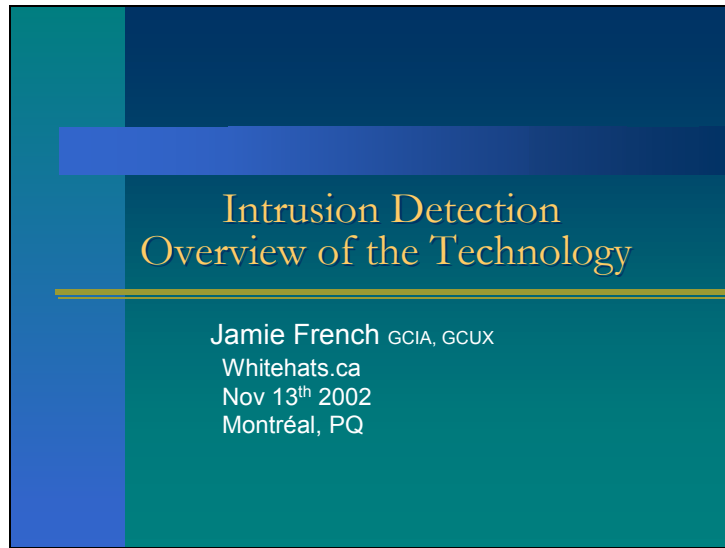


Slide 1

The slide features a dark blue background with a teal vertical bar on the left and a horizontal blue bar at the top. The title is centered in a serif font, and the author information is in a sans-serif font below a thin yellow line.

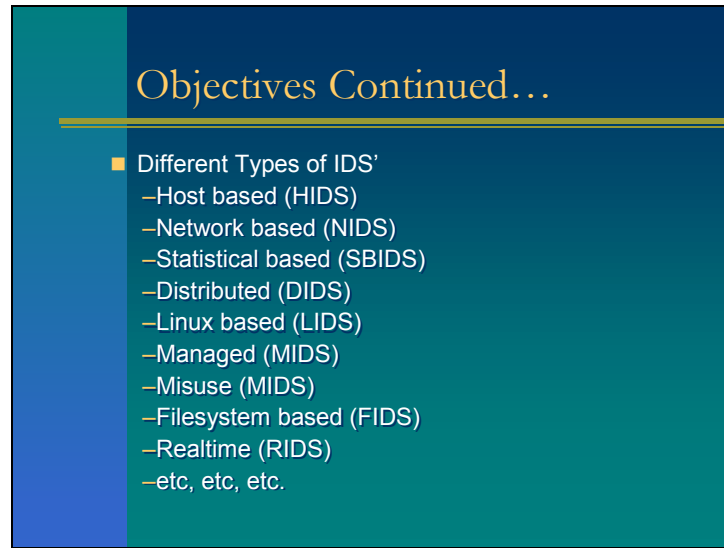
Intrusion Detection
Overview of the Technology

Jamie French GCIA, GCUX
Whitehats.ca
Nov 13th 2002
Montréal, PQ

Slide 2

Objectives

- Why should we spend \$\$\$ on IDS?
 - Threat vs Risk vs Cost Assessment
 - Impact and damage assessments
 - Defense-in-depth
 - Redundancy
 - Evidence trail



Objectives Continued...

- Different Types of IDS'
 - Host based (HIDS)
 - Network based (NIDS)
 - Statistical based (SBIDS)
 - Distributed (DIDS)
 - Linux based (LIDS)
 - Managed (MIDS)
 - Misuse (MIDS)
 - Filesystem based (FIDS)
 - Realtime (RIDS)
 - etc, etc, etc.


It isn't hard to see a trend here. Everybody and their uncle is deciding to coin a term for an Intrusion Detection System that does something slightly different or implements its detection mechanisms with a different approach. Really, this is a market where a lot of different solutions appeared almost overnight, each trying to differentiate their product into the niche in an attempt to succeed. Don't be overly concerned about the terminology. There are only really a few types of IDS' and we'll cover them in the coming slides.

