TAG | KPMG

A SPECIAL PUBLICATION

REINVENTING CYBER BUDGETING

FRESH IDEAS FOR ENTERPRISE SECURITY IN AN ERA OF AI AND GROWING RISK

CONTENTS

INTRODUCTION

Enterprise security enters the AI era seeking new ideas to counter mounting risks.

DR. EDWARD AMOROSO,
CEO TAG INFOSPHERE INC., RESEARCH PROFESSOR, NYU

LAURENT GOBBI,
GLOBAL CYBER AND TECH RISK COE LEADER,
KPMG INTERNATIONAL

It is no secret that budgets for enterprise security teams have flattened considerably in 2026. This is the normal progression as any new discipline matures. And after two decades of often unbounded spending, chief information security officers (CISOs) have come to recognise that the budget music has stopped, so to speak. They must now learn to operate in a manner that their IT counterparts accepted many years ago.

The insight they are seeking—and the general topic that we address in this joint publication from the teams at KPMG and TAG Infosphere—is how these changes in budgeting should be managed. Some adjustments should be easy, such as renegotiating vendor contracts or focusing on automating manual tasks, but these are the obvious measures. The question is whether the cyber budgeting process can be truly reinvented?

We believe the answer is a definitive yes. And this publication provides our specific, detailed, and actionable guidance on how this should be done.

The articles we've included take several different approaches, each of which we believe provides useful guidance for CISOs and cyber budget managers. The approaches are derived from our collective experiences in the industry while negotiating our own budgets. But our insights also stem from decades of review of client security budgets for organisations of every size, shape, scope, and sector. This collection is a culmination of that experience and practice.

OVERVIEW OF THE ARTICLES

Our first article is from one of us (Ed Amoroso), offering a creative take on zero trust. Usually applied to how architectures deal with failed perimeters, it can also be used to rethink budgets. In essence, he shows how each budget line item must justify itself, with no help from any assumed carry-over. This is a truly different way to address budgets. You'll want to see how it works.

The second article is from Akhilesh Tuteja, former global cybersecurity practice lead at KPMG, and partner and national leader, Clients and Markets, for KPMG in India. His work outlines how risk can and must be the primary driver of modern enterprise security budgeting. He offers practical insights and a framework that can be used to develop a roadmap for how this can be done. His experience working with KPMG clients has been invaluable in the cyber budgeting process, and you will benefit by reading and applying his list of tenets.

The third article is from David Neuman, former CISO of Rackspace, former commander in the United States Air Force, and professor at the University of Texas at San Antonio. Neuman outlines the growing problem of risk compression and how it is becoming increasingly difficult to differentiate real risks from ones of lesser impact. Neuman will help you do a better job using cyber risk to justify and support future budget requests.

Dr. Jayne Goble, who leads KPMG UK's Operational Technology (OT) Cyber Security Services, wrote the fourth article. She notes that while the convergence of IT and OT has had benefits, it has also drastically increased the dangers of cyberattacks in OT environments. Goble sets out how tooling helps companies invest in defending their areas of greatest vulnerability and helps them avoid pouring funds into protecting operational assets not in danger. As she patiently explains, a realistic defence requires an approach that embraces both.

The fifth article is from Raj Cheema, KPMG's head of Cyber Healthcare in the UK and EMA. Everyone in this sector understands that healthcare consistently ranks as one of the industries most frequently targeted by ransomware attacks. Cheema will help you understand why, and he also offers advice on how to invest in cybersecurity in a way that will be most likely to protect the valuable patient data that attackers go after.

I (Ed Amoroso again) return with an article that addresses the provocative, but perhaps inevitable, question of whether artificial intelligence (AI) can be used in the budget process. We expect that many budget managers will be tempted to ask ChatGPT (or an equivalent) for guidance on budget optimisation. Amoroso will help you understand the pros (few) and cons (many) of this approach.

The seventh and final article in this publication is a panel discussion among experts and practitioners to discuss the modern budgeting process. We asked Martin Tyley from KPMG CRI, Andreas Wuchner of Wuchner Securities, Phil Huggins, director of the UK's Department of Health & Social Care, and Patty Ryan, CISO at QuidelOrtho, to participate in a roundtable to share their insights into the budget process. You won't want to miss what these cyber luminaries have to say.

DR. EDWARD AMOROSO

On first glance, you might be expecting that we will address how an enterprise security team obtains funds for zero trust design. While we hope such an effort would be successful, our focus here is different. We aim to leverage the general principles of zero trust design to rethink the cybersecurity budgeting process. And we will do so from the perspective of the security practitioner, as opposed to the finance team.

An optimisation problem has emerged for security managers. Specifically, they must find balance between two business metrics that are traveling in different directions: Budgets are flattening and perhaps shrinking, and cyber risk is increasing dramatically. Our view is that incremental changes to budgeting will not work to find the right balance. Instead, a truly different approach is needed.

WHAT DO WE MEAN BY ZERO TRUST?

The concept of zero trust emerged as a consequence of a flawed perimeter-centric design model for enterprise security. At its core, it rejects the assumption that any user, system, device, or process should be inherently trusted, regardless of whether it sits inside or outside the traditional perimeter. This design concept relies on continuous verification, least privilege enforcement, segmentation, and strong identity.

Zero trust emphasises breaking larger entities into smaller and more verifiable microdomains of risk and trust. That same approach—disaggregating a complex system into observable, measurable units—is what we should be applying to the cyber budgeting process. Rather than treating cybersecurity budgets as large, opaque categories that get rubber-stamped each year, we suggest slicing these budgets into functional components.

Dr. Edward Amoroso
CEO & Founder, TAG,
Research Professor, NYU

HOW ARE BUDGETS DONE NOW?

Our work advising security teams on cybersecurity budgets, as well as our own experience managing cyber budgets, gives us useful insight at TAG into how enterprise teams work the process. To start, we must acknowledge that a decade of robust growth in security, from roughly 2012 to 2022, involved many double-digit growth objectives year over year, with some organisations reporting essentially a blank check for cyber.

This has changed as our security discipline has matured into more of an IT-orientation towards budgets and return on investment (ROI). Now, more typically, budgeting involves incrementally adjusting the previous year's allocations, a process often led more by inertia and politics than by risk-based decision-making. Readers with experience in areas adjacent to cyber, such as IT operations and supply-chain management, will recognise this approach.

More specifically, we now see operating expenses, licensing, headcount, and professional services being rolled up into large line items, with little or no meaningful correlation to their actual impact on cyber-risk reduction. In most cases, if a tool or platform has been around for several years, it receives continued funding without much scrutiny. This has kept many cybersecurity vendors in business long after their usefulness for risk reduction.

At TAG, we believe that this type of budget planning, while straightforward, lacks transparency and flexibility, which are two attributes that seem quite foundational to modern cybersecurity design. Worse, when tough trade-offs must be made, the process often becomes even more arbitrary, perhaps favoring the most vocal managers and teams or the largest commercial vendors.

The result of all this is that we see teams now perpetuating a cycle where their legacy spending is preserved at the expense of innovation, and where budget justification becomes an exercise in narrative-building, rather than data-driven prioritisation. If we had to summarise, we'd say that today's cyber budgets are built on implicit trust: trust in past decisions, in vendor promises, and in muscle memory. This led us to consider zero trust as a solution.

HOW WOULD ZERO TRUST DESIGN APPLY TO BUDGETS?

Zero trust can offer security teams a useful framework to rethink their cybersecurity budgeting as a continuous process of verification and adaptive response. Obviously, this would be used as a back-office methodology, since chief financial officers and their teams would be largely unfamiliar with the model or concept (which might not be a terrible thing to rectify, by the way).

Let's see how this would work. Imagine each budget item, whether endpoint security license, a managed service, or a cloud security tool, is forced to justify its access to funds. Instead of assuming that existing tools deserve next year's investment, we demand fresh validation, just as we would require a device to reauthenticate before accessing sensitive data. This would transform budgeting from a one-time annual process into a living, evidence-driven investment cycle.

Furthermore, zero trust budgeting would encourage micro-segmentation of funding decisions. Rather than lumping entire categories together—as in, say, "network security" or "cloud security"—each micro-investment would be measured in terms of marginal risk reduction and business alignment. The security team could be engaged to create these categories based on some meaningful local attribute (e.g., business unit initiative or functional control).

As a simple illustration, a security team might isolate the 200K USD spent on data loss prevention (DLP) for remote workers and evaluate its effectiveness separately from the 500K USD spent on cloud access security brokers (CASBs) for SaaS platforms. Just as in zero trust architecture, this approach would support

continuous improvement, detect misaligned spending early, and support automation and orchestration for financial governance.

CASE STUDY: APPLYING ZERO TRUST TO A BANK'S 2026 CYBER BUDGET

Developing a meaningful case study for this new security budgeting concept requires that we think notionally, since we have not seen zero trust used in this manner. But with so much experience managing budgets, it is not difficult for us to create a realistic scenario. Readers are encouraged to modify or adjust what we include below to match up with their own local conditions.

"IF THE DECISION IS MADE TO UTILISE ZERO TRUST FOR BUDGETING, THEN WE ENVISION THE FIRST STEP AS REQUIRING A DISMANTLING (GULP) OF THE NOTION OF AUTOMATIC ENTITLEMENT FOR ANY LEGACY SYSTEM OR TOOL."

Let's start with a typical retail bank in the U.S. preparing its 2026 cybersecurity budget. During the past two or three years, we would expect that this bank's CISO would have followed a traditional budgeting approach—namely, rolling over the previous year's plan with modest adjustments. But facing flat budget growth and an evolving threat landscape, the CISO reads our article and decides to apply zero trust principles to budgeting.

The first step that would be taken here would involve breaking the security budget into smaller, verifiable units, each mapped to specific risk reduction goals, such as reducing phishing-related fraud or improving third-party risk visibility. Members of the security team would be required to present evidence for each investment, including but not limited to telemetry from prior years, red team findings, coverage gaps, and expected KPIs.

The result would be a fundamentally different internal dialogue with finance and executive leadership. Instead of lobbying for dollars based on historic spend, the CISO would present a trust-but-verify type plan that would emphasise adaptive funding tied to observed outcomes. For example, instead of a blanket 3M USD for endpoint security, the plan would allocate targeted micro-budgets for specific use cases like laptop hardening or DevOps controls.

Admittedly, this could initially introduce additional complexity to the budgeting process, and might even lengthen the planning, development, and negotiation periods. But with excellent tools, including artificial intelligence (AI) support, we are convinced that this would not only work but would quickly become an embedded practice. And over time, it would likely result in ancillary benefits. At a minimum, it would train a new lens on the budget line items and would help CISOs understand the components of their budgets in more granular ways.

ACTION PLAN FOR SECURITY BUDGET MANAGERS

Our advice to readers who might be interested in trying out this approach would be to first discuss the possibility with their local teams. If you are a CISO reading this article, gather your direct reports and discuss the idea. Our teams at TAG and at KPMG are more than happy to participate in this process, either as presenters of the idea or perhaps as participants who can offer external context.

If the decision is made to utilise zero trust for budgeting, then we envision the first step as requiring a dismantling (gulp) of the notion of automatic entitlement for any legacy system or tool. Budget managers would have to adopt a verification model for funding, requiring each line item to pass a series of checkpoints—and these would have to be developed based on the local environment.

Checkpoints might include answers to the following types of questions: What measurable risk does a given line-item address? What supporting telemetry or data exists to justify continued investment? What is the delta in protection versus alternative options? These questions should be asked with the flexibility to adjust, since this is your first time through this process.

The next step would involve the CISO and team identifying budget categories, as well as aligning funding units with risk domains and business functions. This would enable creation of a security investment map, where visibility and agility replace tradition and trust. Just as firewalls gave way to zero trust architecture, spreadsheets and silos must give way to adaptive budgeting platforms.

Furthermore, zero trust budgeting would encourage microsegmentation of funding decisions. Rather than lumping entire categories together—as in, say, "network security" or "cloud security"—each micro-investment would be measured in terms of marginal risk reduction and business alignment. The security team could be engaged to create these categories based on some meaningful local attribute (e.g., business unit initiative, functional control).

As a simple illustration, a security team might isolate the 200K USD spent on DLP for remote workers and evaluate its effectiveness separately from the 500K USD spent on CASB for SaaS platforms. Just as in zero trust architecture, this approach would support continuous improvement, detect misaligned spending early, and support automation and orchestration for financial governance.

