Decisions & Disruptions 2: Decide Harder
A custom cyber security incident response exercise

Benjamin Shreeve
University of Bristol

Joseph Gardiner
University of Bristol

Joseph Hallett
University of Bristol

David Humphries
City of London Police

Awais Rashid
University of Bristol

Abstract

Cyber incident response is critical to business continuity—we describe a new exercise that challenges professionals to play the role of Chief Information Security Officer (CISO) for a major financial organisation. Teams must decide how organisational team and budget resources should be deployed across Enterprise Architecture (EA) upgrades and cyber incidents. Every choice made has an impact—some prevent whilst others may trigger new or continue current attacks. We explain how the underlying platform supports these interactions through a reactionary event mechanism that introduces events based on the current attack surface of the organisation. We explore how our platform manages to introduce randomness on top of triggered events to ensure that the exercise is not deterministic and better matches incidents in the *real world*. We conclude by describing next steps for the exercise and how we plan to use it in the future to better understand risk decision making.

1 Introduction

Major cyber security incidents regularly disrupt businesses, and in extreme circumstance have even bankrupted them. We have created a major new incident response exercise to help businesses. We have worked with specialists from law enforcement and major financial organisations to create an exercise that challenges teams to handle a major, responsive cyber security incident. The aim of the exercise is to expose senior managers and incident response teams to the time, resource and political pressures they will encounter whilst handling a major crisis whilst at the same time gathering granular decision-making data to inform cyber response handling research.

This work builds upon a freely available previous game released under a CC-BY-NC license: Decisions & Disruptions (D-D) [1]. D-D is a highly successful tabletop exercise utilised by police forces across the UK and businesses across the world. It was designed to explore how people make risk decisions around cyber-physical infrastructures. Whilst D-D has provided valuable practical and research insights [5, 12–14], we argue that it is limited by a deterministic mechanism, fixed and tied to one sector. We, therefore, propose a new game: Decisions & Disruptions 2: Decide Harder (D-D 2) which provides an extensible engine for risk decision making exercises that incorporates randomness and more complicated threat relationships that can be targeted towards any industry, rather than just for critical national infrastructure. This paper introduces our proposed new game D-D 2 and the new features it incorporates.

2 Related Work

Cyber security exercises are a popular way of raising awareness of the subject matter. A number of exercises have been developed specifically as part of University courses (e.g, [3, 10]). However, such exercises tend to have a relatively narrow scope related to the content of specific University courses and they are rarely validated with or used outside of academia.

Other existing exercises have been created to help raise awareness of cyber security issues in industry (e.g., [2, 4–6, 8, 11]). All of these exercises are tabletop exercises and with the exception of Frey et al. [5] all use card-based mechanisms. Whilst the card-based mechanisms are valuable for raising awareness they often limit how well exercises can reflect real-world scenarios. For example, such mechanisms rely on a mixed deck of cards (or several decks) which are then drawn from at random to provide game events. As such

Workshop on Deconstructing Gamified Approaches to Security and Privacy, located at USENIX Symposium on Usable Privacy and Security (SOUPS) 2023.
August 6–8, 2023, Anaheim, CA, United States.

[1] decisions-disruptions.org

it is hard to such mechanisms to capture the way that events in the real-world may be related with one event causing another to occur. This is not to say they are not of value, but they often emphasise learning of specific aspects rather than emulation of scenarios. For example, Hart et al. [8] introduce the Riskio serious game—their exercise is aimed at non-technical participants and challenges them to consider potential threat vectors and then identify possible countermeasures. This provides a no doubt valuable learning experience, but is not the same as exposing participants to an emulation of decision-making under pressure which we is the aim of our exercise. Hart et al. [7] have taken lessons learned from the development of Riskio [8] and used it, along with a careful analysis of a wide range of cyber security games (including D-D), to create the MOTENS design model for serious games. This model suggests that the most effective cyber security games include: **M**ultiple modes of learning—exposing players to a wide range of cyber security aspects; **O**wnership Self-Learning—Providing a range of options to help meeting learning objectives;**T**heory—that supports the design;**E**nvironment—creating an appropriate environment where people feel they can learn;**N**egotiation—moving toward a coaching and problem-based learning style;**S**elf learning—enabling participants to build upon their base knowledge.

