

An Empirical Approach to Measuring Defense-in-Depth (as a Cloud-based Service)

IFIP WG 10.4 - 65th meeting
Sorento, Italy
January 25, 2014



[Salvatore J Stolfo](#)
Columbia University
[Allure Security Technology](#)
[Red Balloon Security](#)

Sponsored by...

- Information Operations and Security Air Force Office of Scientific Research (AFOSR/RSL)
- AFOSR Contract FA9550-12-1-0162
- Title: *Designing for Measurable Security*
- "Any opinions, findings, and conclusions or recommendations expressed in this publication/presentation are those of the author(s) and do not necessarily reflect the views of the AFOSR."

Engineering Disciplines have measurement problems, too...



May 16th, 2013

Candidacy Exam

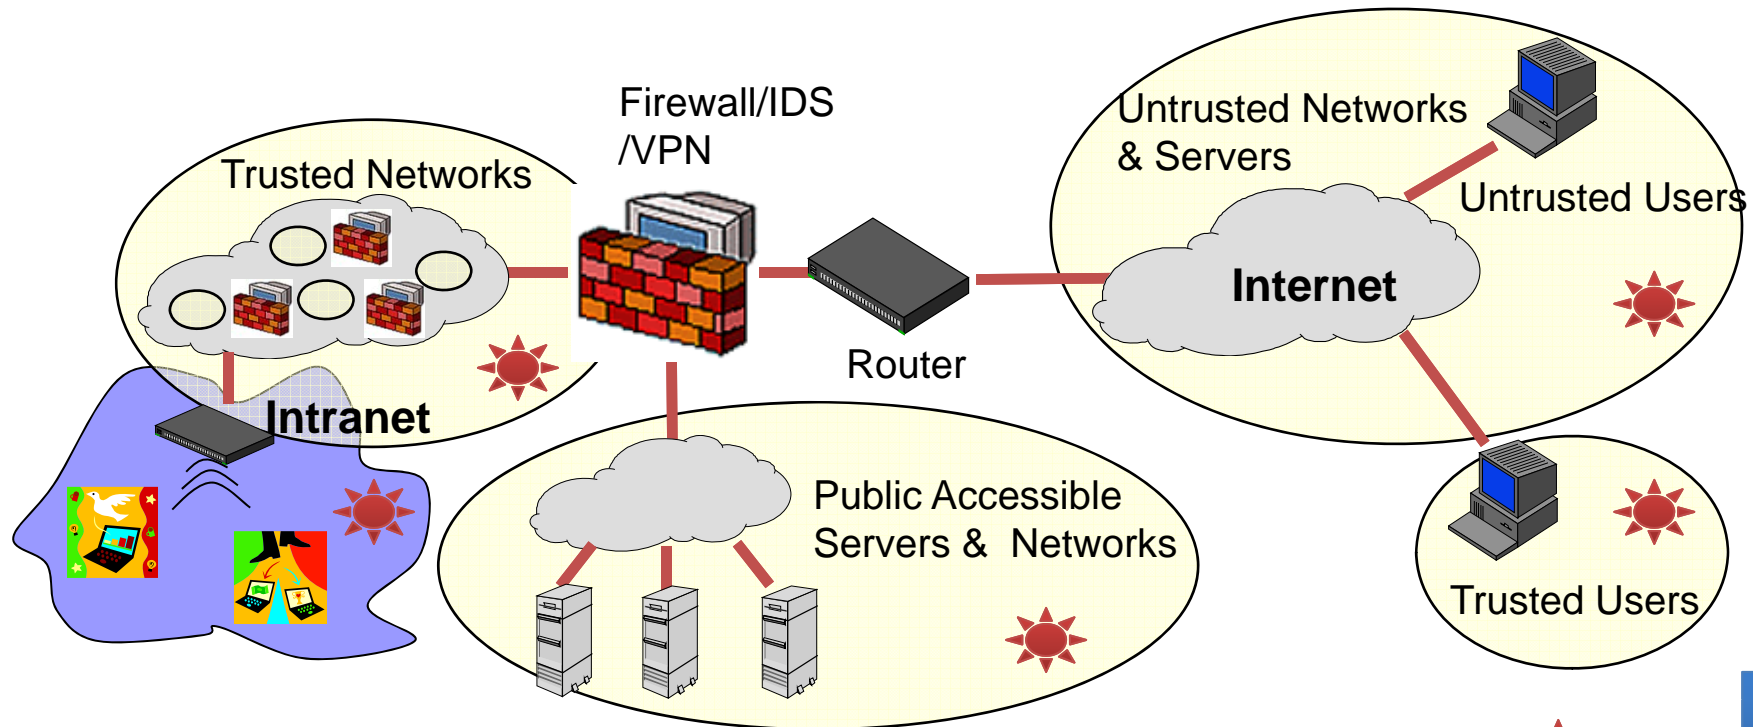
3

Secure Me


Many organizations have heterogeneous and distributed networks

What does security mean?

What are the challenges in *measuring* security properties?




Defense-in-Depth Principle... Layered Security



Prevent


Filter

Move the Attack Surface
Diversity/Randomization



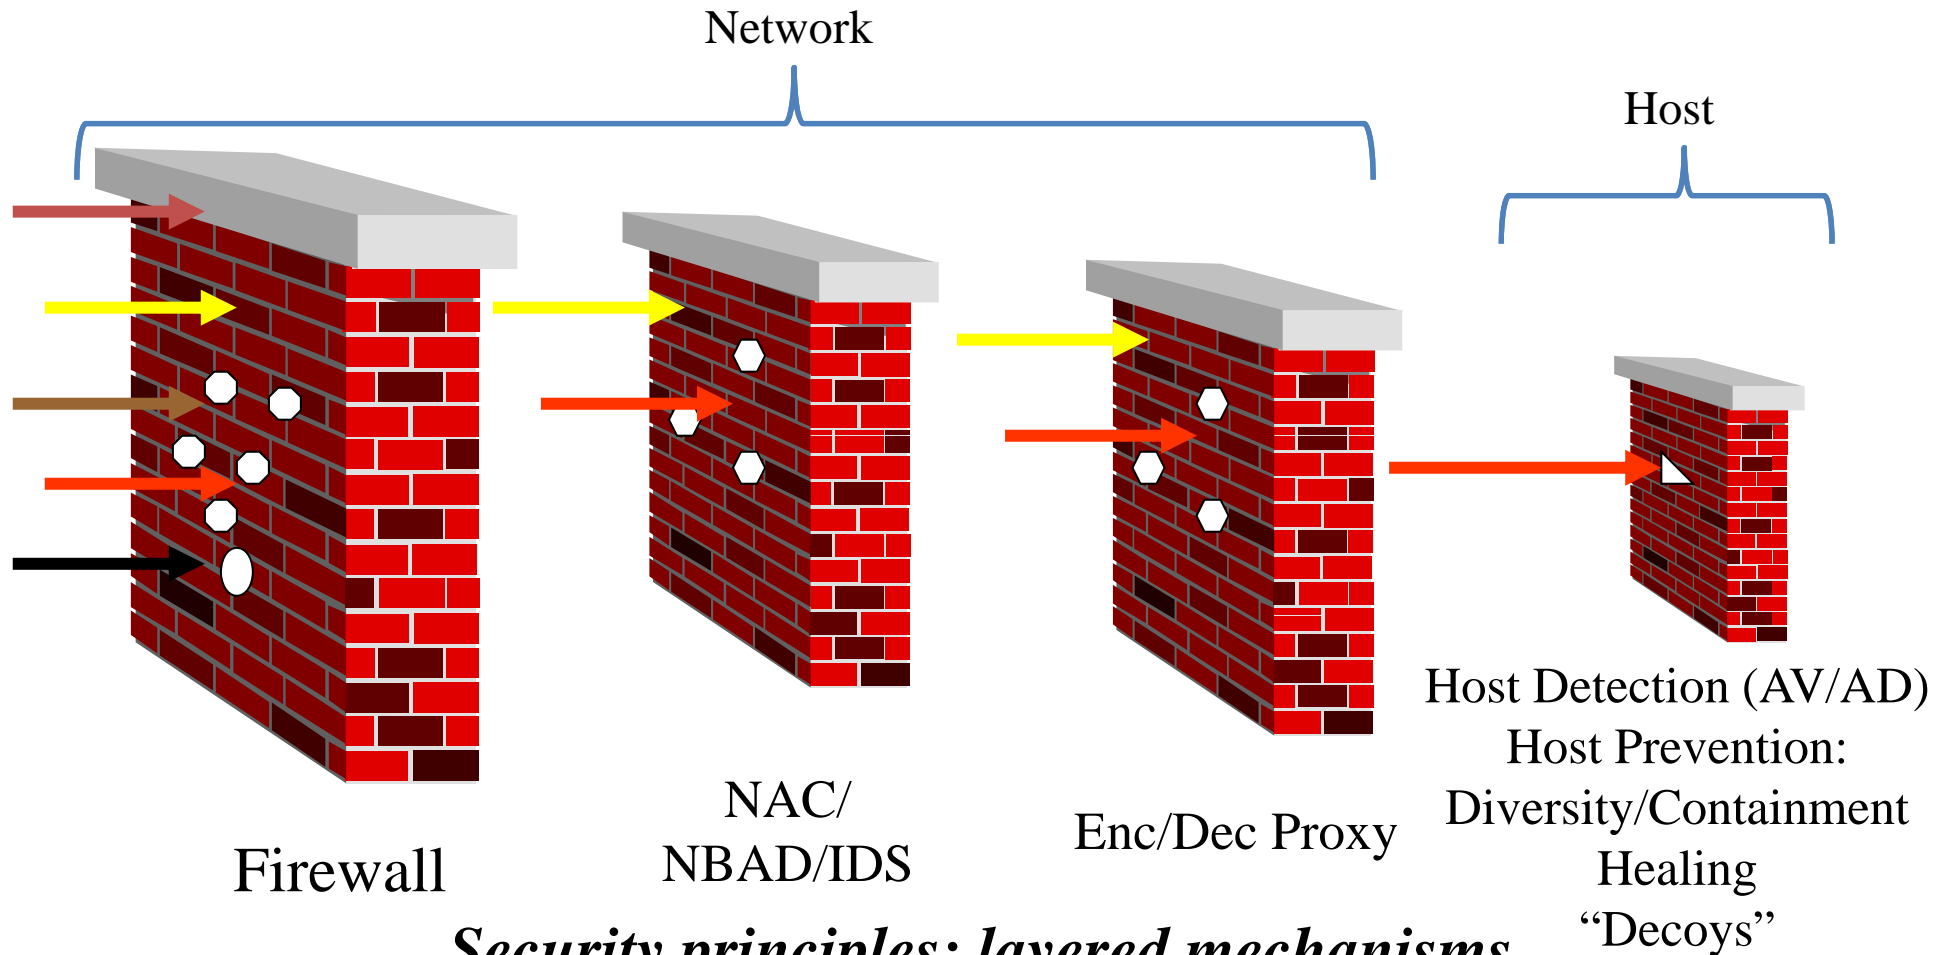
Detect

When Prevent fails
Sigs/Behavior



React/
Self-Healing
Survival of Mission

Defense-in-Depth In Practice



*Security principles: layered mechanisms
Can we measure effort to traverse each?*

Choosing a New Security Control

- Consider an organization contemplating investing in a new security control
- Basically the CISO needs to answer this:
 - *How secure* is my security architecture?
 - *Do I need another* security product?
 - Does a new security control *cost less* than the damage we would suffer from attacks that it may block?
- First step: see what the vendors say ...