Objectives Continued...

- So, what kind of IDS is right for me?
 - Feature comparison
 - Future considerations
 - Weighted Comparison

Objectives Continued...

- The fortune tellers crystal ball



What's here today and coming tomorrow?

Threat vs Risk vs Cost

- Who are our threats (enemies)?
- Do we know where we stand right now?
- Use of our resources to attack others (liability).
- Security of the country???
- Can we accept the risk?
- Figure out how much your assets are worth.
- If asset is lost, can it be replaced?
- Do we have the resources to deal with the threat?
- Insurance on intellectual property.
- Choose your high pri assets and protect them.

Threats:

You have to figure out where you stand right now and benchmark yourself in a self assessment. Hiring a third party to do this is a common practice by large organisations and it would likely involve a security audit. The threats are real. Computers are compromised regularly and real damage occurs. For a small snapshot, take a look at defacement mirrors like <http://www.zone-h.com/>. Are you in a situation where cyber-attack threats are elevated? Are you in a position where there are protesting parties who disagree with your business practices and may wish to cause your company harm? A good resource on the topic of threats is located at: <http://project.honeynet.org/papers/enemy/>

Risks:

The risks generally revolve around acknowledgement of your own vulnerabilities. If you have a system that requires maintenance, this asset has a risk value associated with it. Just by connecting to a network, risk is introduced. Many people have difficulty quantifying risk. On the proceeding slide I will present a formula that you may choose to use to help you quantify your risk.

Cost:

We are not just talking about how much a compromised computer is worth at resale. Your most valuable assets will be the people and their skills, as well as proprietary or intellectual property and trade secrets. The cost of compromise will vary. A file and application server with corporate accounting and payroll information might be very valuable while a customer service representatives terminal might be of low value. Consider what resources are accessible and what the impact or cost might be if this were to be compromised or destroyed.

When performing a Threat Risk Assessment (TRA) it is often helpful to categorize threats, risks, and costs. The following URL provides an example:

<http://www.acusafe.com/Security/AcuTech%20Intentional%20Threat%20Matrix%201-8-02.PDF>

Insurance on intellectual property:

<http://www.ipsearchengine.com/index.asp?from=&id=05&opt=17>

Use of our resources to attack others:

http://www.simpsonthacher.com/FSL5CS/articles/articles711.asp#_ednref5

Slide 7

Quantifying Risk

$$\text{Action Priority} = \frac{\text{Solution Modifier} = 1000 \times \% \text{Risk Tolerance} \times \% \text{Solution Ease}}{[\text{Risk} = \text{Threat} \times \text{Vulnerability} \times \text{Cost}]}$$

- action priority = lowest point value should be actioned first (higher point = lower priority)
- risk tolerance = constant for the organization (higher points = acceptance of more risk)
- solution ease = cost to implement solution (higher points = more man-hours, money etc.)
- cost = cost to recover from situation (higher points = higher cost)
- threat = odds of someone targeting this (higher points = greater threat)
- vulnerability = how dire is the vulnerability (higher points = bigger problems resultant in vulnerability)

Not every risk can be mitigated. You will have to identify your risk tolerance and probably accept some of the risks.

Reference and example:

http://www.whitehats.ca/main/members/Malik/malik_gcux_practical/malik_gcux_practical.html



Implementing a successful security strategy requires layers of security. There is only one silver bullet in computer security (<http://www.ranum.com/pubs/a1fwall/>). Aside from this, the more layers there are the more obstacles a potential miscreant has to overcome to achieve their goals. Defenses rely upon each other and can be compared in a “Siege the Castle scenario”.

Castle #1 only has a moat protecting it because most attackers cannot swim, therefore the castle is safe. Right?... Wrong!

Castle #2 has many defense-in-depth solutions implemented. It is not nearly as easy to breach because it has been built on high ground, with a moat, a sheer cliff on three sides, guard towers, low walls build around the exterior compound that require breach before arriving at the main castle gates and 30 foot stone walls etc.

