

ANALYSIS OF CYBER METRICS WORKSHOP

London, 23 May 2018

INTRODUCTION

The Research Institute in Science of Cyber Security (RISCS) is the UK's first academic Research Institute to focus on understanding the overall security of organisations, including their constituent technology, people and processes. RISCS is focused on giving organisations more evidence, to allow them to make better decisions, leading to the development of cybersecurity as a science. It collects evidence about what degree of risk mitigation can be achieved through a particular method – not just the costs of its introduction, but ongoing costs such as the impact on productivity – so that the total cost of ownership can be balanced against the risk mitigation that has been achieved. RISCS main goal is to move security from common, established practice to an evidence base in much the same way as happened in sectors like medicine.

On May 23, 2018, RISCS held a workshop in London that looked at the utility of cyber security metrics. The purpose of the workshop was to develop a deeper understanding of the ways in which cyber security metrics are used in decision-making more generally, and also to raise questions about how data is best presented to the board and the policy community more specifically. We wanted to explore the potential for metrics to help but we also want to take a critical approach to the underlying values that can shape metrics – and consequently, decisions.

For further information about this report or about other RISCS initiatives, please contact Emma Bowman at riscs.administrator@ucl.ac.uk.

METHODOLOGY:

To investigate the utility of cyber security metrics in the decision making process of industry and the policy community, we gathered a group of 70 people from academia, industry, the policy community and the technical community. We asked these people to self-identify themselves as ‘providers’ or ‘consumers’ of metrics and to individually or collaboratively record their responses to four questions that we asked of their group. For the cyber security metrics provider group, we asked them to populate the following table:

		PROVIDERS OF CYBER METRICS	
		INCLINED	DISINCLINED
FEASIBLE	FEASIBLE	<i>I can/do give you this</i>	<i>I could but I'd rather not</i>
	NOT FEASIBLE	<i>I'd like to but...</i>	<i>We need to talk</i>

We asked those people who identified as the consumers of cyber security metrics (decision makers about investment, policy etc) to respond to the questions on a separate table:

		CONSUMERS OF CYBER METRICS	
		USEFUL	NOT USEFUL
RECEIVING	RECEIVING	<i>Keep it coming!</i>	<i>Stop!</i>
	NOT RECEIVING	<i>Wish list</i>	<i>Thanks for the offer, but....</i>

A complete transcript of the recorded responses is included at Appendix A. On the following pages, we present our analysis of the findings.

SUMMARY OF FINDINGS

Providing tailored cyber metrics is an opportunity to engage with leaders and shape their perceptions of information risk. The outputs of this workshop suggest that success depends on providers delivering material that:

- genuinely reduces uncertainty;
- addresses specific questions; and,
- uses the language of business.

The table below brings out the trends from our workshop in terms of what is and is not required by decision-makers. Below are some of the most interesting conclusions that we have drawn from the data.

- ➔ We saw that some requirements may be met with reticence from metrics providers, perhaps because they are incendiary or embarrassing (e.g. rankings, disclosures, overly ambitious information sharing regimes). Providers tend not to want to deliver unwelcome news, for example proof that past investments have delivered little benefit. The most significant tension appears to be between the need to inform financial decisions and the reluctance by some providers to deliver metrics which over-promise on that front. The extent to which this something to do with trust and liability, or a lack of mutual understanding between consumer and provider would require further study.
- ➔ The responses also revealed a sense of mistrust in metrics delivered by some commercial providers. While these metrics would be increasingly useful as more services are outsourced, there was a feeling among some participants that commercial service providers hadn't necessarily the vested interest to provide accurate or timely metrics, especially if Service Level Agreements were being breached.
- ➔ Some consumer requirements – while valid – are difficult to achieve, perhaps because of financial constraints, frailty of commercial products or lack of quality data. For example, we saw a few comments which suggested that common vulnerability scanning tools lacked accuracy. Other examples of metrics which could be difficult to produce included: protective monitoring and alerting capabilities, assessing the true cost of an incident, gaining confidence in cyber insurance policies, understanding the overall costs from cybercrime, metricizing the blockers to successful GDPR compliance, and delivering “micro-narratives” to decision-makers.
- ➔ There were some notable instances where providers of metrics could offer new ideas. For example, how much systems downtime had been encountered due to rogue events or activities which were beyond the control of network managers. Some providers saw value in engaging in a dialogue with leaders to refine priorities and build consensus, for example where excessive cost had been incurred due to not investing earlier in security. There was also the suggestion to obtain accurate impact and vulnerability scoring through table-top exercising. Some providers were inclined to deliver quite revelatory metrics – conditional on feeling safe to do so – including the real “value” of extended support contracts and evidence of sensitive compromises and vulnerabilities.