“Industry leader” and the “world’s only.” Sounds like an easy choice.

WHY FIREEYE?

FireEye is the leader in stopping the new breed of cyber attacks, such as advanced malware, that easily bypass traditional signature-based defenses and compromise the majority of enterprise networks. FireEye's next-generation threat protection complements these defenses with the world's only signature-less solution that protects across all major threat vectors.

BUZZ

IANIS conducted a Return on Security analysis of the FireEye Web Malware Protection System (MPS). The estimated average net total return is \$16.5M. This represents a positive three-year return on security for FireEye, including the


NEWS EVENTS RESOURCES


4/30/13 - Dark Reading
[Chinese Cyberespionage: Brazen, Prolific, And Persistent](#)

4/26/13 - Public Service Europe
[Hacking All Over the World - Global](#)

What numbers?

View this email [online](#).

 McAfee
An Intel Company



Numbers Don't Lie

#1 in Exploit Protection
#1 in Evasion Protection

McAfee Leads in Protection Against Most Advanced Threats

McAfee is #1 in exploit protection and evasion protection announced by NSS Labs in a recent publication.

NSS Labs is one of the most respected third party test labs in the security industry. They are best known for their research and testing against modern threats.

Combined Detection Effectiveness

	Exploit	Evasion	Combined
McAfee	97%	100%	99%
Symantec	91%	100%	96%
Sophos	88%	97%	93%
Kaspersky	92%	92%	92%
F-Secure	79%	88%	84%
Microsoft	65%	100%	83%
AVG	76%	88%	82%
ESET	71%	92%	82%
Trend	73%	53%	63%
Norman	47%	75%	61%
Panda	41%	75%	58%

Jan 27, 2014

IFIP Meeting

9

These numbers make it easy. Just buy McAfee!

View this email [online](#).



Numbers Don't Lie



#1 in Exploit Protection
#1 in Evasion Protection

McAfee Leads in Protection Against Most Advanced Threats

McAfee is #1 in exploit protection and evasion protection announced by NSS Labs in a recent publication.

NSS Labs is one of the most respected third party test labs in the security industry. They are best known for their research and testing against modern threats.

Combined Detection Effectiveness

	Exploit	Evasion	Combined
McAfee	97%	100%	99%
Symantec	91%	100%	96%
Sophos	88%	97%	93%
Kaspersky	92%	92%	92%
F-Secure	79%	88%	84%
Microsoft	65%	100%	83%
AVG	76%	88%	82%
ESET	71%	92%	82%
Trend	73%	53%	63%
Norman	47%	75%	61%
Panda	41%	75%	58%

Jan 27, 2014

IFIP Meeting

10



But wait, Symantec is better!

Symantec. | Enterprise United States Shopping

Products & Solutions | **Support & Communities** | **Security Response** | **Try & Buy**

Products & Solutions / Symantec Endpoint Protection Family

Symantec Endpoint Protection Family

Computer Protection Software

Unrivaled Security. Blazing Performance.
Built for Virtual Environments.
Symantec Endpoint Protection.

Product	Relative Performance
Microsoft System Center Endpoint	Low
Trend Micro OfficeScan	Low-Mid
McAfee VirusScan	Mid
Kaspersky Endpoint Security	High
Symantec Endpoint Protection 12	Very High

OTHER RESOURCE

[Symantec Positioned Highest in Vision & Execution In Gartner's Magic Quadrant for Endpoint Protection Platforms](#)

REVIEW

[Dennis Technology Labs: Enterprise Anti-Virus Protection July-September 2012](#)

How can we evaluate which control is better?

The screenshot shows the Symantec Enterprise website. The main heading is "Symantec Endpoint Protection Family Computer Protection Software". Below this is a promotional banner with the text "Unrivalled Security. Blazing Performance. Built for Virtual Environments. Symantec Endpoint Protection." and a bar chart comparing Symantec Endpoint Protection 12 to other products: Microsoft System Center Endpoint, Trend Micro OfficeScan, McAfee VirusScan, and Kaspersky Endpoint Security. The Symantec bar is the tallest and is yellow, while the others are grey. Below the banner are two resource boxes: "OTHER RESOURCE" and "REVIEW".

OTHER RESOURCE

[Symantec Positioned Highest in Vision & Execution In Gartner's Magic Quadrant for Endpoint Protection Platforms](#)

REVIEW

[Dennis Technology Labs: Enterprise Anti-Virus Protection July-September 2012](#)

Measure What (1)?

- Security properties of a system?
 - *Absolute* metrics are unlikely – how do you count bugs?
How do you account for threats you aren't aware of?
- *Relative* security properties of two systems?
 - Differential metrics are feasible:
 - Penn testing makes sense if you know what to test for
 - Evaluate the severity of attacks each is not able to defend against
but
- The threat is dynamic and relentless so constant measurement is required

Some security metrics...see Bibliography

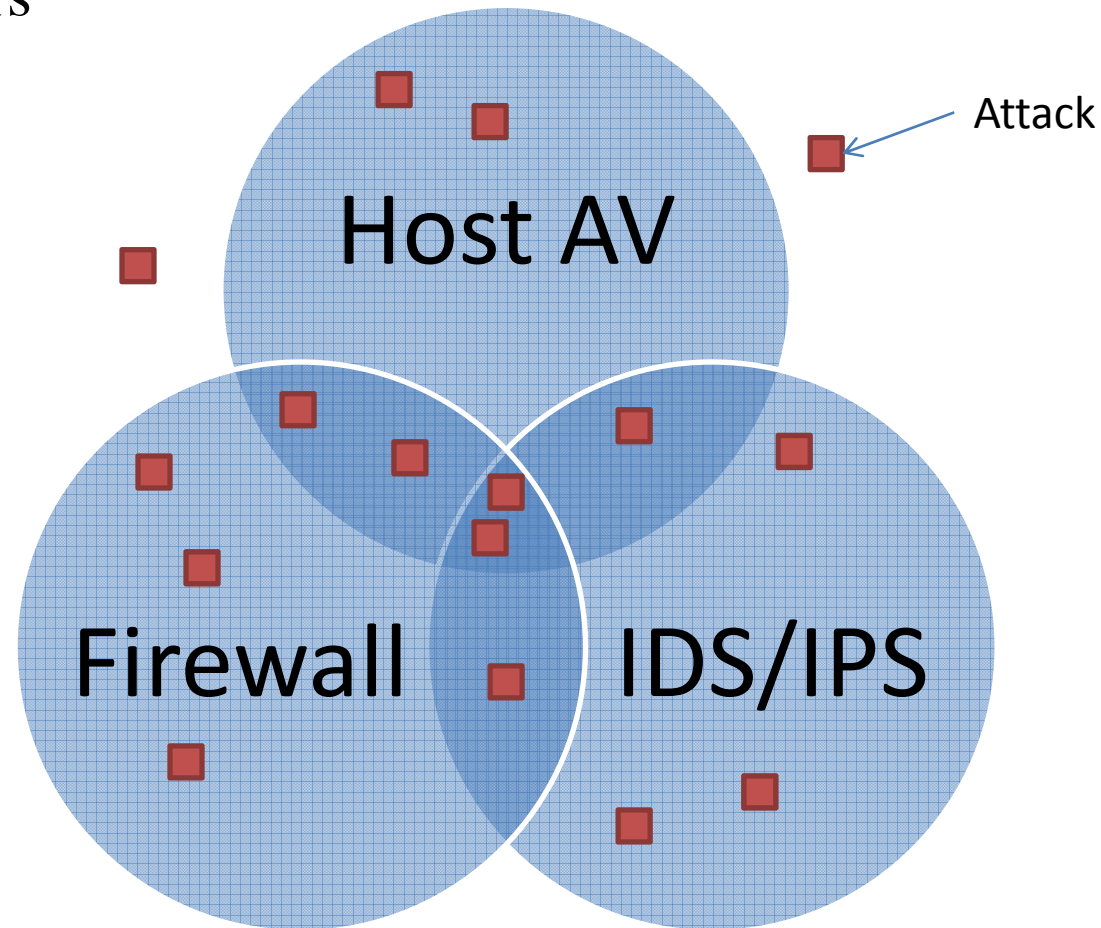
- Computational complexity and Entropy
 - Cryptology and hardness of problems
 - Automated diversity (ASLR, ISR, ...)
 - Warning: difference between mathematical abstraction and system implementation
- Economic/biological metrics
 - Cost-based IDS (stop loss/accuracy)
 - Polymorphic engine strength (variation/propagation)
- Empirical (structured experiments and testing of coverage)
 - Adversary models

How to Measure Coverage

- Goal: Measure total detection rate and find which security products best compliment each other rather than estimate the best overall or individual score
- Which attacks? Attacks change
- Ground truth issues
- Organizations have different management requirements


Is this the coverage of my Defense in Depth Architecture?

We assume layers
provide broader
coverage, better
security.




Or this?

What if they look
more like this?
We measure overlap
between products and
total coverage!



