Microsoft | Trustworthy Computing

Threat Genomics

An evolution and recombination of best-available models and techniques for characterizing and understanding network threats

Jon Espenschied, TwC Network Security Advanced Analytics
Angela Gunn, TwC Security Response Communications

Microsoft Trustworthy Computing (TwC)

1:15

Four Brief Talks

About 10-15 min each:

- Bases
  - Defining a sufficiently complete picture of actions in attacks
  - Build and borrow from the best available theories
- Structure
  - Connecting base types of action to find patterns
  - Focusing on history; weeding out subjective attributes
- Genome
  - Recognize the current state of an attack
  - Predict missing pieces of incomplete patterns from hard markers
- Phenome
  - Improve detection solely from observable behavior
  - Differentiate groups of actors from concurrent data

…all really about this thing we call "Threat Sequences"

Introduction - Threat Sequences

Introduction

- Many threat models seek to determine the nature of an online attack
  - From discrete events or technical artifacts
  - Useful in detections and identification
  - BUT limited by lack of behavioral and timing data
- We seek to determine the nature of an attack from recognizable states and transitions between those states
  - Predictive of likely next activities (faster detection and response)
  - Predictive of likely types of targeted assets
- Signals vs intel
  - Metric-based (signal) analysis eliminates speculation, trust issues, FUD
  - Leverages data collection tactics and protections you already use
  - Eliminate the subjective and focus on facts and remediation

Today's talk: Terminology

- Focusing on serious adversaries (DHA, APT)
  - Attackers tend to take similar (but not identical) steps to reach similar ends
  - A tool (an 0day, a piece of malware) is only a tool – focus on the actions
- Threat Sequence model
  - Qualitative labeling and characterization of security events
- Bases
  - Ten basic forms, types or stages of action
- Structure or Sequence
  - Common patterns in the base types of action; a library of groups of actions
- Genome
  - Using structural / sequence evidence to recognize or detect new activity
- Phenome
  - Observable characteristic behaviors in or at transition points between base types

Today's talk: Quick Preview

- A quick preview of Threat Sequences

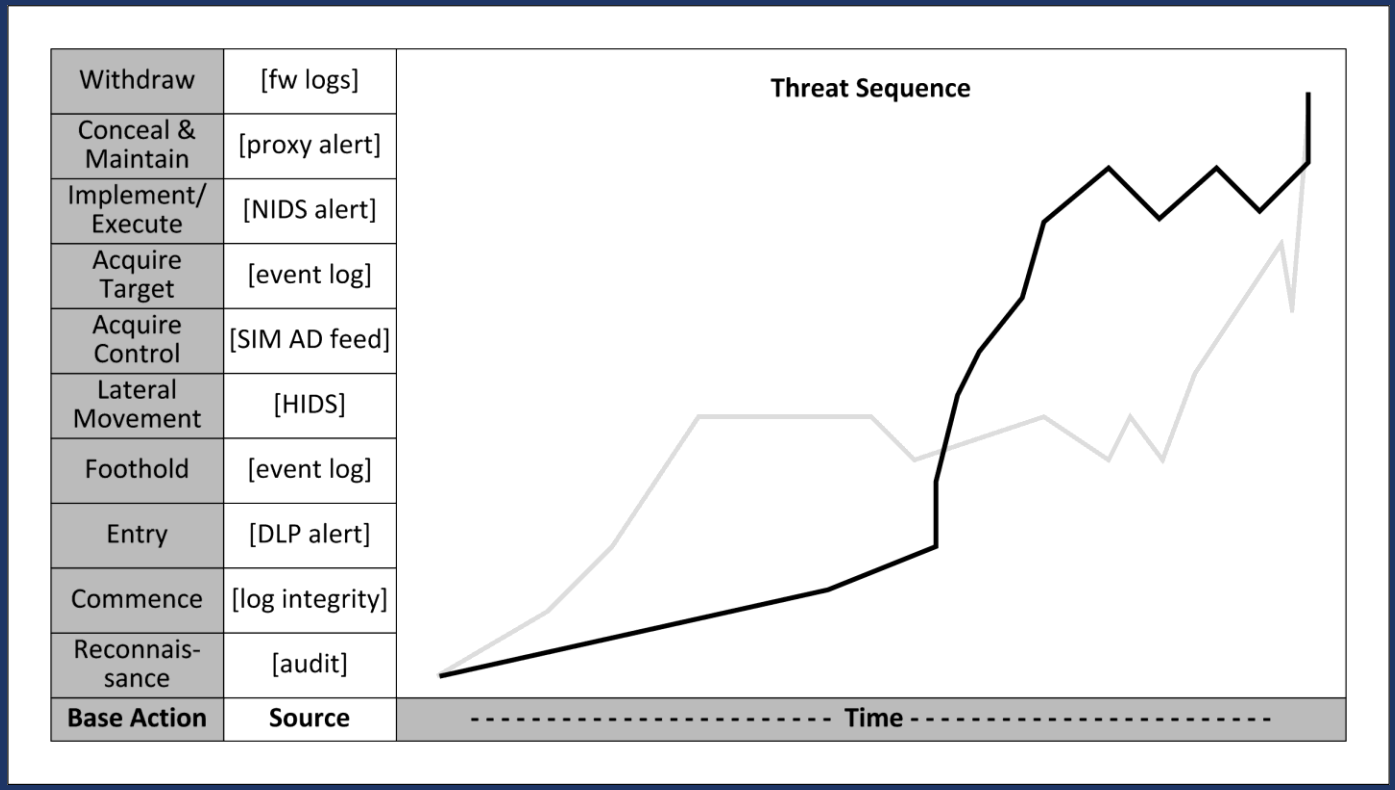

- Qualitative characterization of actions over time, with attributes useful for detection

Today's talk: Threat Sequences

- Foundations of our approach come from technical and military ("cyber") theory as well as from empirical observations
  - McRaven: Relative superiority (RS50)
  - The Cyber Kill Chain (intel-driven)
  - …and many others
- Addressing persistence, movement, transitions and graph theory
  - Any specific tool (an 0day, a piece of malware) is only a temporary means to an end – focus on the behavior of the actor
  - Use historical patterns as a detection template

Today's talk: Threat Sequences (2)

