# **Data Mining for Improving Intrusion Detection**

presented by: Dr. Eric Bloedorn

Team members: Bill Hill (PI) Dr. Alan Christiansen, Dr. Clem Skorupka, Dr. Lisa Talbot, Jonathan Tivel

12/6/00



### **Overview**

- Background and Motivation
- Our Focus:
  - Reduce burden on human analysts
- Our Approach
  - Anomaly Detection
  - Classification
    - Data Issues
    - Modeling Issues
- Lessons Learned



### **The Problem**

 Medium to large organizations are subject to constant attack by outsiders.

MITRE Data	Volume (every week)
6,000,000	sensor alarms.
125,000	"priority 1" alarms.
300	intrusion incidents.

- Just detecting that you are under attack requires significant effort from employees.
- Can data mining help with this problem?



## **Data Flow in the MITRE network**



### **Goal: Reduce Burden on Human Analyst**

- MITRE currently makes heavy use of human analysts in identifying real attacks from the large amount of data collected
- Standard procedure is to review the previous days sensor events in the morning
  - The large numbers of raw sensor events (most of which are uninteresting) make detecting real attacks or potential problems difficult
- Data Mining is not used to replace human analyst, but to reduce burden by allowing him to focus his expertise on those alarms most likely to be cause for real worry



### Approach

- Representational Issues
  - What is an incident?
  - How do we get incident labels?
  - How is an incident described?
- Data Mining Analysis
  - Anomaly Detection
    - Clustering-based
    - Rulequest's Gritbot
  - Classification
    - Identification of false alarms
    - 7-way classification





## What is an incident? The problem of Aggregation *and* Classification

- Problem: Data consists of individual sensor events (sensorlog database records) which need to be both aggregated into an incident and classified, but which do we do first?
- Approach:
  - Construct features for individuals that capture relationship to aggregate
    - how many other records have the same srcip as this record?
    - how many other records have the same srcip and dstport as this record?



# How is an incident labeled?

47	/iew <u>G</u> o	Comn	nunicati	or <u>H</u> e	lp a T		-	- (10) 20									
🜒 Back	Forward	Reloa	1 H H d	ne me	🧀 🟳 Search Guide	Strint Strint	ے Securi	ty Stop									N
🀠 🕻 Bo	okmarks	4	Netsite:	https:	//prowler.mitre.org/c	letection_(	data/t	agging/inci	dent_rep	ort.cfm							•
🖳 Alta\	/ista: Mair		rrs-w	eb Use	a Gu 🖳 MITRE I	nformati	🖳 M	lorningstar.	Net 🖳	) SmartUp	odate [	y wsec					
								INC	IDE	NT	REF	ORT					
		_															
ncident	ID: <mark>22 o</mark>	f 21			Date: 2000/0	1/28 00:44	:01	<b>⊡  </b>	svent Type	Informa	ation Gat	hering-Host Map	ping 🔽		Report Date: 2	2000/01/	28
										Informa	ition Gati	ering-Host Map ering-Service P	robing				
									Incid	10 Unauth Denial-	orized Ac of-Servic	cess Attempt e Attack					
										Relay/E Anamo	)ounce A Ious	ttack					1
																	1
			_						T.								
perator	Forrest	Carr	-		nvestigator: For	est Carr	-		Investi	igation S	tatus:	Jpen		cesolution D	ate: 2000-01-	-28	
										Resoluti	.on						
																	_
								Save	View I	ncident		eny Screen					
								RI	ELAT	ED Q	UER	IES					
								Incl	ude? 🛛	)ueryD	Query	Name					
								N	o Candi	idate Que	eries Fou	and!					
		RELATED RECORDS															
								Event	Proto	Src Port	Dst Port	Src IP	Src DNS	Dst IP	Dst DNS	Info Link	
	Include?	Query	Sensor	Туре	Start Date	End Date	Pn		-	22	1009	199.94.100.49	extranet-ca.mitre.org		duna mita ana		
	Include?	Query	Sensor bt1	<b>Type</b> RS	Start Date 2000/01/28 00:44:01	End Date	1	Port_Scan	TCP	44			-		Graim mile.org		
	Include?	<b>Query</b> 1 1	Sensor bt1 bi1	Type RS RS	Start Date 2000/01/28 00:44:01 2000/01/28 00:44:02	End Date	1 1	Port_Scan Port_Scan	TCP TCP	22	1009	199.94.100.49	extranet-ca.mitre.org		chann mitre.org		
	Include?	<b>Query</b> 1 1 1 1	Sensor bt1 bi1 we1	Type RS RS RS	Start Date           2000/01/28 00:44:01           2000/01/28 00:44:02           2000/01/28 03:33:42	End Date	1 1 1	Port_Scan Port_Scan UDPBomb	TCP TCP UDP	22 22 62064	1009 8	199.94.100.49 158.152.1.58	extranet-camitre.org cache-1 ns.demon.net	 	chann mitre.org		
	Include?	<b>Query</b> 1 1 1 1	Sensor bt1 bi1 we1	Type RS RS RS	Start Date 2000/01/28 00:44:01 2000/01/28 00:44:02 2000/01/28 03:33:42	End Date	1       1       1	Port_Scan Port_Scan UDPBomb	TCP TCP UDP	22 22 62064	1009	199.94.100.49 158.152.1.58	extranet-ca mitre.org cache-1 ns.demon.net		chann mitre.org mwunix mitre.org		