WHAT CONSUMERS WANT	WHAT CONSUMERS DON'T WANT
<p>THE COMPETITION:</p> <ul style="list-style-type: none"> • How am I performing against my competitors? • What happened to my opposite numbers in other companies after a breach? • What was the impact/cost of their mistakes? • Benchmarking: how am I doing compared to my peer group? What is the probability of a breach compared to my competitors? 	<ul style="list-style-type: none"> • Made-up numbers. • Feeling bombarded by meaningless or nitty-gritty, technical data. • Snapshots with no temporal context. • RAG (Red, Amber, Green). • Unqualified opinion. • Bias or excessive subjectivity, especially wrapped up in something succinct or scientific-looking. • Fear-mongering. • Blame. • Sales pitch for snake-oil or magic bullets. • Unstable or unrepeatable stats. • Style over substance. • Jargon. • Old data. • Inappropriate or ineffective visualisation techniques. • Spin.
<p>IT'S ALL ABOUT THE MONEY</p> <ul style="list-style-type: none"> • Relevance: what is my return on investment? • How do my cyber risks affect my ability to raise capital? • How much risk am I unable to transfer (e.g. through cyber insurance)? • Am I carrying criminal liability that I can't transfer? • Getting to root cause: solve many problems with one fix. • Cost of recovery vs cost of control (proactive vs reactive posture on cyber security, and comparative costs). • What's the risk/cost associated with doing nothing? 	
<p>CYBER SECURITY AT A GLANCE</p> <ul style="list-style-type: none"> • Some comments pointed towards a basic Security Operations capability • Dashboarding, e.g. network boundary activities 	
<p>THREATS & VULNERABILITIES</p> <ul style="list-style-type: none"> • Past breaches are a good indicator of future vulnerabilities. And metrics of real incidents are more informative than potential ones. • Is threat intelligence used? How useful is it? Where else is it used? Where is threat intel real and valuable? 	

<ul style="list-style-type: none"> • What are my priorities, rather than just looming threats? Don't scare me, inform me. • Are there quick wins available? How can I make some progress fast? Help me lift opportunities out of the noise – what can I do right now? • Tracking capability of low-capability threat actors. • Impact, rather than quantity, of incidents is important. 	
<p>THE USER</p> <ul style="list-style-type: none"> • Are users absorbing training? E.g. are they forwarding suspicious emails on to the security team? • Security vs productivity: spotting where security policy is fatiguing people and impeding productivity. • Culture and indications of user wellbeing and behaviour (and the risk that presents). 	
<p>THIRD PARTIES</p> <ul style="list-style-type: none"> • Where is sensitive data going? Whom is it being shared with? • How do you trust the stats provided to you by third party service providers? 	

APPENDIX A: Transcription of Cyber Metrics workshop responses

Part One: Responses from consumers of cyber security metrics


“Keep it coming!” (Already receive this and find it useful)



- [Common Vulnerabilities & Exposures \(CVE\)](#) vulnerability database
- Results of cyber defence maturity assessments
- Details of past incidents
- Timing series analysis of past breaches
- Risk assessment of IT change projects
- Patching status
- [Cyber Essentials](#)
- Consistent time series metrics
- Organisations' strategy: relevant to how this business makes its money
- Incidents, but by dept/function to find risk areas: simple total number is not that useful
- Risk assessments of IT change projects: base, with controls, costs and options
- Return on investment
- Barriers to uptake of behaviour
- Real incident breaches, rather than potential ones
- Punishments received by other boards' directors
- Overview of attack attempts, especially over time
- Live dashboard, e.g. AV levels/status
- Metrics in terms we understand, e.g. business language
- Understanding perceptions of risk
- Do organisations understand what is critical for them, i.e. what to protect?
- Malware notifications
- Qualitative analysis of what's useful
- What is my uninsured risk?
- Data that enables me to do something, what changes can the organisation do to respond?
- What incidents have affected organisations and what is the trend?
- Short, sharp and to the point
- How does my level of uninsured risk affect my ability to raise capital?
- Metrics aligned to solutions
- Level of real valuable knowledge sharing within a sector
- Do organisations receive any used threat intelligence (understand)?