- Where are we going with this? What can we do?
- From labels and patterns to sequences
  - Characterize an attack, not a tool
  - Tools/signatures often transient – mitigations improve, vulnerabilities are addressed, malware detections propagate
- From sequences to genomes and phenomes
  - Match patterns using hard and soft markers
  - The behavior of actors moving through a sequence of base actions often remarkably consistent over time / through multiple campaigns
- Note: Certain additional methods & tools information under NDA

Microsoft | Trustworthy Computing

Bases

Bases

- How we got here: McRaven's Relative Superiority Model

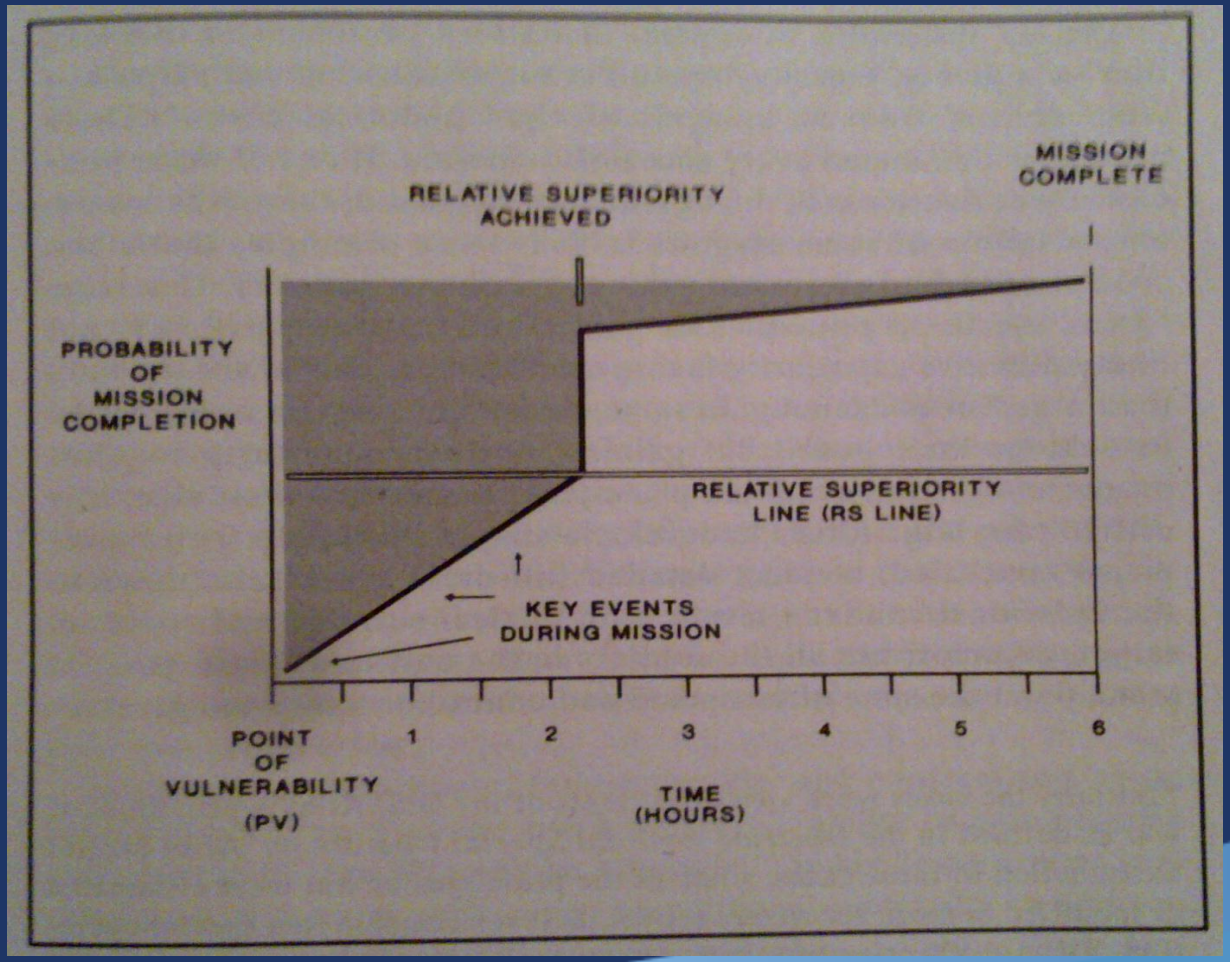

Bases