Frey et al. [5] provide a different approach—their exercise *Decisions & Disruptions* provides teams with a Lego representation of a hydro-electric company with an plant (operations) site and separate office site. Teams play through 4 rounds, investing a budget each round to implement security controls on these sites. At the end of each round they discover what cyber attacks have befallen the business as a result of their choices. The game mechanism makes it clear to players that there is a direct correlation between the investment choices they have made and the events they have suffered. This is a significant improvement in terms of realism, enabling more complex attack scenarios to be developed. However, there are limitations to the mechanism created by Frey et al. [5], the game mechanism has hard-coded paths through it (e.g., event A will occur if security control B has not been purchased by round 3). This means that teams cannot replay the exercise as they will always experience the same set consequences for their actions.

We seek to build on the success of D-D by developing an exercise where the players affect not only the landscape in which they play but also the consequences of their choices. An exercise where events that have occurred and choices that have been made increase the likelihood of (or even directly cause) specific events to occur, and where events in turn can change the landscape. We work closely with CISOs and Cyber Security Incident Response Teams (CSIRTs) to develop a unique, replayable exercise, inspired by D-D that challenges teams to make cyber security decisions under the time, resource and political pressures of a series of unfolding cyber incidents.

3 The Exercise

The exercise challenges teams to help the CISO of a fictitious financial organisation handle a *very bad* cyber security week—each day of the week is represented by a 20 minute time-limited round. Each day, a series of cyber-related events will occur, which can be handled in various ways. Players have to decide which tasks the CISO's team will tackle each day to stay on top of what's happening to their company.

Teams have the ability to affect the overall attack-surface of the business and therein the type and effectiveness of possible attack vectors used against the business by updating assets through actions such as patching and staff training. How attacks (or symptoms) are handled affect whether attack(s) are prevented, continue or whether new related attacks are triggered. The exercise is designed to be played by CSIRT, CISOs and senior executives.

3.1 A Reactionary Event Mechanism

In order to be able to represent the complexities of real-world decision-making a complex and reactionary event mechanism was developed (see figure 1). Teams are tasked with protecting the EA of the business (see figure 2). They are able to affect the company's attack surface by investing staff hours each day into a wide range of possible upgrades (and in some cases by sacrificing profit). The attack surface is evaluated at the start of each round and used to identify which events can occur (out of a library of 120 possible events). Of those that can occur, 5 are randomly selected and added to the event list for the next round. Events can be specified to only occur given a particular set of criteria—such as a particular EA upgrade having been purchased or a previous event having occurred. Each choice that a team makes has an associated cost in terms of staff hours and company profitability. Events have been carefully designed in tandem with cyber security law enforcement officers to provide a realistic representation of the threats facing financial organisations. The range of events reflect the major MITRE ATT&CK [2] areas highlighted through interviews with Chief Information Security Officers (CISOs) in major financial organisations as problematic.

For many choices there are also consequences. These are either explicit feedback to teams as to the impact of their choices—including additional impact to hours/profitability/shareprice—or an in-game consequence which can include triggering other events in the future. These performance indicators were flagged up through interviews with CISOs and board members of major financial organisations as vital indications of performance during a cyber incident [1, 15]. These triggered events are then added to the event list for the next round (possibly bypassing any criteria evaluation). Events can also trigger other events to occur if they are ignored; for example if a team decides an event is

[2] https://attack.mitre.org/matrices/enterprise/

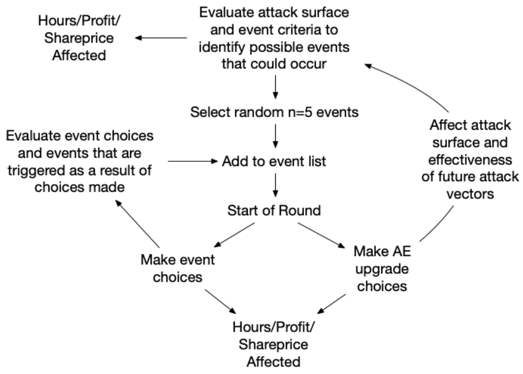

Figure 1: Summary of the different mechanisms in the game and the outcomes that can happen.

not important in one round it can still have an impact on their next round by queuing up an event to penalise their neglect. Some events are designated *'on-draw'* events—they affect the round as soon as they are drawn, deliberately introducing variance to the game. For example, an event may tell a team that a member of the CSIRT is ill that day; they therefore will start the round 8 hours down, with only 72 hours available to them to utilise. These *'on-draw'* events can force hours, profitability, shareprice and even make EA upgrade choices for teams before a round starts.