Castle #2 would be a lot harder to breach than Castle #1 because all the defenses compliment each other. We try to follow the same logic with computer security, and this is the main supporting factor for why you would justify having an Intrusion Detection System.

Notice that Intrusion Detection Systems are only one of the many defense-in-depth layers!!!

Security policy is the first step in securing a network. This document should theoretically be created and adhered to prior to the design and implementation phases of a network. In reality this doesn't happen though because people don't realize they need a policy until they already have their network up and running. This is generally how it works. A good resource for some policy documents and

templates are available at:

<http://www.sans.org/newlook/resources/policies/policies.htm>

Many common vulnerabilities are identified in the SANS/FBI Top 20 List. This list is not necessarily technical in nature, but it does identify threats, hints at the risks your organisation might share, and provides some ideas of how to mitigate these risks. These ideas on how to mitigate or minimize risk are dealt with mostly through technical defenses which add to the defense-in-depth posture. The SANS/FBI Top 20 list is available at: <http://www.sans.org/top20/>

Access control incorporates proper ingress/egress filtering, segmentation of networks both internal and external for control, account based access, address or resource based access, physical access etc. These are implemented through Firewalls, router ACL's, Account management, hiring a security guard or using smart token authentication etc.

Time Based Security conceptually revolves around increasing your reaction time through defense-in-depth so that you may counter a threat before it succeeds. Mr. Winn Schwartau has written a book on the topic (not overly lengthy) and I recommend reading it if you're having issues understanding why defense-in-depth is important. Interpact Press; ISBN: 0962870048 and available on Amazon.com (search for the ISBN using the "More Search Options")

Redundancy is important. There are those who succeed and those without backups. If you have the resources to implement a few redundant systems you will very likely come to rely upon them at some point. Evaluating the value of the resource first is important. You wouldn't want to spend 80% of your companies capital installing diesel generators to cover extended periods of power outages unless there were very severe implications of such a failure happening (like in the health industry or financial industry with large institutions).

Evidence/Audit trail. Almost all the defense-in-depth procedures followed can produce some type of log or audit trail. These are very useful in troubleshooting as well as identifying either attempts to breach security or actual successes. If the evidence trail is handled properly it can also be very useful in prosecuting an offender and trying to recover damages (though most organisations find this is not actually an economical scenario to follow through with as it ends up costing them more in the long run than they recover).

Different types of IDS'

- Many different types
- Vendors try to differentiate product
- Underlying fundamentals
 - Detect that which is known (signature based)
 - Detect that which is unknown (anomaly based)
 - Employed on Network Traffic
 - Employed on a Host
 - A mix of these commonly referred to as a *HYBRID*
- Most other differentiation has to do with management of the IDS or other features

Signature Based IDS

- Match a string or pattern
- An engine that can apply regular expressions
- Applied by network based IDS' against network traffic
- Prone to many false positives
 - Triggers on a pattern detected that is not related to the signatures intended purpose
- Signature detection engines may perform more efficiently depending upon vendor
- Signatures can also be used in a more loose sense
 - Comparison done on traffic and compared to expected result (some vendors consider this to be protocol analysis)

Signatures - Pros

- Attack is known and documented
- Support and correlation available, aids the analyst
- Can perform comparisons efficiently
- Not difficult to understand underlying detection methodology
- Signatures can be created quickly
- Signatures can be written for many NIDS products by the customer
- Can compartmentalize IDS infrastructure based upon signature types
 - Useful in load balancing
 - Useful for management of sensors

Signatures - Cons

- Signature management can be very daunting
- False positives are common, waste time
- Not all vendor products allow customer to generate custom signatures
- Reliance upon vendor adds a middleman between you and the protection you require
- Many methods of evading known signatures
 - Polymorphic shell code, Insertion, Exclusion, Fragmentation, TCP un-sync, Low TTL, Max MTU, trust exploitation etc.
- Poorly written signatures can overwork NIDS
- False negatives occur
 - Malicious traffic passes the NIDS for which there is not a signature and it is not identified