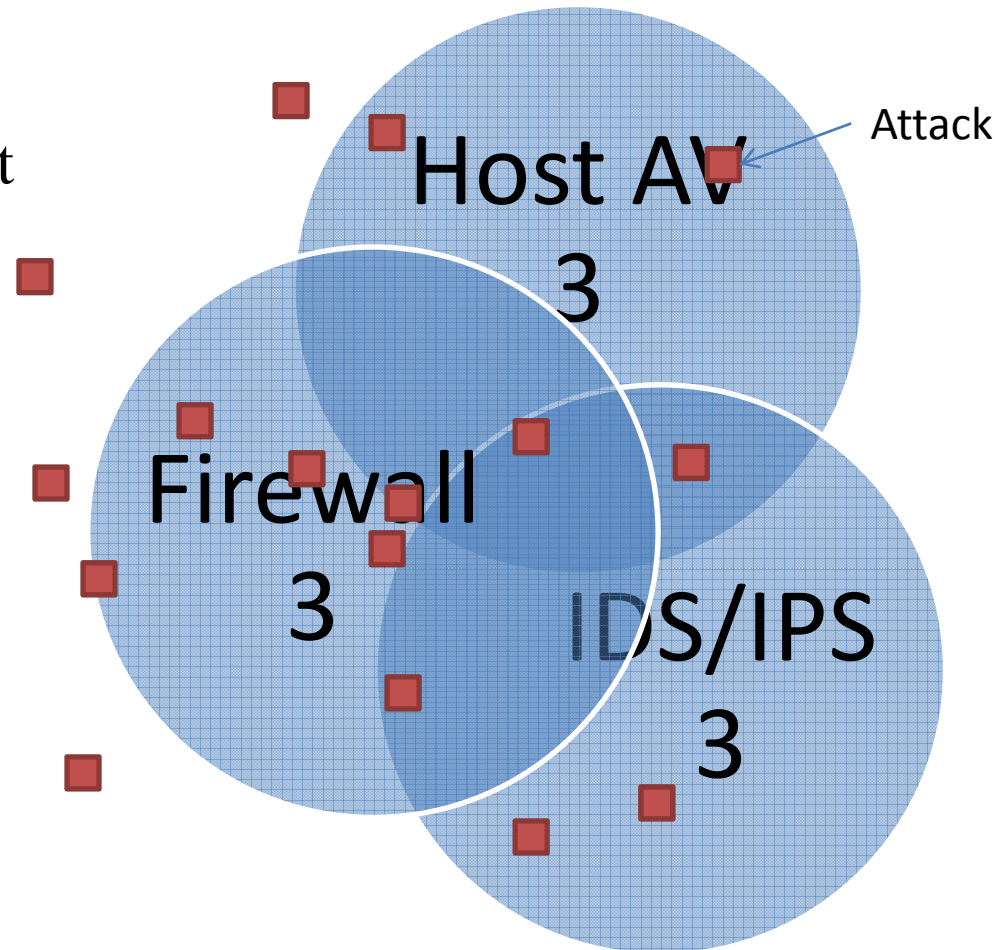
Or better yet, this?

Mutual coverage might vary significantly depending on which AV, IDS, Firewall is deployed.



Does it shift to this in time?

Best AV + Best
Firewall + Best IDS
might not be the best
mutual coverage
when combined!



Measure What? (2)

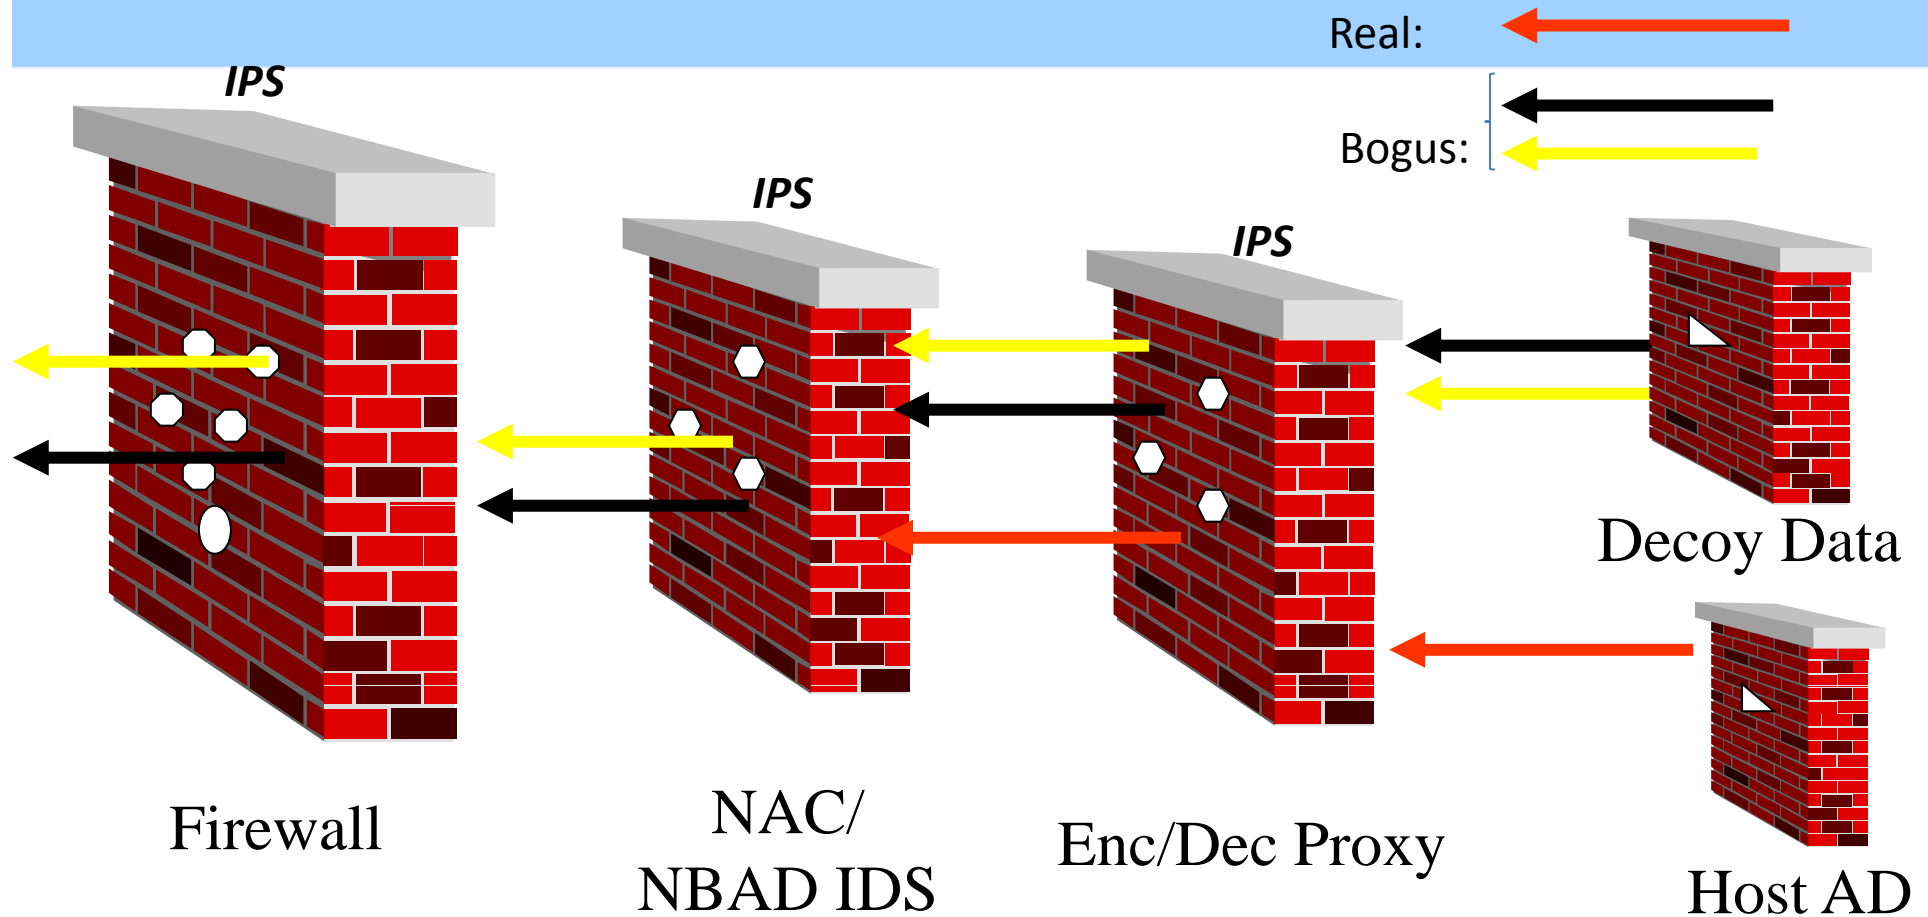
- Measure adversary effort (“invite” them and observe their success) to cross security layers?
 - Adversary effort is generally linear in number of layers
 - How do we design layers so that breaking through two layers is proportional to the product, not the sum of the adversary’s effort?

Solve this and WE win!

Measure What? (3)

- Measure adversary effort to exfiltrate data through each layer
 - Measuring amount of egress information conditioned on input
 - Reply with Decoy Data when threshold triggered between layers to “poison” exfiltrated data

Quantifiable Egress – Managed by layered IPS's

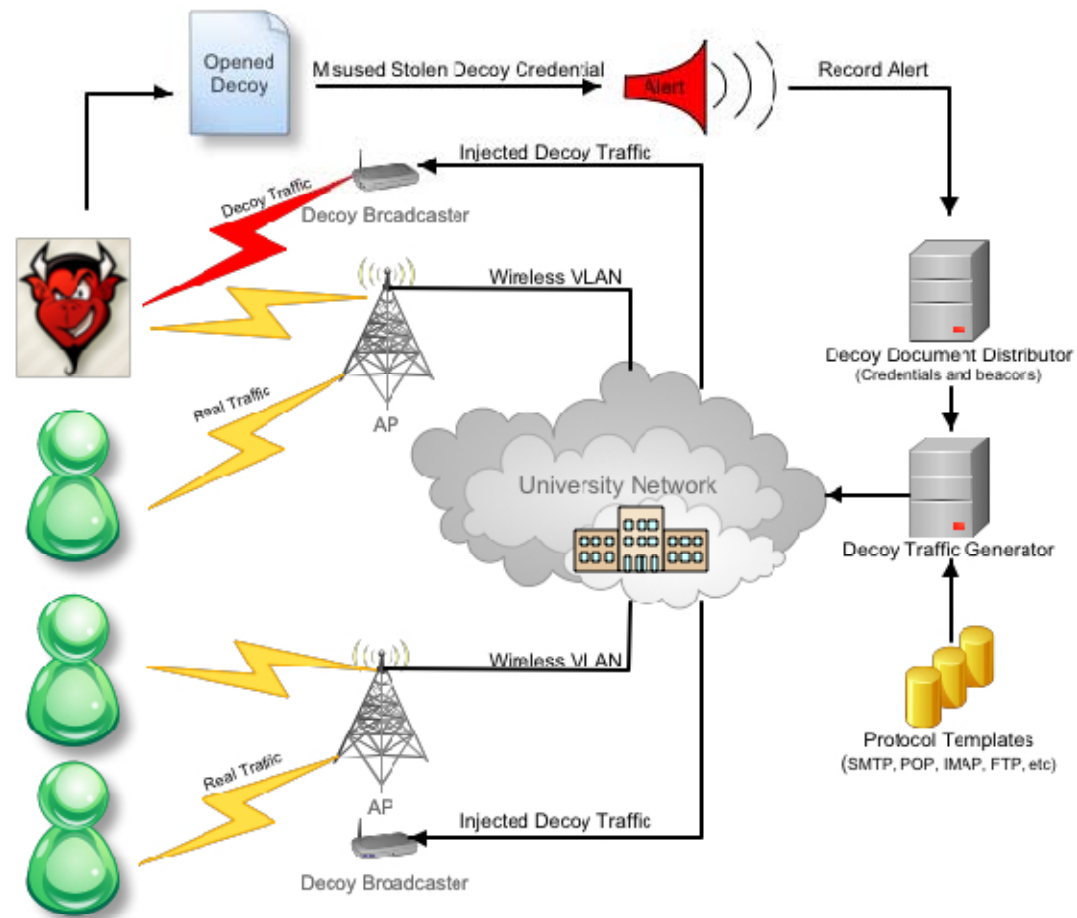


*Each Layer measured by an IPS for exfiltrated data:
Threshold logic rate limits real data and generates
DECOY DATA.*

