

Information Security: What the Market Wants and Why

Dan Geer

geer@stake.com
+1.617.768.2723

Trends that matter

- Risk management has won
- Anticipate failure or be damned
- Peaking demand for security expertise maxes charlatan fraction

But most importantly,

- The future belongs to the quants

Does the rubber actually meet the road?

- Security is a “top management priority”
 - For CEOs, 7.5 out of 10 (Booz-Allen Hamilton 2002)
- But funding levels don’t match the rhetoric
 - On average, companies spend 0.047% of revenue on security
 - \$1.1 million average annual budget (excluding staff)
 - \$196 per person (IDC 2001)
- Why the disconnect?
 - The case for security is rarely made in business terms

Science is calling, will we answer?

When you can measure what you are speaking about, and express it in numbers, you know something about it; but when you cannot measure it, when you cannot express it in numbers, your knowledge is of a meager and unsatisfactory kind: it may be the beginning of knowledge, but you have scarcely, in your thoughts, advanced to the stage of science.

-- William Thomson, Lord Kelvin

The demand is there, what will we say?

- How secure am I?
- Am I better off than I was this time last year?
- Am I spending enough on security?
- How do I compare to my peers?
- What risk transfer options do I have?

The wisdom is timeless, but how do we operationalize?


- Know thyself
 - Can't manage what can't measure
 - Basel II brings "carrier grade" measurement to banks
 - The demand for certification, standards, liability clarity
- Nothing to excess
 - Provable security is never affordable
 - Affordable security is never provable

Gather ye numbers where ye may


- All data has bias, only question is can you correct for it
 - Poor estimators, if nonetheless stable, expose trends
 - Build models and calibrate them with what you *can* measure
- Good artists create, great artists steal
 - QA literature
 - Public health terminology and reporting structure
 - Portfolio management
 - Accelerated failure testing
 - Insurance

We can already measure wind speed

- Moore's Law, 18mo doubling
- Storage, 12mo doubling
- Bandwidth, 9mo doubling



We can already measure precipitation



We can already feel global warming

- $\text{Cost}(\text{Access_Control}) \propto \{ N(\text{people}) \times N(\text{functions}) \}$
 - Grows faster than linear hence unscalable
- Accountability only alternative
 - Begs question of anomaly detection, not intrusion detection
 - Consistent with dissolved perimeter (inside=outside)
 - Defers costs to times of forensic necessity
- Selective data deletion more expensive than complete retention
 - cf. Privacy, limiting discoverability

Good artists create, great artists steal

- QA literature
- Public health terminology and reporting structure
- Portfolio management
- Accelerated failure testing
- Insurance

Steal from the QA literature

Relative cost to fix issues, by stage

Design	1
Implementation	6.5
Testing	15
Maintenance	100

Source: *Implementing Software Inspections*,
IBM Systems Sciences Institute, IBM, 1981

Software development costs, by stage

Design	15%
Implementation	60%
Testing	25%

Source: *Architectures for Software Systems*,
course Notes, Garlan & Kazman, CS, CMU, 1998

Steal public health terminology and reporting structure

- Herd immunity vs NIMDA
 - We have no perimeters worth talking about anyway
- Asymptomatic carriers vs DDOS zombies
- Centers for Disease Control vs ISACs
 - Mandatory reporting of communicable diseases
 - Statistical identification of excess incidence
 - Longitudinal trending to calibrate epidemiologic models

<http://www.usenix.org/events/sec02/staniford.html>

<http://www.caida.org>

<http://average.matrixnetsystems.com>

<http://www.fsisac.com>

<http://www.cdc.gov/mmwr>

Steal from portfolio managers

- Redundancy and diminished operation vs hedging
 - So what knobs do you turn when DHS says “Orange”?
- Everyone in finance has this skill in-house and finance always leads security investment
- GIGA Group already touting portfolio management for IT
http://www.cio.com/analyst/012502_giga.html

Steal from accelerated failure testing

- Measurement is what drives reproducibility
 - What is the difference between a pen test and UL?
 - The most important calibrator is level-of-effort to subvert
- Assume failure, build in rollover
 - Mandatory upgrade, anti-retention in Windows Media Player
 - Arbaugh, Fithen & McHugh, “Windows of Vulnerability”
<http://computer.org/computer/co2000/rz052abs.htm>

