

Jail, Hero or Drug Lord? Turning a Cyber Security Course Into an 11 Week Choose Your Own Adventure Story

Chothia, Tom; Holdcroft, Sam; Radu, Andreea; Thomas, Richard

License:
Creative Commons: Attribution (CC BY)

Document Version
Publisher's PDF, also known as Version of record

Citation for published version (Harvard):
Chothia, T, Holdcroft, S, Radu, A & Thomas, R 2017, Jail, Hero or Drug Lord? Turning a Cyber Security Course Into an 11 Week Choose Your Own Adventure Story. in *2017 USENIX Workshop on Advances in Security Education (ASE 17)*. USENIX Association, Vancouver, BC, 2017 USENIX Workshop on Advances in Security Education, Vancouver, British Columbia, Canada, 15/08/17.
<<https://www.usenix.org/conference/ase17/workshop-program/presentation/chothia>>

[Link to publication on Research at Birmingham portal](#)

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

Jail, Hero or Drug Lord? Turning a Cyber Security Course Into an 11 Week Choose Your Own Adventure Story

Tom Chothia, Sam Holdcroft, Andreea-Ina Radu and Richard J. Thomas
School of Computer Science, University of Birmingham, UK

Abstract

In this paper we argue that narrative and story are important elements of gamification, and we describe a framework that we have developed which adds a story to an 11 week cyber security course. The students play the part of a new IT security employee at a company and are asked to complete a number of security tasks, for which they receive flags. The students can send the flags they find to a number of different characters to move the story along in different ways. As the story unfolds they find deceit, corruption and ultimately murder, and their choices lead them to one of three different endings. Our framework for running the story and the exercises is completely contained in a single VM, which the students each download at the start of the course. This means that no backend or cloud support is needed. We report on the results of qualitative and quantitative evaluations of the course that provides evidence that the story increased student engagement and results.

1 Introduction

Gamification, in the form of point scoring, competing, team events, freedom to fail, and rapid feedback, has proved popular in cyber security education [LC05, SN13, And15]. However, many of the sources for gamification in education (e.g. [Kap12, She11]) highlight the importance of story, narrative and character development in gamification, but so far, this seems to have been overlooked in any existing approach to cyber security education.

With the aim of increasing student engagement, we have added a fictional story to an 11 week introduction to security course. After completing each exercise, the student will find some flags, which can be sent to one of a number of different characters to move their individual story along in different ways. This gives the students control over what happens next and how the story

evolves. Such an approach is common in large scale computer games in which, after completing a section of the game, the player can make a choice that affects a story that is told using cutscenes. This approach has repeatedly been shown to increase a player's engagement with a computer game, therefore, it is logical to suspect that a similar approach would work for a taught course. We also note that the time players are expected to spend to complete a large computer game is similar to the 100 hours that our students are expected spend on a single course.

Our story framework builds on past work which developed a stand alone capture-the-flag style VM for cyber security education [CN15]. Students each download their own copy of the VM at the start of the course. When they first boot the VM, a set up program runs which creates a unique set of flags for the VM, and then deletes itself. This VM contains a number of vulnerable services and configuration mistakes that support the material taught in the course. For each vulnerability the students discover and exploit, they receive a flag that they then submit to a website for marking. The exercises include misconfigured access control policies, confused deputy attack, a vulnerable website, reverse engineering and buffer overflow attacks. The students do not obtain root on the VM until the end of the course.

Each flag is an AES encryption of a unique VM identifier and a exercise ID, meaning that the flags are all unique but can be checked by the marking website. We note that while this VM is used for an introductory course, advanced students may be able to recover and reverse engineer the start up script, finding out how the flags are generated. We aim to make finding the flags in this manner significantly harder than solving the exercises. A key benefit of our framework is that it does not require any backend cloud support, apart from a lightweight flag submission site. Once the VMs are handed out in the first week of term, course TAs can focus on supporting students and teaching, rather than

technical maintenance. Giving each student unique flags also helps to combat plagiarism.

In this paper, we add a mail server to this VM and a new “story engine”, which runs as a cronjob. This story engine will send e-mails to the students that appear to come from a number of different people. After completing each exercise, the student needs to pick a story character to send the flags they find to. When the story engine sees e-mails from the student, it scans them for flags and looks at the address the student sent them to. Based on this, the engine will proceed the story in different ways by sending new e-mails to the student and updating news stories on the VMs website. The story engine and mail server are internal to the VM, meaning that no backend support is needed for this and each student gets their own version of the story.

The story is told using e-mails and news updates to a website running on the VM. Which e-mails and news stories are sent, and when, is controlled by a single XML file which specifies the story logic. This file maintains the story state, and specifies which e-mails should be sent in response to which events. Our XML uses an expressive logic that allows us to enforce complex conditions, for example, mutual exclusive events and different responses to different orders of actions. This framework also makes it easy to change and update the story without changing the mechanics of the VM.