Anomaly Based IDS

- Identify what is normal, everything else determined to be abnormal
- Many different methods of implementation
 - Protocol usage
 - Time of day usage
 - Association usage
 - CPU usage
 - Anything that can be measured as a deviation from a known state

Anomaly Based - Pros

- Find new attacks and methodologies
- Can find attacks that are intended to evade signature based IDS'
- Great freedom in definition of anomaly
- No signatures (per say)
- Can be measured using mathematics
- Can identify new hardware on network or network configuration changes quickly

Anomaly Based IDS - Cons

- How do you define what is normal?
- Prone to false positives
 - Tune anomaly definitions too tightly and it will identify non-anomalous activity as anomalous
- Prone to false negatives
 - Tune anomaly definitions too loosely and it will fail to identify truly anomalous activity
- Difficult to manage outside the realm of statistical analysis
- Learning curve for administrator is steeper
- Possible to evade anomaly based IDS by exploiting its known anomaly learning mechanisms or configurations

Host Based IDS

- Works beyond the network layer, actually implemented on the host and monitors state changes
 - File system
 - I/O calls (kernel hack or shim)
 - Application layer
- One-to-one relationship (covers only the host it is employed upon)
- Generally intended to be **Active**

Host Based IDS - Pros

- Application layer monitoring allows for inspection of data after decryption
- Can be finely tuned
- Much lower false positive rate
- Active response actually tries to stop attack
- Management is usually less complicated at host level

Host Based IDS - Cons

- Deployment and centralized management issues
- Cost can be prohibitive
- May not be deployed on every host (i.e. OS390 Mainframe)
- Self imposed denial of service
- Learning cycle can introduce complexity

Network Based IDS

- Monitors network traffic traveling on the network
- Works between layers 2 through 7 of OSI model (yes but not really – more like an emulation of coverage of the higher layers)
- Numerous methods of deployment
- One-to-many relationship (one NIDS can cover many hosts)
- Generally intended to be **Passive**

Network Based IDS - Pros

- More cost effective than HIDS
- Greater scope of coverage than HIDS
- Identifies attacks that might be dropped before reaching the target – good for quantifying threat
- Generate meaningful statistics
- Can help troubleshoot network problems
- Promiscuous in most deployment scenarios, not readily identifiable by the attacker

Network Based IDS - Cons

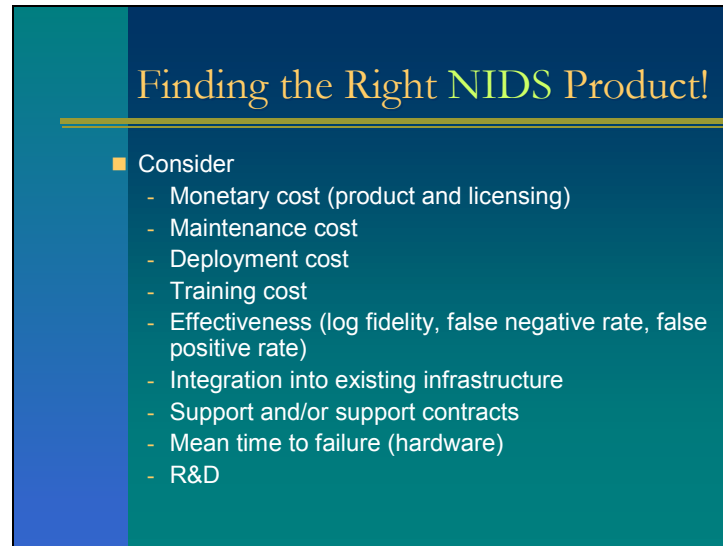
- Unauthorized activity local to a host will not be monitored
- Network load may cause fail-open conditions
- Can be difficult to manage
- Poor management yields large numbers of false positives
- Requires analysts to have deeper understanding of exploits – costly and time consuming

Finding the Right HIDS Product!