CASE STUDY: APPLYING ZERO TRUST TO A BANK'S 2026 CYBER BUDGET

Developing a meaningful case study for this new security budgeting concept requires that we think notionally, since we have not seen zero trust used in this manner. But with so much experience managing budgets, it is not difficult for us to create a realistic scenario. Readers are encouraged to modify or adjust what we include below to match up with their own local conditions.

Let's start with a typical retail bank in the U.S. preparing its 2026 cybersecurity budget. During the past two or three years, we would expect that this bank's CISO would have followed a traditional budgeting approach—namely, rolling over the previous year's plan with modest adjustments. But facing flat budget growth and an evolving threat landscape, the CISO reads our article and decides to apply zero trust principles to budgeting.

The first step that would be taken here would involve breaking the security budget into smaller, verifiable units, each mapped to specific risk reduction goals, such as reducing phishing-related fraud or improving third-party risk visibility. Members of the security team would be asked to present evidence for each investment, including telemetry from prior years, red team findings, coverage gaps, and expected KPIs.

The result would be a fundamentally different internal dialogue with finance and executive leadership. Instead of lobbying for dollars based on historic spend, the CISO would present a trust-but-verify type plan that would emphasises adaptive funding tied to observed outcomes. For example, instead of a blanket 3M USD for endpoint security, the plan would allocate targeted micro-budgets for specific use cases like laptop hardening or DevOps controls.

Admittedly, this could introduce considerable complexity to the budgeting process, and might even lengthen the planning, development, and negotiation periods. But with excellent tools, including artificial intelligence (AI) support, we are convinced that this can work. At a minimum, it would train a new lens on the budget line items and would help CISOs understand the components of their budgets in more granular manners.

ACTION PLAN FOR SECURITY BUDGET MANAGERS

Our advice to readers who might be interested in trying out this approach would be to first discuss the possibility with their local teams. If you are a CISO reading this article, gather your direct reports and discuss the idea. Our teams at TAG and at KPMG are more than happy to participate in this process, either as presenters of the idea or perhaps as participants who can offer external context.

If the decision is made to utilise zero trust for budgeting, then we envision the first step as requiring a dismantling (gulp) of the notion of automatic entitlement for any legacy system or tool. Budget managers would have to adopt a verification model for funding, requiring each line item to pass a series of checkpoints—and these would have to be developed based on the local environment.

For example, checkpoints might include answers to the following types of questions: What measurable risk does a given line-item address? What supporting telemetry or data exists to justify continued investment? What is the delta in protection versus alternative options? These questions should be asked with the flexibility to adjust, since this is your first time through this process.

The next step would involve the CISO and team identifying budget categories, as well as aligning funding units with risk domains and business functions. This would enable creation of a sort of security investment map, where visibility and agility replace tradition and trust. Just as firewalls gave way to zero trust architecture, spreadsheets and silos must give way to adaptive budgeting platforms.

TAKEAWAYS

- **Today's cyber budgets are built on implicit trust: trust in past decisions, in vendor promises, and in muscle memory.**
- **Zero trust emphasises breaking larger entities into smaller and more verifiable microdomains of risk and trust. That same approach—disaggregating a complex system into observable, measurable units—is what we propose as applying to the cyber budgeting process.**
- **Instead of assuming that existing tools deserve next year's investment, we demand fresh validation, just as we would require a device to reauthenticate before accessing sensitive data. This transforms budgeting from a one-time annual process into a living, evidence-driven control cycle.**

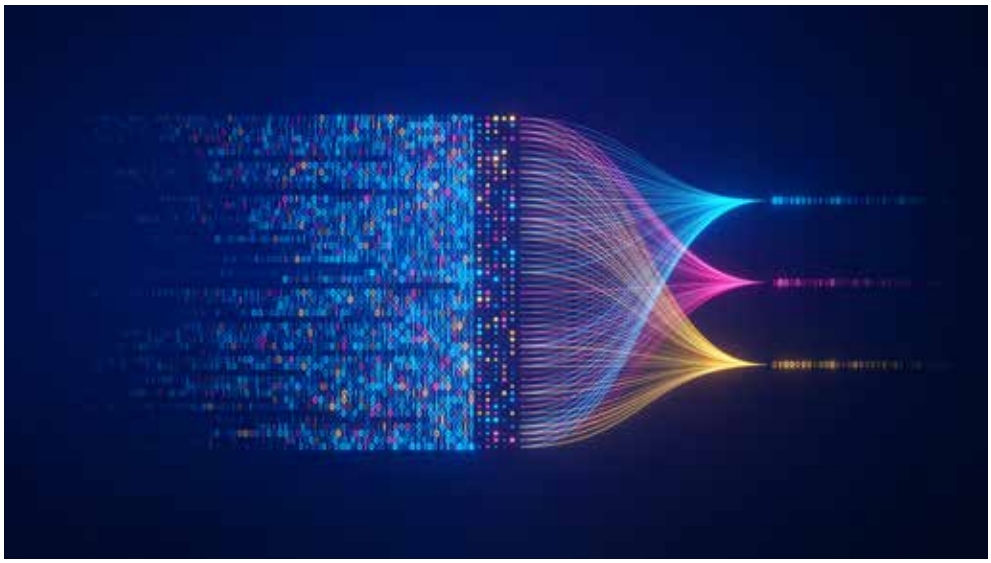

A RISK-BASED APPROACH TO CYBER BUDGETS

AKHILESH TUTEJA

Cybersecurity budgets are often poorly aligned with the actual level of risk to the organisation. Such misalignment can be driven by local challenges measuring and quantifying cyber risk, but it is compounded by the challenge of mapping perceived risk levels—accurate or otherwise—to security staff levels, controls, and approaches to risk mitigation.

The result is a budgeting process that is often inconsistent with the ultimate purpose of cybersecurity investment: namely, to reduce risk. Instead, enterprise security managers silently accept whatever they're allocated, or they distribute resources based on inertia rather than real exposure. I believe, rather, that they should tie their enterprise budgets to quantifiable cyber risk.

Admittedly, this is easier said than done. But this article suggests a practical framework for how risk can become the driver of budgeting decisions. Our experience of working with KPMG clients globally has consistently shown that when budgets are mapped to risks, with measurable business outcomes, organisations achieve greater resilience, better board alignment, and higher returns on their security investments.

PRINCIPLES OF RISK-DRIVEN BUDGETING

Our framework is based on three cybersecurity management principles. The principles provide a basis for a three-step process that will help companies better manage the risks in their budgets.

Akhilesh Tuteja
Partner & National Leader,
Client and Markets, KPMG in India

The first principle is alignment where every budget category is mapped to a documented and quantified risk. If a given mapping is unclear, then the spend should be challenged. For example, if an AI security platform is

proposed for deployment, but the purpose of the solution is unclear, then our framework would suggest that the investment be delayed until a real threat can be identified.

The second principle is adaptability. As should be evident to any cybersecurity practitioner, cyber risk evolves at varying speeds and with varying outcomes. Therefore, budgets must be flexible, adaptive, and continuously updated rather than set once a year in stone. We understand that this is not the typical approach. Managers are often handed rigid year-over-year, carry-over budgets with little room for change.

The third principle is transparency. A properly designed, risk-driven budget should allow all stakeholders to see the link from money spent to risk reduced. The implication, of course, is that the security function understands the risks that apply to the organisation. Transparency will be of little use if the risks are poorly identified or exist in some ad hoc format or representation.

"THIS CHANGE REPRESENTS A FUNDAMENTAL REORIENTATION THAT WOULD FORCE EVERY BUDGET ITEM TO BE JUSTIFIED BASED ON RISK REDUCTION."

To put these principles into practice, CISOs will need to consider moving away from line-item categories like endpoint or network security and toward categories such as "ransomware disruption risk" or "third-party access exposure." This change represents a fundamental reorientation that would force every budget item to be justified based on risk reduction. Let's see how this would work in practice.

A PRACTICAL FRAMEWORK

Our framework should be guided by the following core management steps:

Step 1 Identify Top Risks: Everything starts with risk identification and quantification. You may choose to build your profile from the organisation's risk register, or there might be a preferred quantification process or methodology. It is likely that for many enterprises, ransomware, insider misuse, third-party dependencies, and emerging categories like AI misuse will bubble up as the greatest risks.

Step 2 Map Current Spend: Every existing budget line included in the current spending plan should be mapped or recast in the context of one of these risk categories. Inevitably, you will find mismatches. That is, large allocations might emerge that are supporting risks that are no longer material, and gaps might emerge where critical risks lack adequate investment.

Step 3 Resolve Gaps and Coordinate with Procurement: Clean up tasks where gaps need to be resolved—either with proposed changes to vendor spend or changes to the staffing plan. The goal is to rebalance the portfolio. This will demand working with procurement, because purchase plans are typically driven by vendors, not by risk. So, coordination will be required to map the risk-based plan to an actual purchasing plan for vendors.

This approach requires some discipline and collaboration, because it introduces the new step of mapping budget to risk rather than the easier (but less effective) approach of buying into the usual categories, like endpoint, SIEM, MFA, and so on. We recommend this process, because it allows security budgets to function more like investment portfolios, which are constantly rebalanced based on where the greatest exposures exist.

TOOLS AND METRICS

Despite the fact that different teams will have different means for identifying risk, we strongly expect that cyber risk quantification (CRQ) will ultimately be required. Frameworks such as FAIR (**factor analysis of information risk**) and other quantitative methods might also offer the ability to assign financial values to risks. Useful metrics include risk reduction per dollar spent and a residual risk index, which will allow the CISO to demonstrate efficiency and effectiveness of spend.

These types of tools should also help to shift the narrative with senior leadership. Instead of making claims such as "we need another 2M USD for monitoring," the improved discussion would become something more like "for 2M USD we will reduce expected annualised loss from ransomware by 10 million." That is a language business leaders understand, and it reinvents how security investments are proposed, justified, and approved.

We understand that this approach is not without obstacles. Organisations often resist moving away from legacy spend, and legacy vendor relationships can distort priorities. Risk modeling also requires asking questions and collecting data that organisations are not used to answering or providing such as "if we experience a ransomware attack, how long are critical services likely to be unavailable?" Finally, this process requires collaboration across security, finance, and enterprise risk functions—a cultural shift that some enterprises find challenging.

THE BOTTOM LINE: AN ACTION PLAN FOR CISOS

We strongly recommend that every CISO considers a risk alignment audit of their current budget. They should ask which budget allocations map to which risks—establishing a joint working group with finance and enterprise risk managers. Finally, CISOs should consider presenting budgets in risk-justified terms. Done well, this transition creates budgets that adapt to threats, speak the language of management, and deliver reductions in enterprise risk.

TAKEAWAYS

- **An effective cybersecurity programme should reduce an organisation's risk. But often that's not reflected in the cyber budget, because that risk is hard to measure.**
- **Companies sometimes beef up their security staff, believing that this will reduce risk. But they do so without confirming that the two are aligned.**
- **Realigning budgets to address actual risks requires collaboration across security, finance, and enterprise risk functions—a cultural shift that some enterprises find challenging.**

DIFFERENTIATING CYBER RISKS: A STRATEGIC FRAMEWORK FOR ENTERPRISE RISK MANAGEMENT

DAVID NEUMAN

Chief information security officers (CISOs) face an increasingly untenable reality: Everything appears critical, yet nothing can truly be prioritised. As digital transformation accelerates and geopolitical tensions reshape the threat landscape, security leaders find themselves drowning in an ocean of seemingly existential risks—from AI-powered deepfake attacks and quantum computing threats to supply chain vulnerabilities and nation-state infiltration campaigns.

This phenomenon, which we term "risk compression," represents one of the most significant strategic challenges facing enterprise security today. When every vulnerability assessment flags hundreds of "high priority" issues and every threat intelligence briefing introduces new vectors requiring immediate attention, security organisations lose their ability to make meaningful distinctions between risks that could bring critical services to a standstill and those that merely make headlines.

What is the solution? The key to addressing the myriad cyber risks we face begins with diffentiating them. We have developed a four-dimensional assessment framework that enables organisations to move beyond vulnerability scoring, which can exacerbate the feeling that you're overwhelmed, toward impact-driven risk differentiation. Organisations implementing this approach focus instead on the effect of various risks on the business. When they follow the framework we lay out below, they demonstrate measurable improvements in resource allocation—with ROI improvements for systemic risk mitigation.

