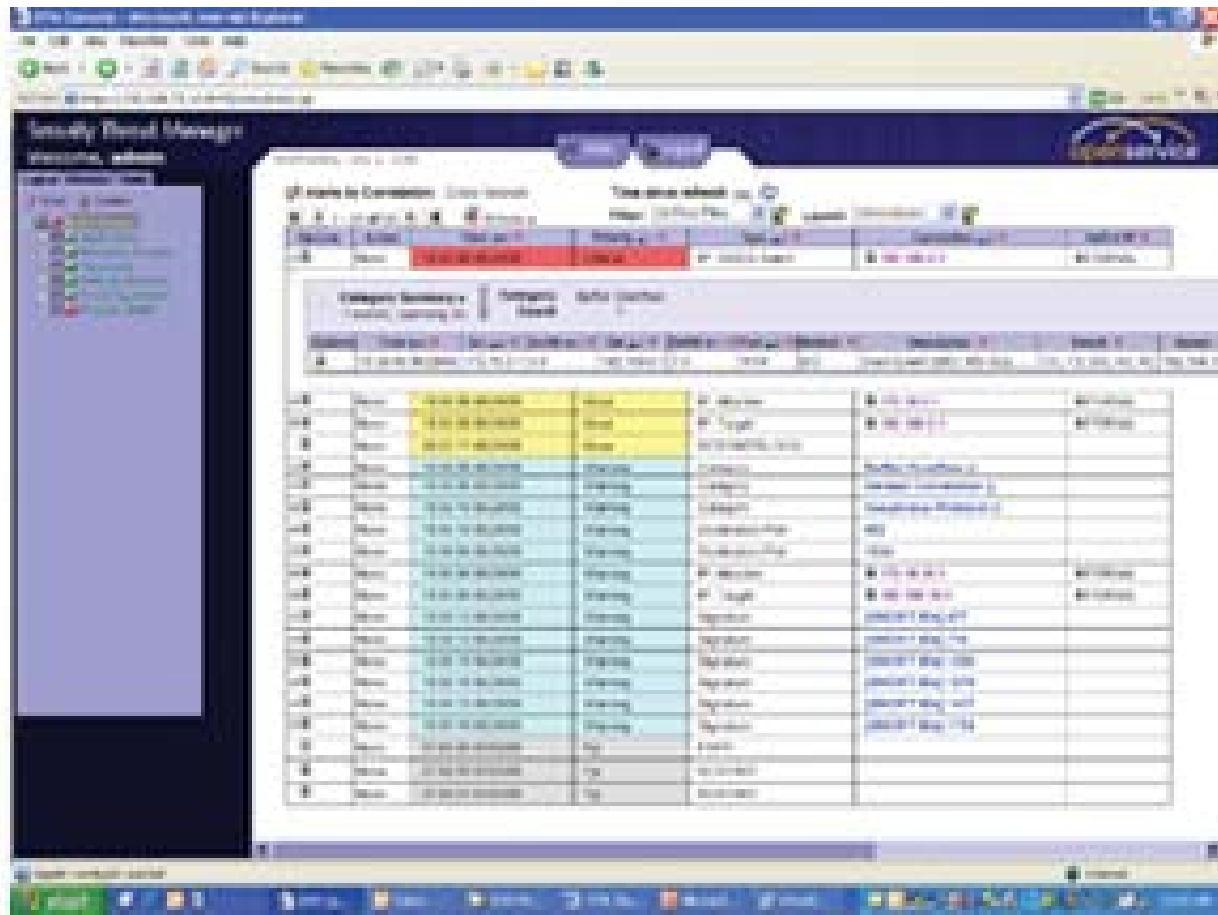
HTTP-S or SSH



SIEM tool



The graphical user interface on one SIEM server





**IX JORNADA
de SEGURIDAD
INFORMÁTICA**
Monitoreo y Evolución de
la Inseguridad Informática
Junio 17, 18 y 19 de 2009

Outline

Introduction

SIEM functionality

How SIEM tools generally work

How to select the right SIEM tool for you

Conclusion



A tragic mistake

- One of the biggest critical success factors in SIEM technology is buying the right SIEM product for an organization's security needs