Our story follows the classic Hollywood story arc of a trigger event, crossing a threshold, overcoming obstacles, a set back and then a final push. Each of these five stages of the story correspond to a two week exercise. The students start out as a security employee at a new company, and are asked by their line manager to complete some decryption code and find some flags, being warned to keep the flags secret. Meanwhile, the students receive a second e-mail from a mysterious stranger, telling them that there is something wrong at the company and they should instead send them the flags. The game engine ensures that the students can only take one of the two choices on a single VM. As the story and exercises move on, they discover that the company is being used as a front for a black market website (The “Cotton Highway”) and, after the police become involved, they must decide what to risk and who to side with.

As a trial, we gave the students the option of following the story or not. A statistically insignificant bonus mark was offered for starting the story and we described how the students could send the first e-mail in lectures. We carried out a survey of the students that did, and did not, follow the story, and this showed that students felt engaged when following the story, and for those who did not do the story, the story concept was good, with positive feedback. We have carried out an evaluation of the difference between the student’s marks, which showed

that students who followed the story achieved significantly higher marks on average than those who did not. To rule out the possibility that stronger students decided to follow the story and weaker students decided to ignore it, we compared the marks our students achieved on this course vs the marks they achieved on other courses. We found that students who followed the story did much better than their marks from other courses would predict, whereas the marks of students that didn’t follow the story were inline with their marks in other courses. This suggests that following the story did increase student engagement.

It is common for CTF competitions to have themes or simple linear stories (for instance PicoCTF has used a story about helping a broken robot to get home with the aim of engaging students [CBB14]). Cyber Security Camps have also used stories to help engage participants, (e.g. Feng [Fen16] reports on a camp with a Divergent theme story line and the company MWR runs HackFu¹, in which professional actors are hired to play characters who will move the story along. Unlike this previous work, our aim is to integrate a story into a full 11-week cyber security course, rather than to use it for a single event. Flushman et al. have developed a full course, based on CTFs and linear alternate reality games [FGP15] which engage students by providing a narrative to exercises. Our aim in this paper is to provide a compelling narrative, following best practice (e.g. [Tro98]), complete with characters which will develop and students will become attached too. Our story is non-linear, meaning that student’s choices effect what happens, and it runs as a self-contained VM that is highly configurable and, once issued to students, does not need any support from course staff.

In the next section, we describe how the mechanics of our story infrastructure works. In Section 3 we give an outline of the story itself and the choices the students can make in it. In Section 4 we provide an evaluation of the students marks and feedback from a survey on the story. We conclude in Section 5.

A website with additional information, our VM and information about how to use and develop the VM can be found here: <http://www.cs.bham.ac.uk/internal/courses/comp-sec/story>.

2 The Story Engine

In order to convey a story to the students, we required an immersive method of communicating with them. We felt it was necessary to make the story dynamic, where it needed to react to the decisions that the students had taken so that they would get a personalised experience.

¹<https://hackfu.mwrinfosecurity.com>

```

<story_map>
  <exercise>
    <event>
      <tasks>
        <email/>
        <news/>
      </tasks>
      <required>
        <!-- combination of conditions -->
        <token/>
      </required>
    </event>
  </exercise>
</story_map>

```

Figure 1: Story XML High-Level Definition

In order to achieve this, we needed to develop a story engine that could not only read the decisions the students had made, but also tailor the emails that were sent, based on these decisions. To improve the immersion, we also added a company website (based on a local web server in each VM). This website features a news section, where the story engine would add new stories to this as the player progressed.

We wanted to make the story easy to create, where no knowledge of Java would be needed to create a new story, or edit an existing one. We decided to use an XML-based configuration file (which we referred to as a ‘story map’). This would represent the different sections of the story and the objectives that would have to be completed to progress further. This has the additional advantage of making the engine self-contained – it does not need to be edited, and can be shipped out as a jar file. In order to further reduce complexity, state fields in the the configuration file are used to track the progress of a player through the story.

In the event that a student decided to abandon the story part way, they would be able to return at any point. The story is self-contained, where exercises handed out to the student are reflected in the story, where the assignment is presented as part of the plot and instructions given to the student, therefore, it is straightforward to map assignments to the engine.

The Story Map Configuration File. The story map served two purposes: first it needed to describe the story to the story engine (so that the correct emails could be sent at the correct times) and, second, it needed to keep track of the progress through the story (so that the story could be resumed after the VM is restarted). To solve the first task, we devised the following format for the story map, defined in Figure 1.

The `<exercise>` tag provides a logical separation be-

tween the different parts of the story. As the course was already written around a series of 5 exercises, we chose to write the story in 5 parts, each matching up to a particular exercise. Inside of each `<exercise>` are several `<event>` tags. These tags hold the actual story content, split up into `<tasks>` and `<required>` conditions for execution.

