Adversarial Analytics 101

Strata Conference & Hadoop World

October 28, 2013

Robert L. Grossman

Open Data Group

Examples of Adversarial Analytics

- Compete for resources
  - Low latency trading
  - Auctions
- Steal someone else's resources
  - Credit card fraud rings
  - Insurance fraud rings
- Hide your activity
  - Click spam
  - Exfiltration of information

Conventional vs Adversarial Analytics

	Conventional Analytics	**Adversarial Analytics**
Type of game	Gain / Gain	Gain / Loss
Subject	Individual	Group / ring / organization
Behavior	Drifts over time	Sudden changes
Transparency	Behavior usually open	Behavior usually hidden
Automation	Manual (an individual)	Sometimes another model

1. Points of Compromise

Source: Wall Street Journal, May 4, 2007

Case Study: Points of Compromise

- TJX compromise
- Wireless Point of Sale devices compromised
- Personal information on 451,000 individuals taken
- Information used months later for fraudulent purchases.
- WSJ reports that 45.7 million credit card accounts at risk
- "TJX's breach-related bill could surpass $1 billion over five years"

Source: Wall Street Journal, May 4, 2007

UNITED STATES DISTRICT COURT
DISTRICT OF MASSACHUSETTS

UNITED STATES OF AMERICA	)	Criminal No.
	)	
v.	)	VIOLATIONS:
	)	18 U.S.C. § 371 (Conspiracy)
ALBERT GONZALEZ,	)	18 U.S.C. § 1030(a)(5)(A)(i) (Damage to Computer
a/k/a cumbajohny, a/k/a cj,	)	Systems)
a/k/a UIN 201679996, a/k/a	)	18 U.S.C. § 1343 (Wire Fraud)
UIN 476747, a/k/a soupnazi,	)	18 U.S.C. § 1029(a)(3) (Access Device Fraud)
a/k/a segvec, a/k/a k1ngchilli,	)	18 U.S.C. § 1028A (Aggravated Identity Theft)
a/k/a stanozololz,	)	18 U.S.C. §§ 1029(c)(1)(C), 982(a)(2)(B), 981(a)
	)	(1)(C), 28 U.S.C. §2461(c) (Criminal Forfeiture)
Defendant.	)	

INDICTMENT

COUNT ONE
(Conspiracy)
18 U.S.C. § 371

The Grand Jury charges that:

1. From approximately 2003 through 2008, in the Southern District of Florida, the District of Massachusetts, Eastern Europe and elsewhere, ALBERT GONZALEZ, Christopher

The Grand Jury in and for the District of New Jersey, sitting at Newark, charges:

COUNT 1
(Conspiracy)
18 U.S.C. § 371

1. At various times relevant to this Indictment:

The Defendants

a. Defendant Albert Gonzalez, a/k/a "segvec," a/k/a "soupnazi," a/k/a "j4guar17" ("GONZALEZ"), resided in or near Miami, Florida.

b. Defendant HACKER 1 resided in or near Russia.

c. Defendant HACKER 2 resided in or near Russia.

Coconspirator

d. P.T., a coconspi

defendant herein, resided in or

in or near Miami, Florida.

The adversary is often a group, criminal ring, etc.

It’s Not an Individual…

Source: Justice Department, New York Times.

Gonzalez Tradecraft

1. Exploit vulnerabilities in wireless networks outside of retail stores.
2. Exploit vulnerabilities in databases of retail organizations.
3. Gain unauthorized access to networks that process and store credit card transactions.
4. Exfiltrate 40 Million track 2 records (data on credit card's magnetic strip)

Gonzalez Tradecraft (cont'd)

5. Sell track 2 data in Eastern Europe, USA, and other places.
6. Create counterfeit ATM and debit cards using track 2 data.
7. Conceal and launder the illegal proceeds obtained through anonymous web currencies in the US and Russia.

Source: US District Court, District of Massachusetts, Indictment of Albert Gonzalez, August 5, 2008.

Data Is Large

- TJX compromise data is too large to fit into memory
- Data is difficult to fit into database
- Millions of possible points of compromise
- Data must be kept for months to years


Maximum Security For Your Wine
For Stars, A Sticky Issue Of Sequels
Bike-Friendly Cities Gain A Following

THE WALL STREET JOURNAL.

What's News—

As Market Cools, Home Buyers Seek a Way Out

Builders Face Lawsuits, A Rash of Complaints; The $266,000 Refund

How Credit-Card Data Went Out Wireless Door

In Biggest Known Theft, Retailer's Weak Security Lost Millions of Numbers

Lawless Legislators Thwart Social Progress in India

Malnutrition, Polio; A 'Superior' Jail Cell For Mukhtar Ansari

'Friends' Like These Help Candidates Woo the Young

Some nice guys do finish first.

Saturn AURA. North American Car of the Year.

Source: Wall Street Journal, May 4, 2007

2. Introduction to Adversarial Analytics

Conventional Analytics

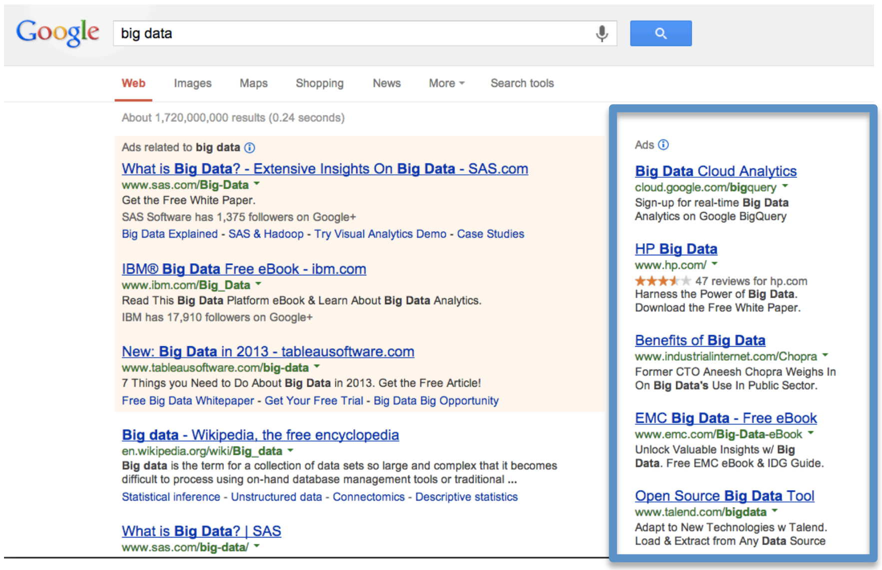

- Example: predictive model to select online ads.
- Example: predictive model to classify sentiment of a customer service interaction.

Adversarial Analytics

- Example: commercial competitor trying to maximize trading gains.
- Example: criminal gang attacking a system and hiding its behavior.

What's Different About the Models?

	Conventional Analytics	**Adversarial Analytics**
Data	Labeled	Unlabeled
Data size	Small to large	Small to very large
Model	Classification / Regression	Clustering, change detection
Frequency of update	Monthly, quarterly, yearly	Hourly, daily, weekly, …