Measure What? (4)

- Security posture of an organization
 - User violation testing!
 - Quantify number of violations
 - Longitudinal analysis: Am I getting better?
 - Measure repeat (and repeat-repeat) offenders after “training”

A “decoy” generation system for measuring user violations



Summary of Measurements

- Relative coverage of existing/changing architectures
- Adversary effort and cost to evade and penetrate layers
- The propensity of an architecture to leak data and how much
- *De Facto* policy violations by legitimate users (forget *De Jure*)

Some Literature on Measurement

Qualitative	Empirical	
	Cost to Evade	Detection Rate
Chapin05	Stolfo11	Cavusoglu04
Mateski12	Weir10	Ingham07
	Kelley12	Boggs11

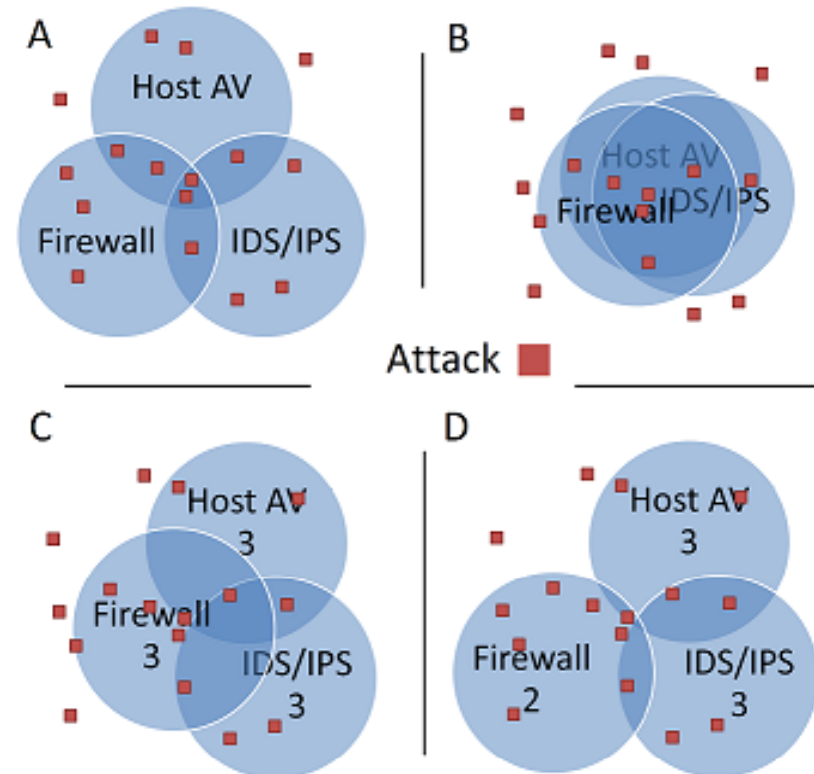
See Back up and Bibliography

Measure Layers – Boggs11

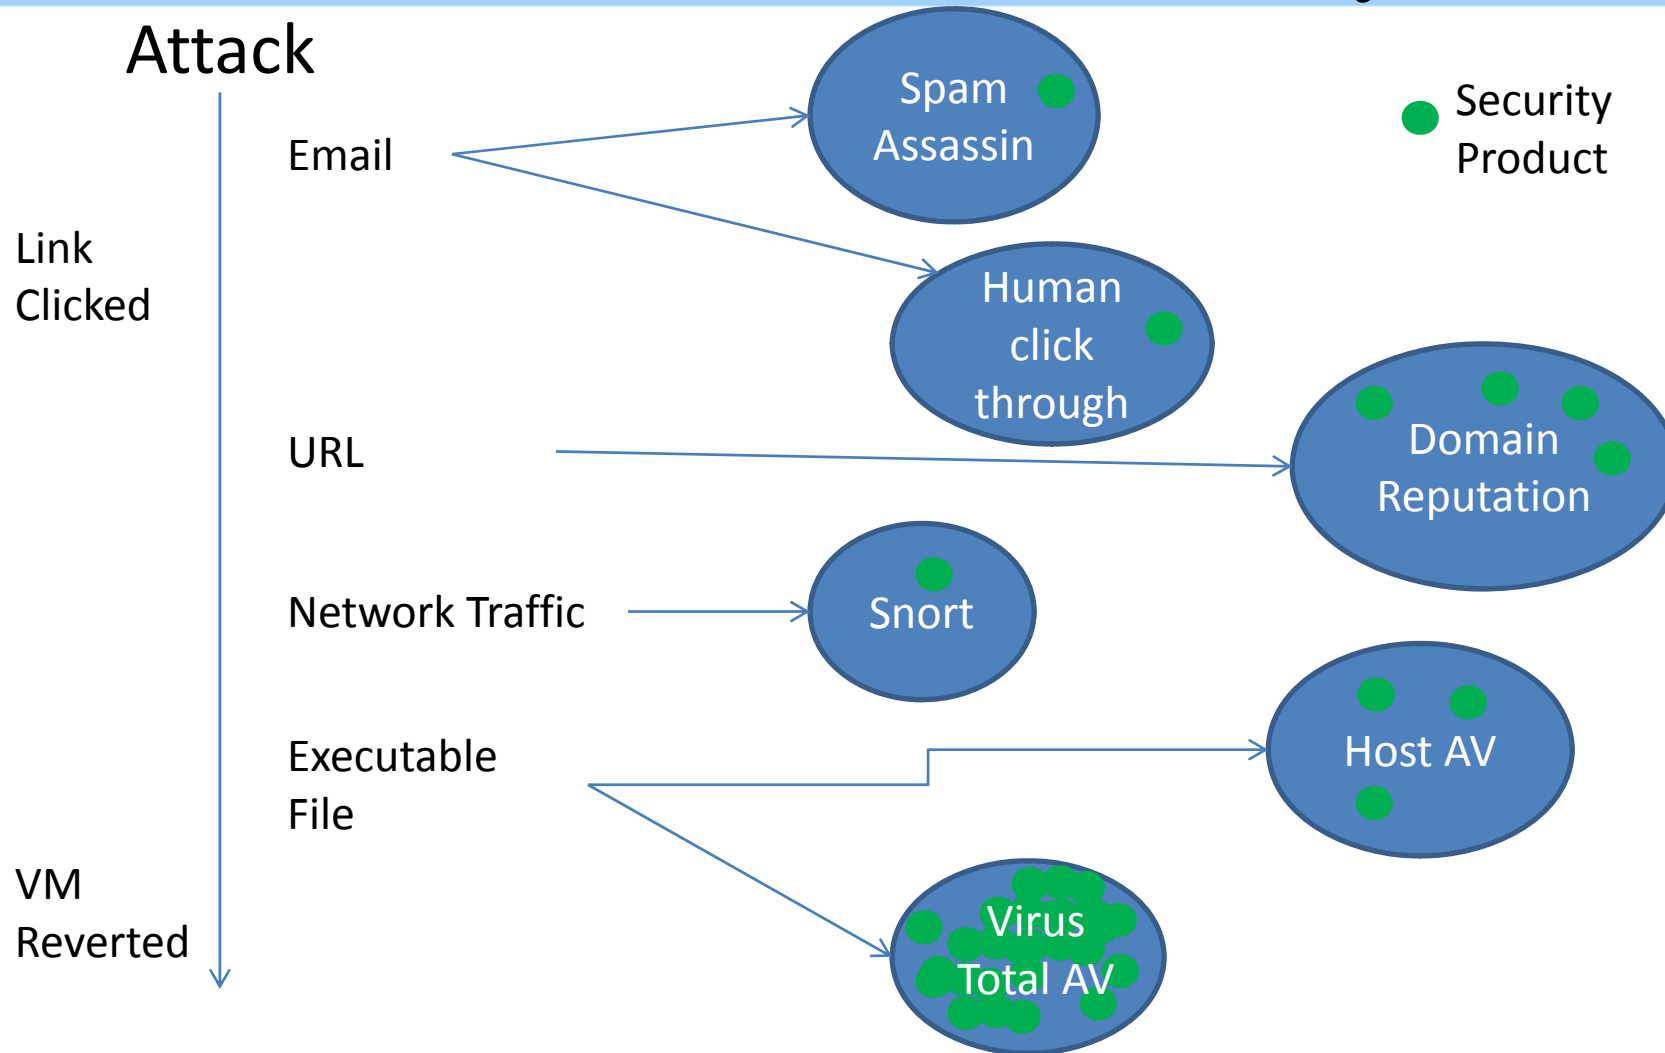
- Nathaniel Boggs, Salvatore J. Stolfo; "ALDR: A New Metric for Measuring Effective Layering of Defenses;" Layered Assurance Workshop; 2011/12/06

Measure Layers – Boggs11

- Empirical measurement of defense in depth
- Capture each attack across layers (record email, URL, network traffic, dropped executable file, etc.)
- Test multiple layers against same attack and track which controls detect each attack and at what layer
- Union the sets of attacks detected by a group of security products to determine total detection rate



Attack Data Scanned by Real Security Products at Different Layers



Collecting/Creating Attack Data

- Initial compromise attack vectors
 - Exploit clients, ex. drive-by download
 - Service exploits, ex. SQL injection
- Honeypots
- Use the same exploit kits
- Academia at a great disadvantage

Layers Tested

- Over 40 security products tested
 - Spam Assassin
 - 4 Domain reputation systems
 - Snort - Emerging Threat rule set
 - 3 Stand alone AV
 - 40 AV engines from VirusTotal (online scanning service)
 - Human click through