- Consider
 - Monetary cost (product and licensing)
 - Maintenance cost
 - Deployment cost
 - Training cost
 - Effectiveness (log fidelity, false negative rate, false positive rate)
 - Integration into existing infrastructure
 - Support and/or support contracts
 - R&D

Finding the Right HIDS Product!

- Application inspection (how is it implemented)
- Security of HIDS product (is the product itself secure)
- Signature maintenance and updates
- Back end/front end intuitiveness
- Produces Eye Candy for management
- Management mechanism (local or centralized?)
- Notification mechanisms (after hours)
- Active response mechanisms
- Plugins for in-house mods
- Session reconstruction
- Visualization tools



Finding the Right NIDS Product!

- Consider
 - Monetary cost (product and licensing)
 - Maintenance cost
 - Deployment cost
 - Training cost
 - Effectiveness (log fidelity, false negative rate, false positive rate)
 - Integration into existing infrastructure
 - Support and/or support contracts
 - Mean time to failure (hardware)
 - R&D

Monetary cost:

Simply – how much money do you have to spend?

Effectiveness:

Does the NIDS add value or is it a value sucker? Make sure you can do something with the logs it generates and that it is managed and configured properly or you'll be wasting your time. Some NIDS have low fidelity logs that do little more than tell you the TCP quad (source IP and port, destination IP and port) with a timestamp. Strongly consider another product if this is the case!

Integration into existing infrastructure:

Can this product integrate into what your already employing. An example might be Snort and Checkpoint FW1 or CISCO Secure IDS and CISCO Routers. This might also lower the training curve and make the product more useful.

Support:

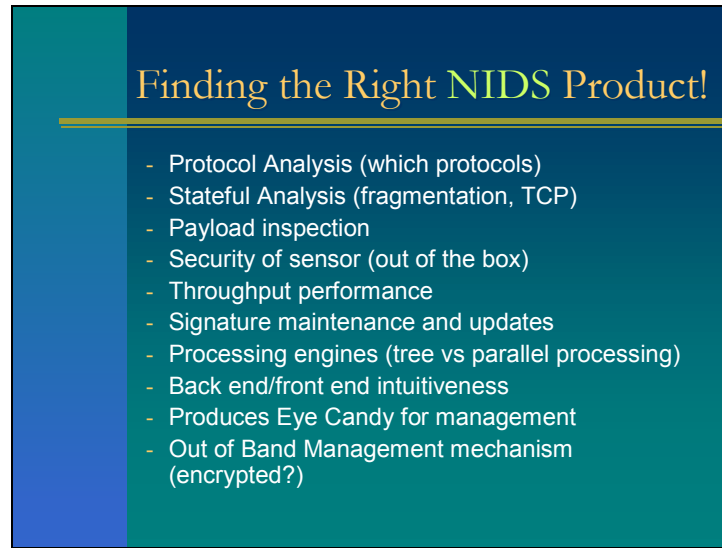
Call up the support number with a question (either a real one or a bogus one) and test their skills. If they can't help you or are reading from the manual then value this accordingly. If support is important to you, consider another vendor or discuss the problem with the vendor, offering them a chance to explain what transpired (and then test them again when they least expect it...).

Mean time to failure:

Some NIDS solutions come packaged up in a nice shiny box that your not supposed to open up. If this is the case, what type of hard drives to they use? What production series were these from? How many defects were noted? What is the operational life expectancy of the product before it fails. Same goes for other components like the power supply, cooling fans etc. How long before the vendor can have a replacement on site and do they cover the shipping, etc.

R&D:

Research and development – Does the company or vendor seem to be riding the wave or are they looking forward and coming out with new and exciting advances that you may be able to benefit from. What percentage of revenue is spent on R&D? How much capital does the company have? How big is the user base. Basically, is the company going to be there tomorrow? If you have doubts, consider another product.



Finding the Right NIDS Product!