Types of Adversaries

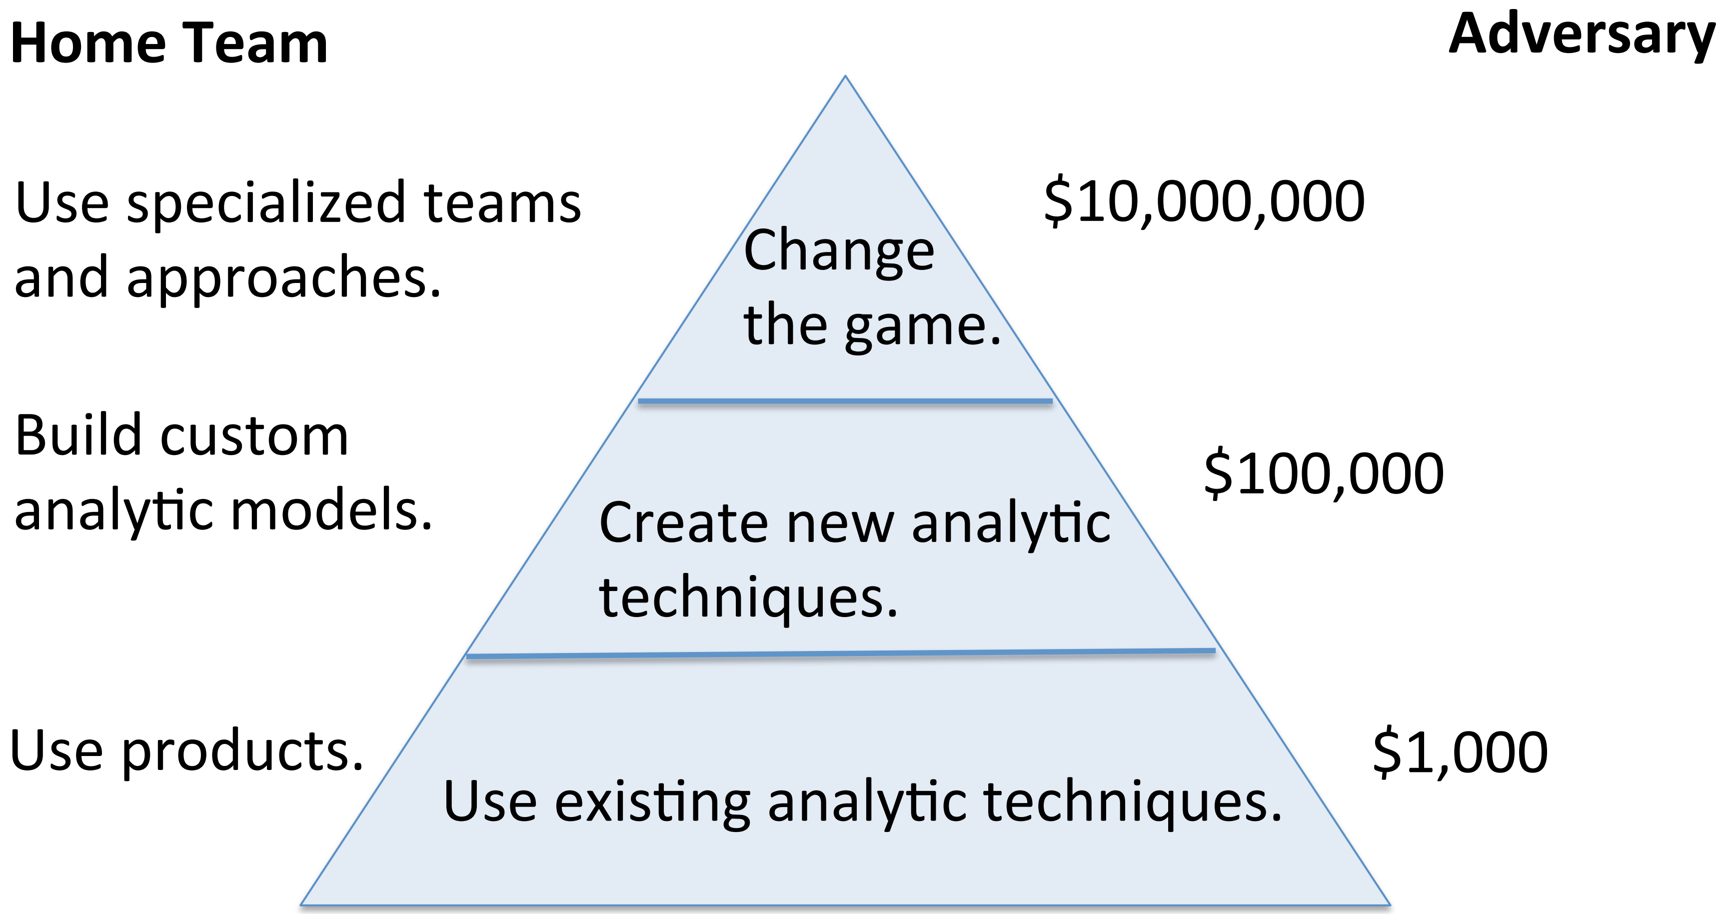

Source: This approach is based in part on the threat hierarchy in the Defense Science (DSB) Report on Resilient Military Systems and the Cyber Threat, January, 2013

What is Different About the Adversary?

	Conventional Analytics	**Adversarial Analytics**
Entity	Person browsing	Gang attacking a system
What is modeled?	Events, entities	Events, entities, rings
Behavior	Component of natural behavior	Waiting for opportunities
Evolution	Drift	Active change in behavior
Obfuscation	Ignore	Hide

Updating Models

	Conventional Analytics	**Adversarial Analytics**
When do you update model?	When there is a major change in behavior	Frequently, to gain an advantage.
Process for updating models	Automation and analytic infrastructure helpful.	Automation and analytic infrastructure essential.

3. More Examples

Click Fraud

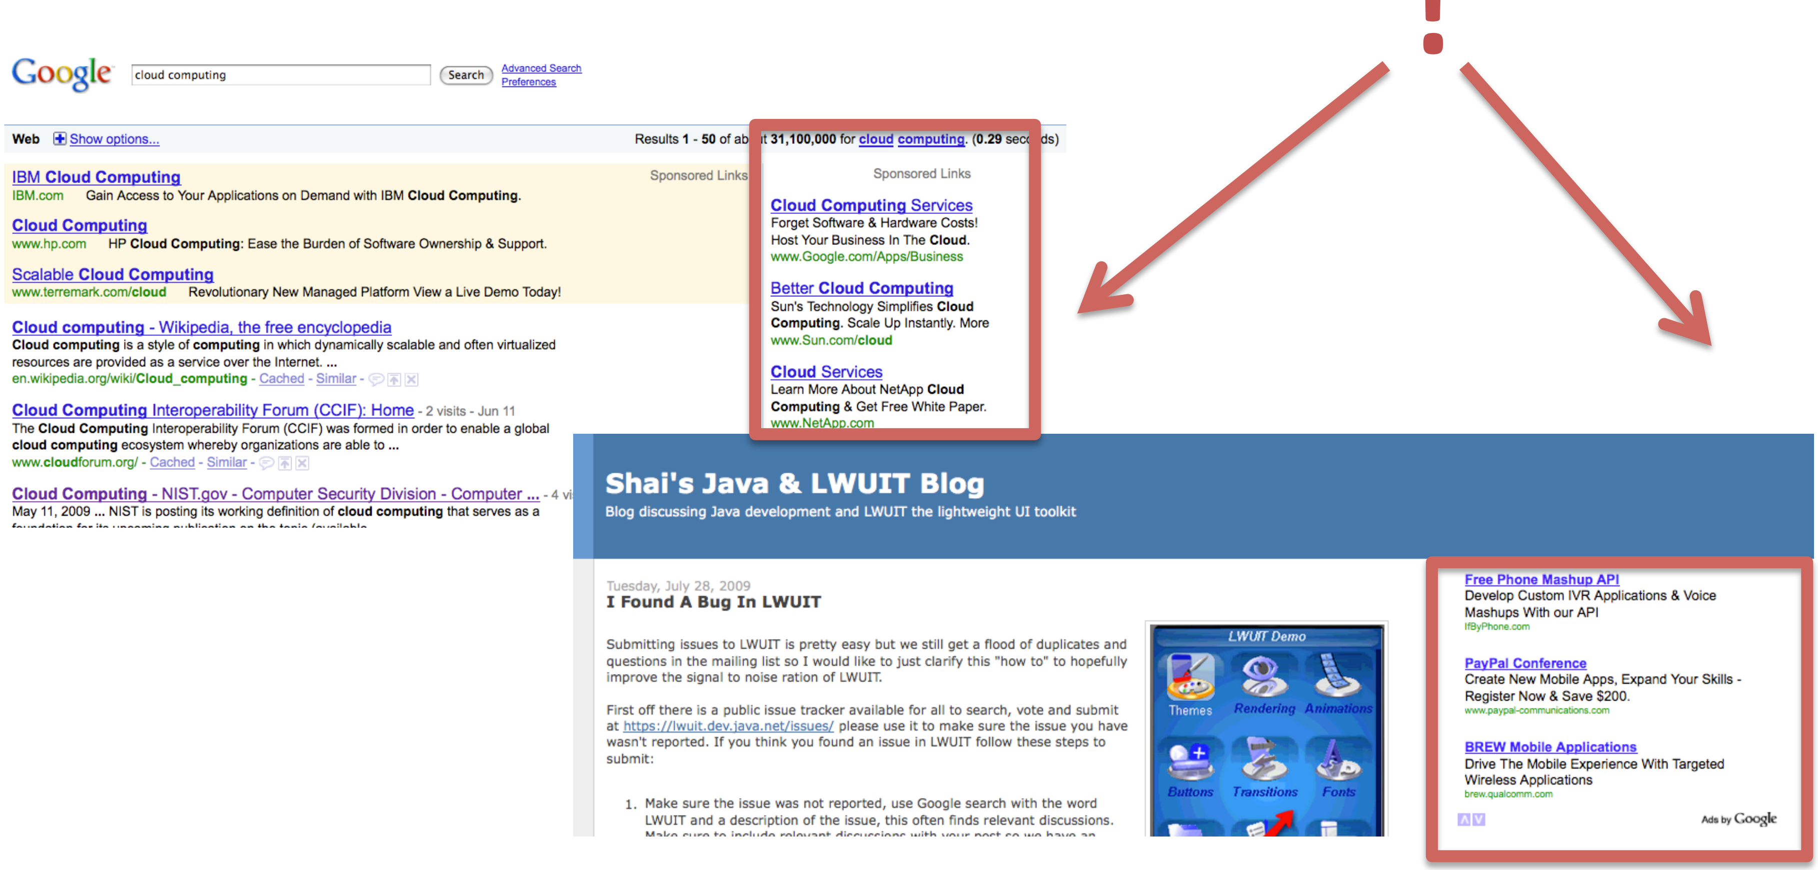

- Is a click on an ad legitimate?
- Is it done by a computer or a human?
- If done by a human, is it an adversary?