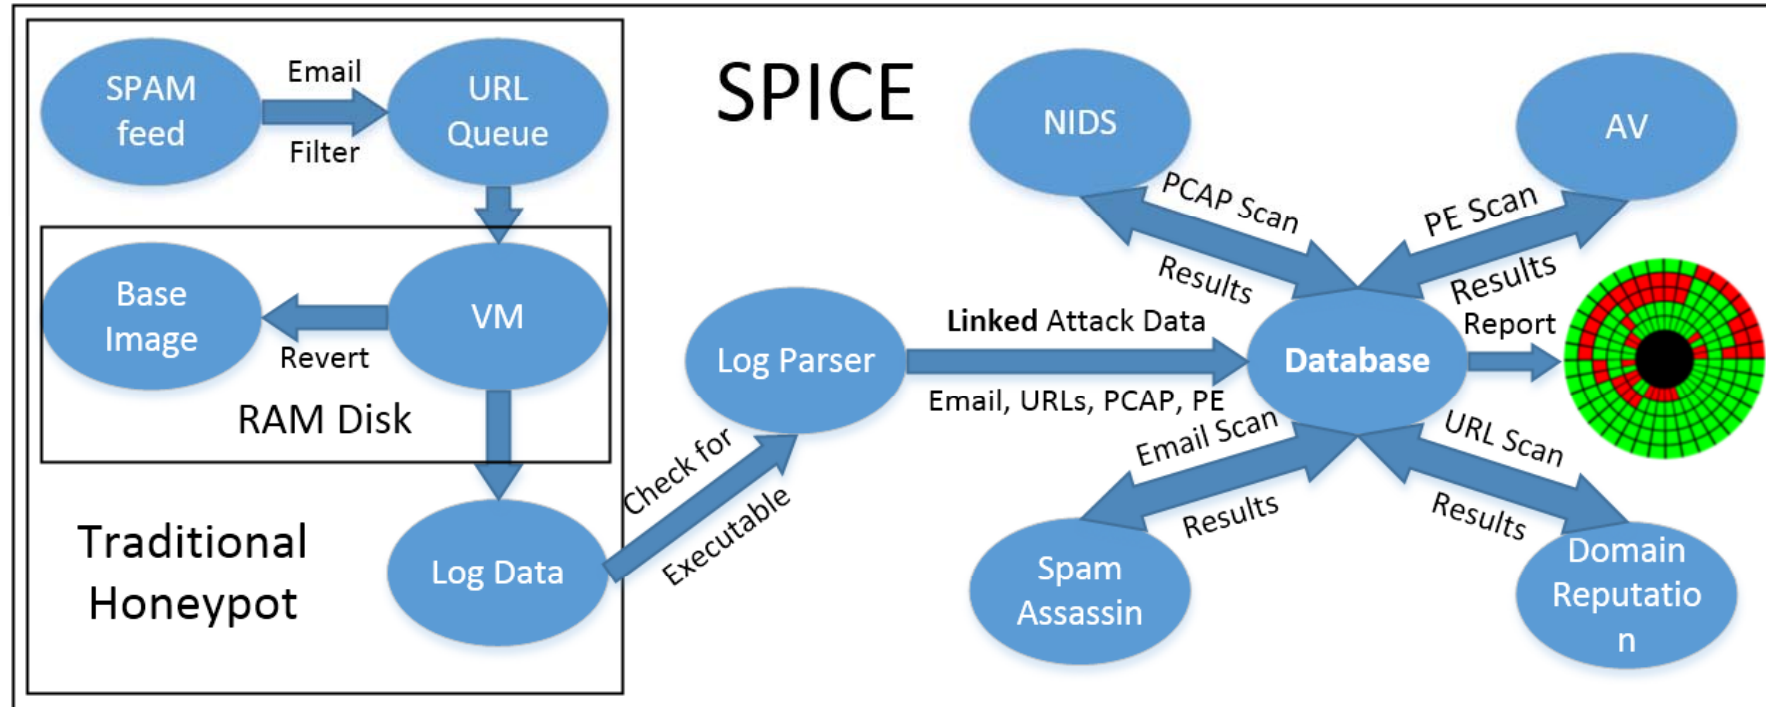
Experiment Details

- 4 VM setups (Acrobat Reader, Flash Player, Java, Firefox)
- Cuckoo Sandbox
- Scan emails (1 million/day)
- Send VMs to 1 link per unique domain
- Each VM setup visits link 3 times (mean of ~2 infections per malicious link)

Attack Data Collected

- 1463 malicious site visits by VMs ending in compromise
- 730 unique malicious emails
- 576 unique executables
- 36 clusters of distinct email content

SPICE Architecture




Security Posture Integration and Correlation Engine (SPICE)

- Practical prototype
- Report shows which layers detect which attacks


Real Results – Initial Detection

- Green is detected attacks
- Each concentric circle is one layer of defense
- Each arc is a different attack cluster



Real Results – Eventually Detected

- Green is detected attacks
- Each concentric circle is one layer of defense
- Each arc is a different attack cluster



Results – Findings

- Most security products are horrible
 - Mean detections: 11.3/36 clusters
- No security product is perfect
 - No single product detected all clusters
- With time most products can detect attacks
 - Eventually detected mean: 27.3/36 clusters

Example Using Real Data

- Assume a small organization with the best AV and best domain reputation seen in our experiment
- AV: detects 29/36 attack clusters
- Domain reputation detects 22/36
- Current state of the art

Example Using Real Data

- With our data we can go further!
- Together detect: 33/36
- What products detect the last 3 clusters?
- Snort detects 27/36 but more importantly 2/3 of the previously undetected attacks
- Spam Assassin detects 31/36 total and 3/3 of the previously undetected attacks
- Imagine zero day attacks, more layers, more security products tested, are you secure?

Our Approach's Key Attributes

- Products tested individually
- Expandable framework
 - Measure education benefit
 - Social engineering attacks
 - Any 'attack' representable
- Evaluate products in the context of existing layers of security rather than in absolute/isolated terms

Reducing High Cost of Acquiring Attack Data and a Set of Products

- Cloud service could amortize cost across many organizations
- Develop attack data sets centrally
- Test all security products centrally
- Custom report per organization based on their current security products
- Optional false positive reports based on their real data

Measure What? (5)

- Measure (empirical testing) the capacity for self-healing?
 - Time and accuracy of automated patch generation
- Measure the total amount of data possibly exfiltrated?
 - What is the adversary's effort to read the entire store and “cut CD's”

Summary

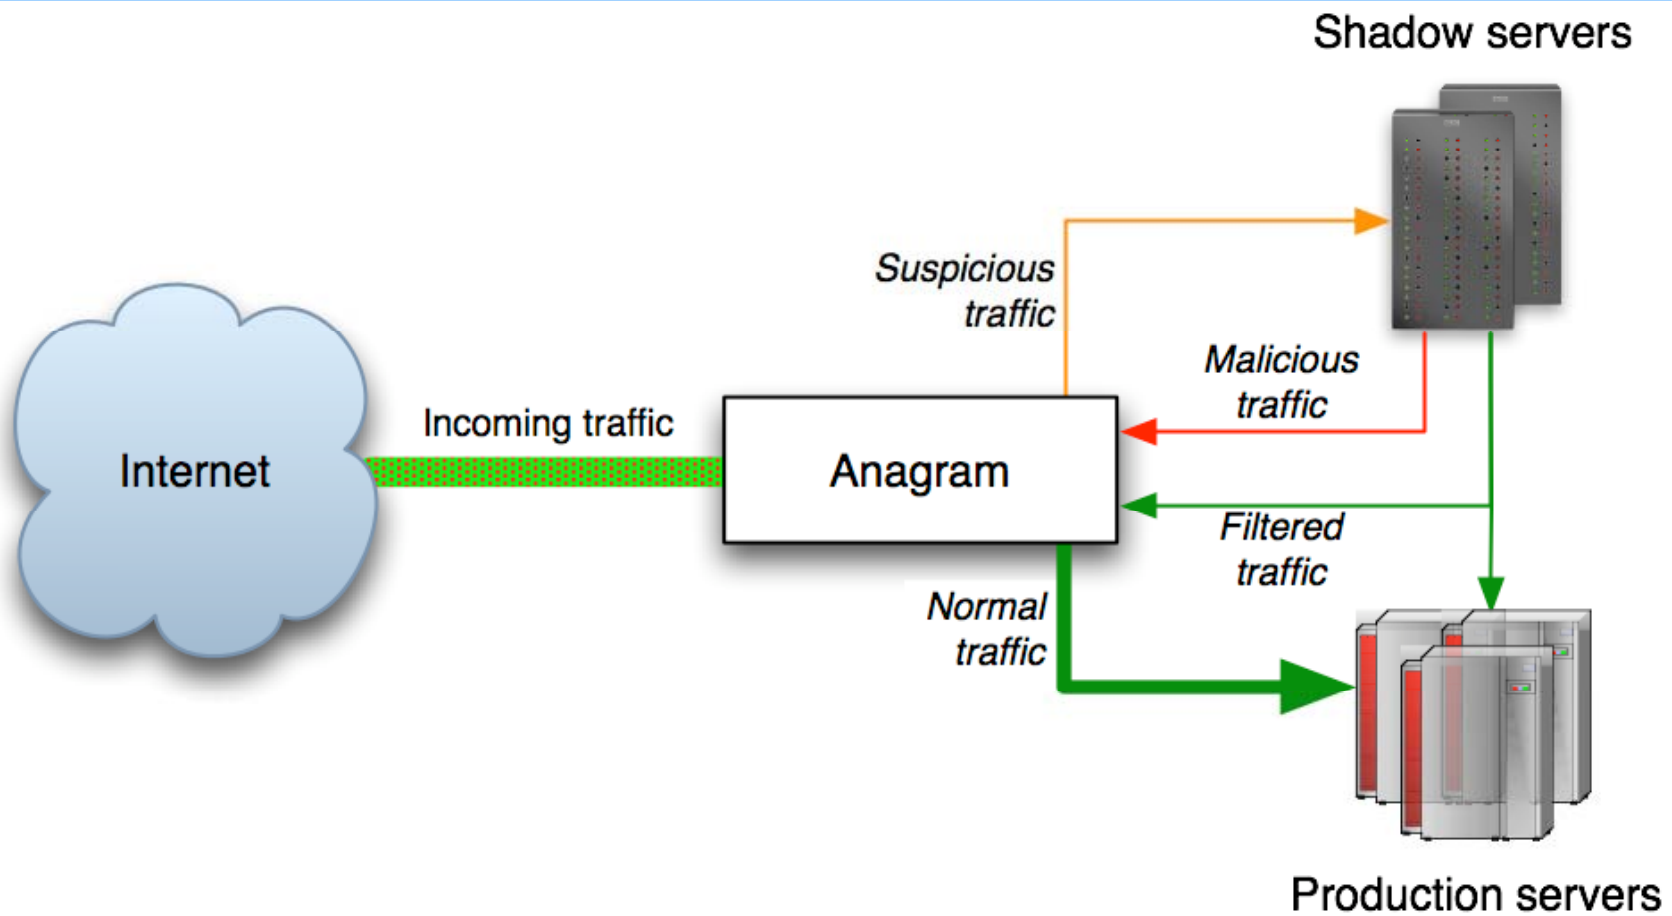
- Absolute metrics unlikely
- Relative metrics are likely feasible
 - Longitudinal analysis to compare one SUT
 - Compare two SUT's side by side using empirical tests
- Design systems/layers for Measurable Defense in Depth using relative metrics


Concluding Remarks – Continuous Measurement

- A cloud measurement service could answer a number of questions:
 - are we secure with certain products against different classes attacker?
 - false positives of those products given samples/real-time feeds of an organization's data
 - most complimentary security products and suggest additional redundancy to increase evasion cost to attacker
- Amortize costs across many organizations

Designing for Multiplicative Adversary Effort...