David Neuman
Senior Analyst, TAG

FOUR-DIMENSIONAL RISK ASSESSMENT FRAMEWORK

Our model differs from the traditional two-dimensional risk assessment approach that focuses on likelihood and impact. We believe that cyber risks cannot be adequately understood through simple probability and consequence calculations. They require deeper analysis of how threats materialise and propagate through interconnected business ecosystems, So we've expanded the model into a comprehensive four-dimensional framework that captures the dynamic nature of modern cyber threats.

Dimension 1: Enhanced Likelihood Assessment moves beyond probability calculations to encompass threat actor motivation, capability, and opportunity within specific industry and organisational contexts. This recognises that a vulnerability affecting industrial control systems carries fundamentally different likelihood profiles for a manufacturing company versus a financial services firm—not because the technical flaw differs, but because the threat actor ecosystems and organisational attack surfaces vary dramatically. Nation-state actors targeting intellectual property prioritise different industries than cybercriminals focused on financial fraud. They present distinct risk profiles that generic vulnerability databases are unlikely to capture.

Dimension 2: Comprehensive Impact Assessment reconceptualises impact measurement from technical damage assessment to comprehensive business disruption modeling. Operational continuity assessment translates cyber scenarios into specific business process disruptions. For manufacturing organisations, this requires modeling how different attack vectors could affect production schedules, supply chain coordination, and customer delivery commitments, moving beyond generic "system downtime" to quantified business consequences.

Dimension 3: Velocity Analysis measures both threat materialisation speed and available organisational response timeframes. This temporal dimension directly affects resource allocation between proactive monitoring and rapid response capabilities. Nation-state actors may establish persistence—that is, a long-term, undetected threat actor in an environment—over months or years, providing extended detection windows and allowing for strategic defensive positioning. In contrast, automated ransomware can encrypt critical systems within hours of initial compromise, demanding immediate response capabilities and compressed decision-making timeframes.

Dimension 4: Interdependency Reach Assessment quantifies how cyber events propagate beyond their initial technical footprints through connected business processes, third-party relationships, and stakeholder ecosystems. Some vulnerabilities remain contained within specific technical systems, requiring localised remediation efforts. Others cascade across supply chains, regulatory environments, and competitive landscapes simultaneously, creating complex recovery scenarios that extend far beyond the originally compromised systems. Understanding interdependency reach is critical for resource allocation decisions. A vulnerability in an isolated legacy system may score highly on technical severity metrics while having minimal interdependency reach, suggesting different investment priorities than a lower-severity vulnerability in a system central to business operations, customer interactions, and regulatory compliance.

DYNAMIC PRIORITISATION IMPLEMENTATION

The four-dimensional framework provides data based on these assessments. But the data still needs to make its way into the cyber budget. That happens through operationally valuable dynamic prioritisation mechanisms that adjust risk assessments based on business contexts and evolving threat landscapes.

"THIS SHIFT DEMANDS SYSTEMATIC INTEGRATION OF REAL-WORLD VALIDATION DATA INTO RISK SCORING PROCESSES."

For example, contextual threat alignment adjusts likelihood assessments based on current threat actor campaigns targeting specific industries or organisational profiles. When sophisticated actors launch coordinated attacks against particular verticals—as occurred with healthcare organisations during the pandemic or energy companies following geopolitical conflicts—relevant organisational vulnerabilities automatically receive elevated likelihood scores based on demonstrated targeting patterns.

Another is business calendar integration, which modifies impact assessments based on operational cycles, strategic initiatives, and stakeholder events. During earnings seasons, financial data protection receives enhanced impact weighting. During product launches, intellectual property and operational continuity risks demand increased focus. During merger and acquisition activities, competitive intelligence protection becomes paramount. This dynamic adjustment ensures risk prioritisation remains aligned with current business realities rather than static threat assessments.

Capability gap weighting adjusts both likelihood and impact scores based on organisational detection, response, and recovery capabilities for different threat categories. Organisations with sophisticated email security programmes may appropriately reduce phishing-related likelihood scores while increasing impact scores for supply chain security concerns where capabilities remain limited. This approach reduces the likelihood of both overinvestment in areas where strong controls already exist and underinvestment in areas where organisational vulnerabilities remain significant.

RISK SCORING AND VALIDATION

These capabilities stand in stark contrast to traditional risk scoring systems, which rely on standardised frameworks that treat organisational risk profiles as static entities, generating consistent scores across similar organisations but failing to capture the dynamic interplay between evolving threats, changing business priorities, and organisational capability maturation. The older approaches create false precision around inherently uncertain estimates, leading executives to make resource allocation decisions based on artificial mathematical certainty rather than informed uncertainty management.

The failure of checklist-based approaches becomes particularly apparent when organisations with identical technical configurations face dramatically different business consequences from similar cyber events. For instance, a financial services firm and a manufacturing company may receive identical scores for the same technical vulnerability, despite facing fundamentally different threat actor motivations, business impact scenarios, and regulatory consequences.

Strategic validation requires moving beyond theoretical risk calculations toward evidence-based assessment methodologies that ground risk scores in demonstrated threat capabilities, actual business impact relationships, and validated organisational response effectiveness. This shift demands systematic integration of real-world validation data into risk scoring processes, creating feedback loops that improve assessment accuracy over time.

REAL-WORLD VALIDATION METHODOLOGIES

There are other ways to validate risk data. Red team integration provides critical reality testing for theoretical risk assumptions. It can also be fine-tuned to focus on specific vulnerabilities. Rather than conducting red team exercises as isolated security tests, strategic validation integrates red team findings directly into risk scoring adjustments. When red teams successfully exploit supposedly low-risk vulnerabilities to achieve high-impact business disruption, risk scores require immediate recalibration based on demonstrated rather than theoretical attack pathways.

Effective red team validation extends beyond technical penetration testing to include business process disruption simulation. Red teams receive specific mandates to achieve business impact objectives—such as disrupting customer onboarding processes, compromising competitive intelligence, or triggering regulatory reporting requirements—rather than simply demonstrating technical system compromise. These exercises reveal gaps between theoretical risk models and actual organisational vulnerabilities, including those around human factors and process interdependencies that automated scanning cannot detect. They allow organisations to calibrate the data.

Threat intelligence calibration grounds risk scores in current adversary capabilities and demonstrated targeting patterns rather than generic vulnerability databases. Strategic threat intelligence integration tracks which specific threat actors actively target organisational industry verticals, geographic regions, and technology platforms, adjusting likelihood assessments based on observed rather than theoretical threat actor interest. When sophisticated actors demonstrate new attack techniques against similar organisational profiles, relevant risk scores automatically reflect elevated likelihood based on proven rather than speculative threat capabilities.

Security events that come close to harming a business may also provide useful information. Near-miss incident analysis can capture organisational vulnerability insights that include data about actual attack progression paths, detection capability gaps, and response timeline constraints that theoretical risk models cannot anticipate. They focus on business impact correlation rather than purely technical incident characteristics, revealing both vulnerabilities that theoretical models underestimate and controls that provide more effective protection than anticipated.

CONFIDENCE LEVEL ARCHITECTURE

HIGH CONFIDENCE SCORES	LOW CONFIDENCE SCORES
Extensive Threat Intelligence	Limited Threat Intelligence
Proven Attack Demonstrations	Theoretical Attack Scenarios
Robust Business Impact Modeling	Uncertain Business Impact Relationships

There are also ways to gauge the uncertainty inherent in these processes by using a scoring system that explicitly quantifies confidence levels for different risk assessments. High confidence scores derive from extensive threat intelligence, proven attack demonstrations, and robust business impact modeling. Low confidence scores reflect limited threat intelligence, theoretical attack scenarios, or uncertain business impact relationships. This framework prevents both underinvestment in proven threats and overinvestment in speculative scenarios.

Confidence scoring enables more sophisticated resource allocation decisions by allowing organisations to invest appropriately across different risk categories based on assessment certainty. High-impact risks with high-confidence assessments merit immediate, substantial investment. High-impact risks with low-confidence estimates require measured approaches that balance potential consequences against assessment uncertainty. This approach prevents decision paralysis while ensuring resources flow toward areas where risk understanding supports confident investment decisions.

Evidence-based confidence building establishes systematic processes for improving risk assessment accuracy through targeted intelligence collection and validation exercises. When risk assessments carry low confidence scores due to limited threat intelligence, organisations can prioritise specific intelligence collection activities to improve assessment reliability. When low confidence stems from uncertain business impact relationships, targeted business process analysis and impact modeling exercises can enhance score reliability over time.

The following case study demonstrates how the four-dimensional framework enables organisations to make defendable resource allocation decisions far better than those based on the traditional model. That's because they're based on quantified business impact analysis rather than technical vulnerability severity alone.

CASE STUDY: MULTIFACTOR AUTHENTICATION EXPANSION VS. DATA LOSS PREVENTION UPGRADES

Traditional assessment approaches struggle to differentiate between these competing investments, both of which address high-severity technical vulnerabilities. MFA expansion targets authentication weaknesses affecting multiple systems, while DLP upgrades address data exfiltration risks across enterprise applications. Standard risk frameworks often recommend parallel implementation despite resource constraints, failing to provide clear guidance.

By contrast, the four-dimensional framework reveals critical distinctions that enable confident prioritisation decisions. Likelihood analysis demonstrates that credential compromise represents a higher probability threat due to extensive attack tooling availability and proven social engineering techniques targeting users across all organisational levels. Data exfiltration, while serious, requires more sophisticated capabilities and targeted execution, typically following initial access achieved through other vectors, such as compromised credentials.

Credential compromise creates immediate, cascading impacts across operations, compliance, competitive position, and stakeholder trust—with business consequences materialising within hours. DLP bypass primarily affects data protection without immediate operational disruption, typically allowing detection and response time before material impact. This velocity difference drives effective resource deployment to target faster credential threats first.

FINANCIAL IMPACT ANALYSIS

Implementation Costs (Year 1)

- MFA expansion: 450K USD implementation + 180K USD annual maintenance = 630K USD
- DLP upgrades: 1.2M USD implementation + 300K USD annual costs = 1.5M USD

Risk Reduction Valuation (IBM Cost of Data Breach Methodology)

- MFA prevents an estimated 12M USD in potential breach costs. This figure applies a cascade multiplier of approximately 2.7 x to the **IBM-reported 4.45M USD average breach cost**, reflecting the compounding impact of credential-led incidents based on longer detection lifecycles, multisystem lateral movement, and broader downstream exposure.
- DLP prevents 3.2M USD in data exfiltration costs (2.8M USD average incident cost, adjusted for delayed

ROI Comparison

- MFA expansion: 19:1 return • DLP upgrades: 2:1 return

This analysis provides clear justification for immediate MFA prioritisation. The 850K USD saved from deferring DLP implementation can fund additional high-velocity security controls while maintaining superior overall risk reduction.

Of course, not all assessments result in a message this clear. But they are often quite helpful.

THE CISO'S TACTICAL IMPLEMENTATION PLAYBOOK

Successful implementation of strategic risk differentiation requires systematic organisational changes that embed the four-dimensional framework into existing planning and decision-making processes. The most critical transformation involves establishing quarterly risk calibration cycles that precede annual budgeting processes by 90 days. This provides sufficient time for comprehensive analysis, validation, and organisational alignment before budget submissions require finalisation.

These calibration cycles begin with comprehensive threat landscape analysis, requiring security leaders to work closely with threat intelligence teams to update likelihood assessments based on demonstrated rather than theoretical threat actor capabilities. This process typically requires two weeks and should focus on current adversary campaigns, emerging attack techniques, and industry-specific targeting patterns that affect organisational risk profiles. The analysis should specifically identify changes in threat actor motivation, capability development, and targeting priorities that could affect existing risk assessments.

"EVERY DOLLAR IN YOUR CYBERSECURITY BUDGET REPRESENTS A STRATEGIC CHOICE ABOUT WHICH RISKS YOUR ORGANISATION WILL ADDRESS AND WHICH IT WILL ACCEPT."

Business impact recalibration follows, requiring direct consultation with business unit leaders to validate impact assessments against current operational dependencies, strategic initiatives, and stakeholder relationships. This process typically requires three weeks and should include structured interviews with operations teams, regulatory compliance leaders, customer relationship managers, and strategic planning teams to ensure that impact modeling reflects actual business realities rather than security team assumptions about business consequences.

Cross-functional risk workshops conclude each calibration cycle, bringing together security, business, and executive stakeholders to review updated risk prioritisations and associated budget implications. These half-day sessions focus on validating risk differentiation logic, identifying potential gaps in assessment methodology, and building organisational consensus around resource allocation decisions that may require difficult trade-offs between competing security priorities.