Types of Adversaries

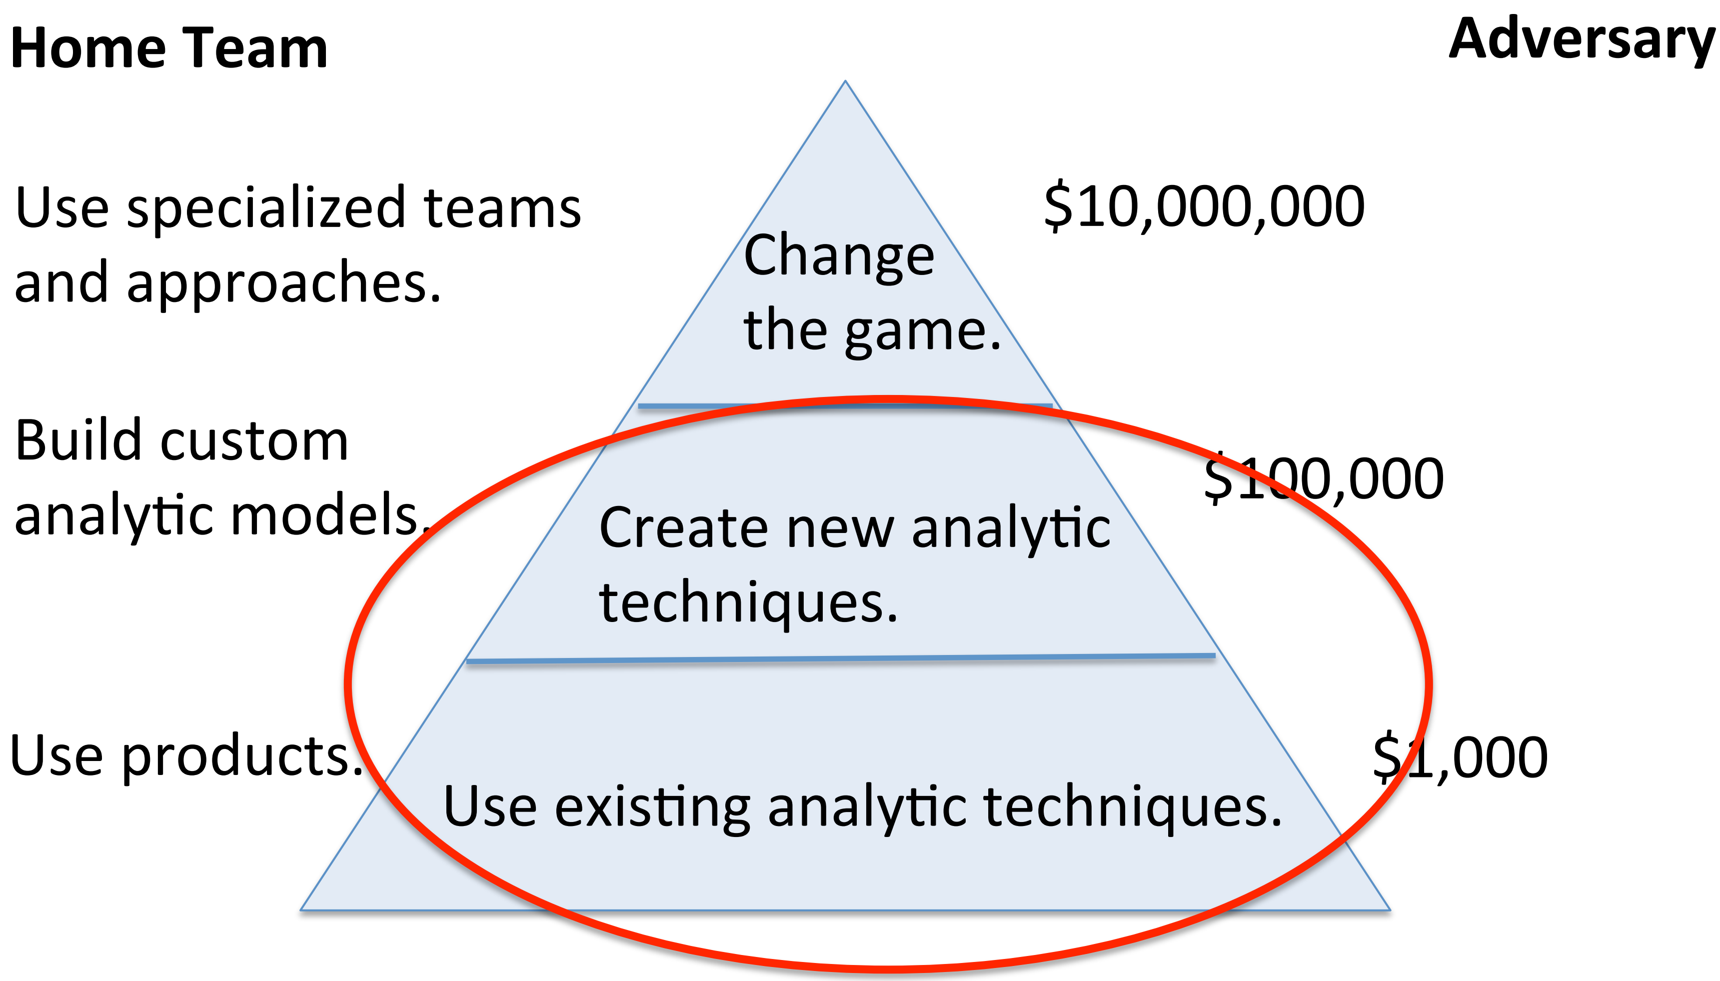

Source: This approach is based in part on the threat hierarchy in the Defense Science (DSB) Report on Resilient Military Systems and the Cyber Threat, January, 2013