- How do we design layers so that breaking through two layers is proportional to the product, not the sum of the adversary's effort?
- An initial idea:
 - Evasion tactics to thwart content AD that require concurrent shaping and padding across multiple incongruent features may raise the adversary cost if
 - the features are independent – guessing for one provides no information about the other
 - Features are chosen randomly unbeknownst to the adversary – multiple guesses would be needed to succeed



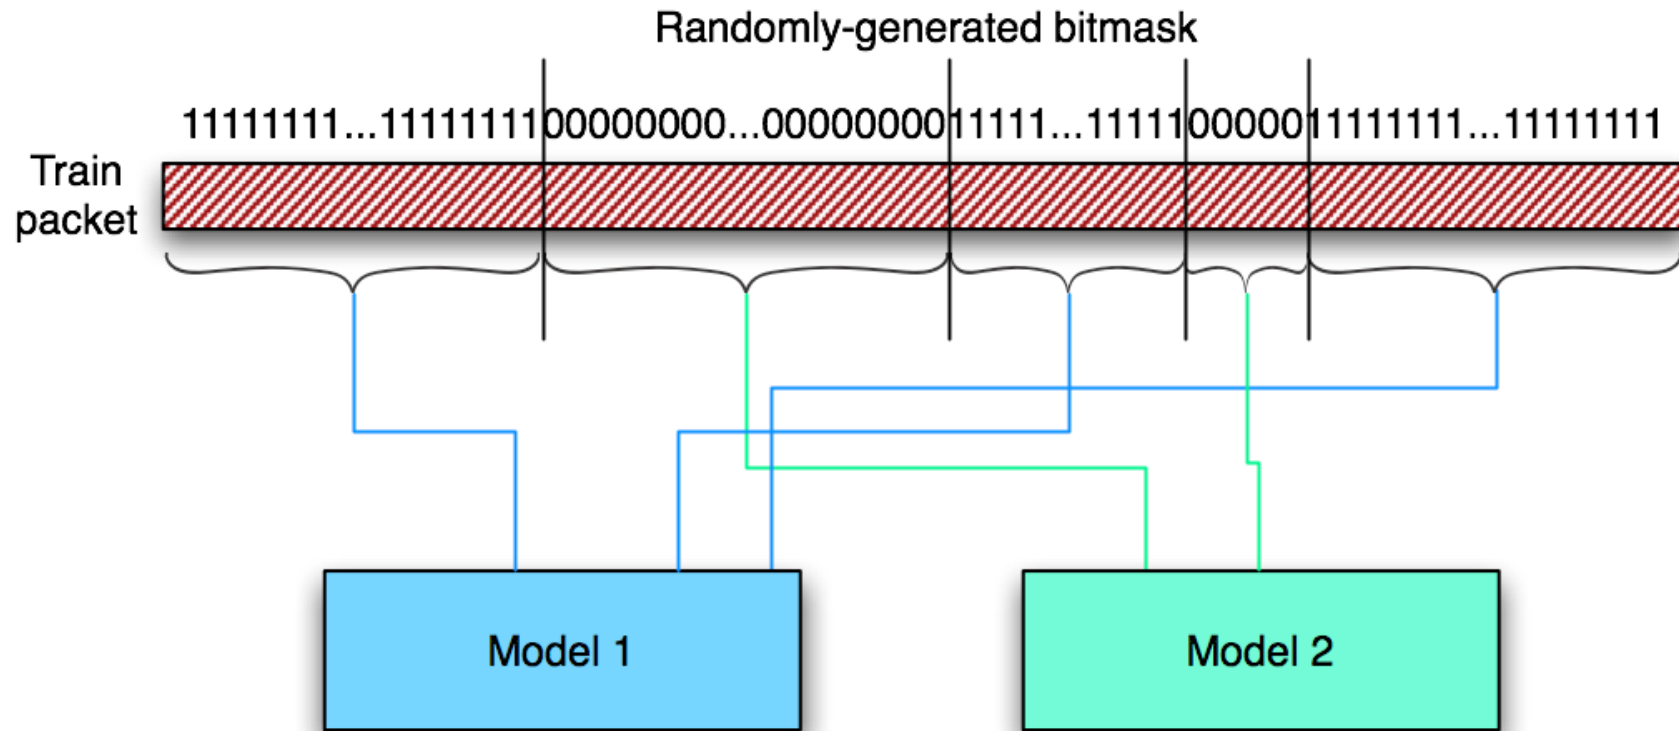


Randomization against mimicry attacks

- The general idea of payload-based mimicry attacks is by crafting small pieces of exploit code with a large amount of “normal” padding to make the whole packet look normal.
- If we *randomly choose the payload portion for modeling/testing*, the attacker would not know precisely which byte positions it may have to pad to appear normal; harder to hide the exploit code!
- This is a **general** technique can be used for both PAYL and Anagram, or any other payload anomaly detector.

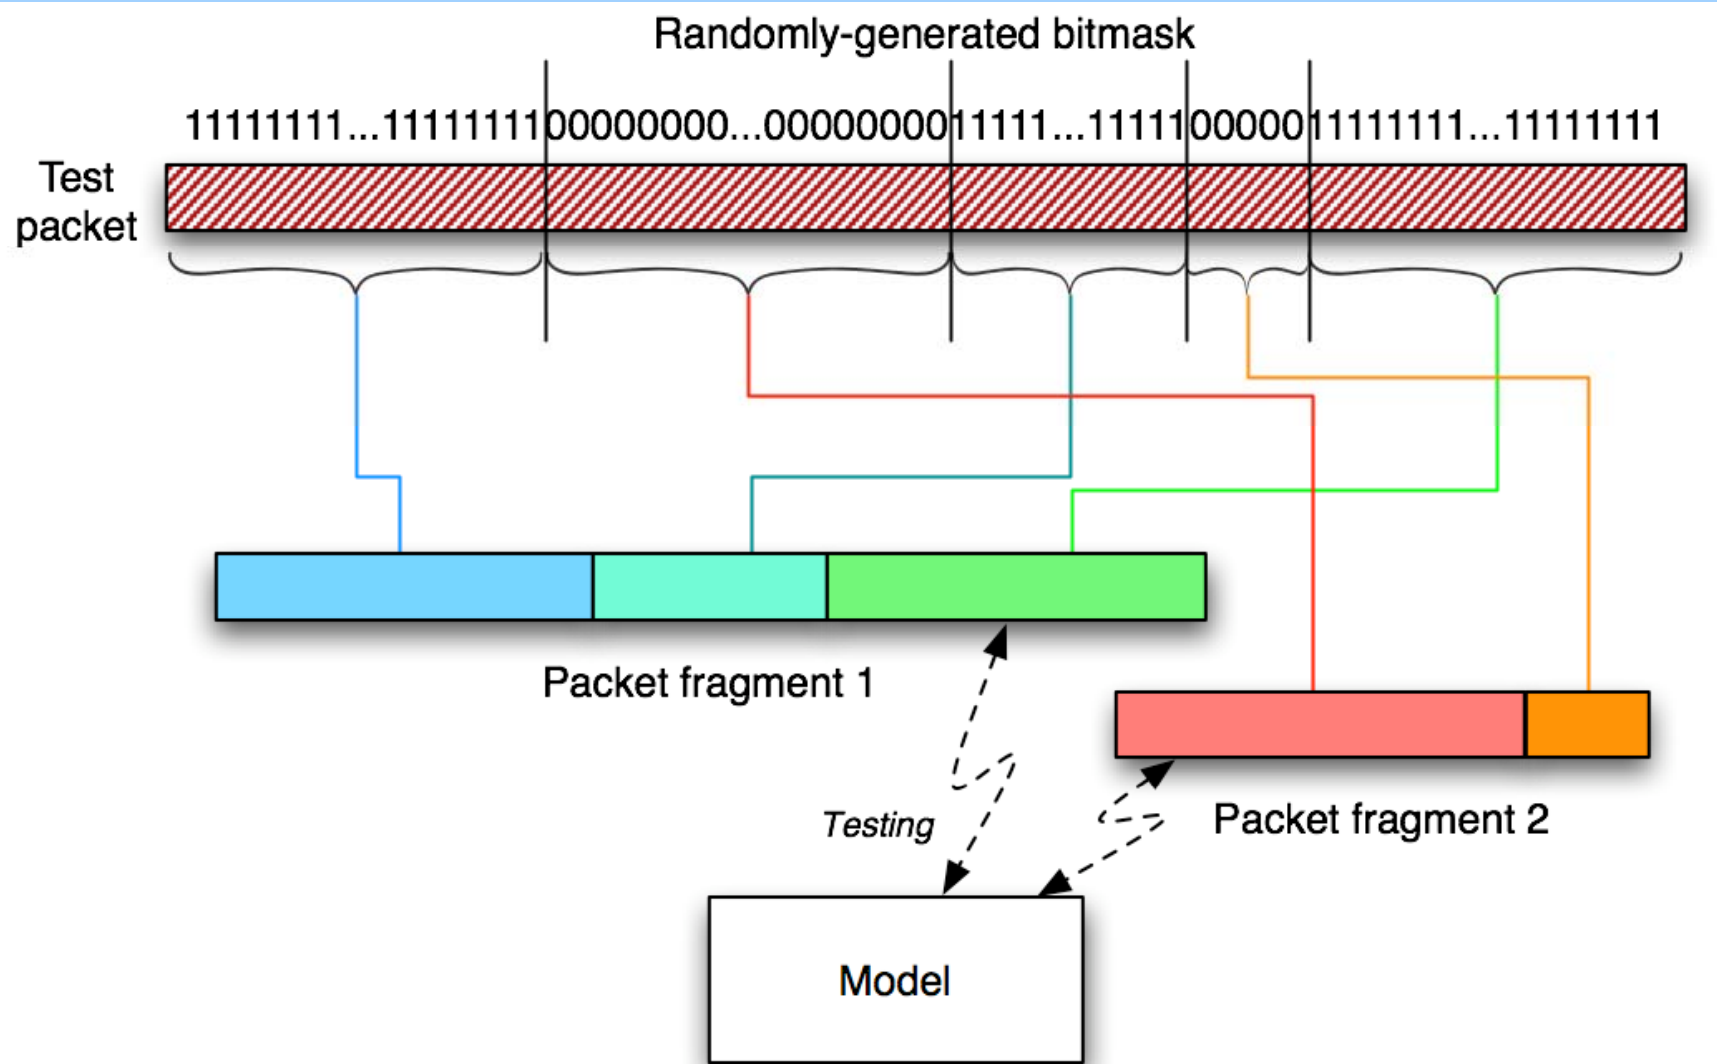
Randomization techniques

- **Randomized Modeling:** Separate the whole packet randomly into several (possibly interleaved) substrings or subsequences: S_1, S_2, \dots, S_N , and build one model for each of them
- Test packet's payload is divided accordingly
- Simpler strategy that does not incur substantial overhead: randomized testing.
 - build one model for whole packet, randomize testing portions