The `<tasks>` tag contains a list of tasks that the engine might need to perform. These are either `<email/>` tasks – instructing the engine to send an email to the player, or `<news/>` tasks – instructing the engine to update the website with a new story. The exact content of the email and news stories are stored in a separate XML file, in a location the student does not have read access to. The task contains a path to that file which the engine follows to get the content.

The `<required>` tag contains a list of requirements that must be met before the relevant `<tasks>` are executed. These would either be `<token/>` – a specific flag must have been sent to a specific email address, or `<finished/>` – a specified event must have already been completed. We allow for a series of logical operators inside of the `<required>` tag, namely: AND, OR & NOT. These operators can be nested, allowing for arbitrarily complex requirements.

In order to meet the second goal for the configuration file, a `complete` attribute was added to the `<event>`, `<email/>`, `<news/>` and `<token/>` tags. These are initially set to `false`, then updated to `true` as the story engine progresses through the story map.

Reading Decisions. The students move the story along by e-mailing the flags they find when solving exercises to one of the story characters via a mail server running on the VM. The mail client on the VM is preconfigured to use this server.

The story engine logs into each of the email addresses and checks for any emails in the inbox (all of the account names and passwords are stored in the story engine). These emails are then searched for any 32 character long hex strings (all of the flags were 32 hex characters). Any strings found are checked, using the course AES key, to see if they are a valid flag, and if so, which exercise and question they correspond too. After analysing the email, it is deleted from the inbox.

When a correct flag is detected, the story engine searches through the story map file and identifies any requirements that involved this flag. Part of the requirement needed the flag to be sent to a specified user, so we compare the recipient of the flag email to the `sent_to` attribute of the requirement. If this matches, then the requirement’s `complete` attribute is set to `true`.

Sending Emails. The most common reaction to events is for the story engine to send an e-mail, which is done via the VM mail server using the javax mail library. When the story engine reaches an `<email/>` task in the story map it extracts the `path` attribute from this tag. This is a location to an XML file which describes the email containing the body of the email, the subject, the sender and the recipient. The story engine uses this information to create and send an e-mail, which it does by using a given login to connect to the mail server, then altering the sender address.

By taking advantage of the logical operators available in the `<required>` tag, we are able to send an email depending on choices taken. A common pattern for this follows the logic of “send email ‘w’ if flag ‘x’ was received by user ‘y’, but not if event ‘z’ had been completed”.

Posting News Stories. The website uses in one of the final exercises contains a news story section. These news stories are filled in by the story engine as the story progressed. This is achieved by using JavaScript to propagate the website HTML source code with all news stories located within a specified folder. To put a new story on the website the engine moves an XML file (the path of which was provided in the `path` attribute of the `<news/>` tag) into this folder. To make the news stories react to the players actions we use the same technique described in the above section.

As well as dynamic stories, we wrote several static stories. These were not related to the overall narrative, and were meant to give a sense of realism to the website. The method for placing them on the website is the same as before, but instead of being moved based on story progress, the trigger to display these stories on the website depended on the date. This mechanism provides a regular update of news stories, giving the students a reason to keep checking the news site.

Logging E-mails Sent. In order to gather data on the students choices, we log the emails they sent. When the story engine parses an email to check for flags, it also sends a copy of it to a logging server running on an external machine. This logging message includes all of the data from the email, and an identifier unique to each VM instance.

These logged messages may be used as the formal flag submission method, therefore forcing students to take part in the story. However, for this iteration of the course, we wanted to make following the story optional – we had the student submit the flags to a website on our school’s web server.

3 The Story

Progression and storytelling, as identified by Stott and Neustaedter [SN13], are two key concepts of game design which can be successfully applied in learning environments. However, this approach does not appear to be prevalent within the area of cyber security education. With the purpose of increasing student engagement, we have introduced a fictional story to our introduction to computer security course. In order to be able to create an exciting, alluring and believable story we have investigated the three-act structure [Tro98] screenwriting model, commonly used in Hollywood scripts.

Our story follows the classic Hollywood screen writing structure. Each exercise corresponds to a stage within the story arc: a call to adventure, crossing a threshold, overcoming obstacles, a set back and then a final push. This ensures the story arc will keep the students engaged and excited to progress through the story.

The story takes place at the fictional company Sensible Furniture. It is a furniture company, with a dark secret, which the student has to discover along the way. The initial setup has the student starting off as employee **427**, the new cyber security advisor of the company. Below, we introduce each character of the story, then explain the plot arc and possible paths through the narrative.

When developing the story plot, one of the considerations when writing the content was the target audience and whether events in the story would be appropriate. We therefore applied the BBFC (<http://www.bbfc.co.uk/>) and ESRB (www.esrb.org/) ratings scheme against the proposed story. Under these rating schemes, the story written would be given a ‘12’ or ‘Teen’ rating respectively, as it has mild references to drugs and violence, and moderate threat, making it appropriate for undergraduate students. That said, the story can be changed to suit alternative audiences, for example younger teens and children.