- Data that has a supporting narrative
- Surge capacity, used space, e.g. when it's > 80%, it should give you an alert
- Usage patterns, responses in short, medium and long term
- AV / malware coverage
- No. of installs up to date
- Detections alerts: types and propagation
- Actions: cleaned and quarantined
- Network capacity: e.g. used bandwidth/time > x%, it triggers an alert
- Position metrics as a tool to enhance the business, not to present hurdles
- Phishing tests
- Governance of cyber security, see [Cyber Assessment Framework \(CAF\)](#)
- How does an expenditure affect my uninsured risk?
- [Boston consultancy matrix](#): I want to know more about my "star" and "cash cow" areas than my "dog" or my "question marks"
- Areas where we are not compliant
- How do we have the confidence to discuss cyber once a year, rather than once a month?
- How is my cyber security posture affecting my defence with respect to criminal liability that I cannot convert?
- Small & Medium-sized Enterprises (SME) with [Cyber Essentials](#) (Basic & PLUS), [IASME](#) or other level of security management

Wish list (Don't currently receive this but would like to)



- Configuration status of the corporate IT: no of devices, software running, known vulnerabilities mapped from CVE database, plus severity
- Sensor data: types of attack, access vectors, effects, impact
- Effectiveness of training and awareness packages, e.g. how many people clicked the link during a spear-phishing campaign?
- Suitable quantitative metrics for human factors (not sure if this is possible)
- Physical security: not seen much in this area
- Risk assessment and associated investment plan
- Secure behaviours vs target/norm, e.g. phish is reported
- Supporting analysis/data that give more details when needed
- Awareness, engagement: minus done + how well done, e.g. linger time
- Metrics aren't just about numbers
- Days lost per year due to security 'features'
- Quantitative and qualitative data (holistic view)
- User awareness level

- Robustness of change management process in organisation
- Time to fix by the ISP/System Integrator
- IT projects with security designed as a proportion of mere existence of secure design practice
- Likelihood of the breaches (supported by robust model, e.g. [couldn't read it])
- Insecure behaviours: clicking dodgy links, use of USBs, use of dropbox, webmail
- Doing training too quickly, ignoring awareness material
- Real incident data
- Impact of breaches
- Metrics that show impact of incidents not just number of incidents
- Behaviours that could be insecure, who clicks lots of links, who browses a lot?
- Anything that demonstrates impact on financial accounting metrics
- Third parties
- Good benchmarks
- Threat actor activity
- Relevant business activity
- Putting information in the corporate context, think about our annual report
- Sector-by-sector metrics with uniform methodology
- Qualitative, outcomes-based data
- Network maps of how different types of attack are maybe caused by the same root causes
- A quick summary for high level executive pack
- Identified cyber security requirements to be implemented by users
- Quantitative consequences from case studies of similar companies
- When we are in mergers and acquisitions: overview of system integration risk as part of bid cost vs bringing data onto our existing system
- Number of silent connections to my phone/device, their threat level and actionable steps to reduce risks
- Resilience of my devices to different types of attack
- McGraw/BSIMM data
- Number of passwords re-used/repeated across websites and services
- Metrics must have context, otherwise they're just stats
- Mandatory criteria to benchmark cyber security status of an organisation
- Security actions I have done right
- Notification of personal impact, not organisational or technical impact
- Quick wins and longer-term solutions
- Real time dashboard
- Return on investment / cost-benefit
- Please put system risks in business context, productivity, cost etc
- Productivity, cost-benefit of controls, e.g. time training vs value
- Develop secure coding capability
- How our competitors are doing: not seeing much
- What a good process looks like rather than an outcome
- If we have outsourced our IT, what information should we contract our provider to report? How can we trust them?
- Understanding of effectiveness of a security control
- Quantifiable risk of doing nothing differently