Randomization techniques (2)

- **Randomized Testing:** Simpler strategy that does not incur substantial overhead:
- Build one model for whole packet, randomize tested portions
 - Separate the whole packet randomly into several (possibly interleaved) partitions: S_1, S_2, \dots, S_N ,
 - Score each randomly chosen partition separately

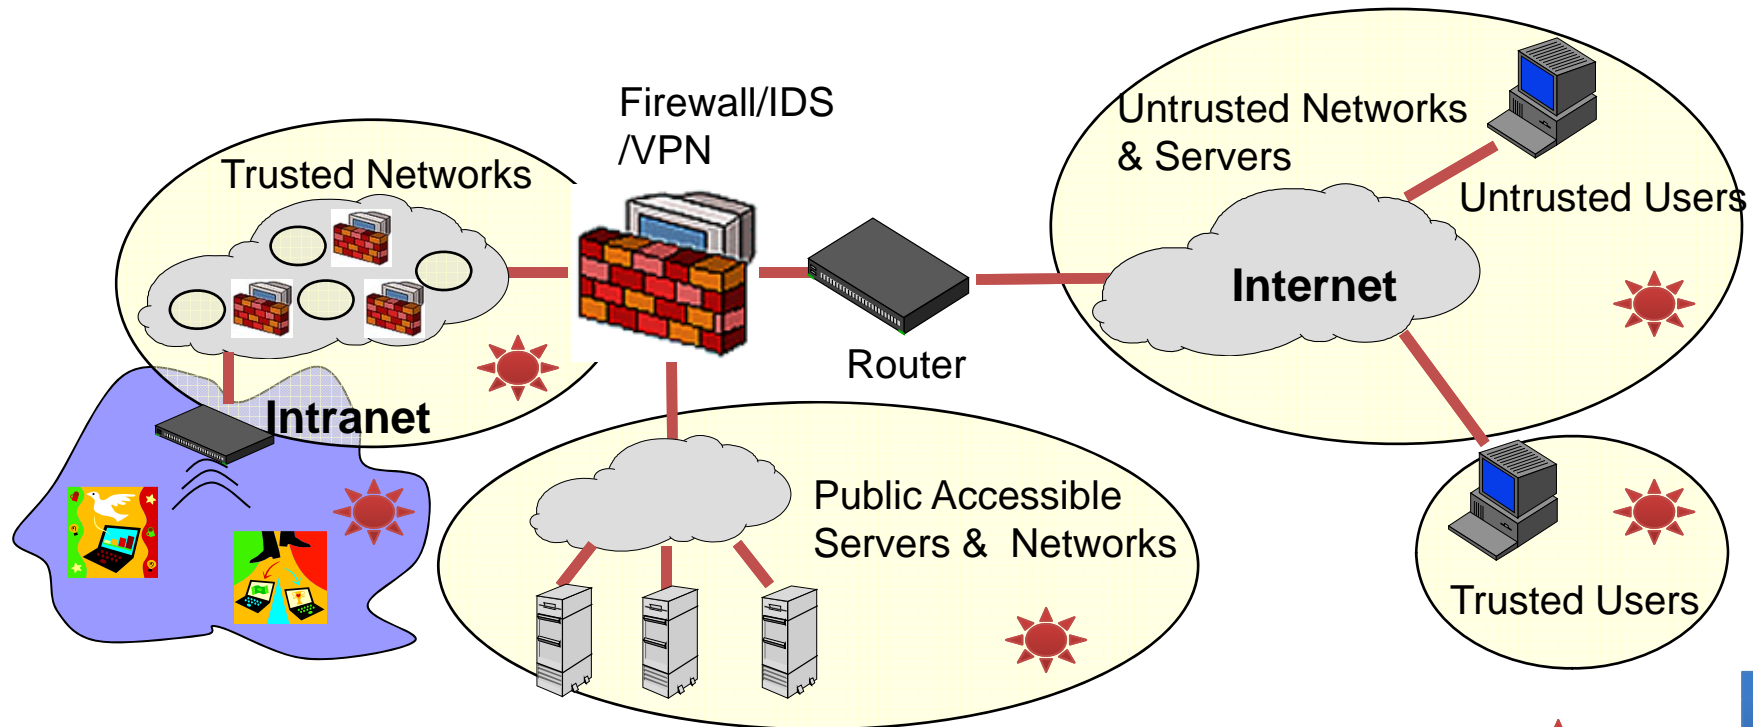



Secure Me

Many organizations have heterogeneous and distributed networks

What does security mean?

What are the challenges in *measuring* security properties?






Billions of Embedded Systems with no Anti-Virus (We have to break them to learn how to fix them)



50 Million - [Hacked](#)



[Hacked](#)



3? [Hacked???](#)



100 Million/year - [Hacked](#)



How many? [Hacked](#)



Thank you [Stuxnet](#)



Millions? (Unknown due to HIPPA)

Jan 27, 2014



70 Million - [Hacked](#)



Runs the internet! [Hacked](#)



- Back Up and Bibliography...

Security Maturity Model – Chapin05

- Chapin, David A., and S. Akridge. "How can security be measured." *Information Systems Control Journal* Vol 2 (2005): 43-47.




Security Maturity Model – Chapin05

- Goal: Provide metric for entire security architecture, lead to best practice, ability to compare organizations' security programs
- Uses baseline of complete security program - ISO 17799
- Measure the existence of each element for an organization
- Use this existence as a baseline for comparison between organizations
- Measure quality of each element via expert opinion (can decide on common features required to add some objectivity)
- Aims to measure broad improvement over time

Security Maturity Model - Discussion

- Feasible
- Assumes that ISO 17799 is a good baseline
- Hard to say how much improvement correlates to actual reduction in compromises

Table 9—Simulated Example Showing a Management Dashboard Comparison of Security Performance by Department

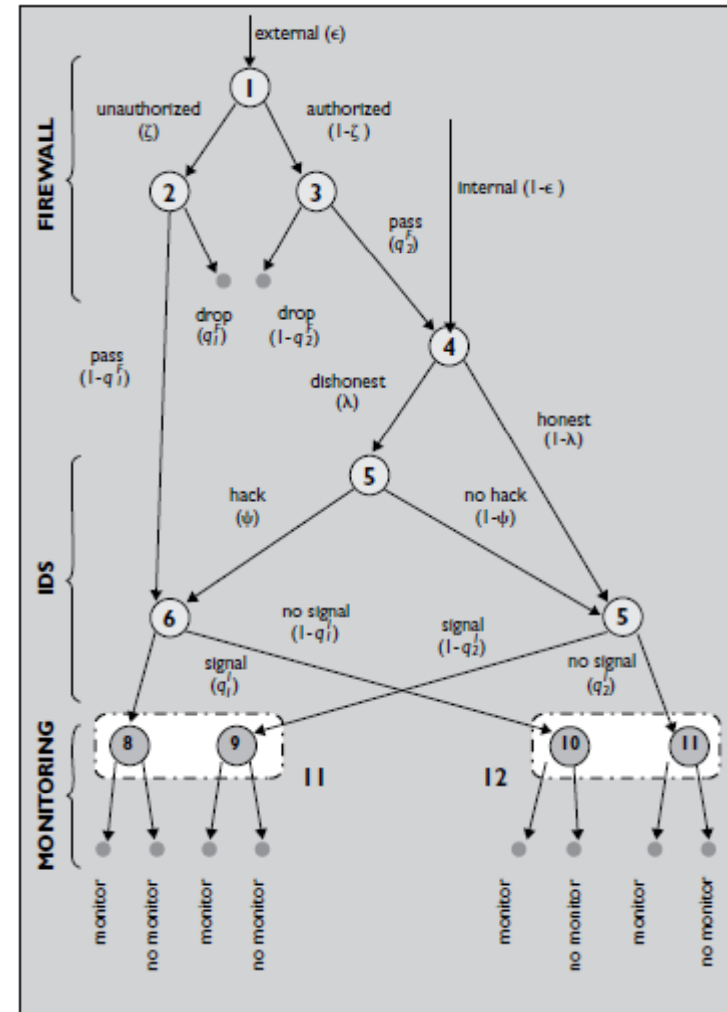
Department	Maturity Elements Owned (implemented elements in bold)	Maturity Level	Quality of Implemented Security Elements
			<input type="checkbox"/> high <input type="checkbox"/> medium <input type="checkbox"/> low
1	1.1, 3.2, 4.1, 7.5	Three (3) of four (4) implemented—75 percent	
2	7.1, 7.2, 7.3, 7.4, 7.6, 7.7, 7.8, 7.9, 7.10, 7.11, 7.12, 7.13, 7.14, 7.15, 7.16	Twelve (12) of 15 implemented—80 percent	
3	4.2, 4.3, 4.4, 4.5, 4.6, 4.7, 4.8, 4.9, 4.10, 4.11, 4.12	Eight (8) of 11 implemented—73 percent	

Cavusoglu04 – Analytical Model for Security Investments

- Huseyin Cavusoglu, Birendra Mishra, and Srinivasan Raghunathan. 2004. A model for evaluating IT security investments. Communications of the ACM Vol. 47, issue 7 (July 2004), 87-92.