3.1 Characters

The story introduces employee **427**, the main character, and five other characters. They evolve as the story progresses and their roles are to sway employee 427 towards their goals.

Employee 427. The student plays this character. We chose a number, instead of a name, in order to (a) ensure each student can identify with the character and (b) introduce a sense of impersonal coldness within the environment the story takes place in.

Jak Kinkade. Jak is the CEO of Sensible Furniture, Inc. He introduces employee 427 to their new working environment, and reappears towards the end of the story, in order to give 427 a last push towards a dangerous path.

Nik Adler. Nik is employee 427’s Line Manager at Sensible Furniture, Inc. His main role is giving employee 427 their daily tasks, and keeping them on track.

Charle Garcia, aka Chimp. Charle is also an employee of Sensible Furniture, Inc. At the start of the story, Charle goes

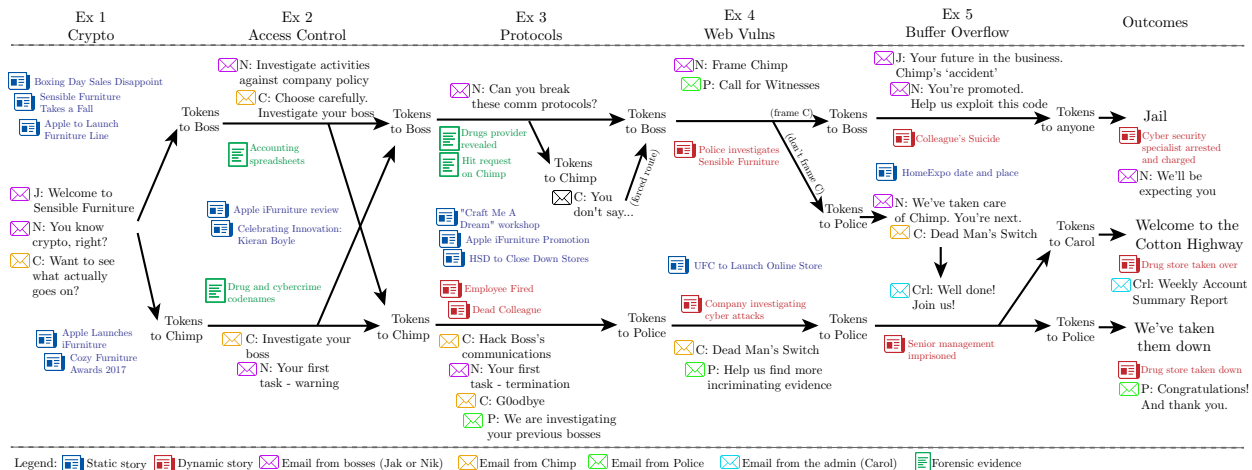


Figure 2: Story flow, as it progresses based on timeline (left to right).

by the alias Chimp and approaches employee 427 in order to convince them to join their side, hinting that the bosses are suspicious characters. Chimp’s character represents the outcast, the undercover potential ally who will guide 427 on the path of good.

Thomson Gazal. Thomson is a police officer, they come into the story when the main character appears to be going through a crisis and options seem to be limited. Thomson represents a new opportunity, a door being open for the possibility that the main character will take the path of righteousness.

Carol Miller. Carol is the IT administrator of Sensible Furniture, Inc by day, and a questionable character by night. She appears in the story towards the end, in order to give it a new twist. Carol offers 427 a different perspective on the events that have unfolded and a new opportunity: a role in the underworld.

3.2 Plot Arc

The overview of the story flow can be seen in Figure 2. This shows the mapping between the exercises and the story progress. It contains all information which is sent or exchanged with the student, and all characters involved. The arrows show the possible paths the main character can follow, together with the outcomes. The figure also shows which e-mails will be sent, when, and what news stories will be posted at what points.

The student plays the part of the employee 427, a new cyber security advisor. They are welcomed by Jak Kinkade, the CEO, and introduced to Nik Adler, their line manager. Nik provides 427 with their first task – demonstrating an understanding of encryption methods by writing code to decrypt a range of files, and asks 427 to submit the resulting plaintexts (flags) to them. 427 also receives an email from a person named Chimp (see Appendix B.1), which suggests the manager is suspicious, in “cahoots with the underworld”. Chimp requests that 427 sends him the tokens instead, and collaborates to take down the managers. This opens two paths: they can either choose to be a good employee or to trust Chimp.

The second task is to investigate others’ home directories. This matches the plot point, in which the students need to learn

more about the company’s other employees. This requires the student to demonstrate an understanding of how access control works on Linux and common attacks against this. The contents of the directories are provided with some clues that there is illegal activity within the company, but no indication of who is involved is given. This makes the students question the paths they have chosen. They can find accounting spreadsheets in Nik’s directory with some suspicious entries, code names, for drugs and cybercrime. In Chimp’s directory, some files containing information about drugs and cybercrime transactions can be found. At this point in the story, 427 can change paths depending on who the player sends the flags to.