- How we got here: McRaven's Relative Superiority Model

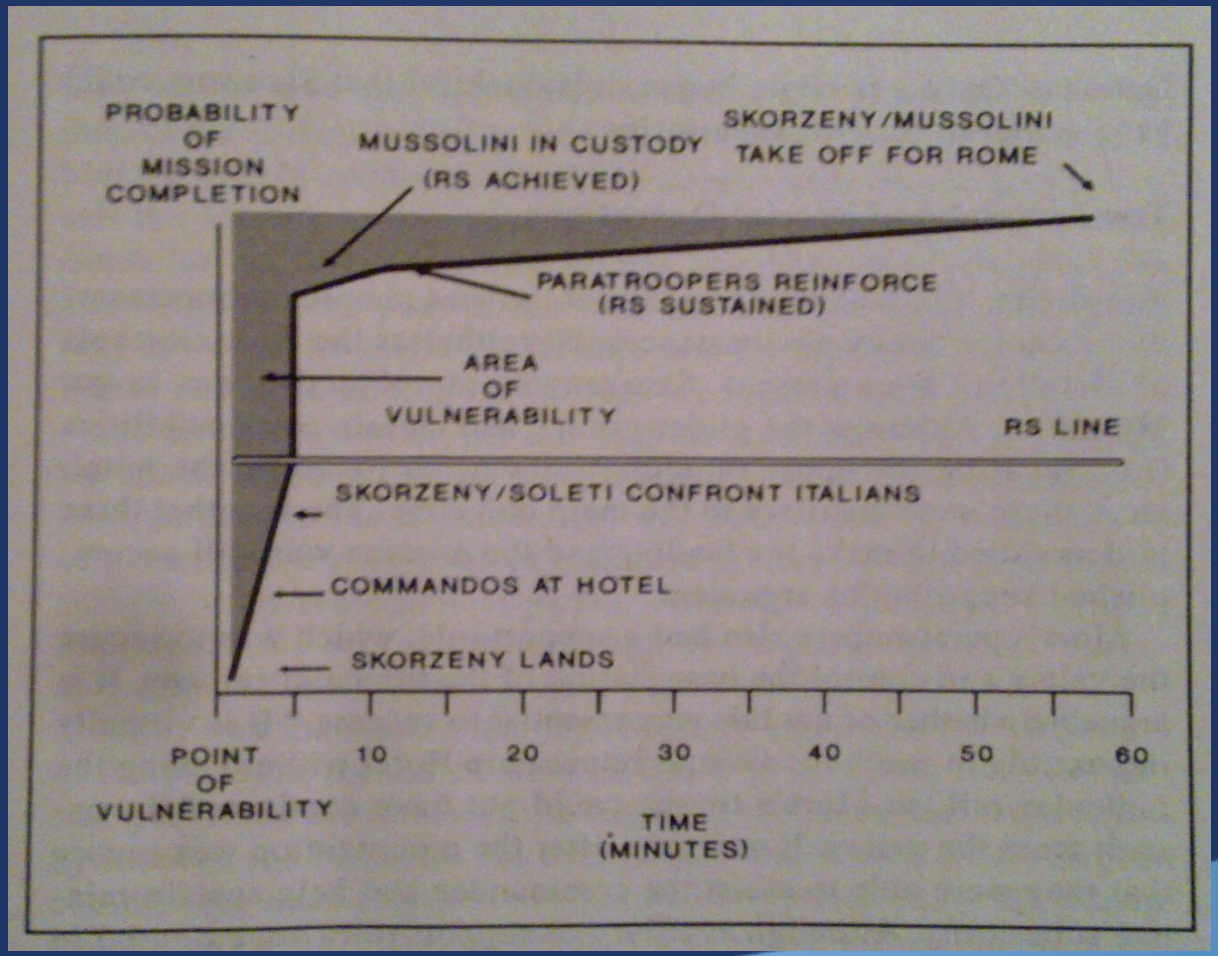

Bases

- How we got here: McRaven's Relative Superiority Model

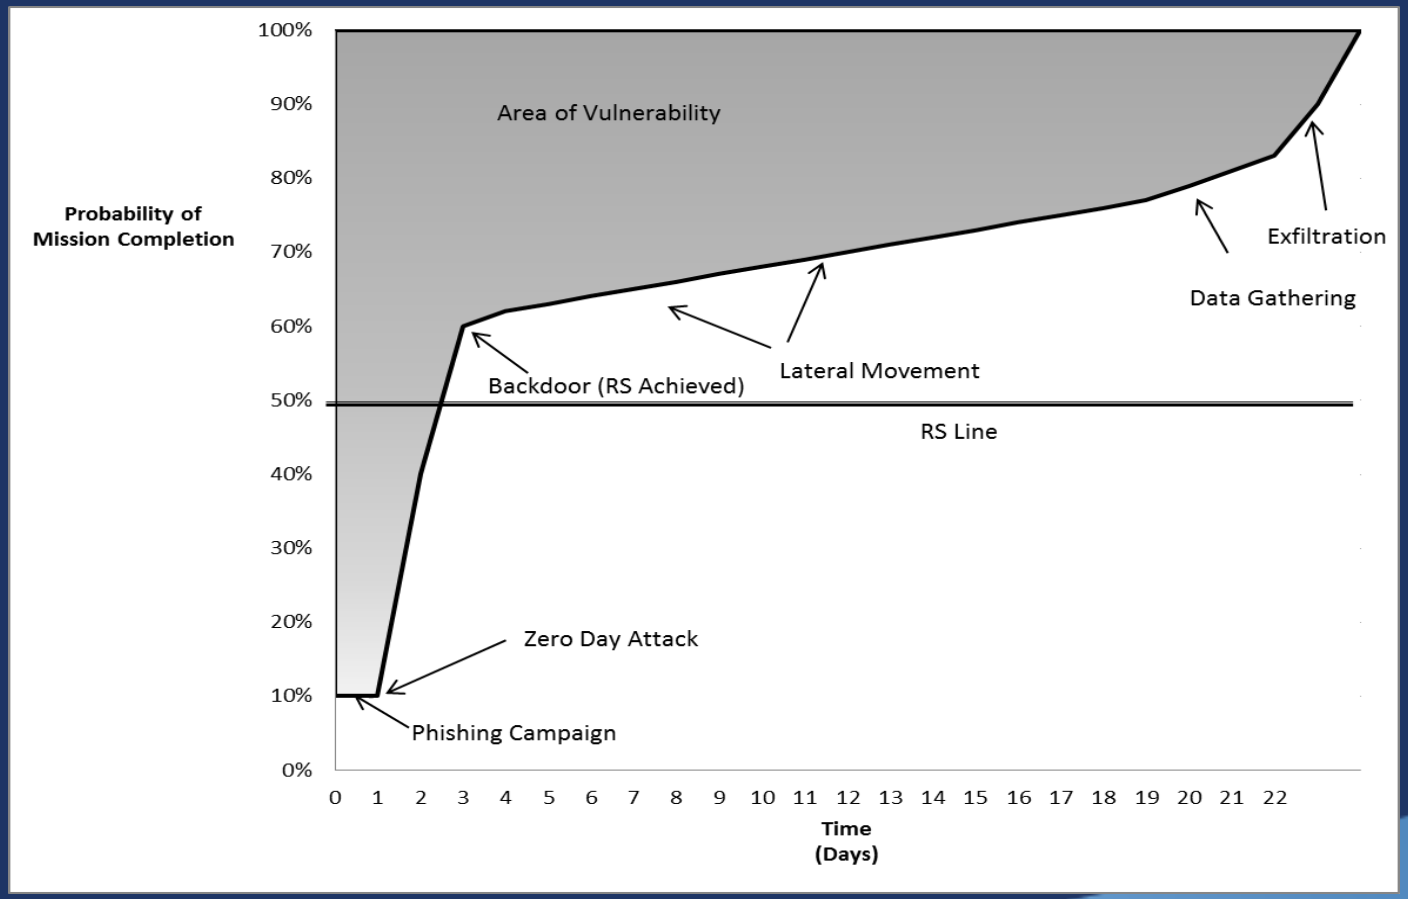

Bases

- How we got here: Kill Chains and other categories of action

Bases

- Ten base types of action
  - Reconnaissance
  - Commencement
  - Entry
  - Foothold
  - Lateral movement
  - Acquire control
  - Acquire target
  - Implement / execute
  - Conceal / maintain
  - Withdraw
- In practice, it is rare for an attack to include all base types sequentially, and some may not be present at all
- Also highly unlikely that an attack will be detected during initial reconnaissance