### How is an <u>event</u> described?

- Base collected by network sensors
  - examples: date, type of sensor, protocol, srcip, dstip, srcport, dstport
- Incident relationship to known security incidents
  - example: has this srcip/dstip been listed in an incident recently?
- Record data lookups specific to a single record
  - example: duration, endtime, starttime, highport, srczone, hostsrcip
- Host data related to the source or destination host
  - example: #alarms with same srcip &dstip, #other alarms with same srcip
- Time Window statistics gathered over time
  - example: avg. time between connections for a srcip or dstip

### **Clustering - Anomaly Detection**

- Preprocessing and clustering: SPSS's Clementine
  - Aggregate by Julian date, source port, source IP
  - K-means with many K's.
- Outlier criteria:
  - 1. Maximum sum of distances to cluster centroids.
  - 2. Records from clusters containing few records.
- Analysis and presentation of results:
  - Present cluster model as a table with highlighted attribute importance, α,

$$\alpha = x_1 - \sum_{j=2}^{J} (.5)^{j-1} x_j$$

Extract potentially anomalous records, append clustering results.

### **Clustering - Anomaly Detection : Maximum Total Distance**

Idea: Records with high total distance are likely outliers. Top 5 records in terms of maximum total distance.

	recid	325500322	305979032	312114228	315380855	312921360
	Cluster	cluster-2	cluster-2	cluster-2	cluster-2	cluster-5
K=5	Distance	3.03531	3.37275	3.20283	3.18285	2.99647
	Rank	31	0	1	16	35
	Cluster	cluster-2	cluster-7	cluster-7	cluster-7	cluster-1
K=8	Distance	3.08867	3.07683	2.95816	2.94178	2.90891
	Rank	0	1	2	9	22
	Cluster	cluster-1	cluster-1	cluster-1	cluster-1	cluster-1
K=12	Distance	2.48668	2.06143	2.17984	2.21128	2.37171
	Rank	7	247	127	95	26
	Distance	8.61	8.51	8.34	8.34	8.28



## **Clustering - Anomaly Detection : Low-record Clusters**

Idea: An entire cluster may be an outlier. Such a cluster may contain few examples and be distant from other clusters.

	Number of examples assigned to clusters											
Cluster Number												
	1	2	3	4	5	6	7	8	9	10	11	12
K=5	359	477	150	30	1400							
K=8	849	327	8	24	550	359	149	150				
K=12	26	186	1	24	12	359	130	141	29	1382	6	120
Mi	Minimum distance from cluster centroid to neighboring cluster centroid											
					Clu	ster N	lumbe	er 🛛				
	1	2	3	4	5	6	7	8	9	10	11	12
K=5	2.55	2.03	2.33	2.21	2.03							
K=8	1.11	1.90	2.41	2.40	1.11	2.28	1.90	2.40				
K=12	1.85	1.18	2.62	2.43	1.79	2.55	1.86	2.20	1.79	1.85	2.52	1.18

### **Clustering - Finding Classes in Suspicious Records**

# Idea: Compact, well-separated clusters may define additional classes of interest in anomalous record data.

From the histogram: Records in cluster 3 (for K=5) appear to be compact. From the table: Clusters 2, 3, and 5 are most distant.