ALIGNING INVESTMENTS WITH REALITY

Every dollar in your cybersecurity budget represents a strategic choice about which risks your organisation will address and which it will accept. The question is whether those choices are being made deliberately, with full visibility into business impact, or by default through the accumulated weight of vendor pressure, compliance checklists, and headline-driven fear. The four-dimensional framework transforms budget conversations from defensive justifications into strategic investment discussions. This is where security leaders can demonstrate precisely why a 630K MFA investment delivers 19:1 returns while a 1.5M DLP upgrade—despite addressing legitimate concerns—yields only 2:1. That clarity changes everything about how executives perceive security spending.

Begin your next budget cycle differently. Before defending line items inherited from previous years or responding to the latest threat du jour, apply rigorous risk differentiation to your top ten spending categories. Quantify the business impact each investment protects against, validate your assumptions through red team exercises and near-miss analysis, and present your board with investment recommendations grounded in demonstrated rather than theoretical value.

Security budgets will always face constraints. The organisations that thrive will be those that allocate constrained resources against their highest-consequence, highest-velocity risks rather than spreading investments thin across every vulnerability that earns a "critical" label.

TAKEAWAYS

- **When every vulnerability assessment flags hundreds of "high priority" issues and every threat intelligence briefing introduces new vectors requiring immediate attention, security organisations lose their ability to make meaningful distinctions between risks that could impact critical business processes and those that merely make headlines.**
- **Strategic validation requires moving beyond theoretical risk calculations toward evidence-based assessment methodologies that ground risk scores in demonstrated threat capabilities, actual business impact relationships, and validated organisational response effectiveness.**
- **Every dollar in an organisation's cybersecurity budget represents a strategic choice about which risks the organisation will address and which it will accept. The question is whether those choices are being made deliberately, with full visibility into business impact, or by default through the accumulated weight of vendor pressure, compliance checklists, and headline-driven fear.**

DEFENDING AGAINST ATTACKS THAT CAN SHUT DOWN POWER GRIDS

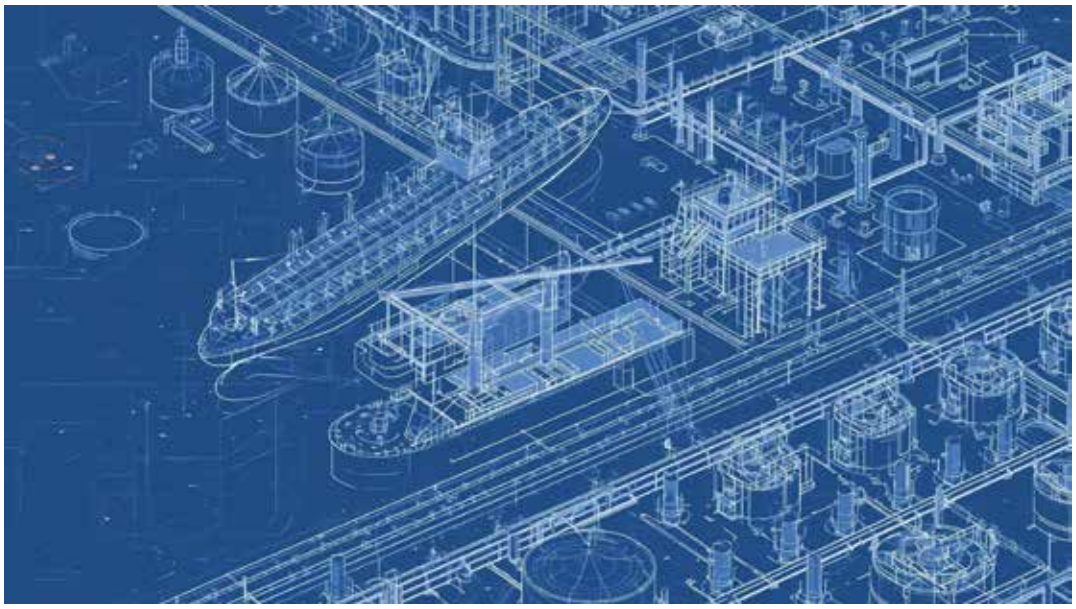

DR. JAYNE GOBLE

In the world of operational technology—the hardware and software that control physical equipment—these systems are becoming increasingly connected, which is shifting the stakes for a cyber event in many organisations from digital loss to physical devastation.

When cyber risk is inseparable from physical harm, a line of code is no longer just data; it is the command that opens a dam, shuts down a power grid, or overrides the safety sensors in a chemical plant. Anticipating the security and budgetary needs of operational technology equipment is a big part of my job as the KPMG partner in the UK who heads OT security. We work with oil and gas majors, utility providers, and industrial manufacturing companies around the globe.

Dr. Jayne Goble
Head of OT Cyber Security Services, KPMG in the UK

THE CONVERGENCE OF TWO WORLDS

Historically, OT like industrial control systems (ICS) were "air-gapped," meaning they were physically isolated from the internet. Today, the drive for efficiency has bridged this gap, connecting factory floors to corporate networks. While this

confers benefits—it allows for real-time monitoring of equipment—it also creates pathways for attackers to reach high-stakes physical assets. The result is that for OT, the priorities have to be safety and availability. A cyberattack can manifest in several dangerous ways:

- Malware can freeze a human-machine interface (HMI) so it displays "normal" readings while a process is actually reaching critical failure levels.
- Attackers can hijack programmable logic controllers (PLCs) to change chemical ratios in water treatment or increase pressure in gas pipelines to the point of rupture.
- A compromise in one sector, such as gas transmission, can cause a domino effect that collapses an electricity grid within minutes.
- In events like the **2017 TRITON malware cyberattack,** hackers targeted safety instrument systems (SIS) designed as the last line of defence to protect Saudi Arabian oil refineries. Fortunately, a coding error meant that only one facility was shut down temporarily. Had the attack succeeded, it could have caused explosions or led to the release of toxic hydrogen sulphide gas.

A DANGEROUS TREND

Cyber-physical system (CPS) and OT attacks have surged in recent years, shifting from simple digital disruption to large scale physical consequences. For example, the number of physical sites suffering operational impairment due to cyberattacks increased by 146% in 2024, rising from 412 sites in 2023 to 1,015 sites the following year. Ransomware attacks targeting industrial organisations spiked by 87% over the same period. By 2025, ransomware was present in 44% of all breaches, up from approximately 32% the previous year.

A **recent example** of the inseparable link between cyber and physical risk occurred in late December 2025, when Russia-linked threat actors (attributed by some experts as the state-sponsored Sandworm group and by others as another group known as Berserk Bear), targeted Poland's energy infrastructure.

The attack focused on approximately 30 distributed energy sites, including two combined heat and power plants and systems managing renewable energy sources like wind and solar farms. Polish officials alleged the goal was a deliberate attempt to cause a widespread blackout and destabilise the country during a period of extremely cold weather, which could have led to lethal consequences for the civilian population.

The attackers deployed a novel data-wiping malware (called DynoWiper) and gained access to OT systems by leveraging exposed network devices and vulnerabilities. They successfully disabled communications equipment, in some cases "bricking" the devices beyond repair. Poland's robust defence systems and effective incident response helped the country prevent the attack from causing any major power outages or widespread physical harm. However, the incident highlighted the significant vulnerability of modern, interconnected energy systems. It should be viewed as a major warning shot for all NATO members.

A NEW APPROACH FOR RESILIENCE

IT-OT convergence is a popular buzzword these days for driving efficiency. But connecting the two worlds is often dangerous. Convergence focuses on connectivity, often at the expense of security. And it's a central reason for the increase in OT attacks. Typically, these occur when a threat actor gains entry through a standard corporate IT network and "pivots" into the industrial control systems that manage physical machinery. Sometimes our post-incident forensic reviews reveal that the attackers didn't seem interested in the safety angle at all—until they accidentally stumbled into the OT side.

To address these escalating risks, we take a defensible approach which moves beyond a simple connection of IT and OT to a structured, resilient architecture. Organisations should focus on proactive defence, asset protection, and rapid recovery. This evolution treats operational resilience and safety as the foundations for defence rather than secondary concerns.

To move from a risky connected state to a defensible operating model, we provide specialised services that bridge the gap between digital security and physical safety. And we do this in a variety of ways.

Instead of just connecting devices, a defensible model uses tools to discover and monitor all OT assets in real-time, identifying anomalies before they cause physical harm. For instance, we deploy nonintrusive, passive monitoring tools (like **Armis** or **Radiflow**) to map the entire OT ecosystem without disrupting industrial processes.

We've supported many organisations in developing a "live" digital twin of their hardware. We don't just hand over a list of IP addresses; we provide a prioritised risk register that tells clients exactly which PLC is vulnerable and which anomaly requires immediate physical inspection.

"A KEY ELEMENT IN MARSHALLING RESOURCES IS DECIDING HOW TO INVEST. THAT'S WHERE CYBER RISK QUANTIFICATION (CRQ) COMES IN."

Network segmentation helps ensure that if a hacker breaches the IT corporate email, a defensible architecture prevents lateral movement to prevent a jump to the OT controls. Using O-PAS and OTSynapse (KPMG's proprietary tool), we work with clients to develop a blueprint for a secure, modular future. This is what they need to achieve a defensible, micro-segmented architecture.

HOW TO BUDGET FOR LONGTERM DEFENCE

A key element in marshalling resources is deciding how to invest. That's where cyber risk quantification (CRQ) comes in. When juggling a large portfolio of OT security projects, CRQ can be the difference between guessing and investing. In the physical world of OT, "high risk" is too vague; you need to know if a project will prevent a £10M outage—or a catastrophic safety event.

CRQ is helping leaders think strategically about risk reduction. It can help them understand their critical assets by decomposing losses against specific risk scenarios aligned to value streams within the organisation. It can also help them present the investment case to company decision-makers by explaining that the investment will enhance how the technology is networked and ensure that it's properly protected. And the leaders are not just thinking about the enterprise's backend offices. They're starting to think about using this approach to understand where to spend on equipment like wind turbines, gas turbines, and pipelines.

CRQ points out how vulnerable those assets would be if they weren't upgraded and maintained. It can show the difference between a default configuration that isn't monitored, and a hardened environment with 24/7 monitoring, in both reducing the likelihood and impact of any cyber event. This approach is threat-driven. It looks at threats in financial terms and can help leaders make sensible decisions on where to enhance security in the client's operational environment.

Equally important is what it *doesn't* require. Instead of looking at the technology and, out of fear, trying to determine how to secure *everything*, CRQ helps leaders prioritise maybe 20% of their operational assets where investment is most needed, and then leave off other equipment because it isn't critical or as exposed as they'd thought.

What makes this easier in the UK than it might be in other places is that money is sometimes readily available. One of my clients, a water company, is going to get a 100% asset uplift for maintenance and performance every four to five years. In the UK, we get the funding through the government. (It's also subsidised through a comparable format in the U.S. and Australia.) What happens as part of this multimillion-dollar upgrade, in addition to checking the waterways and the wastage, is considering the company's security needs.

Whether or not the funding is available, how we embed security should be part of any uplift. What we *don't* want to be doing is delivering the updated programme, and then thinking two years later about strapping on a load of security.

In this context, it's not a hard sell. Security by design is just common sense. The asset owners and operating officers who are in charge require this perspective already. All we should need to do is suggest, encourage, and validate.

TAKEAWAYS

- **The convergence of IT and OT has had benefits, but it has also drastically increased the dangers of cyberattacks in OT environments.**
- **A defensible operating model uses tools that monitor all OT assets in real-time, identifying anomalies before they can cause physical harm.**
- **Cyber risk quantification is a valuable tool that allows companies to take a threat-driven approach to directing their cyber budgets. This allows them to invest in defending areas of greatest vulnerability instead of embracing the unrealistic and unnecessary quest of protecting *all* of their operational assets.**

HEALTHCARE IN THE CROSSHAIRS: WE'VE COME A LONG WAY

RAJ CHEEMA

Hospitals are increasingly finding themselves in ransomware's "blast radius," and that risk is still rising. In Sophos' latest sector survey, 67% of healthcare organisations reported a ransomware hit in the past year, with recovery taking longer as attacks grow more complex. Meanwhile, IBM reports the average cost of a healthcare data breach is 10.93 million—the highest of any industry.

I've watched the growth of this onslaught from the frontline. I've spent over 15 years helping national bodies and healthcare providers defend themselves and their patients. The challenges have been growing year-on-year, and today the threat is even more substantial.