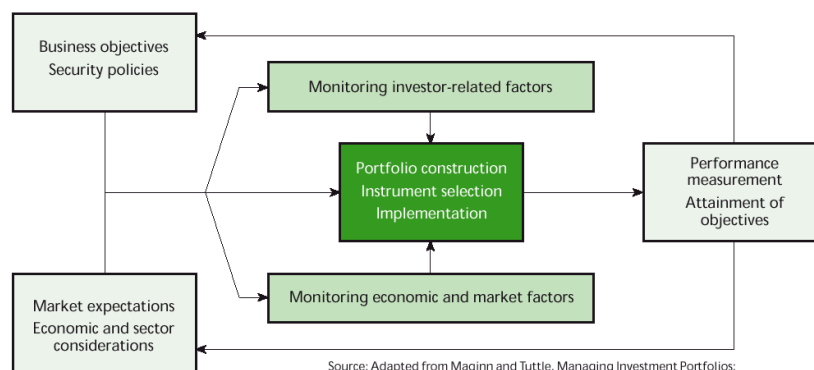
Steal from insurance

- Actuarial data impossible
 - Left censoring to maintain currency vs high change flux
 - Configuration complexity
- The real question: Risk Aggregation
 - All major events are cascade failures
 - Absence of a loss history not necessarily comforting
 - Monoculture risk vs management complexity
- Biggest leverage: liability tuning

Portfolio theory and risk management

- Portfolios balance the risk of multiple investments
- Risk is a commodity that can be
 - Classified
 - Measured
 - Priced
 - Traded
- How do financial portfolio managers classify risk?
 - **Unique risk:** Company- and sector-specific risks that can be minimized through countermeasures; also known as *residual risk*
 - **Systematic risk:** Macro factors – economic, political, cascade failures that cannot be eliminated; also known as *market risk*

Applying portfolio risk management to security



Implications of a portfolio approach

- Balanced security portfolios diversify away unique risk
 - Multiple protection layers
 - Complimentary countermeasures
- Systematic risk requires attention, too
 - Apply countermeasures to what you can (antivirus, DDOS,...)
 - Insure or bear the rest
- Security analytics measure portfolio performance
 - Drive return on security investment (ROSI) calculations
 - Feed back into risk quantification

Getting serious about risk analytics

- Why measure?
 - Current state assessment
 - Planning
 - Ongoing measurement
- What's to measure?
 - Application defects
 - Network vulnerabilities
 - Password cracks
 - Intrusions
 - Patch costs
- How to compare?
 - Over time
 - Against self, peers, industry
- How to quantify benefit?
 - NPV, revenue gains
 - Reduced rework, avoided costs

Where does data sharing come in?

- The data to share is *normal data*, not *exception data*
 - Can't do own anomaly/outlier detection w/o it
 - Can't differentiate "target of chance" from "choice" w/o it
- Recognized practices vastly better than mandated ones
 - Can't recognize without data to do bias correction
- Retaliation requires evidentiary-grade forensic readiness
 - Pool data automatically else interfere with operational recovery
- Immunity and virulence are large number concepts
 - What fraction of infected machines is tolerable?

So, put up or shut up...

Applications are where the action is now

- Security trends say so
- Business realities say so
- “Crime” statistics say so
- Risk management means quantitative decision support

Contributing factor 1 - *Applications are federating*

- Distributed applications have multiple security domains
 - **The firm**: client service & administrative functions
 - **External providers**: front-end Web farms and application hosting
 - **Partner interfaces**: data streams (inventory, payment, real-time feeds)
- Applications get ever more moving parts
 - Mainframe □ client-server □ *n*-tier □ Model 2 (J2EE and .Net)
- Network service stratification
 - Bandwidth, hosting, provisioning, delivery

Contributing factor 2 - *Perimeter defense is diseconomic*

- “Shared wire” supplants “shared model”
 - XML is the great equalizer
 - SOAP and XML-RPC specifically designed to go through firewalls
 - Emerging web services
- Firewalls stop nuisance attacks, not application traffic
 - Everyone leaves ports 80 and 443 open
- As a result, the threat model mutates
 - More attacks through HTTP, at application level
 - More attacks targeted at specific application components
 - Attacks on applications require lower skill levels