- Protocol Analysis (which protocols)
- Stateful Analysis (fragmentation, TCP)
- Payload inspection
- Security of sensor (out of the box)
- Throughput performance
- Signature maintenance and updates
- Processing engines (tree vs parallel processing)
- Back end/front end intuitiveness
- Produces Eye Candy for management
- Out of Band Management mechanism (encrypted?)

Protocol Analysis:

Can be as simple as looking at protocol field values for inappropriate values (crafted) or looking at field sizes to see if they are inline with what is expected. Basically, working with the RFC's and comparing captured traffic matching that protocol with the protocol specifications to look for non-conformance and anomalous activity.

Stateful Analysis:

Putting a session back together and analyzing the sequence numbers, payload lengths, as well as state information in the layer 3 and 4 headers (generally) prior to piece together a complete message before analyzing it.

Security of sensor:

Make sure you don't connect a machine to your network that has security vulnerabilities itself. Having your security infrastructure compromised is a bad thing!!!

Throughput performance:

If you have a Gigabit Ethernet connection, you must make sure the NIDS can handle the traffic throughput under full load. If not, when you come under attack at full load, your NIDS will not be reliable and depreciates in value very quickly.

Signature updates:

Optimally, make sure you can write your own signatures and employ them. If not, consider how long it takes for the vendor to release updates and decide if this is an acceptable risk to take (waiting for them to produce something and distribute it) in the face of a zero day exploit that affects your organization.

Processing engines:

Complicated topic. There are different methods of processing. Some will check packets against a flat table of signatures. Slow because the last match might be hit but every rule had to be checked. Some categorize based upon protocol or other criteria. This is faster. Others branch off using logic trees and exponentially increase the processing speed. Consider also whether the underlying mechanism uses something such as PERL or is written and compiled in C or C++. Still even further, does it use shared libraries or static libraries, still further yet, can the process or daemon be optimized to spawn more threads as activity increases or assign a higher priority to the process?

Back end/front end intuitiveness:

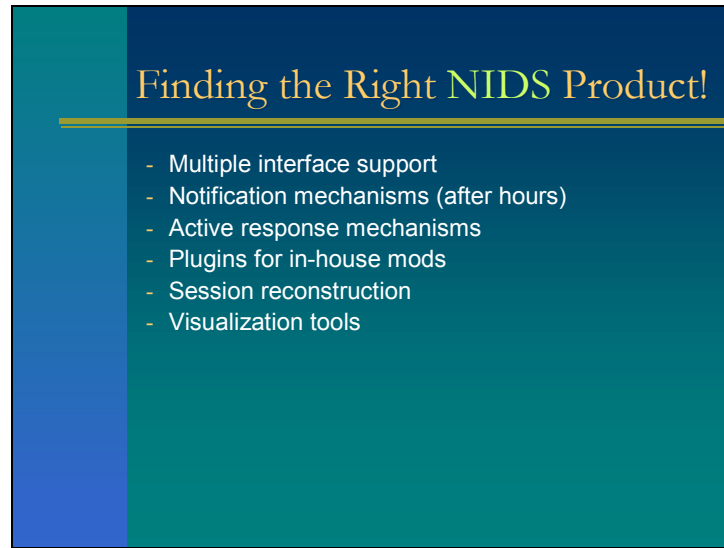
Is there a user friendly GUI or a logical method of retrieving the information you collect? Can you use it effectively or is it inefficient and you are left constantly asking "why am I wasting my time with this? If they (the vendor) would just do it this way...". Conversely, is it easy to manage the sensor or is it all based upon flat text files with very little documentation? Some people prefer to operate this way so this might be considered a benefit (deeper understanding without obscurity)

Produces Eye Candy:

Can the product generate some reports that can be provided to management or decision makers. Can management understand the information presented? A feature that is very important! If not, do you have the technical knowledge to generate some yourself?

Out of Band Management:

How do you remotely manage sensors? Encrypted via SSL and a web browser, via SSH, or a proprietary method? Using a standards based approach with publicly released algorithms is generally considered to be the best approach. This means the algorithm has withstood a lot of testing by the brightest cryptanalysts in the field and has survived! Hint – this means it is strong and you should consider using it.