For example, as part of the game an event may occur where a staff member reports finding a USB stick on the ground. The CISO can decide whether to ignore it, to forensically analyse it (at a cost of time) or to destroy the device. Each decision will have impacts (and could in turn lead to more events occurring. To visualise these kinds of decisions we created an automated tool to create attack trees from the engine's database of events and explore what happens. This gives a quick pictorial guide to the consequences of any decision in the game and the threat landscape of any particular configuration.

3.2 Resources

Teams have to negotiate various resource constraints. Firstly, they have to contend with the challenge that a day has a finite length—in this case a configurable 20 minute limit. If they fail to utilise their resources effectively in that time then the game automatically moves forward to the next day. The primary unit of daily resources are the number of *'hours'* available to the team each day—that is, how many of staff in the CSIRT team that they are managing are available that day (the exercise assumes that there are 10 staff members, each with a maximum 8 hours per day, 80 hours in total per day). These hours are then allocated to event/EA actions.

The businesses performance is represented by two broad financial metrics that are both affected by choices made and consequences. Firstly, the team have to manage the overall long-term *'projected profit'* for the business—this starts at £500,000. However, it can be reduced through loss of business or fines that occur as a result of choices made. Teams can also choose to reinvest some of this projected profit into the business through certain EA upgrades or event choices which have associated financial costs. Secondly, teams have to consider the short-term performance of the business in the form of its *'share price'*—which can be affected both positively and negatively as a result of choices made from an initial value of £100. If the players actions result in the *'projected profit'* or *'share price'* reaching zero then the game ends.

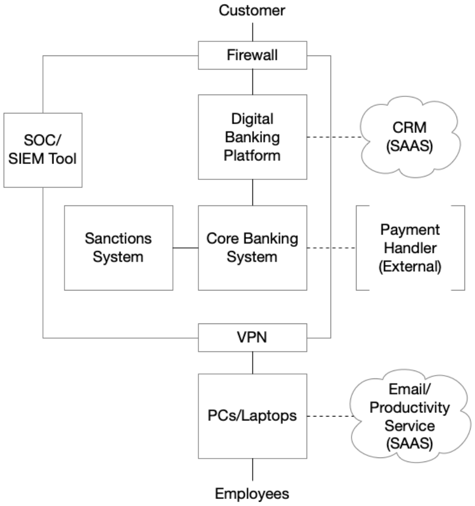

Figure 2: Financial Organisation enterprise architecture

3.3 Exercise Management

The exercise comprises two parts: a game master tool that tracks and evaluates the state of play and an associated physical card deck (see figure 3). The game master utilises a digital tool, written in Java, that reads a database containing details on the possible events (and their relationship to one another) as well as the assets and upgrades that possible for the EA. Each event/choices and asset/upgrade are also printed to a deck of cards providing session flexibility.

3.4 Typical Play Through

At the start of each *day* (round) the exercise interface updates to reflect the status of the business (see example in appendix B).

This right hand panel consists of a list of possible actions that can be undertaken to help improve the cyber security of their assets. The left hand panel shows a list of the most pertinent cyber events that have hit the business in the last 24 hours and that need resolving (a minimum of 5 occur each day) plus any that have yet to be actioned in some way that are remaining from previous rounds.

The EA upgrades (see example in appendix B) can be purchased at any point during a round and affect the attack surface of the game in the next round.

Each event presents presents teams with several choices about how they might choose to resolve each of these events (see example in appendix B). All events by default include ‘ignore this round’ which defers making a choice for that event to the next round. All choices are final—once a choice has been made (other than ignore), then it cannot be undone.

Teams work together to identify which events and EA upgrades they should dedicate resources towards in the round. As they make each choice the game provides them with instant feedback which may result in their hours, profitability or share price being immediately hit, which may in turn then limit the choices they had planned. When the round concludes the game utilises the reactionary event mechanism logic (see section 3.1 and figure 1) to evaluate which events will occur in the next round (as a result of changes the team have made to the businesses attack surface), or that the team have directly triggered (because of choices made in during the round).

The exercise continues this way through multiple rounds with teams starting each round with at least 5 events to address and a starting number of hours for the day affected by ‘on-draw’ events. The exercise concludes either when the round limit is reached (7 rounds) or when one of the failure criteria are triggered. The game has 3 failure criteria: the profitability of the business can reach zero, the share price can reach zero or the team can have 10 or more concurrent events open. We consider a team that has failed to resolve (or has triggered) 10 (or more) events to have reached a point of event saturation that could never be resolved in the real world. The extent to which consequences of event choices affect the profitability and shareprice of the organisation vary based on the magnitude of the event. Some consequences may have a positive impact whilst others may be severe, or even sufficient to bankrupt the organisation.