Low Latency Trading

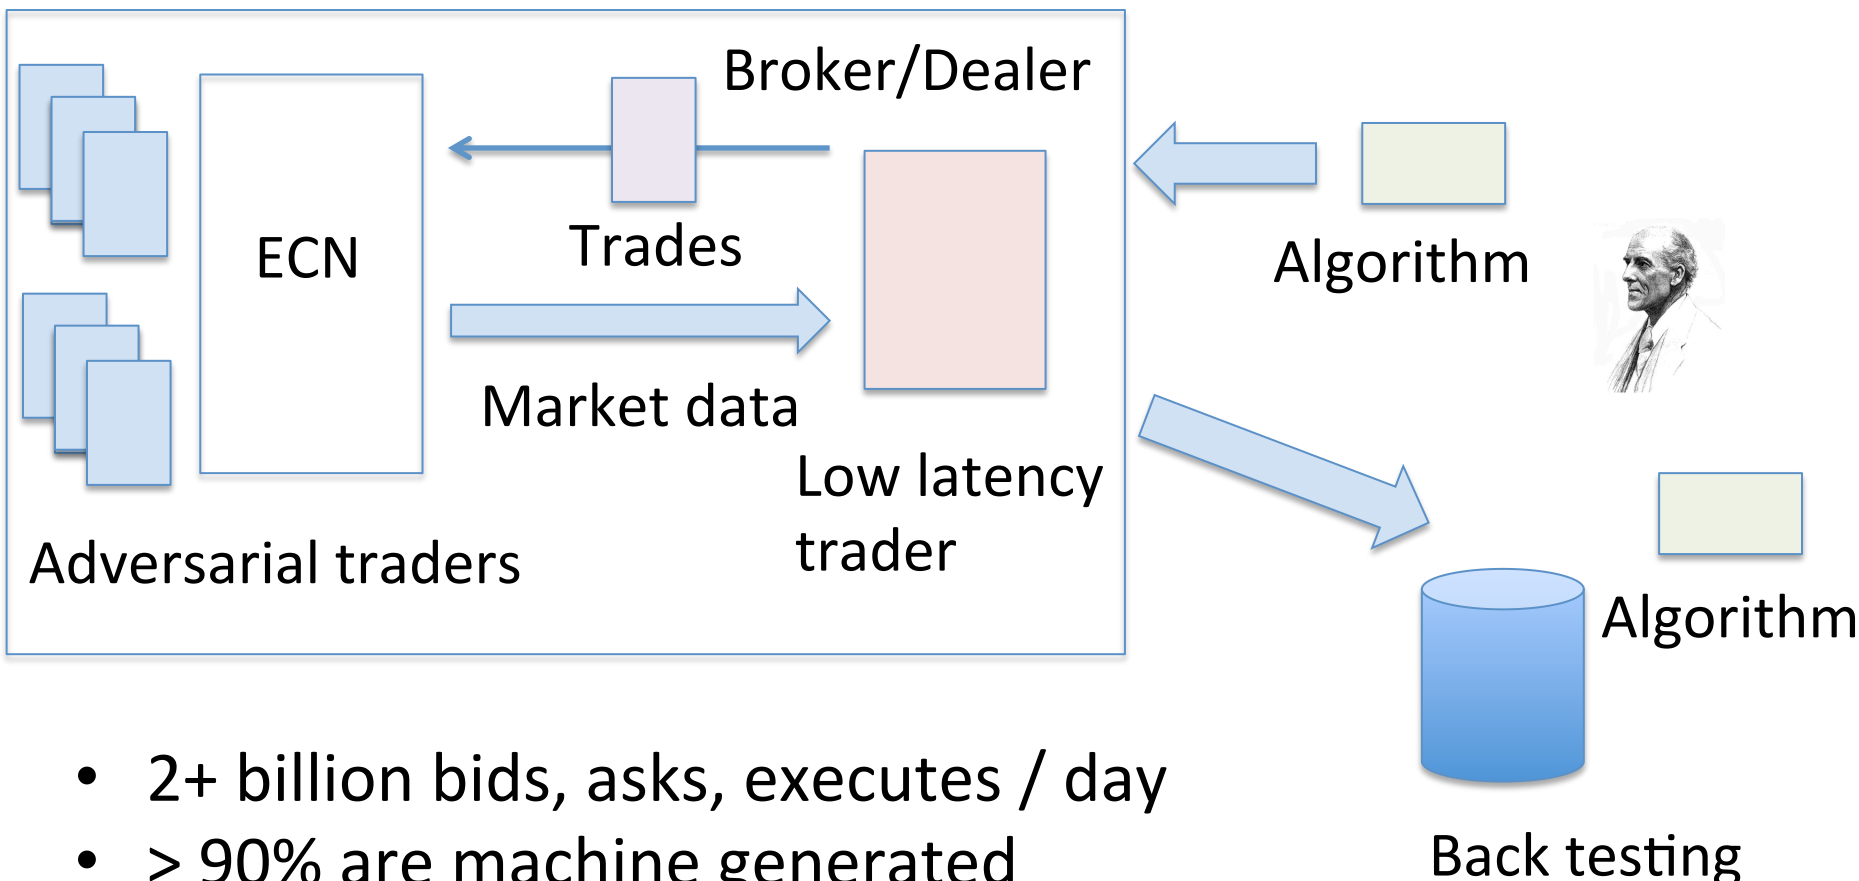

- 2+ billion bids, asks, executes / day
- > 90% are machine generated
- Most are cancelled
- Decisions are made in ms

Connecting Chicago & NYC

Technology	Vendor	Round Trip Time (ms)
Microwave	Windy Apple	9 .0
Dark fiber, shorter path	Spread Networks	13.1
Standard fiber, ISP	Various	14.5+

Source: lowlatency.com, June 27, 2012; Wired Magazine, August 3, 2012

Types of Adversaries

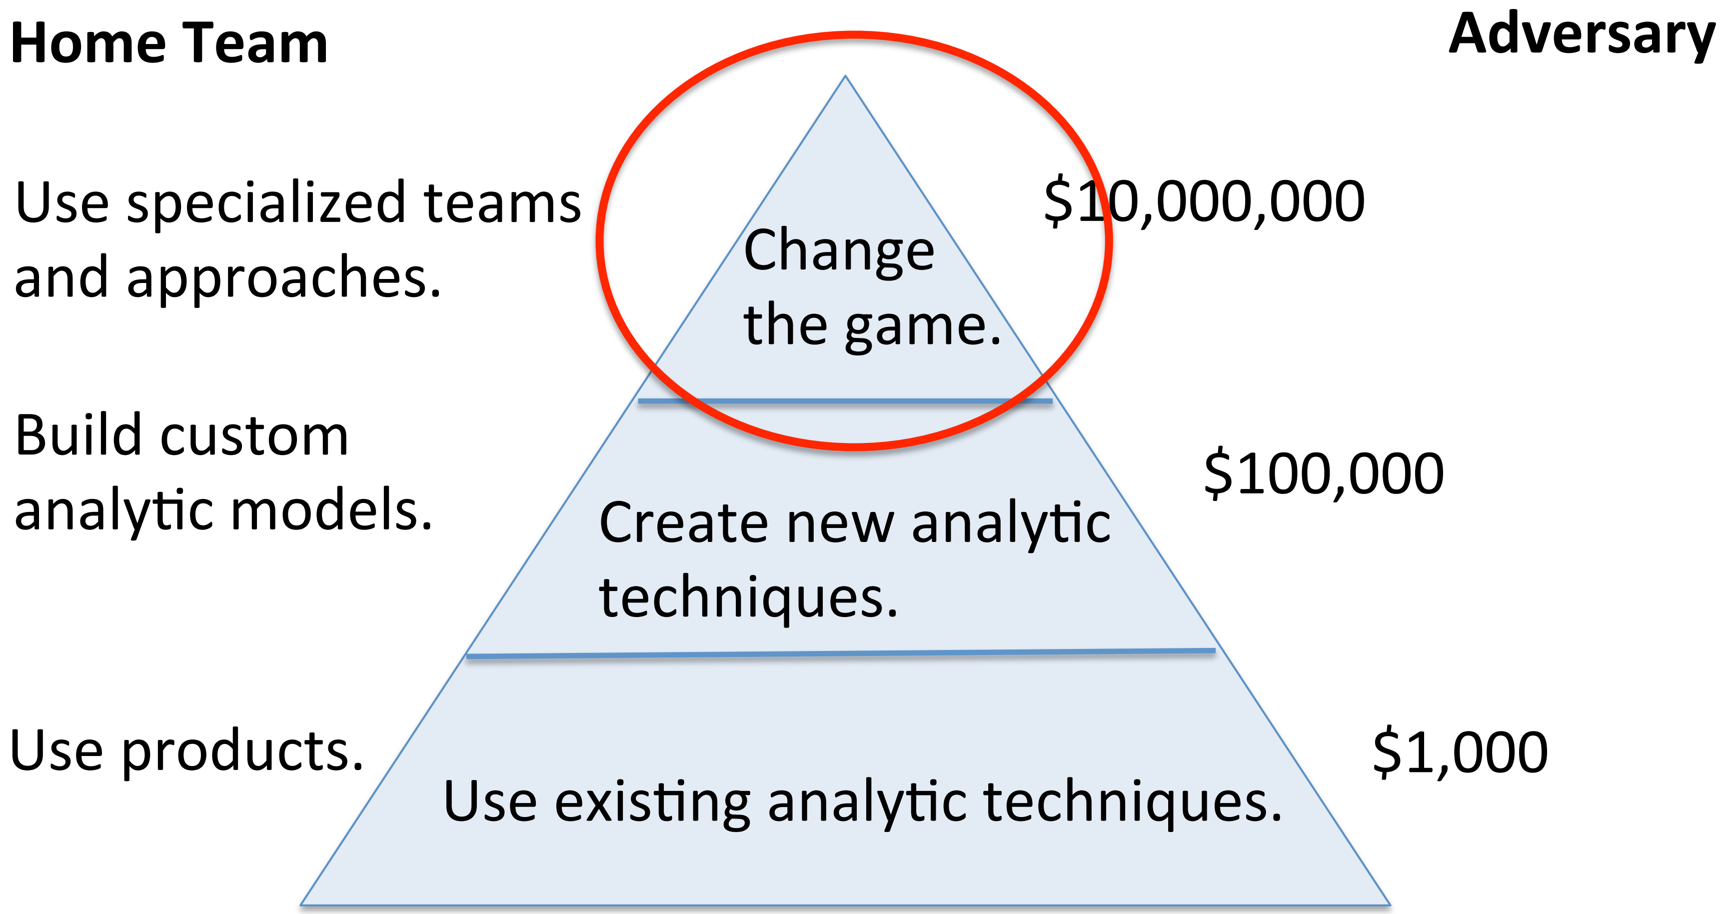

Source: This approach is based in part on the threat hierarchy in the Defense Science (DSB) Report on Resilient Military Systems and the Cyber Threat, January, 2013

Example: Conficker

Numbers of infections:

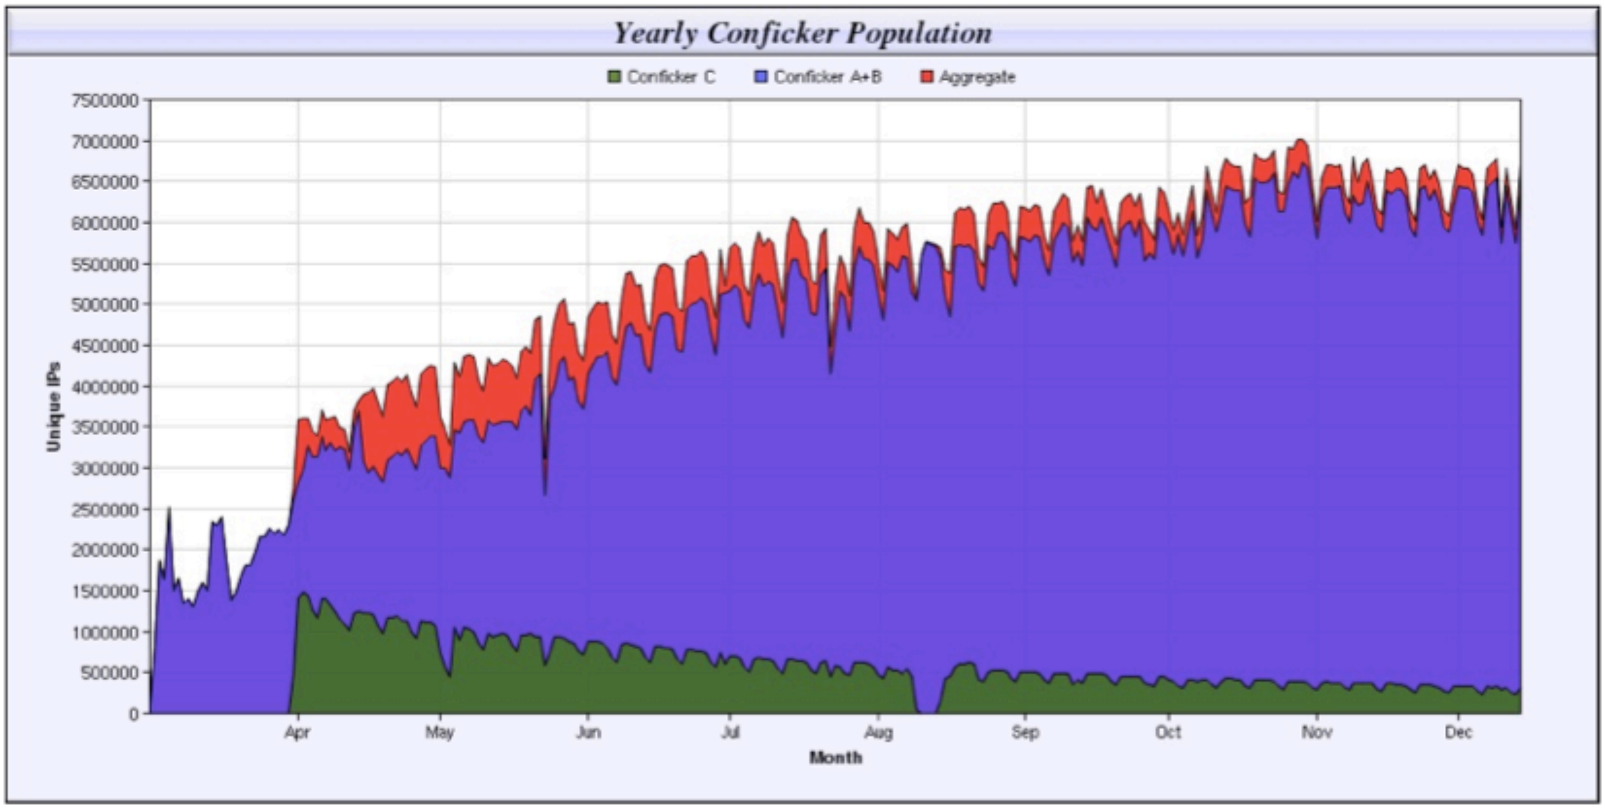

http://www.shadowserver.org/wiki/uploads/Stats/conficker-population-year.png

Source: Conficker Working Group: Lessons Learned, June 2010 (Published January 2011),www.confickerworkinggroup.org

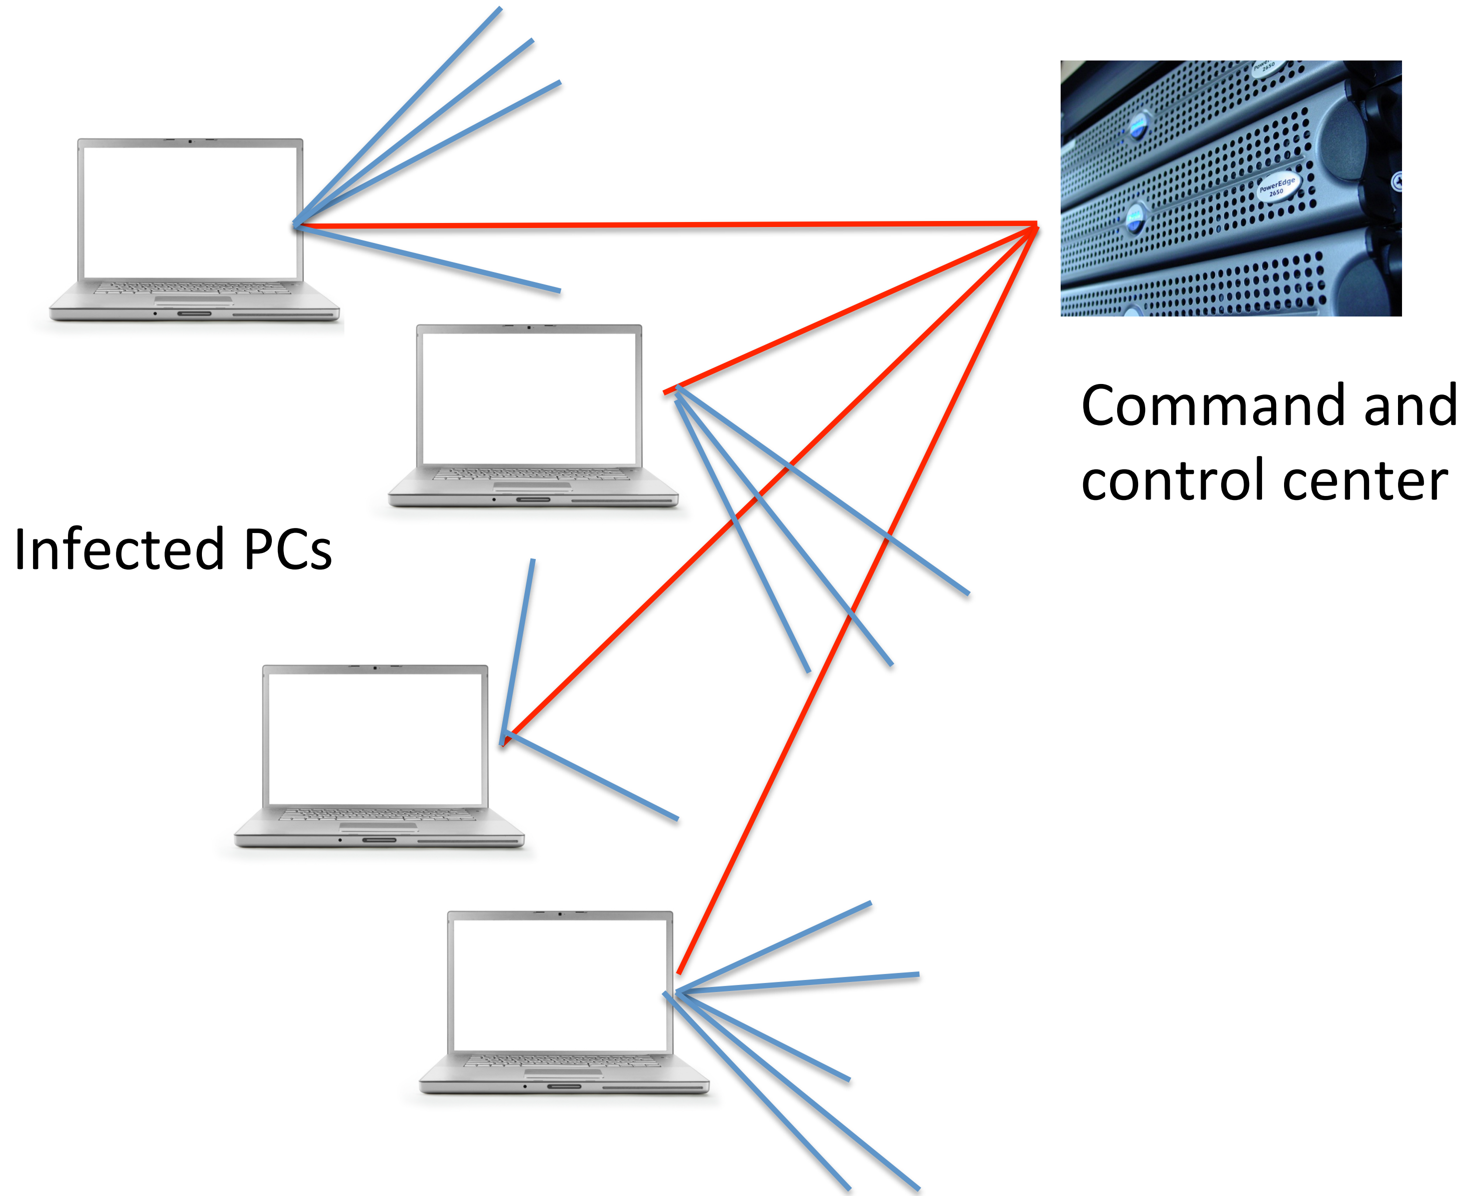
Command and
control center
Infected PCs

Variant	Detection Date	Infection Vectors	Update Propagation	End Action
Conficker A	21-Nov-08	Net BIOS; Exploits MS08-067 vulnerability in Server service	HTTP pull; Downloads from trafficconverter.biz; Downloads daily from any of 250 pseudorandom domains over 5 TLDs	Updates self to Conficker B, C or D
Conficker B	29-Dec-08	NetBIOS; Exploits MS08-067 vulnerability in Server service; Creates DLL-based AutoRun trojan on attached removable drives	HTTP pull; Downloads daily from any of 250 pseudorandom domains over 8 TLDs; NetBIOS push	Updates self to Conficker B++ or E
Conficker B++	20-Feb-09	NetBIOS: Exploits MS08-067 vulnerability in Server service; Creates DLL-based AutoRun trojan on attached removable drives	Blocks a selective list of DNS lookups to prevent remediation; Disables AutoUpdate	Updates self to Conficker C