- Robust stats about behaviour rather than thinking/intention/awareness
- Risk managers report
- Option/negotiation: information/threat alert suggest/requires actions/response, but what are the alternatives? The in-between options and consequences?
- Meaningful connection: what do these numbers/percentages mean in terms of required or recommended actions?
- Indication of possible employee abuse that might actually indicate an HR issue, e.g. stress or poor management
- Analysis couched in terms of business continuity
- Security culture (see [CPNI](#) tools)
- True cost of recovery vs cost of control
- Want more! 80% say.... [reference to Madeline's presentation]
- Cost of measure implemented this FY
- Formula to turn threat intelligence into risk profile
- Capability level of attacker you can defend against ([STIX/TAXII](#))
- How much productivity lost to stupid security policies?
- Subjective assessment of risk of attack, threat vs measures
- Benchmarking against a similar group of organisations
- Capability to achieve recovery time and point [sic] objectives
- Less stats, more info
- Changes of network traffic following new security policy implementation
- Incidents traced to root cause
- How many other people forward phishing emails for analysis?
- Probability of a breach in my industry for my application
- How services are interrogating my data, e.g. how is my email being read?
- Onward transmission of my personal data
- Large companies are taking an active leadership role in supporting their industry sector
- How do you capture/represent the intent of the action in a metric form?
- Performance of trained users three months after training?!
- Do you have an incident management plan and have you exercised it?

“Stop!” (Currently receiving this but don't find it useful)

- Making up numbers
- Numerical stats overload
- Single reports, i.e. not in relation to past or future metrics
- Number of hits on my firewall
- Number of employees clicking on fake phishing emails

- Privacy statements and password strength
- RAG (Red Amber Green) ratings: doesn't mean the same thing to everyone
- "Expert opinion"
- How many orgs have a security policy?
- Number of incidents or detection events with no base rate of occurrence
- Cherry picking data, i.e. biased analysis
- Bullying with threats of "bad things"
- Machines patched
- Network monitoring stats
- Blaming users by telling us what they've done wrong
- Users trained
- Magic bullet solutions
- Traffic lights
- Anything that you cannot prove to me will be stable enough to invest in measuring over time
- 3D pie charts or bubble charts
- Uncontextualized numbers
- Non-contextual stats
- Metrics full of jargon without explanations
- Drowning me in data
- "indexes" that hide complex, subjective methodologies
- [SIEM](#)
- AV stats
- Patched percentages: I'm only interested in effects
- Nitty-gritty detail of system patching
- Tick-box process confirmation
- Outdated data
- Anything qualitative
- Historic events, e.g. paste bin
- Details of individual lower-significance incidents/issues
- Progress against compliance requirement
- Irrelevant data where the results are not significant
- Phishing test stats: hugely variable if compared one to the next, but useful to compare between depts to find risk areas
- Poor presentation of the graphics, e.g. poor choices of colour, inappropriate chart types
- Incident numbers without context, e.g. 6,000 incidents across sector X in 2015

“Thanks for the offer, but....” (Don’t receive it and wouldn’t want to if it were offered)



- Metrics for a fee
- Risk metrics without solutions
- Metrics that prompt more questions than answers: don’t give me problems without solutions, I’m busy enough already!
- Data that has been sanitised by middle management
- Unsolicited sales pitch: information gathering, social engineering

Part Two: Responses from providers of cyber security metrics

“I can/do provide this” (Currently delivering and plan to continue)

- Total cyber security spend as a percentage of revenue/profit
- Existence and rehearsal of incident response plan
- [BGP](#) routing tables
- DNS trace data above recursive resolver
- Likelihood of future breach
- Number of phishing attacks in my organisation
- Time to resume service delivery post-incident
- Impact of incident on service delivery / BaU processes
- Time to incident/compromise detection
- Max number of rogue change days on my network
- Evidence of network compromise
- #comment: categorise prevention metrics, detection metrics, response metrics, and recovery metrics
- Total number of known vulnerabilities on the network
- Number of detected network intrusions
- Number of breaches as a result of untargeted and unsophisticated attacks
- Risk register entries, i.e. likelihood of a breach of customer personal data
- Motivation of threat
- Software inventory via Software ID (SWID) tags
- Case studies/scenarios
- All metrics are proxies and subject to calibration errors
- Number of employees attending/completing training, infosec, phishing etc (doesn't show how effective it's been)
- Evidence of compromise
- Number of staff without adequate security training
- Metrics/data for the sake of it = comforting
- “value” of extended support
- Worst case scenario, e.g. days website would be down, systems that would need to be rebuilt
- Uptake of / barriers to password behaviour (maybe)
- Cost of past breaches
- Total cost to insure