4 Design Choices

We have created a mechanism that can identify if an attack is viable based upon not only the current state of the organisations attack surface, but also on whether specific previous events have occurred and particular choices made. This means that events that are presented are far more representative of real-world scenarios.

We have taken this further by treating certain defensive approaches as *‘resistances’*. Defensive mechanisms, like *firewalls* and *antivirus*, can never be 100% effective. Instead, our game mechanism captures the current success rate of these mechanisms (for example the Firewall is only 60% effective) and then uses these values to establish if an attack that could be prevented by a firewall will occur by incorporating chance. We also provide teams with the ability to improve these resistances by investing in specific EA upgrades—and punish *negligence* for not doing so in a timely manner. In doing so we create an exercise where there is no longer a 1:1 relationship between events and defences and instead there are combinations of defences that can limit the likelihood of specific attack vectors being exploited—just like in the real world. The tool is itself also extensible. The events for a given game are taken from a database and can be rewritten for different sectors or markets. Penalties for different events and how events interlink is configurable allowing for a wide range of variations of the game to be created with relative ease.

Our design choices fulfil Hart et al.’s [7] MOTENS pedagogical design framework for serious cyber games: **M**ultiple Modes of Learning: The mechanism enables participants to experience not only a range of events, but also to explore the relationship between EA and system security and attacks. **O**wnership and Self-Learning: The underlying platform is flexible, allowing sessions to be tailored to specific learning objectives. **T**heory: The overarching game principle is informed by experiential learning theory [9]. **E**nvironment: The exercise environment is designed so teams have more forgiving opening rounds to help them familiarise themselves with the exercise mechanics and expectations. **N**egotiation: The promotion of shared decision-making amongst participants and the immediacy of feedback through the reactionary event mechanism move learning from static presentation to an immersive exploratory experience. Finally, **S**elf-Learning: the facilitation of these exercises enables participants to ask for clarifications when needed, enabling each participant to start from their own knowledge baseline.

5 Next steps and Future Work

The exercise is ready for testing with real-world CSIRTs and executives. This stage will be used to identify UI bugs, and refine the content of the event and asset database. Once this is complete the exercise will be made freely available under a Creative Commons License for organisations to use. We will continue to work with our industry partners using the exercise to gather data around how organisations prioritise and respond to different incidents and triggers.

Future work will focus on the creation of new asset and event databases for different sectors, making it possible to explore how different sectors and different technology stacks handle incident scenarios. Decisions and Disruptions has always helped organisations better understand how they make risk decisions: but now they have to *decide harder*.

Acknowledgments

We would like to thank Cyber Griffin, the City of London Police and the City of London Corporation for their funding and support.

References

[1] Ioannis Agrafiotis, Jason R C Nurse, Michael Goldsmith, Sadie Creese, and David Upton. A taxonomy of cyberharms: Defining the impacts of cyber-attacks and understanding how they propagate. *Journal of Cybersecurity*, 4(1), 10 2018. tyy006.

[2] Kristian Beckers and Sebastian Pape. A serious game for eliciting social engineering security requirements. In *2016 IEEE 24th International Requirements Engineering Conference (RE)*, pages 16–25. IEEE, 2016.

[3] Kevin Bock, George Hughey, and Dave Levin. King of the hill: A novel cybersecurity competition for teaching penetration testing. In *2018 USENIX Workshop on Advances in Security Education (ASE 18)*. USENIX Association, 2018.

[4] Tamara Denning, Adam Lerner, Adam Shostack, and Tadayoshi Kohno. Control-alt-hack: The design and evaluation of a card game for computer security awareness and education. In *CCS*, pages 915–928. ACM, 2013.

[5] Sylvain Frey, Awais Rashid, Pauline Anthonysamy, Maria Pinto-Albuquerque, and Syed Asad Naqvi. The good, the bad and the ugly: A study of security decisions in a cyber-physical systems game. *IEEE Transactions on Software Engineering*, 45(5):521–536, 2019.

[6] Mark Gondree and Zachary NJ Peterson. Valuing security by getting [d0x3d!]: Experiences with a network security board game. In *Presented as part of the 6th Workshop on Cyber Security Experimentation and Test*. USENIX, 2013.