Bases

- Base Types of Action are qualitative labels for adversary activity
- Base types have comparable attributes; e.g.:
  - Mode or time of detection
  - Duration of action
  - Number of concurrent, sequential, or periodic actions
  - Source, destination, and/or vector
- Metrics mapped to base types help make clear the difference between the forest and the trees
  - Avoiding subjective attribution and judgments
  - Sufficient characterization is possible with objective attributes applied to data

Bases

- Building a superset from Kill Chain and other sequential models
- Relative Superiority Model redone as qualitative buckets

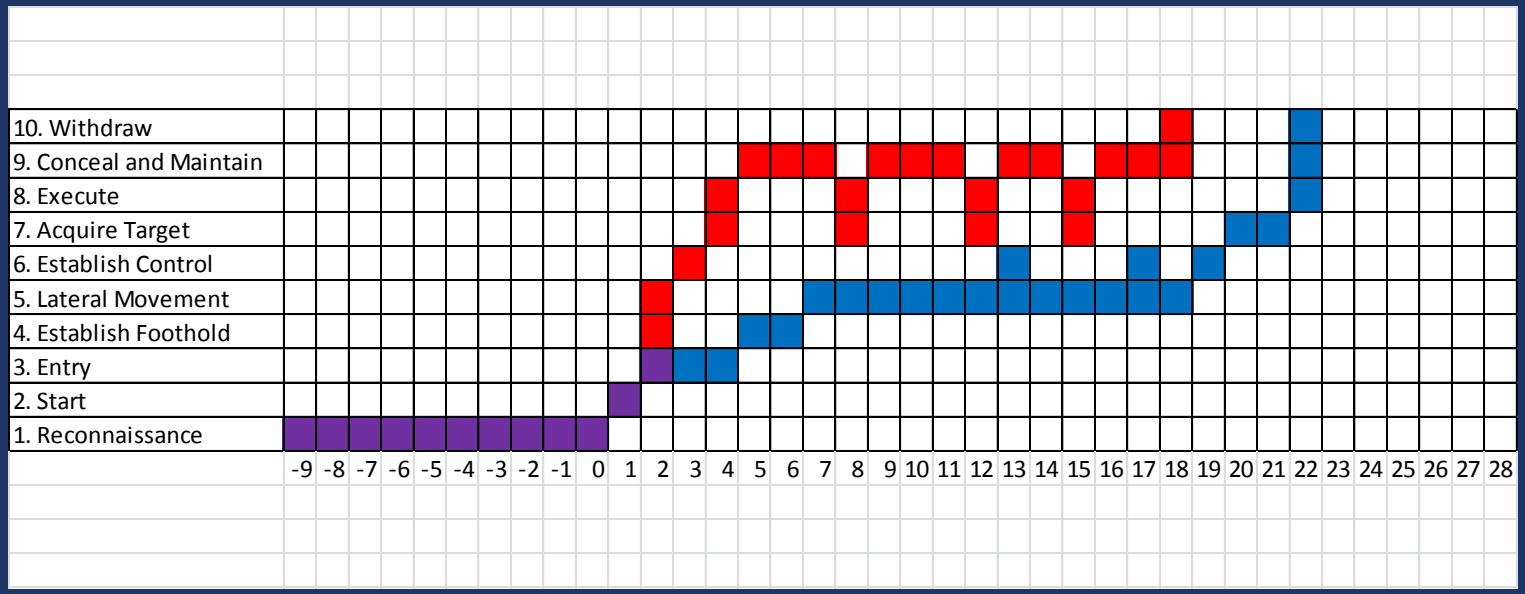

Microsoft | Trustworthy Computing

Structure

Structure

- We need a way of associating (structuring) the base events and pinpointing what's relevant in "Threat Sequences"
  - What detection indicators apply to each base type of action?
  - What attributes of actions tell us to go back and look for other detection indicators?
- Finding and defining attributes for each base type allows us to connect, group, and ultimately correlate activities that might otherwise appear unrelated
- Relevant attributes will vary by organization and environment
  - Data sources and feeds
  - Direct indicators
  - Complex and correlated indicators
  - Extended attributes

Structure

- Mapping detection to a qualitative metric (the base type of action)

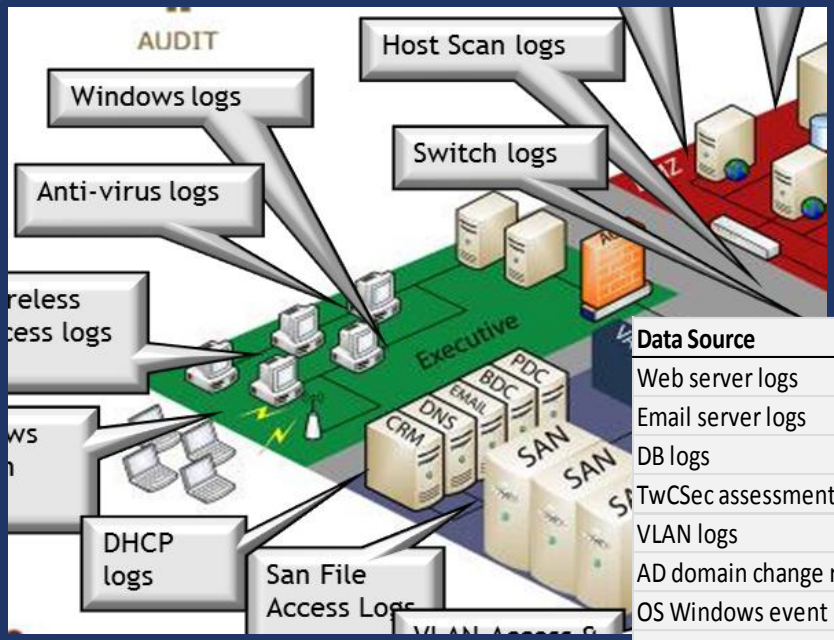

Data Source	1	2	3	4	5	6	7	8	9	10
Web server logs		■	■							
Email server logs		■	■	■						
DB logs				■			■	■		
TwCSec assessment				■		■	■	■		
VLAN logs					■				■	■
AD domain change reports					■	■			■	
OS Windows event logs		■	■							
OS other desktop logs			■							
AV logs		■	■	■		■				■
Host scan logs				■		■				
HIDS		■				■				
ACS/FEP event logs			■				■			
Web proxy logs	■	■		■		■		■	■	
IDS/IPS logs	■	■	■	■	■	■		■	■	■
Firewall logs		■								