The third task starts sealing in 427’s fate. To continue the story the student needs to intercept secret messages been passed out of the company, therefore this exercise involves breaking some encrypted communication protocols (supporting the teaching of secure protocols on the course) and find out what is going on. The communications will reveal who the drugs provider is and, concerningly, a request for a hit on Chimp. If 427 is collaborating with Nik, they cannot, at this point, switch sides. If they try to email Chimp, he will bitterly refuse and tell them to carry on down the route. If 427 is collaborating with Chimp, a setback appears. They are fired, and they receive an email from Chimp that he has been found out. Chimp redirects 427 to Police Constable Thomson Gazal. The company website will display two pieces of news: employee 427 has been fired, and that Charle Garcia has committed “suicide”.

The fourth task thickens the plot on both story paths. On the route where 427 works together with Nik, they receive a surprising request: investigate, and attack, the company website (supporting the teaching of web security on the course) and if illegal activity is found, frame it on Chimp (as Chimp is alive on this path). At the same time, PC Thomson Gazal asks 427 to become a confidential informant. The situation seems very grim on this path and Thomson represents the chance of redemption. The student can choose to continue down what seems like a dark road and frame Chimp, or turn to the police. On the other path, 427 receives a ‘dead man’s switch’ email, set by Chimp to

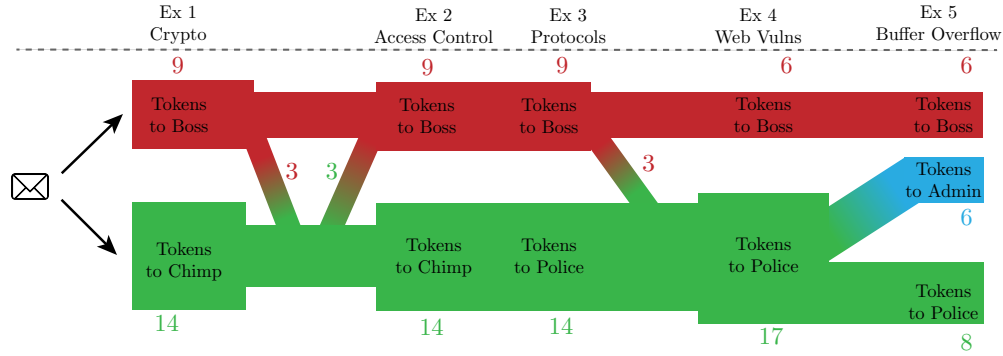


Figure 3: Paths and number of students flow.

trigger in case something happens to him. The email contains information which incriminates Nik and Jak as drug lords and hitmen. At this point in the story, only two endings for the story are foreseeable: either continue working with Nik and Jak, and turn to illegal business, or cooperate with PC Thomson Gazal and hope to be exonerated.

The final task is introduced by the company IT admin Carol, who has been a background character until this point. Carol reveals that she has been running the company’s hidden black market site, and gives 427 one last choice: join the underworld team, replacing Nik and Jak, and take over as drug lord. To do this 427 must get root on the VM by reverse engineering binaries using the IDA tool, and writing code for buffer overflow attacks.

With the police closing in on Nik and Jak, Carol’s offer represents the last twist of the story. The option to collaborate with Carol is given to those who have chosen to collaborate with PC Thomson Gazal in the previous task. The news of Charles’ “accident” appears on the company website. However, it is revealed that the “accident” was setup by Nik and Jak.

If 427 has stuck with the bosses, at this point they have no escape and, regardless of who they chose to turn to, the outcome is that they are sent to **jail** along with Nik and Jak. For those that sided with Chimp and turned the last set of flags over to the police, a happier ending is in sight: the management is imprisoned, and the drug store is taken down. PC Thomson Gazal also reveals Chimp was their fiancée, and thanks 427 for all their help in catching the murderers. On the path of collaborating with Carol, if 427 has previously betrayed the bosses, they receive a threatening email from Nik, letting them know Chimp has been “taken care of”, and that 427 is the next target on the list. Working with Carol means 427 takes over the dark net website **Cotton Highway**, becoming the new drug lord.

3.3 Student choices

Figure 3 shows a ‘roadmap’ diagram of the choices students have made throughout the story. The roadmap is based on 23 story maps (the XML files) provided by the students who have chosen to follow the story. The paths are colour-coded as follows: **red** paths mean the students were collaborating with Nik, **green** paths mean the students were working together with

Chimp and the Police, and finally, **blue** paths represent collaboration with Carol.