[7] Stephen Hart, Basel Halak, and Vladimiro Sassone. Motens: A pedagogical design model for serious cyber games, 2021.

[8] Stephen Hart, Andrea Margheri, Federica Paci, and Vladimiro Sassone. Riskio: A serious game for cyber security awareness and education. *Computers & Security*, 95:101827, 2020.

[9] David A Kolb. *Experiential learning: Experience as the source of learning and development*. FT press, 2014.

[10] John R. Morelock and Zachary Peterson. Authenticity, ethicality, and motivation: A formal evaluation of a 10-week computer security alternate reality game for CS undergraduates. In Wu-chang Feng and Ashley L. Podhradsky, editors, *2018 USENIX Workshop on Advances in Security Education, ASE 2018, Baltimore, MD, USA, August 13, 2018*. USENIX Association, 2018.

[11] Adam Shostack. Elevation of privilege: Drawing developers into threat modeling. In *2014 USENIX Summit on Gaming, Games, and Gamification in Security Education (3GSE 14)*, 2014.

[12] Benjamin Shreeve, Catarina Gralha, Awais Rashid, João Araujo, and Miguel Goulão. Making sense of the unknown: How managers make cyber security decisions. *ACM Trans. Softw. Eng. Methodol.*, aug 2022. Just Accepted.

[13] Benjamin Shreeve, Joseph Hallett, Matthew Edwards, Pauline Anthonysamy, Sylvain Frey, and Awais Rashid. “So if Mr Blue Head here clicks the link…” Risk thinking in cyber security decision making. *Transactions on Privacy and Security*, 24(1):29, November 2020.

[14] Benjamin Shreeve, Joseph Hallett, Matthew Edwards, Kopo Marvin Ramokapane, Richard Atkins, and Awais Rashid. The best laid plans or lack thereof: Security decision-making of different stakeholder groups. *IEEE Transactions on Software Engineering*, 48(5):1515 – 1528, 2020.

[15] Onur Kemal Tosun. Cyber-attacks and stock market activity. *International Review of Financial Analysis*, 76:101795, 2021.