Structure: Common attributes

- Connecting and finding historical patterns between types of action
  - Can use IoCs, VERIS, other models as long as they connect bases
- Typical set of attributes useful for building a library of patterns:
  - An identifier (ID) and optional name for automatable reference
  - Time detected, usually a marker of first detection set by an IOC
  - Duration start, Y/M/D/H (or Δ- D/H time in retrospect) if != detected
  - Duration end, or last known/confident detection
  - Source of alert or detection, specific or in aggregate with ID that allows traceback
  - Targeting, including evidence of randomness or selection by opportunity, area, sequence, or specific point
  - Indicator of Compromise (IOC) record, if available
  - IDs of all involved source/destinations, whether system, account, or application
  - Vector, showing incoming, outgoing, stasis, or lateral movement; avoid intermediate guesses of victimhood or attribution

Structure: Extended attributes

- This information is commonly available but tends to focus on a specific vulnerability or incident and should be adjusted for the environment:
  - Base type of action, usually estimated by analyst or normalization rules
  - Time in relation to potentially related base actions
  - Evidence of human behavior, including parallel or sequential actions, decisions, escalation, coordination, defacement or other markers, and other behavioral attributes
  - IOC or other alert record
  - Alert source and type
  - IPv4/6 and any DNS records for involved entities
  - IP flow or trace data, or other captured data in the alert
- (but wait, there's more…)

Structure: Extended attributes (2)

- Target asset sensitivity or entity access level; a suggested basic nomenclature is:
  - *Low*: Public or low business impact data for which integrity outweighs confidentiality
    Upper range: Negotiable assets (money/financial assets which may be insured)
  - *Medium*: Confidential or medium business impact data
    Upper range: Tools, code, credentials, or data which allows elevation
  - *High*: High business impact data, such as critical trade secrets and classified data
    Upper range: Assets affecting human life and safety, or classified compartmentalized information

(…and still more…)

Structure: Extended attributes (3)

- Organization type, usually by industry, size, or business relationship, such as:
  - General populace/individuals
  - Education, research, and other independent nonprofits
  - Technology and telecom organizations including software, hardware, integrators, and operators
  - Industries including service, retail, manufacturing, and materials producers
  - Infrastructure and transport including all utilities
  - Finance including banks, CU, credit, transaction processors, and financial NGOs
  - Government including all federal/state/local civilian agencies, domestic intelligence, and law enforcement
  - Military including geopolitical actors, international intelligence, and some NGOs

Structure: Preparing for analysis

- Evaluate only the history and existing data
  - Long-term projections are inappropriate at this stage
- Think nodes and vectors (direction)
  - Avoid speculating on the identities of attackers – or even of victims
- Even at this stage, notable patterns may begin to emerge
  - Persistence after initial success correlates with different targeting
  - "Fast," wildly successful attacks usually preceded by long reconnaissance
  - Operational sophistication does not correlate with applied technical sophistication
  - Historical similarities between attacks are useful for pattern matching
- Common fallacies
  - Code reuse means common actor; 0-day is APT; etc.

Microsoft | Trustworthy Computing

Genomes

Genomes: A useful analogy

- Using structure/sequence techniques and “hard markers” to detect and recognize attacks
- Similarities to process of genome analysis
  - Manual or automated “sequence walking” compares active event to known patterns
  - Libraries of full and fragmented threat sequences can be assembled
  - Patterns can be isolated and correlated
  - An alert about a detected pattern can serve to lower detection thresholds for related (“chained”) patterns – if X, usually Y, so prepare for Y
- Setting up and using detection metrics
- Finding recognizable patterns in common sequences

Genomes: Construction

- Setting up and using detection metrics
  - Indicators or inputs should exist for each base type
    - HIDS alert thresholds may indicate Foothold or Acquire Target events
    - Analytics and border-firewall logs may indicate Reconnaissance
    - Internal security log analytics servers may spot Conceal & Maintain activity
    - And so forth
- For the purposes of characterizing an adversary, focus should remain on the output of detective controls, with the goal of formulating a matrix in which each base type is addressed, as well as a baseline from which to detect changes from normal activity patterns
- Finding recognizable patterns in common sequences
  - The datasets may be quite large
  - Adjacent activity is a helpful indicator when logs/alerts are correlated with base types

Genomes: Visualization

- Visualization can be a powerful aid to understanding and response, and when communicating your findings to others

- Three examples of visualized threat sequences:
  - Simple sequencing to characterize a maneuver or single attack
  - Differentiating between complex attacks with threat sequences
  - Advanced sequencing to recognize campaigns or multiple actors

Genomes: Simple visualization

- A simple sequence visualization, characterizing a single maneuver or an attack in which just one tool is deployed (eg., a malware infection)