But it wasn't always like this. Not so long ago there was no bullseye on our back in healthcare. Historically, it was seen as above the fray—almost a sacred domain. After all, who would want to attack the industry that saves lives, the industry that, sooner or later, we all depend on?

Raj Cheema
Partner and Head of Cyber Healthcare for the UK and EMA at KPMG in the UK.

For decades, the perception that we were safe led to what I saw as a chronic underinvestment in cybersecurity, which in turn led to a myriad of vulnerabilities. During this time, the attackers have become less ethical, better skilled, and more strategic in their attacks.

The WannaCry ransomware attack in 2017 served as a pivotal moment for the industry, affecting both older computers running Microsoft Windows and newer systems lacking essential security updates. Approximately 70,000 devices were compromised during this incident. The United States government attributed the attack to North Korea, a conclusion supported by multiple other nations.

Among the impacted organisations, the National Health Service (NHS) in England and Scotland experienced significant disruptions. The cost of this attack to the UK healthcare sector alone was estimated to be £92 million. Over 19,000 appointments/surgical procedures were cancelled, and patient records reverted to paper formats that required subsequent restoration. That alone took months to put back onto IT systems. Was this event preventable? Possibly not. Did this event highlight inefficiencies in preparedness? Definitely. However, this event also acted as a catalyst and the genesis of cyber risk quantification (CRQ) playing a pivotal role in translating healthcare outcomes and impact into quantifiable loss.

"ATTACKERS TARGETING THE HEALTH AND CARE SECTOR HAVE RAPIDLY EVOLVED FROM SCATTER-GUN OPPORTUNISTS TO HIGHLY CAPABLE, ORGANISED, AND INCREASINGLY AUTOMATED ADVERSARIES."

One thing everyone agreed on was that it could have been far worse if the attack had struck at a different point in the week, or during winter when pressures on health services are intensified. It served as a wake-up call, especially as our reliance on technology and interconnected environments in the NHS continues to grow, and attackers become more sophisticated.

CHALLENGES ON ALL FRONTS

Challenge One: Leadership

We know the issues are real, and we know the impact poor cybersecurity can have. But can we measure it effectively in a way that allows us to quantify the impact, so that CEOs and C-Suites understand and can effectively help prioritise activities? Cyber was seen as a dark art—one often poorly understood by leaders. So, the first challenge the health and care sector faced was leadership becoming au-fait and informed in decision-making in this area.

Challenge Two: The Human Factor

Health and care is a people business. We care for people, ensure they live longer and healthier lives with their loved ones. But the challenge here is balancing what staff and patients want, which is timely access to great care, against the need to build safeguards to protect security. So, the second challenge that needed to be addressed is how do you build security design into a huge people business where cybersecurity is not their core skillset.

Challenge Three: External Danger

Attackers targeting the health and care sector have rapidly evolved from scatter-gun opportunists to highly capable, organised, and increasingly automated adversaries. They now exploit supply-chain weaknesses, harvest credentials at scale, and weaponise zero-day vulnerabilities. This can be seen in incidents like the Synnovis ransomware attack and widespread credential theft via infostealer malware. These threat actors operate with speed, precision, and commercial discipline, overwhelming legacy systems and fragmented defences. So, our third challenge was clear: figure out how the health and care sector can outpace attackers who are now smarter, faster, and far more sophisticated than ever before. At the same time, the industry must also make sure spend is optimised to create the biggest impact in reducing cyber risk, since funding is constrained and value needs to be demonstrated against other public service demands.

In the UK, leaders of the health and care system took the threat seriously, but also wanted to know the scale of the issue. I saw first-hand how one of the biggest hospital providers in the UK configured a mock phishing email attack to breach patient data—addressing each of the aforementioned challenges in the process. The results of this test were concerning but also insightful: It took only six hours from launching a phishing attack to having access to patient data. Notably, this wasn't accomplished using advanced technology or sophisticated methods, but rather through a straightforward attack scenario matching the abilities of a novice. This finding highlighted the magnitude of the problem and prompted a collaborative effort with the UK government to counter the threat.

WHERE WE ARE TODAY: CYBER AS A DETERMINANT OF CARE

Across global health systems, every pound of investment is under scrutiny. Public healthcare is serving larger, older populations with rising expectations, while operating under relentless financial pressure. Technology has become essential to meeting that challenge—improving productivity, enabling new models of care, and shifting activity closer to patients.

But this dependence on digital infrastructure has quietly changed the nature of risk.

Appointment platforms, digital diagnostics, shared care records, remote monitoring—these are no longer "IT systems." They are clinical enablers. When they fail, care slows. When they are compromised, patients are put at risk. Cybersecurity, therefore, has become inseparable from service continuity and patient safety.

The health and care sector has recognised this shift. In the UK, technology adoption is accelerating across clinical pathways and patient engagement, with security and privacy increasingly embedded into solution design. Secure-by-design principles are now being applied earlier, and new delivery models are assessed not just for clinical benefit, but for resilience and trustworthiness.

At the same time, the sector is addressing structural constraints. The shortage of cyber specialists remains acute, but regional capability models and a growing professional cyber cadre are beginning to emerge. Mandatory annual training has improved baseline awareness, while more advanced education is developing specialist expertise within trusts and integrated care systems.

Progress is real, but it is uneven. And it has exposed a deeper issue.

THE REAL PROBLEM: CYBER RISK IS STILL HARD TO GOVERN

Despite increased investment and awareness, many health leaders still struggle with a fundamental question: how much cyber risk are we actually carrying, and is our spend reducing it?

Traditional cyber reporting, which includes maturity scores, control compliance, heat maps, etc., rarely answers that. It tells boards what exists, but not what it means. It doesn't translate outages into cancelled procedures, or data breaches into recovery cost, or control gaps into patient harm. As a result, cyber risk has historically sat uncomfortably in executive conversations: acknowledged as serious, but difficult to prioritise against visible operational pressures. This is where many programmes stall—not because leaders don't care, but because they lack decision-grade insight.

THE TURNING POINT: QUANTIFYING CYBER IN HEALTHCARE TERMS

High-profile incidents have forced a shift. WannaCry showed how systemic vulnerabilities could cascade across an entire health service. The 2024 Synnovis ransomware attack reinforced the lesson, disrupting pathology services for months, cancelling procedures, and exposing sensitive data. These were not abstract IT failures. They were care delivery failures, with real human consequences.

What changed after these events was not just investment levels, but the way risk was being discussed.

Cyber risk quantification is a critical enabler. By using established quantitative risk approaches, health leaders are able to express cyber risk in financial and operational terms: annual loss expectancy, service disruption cost, recovery time, and the likelihood of patient harm. For the first time, cyber risk could be compared directly with other enterprise risks, and governed accordingly. In the NHS, this enabled a step change. CRQ is currently being used in many guises: from helping quantify the benefits of cyber transformation initiatives to helping over 200 organisations prioritise investment in cyber. This means that leaders can now ask: "Where will our next pound reduce the most risk to patients and services?"

FROM SPEND TO IMPACT: WHY CRQ MATTERS

Perhaps the most powerful benefit of CRQ is not prioritisation, but feedback.

By quantifying risk before and after interventions, organisations can finally assess effectiveness. If £10,000 invested in multifactor authentication reduces expected annual loss by only a marginal amount, that becomes visible. If an alternative control materially reduces outage risk or data compromise, that too becomes clear.

This shifts cybersecurity from a faith-based exercise to an evidence-based one—essential in a sector where resources are finite and opportunity cost is real. More importantly, it allows cyber leaders, clinicians, and executives to align around a shared objective: protecting care outcomes, not just systems.

THE BOTTOM LINE

Healthcare has come a long way since WannaCry. But the journey is far from over.

The sector now understands that cyber threats are not moral outrages or technical nuisances. They are real dangers that are persistent, well-resourced, and indifferent to patient impact. Adversaries optimise relentlessly. Healthcare cannot afford not to.

What gives cause for cautious confidence is not technology alone, but a change in mindset. Cybersecurity is being reframed as a core component of operational resilience and patient safety. Risk is being discussed in terms leaders understand. Investment decisions are becoming more disciplined, more transparent, and more defensible.

This is not a moment for complacency or self-congratulation. But it is evidence that healthcare is learning how to fight a modern threat on modern terms—by making cyber risk visible, measurable, and governable.

And in a system built to save lives, that matters more than ever.

TAKEAWAYS

- **For decades, the healthcare industry believed we were safe, which led to a chronic underinvestment in cybersecurity. At the same time, the attackers grew less ethical, better skilled, and more strategic in their attacks.**
- **By using established quantitative risk approaches, health leaders are able to express cyber risk in financial and operational terms: annual loss expectancy, service disruption cost, recovery time, and the likelihood of patient harm. For the first time, cyber risk could be compared directly with other enterprise risks.**
- **What gives cause for cautious confidence is not technology alone, but a change in mindset. Cybersecurity is being reframed as a core component of operational resilience and patient safety.**

PROSPECTS OF USING AI FOR CYBER BUDGETS

EDWARD AMOROSO

Let me start by asking you to be honest: If I suggested to you that by entering last year's cybersecurity budget numbers into ChatGPT, along with some general guidance on your shifting work priorities and business objectives, would you at least be *tempted* to go ahead and see what you get back? Wouldn't you want to learn what ChatGPT, or some other generative AI tool of choice, might have to say?

Let's ask this question another way: Should it be considered irresponsible for human executives managing cybersecurity programmes to *not* utilise AI tools for their budget optimisation, given that such tools might add useful insight or direction? **Or more to the point: Will AI tools soon emerge as a required component of financial due diligence?**

If your answer is yes, then I suspect you are part of a growing number of business executives, both in cybersecurity and other disciplines, who are feeling more comfortable using generative AI tools with LLM interfaces to provide input to decision-making.

And since this method simply produces input to the financials, the entire process can be viewed as passive and advisory in nature, since any suggestions from ChatGPT can be absorbed, casu-ally perused, or just tossed out. This makes the use of AI for budgeting a much easier pill to swallow for many traditional managers.

In this article, we outline the prospects for using AI tools to generate, optimise, maintain, and otherwise manage the enterprise cybersecurity budget. We will be honest in our assessment, setting aside any biases for or against continued human control or involvement in the process.[1]

[1] *Note that with this article penned by a human, your author is obviously biased toward, and on the side of, the biological entities.*

AI BUDGETING MODEL

To describe the various options for using AI tools in the budget process, we first create a simple model—not based on maturity, but rather on the degree to which the budgeting team chooses to make use of AI. We see three levels in the model, each with increasing intensity of dependence.

1. **Research Mode:** This first level uses generative AI to investigate issues and answer specific questions about line items.

2. **Advisory Mode:** This second level uses generative AI to make recommendations and offer suggestions on the manual budgeting process.

3. **Automation Mode:** This third level tasks generative AI with ingesting data and creating a proposed budget that can be reviewed by humans—or not.

While these levels might give the impression that tools such as ChatGPT are capable now of performing the full budgeting task in automation mode, we believe that this is highly premature. Our view instead is that the current state of AI in this domain is not ready to own or drive the process but can certainly help to inform it. That said, let's review the pros and cons of each level in the model.

LEVEL 1: RESEARCH MODE

The appeal of using AI in research mode for security budgeting is simple. LLMs like GPT-4 have been trained on troves of financial, operational, and cybersecurity data. They're adept at pattern recognition and synthesis. They can help surface overlooked categories, suggest risk-aligned investments, and even generate budget narratives that tie spend to outcomes in language that a board or CFO can understand.

To that end, we believe that **the use of generative AI to assist in the cybersecurity budgeting process is highly recommended.** Some readers might protest that this will lead to sloppy usage and premature dependence on these tools, but so long as the process is properly managed, we see mostly upside potential for using generative AI to research and assist with aspects of the cyber budgeting process—especially in times of limited budget growth.

LEVEL 2: ADVISORY MODE

The decision to use generative AI tools to actually propose draft budgets is a natural progression from research mode. In fact, it seems unlikely that budget teams grappling with the complexity of rationalising security investments would not be at least curious to see what tools like GPT-4 would provide. Again, if the process is properly managed, then this might be a harmless exercise—and perhaps even an enlightening one.

The challenge is that at a time when security leaders are expected to quantify every dollar in terms of risk mitigation or business enablement, AI assistance might actually provide insights that humans would not see. When pressed for time, or facing an unfamiliar line item—say, budget for AI red-teaming or automated agent authentication—it might be helpful *and extremely tempting* to let the AI system use its knowledge (e.g., having read all the documentation) to propose a budget.