Source: Conficker Working Group: Lessons Learned, June 2010 (Published January 2011),www.confickerworkinggroup.org

Variant	Detection Date	Infection Vectors	Update Propagation	End Action
Conficker C	4-Mar-09	HTTP pull; Downloads daily from any 500 of 50000 pseudorandom domains 110 TLDs; P2P push/pull; Uses custom protocol to scan for infected peers via UDP, then transfer via TCP	Blocks DNS lookups; Does an in-memory patch of DNSAPI.DLL to block lookups of anti-malware related web sites; Disables Safe Mode; Disables AutoUpdate; Kills anti-malware; Scans for and terminates processes with names of anti-malware, patch or diagnostic utilities at one-second intervals	Downloads and installs Conficker E
Conficker E	7-Apr-09	NetBIOS; Exploits MS08-067 vulnerability in Server service	NetBIOS push; Patches MS08-067 to open reinfection backdoor in Server service; P2P push/pull; Uses custom protocol to scan for infected peers via UDP, then transfer via TCP	Updates local copy of Conficker C to Conficker D; Downloads and installs malware payload: Waledac spambot; SpyProtect 2009 scareware;

Source: Conficker Working Group: Lessons Learned, June 2010 (Published January 2011),www.confickerworkinggroup.org

Types of Adversaries

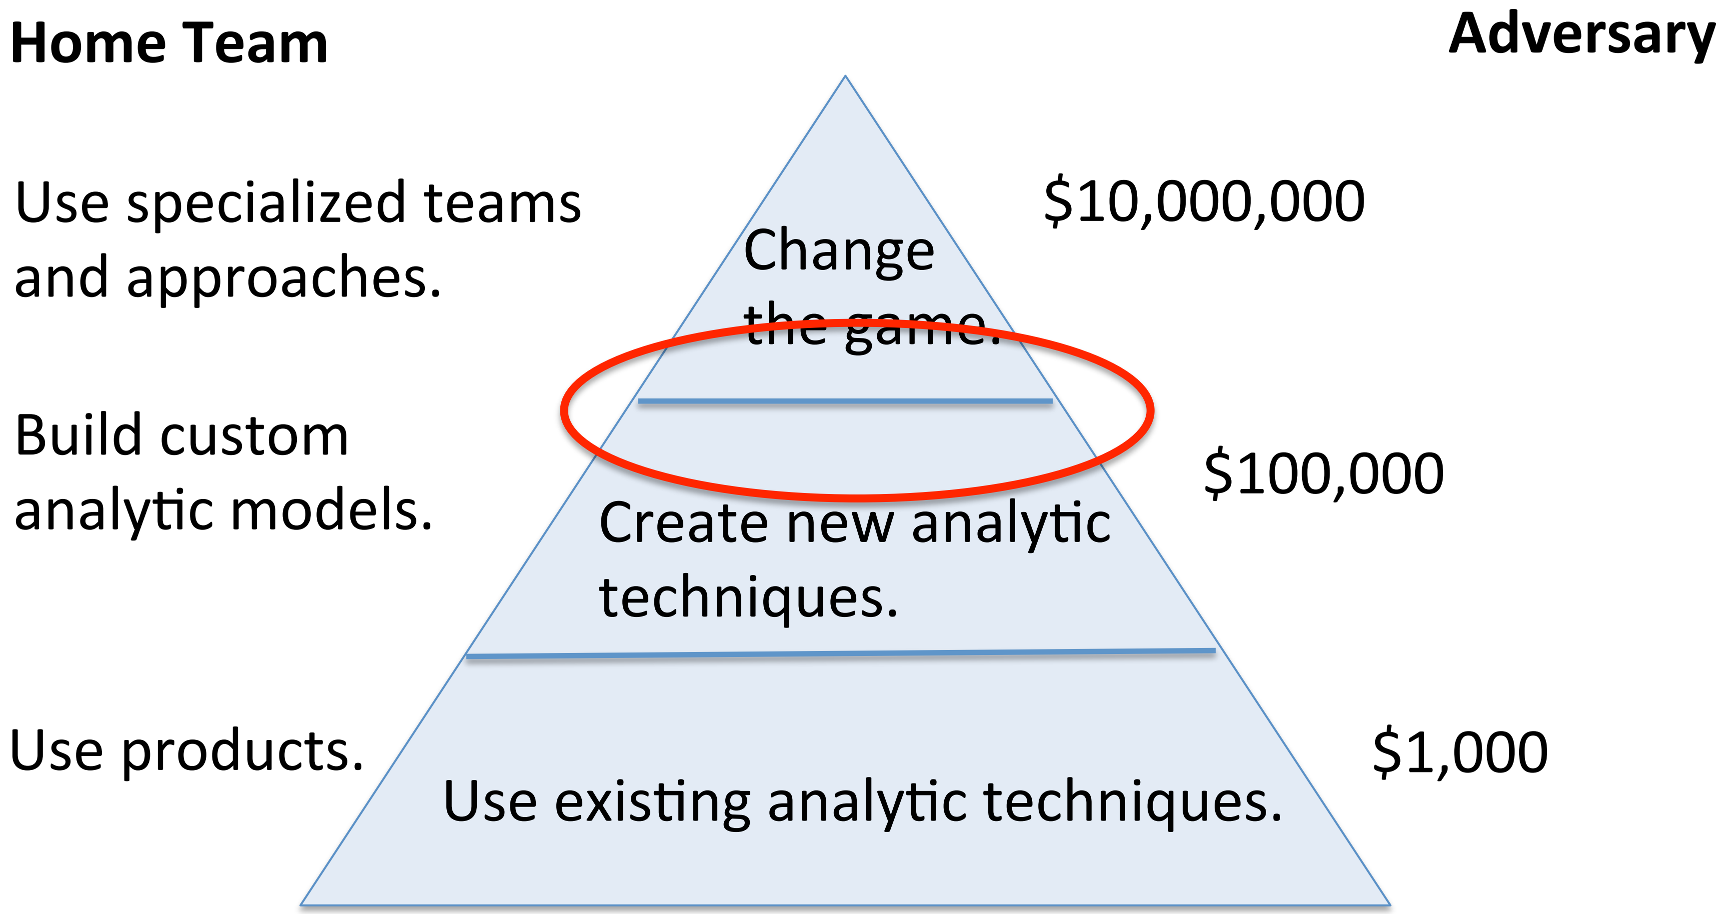

Source: This approach is based in part on the threat hierarchy in the Defense Science (DSB) Report on Resilient Military Systems and the Cyber Threat, January, 2013

4. Building Models for Adversarial Analytics

Building Models To Identify Fraudulent Transactions

Building the model:

1. Get a dataset of labeled data (i.e. fraud is labeled).
2. Build a candidate predictive model that scores each transaction with likelihood of fraud.
3. Compare lift of candidate model to current champion model on hold-out dataset.
4. Deploy new model if performance is better.

Scoring transactions:

1. Get next event (transaction).
2. Update one or more associated feature vectors (account-based)
3. Use model to process updated feature vectors to compute scores.
4. Post-process scores using rules to compute alerts.