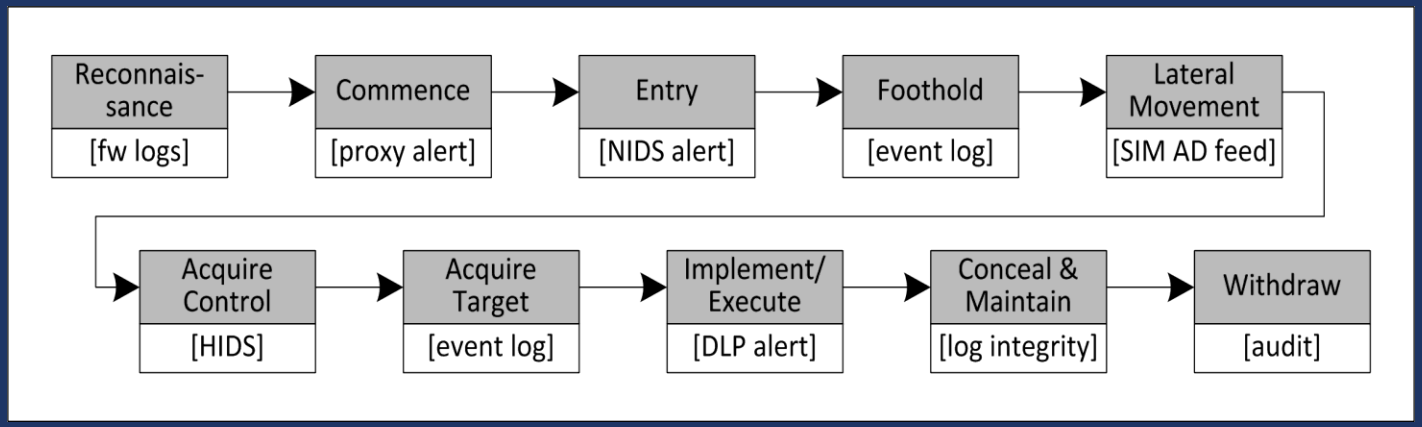

- Examples include Cyber Kill Chain, Common Attack Pattern Enumeration and Classification Model, many others

- Can pinpoint control failure

Genomes: Differentiated visualization

- An attack-curve visualization differentiating between complex attacks with threat sequences, mapped over time

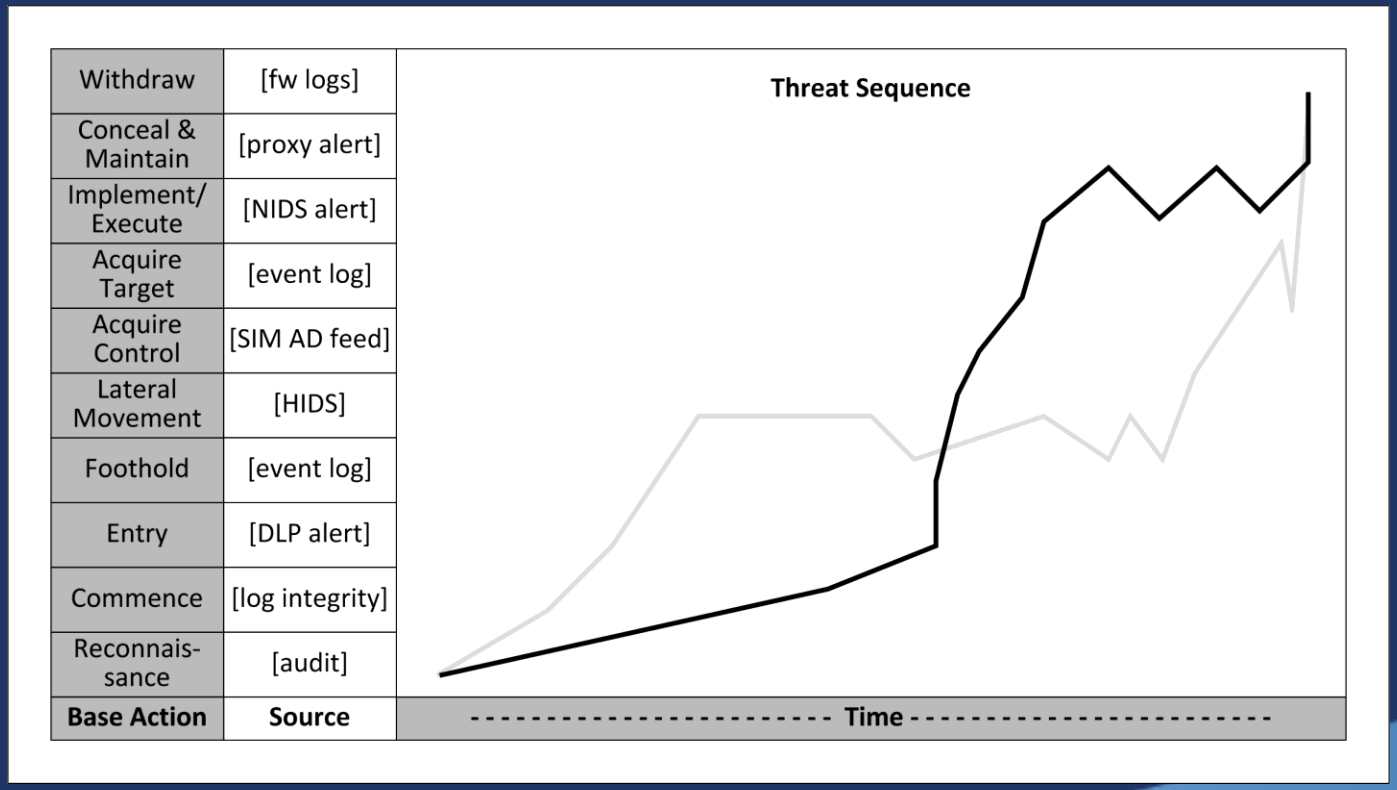

Genomes: Advanced visualization

- Advanced visualization showing a source-mapped threat sequence against a background of other activity, including a DDoS (lower right)