We can see that, at the start, 9 (40%) students chose the path of the bosses, and 14 (60%) chose to trust Chimp. However, for the second exercise, where they are given some clues that there is illegal activity within the company, 3 students from each path chose to switch sides. This behaviour corresponded with our expectations, where we would expect students to not follow a specific path, rather doubting their choices, or switching paths to see what the happened. Reviewing the paths students took, our expectation was confirmed, showing that these students doubted their choices. Numbers remain steady for the third exercise. The fourth exercise sees 3 more students shifting from the bosses path, in order to collaborate with the police. This corresponds to students refusing to frame Chimp for the illegal activity found on the company website. At the end of this exercise we see only 6 (26%) students still committed to their bosses, whereas 17 (74%) of them are working with PC Thomson Gazal in order to find incriminating evidence on Nik and Jak. The appearance of Carol within the story seems to divide the students on the green path, with 8 choosing to continue their work, and 6 choosing to become a drug lord. We also note that 3 students did not progress the story from exercise 4 to the final one. The even breakdown of endings shows that students had genuine choices, validating the writing and design quality of our story.

4 Evaluation

In this section, we carry out an analysis of the effect of the choice-based exercises on student engagement and the impact on student marks. As mentioned earlier, we made the story optional to follow. This was done so that we could compare the effect of the story, and, because we did not want to make a new experimental system compulsory for a large class. A self-selected sample of 23 out of 144 students chose to follow the story, and, while less than we would have liked, proves enough data for an interesting analysis. As the story was not compulsory for this year, it is possible students may have been less inclined to try it. We did not promote the story regularly, which may have also reduced the number of participants, which can be addressed in the next iteration of the course. Furthermore,

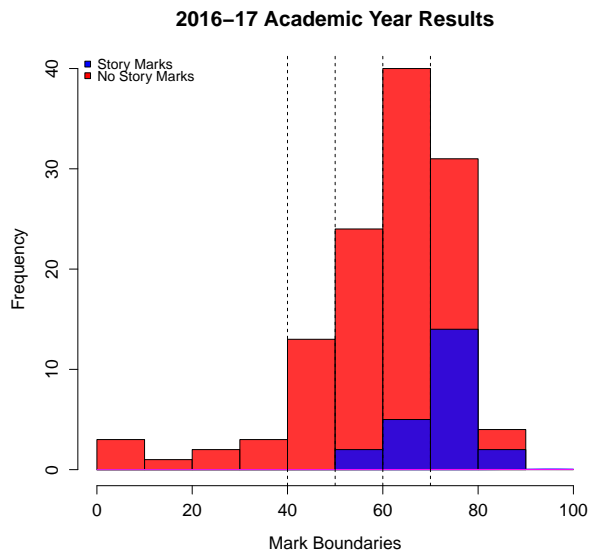


Figure 4: Histogram of Student Marks

as the story was optional for the students, those who followed the story were not at an advantage to their peers, as no further information, including hints, were given in the story.

Analysis of Marks vs. Story Engagement. We first compare the final marks awarded to the students that did and did not follow the story. Table 1 shows a summary of the final marks for this course, and a histogram of these marks is given in Figure 4.

students	course avg.	avg. story	avg. no story
144	63.04	72.35	61.27

Table 1: Breakdown of 2016-17 Academic Year Results

This table shows a strong correlation between students achieving high marks and those who did the story. The histogram also shows that all students who did the story achieved a minimum mark of 50%, with the majority achieving a mark between 70-80%. When analysing the marks, we also noted that, of the top 20 students in the course based on coursework (CA) marks, 15 students were involved with the story.

Analysis of Student Engagement in the Course. To look for evidence of student engagement, rather than just high marks, we analysed reports written by the students as part of a reading week assignment, and the emails that were captured from the virtual machines when they progressed in the story.

In the case of the reading week reports, we found that for students who completed the story, their reports contained, on average, 74% more words than the reports submitted by their peers who did not complete the story. The marks allocated to this assignment, were simply for completion of the report. Therefore, we believe that writing more for this report shows greater engagement in the course.




Figure 5: Student Marks vs. their Previous Academic Year

As part of telemetry data sent by the Virtual Machines used in the course, any emails sent by the students as part of the story were captured and forwarded to a logging server for further analysis. In the general case, the students submitted the flags to the website prior to progressing the story via email. However, seven students were more involved with the story than others. These students interacted in conversation with the story engine, as shown in Appendix A. This suggests that the story was believable and engaged the students in the course, making them spend more time on course-related activities.

Comparison of Marks with Other Courses. The data points could be explained by the best students being the most involved with all aspects of the course. To control for the students ability we looked at their marks from other courses taken as part of their degree in the previous year. We assessed the difference in the course mark given for this course against the student year averages, removing students who had either dropped out of the course, or had not completed any of the assignments.

Students that did not do the story averaged 5% higher in the CA and 10.6% lower in the course marks, compared to their previous academic year exam marks (as our course marks tend to be lower than first year marks, this difference is unsurprising). In comparison, the students following the story improved on average, in their CA by 18%, and had only 5.1% decrease in course marks, clearly showing that the students that followed the story performed better than those who did not.