LEVEL 3: AUTOMATION MODE

The decision to allow AI tools to completely own and 100% manage the cybersecurity budget is obviously not recommended today, if only because the tools are not sufficiently ready to handle such a task accurately. But we warn readers who might be negative on this level in our model that your bank

balance and many other aspects of your financial life are fully automated. We understand that AI is different from automation, but you get the idea.

One advantage of allowing AI to eventually own this task involves benchmarking. By feeding sanitised and anonymised security budget categories into an internal model trained on historical data from hundreds of enterprise clients, AI tools can generate comparative assessments, such as: "Company X's identity spend is 25% higher than sector peers when normalised by employee count." This can be spectacularly useful.

"THE CHALLENGE TODAY, AND THE REASON LLMS ARE NOT RECOMMENDED IN AUTOMATION MODE, INVOLVES THE SUBTLE NUANCES OF BUSINESS CONTEXT."

THE PROBLEMS START WITH CONTEXT

The challenge today, and the reason LLMs are not recommended in Automation Mode, involves the subtle nuances of business context. If a security team asks an LLM for help with the cyber budget, this will result in a generic set of recommendations. For instance: rationalise overlapping tools, invest in training, double down on identity, and so forth. These will likely be mechanical, even predictable, recommendations that any observer can easily identify.

The problem is that mechanical formulas for budgeting based on familiar themes fall apart when applied to the nuanced, political, and organisational realities of a specific enterprise. AI doesn't know, for example, that your biggest contract is up for renegotiation next quarter. Or that your top vendor is out of favour with procurement. Or that your cloud migration has hit a bureaucratic snag. Or that two different group leaders just do not get along.

What's worse, AI models are still known to hallucinate. Ask for a 2026 security budget roadmap tailored to a financial services company with 10,000 employees, and you might get a perfectly structured document, complete with citations to frameworks and tools that don't exist. That's dangerous, especially in a budgeting context where decisions cascade into real spending, real hires, and real gaps if misaligned.

To be clear, this is not a criticism of the technology itself. Generative AI does what it was trained to do: produce fluent, plausible text. But for something as sensitive as enterprise cybersecurity budgeting, where the margin for error is slim and the consequences are steep, "plausible" is not good enough. So, we will remain committed to recommending that CISOs not consider use of AI in Automation Mode for budgeting.

TOWARD A MIDDLE GROUND

Where AI can add value, and where TAG encourages experimentation, is in the pre-budgeting phase: the planning, the exploration, the idea generation. This implies use of AI in Advisory Mode. For example, here are some scenarios where we can imagine considerable value from using AI in the budgeting process.

- A security leader might use AI to summarise historical spend reports, identifying patterns or year-over-year anomalies.
- AI could help draft budget narratives that frame investments in terms of risk, compliance, and business enablement, improving communication with finance teams.
- Teams might prompt AI to suggest tooling for new initiatives, such as securing AI work-loads or tightening OT-IT integration, helping to populate new budget categories.

These scenarios should highlight our view that **the time has not yet occurred where CISOs will be handing over their budget process to a machine**. Instead, we recommend that CISOs think of AI as an assistant, one that is good at surfacing ideas, fast at parsing language, and tireless when the need arises for framing or structure. The implication is that the accountability will stay with the human decision-maker for budgeting. At least for now.

That said, **we do believe the time will come, and it might be sooner than most readers might find comfortable, when an AI platform will ingest all the relevant factors and provide a carefully reasoned budget proposal.** This AI-generated output will serve more as a baseline than as advisory data, and it will lead to a new ecosystem where AI plays a more fundamental role in corporate finance in general. This will represent a *profound change* in business.

BOTTOM LINE FOR CISOS

We encourage CISOs to begin experimenting with AI for budget analysis—under controlled conditions, of course. CISOs should use their approved local AI assistant (*not public ChatGPT, obviously*) to review and assess their current spend. This can include testing out some "what if" budget scenarios to see what the AI returns. There might be a few good ideas generated. And CISOs might feel better informed.

But again, just remember this: AI is currently a useful tool, but most definitely not a replacement for human judgment. If you're unsure whether a budget recommendation makes sense from your AI assistant, then do what you have always done. Do the background research, talk to your people—including your expert advisors—and trust your human instincts. This remains the best approach to budgeting . . . for now.

TAKEAWAYS

- **As of this writing, AI is not sufficiently advanced to automate the full cyber budgeting task.**
- **AI right now can help in a pre-budget advisory capacity by helping draft narratives that frame investments in terms of risk, compliance, and business enablement.**
- **Will AI tools soon emerge as a required component of financial due diligence? We don't know yet, but they could.**

IN CONVERSATION

TALKING AROUND THE TABLE TO WRESTLE WITH CYBER BUDGETS

Ed Amoroso

Martin Tyley

Andreas Wuchner

Patty Ryan

Phil Huggins

Ed Amoroso: I'm Ed Amoroso from TAG and from New York University. I want to welcome you to our discussion. If you do cybersecurity, you're involved in day-to-day protection and curation of infrastructure and systems and apps. And you know that budgeting is really the most—I'm going to go ahead and say it—may be the most important thing that you can do during the year to make sure that you have proper resources to meet your objectives. I have a blue ribbon panel here. I'm going to ask each of them to say hello.

Patty Ryan: My name is Patty Ryan. I'm the CISO of QuidelOrtho, which is a medical device manufacturer. I've been in the C suite and security for about 20 years, and budgets continue to be a problem every one of those 20 years.

Amoroso: It's always a new experience. Just when you think "now I know it," there's the next year.

Phil Huggins: I'm Phil Huggins. I'm a civil servant in the UK government, and I'm the national CISO for the health and care sector, which is about 12% of our GDP—which, as you imagine, makes my job almost impossible. Financially, managing budgets has always been a challenge within an organisation. Managing budgets across a sector is crazy. During the 30 years I've had in cyber, I've moved from being very technical to sort of standing in front of whiteboards, spending my entire life just writing about money. Which is the joy of being a CISO these days, I think.

Amoroso: Some of the folks in our audience who live and work in the public sector are going to be interested in your guidance. And for them our goal is: We want to help you. We want you to pick up one or two tips that might help you as you're making your case with budgeting.

Andreas Wuchner: I'm Andreas Wuchner. I had the pleasure to run security and risk organisations for the last 30 years or so. I'm based in Switzerland, and I'm now an angel investor. I help VC [venture capital) and PE [private equity) firms and family offices with cybersecurity decisions. So this topic of budgeting to get money from the other side is the same kind of approach and same challenge as being a CISO. Therefore, I love the topic.

Amoroso: And now my co-conspirator here, Martin Tyley. He and I sort of dreamed this up together. My good friend from KPMG, Martin, say a couple of words and thanks again for planning this with me.

Martin Tyley: Not at all. My name is Martin Tyley. I'm a KPMG partner and co-leader of our cyber risk insights (CRI) solution. It's a global B2B SaaS solution for quantifying cyber risk in financial terms. We've got over 100 clients across 15 different industries, so I've had the privilege of seeing a wide range of risk profiles and the different ways that those boardroom conversations can go around budget.

Amoroso: Martin, let's stay with you on our first topic. We're making the presumption that cyber budgeting—both the tactics and strategy—could be improved. Maybe you could say a couple of words about some things that you've seen and how in general it could stand an upgrade.

"LOTS OF INFORMATION MAY BE AROUND THE TOOLS WE'RE GOING TO USE, OR THE THREATS THAT EXIST, BUT NOT ENOUGH ON WHAT THE FINANCIAL IMPACT WOULD BE IF AN EVENT WAS GOING TO OCCUR."

Martin Tyley

Tyley: I think it is an area where there's room for major improvement. There's broadly three, issues. First, decisions are often made with partial data. Lots of information may be around the tools we're going to use, or the threats that exist, but not enough on what the financial impact would be if an event was going to occur. Second, we've seen incidents that have really caused a shock wave with tens of millions or even hundreds of millions of losses as a result of different cyber incidents. But we're not bringing those numbers into the conversation when budgets are being set. And the other one, which we're probably all familiar with, is sometimes the chief information security officer [CISO] might be expected to own the risks, but really those risks belong to the leaders of the organisation. The ones who understand how revenue, operations, customers are going to be impacted. So fixing it, connecting those dots, valuing what you're protecting, being honest about the tradeoffs that you need to make—there has to be that top-down ownership as well of risk rather than, at its worst, treating the budget as kind of an annual negotiation around headcount and tools.

Amoroso: That's a very useful framework. Patty, do those points make sense to you?

Ryan: Yeah. I actually consider budgeting cyber an art not a science. There is that whole idea of the unknown. Even when you're doing a keep-the-lights-on budget, you don't know when the next zero day is going to come out that's going to claim an incredible amount of overhead to your organisation to fix or remediate. I really, really love the idea that you said, Martin, about the conversations, because the conversations are not about budget. Conversations are about the underlying risks the budget is mitigating, and the more that the leadership understands their role in that risk conversation, hopefully the budgets become more of an ongoing, iterative process, not a "we're going to plan what's going to happen 18 months from now," because that's impossible.

Amoroso: Sounds like you're someone who's been through this. Phil, what's been your experience? Can you bring that public sector view? Are these points making sense to you?

Huggins: Absolutely. The conversations is an important point, and I am somewhat critical of our discipline because we don't really properly consider benefits. Everybody else is walking into those budget conversations and saying what the executives are going to get for the investment they're making. We walk in and we say, "Well, the world's not going to catch fire if you give me this money. If you don't give me this money, the world will catch fire." And unfortunately, that only works once, and it only works when they've recently been burnt on something. What we tend to be pretty bad at in the discipline is really understanding the benefit of what we're doing. So it's: How much money is on the table in terms of risk? How many patients could get harmed, how many patients could die? How much trust could we lose from our patients for use of their data? I happen to sit in the public sector. It's the economic value of the data we hold. Will we lose if our strategic competitors get hold of it? If I'm talking to a clinician, they care about patient harm. If I'm talking to a CFO, they care about the financial harm, But what we don't do—and I think we don't help ourselves here—is we will do a huge amount of analysis to justify why we should be given the money. We'll go off and deliver the projects we set out. We very rarely come back and do a similar amount of analysis to prove why that was a good thing. And over time, people get suspicious that they're just piling money into a bucket, and they have no idea where it's going. So we as a discipline need to be better at understanding the economic analysis, the risk reduction benefits, and we need to be able to articulate them. And I think my journey with some very big budgeting conversations over the last few years has very much been: How do I make it easy for a cynical and skeptical audience to give me the money?

Amoroso: Certainly you've never encountered a cynical and skeptical audience in asking for funding, right? I'm just kidding. That's a very good point, Phil. Andreas, any thoughts on this topic?

Wuchner: Plenty of thoughts. The security area in itself is a very young area. And as long as the industry is still measuring IT security and cybersecurity spent in percentages of IT spent, we will never get to this point that Martin rightly said that we have a clear decoupling from risk and value. And at the same time, everything that was said I totally agree with. On the other side, if you discuss with CISOs budgeting, and you ask them to articulate what they are asking for, compared or benchmarked against the organisation's actual loss exposure, you see very quickly that a lot of CISOs are far away from the maturity which we are talking about here. It's so nice and easy to hide behind the IT budget, and go back and say, "Gartner says in my industry, it should be 13.2% of the IT spend. You don't give me enough money." But if you sit on the board, you don't care about that. You want to say, "OK, how much is at risk? What is the actual loss?" And you will be really, really surprised about the small amount of CISOs who can envision what you're talking about when you ask this question.

Amoroso: Very good point. You know, in a minute, I'm going to ask each of you about the delta between rising risk and flat budget. But I want to take minute and share with you a little trick that I offer to my coaching clients, because I coach a lot of CISOs on this. And I always ask, "When you're doing your budget discussions, are they ever in person with the decision-makers?" And most of them tell me yes, meaning they'll be in a conference room, and there'll be this discussion. I tell them. "Memorise your numbers. I mean literally memorise them, and then as you're talking, going through the material and showing the numbers, turn away from your PowerPoint, look at the audience, and continue to accurately state each number." I know that sounds ridiculous, but it seems like that kind of thing demonstrates in a performative way that you've prepared. If you've prepared, they do notice, Now we're going to start with Phil on this one: There is a delta that I think all of us would agree is happening. If your budget has leveled off, I don't think the offensive actors have said, "Oh, you have a leveling off budget. Wow. I'll back off a little bit to give you some space." I don't think that's happening. I think we see attacks getting more intense, more automated, more relentless. How do we deal with that?