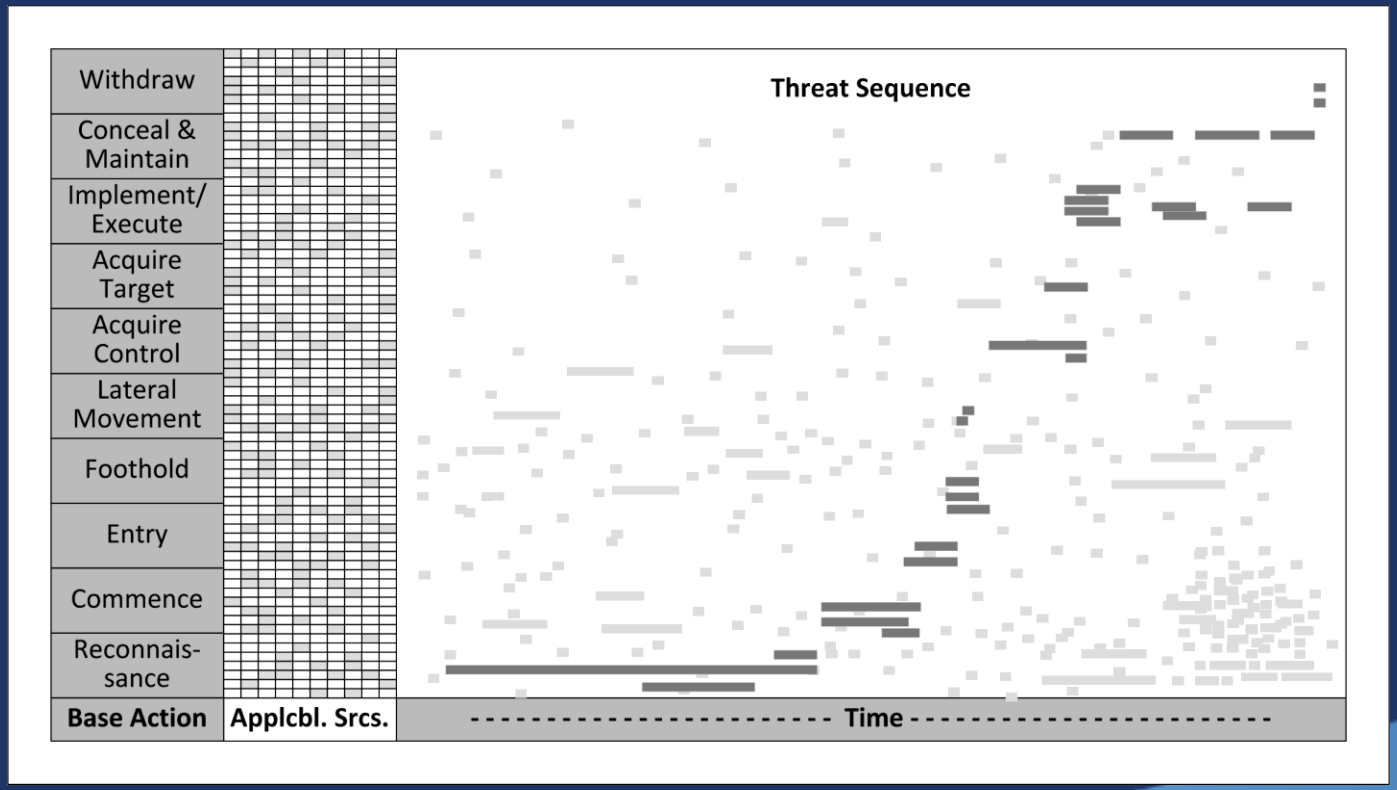

Microsoft | Trustworthy Computing

Phenomes

Phenome

- The genome analogy extends to "phenome"
  - A threat genome maps a particular threat's structure, while
  - Phenome indicates an observable expression of its nature
  - In this case, behavior as an attack transitions between base types
- The "soft markers" of our model
  - Used to refine differentiation
  - A method to find serious actors in background noise
- Though our model excludes attribution*, this is the point at which certain kinds of human intel can play a useful part in the process

*Char Sample (US-CERT) is taking this further and doing great work in this area (by EOY).

Phenome: Expressions

- Observable expressions of behavior
  - Two actors attempting the same attack may differ in their approach due to cultural / organizational differences between the two
- Differentiation vs attribution
  - Possible to richly differentiate among attacks with similar goals (genome) by noting differences in behaviors (phenome), while still avoiding attribution
  - Data sources for analysis of behavioral attributes include speed, use of surprise, management, evidence of group activity, apparently preceding action (planning and/or duration), contingencies, timing around holidays, evident creativity in problem-solving, and many more
  - McRaven called out six similar principles in his Relative Superiority work – simplicity, security, repetition, surprise, speed, and purpose – many of these are key indicators of successful attacks, and the visualization tools are meant to show them.
  - These often manifest in fairly specific, measurable ways.

Phenome: Behavioral metrics

- Variations in phenome expression may have a cultural basis, an organizational basis, or a combination of the two
- Cultural values map poorly to individual actors, but they are firm underpinnings for how people act within groups; 40+ years of research by Hofstede/IBM show measurable differences exist in:
  - Power Distance Index (PDI)
  - Uncertainty Avoidance Index (UAI)
  - Individualism vs collectivism (IDV)
  - Aggression (masculinity) (MAS)
  - Long-term vs short-term orientation (LTO)
  - Indulgence vs restraint (IVR)
- By excluding inherent behavioral indices from bases where they are expectable, cultural artifacts begin to appear

Phenome: Behavioral metrics

- Others can (and are) pursuing specific attribution; we're content to borrow metrics tested to show statistically significant differentiation
- Sample of "6 dimensions" metrics by country:

country	pdi	idv	mas	uai	ltowvs	ivr
Africa East	64	27	41	52	32	40
Africa West	77	20	46	54	9	78
Argentina	49	46	56	86	20	62
Australia	36	90	61	51	21	71
Austria	11	55	79	70	60	63
Canada	39	80	52	48	36	68
Chile	63	23	28	86	31	68
China	80	20	66	30	87	24
Colombia	67	13	64	80	13	83
India	77	48	56	40	51	26
Indonesia	78	14	46	48	62	38
Iran	58	41	43	59	14	40
Ireland	28	70	68	35	24	65
Israel	13	54	47	81	38	#NULL!
Italy	50	76	70	75	61	30
Philippines	94	32	64	44	27	42
Poland	68	60	64	93	38	29
Sweden	31	71	5	29	53	78
U.S.A.	40	91	62	46	26	68
Venezuela	81	12	73	76	16	100
Vietnam	70	20	40	30	57	35