BUT

- Many organizations buy SIEM products that do not really meet their needs very well
 - They often end up scrapping the product that they have bought, forcing them to look for and eventually buy another one—a gigantic waste of time and money
 - Additionally, some SIEM products have much better functionality than others (to be explained shortly)



Selection criteria for SIEM tools

- Setup and installation
- Reliance on agents
- Role-based access
- Auto discovery of
 - Devices and hosts
 - Changes in the network
- Granular privileges--so that you can enforce the "least privilege principle"
- User friendliness of user interface
- Quality of event correlation
- Performance—faster is better
- Encryption of all network traffic

Continued on next slide





Selection criteria for SIEM tools

- Encryption of all data at rest
- Log archival
- Log auto-deletion at selected time(s)
- Huge storage capacity
- Ability to create custom reports
- Ability to create custom rules
- Quantity and quality of reports
- Availability of compliance reports
- Ability to send alerts and alarms based on selected criteria
- Ability to easily obtain detailed information about attacks

Continued on next slide





Selection criteria for SIEM tools

- A layered set of defenses to protect the SIEM itself
- Audit logs that capture all actions performed on the SIEM tool itself
- Ability to interface with a wide range of data collectors (show the basic set up--client-server)
- Forensic preservation of all collected data
- Incident response facilitation
- Intrusion prevention functionality
- Hierarchical functionality
- Ability to perform all administration activity on one SIEM
- Degree of reliance upon agents



Twelve pitfalls to avoid

1. Extremely limited functionality--some SIEM tools do little more than collect and archive log information
2. A complex, drawn-out and expensive installation process
3. Complex and customer-adverse license and maintenance contracts
4. Overreliance on any one source of security event information (e.g., Snort)
5. Ability to interface with only a limited set of data collectors
6. Low-level event correlation
7. Ability to cover only a relatively small portion of a network
 - Necessitates having to purchase more SIEM tools



Twelve pitfalls to avoid

8. Bad user interfaces--some SIEM tools
 - Have very little graphical user interface (GUI) functionality—they mostly present mostly long, tedious-to-read lists of data
 - Difficult-to-use administrative functionality
9. Limited reporting capability
10. Bad disk space management—disk space management must be
 - Seamless
 - Reliable
11. Lack of ability to customize SIEM functionality
12. All kinds of dashboards, colorful icons, etc. that have little meaning or significance



Some final advice

- Each SIEM tool is available as an appliances or software product (or sometimes both)
 - Which is better? It depends on your particular needs and preferences...
 - Never underestimate the importance of properly tuning the parameters and functionality of a SIEM tool
 - Usually requires considerable time and effort if it is to be done right
 - Be sure to consider Total Cost of Ownership (TCO) when deciding on a SIEM tool
 - Purchase price is only one of many total costs
 - If your organization does not have enough money to buy a commercial SIEM tool, consider OSSIM* (www.ossim.net)
 - An open source product
 - Supported and maintained remarkably well
- * - Open Source Security Information Management



**IX JORNADA
de SEGURIDAD
INFORMÁTICA**
Monitoreo y Evolución de
la Inseguridad Informática
Junio 17, 18 y 19 de 2009

Outline

Introduction

SIEM functionality

How SIEM tools generally work

How to select the right SIEM tool for you

Conclusion



Conclusion

- SIEM tools do not by any means solve all security problems, but they are an effective defense method in a “defense-in-depth” strategy
 - For example, SIEM technology can identify and respond to attacks that might otherwise be missed by IDSs and IPSs
- Additionally, SIEM tools solve many practical problems
 - Log management and archival
 - Compliance reporting
- Be careful, though, as many SIEM products appear to be much better in functionality than they really are
 - Use the criteria presented earlier in this presentation to evaluate and select each candidate SIEM product
 - Always thoroughly test and evaluate any product before buying it
- SIEM technology is still very young—it is bound to only get better over time



Questions?

Emagined Security
2816 San Simeon Way
San Carlos, CA 94070
(650) 593-9829
eugeneschultz@emagined.com
web: www.emagined.com

See blog.emagined.com

For a PDF copy of these slides send
email to:

YvonneVega@emagined.com