Building Models to Identify Fraud Rings

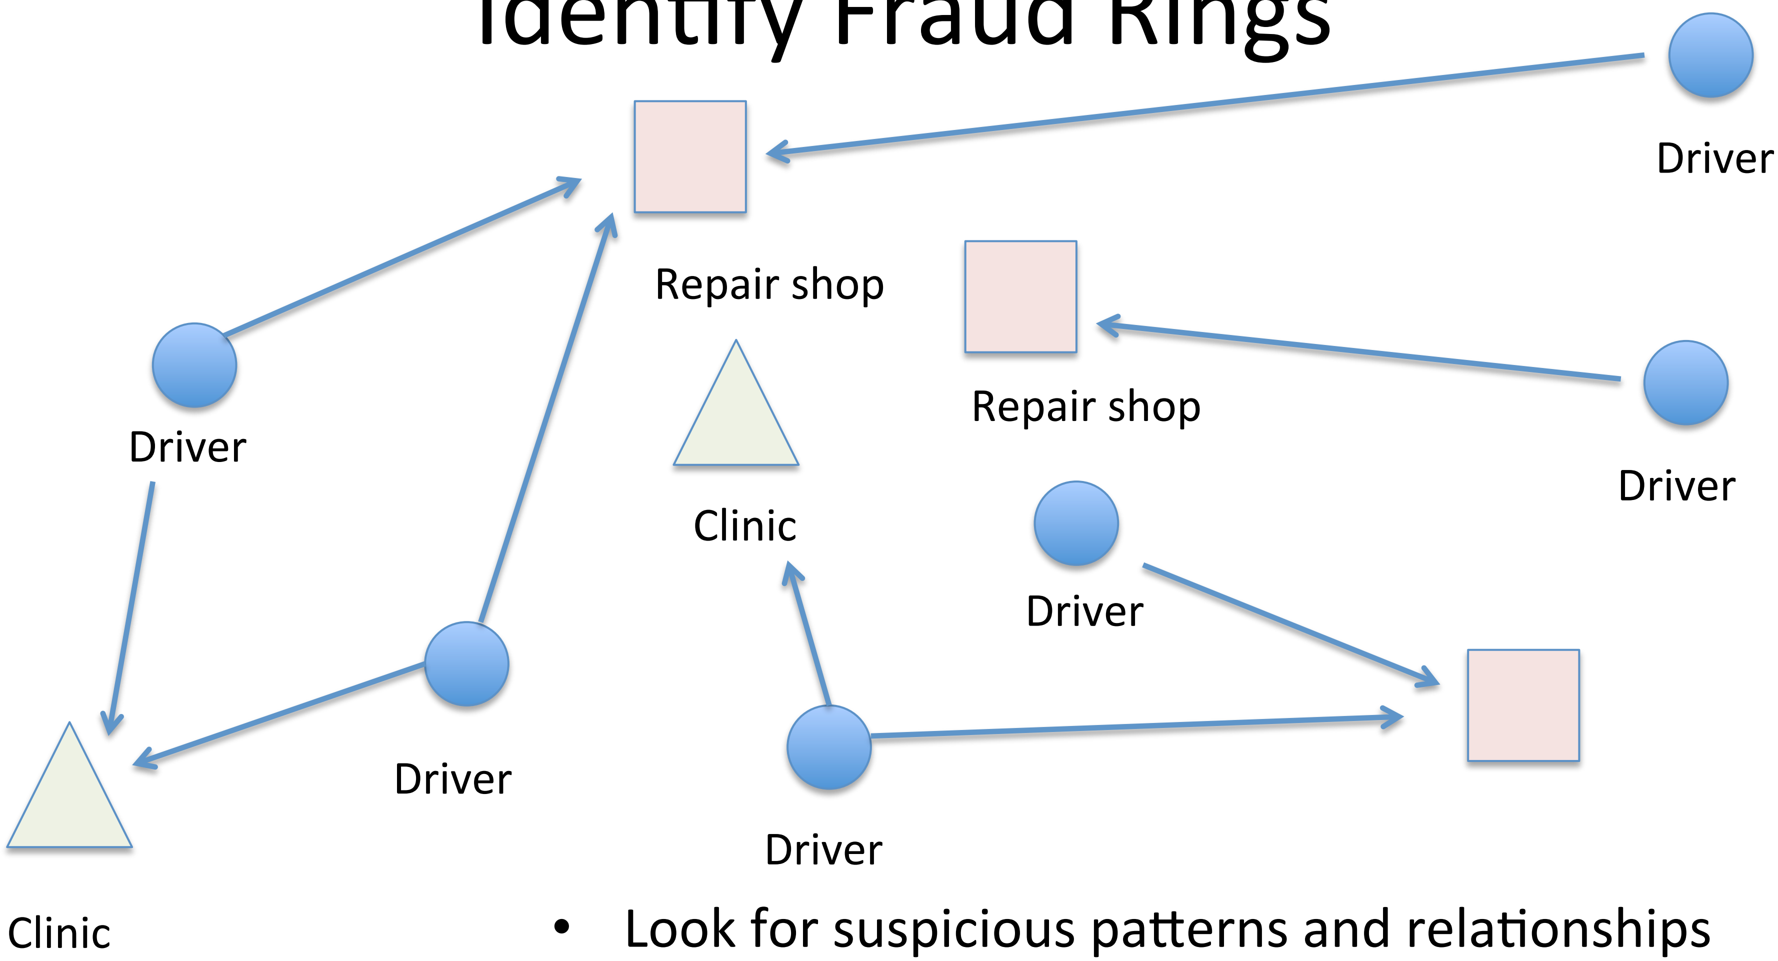

- Look for suspicious patterns and relationships among claims.
- Look for hidden relationships through common phone numbers, nearby addresses, etc.

Method 1: Look for Testing Behavior

- Adversaries often test an approach first.
- Build models to detect the testing.

Method 2: Identify Common Behavior

- Common cluster algorithms (e.g. k-means) require specifying the number of clusters
- For many applications, we keep the number of clusters k relatively small
- With microclustering, we produce many clusters, even if some of the them are quite similar.
- Microclusters can sometimes be labeled

Microcluster Based Scoring

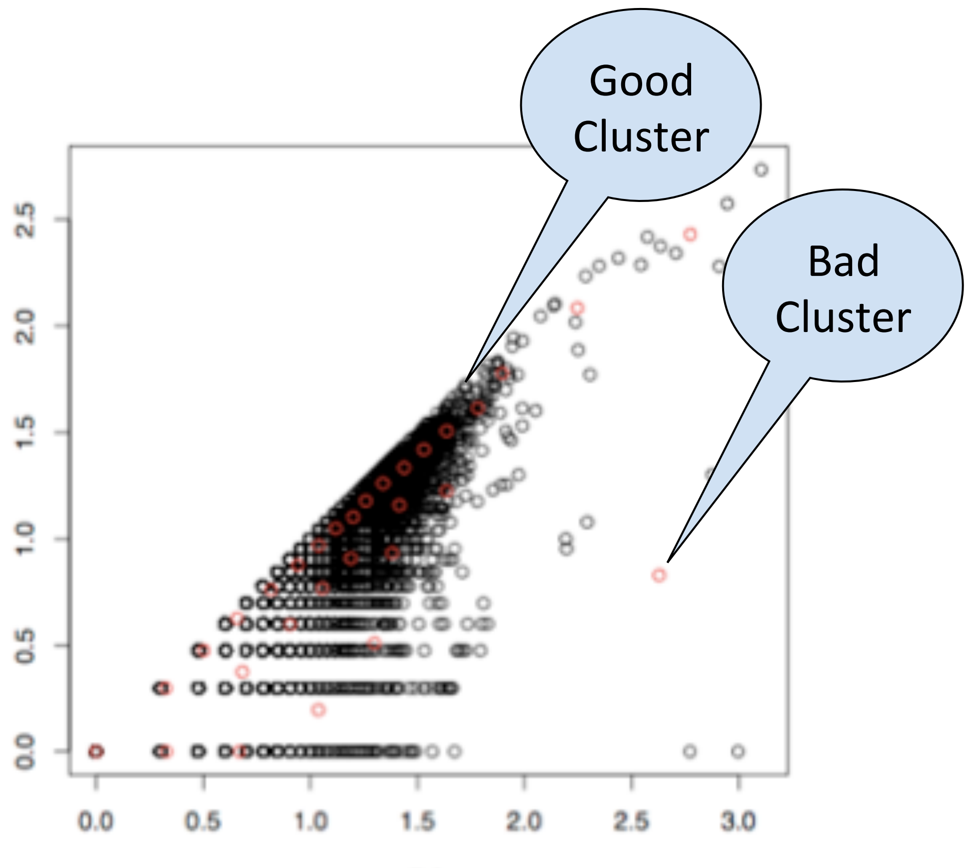

- The closer a feature vector is to a good cluster, the more likely it is to be good
- The closer it is to a bad cluster, the more likely it is to be bad.

Method 3: Scaling Baseline Algorithms

The Ghost In The Browser
Analysis of Web-based Malware

Niels Provos, Dean McNamee, Panayiotis Mavrommatis, Ke Wang and Nagendra Modadugu
Google, Inc.
{niels, deanm, panayiotis, kewang, ngm}@google.com

Drive by exploits

Source: Wall Street Journal

Points of compromise

What are the Common Elements?

- Time stamps
- Sites
  - e.g. Web sites, vendors, computers, network devices
- Entities
  - e.g. visitors, users, flows

- Log files fill disks; many, many disks
- Behavior occurs at all scales
- Want to identify phenomena at all scales
- Need to group “similar behavior”

Examples of a Site-Entity Data

Example	Sites	Entities
Drive-by exploits	Web sites	Computers
Compromised user accounts	Compromised computers	User accounts
Compromised payment systems	Merchants	Cardholders

Source: Collin Bennett, Robert L. Grossman, David Locke, Jonathan Seidman and Steve Vejcik, MalStone: Towards a Benchmark for Analytics on Large Data Clouds, The 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 2010), ACM, 2010.