Phenome: Behavioral metrics

- Mapping base types of action and transitions between bases in an incident, this time with qualitative behavioral metrics:
  - Power Distance Index (PDI):
    Do parallel actors take the same actions? Are they using a playbook?
  - Uncertainty Avoidance Index (UAI)
    Are attackers pragmatic? Do they adapt or keep trying failed attacks?
  - Individualism vs collectivism (IDV)
    Is there an aversion to using NIH tools? Tendency to follow group activity?
  - Aggression (masculinity) (MAS)
    Is there direct reaction to being blocked or removed from a system? Are there markers for ownership or entitlement? Hostility toward remediation?
  - Long-term vs short-term orientation (LTO)
    Is there an investment and intent to stay resident? Active maintenance or observation (not just time in a botnet)?
  - Indulgence vs restraint (IVR)
    Is there defacement? Flair? A distinctive style or tendency to leave cryptic clues?
- Qual rating (0-lo-med-hi-1) appears sufficient to differentiate.
- Exclude inherent behavioral markers (throw out MAS in some cases)

Phenome: Behavior of organizations

- In addition, organizations often express their own corporate cultures in dimensions that can measured relative to other organizations:
  - Means vs goals (results orientation)
  - Internally or externally driven (attitude toward customers)
  - Easygoing vs strict work discipline (internal structuring)
  - Local vs professional (identification)
  - Open system vs closed (organization accessibility)
  - Employee focus vs work focus (management philosophy)
- As with cultural measures, some expressions are intrinsic to certain bases; other behaviors indicate specific value systems below the surface: *Free actors? Corporate? Military?*
- Both cultural and org-culture behaviors can be applied throughout the base-type sequence, as long as activity can be detected and recorded

Phenome: Indicators and bases (examples)

- Site defacement is often an impulsive act attackers use to assert their dominance over a network – until the defacement is taken down. It conforms to the base type Implement / Execute and indicates that the attacker has low Long-Term Orientation and a tendency toward Indulgence over Restraint
- Some cultures are more aggressive (high-MAS) than others. When detected and thrown off the network by a canny sysadmin, some attackers may simply leave, while others may attempt to retaliate against the system off which they're being thrown
- A full mapping of base actions to cultural indicators, though a fruitful endeavor, is beyond the scope of this presentation

Conclusion

Conclusion: Four Parts

- The Threat Sequence model allows qualitative characterization and labeling of security events so that they may be normalized and correlated into a coherent whole
- The 10 base types used to construct a threat sequence are labels or categories into which an attacker's actions can be sorted, with sufficient precision to distinctly characterize the attack for analysis
- Genome analysis uses structural / sequencing techniques to recognize or even detect patterns
- Phenome analysis examines characteristic behaviors at decision points to further refine recognition, detection, and response options

Conclusion: Anton's Questions

- How you achieved "quick wins" with security metrics?
  - Not a quick win in this case – a few years of thinking & observation
- How you define useful metrics, whether risk or operational?
  - Operational: "finding bad guys" is a driving force
- What metrics you track are the most useful?
  - The whole is more than the sum of the parts
- How did you solve a particular challenge in security metrics area?
- How your tool helps (not "can help"!) with collecting and analyzing security metric data?
  - Buy or build a correlation engine; when you write normalization and correlation rules, think about human behavior, not tools and code
- What metrics you use to determine that security controls are effective?
- How to track that your security is improving using metrics?

Sources and Further Reading

- Scott D. Applegate, "The Principle of Maneuver in Cyber Operations", George Mason University. http://gmu.academia.edu/ScottApplegate/Papers/1486050/The_Principle_of_Maneuver_in_Cyber_Operations
- Mike Cloppert, "Attacking the Kill Chain," SANS CFIR, 19 Oct 2009. http://computer-forensics.sans.org/blog/2009/10/14/security-intelligence-attacking-the-kill-chain/
- Jon Espenschied, "A Discussion of Threat Behavior: Attackers & Patterns" Microsoft Corporation and NATO CyCon, June 2012. http://www.microsoft.com/downloads/details.aspx?FamilyID=8fbbe2a9-a548-4c69-a6d3-0b04a39574ea or http://www.ccdcoe.org/cycon/doc/20120501-Espenschied-ThreatPatterns-public.pdf
- Matt Frazier, "Combat the APT by Sharing Indicators of Compromise," Mandiant Corporation. https://blog.mandiant.com/archives/766
- Geert Hofstede, "Culture's Consequences: Comparing Values, Behaviors, Institutions and Organizations Across Nations," Thousand Oaks CA: Sage Publications, 2001 (second edition).
- Geert Hofstede, Gert Jan Hofstede, Michael Minkov, "Cultures and Organizations: Software of the Mind," New York: McGraw Hill, 2010.
- Eric Hutchins, Michael Cloppert, Rohan Amin, Ph.D.; Lockheed Martin, "Intelligence-Driven Computer Network Defense Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains" http://papers.rohanamin.com/wp-content/uploads/papers.rohanamin.com/2011/08/iciw2011.pdf
- William J. Lynn III, "Defending a New Domain: The Pentagon's Cyberstrategy," Foreign Affairs magazine, Sept/Oct 2010. http://www.foreignaffairs.com/articles/66552/william-j-lynn-iii/defending-a-new-domain or http://www.dtic.mil/cgi-bin/GetTRDoc?AD=ADA527707
- William McRaven, "Spec Ops: Case Studies in Special Operations Warfare Theory & Practice," New York: Presidio Press (Random House Publishing Group), 1995.
- Mandiant Corporation, "M-Trends: The Advanced Persistent Threat" http://www.princeton.edu/ ~yctwo/files/readings/M-Trends.pdf
- Mitre Corporation, "Common Attack Pattern Enumeration and Classification (CAPEC)." http://capec.mitre.org/about/documents.html
- Char Sample, "Using Soft Markers in Attack Attribution," presented at Shmoocon 2012.
- Adam Shostack, "Security Breaches are Good for You," presented at Shmoocon 2007. http://www.homeport.org/~adam/Security%20Breaches%20are%20good%20for%20you.pdf
- Nart Villeneuve, "Trends in Targeted Attacks," Trend Micro, 2011. http://www.trendmicro.com/ cloud-content/us/pdfs/about/wp_trends-in-targeted-attacks.pdf
- Paul Wright, "Ten Stages of a Network Attack" (excerpt from "Oracle Forensics: Oracle Security Best Practices"). http://www.dba-oracle.com/forensics/t_forensics_network_attack.htm

Microsoft

Be what's next.™