Contributing factor 3 - *Data, data everywhere*

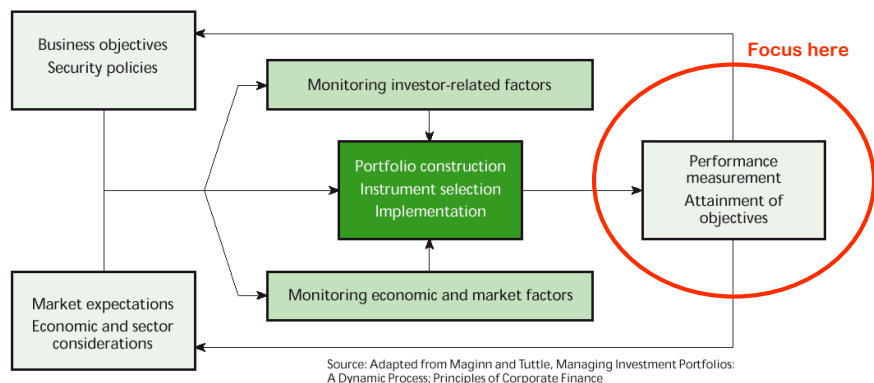
- Data storage needs increasing exponentially
 - More new data produced in next 3 years than in all of human history
 - Corporate IT spending on storage:
4% in 1999 v. 17% in 2003 (Forrester)
- Form factors proliferating
 - Local storage
 - Storage arrays
 - Appliances/network-attached storage
 - COTS: <\$1/GB, >100TB/rack

A little example of pooled data

Security evaluation of major applications treated as a source of summary numbers and shared intelligence

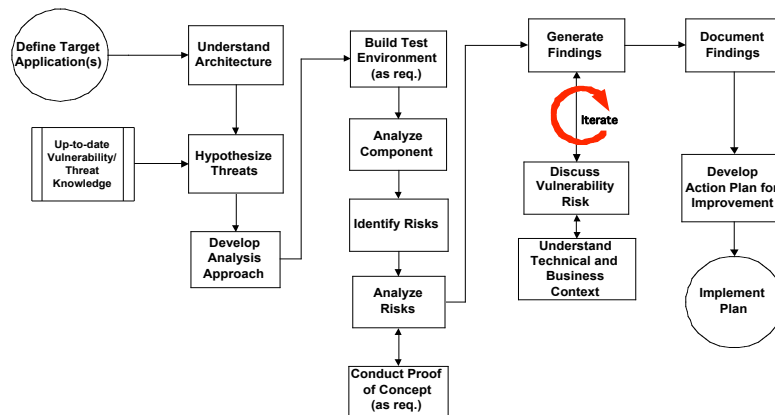
All data are real, pooled and hence anonymized within a trust relationship, and modeled as normative

Applying portfolio risk management to security



Data acquisition

Application Penetration Testing Approach



Findings (1/4): Security Defects Are Common


Security Defects by Category

Category	Engagements where observed	Design related	Serious design flaws*
Administrative interfaces	31%	57%	36%
Authentication/access control	62%	89%	64%
Configuration management	42%	41%	16%
Cryptographic algorithms	33%	93%	61%
Information gathering	47%	51%	20%
Input validation	71%	50%	32%
Parameter manipulation	33%	81%	73%
Sensitive data handling	33%	70%	41%
Session management	40%	94%	79%
Total	45	70%	47%

*Scores of 3 or higher for exploit risk *and* business impact


Source: 2002 @stake - The Hoover Project (n=45)

Top 10 Application Security Defects



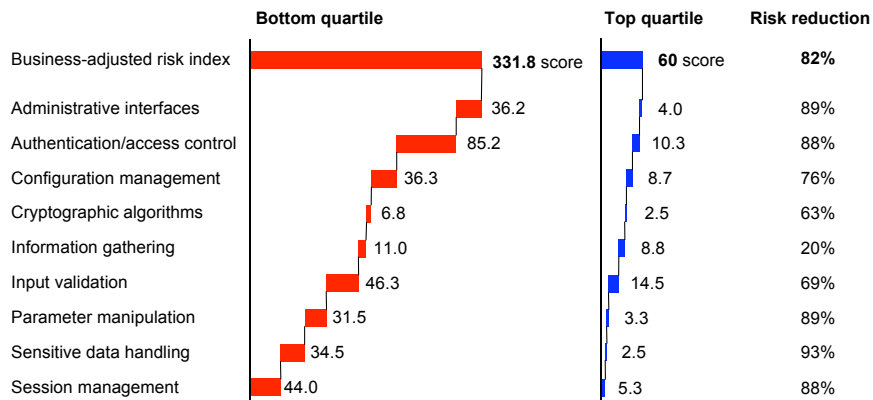
Assessments where encountered, percent

Findings (2/4): Leaders Have Fewer Defects



Source: 2002 @stake - The Hoover Project (n=23)

