MIND THE GAP

Charaka Goonatilake *considers the importance of bridging the divide between GRC and cyber*

Complexity so often breeds uncertainty, and so it is with highly complex IT infrastructures and how they are used and secured. GRC (governance, risk and compliance) teams need answers to manage risk and assure compliance. Answers are invariably found, but seldom quickly, efficiently and accurately. Cyber and GRC teams share the same goals – needing to know what is going on beneath the many overlapping layers and pockets of security – but there is friction between them instead of harmony. Both need to be certain of security controls coverage and to start basing judgements and decisions on truth rather than belief.

Every element of cybersecurity and the wider IT estate produces data. Data that, when cross-referenced and analysed, can reveal whether that organisation is secure, compliant and manages its risks appropriately.

This strikes at the heart of what an organisation's GRC function is for. And yet, with the growing spectrum of internal and external policy and regulatory concerns, managing security and IT data into conclusive evidence is becoming more difficult. Cyber and GRC are separate functions with shared goals. What they have in common is a manual, and often broken, system for understanding if adequate security controls coverage is in place, making risk, compliance and cyber objectives far harder to achieve.

As regulators and policymakers develop more stringent data privacy and protection requirements, internal frameworks for appropriate data governance also increase in scope. This gives GRC teams more cyber-related challenges to focus on, though their toolsets are not geared up to address them.

GRC teams look at more than just cyber and use well-established tools to help execute their role. The GRC approach takes a high-level view of all applicable internal frameworks and external regulatory requirements, how these are distilled into corporate policies and how the organisation's compliance against these policies can be periodically verified.

It is imperative to understand that traditional GRC tools do not adopt a rigorous, quantitative, data-driven approach to establishing truth. Rather, they seek to establish compliance assurance and manage risk, typically by asking questions and sampling data.

The common artefacts employed by GRC teams are qualitative questionnaires. These range from the binary yes/no to the highly detailed. There are clearly a number of limitations with this approach.

Qualitative assessment leads to evidence that is substantially more subjective than objective. Moreover, these questionnaires typically operate on the basis of representative sampling rather than a complete picture, which can skew results. Yet the strategic purpose of the exercise is to validate compliance, and on this measure, it is found wanting. One would not ask the CFO to state the financial health of the company on 'gut feel' or an opinion unsupported by complete, incontrovertible facts.

The product of these manual processes amounts to a 'point-in-time' estimation of compliance posture. This may satisfy interested parties, but will necessitate a ground-up repeat of the same process every time the same verification is sought. Other accuracy concerns surrounding manual processes are the potential for human error, bias or even abuse. This is in the absence of a continuous source of facts from which trustworthy data can be extracted.

The information needed to properly validate searching compliance questions exists, but GRC tools are not oriented towards collecting, storing, analysing and presenting this data. This challenge is often passed to security teams, adding to their workload. As we will examine, this doesn't make it any quicker or easier to determine, especially without an appropriate data-driven tool sitting across all assets.

Some large institutions may employ tens if not hundreds of people to manually undertake qualitative compliance checks. With regulatory requirements increasing all the time, more human resources will be needed. This is neither scalable nor sustainable.

GRC tools lack the ability to isolate and identify applications associated with particular business processes, or the interrelationships between assets (ie the infrastructure supporting the applications such as devices and databases) and the people who interact with them. The disconnected, 'one-at-a-time' nature of qualitative assessment prevents a complete contextual understanding of what risk/non-compliance means to the business, making it incredibly difficult to accurately assess the cumulative risk generated by 'toxic combinations' of risk factors.

With all this in mind it's hardly surprising that GRC tools struggle to meet the needs of measuring cyber risk and compliance, given the complexities of the fragmented cyber domain.

GRC tools are an easy target for criticism because they aren't designed for quantitative assessment. For cyber professionals, it's a case of "welcome to my world", because pinpointing real, data-driven answers to these questions is not at all straightforward. It comes back to the challenge shared by cyber and GRC teams alike: no touchstone for knowing the extent (and detail) of ongoing controls coverage and effectiveness across different asset types.

With the growing spectrum of internal and external policy and regulatory concerns, managing security and IT data is becoming more difficult

Typically, the buck stops with the security operations team. For example, when presented with a need to substantiate regulatory compliance with quantitative (rather than qualitative) data, GRC teams often outsource the manual cycles needed to extract the information to security colleagues. This causes friction inside organisations, particularly when there is time pressure on responding to external requests from regulators and means security teams' time is not spent on security.

Panaseer commissioned a study among CISOs and other security leaders at large US and UK-based financial institutions earlier this year to drill into this very question. It found that, on average, GRC teams were requesting metrics from their security colleagues once every 16 days, and that working on these requests consumed upwards of five days a month.