Cavusoglu04 – Analytical Model for Security Investments

- Find probability of prevention/detection
- Find cost and effectiveness of monitoring
- Find cost of compromise
- Calculate the total costs with particular prevention and detection controls using game theory model to optimize monitoring choices
- Compare total costs to find which set of controls is best



Cavusoglu04 – Analytical Model for Security Investments - Discussion

- Considers layers and costs
- Probability of attack and effectiveness of different attacks is hard to determine

Stolfo11 – Relative Metrics and Defense in Depth

- Stolfo, S.; Bellovin, S.M.; Evans, D.; , "Measuring Security," Security & Privacy, IEEE , vol.9, no.3, pp.60-65, May-June 2011
- Bellovin, S.M.; , "On the Brittleness of Software and the Infeasibility of Security Metrics," Security & Privacy, IEEE , vol.4, no.4, pp. 96, July-Aug. 2006

Stolfo11 – Relative Metrics and Defense in Depth

- Goal: measurable defense in depth
- Is one set of controls more secure than another?
- Measure lower bound on adversary effort required to bypass a set of controls
- Assume: adversary total effort is roughly proportional to sum of lower bounds (assumes baseline of redundancy)
- Can we design ‘bonded’ layers that make adversary effort scale nonlinearly? (Bellovin06)
- Example attackers (design different layers for each):
 - Remote nation state – many layers + rate limiting
 - Inside operator – log and limit authorized users
 - Insider developer – randomize code and layout/implementation

Stolfo11 – Relative Metrics and Defense in Depth - Discussion

- Measuring adversary cost is promising (requires different data)
- Adversary can amortize much of the cost though
 - Example: 0-days normally cost 50k-100k+ but recently 0-days have been present in \$2000 exploit kits
 - 0-days last for a year on average in the wild [Bilge12]
- Complementary super linear combination of security controls are ideal but are difficult to create
- Attacker effort may not be proportional to the sum of the cost to bypass multiple layers if certain evasion techniques bypass multiple layers at once

Anomaly Detector Comparison - Ingham07

- Ingham, Kenneth, and Hajime Inoue.
"Comparing anomaly detection techniques for
http." Recent Advances in Intrusion Detection.
Springer Berlin/Heidelberg, 2007.

Anomaly Detector Comparison - Ingham07

- Goal: Compare HTTP anomaly detector algorithms on same data to test relative performance
- Collected many algorithms
- Four data sets
 - Real web servers for normal data scrubbed of attacks
 - Attack data from public sources for ground truth
 - Not shared
- Difficulty implementing others' algorithms due to poor descriptions

Algorithm	FP/day
Mahalanobis distance	91,524
χ^2 of <i>ICD</i>	∞
Length	∞
6-grams	13
DFA	37
Markov Model (log transform)	39,824
Linear combination	∞

Anomaly Detector Comparison - Ingham07 – Discussion

- Data sets leave open questions
 - Are the attacks representative? (Maybe actual attack data seen in the wild could have been left in)
 - Is one week long enough?
- Shows the need for clear algorithm definition or code sharing

Cost Sensitive Metrics – Lee00

- W. Lee, W. Fan, M. Miller, S. Stolfo, and E. Zadok. Toward cost-sensitive modeling for intrusion detection and response. In 1st ACM Workshop on Intrusion Detection Systems, 2000

Cost Sensitive Metrics – Lee00

- Cost-sensitive IDS – alerts only if
 response cost < damage cost prevented
- Intuitive example: choosing not to respond to low damage attack such as a scan because raising and processing an alert costs more
- Damage cost – damage done by an attack
- Response cost – examples: blocking an attack or manually investigating an intrusion

Cost Sensitive Metrics – Lee00

Table 2: Model for Consequential Cost

Outcome	Consequential Cost $CCost(e)$	Condition
Miss (False Negative, FN)	$DCost(e)$	
False Alarm (False Positive, FP)	$RCost(e') + PCost(e)$	if $DCost(e') \geq RCost(e')$ or if $DCost(e') < RCost(e')$
Hit (True Positive, TP)	0	
Normal (True Negative, TN)	$RCost(e) + \epsilon_1 DCost(e), 0 \leq \epsilon_1 \leq 1$	if $DCost(e) \geq RCost(e)$ or if $DCost(e) < RCost(e)$
Misclassified Hit	$DCost(e)$	
	$RCost(e') + \epsilon_2 DCost(e), 0 \leq \epsilon_2 \leq 1$	if $DCost(e') \geq RCost(e')$ or if $DCost(e') < RCost(e')$

- Costs based on alerts' actual ground truth
- PCost - cost of denying resources to a legitimate user
- Costs measured in relative units (ideally but likely not the same unit for DCost and RCost)
- Computing costs is hard
- Insurance payouts of breaches could be used to estimate certain damage costs [Greisiger12]

Bibliography

Approaches to Defense in Depth

Stolfo, S.; Bellovin, S.M.; Evans, D.; , "Measuring Security," Security & Privacy, IEEE , vol.9, no.3, pp.60-65, May-June 2011

URL: <http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=5772963&isnumber=5772950>

Bellovin, S.M.; , "On the Brittleness of Software and the Infeasibility of Security Metrics," Security & Privacy, IEEE , vol.4, no.4, pp. 96, July-Aug. 2006

<http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=1668014&isnumber=34919>

Geer, D., Jr.; Hoo, K.S.; Jaquith, A.; , "Information security: why the future belongs to the quants," Security & Privacy, IEEE , vol.1, no.4, pp. 24- 32, July-Aug. 2003

URL: <http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=1219053&isnumber=27399>

Chapin, David A., and S. Akridge. "How can security be measured." information systems control journal 2 (2005): 43-47.

<http://www.isaca.org/Journal/Past-Issues/2005/Volume-2/Documents/jpdf052-how-can-security.pdf>

Huseyin Cavusoglu, Birendra Mishra, and Srinivasan Raghunathan. 2004. A model for evaluating IT security investments. Commun. ACM 47, 7 (July 2004), 87-92.

<http://dl.acm.org/citation.cfm?id=1005828>

Bibliography

Experiments

Nathaniel Boggs, Salvatore J. Stolfo; "ALDR: A New Metric for Measuring Effective Layering of Defenses;" Layered Assurance Workshop; 2011/12/06

<http://ids.cs.columbia.edu/sites/default/files/law2011-aldr-final2.pdf>

Peisert, Sean, and Matt Bishop. "How to design computer security experiments." Fifth World Conference on Information Security Education. Springer Boston, 2007.

http://link.springer.com/chapter/10.1007%2F978-0-387-73269-5_19?LI=true

Ingham, Kenneth, and Hajime Inoue. "Comparing anomaly detection techniques for http." Recent Advances in Intrusion Detection. Springer Berlin/Heidelberg, 2007.

http://link.springer.com/chapter/10.1007%2F978-3-540-74320-0_3?LI=true#page-1

Hovav Shacham, Matthew Page, Ben Pfaff, Eu-Jin Goh, Nagendra Modadugu, and Dan Boneh. 2004. On the effectiveness of address-space randomization. In Proceedings of the 11th ACM conference on Computer and communications security (CCS '04). ACM, New York, NY, USA, 298-307. <http://doi.acm.org/10.1145/1030083.1030124>

Matt Weir, Sudhir Aggarwal, Michael Collins, and Henry Stern. 2010. Testing metrics for password creation policies by attacking large sets of revealed passwords. In Proceedings of the 17th ACM conference on Computer and communications security (CCS '10). ACM, New York, NY, USA, 162-175. <http://doi.acm.org/10.1145/1866307.1866327>

Kelley, P.G.; Komanduri, S.; Mazurek, M.L.; Shay, R.; Vidas, T.; Bauer, L.; Christin, N.; Cranor, L.F.; Lopez, J.; , "Guess Again (and Again and Again): Measuring Password Strength by Simulating Password-Cracking Algorithms," Security and Privacy (SP), 2012 IEEE Symposium on , vol., no., pp.523-537, 20-23 May 2012

URL: <http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=6234434&isnumber=6234400>

Bibliography

Data

Ali Shiravi, Hadi Shiravi, Mahbod Tavallaee, Ali A. Ghorbani, Toward developing a systematic approach to generate benchmark datasets for intrusion detection, *Computers & Security*, Volume 31, Issue 3, May 2012, Pages 357-374, ISSN 0167-4048, 10.1016/j.cose.2011.12.012.

<http://nsl.cs.unb.ca/images/stories/publications/iscx2012.pdf>

Tavallaee, M.; Bagheri, E.; Wei Lu; Ghorbani, A.A., "A detailed analysis of the KDD CUP 99 data set," *Computational Intelligence for Security and Defense Applications, 2009. CISDA 2009. IEEE Symposium on*, vol., no., pp.1,6, 8-10 July 2009

<http://www.ee.ryerson.ca/~bagheri/papers/cisda.pdf>

Bibliography

Attacker Modeling

Sheyner, O.; Haines, J.; Jha, S.; Lippmann, R.; Wing, J.M.; , "Automated generation and analysis of attack graphs," Security and Privacy, 2002. Proceedings. 2002 IEEE Symposium on , vol., no., pp. 273- 284, 2002

URL: <http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=1004377&isnumber=21681>

Bruce Schneier. Attack trees: Modeling security threats. Dr. Dobb's Journal, December 1999.

<http://www.schneier.com/paper-attacktrees-ddj-ft.html>

Kyle Ingols; Richard Lippmann; Keith Piwowarski; , "Practical Attack Graph Generation for Network Defense," Computer Security Applications Conference, 2006. ACSAC '06. 22nd Annual , vol., no., pp.121-130, Dec. 2006

URL: <http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=4041160&isnumber=4041139>

Moore, Andrew P., Robert J. Ellison, and Richard C. Linger. Attack modeling for information security and survivability. No. CMU-SEI-2001-TN-001. CARNEGIE-MELLON UNIV PITTSBURGH PA SOFTWARE ENGINEERING INST, 2001.

<http://www.dtic.mil/cgi-bin/GetTRDoc?Location=U2&doc=GetTRDoc.pdf&AD=ADA388771>

Bilge, L., & Dumitras, T. (2012). Before We Knew It. CCS

http://users.ece.cmu.edu/~tdumitra/public_documents/bilge12_zero_day.pdf

Verizon Data Breach Investigations Report

http://www.verizonbusiness.com/resources/reports/rp_data-breach-investigations-report-2012_en_xg.pdf

Mark Mateski, Cassandra M. Trevino, Cynthia K. Veitch, John Michalski, J. Mark Harris, Scott Maruoka, Jason Frye. Cyber Threat Metrics. 2012. Sandia National Laboratories. <http://www.fas.org/irp/eprint/metrics.pdf>

Bibliography

Defender's Costs

Ross Anderson, Chris Barton, Rainer Bohme, Richard Clayton, Michael van Eeten, Michael Levi, Tyler Moore, and Stefan Savage. Measuring the cost of cybercrime. 11th Annual Workshop on the Economics of Information Security (WEIS), 2012.

http://weis2012.econinfosec.org/papers/Anderson_WEIS2012.pdf

Florêncio, Dinei, and Cormac Herley. "Sex, Lies and Cyber-crime Surveys." Economics of Information Security and Privacy III (2011): 35-53.

<http://research.microsoft.com/pubs/149886/SexLiesandCybercrimeSurveys.pdf>

W. Lee, W. Fan, M. Miller, S. Stolfo, and E. Zadok. Toward cost-sensitive modeling for intrusion detection and response.

In 1st ACM Workshop on Intrusion Detection Systems, 2000

<http://ids.cs.columbia.edu/sites/default/files/wenke-acmccs2k-cost.pdf>

Mark Greisiger. Cyber Liability & Data Breach Insurance Claims. October 2012

<http://www.netdiligence.com/files/CyberClaimsStudy-2012sh.pdf>