Findings (3/4): Leaders Carry Less Risk



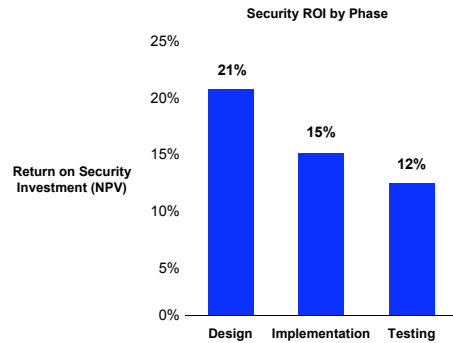
Average business-adjusted risk (BAR) index per engagement, with breakdown by risk category

Source: 2002 @stake - The Hoover Project (n=23).

BAR index = sum of all defects' individual BAR scores, where each defect's score = exploit risk (5 point scale) x business impact (5 point scale).

Findings (4/4): *Fixing security defects earlier pays off*

- Although benefits can be found throughout the lifecycle, earlier involvement is most beneficial
- Vulnerabilities are harder to address post-design
- System-wide changes may be required at later stages
- Enabling improvements can be made at design state



Source: 2002 @stake - The Hoover Project

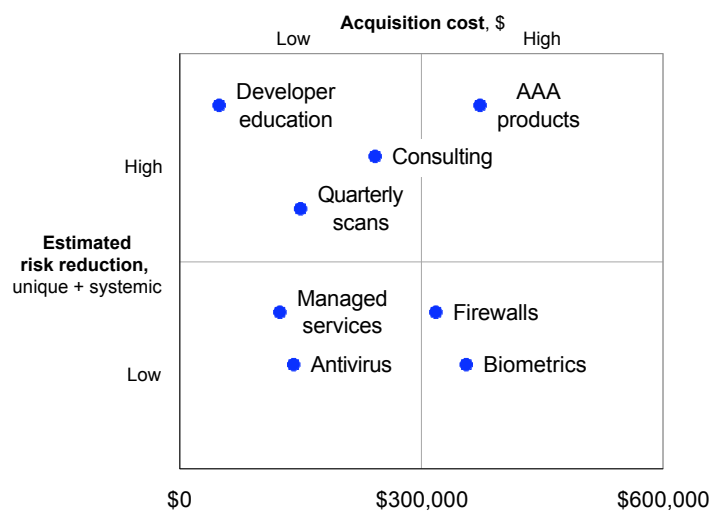
So what's next?

Use what we've learned, learn what else we need

Investment example: *patch management*
Technique: *assessment + rework reduction*

- Leading financial services provider uses application scanner to determine risk profile of COTS products
 - Risk index r_a calculated
- When new patches are released
 - Patched version scanned again, r_b determined
 - Deployment/rollout cost c_r calculated
- Relative benefit calculated
 - Dollars per unit of risk reduction = $(r_a - r_b) / c_r$

Portfolio planning example: *cost effectiveness judgment*
Technique: *2x2 matrices*



What else is needed?

- Systematic v. unique risk classifications
 - Are security incidents like earthquakes or housefires
 - Create classification scheme
 - Decompose security risk to facilitate management
- Risk ratings
 - Similar to credit ratings or stock “betas”
 - Partner with insurance carriers (premium bases)
- Cost effectiveness
 - Amount and type of risk reduction per dollar spend

Repeating: *Trends that matter*

- Risk management means quantitative decision support
- Anticipate failure or be damned
- Charlatan fraction is rising (as is overselling)

But most importantly,

- The future belongs to the quants

Summary

- If not now, when
- If not us, who

Dan Geer
geer@stake.com
+1.617.768.2723
in collaboration with
Andrew Jaquith
ajaquith@stake.com
+1.617.768.2711