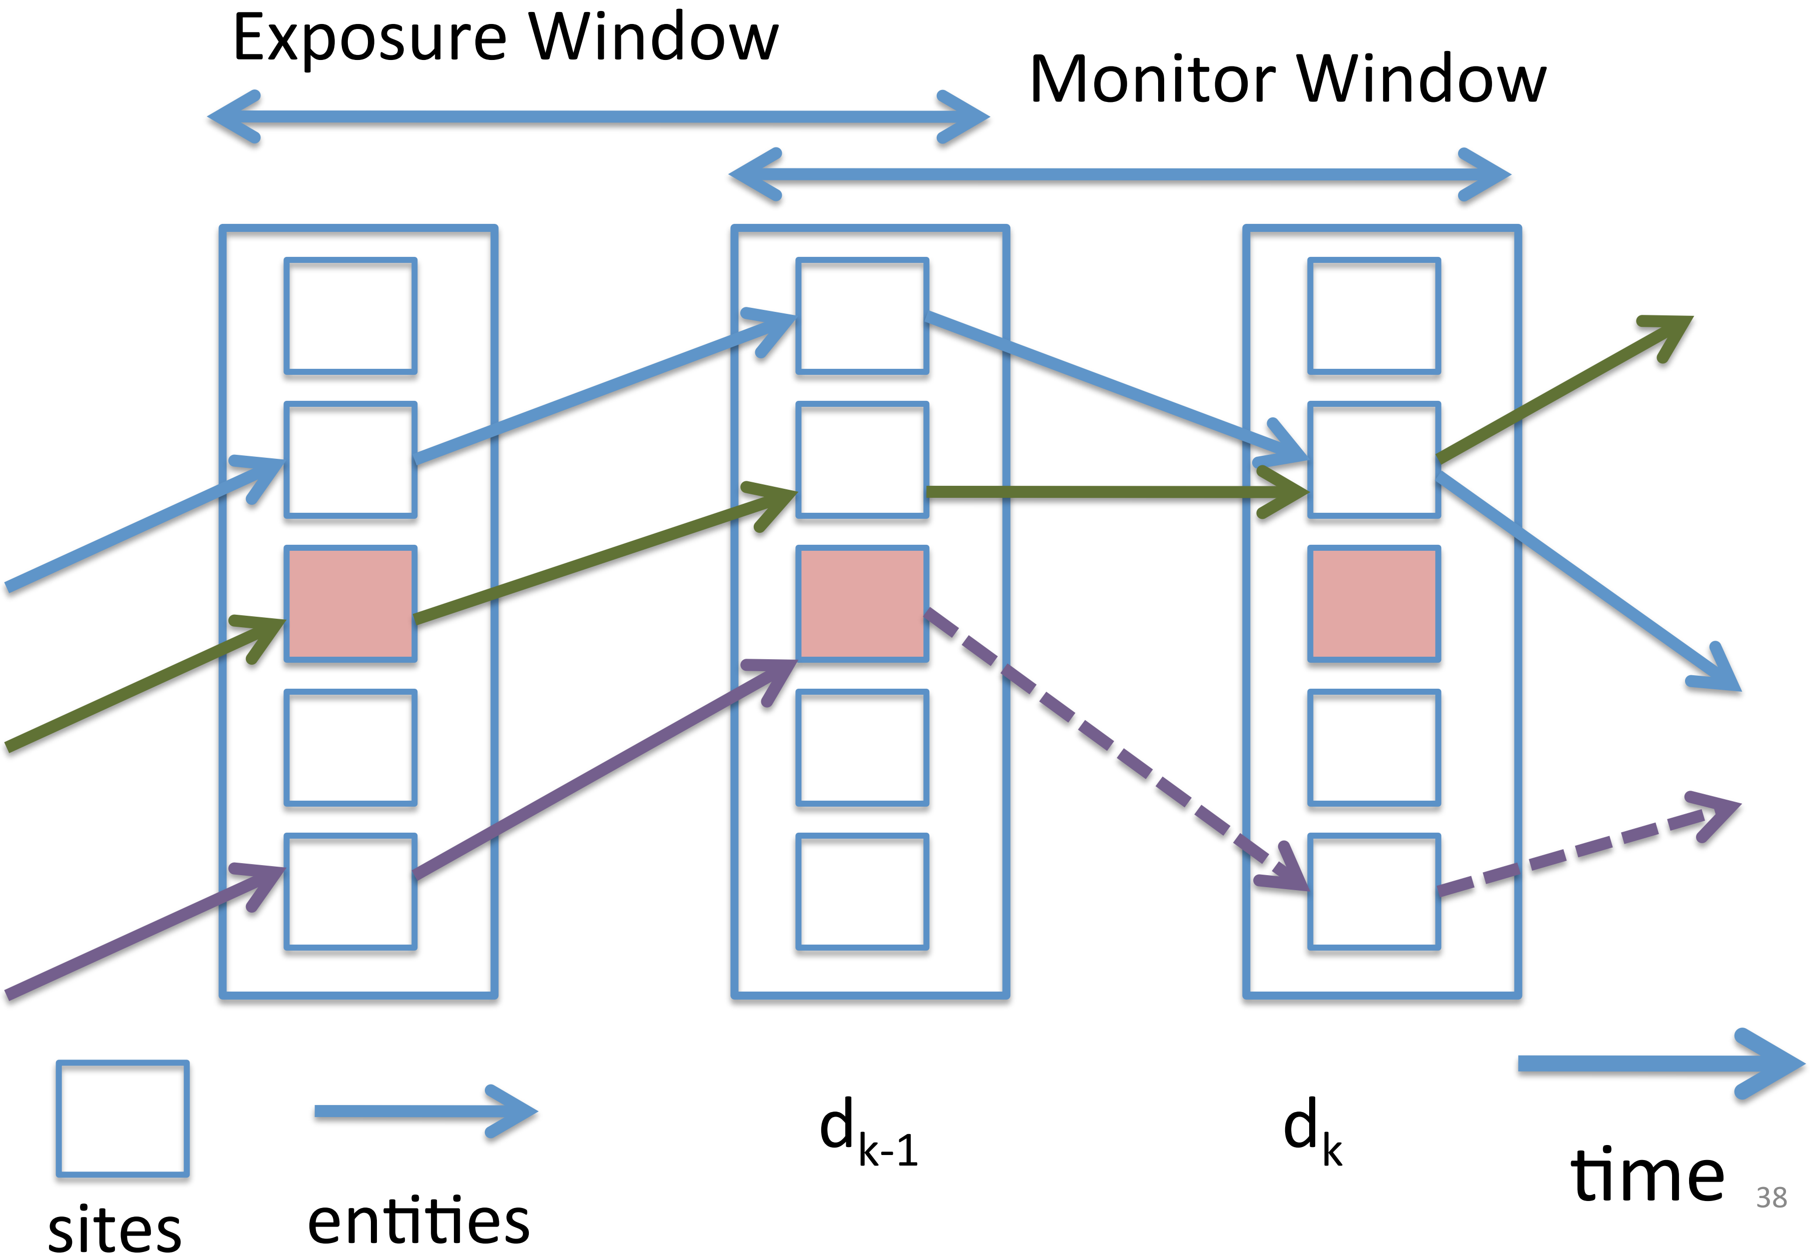

Exposure Window
Monitor Window
sites
entities
d_{k-1}
d_k
time

The Mark Model

- Some sites are marked (percent of mark is a parameter and type of sites marked is a draw from a distribution)
- *Some* entities become marked after visiting a marked site (this is a draw from a distribution)
- There is a delay between the visit and the when the entity becomes marked (this is a draw from a distribution)
- There is a background process that marks some entities independent of visit (this adds noise to problem)

Subsequent Proportion of Marks

- Fix a site s[j]
- Let A[j] be entities that transact during ExpWin and if entity is marked, then visit occurs before mark
- Let B[j] be all entities in A[j] that become marked sometime during the MonWin
- Subsequent proportion of marks is

$$r[j] = | B[j] | / | A[j] |$$

- Easily computed with MapReduce.

Computing Site Entity Statistics with MapReduce

	MalStone B
Hadoop MapReduce v0.18.3	799 min
Hadoop Python Streams v0.18.3	142 min
Sector UDFs v1.19	44 min
# Nodes	20 nodes
# Records	10 Billion
Size of Dataset	1 TB

Source: Collin Bennett, Robert L. Grossman, David Locke, Jonathan Seidman and Steve Vejcik, MalStone: Towards a Benchmark for Analytics on Large Data Clouds, The 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 2010), ACM, 2010.

5. Summary

Some Rules

1. Understand your adversary's tradecraft; he or she will go after the easiest or weakest link.
2. Your adversary will be creating new opportunities; you must create new models and analytic approaches.
3. Risk is usually viewed as the cost of doing business, but every once in a while the cost may be 100x – 1,000x greater
4. The quicker you can update your strategy and models, the greater your advantage. Use automation (e.g. PMML-based scoring engines) to deploy models more quickly.

Questions?

About Robert L. Grossman

Robert Grossman (@bobgrossman) is the Founder and a Partner of Open Data Group, which specializes in building predictive models over big data. He is also a Senior Fellow at the Computation Institute and Institute for Genomics and Systems Biology at the University of Chicago and a Professor in the Biological Sciences Division. He has led the development of new open source software tools for analyzing big data, cloud computing, and high performance networking. Prior to starting Open Data Group, he founded Magnify, Inc., which provides data mining solutions to the insurance industry. Grossman was Magnify's CEO until 2001 and its Chairman until it was sold to ChoicePoint in 2005. He blogs occasionaly about big data, data science, and data engineering at rgrossman.com.

For More Information

open data

www.opendatagroup.com
info@opendatagroup.com