Huggins: One of the key things we need is to be trusted. We are the specialist in our discipline. Our boards, our investment committees, our CFOs are not specialists in our discipline. And when we say to them, "Gobbledygook, 15%," they say, "I don't know what that means. Is that good or bad?" What we need is for them to trust that we, as you say, are on top of our numbers. We understand the problem we're trying to solve, and that we can articulate it to them in terms that they can understand. This is, you know, 15 days of factory production lost per year. Every time this happens, they can work out the rest of that. They know those numbers. It's also about—and this is something I've really learned—if you underspend, give the money back. Give it back really quickly. Don't fudge it, because the trust you get from financial professionals from handling the money well and saying, "I'm not going to spend this. Use it somewhere else. Get some value from it."—those sorts of things get you trust. And that gets you the ability to play a bit fast and loose sometimes, because they know that you're doing the right thing and you understand what you're doing. But when it comes to increasing threat and flat budgets, it comes down to knowing what matters. We, as a discipline, have a lot of frameworks now. Cyber frameworks didn't really exist 15 years ago, and now that we've got them and they're really good, the temptation is to say, "We need all of that." Well, in your organisation is it ransomware that's the problem? What are the key controls around ransomware? If I had to cut some money, where do I *not* cut from? Where are the things I need to turn to the CFO and say, "This is the red line." And again, you've got to understand the threat to your business. You've got to be able to articulate it, and you need to trust from them that when you say this is the red line, they believe you. And unfortunately, that doesn't come from FUD (fear, uncertainty, and doubt). That doesn't come from fear. To me, it all comes down to trust.

"[W]HEN IT COMES TO INCREASING THREAT AND FLAT BUDGETS, IT COMES DOWN TO KNOWING WHAT MATTERS."

Phil Huggins

Amoroso: That makes sense. Patty, you work in healthcare, and I think all of us would agree that that's one of the most intense sort of targeted areas. Did you agree with Phil that the FUD thing, which we probably all did years ago, does that work? How do you deal with that delta? Because you do feel the threat for sure in your sector.

Ryan: I feel like Phil was speaking my language with that particular response. But I'd like to take an extension on that, because when you start talking about risk, it's a cultural thing. You have to have an organisation that's willing to have the conversations about impact and likelihood, because that's the genesis of this. Nothing is going to be impenetrable. Something is going to happen, but something happening might be OK, depending on what happens. Obviously, a small laptop stolen that's already encrypted is different than network compromise that affects our manufacturer. Three different things. But the risk still needs to be discussed. And in some cases, it's accepted, it's mitigated, it's removed, it's transferred. And what I don't like about some of the conversations that are not continuous is the fact that you don't have that opportunity to build that SME relationship, to be the one to teach, to educate—to get away from the spooky "the sky is falling!" Because the sun is going to set and the sun is going to rise, no matter what happens in your office that day. So it really just becomes: How do you present everything and allow the executives to learn. In some cultures, it works very well because the executives want to be part of it. Unfortunately for others, it's dragging them along to get them to understand the responsibility, which is a huge, huge separate topic for your conversations.

Amoroso: That makes sense. I'd ask Andreas and Martin, as you provide your thoughts on this, to include maybe some commentary on the following. I've often had great success by saying, "We don't put seat

belts in cars to save you from a crash. We put them in so you can drive faster." Sometimes executives see that and think, "Oh, I see. This is an enabler, as opposed to an avoider." So Andreas first and Martin, I'd want your general thoughts, but if you could tell me about the enablement as maybe an approach here to demonstrate that we're not just trying to stop bad things, we're trying to enable good things.

Wuchner: We should separate two things. The question here is maturity. A very mature global organisation where people have a clear understanding of risk, where they have a budget which is linked to their control or risk maps, where you have a water line where you can say, "Oh, that's the budget. Everything below we will not tackle, and we will need someone to sign it off." The risk-aware organisation will be ready to accept that. "We don't have more money. We have to do it." But in my experience, there's so many organisations out there where this culture, as Patricia said, is not there. Where this whole risk clarity is not there, where they are not really in a position to say, "OK, we can do that" because they just don't know. So there's a lot of opportunity for this organisation to mature and get better and then have a business-related discussion. Versus the other ones, where they don't talk of risk, but it's all about IT and relationship. It's really important, as a security leader, to think as a businessperson. If I only have 100 bucks, I cannot spend 150. So how can I get the most out of the 100? How can I express that to my manager and say, "Look, we only have 100, but by doing this, we could get 120 out of it, even if you pay 100." You need to see how you can get the most out of the money you have. Sometimes creative budgeting or creative spending, reprioritisation of things—but being very clear about it—enables trust, as Phil said. And you sit in one boat and all of a sudden you're a businessperson. You're not the security geek or the IT person anymore, who's only saying no and bringing problems all the time.

"IT'S REALLY IMPORTANT, AS A SECURITY LEADER, TO THINK AS A BUSINESSPERSON. IF I ONLY HAVE 100 BUCKS, I CANNOT SPEND 150."

Andreas Wuchner

Amoroso: Very useful insight. Martin, you have such a such an expansive impact through your work and your team. I'm dying to hear what you think here.

Tyley: You want to go faster, right? When *can't* we drive faster? I would say we can't drive faster when we don't have the information to hand. So, if we haven't got enough budget, we've got to make some deliberate tradeoffs. We've got to be quite specific around that. I wouldn't walk outside the house into a storm before choosing between an umbrella and a sweater. It's obvious what I need. I've got the information in front of me. It's raining. But that's what can happen if you get into the room and you end up debating around MFA versus endpoint tools versus backups or training. Without a way to compare them, you can't make quick decisions. So, we have the information that can run thousands, even millions of simulations to show which combination of controls are going to deliver the biggest risk reduction for the money you have. To Phil's point around being trusted, a trusted CISO is one that is able to walk into the conversation not to say we need more, but to say, "Given this flat budget situation, these are the three things that are going to genuinely move the needle. And these are the ones that we can safely park for now. And here's why." You do that and you bring that into the conversation, and I think that's the one that allows you to move more quickly.

Amoroso: Your reference to information is such a good one. It's a good segue into the next topic: metrics. But before we get into metrics, I want to point out to our audience our use of metaphors. If you can find good ones, I've found that with CFOs, with executives, definitely with boards, you could have an hour conversation where you've thrown one useful metaphor in, and six months later they bring up the metaphor.

So if you can practice them with your team, if you can go over them with a coach, if you can go over them with your consulting partners, come up with a few good metaphors that you find are useful. And I think what you will notice is that your more difficult points could resonate a little bit more smoothly. Now Patty, metrics. You and I have both watched them evolve over the years from crazy dashboards you didn't understand to things that I think are much better. To what degree do you rely on them? This acronym CRQ comes up a lot: cyber risk quantification. How much do you use them as you're doing your budgeting?

"AS A MEDICAL DEVICE MANUFACTURER, WE DO RISK ASSESSMENTS IN EVERY SINGLE ONE OF OUR PRODUCTS. OUR HEAD OF R&D AND OUR CEO NEED TO SEE THOSE HEAT MAPS."

Patty Ryan

Ryan: I think we've come a long way from: "Let's check how many firewall alerts that we've gotten over the course of a week." Used to be a big one I think 20 years ago. But for me, it goes back to some of the fundamentals. How is the programme working? So if we start looking at our response capabilities, we start looking at the effectiveness in attack scenarios of different toolings to give a flavor of their investment, and at least to give the board and others the fact that what you purchased is working to x degree is, I think, an SME way. Trying to focus far more on that risk conversation. As a medical device manufacturer, we do risk assessments in every single one of our products. Our head of R&D and our CEO need to see those heat maps. Go into manufacturing, the same thing. So it's not just a red, yellow, green, kind of traffic light. The heat map gives you the perspective of probability and impact. The metrics evolution becomes less about point in time what's happening, and more about end to end. What is the life cycle of your programme, and how is the overlay happening?

Amoroso: Very, very good points, Patty. Andreas, in this quantification thing, again, you have a very broad perspective, given your different roles. What have you seen here? Are these quantification metrics essential? Have you seen ones that work?

Wuchner: Again, a maturity question. It's different if you're smaller than bigger. I'm a big fan of cyber risk quantification. But I have the pleasure to sit on four different international boards, and I get often asked, "What should we report?" Because security on a board is still an outlier, and if there's someone sitting on a board who can spell security, that's also an outlier. And I'm always going back to three different categories of metrics. I'm pretty sure, Patricia, in your world of medical devices, the people on the board have more technical understanding than in a bank, for example, where it's all about numbers only. So I look for three different groups of metrics. One outcome-related, where we measure things like impact, incidents, things like that. Then another one, exposure and control health metrics. Do we have control coverage gaps? Which is a completely different thing from the outcome. And last but not least, financial metrics, where there's estimated, analysed loss exposure, all this stuff. But you need to be mature for that, right? And depending on what the board is focusing on, if it's more a numbers board, if it's more a strategy board, and depending on the maturity of the security people doing that, I would go more for the one or the other. But with the three different categories, you are for sure able to address every single board with one or the other more focused flavor.

Amoroso: I think that's sensible. Phil, this all resonate with you?

Huggins: Yeah, absolutely. For my purposes, because I oversee a lot of organisations and they have to report to me. I ask for three ways of cutting the data. And I'll be honest, I'm awash in data. I need less data, more insight. But the three ways I cut it would be, first of all, the risk. What matters here?

"DO METRICS AND CRQ—IS IT KIND OF BINARY, LIKE ARE WE OK? OR IS THAT WAY TOO TRITE? DO WE NEED TO TRAIN THE FUNDING SOURCES THAT THEY HAVE TO THINK MULTIDIMENSIONALLY?"

Ed Amoroso

What are the most important things we should be worried about? The second one is compliance. We've set a baseline. Has everyone met it? If you're not over that bar, we're not having that conversation. But the one that is often missed is performance. And you could call that maturity, or you could call it other things, but knowing the risk is there and it's managed, knowing that you've met your compliance baselines, and then knowing that you're compliant. But how well are you actually using the things you're compliant to? So that could be coverage, that could be effectiveness, that could be outcomes. Really you have to take all three, because it can be really easy for someone to tell me they're massively compliant and be really bad at using the tools. And have no understanding of their risk. So you need to triangulate. The other thing I've found—and we've all found this. I think I heard it in what Andreas was saying. We're all awash in data. Every bit of tooling that we have produces data. There's metric upon metric upon metric available, most of which is of no interest or use. What I prefer doing—and it's something that came out of software development—is an approach called goal question metric. What is the goal we're trying to achieve? Are we getting closer or further away from the goal? And then, what things could we measure that might indicate an answer to the question? Maybe how many hits there are on the firewall is a really good indicator of whether or not we can answer a question around network security. Probably not. But in every case you can start from, "Are we doing better or worse at achieving a business goal?" You can put it in English language for the board members who don't understand the metric, but there's always that data guy on the board who will dive straight into the appendix and say, "Where's this number? This doesn't look right." So you're going to be given a metric as well, and it's being able to tie those together that tells that story.

Amoroso: Martin, I'm going to want to hear your thoughts here. But I want to tell you an anecdote and see what you think. I can't tell you how many times I've gone through 40 minutes of metrics and charts in front of the boss, taking them through everything. Then there'll be a break and we'll step away, and the boss will come over and say, "So, are we OK?" And I realised I was not answering the question, which in their mind is binary. Martin, you're a partner at KPMG. You probably get that a lot. Do metrics and CRQ—is it kind of binary, like are we OK? Or is that way too trite? Do we need to train the funding sources that they have to think multidimensionally?

Tyley: Why don't some of the dashboards work historically? Why do you get asked that question when you come out of the boardroom? Phil, you might be awash with data, but these organisations are awash with risks. Boards are awash with risks. They're having to think about staff risks. They're having to think about budgets, geopolitics, tariffs. There are so many things that they need to try and understand. If we come in and say we've stopped this number of things, or we've completed this number of exercises, how do I compare that as a board member to credit risk, conduct risk, or operational outages? It's really hard. But risk quantification allows you to talk about things in terms of a likelihood and an impact. For me, a good example would be something like, "We currently face a 5% chance of a ransomware event. If that happens, we think we're going to lose three days of sales or productivity." Let's say that's going to cost us 50M USD. That could happen in the next year. Now, if we invest half a million dollars in these three controls, we can bring that likelihood down to about 2%. A one-in-50-year likelihood that will be in line with our peers. At that stage, we're doing a few things. Where are we versus similar organisations, which most people care about. What are we going to do about it? We're going

to do these three things, and that's going to more than halve the likelihood of this event occurring. You may still get the odd question asking "So that's good, right?" Finally, present it in that way that helps the people in the room understand how it ranks and stacks alongside everything else that they're having to think through right now.