A Example Physical Cards

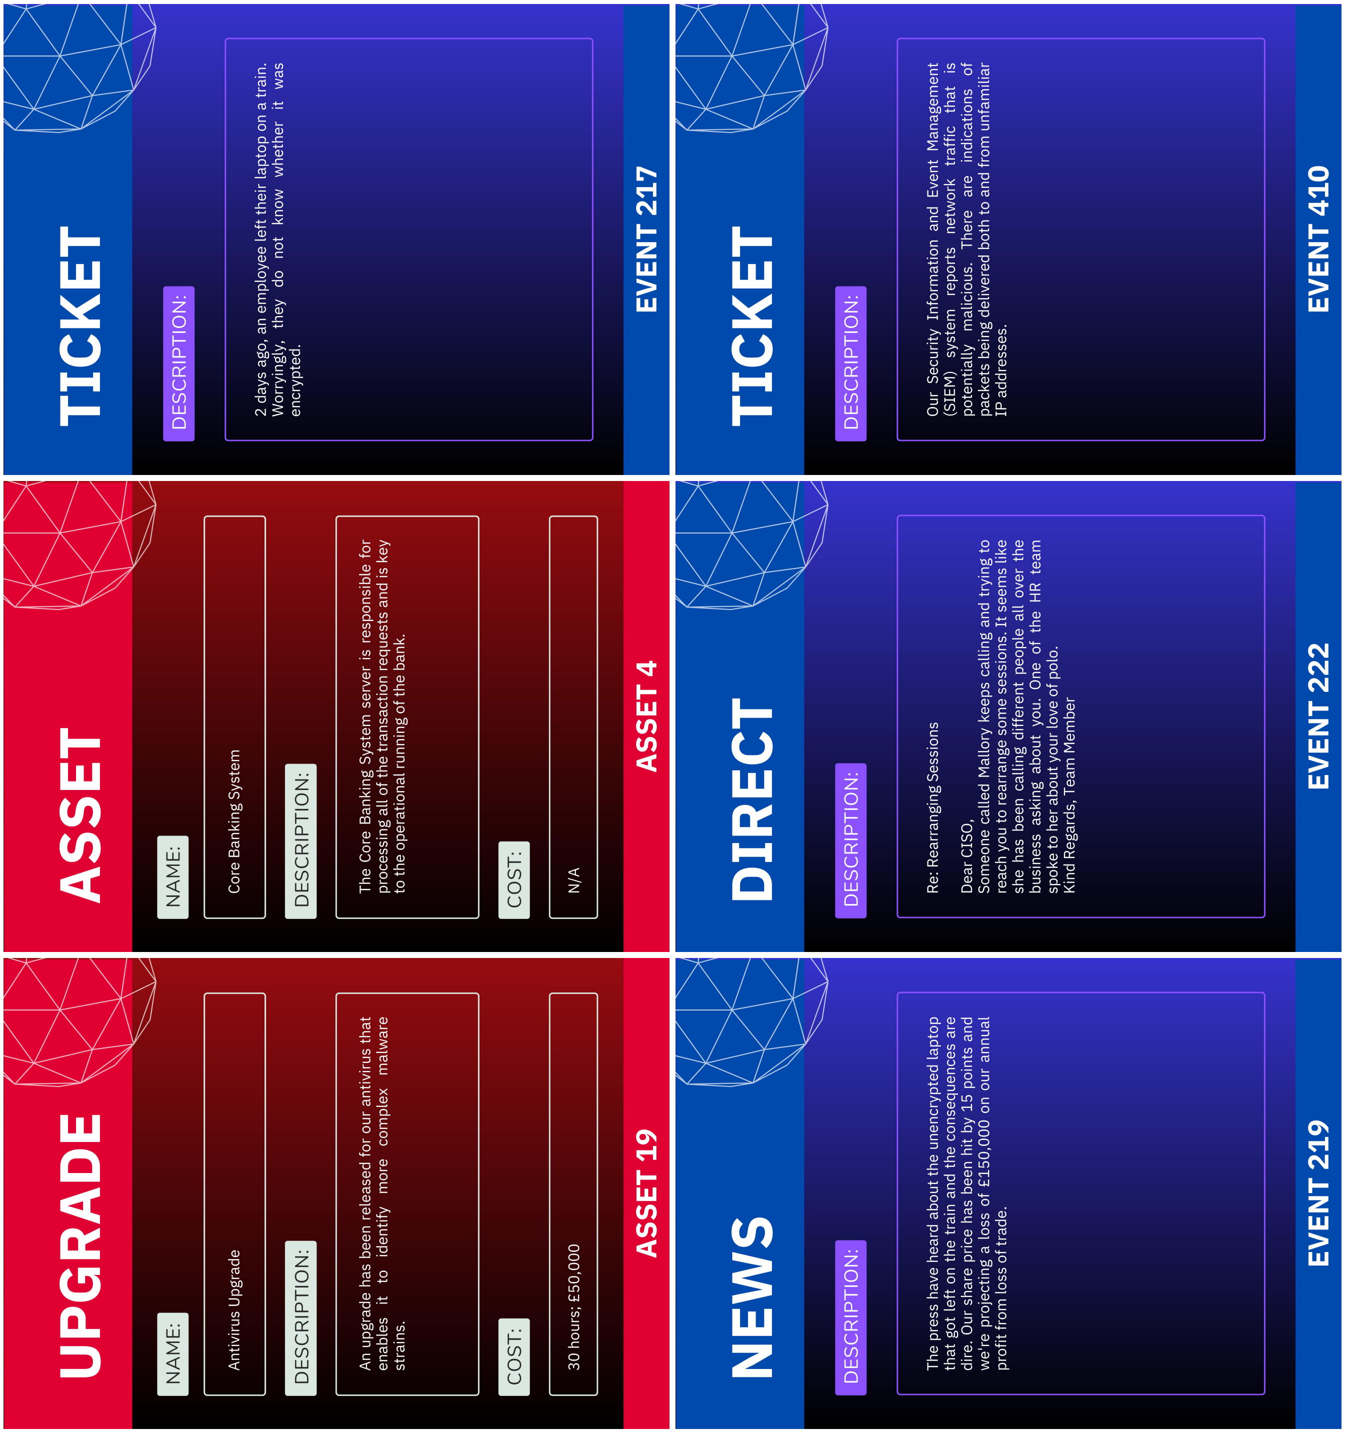

TICKET

DESCRIPTION:

Our Security Information and Event Management (SIEM) system reports network traffic that is potentially malicious. There are indications of packets being delivered both to and from unfamiliar IP addresses.