In Figure 5, marks for our course are plotted against average exam marks in the last academic year. Generally, performance in the course is clustered around 63% for both groups of students, with students who did the story clustering around 72%. The averages for each group are shown with '+', where the average mark for the previous academic year was 71.88% for those who did not do the story, compared to 77.43% for those who did the story. When testing if the differences between the

two groups, with respect to their course marks, are significant, we report a p -value < 0.05 , and we accept the alternative hypothesis that there is a significant difference between the results of the students who did do the story, compared to those who did not do the story, with the latter showing better results.

We see that some students who did not follow the story did very well in their exams for other courses but were not among the top students for our course, suggesting a lack of engagement. We also see some students that did follow the story, and scored lower marks in their exams for other courses, but scored very highly in this course. While this does not prove a definitive causal link, this provides evidence that the story did increase student engagement and attainment.

Lastly, in order to account for individual capability, we computed the differences between last year's results and our course result for each student. We excluded from the dataset 7 students who did not have previous marks, and two students who did not take the module exam or do any CA (outliers). Then, we tested if there are differences between the two groups. We observe the mean of the group of students who did not do the story is -9.22 , which means on average, they obtained a mark lower than their last year's average by 9.22 marks. In contrast, students who did the story performed better, having an average difference of only -5.09 . We report a p -value < 0.05 , and conclude that the differences in means are not likely to be a result of chance, and that engaging with the story did in fact improve student performance.

Survey Results. All students on the course were asked to take part in an online, post-course survey, whether they had tried the story or not. From this survey, we had excellent feedback on how fun the story was, how engaging it was and how much students enjoyed it. We received 52 survey responses. Those who did the story gave an average score of 5.5 out of a possible 10 when asked if the story increased student engagement in the course, with 6.1% of students agreeing it increased their level of engagement a significant amount. When asked if the story made the course fun, the average score was 7/10 with 12.1% rating the course as extremely fun. That said, a score of 5.75/10 was given for engagement, with 6.1% agreeing the story was extremely engaging, where they wanted to find out how the story evolved.

For students who started the story, we also asked how many saw the story to completion, with 38.3% of the students who started the story seeing it through to the end. For the majority of students, it took 1 hour in total to complete the story and it was, predominantly, the curiosity on how the story progressed that made them continue it. Those students who did not start the story had different views on why they did not take part in the story, for example other priorities in other courses. For those who did not do the story, there was a mixed response in the anticipated time required to complete the story, with most suggesting 2 hours to more than 3 hours. There is a clear difference in the time it actually took for the students to complete the story to those who abandoned it and estimated how long it would take to complete it. This suggests a possible misconception on how long it would take. If we had informed the students how long we expected it to take to complete the story, it is possible more students would have participated in it. Some students, however

were unable to complete the story, as they had not finished the first exercise. The issues identified can be addressed in future iterations of the course, with information being provided earlier and being more clear about how to participate in the story.

Overall, the consensus of the students towards the concept of the story was very positive, with 96.88% of students who did the story, and 84.21% of students who did not follow the story agreeing that the story was a good idea, with comments suggesting it would make the course more interesting and would increase engagement, with an alternative reward to what students typically would expect. When asked if it would be a good idea to add a story framework in other courses taught in the department, 62.5% of students following the story, and 58.9% who did not follow the story agreed, suggesting a list of courses which could use the framework.

When asked about the story itself and how happy the students were with the conclusion, the students scored an average of 3.28/5, with 16.7% being extremely happy. Some students commented on its realism and how it developed. A score of 5.85/10 was given by students on the believability of the plot, with 3% reporting it as extremely believable. On how interactive and captivating the story was, scores of 6.16 and 6.31/10 respectively were given, with 12.5% of students agreeing the story was extremely captivating, and were really keen to see how the story progressed. When asked about quality of writing, a average score of 7.49/10 was given, with 18.2% marking as very well written. These factors may have contributed to the level of engagement in the story and the course. Example student feedback is given in Appendix D.

From the results of the student survey, we can now conclude that through the captivation and curiosity of the students engaging in the story, it may have been the story which encouraged them to develop and, therefore, finish the course with higher CA marks. For those who did not follow the story, the students still thought the story was a good idea. If the time requirements and what had to be done to participate in the story were made clear, we would likely have seen more engagement and involvement with the story. The survey shows positive feedback and sheds light on why not many completed the story, and can be easily addressed. The survey gives confidence that the story can be made compulsory in the next iteration of the course.

5 Conclusion

In this paper we have presented a method of adding a story to an 11 week cyber security course. Our framework is highly parameterisable making it easy to update and change the story from term to term. Running the story with a group of students, we have shown that students who followed the story did better on the course than would have been expected from their averages in other courses, whereas those that did not do the story did not do better, therefore presenting evidence that this has successfully increased student engagement. As future work, the addition of social media communication or pre-recorded videos may added to the story telling.