### **Anomaly Detection**

- Motivation: Unusual activity is suspicious
- Goal: To automatically identify *interesting* anomalous behavior
- Approach:
  - Use sensor log events not identified as incidents
  - Filter attributes based on analyst feedback
  - Build web interface for easy viewing of generated anomalies
  - Classify anomalies into incident categories



# **Interface for viewing Anomalies**

🗱 Rulo 1 - No	atecano														
File Edit View Go Communicator Help															
Back Fo	nward Relo	) 🔏	ne S	earch G	42 Juide	de Serie Ser	💕 Security	y Sto	ip						N
😻 Bookmarks 🙏 Netsite: http://nemesis.mitre.org/ids/msr/tasks/gritbot/10.30.run01/rule1.html 🔽															
📔 🖆 IDS 🖆	🖆 IDS 🖆 Netscape defaul 🖳 Welcome to Liqu 🖳 Liquid Music Ne														
Anomaly #1. 8 highdstport = 0 datatype = SN srctld = com	3 case(s). Si 0 (13312 ca IT	gnficance ses, 99.7:	level: 0. 9% `1')	002											
recids	sensorname	datatype	priority	protocol	duration	stdow	stblock	sthour	disctime	srcsubnet	dstsubnet	broadcast	sameport	highdstport	lowsro
315545238 🛛	wel	SNT	1	TCP	?	SAT	nbd	10	?			0	0	0	0
315563818 v	wel	SNT	1	TCP	?	SAT	nbd	10	?			0	0	0	0
315590428 v	wel	SNT	1	TCP	?	SAT	nbd	12	?			0	0	0	0
315590458 v	wel	SNT	1	TCP	?	SAT	nbd	15	?			0	0	0	0
317623161 v	wel	SNT	1	TCP	?	MON	nbd	08	?			0	0	0	0
326956530 v	wel	SNT	1	TCP	?	MON	nbd	08	?			0	0	0	0
317623813 v	wel	SNT	1	TCP	?	MON	nbd	15	?			0	0	0	0
326957182 v	wel	SNT	1	TCP	?	MON	nbd	15	?	210.15.171	120.22.124	0	0	0	0
			lono												

### **Example Anomalies**

### Anomaly #14. 3 case(s). Significance level: 0.015 highdstport = no (281 cases, 98.6% `yes`) synflag = no

130330539,we1,log,2000/02/13,2000,02,13,14,38,46,sun,bus,?,?,?,?,?,3,netbios-ns,tcp,23,137,206.184.139.134,192.47.242.29, r,2451588,in,no,no,no, ?,no,no,no

### • Looks like significant dns queries to all MITRE dns servers, and many MITRE internal hosts. Possible scan attempt to bypass firewall?

# Anomaly #32. 4 case(s). Signficance level: 0.004 srcmitre = no (1692 cases, 99.65% `yes`) dstip = 192.188.104.221

### Looks like scanning for ftp servers



### **Classifying Anomalies**

Classlabel	#Anomalies	Classlabel	#Anomalies
Suspicious	42	Denial of Service	0
IP map	16	Relay 1 bounce attack	0
Port Scan	39	False Alarm	22
Unauthorized Access	0		

• Decision Tree (99% training set accuracy) used here was trained on the same month as the data used for generating anomalies (September)



### **Anomaly Detection Example**

### A "False Alarm" anomaly

Anomaly #3. 22 case(s) Significance level: 0.004 datatype = RS (12116 cases, 99.61% 'SNT') stblock = eve and srczbetween3600 <= -0.18

### • Rule used to classify as a "false alarm":

If the srcip is in the dmz and the average time between connections from that srcip within the last 2 weeks is about 1 min, and the average duration of connections from that srcip in the last week is <= 0.6 sec and the priority is 1 THEN False Alarm