EVENT 410

DIRECT

DESCRIPTION:

Re: Rearranging Sessions

Dear CISO,
Someone called Mallory keeps calling and trying to reach you to rearrange some sessions. It seems like she has been calling different people all over the business asking about you. One of the HR team spoke to her about your love of polo.
Kind Regards, Team Member

EVENT 222

NEWS

DESCRIPTION:

The press have heard about the unencrypted laptop that got left on the train and the consequences are dire. Our share price has been hit by 15 points and we're projecting a loss of £150,000 on our annual profit from loss of trade.

EVENT 219

TICKET

DESCRIPTION:

2 days ago, an employee left their laptop on a train. Worryingly, they do not know whether it was encrypted.

EVENT 217

ASSET

NAME:

Core Banking System

DESCRIPTION:

The Core Banking System server is responsible for processing all of the transaction requests and is key to the operational running of the bank.

COST:

N/A

ASSET 4

UPGRADE

NAME:

Antivirus Upgrade

DESCRIPTION:

An upgrade has been released for our antivirus that enables it to identify more complex malware strains.

COST:

30 hours; £50,000

ASSET 19

Figure 3: Examples of physical cards used to play the game when the UI is unavailable.

B Exercise UI

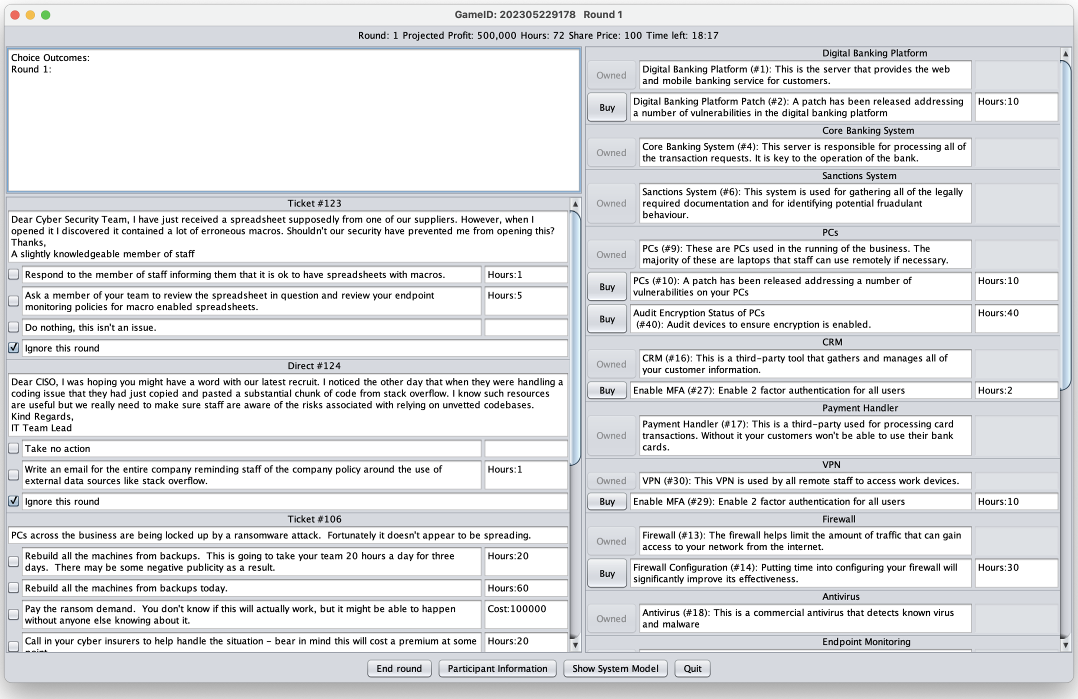

GameID: 202305229178 Round 1

Round: 1 Projected Profit: 500,000 Hours: 72 Share Price: 100 Time left: 18:17

Choice Outcomes:
Round 1:

Ticket #123

Dear Cyber Security Team, I have just received a spreadsheet supposedly from one of our suppliers. However, when I opened it I discovered it contained a lot of erroneous macros. Shouldn't our security have prevented me from opening this?
Thanks,
A slightly knowledgeable member of staff

□	Respond to the member of staff informing them that it is ok to have spreadsheets with macros.	Hours:1
□	Ask a member of your team to review the spreadsheet in question and review your endpoint monitoring policies for macro enabled spreadsheets.	Hours:5
□	Do nothing, this isn't an issue.	
☑	Ignore this round	

Direct #124

Dear CISO, I was hoping you might have a word with our latest recruit. I noticed the other day that when they were handling a coding issue that they had just copied and pasted a substantial chunk of code from stack overflow. I know such resources are useful but we really need to make sure staff are aware of the risks associated with relying on unvetted codebases.
Kind Regards,
IT Team Lead

□	Take no action	
□	Write an email for the entire company reminding staff of the company policy around the use of external data sources like stack overflow.	Hours:1
☑	Ignore this round	

Ticket #106

PCs across the business are being locked up by a ransomware attack. Fortunately it doesn't appear to be spreading.