- Quantified information on personal data, i.e. what is our exposure, email and card details for one million customers
- Money spent responding to preventable incidents, i.e. if there'd been more investment in the first place
- Support status of my estate
- Performance of secure behaviours: reporting incidents, engaging with awareness/engagement duration
- Cyber breaches survey
- Number of password reset requests
- Board engagement with cyber security (FTSE 350 survey)
- Different metrics on the same system for different perspectives
- End user compliance with phishing detection/avoidance rules
- Number of cyber incidents prevented or averted
- Insider threat
- Adam Joinson's obesity map
- Accounting logs (AAA)
- Malware detected and quarantined
- Cyber defence maturity assessments (policy, e.g. NIST Cyber security framework, CD Cat, IA Maturity Model)
- Syslogs: firewall logs, visualisation through graphs
- Patch status
- Desktop build
- Phishes blocked
- [Phishing email life journey in a Sankey diagram](#)
- Capture by security product vendor
- Read teaming: table-top exercises with stakeholders and system owners to agree impact and vulnerability scores for identified attack vectors
- Are these convenient but not useful?
- Level of engagement of employees in good practices on security behaviour, both passive and active
- Statistics on security incident from software, hardware and human analysis (server and client sides)
- Statistics on human security behaviour estimated from software, hardware, surveys, observations, analysis of data and reported incidents etc

“I’d like to but...” (Would be happy to deliver but there are feasibility issues)

- Analyse logs before a problem occurs
- Complete vulnerability management
- Behaviours/security culture
- Assurance levels
- Number of unknown vulnerabilities in the network
- Cost of an incident
- Am I compliant with the terms of my cyber insurance?
- A single metric that can be compared across all organisations
- Cost of future breaches
- Accurate cost of cybercrime to a company
- [GDPR](#) constraints
- Micro-narratives
- Mismatched/ill-fitting data protection rules
- Instruments not available (too expensive)
- Feedback loop on metrics – still relevant?
- How secure are we (with a single percentage value)?
- Qualitative data: resource automate
- Information that inherently represents the task (conceptual rather than the data (computer science))
- Knowing when information is still relevant and not out of date to be of any use – create informed decision

“I could do but I’d rather not” (It’s feasible to deliver but makes me uncomfortable)

- Accuracy of vulnerability assessment tools
- Estimated costs of security incidents and crime in cyberspace, in terms of monetary value
- RAG (Red Amber Green) ratings
- Board engagement by sector

- Rankings of companies by cyber security capacity
- Fix times by supplier/contractor: incendiary or misleading because of Service Level Agreements
- [Cyber Essentials](#) uptake
- Confidentiality constraints to disclosing data
- Numbers of scans on each port on a webserver
- Time to fix website vulnerabilities in a public league table of organisations
- Degree/report of vulnerabilities fixed / recommendations addressed post-penetration test
- Results of testing employee security/awareness
- Days exposed to disclosed vulnerabilities
- Cost of providing/managing technical controls
- Number of times (threat) intel has been shared via [CISP](#) etc
- Employees' digital footprint/corporate information exposed via internet/social media
- Number or % of employees passing formal education/training /certification etc
- Benchmarks against peers in same industry/sector
- Improvement-related statistics on security incidents and behaviours in terms of numbers
- These can be used to game other indicators
- Number of higher privilege accesses
- Too hard to disentangle from value proposition
- Netflow/IPFix at organisation/internet boundary
- DNS Trace data below recursive resolver
- Made up data or misleading data
- Penetration test results
- Provenance
- Ensure coverage / sample size
- Is a metric which of these? Evidence, data, measurements, mathematical sense of distance between two items (I think we mean evidence)?
- Have raw data but difficult to aggregate or visualise
- Number of our back doors

“We need to talk” (This cannot be delivered. Even if it could be, I wouldn’t want to)

- An estimate of the number of cyber breaches prevented
- Anything that shows me in a bad light
- Agreed metrics from UK Government ([NCSC](#)) for scoring impact and vulnerability
- Anything that claims to demonstrate impact on financial accounting metrics (Profit and Loss, balance sheets etc)
- A wider sharing / collaborative network of experts prepared to share information and work together
- Contextualised quantitative data
- Number of competitors’ back doors