What the analyst says: This was an internal scan by our own security folks



## **Data Collection**

### • Seven classes of incidents:

Class	Aug. counts	Sept. counts
Anomalous	2,649	1,888
IP map	2,864	8,667
Port Scan	502	4,652
<b>Unauthorized Access</b>	123	683
<b>Denial of Service</b>	0	0
Relay Attack	0	8
False Alarm	6,762	987
Total	12,900	16,885



### **Need for HOMER**

- IP Mapping Episodes are:
  - frequent (average # / day: 52)
  - large (average sensor events / episode: 6485)
- IP Mapping episodes interfere with manual review of other types of activity
- HOMER detects IP mapping activity automatically
- Humans can focus on more subtle activity



### **The HOMER Heuristic**

Aggregate All Sensor Events by

- Source IP
- Destination Port
- Protocol
- Time Window (currently 1 week)
- Count Number of Distinct Destination Hosts Hit Within Each Aggregate
- Record an IP Mapping Incident When a Threshold is Exceeded (here it is set to 100)



### **HOMER Statistics**

- First 30 days of operations
- Total priority 1 Incidents Discovered: 102
- Total priority 1 events in these incidents: 708,162
- Percentage of total priority one events: 90.2%
- Precision (all priority): 99.95%
- Recall (all priority): 99.5%
- Impact:
  - Analysts can create the 30 remaining incident reports for priority 1 alarms.

### • Future:

All priority 1 events will be classified automatically.
 Analysts will review incident reports for accuracy.



### **Classification Models: False Alarms vs. All other**

- High predictive accuracy for initial model: 96%
- If srczone == boundary and fscan600 == 0 then False Alarm (523, 0.996)
- If the machine is on the boundary to the internet and the srcip has not hit a large number of ports on the dst machine in a 10 minute window then False Alarm
- If srczone == internal and priority==1 and srcstdbetween1209600 > 27423.5 (395.0, 1.0)
- If the src machine is internal and it's a high priority alarm and the standard deviation in times between connections from that src in the last week is > about 8 hours then False Alarm



### **Classification Models: Multiclass**

- Lower predictive accuracy: only 72% on holdout set
- Great deal of confusion between mapping and scanning classes so adding new attributes
- Coincidence Matrix

#### **\$C-classlabelID**

	1	2	3	4	6	7	?
1	2223	342	344	10	0	645	54
2	1408	11645	1308	64	0	62	488
3	733	843	1046	0	0	13	29
4	248	311	410	0	0	2	0
6	1	0	0	0	12	0	0
7	131	74	29	0	0	5082	5



### **Lessons Learned**

- Importance of good representation
  - solves aggregation and classification problem (?)
- Synergy between analysis methods
  - classification aids in finding interesting anomalies
- Goal of Data Mining is to reduce burden on human analysts not to replace them
- Initial work
  - data is based on alarms so although we generalize we do not predict very new types of attacks
  - Need to supplement with host information maybe as a new attribute (unusual use of port)?



# Backup



### Sensors

- JIDS: Looks for "dirty words" in telnet and ftp sessions like 'su root',' '/etc/passwd'
- Real Secure (www.iss.net) Uses one or more pc-based sensor "engines", which are essentially packet sniffers with packet filtering and collection rules, and some alarm logic.
- Snort (www.snort.org) is a rule-based packet sniffer.
- Logger is another packet sniffer Derived from Shadow (developed by the Naval Surface Warfare Center IDS group).
- Syslog: unix activity logger. Used to capture messages from firewalls and other devices.
- TripWire (www.tripwiresecurity.com). Host-based software to monitor accesses and changes to critical system files and changes in user privilege



### **Miscellaneous Comments**

Focus on Network-based intrusion detection (not host)
Based on data from multiple sensors

- CVE (Common Vulnerabilities and Exposures) cve.mitre.org
  - A list of standardized names for vulnerabilities and other information security exposures. CVE aims to standardize the names for all publicly known vulnerabilities and security exposures.

 CIEL now in its infancy, is basically addressing the same problem for IDS. Note that CVE Vunerabilities and Exposures do not all necessarily map directly to attack signatures, but certainly a significant number would.