□	Rebuild all the machines from backups. This is going to take your team 20 hours a day for three days. There may be some negative publicity as a result.	Hours:20
□	Rebuild all the machines from backups today.	Hours:60
□	Pay the ransom demand. You don't know if this will actually work, but it might be able to happen without anyone else knowing about it.	Cost:100000
□	Call in your cyber insurers to help handle the situation – bear in mind this will cost a premium at some	Hours:20

Digital Banking Platform

Owned	Digital Banking Platform (#1): This is the server that provides the web and mobile banking service for customers.	
Buy	Digital Banking Platform Patch (#2): A patch has been released addressing a number of vulnerabilities in the digital banking platform	Hours:10

Core Banking System

Owned	Core Banking System (#4): This server is responsible for processing all of the transaction requests. It is key to the operation of the bank.	

Sanctions System

Owned	Sanctions System (#6): This system is used for gathering all of the legally required documentation and for identifying potential fruadulant behaviour.	

PCs

Owned	PCs (#9): These are PCs used in the running of the business. The majority of these are laptops that staff can use remotely if necessary.	
Buy	PCs (#10): A patch has been released addressing a number of vulnerabilities on your PCs	Hours:10
Buy	Audit Encryption Status of PCs (#40): Audit devices to ensure encryption is enabled.	Hours:40

CRM

Owned	CRM (#16): This is a third-party tool that gathers and manages all of your customer information.	
Buy	Enable MFA (#27): Enable 2 factor authentication for all users	Hours:2

Payment Handler

Owned	Payment Handler (#17): This is a third-party used for processing card transactions. Without it your customers won't be able to use their bank cards.	

VPN

Owned	VPN (#30): This VPN is used by all remote staff to access work devices.	
Buy	Enable MFA (#29): Enable 2 factor authentication for all users	Hours:10

Firewall

Owned	Firewall (#13): The firewall helps limit the amount of traffic that can gain access to your network from the internet.	
Buy	Firewall Configuration (#14): Putting time into configuring your firewall will significantly improve its effectiveness.	Hours:30

Antivirus

Owned	Antivirus (#18): This is a commercial antivirus that detects known virus and malware	

Endpoint Monitoring

Figure 4: Main exercise interface.

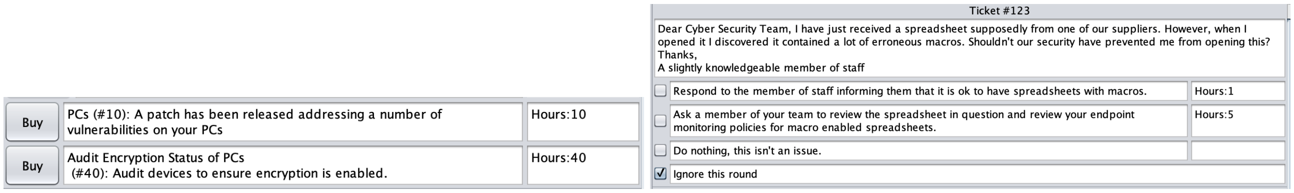

Buy	PCs (#10): A patch has been released addressing a number of vulnerabilities on your PCs	Hours:10
Buy	Audit Encryption Status of PCs (#40): Audit devices to ensure encryption is enabled.	Hours:40

Ticket #123

Dear Cyber Security Team, I have just received a spreadsheet supposedly from one of our suppliers. However, when I opened it I discovered it contained a lot of erroneous macros. Shouldn't our security have prevented me from opening this?
Thanks,
A slightly knowledgeable member of staff

□	Respond to the member of staff informing them that it is ok to have spreadsheets with macros.	Hours:1
□	Ask a member of your team to review the spreadsheet in question and review your endpoint monitoring policies for macro enabled spreadsheets.	Hours:5
□	Do nothing, this isn't an issue.	
☑	Ignore this round	

Figure 5: Examples of an event and an EA.