Amoroso: Really sensible advice. I'm sure your clients appreciate that. For our final topic, we're going to talk about the future and some thoughts on AI. At one end of the spectrum, it becomes a much more thoughtful conversation, human to human, more interaction. At the other we ask ChatGPT, "Generate my optimal budget argument." It spits it out: "Here you go." I'm guessing that neither end of the spectrum would be right. But curious where you guys stand on this. Andreas, what are your thoughts?

Wuchner: I'm sure AI will play a major part in it. For me, the evolution of budgeting for security will make this yearly budget cycle go away. Because we all know if something happens, money is no topic anymore. It's stop the bleeding. And I think for a topic which is risk-heavy, like security, we will sooner or later see a budget cycle which is much more linked to the risk cycles. So not yearly and not tied to IT spends. And if you look at what we know today about AI, it's fantastic when you do the things like control mapping, scenario generation, and things like that. And AI will bring much more data quality, narrative building. What I don't see at the moment is any kind of AI agent-based decision-making. AI agents are very well used in the security world to get rid of level one people and mistakes. But if you look into the business processes, everyone tells you business does really struggles with allowing agents to take decisions. I think the LLMs (large language models) will be there to optimise budgets, information, data, but there will be a strong human element. And maybe I'm just too old to see that in the future there will be agents, but I can't see at the moment any agent-to-agent conversations between finance and IT risk, where they will just say, "Oh, OK. That's nice. Let's spend 100M USD." I can't see that happening at the moment, but maybe the future looks different, and things like hallucination and loss of nuances are not topics anymore. I just doubt that.

Amoroso: I've created mock budgets in front of a group, like a group of students, and popped them into ChatGPT and said, "What do you think?" I've never seen a case where I went, "Aha! I would have never thought of that." When I experiment with that sort of thing, it's very mechanical and obvious. The nuance of knowing the personalities, knowing what your peers are doing, understanding the tendencies, having historical context, all of that is a very human thing. So I think at least in the near future, I can't imagine this *not* being a very human process. Patty, what's been your experience in thinking through where this is all going?

Ryan: I think AI is a huge opportunity to help CISOs understand risk levels across the tremendous amount of data that we're trying to understand. But also some predictive tasks. I hope that AI can help understand a little bit of future trends, which will also help with the budgeting process. But Andreas was right with the idea of no finite cycles or an iterative budget. But maybe something separate, like keep-the-lights-on work, which you can say is steady-state. And then there is the risk-based reactionary work or protection work that maybe goes through a different cycle. But we all have operations that we can say we need to do X amount of tickets and Y amount of time. So that might be something we separate out. But when we start looking at emerging technologies, we start looking at emerging threats. That would be iterative, but I think AI can help lock down your environment as to what you should be focusing on. Human element obviously validating that or agreeing with it. But right now, I can't see it spitting out my budget for me, but I would love to have it normalise some of the activity we're seeing in the environment to kind of like a single pane glass.

Amoroso: Phil, what does your crystal ball say here? What do you see coming?

Huggins: I'm going to take a slightly different approach to my two previous colleagues. I think budget cycles are set up for the convenience of finance colleagues, not for our convenience. And unless it makes their lives easier, they're never going to change. I think they're going to make us jump to their beat. I'd love to say it was different. And in small organisations that have adopted modern technology delivery techniques, you could see there being a little bit more flexibility. But in most big organisations,

they've got a reporting cycle. They've got an accounting cycle. And if you don't fall in line with that, you fall out. We're using LLMs to dive into the data and start giving us some of the analytics that allows us to get behind what's going on in a big old sector with tens of thousands of organisations. But I think there will be a limit. There will be a limit of the use of AI in risk. And the reason is AI requires you to be able to expose how you do something, and quite a lot of the way we do cyber risk is human estimation. So for example, if you're in a big, regulated bank, you've probably got control designs, and you probably have a control testing regime, so you can probably say how good your controls are working. If you're not in a big, regulated bank, I bet you don't have that. In which case when somebody says the inherent risk is a 5, 5 on a 5x5 scale and the residual risk is a 4, 3, and somebody else asks "How you know that?" The answer comes back: "Well, because the controls are working." How do you know that the controls are working? "Well, I thought it over, and that's what I think it would be." And, you can't get AI to go, "Well, I *think* it is." AI has to have some numbers to crunch. And I think what AI in risk is going to expose is how much inherent assumption sits in the risk process for a lot of organisations. It's not necessarily data-driven in the same way it would be if we had control design work and control testing money.

Amoroso: So interesting. You know, Phil, I sat on the board a large bank and found that the control ecosystem was very, very mature. No question. But as a CISO at a telcom, I found that the one thing we worried about, which was availability—making sure your iPhone worked—there we had the best controls available. I mean, that was our business. It's funny how the different industries have different perceptions of risk, and as a result, the controls in those areas become very exaggerated. Anyway, before we give Martin the final word, I want to thank Patty, and Phil, and Andreas for bringing your experience here. Martin and I planned this, and I don't think we could have picked better participants. Martin, why don't you take us home: What's the future, in terms of budgeting and some thoughts on AI?

Tyley: It has been really good. I don't think any of us see a world where AI is going to be setting the cyber budget on a Friday and messaging the board. I do see something more prosaic. And actually, I'm optimistic in this space because of the journey and some of the things I've seen us be able to do in the last few years. I think we're going to be in a situation where the questions around why is our cyber budget that number are actually going to be obvious. And they're going to be obvious because it's going to reflect a risk appetite that we're going to get much better at understanding. It's going to reflect the value of what's being protected and the impact of the controls that the organisation is going to choose to invest in. Every major budget decision is going to increasingly have a clear storyline. This is the scenario we're concerned about. If that happens, this is what it's going to cost us, and this is how that spend changes the odds. And I also am optimistic that it's going to feel more collaborative. I think it's going to be cyber, finance, operations, and product in the same room, looking at the same numbers, not the CISO speculating as to what the impact of this might be. AI and CRQ: Somebody mentioned standards earlier on and said we didn't really have them 15 years ago. I think AI is going to be huge in helping us understand how well controls are operating. For years—certainly the 26 years I've been at KPMG—the best way we've had of doing that is getting two experts or more in a room and then thrashing through the policies and the standards, the configuration of tools, talking about how it's working. And it's good, but that doesn't scale. And as we get more reliant on technology, and our technology becomes more diverse, that's a real challenge. AI, I think, is going to help us cut through that. It's going to do that heavy lifting. CRQ is going to do the heavy lifting on the data and the modeling so that humans can debate the tradeoffs rather than discussing the impacts. So I do think AI and CRQ are going to be really powerful tools, but they're just tools in the toolbox. They're going to be what make it quietly affordable to refresh the data, to run the scenarios through again, and to keep everyone working from a shared view so we have an evidence-based view of risk. And if we get there, I think there's a lot to be optimistic about in the way that we can face into this challenge we've set ourselves today of how to think about cyber budgeting.

Amoroso: Well, a note of optimism is always a good place to end a discussion. I want to thank all four of you. You guys are great.

Dr. Edward Amoroso
CEO Tag Infosphere Inc.,
Research Professor, NYU

Laurent Gobbi
Global Cyber and Tech Risk CoE Leader,
KPMG International

Dr. Edward Amoroso

(Using Zero Trust for Cyber Budgets 5, Prospects of Using AI for Cyber Budgets 28, Moderator: Talking Around the Table to Wrestle with Cyber Budgets 32)

Ed is the founder and CEO of TAG and a research professor at NYU. A former CSO of AT&T, his history at the company dates back to Bell Labs. Ed holds advanced degrees from Stevens Institute of Technology and Columbia Business School.

Akhilesh Tuteja

(A Risk-Based Approach to Cyber Budgets 10)

Akhilesh is Partner & National Leader, Client and Markets, KPMG in India. Akhilesh is passionate about developments in the area of information technology and how these can help businesses drive smart processes and effective outcomes. He has advised over 200 clients on matters relating to cybersecurity, IT strategy, selection of technologies, and helped them realise the business benefits of technology. He possesses good knowledge of behaviour psychology and is enthusiastic about addressing the issues of IT risks in a holistic manner, especially through application of user behaviour analytics..

David Neuman

(Differentiating Cyber Risks 13)

Dave is Lead Analyst at TAG. He previously led security strategy for a 60 billion global supply chain as VP & Business Security Officer at a Fortune 50 company and served as CISO at iHeartMedia and Rackspace Hosting. His career also includes leadership roles at EY and in the U.S. Air Force, where he commanded the first cyber hunting unit. Dave is an Adjunct Professor at University of Texas, San Antonio.

Dr.Jayne Goble

(Defining Cyber Budgets and Managing Cyber Risks in OT Security 20)

Jayne leads KPMG UK's OT Cyber Security Services and has over 20 years' experience working with a range of global clients to oversee and deliver a variety of capital projects. These range from responding to critical security failures of national infrastructure to deploying interception and intelligence platforms.

Raj Cheema

(Healthcare in the Crosshairs: We've Come a Long Way 24)

Raj serves as the head of cyber healthcare for the UK and EMA. As a Senior Partner at KPMG in the UK, he collaborates with governments worldwide to implement secure and resilient technology modernisation in the healthcare sector.

Discussion Participants

(Talking Around the Table to Wrestle with Cyber Budgets 32)

Phil Huggins

Director of the UK's Department of Health & Social Care
As National CISO for Health and Social Care in the Joint Cyber Unit in NHS England's Transformation Directorate, Phil is a proven security leader with over 25 years of experience in security and technology roles. Phil has designed and operated security for critical national infrastructure and sensitive government. Phil has advised and managed global financial services organisations and advised national regulators on cyber resilience and cybersecurity.

Patty Ryan

Chief Information Security Officer at QuidelOrtho
As CISO, Patty is responsible for defining the firm's global information security strategy, roadmap, and operating infrastructure. With over 35 years of IT experience and more than 20 in information security executive positions, she has worked in financial services (Bankers Trust, Citi, CitiStreet), life sciences (Johnson & Johnson), and legal (Fragomen, Del Rey, Bernsen, and Loewy LLP).

Martin Tyley

Global Lead Partner, KPMG CRI
Martin is a cyber leader, with over 25 years of experience. He helps cyber and business leaders understand the real risk of cyber events so they can invest where it matters most. Martin leads KPMG's global cyber risk quantification practice and its Cyber Risk Insights platform, which combine data-driven models and hard-won experience to show the likelihood and impact of attacks in plain, decision-ready terms. He works across multiple industries and is passionate about changing the way we talk about cybersecurity.

Andreas Wuchner

Founder, Wuchner Securities
Andreas is an experienced and recognised cyber & risk expert, business owner, board advisor and investor who has experience in operating within complex global business environments. He has lived in several international business hubs, including London, New York, Frankfurt, Sydney, and Zurich. Andreas has more than 25 years of experience and knowledge in his areas of expertise.

Lester Goodman: Director of Content, TAG
David Hechler: Editor, TAG:
Billy Lawrence: Global Cyber COO, KPMG International
Leonidas Lykos: Global Cyber Assistant Manager, KPMG International
Nabela Ahmed: Head of Marketing, KPMG CRI

REINVENTING CYBER BUDGETING

KPMG

ABOUT KPMG UK

KPMG is a leading professional services firm. It operates from 17 offices across the UK with approximately 17,000 partners and staff and is part of a global network across 138 countries.

Whether it's steering a business through a cyber crisis, sharing guidance on governance responsibilities, making introductions to industry specialists from across our global firm, or helping to tackle the inevitable challenges that come with growth—we help our clients overcome challenges big and small.

We can walk businesses through tried and tested means of getting to the end destination, while exploring alternative routes for businesses that are anything but ordinary.

kpmg.com/uk

TAG

ABOUT TAG

Recognised by *Fast Company*, TAG is a trusted next generation research and advisory company that utilises an AI-powered SaaS platform to deliver on-demand insights, guidance, and recommendations to enterprise teams, government agencies, and commercial vendors in cybersecurity and artificial intelligence.

tag-infosphere.com