The slide has a dark blue background with a vertical teal bar on the left. The title 'Finding the Right NIDS Product!' is written in a gold-colored serif font. Below the title is a horizontal gold line. A list of features is presented in white text with bullet points.

Finding the Right NIDS Product!

- Multiple interface support
- Notification mechanisms (after hours)
- Active response mechanisms
- Plugins for in-house mods
- Session reconstruction
- Visualization tools

Notification mechanisms:

Pager, email, auditory, 24/7 operations Security Operations Centre (SOC) etc.

Active response:

Session sniffing, firewall rule modification, run an application of your choice etc.

Plugins:

Redirect sniffed traffic into an application for further processing, run custom signatures against traffic, compare it against something you've defined etc.

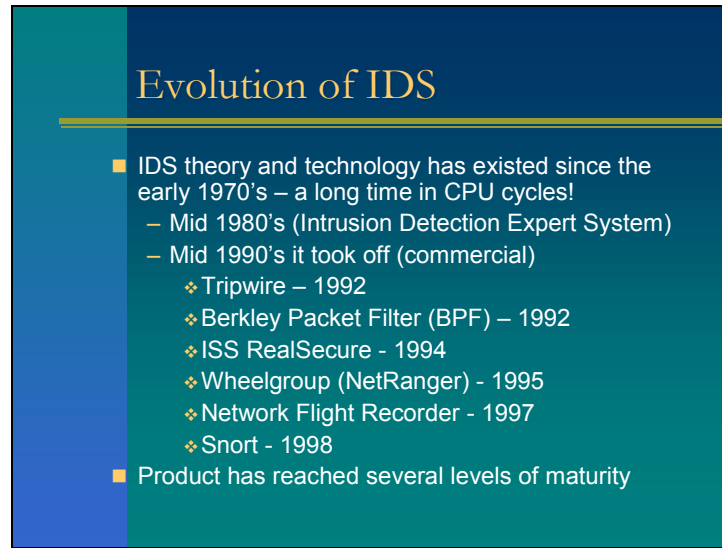
Session reconstruction:

Put the traffic back together and replay it to see exactly what transpired (eg. Reconstruct an HTTP session and see exactly what the web server presented to the suspect of anomalous use)

Visualization tools:

http://www.whitehats.ca/main/publications/link_list/link_list.html

Tools -> Graphical Data Mining/Modeling



Evolution of IDS

- IDS theory and technology has existed since the early 1970's – a long time in CPU cycles!
 - Mid 1980's (Intrusion Detection Expert System)
 - Mid 1990's it took off (commercial)
 - ❖ Tripwire – 1992
 - ❖ Berkley Packet Filter (BPF) – 1992
 - ❖ ISS RealSecure - 1994
 - ❖ Wheelgroup (NetRanger) - 1995
 - ❖ Network Flight Recorder - 1997
 - ❖ Snort - 1998
- Product has reached several levels of maturity

[ANDERSON] James Anderson. Computer Security Technology Planning Study, October 1972. <http://seclab.cs.ucdavis.edu/projects/history/papers/ande72.pdf>

[DENNING] Dorothy E. Denning. An intrusion-detection model. IEEE Transactions on Software Engineering, 13(2):222-232, February 1987.

[BRUNEAU] Guy Bruneau. The History and Evolution of Intrusion Detection, October 2001. <http://rr.sans.org/intrusion/evolution.php>

Maturity Milestones:

Packet Filters – Promiscuous packet level inspection

Berkley Packet Filters – Method of breaking data up, reference pointers within a packet for further action

Signatures – Inspection of packet headers

Signatures – Inspection of packet payload

Signatures – Some fuzzy logic, metacharacters and wildcard searches using regular expressions

Protocol Analysis – Check traffic against known normal protocol behaviour

Active Response – Do something when anomalous traffic is noticed

Centralized management – Gather all logs together in one place

Centralized management – Manage signatures or anomaly criteria from a central location