Acknowledgement This work was supported by the grant GEN1214 from "The Higher Education Academy".

References

- [And15] Andreea-Ina Radu and Sam L. Thomas. Organising monkeys or how to run a hacking club. In *Workshop on Cybersecurity Training & Education (VIBRANT15)*, 2015.
- [CBB14] Peter Chapman, Jonathan Burket, and David Brumley. Picoctf: A game-based computer security competition for high school students. In *2014 USENIX Summit on Gaming, Games, and Gamification in Security Education (3GSE 14)*, San Diego, CA, 2014. USENIX Association.
- [CN15] Tom Chothia and Chris Novakovic. An offline capture the flag-style virtual machine and an assessment of its value for cybersecurity education. In *2015 USENIX Summit on Gaming, Games, and Gamification in Security Education (3GSE 15)*. USENIX Association, 2015.
- [Fen16] Wu-chang Feng. A "divergent"-themed ctf and urban race for introducing security and cryptography. In *2016 USENIX Workshop on Advances in Security Education (ASE 16)*, Austin, TX, 2016. USENIX Association.
- [FGP15] Tanya Flushman, Mark Gondree, and Zachary N. J. Peterson. This is not a game: Early observations on using alternate reality games for teaching security concepts to first-year undergraduates. In *8th Workshop on Cyber Security Experimentation and Test (CSET 15)*, Washington, D.C., 2015. USENIX Association.
- [Kap12] Karl M. Kapp. *The Gamification of Learning and Instruction: Game-based Methods and Strategies for Training and Education*. Pfeiffer & Company, 1st edition, 2012.
- [LC05] Patricia Y. Logan and Allen Clarkson. Teaching students to hack: Curriculum issues in information security. *SIGCSE Bull.*, 37(1):157–161, February 2005.
- [She11] Lee Sheldon. *The Multiplayer Classroom: Designing Coursework As a Game*. Course Technology Press, Boston, MA, United States, 1st edition, 2011.
- [SN13] Andrew Stott and Carman Neustaedter. Analysis of gamification in education. *Surrey, BC, Canada*, 8, 2013.
- [Tro98] David Trottier. *The screenwriter's bible: A complete guide to writing, formatting, and selling your script*. Silman-James Press, 1998.

C Example Exercise XML Definition

```
<exercise id="exercise_2">
  <event id="first_boss_token" complete="true"
    >
    <tasks>
      <email complete="true" path="/root/
        Stories/Emails/Ex2/Start_Boss_Path.
        xml" />
      <email complete="true" path="/root/
        Stories/Emails/Ex2/
        ChimpStart_Boss_Path.xml" />
    </tasks>
    <required>
      <AND>
        <token complete="true" sent_to="
          nikadler@sensiblefurniture.com"
          plaintext="Ex1" />
        <NOT>
          <finished id="first_chimp_token" />
        </NOT>
      </AND>
    </required>
  </event>
  <event id="first_chimp_token" complete="
    false">
    <tasks>
      <email complete="false" path="/root/
        Stories/Emails/Ex2/
        ChimpStart_Chimp_Path.xml" />
      <email complete="false" path="/root/
        Stories/Emails/Ex2/
        BossStart_Chimp_Path.xml" />
    </tasks>
    <required>
      <AND>
        <token complete="false" sent_to="
          monkey.see482@imeverywhere.
          sensiblefurniture.com" plaintext="
          Ex1" />
        <NOT>
          <finished id="first_boss_token" />
        </NOT>
      </AND>
    </required>
  </event>
</exercise>
```

D Survey Feedback

I loved the story but it seemed to finish abruptly, and it wasn't long enough!

More emails would have been nice too, as we only got to interact with the story five times (one for each exercise).

I did like the complicated underground manoeuvres of the Sensible Furniture crowd. The Charles Garcia reveal and discovering the message that led to his demise was also a big moment in the tale.

RIP Chimp, may he never be forgotten

Was genuinely upset when Chimp died. RIP.

Loved it. Great idea from start to finish!

I found the story more enjoyable after finishing all of the exercises because then the story could be retried and different endings could be found.

I liked the opportunity to choose a path, but also be able to change at certain points. Felt involved with the characters and had a fitting ending.

The bad guys got what they deserved! Justice yay!

Interesting to see how the story developed from certain situations.

I didn't lose (end up in jail), the taste of victory is sweet.

It's fun and enjoyable and definitely sets the exercises apart from other courses

Gives context to the exercises, bit of fun to make people want to do them.

It gives the exercises meaning, rather than doing them for the sake of doing them

Engages students to pay attention to the exercises, gives them a little real-life context (which often aids understanding) and instils confidence in students that the course is being very well-managed.