Large financial organisations are worth studying because of their often-superior technological maturity and extensive internal resources relative to other businesses. They are also highly accustomed to coping with policy frameworks and regulators regularly posing searching cyber-related compliance questions. If *they* are struggling to bridge the gap between cyber and GRC, then the issues could be more profound elsewhere.

The follow-up study focused on exactly that, asking GRC leaders their perspective on the problem. It found less than half (41 percent) of the most senior risk and compliance professionals at these businesses were 'very confident' in their ability to fulfil the security-related requests of regulators in a timely manner. Even more worryingly, only 27.5 percent were 'very satisfied' their organisation's security reports align to regulatory compliance needs like GDPR and CCPA.

CONTINUOUS CONTROLS MONITORING LETS ORGANISATIONS BRIDGE THE CYBER AND GRC DIVIDE

The lack of confidence goes on. Only 39 percent were "very confident" in the accuracy of security data provided to regulators on request, with another seven percent "neither confident nor unconfident" – a fairly damning indictment for risk and compliance functions endeavouring to satisfy the statutory obligations of large financial institutions.

The survey also found the top two most significant security-related challenges for GRC leaders were "access to accurate data" (35 percent citing it top) and "number of report requests to deal with" (29 percent).

The bad news is that, without addressing the underlying challenges of incomplete insight into the current status of security controls coverage, performance and gaps, the escalating regulatory climate is set to add more pressure.

The jurisdictions with data privacy laws are growing (120 countries as at 2020), as is the depth and coverage of specific regulatory requirements. Remember that an organisation has only to register a presence in one of these jurisdictions to come under the auspices of its regulatory control, compounding the overload on GRC teams and – inevitably – their cyber colleagues.

We've seen regulatory requests becoming more time sensitive too, acting on the expectation that properly functioning governance processes will have little difficulty responding. For example, the new Monetary Authority of Singapore Notice 655 on Cyber Hygiene requires banks to attest to having endpoint detection and response (EDR) software deployed and operational on every asset. This might appear a slam-dunk for the traditional qualitative approach where someone's "informed opinion" effectively represents the threshold, but this would clearly not be a confident, data-driven position to give to a regulator with the power of life or death over your trading licence.

Such questions warrant a detailed examination of both assets and security controls. But is it good enough to undergo a painful manual effort to arrive at an inferior 'point-in-time' assessment rather than an up-to-date assessment accurately validated at all times? Team leaders and executive leadership must confront whether they are putting a tick in a box (somewhat unconvincingly) or achieving the goal behind the regulation which, in this case, is to ensure that endpoints are actually protected from threats – something that the organisation's cyber function would do well to continuously monitor anyway.

Squaring this circle requires a rethink on data integrity, to a place where GRC tools can harness accurate data that's automated rather than manual, automatically access the required information and easily transform it into the formats different regulators demand.

With a consistent up-to-date view of control deployments, accuracy and confidence is improved since assessments will be derived from instrumentation instead of subjectivity.

The Gartner Hype Cycle for Risk Management, 2020 has identified a valuable and emerging technology that delivers on this – Continuous Controls Monitoring (CCM). The report defining CCM for the first time came out in July 2020, but the technology itself has been developing for years and in live deployments since 2017.

Gartner defines CCM as: "a set of technologies that automates the assessment of operational controls' effectiveness and the identification of exceptions". Using CCM, organisations can bridge the cyber and GRC divide by: creating a comprehensive asset inventory including devices, applications, people, accounts and databases – that can be easily reflected to regulators/stakeholders; uncovering gaps in security controls deployment coverage wherever they are located – whether on-premise or in the cloud; adhering to internal policy compliance; isolating risks to mission-critical parts of the business; integrating with GRC tools to automatically populate them with security controls assurance data; going beyond qualitative analysis to quantitative – gaining fast access to facts that can be substantiated with data instead of subjective questionnaires; and mapping controls data to regulatory frameworks such as CIS or NIST.

Legitimate 'CCM' tools sit on top of existing tooling, ingest data from across security, IT and business tools, and use an entity resolution process to clean, normalise and de-duplicate data before correlating aggregated data to individual assets.

ONE OF THE BIGGEST CHALLENGES FOR GRC LEADERS IS ACCESS TO ACCURATE DATA

By being able to align security controls with framework standards, GRC teams can track and report adherence to best-practice standards and regulatory mandates. CCM self-service reporting capabilities will also enable them to access data from a common, real-time data repository and build custom reports in minutes without burdening the security team.

CCM can support the ultimate objective of establishing common data and information to make GRC and cyber team tasks seamless and interwoven. This common language will still require a change in communications and approach, but will enable harmony and help bridge the divide ●

Charaka Goonatilake, CTO Panaseer, has built and maintained a broad spectrum of big data infrastructure, systems and applications for BAE Systems AI services and products. He's spent the last five years engineering and building Hadoop-based security analytics applications to detect Cyber threats.

Cyber and GRC teams share the same goals, but there can be friction between them instead of harmony