Anomaly detection – Engines using basic statistical approaches to identify abnormal patterns of usage

Expert systems – Employ some learning intelligence to help the filters make decisions on what to do with new traffic

Visualization – Plot, chart, graph, and vector data using different identifying features to help identify patterns of misuse or anomalies quickly

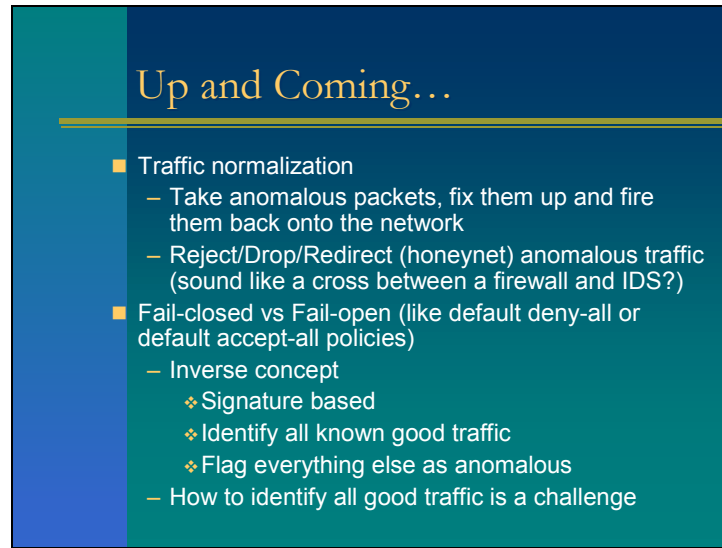
Standardization – Come up with a standard data format that is interoperable with other vendors formats

Classification systems – Classify exploits and anomalous activity under a common name for easy reference

Training – A more informed base of people who understand the issues surrounding Intrusion Detection

Up and Coming...

- Standardization
 - Log format
 - Interoperability of data/logs
 - IDMEF/IDWG
 - ❖ XML tags and objects
- Intrusion Prevention Systems (IPS)
 - Not new idea (buzz acronym for active response)
 - Integration with Firewall/Packet Filter (layer 3)
- Enterprise Security Management
 - Collecting logs from various devices in one spot
 - Normalizing logs (pre-standardization)
 - Centralized security management console



Up and Coming...

- Traffic normalization
 - Take anomalous packets, fix them up and fire them back onto the network
 - Reject/Drop/Redirect (honeynet) anomalous traffic (sound like a cross between a firewall and IDS?)
- Fail-closed vs Fail-open (like default deny-all or default accept-all policies)
 - Inverse concept
 - ❖ Signature based
 - ❖ Identify all known good traffic
 - ❖ Flag everything else as anomalous
 - How to identify all good traffic is a challenge

<http://www.icir.org/vern/papers/norm-usenix-sec-01-html/>


http://sourceforge.net/project/showfiles.php?group_id=22071&release_id=52521

Reject/Drop/Redirect – Why let the malicious or anomalous traffic through the front door. Tell them to go away. Similar concept to Firewalls except we are looking more deeply at the packet (see the emergence of a new acronym in the future? Something to cover what a product that meets a firewall and NIDS at ground zero will be called)

Up and Coming...

- Encryption
 - IPSEC et al.
 - Shared secret key (i.e. tcpdump with -E)
 - SSL proxy
 - Multiple gateways with NIDS between
 - More HIDS deployment
 - Statistical and aggregate scientific methods employed through analysis
- Covert Channels
 - Almost impossible to detect a well constructed one
 - Best indicator is an analyst that knows the network well (right now this is my feeling)

Slide 31



Conclusion

- Brief is available online at:

http://www.whitehats.ca/main/members/Malik/malik_ids_overview.htm

or

http://www.whitehats.ca/downloads/malik/malik_ids_overview.pdf

Thank you, merci!

Jamie French
j.french@whitehats.ca
Jamie.French@nms.ca