

# A Case Study of Phishing Incident Response in an Educational Organisation

KHOLOUD ALTHOBAITI, University of Edinburgh, United Kingdom, Taif University, Saudi Arabia  
ADAM JENKINS, University of Edinburgh, United Kingdom  
KAMI VANIEA, University of Edinburgh, United Kingdom

Malicious communications aimed at tricking employees are a serious threat for organisations, necessitating the creation of procedures and policies for how to quickly respond to ongoing attacks. While automated measures provide some protection, they cannot completely protect an organisation. In this case study, we use interviews and observations to explore the processes staff at a large University use when handling reports of malicious communication, including how the help desk processes reports, who they escalate them to, and how teams who manage protections like the firewalls and mail relays use reports to improve defences. We found that the process and work patterns are a distributed cognitive process requiring multiple distinct teams with narrow system access, and tactic knowledge. Sudden large campaigns were found to overwhelm the help desk with reports, greatly impacting staff's workflow and hindering effective application of mitigation's and the potential for learning. We detail potential improvements to the current ticketing system, and reflect on ITIL, the framework of best practices that informed the full process.

CCS Concepts: • **Security and privacy** → *Phishing; Social aspects of security and privacy.*

Additional Key Words and Phrases: Phishing Management, organisation trade-offs, phishing, Distributed Cognition, ITIL

## ACM Reference Format:

Kholoud Althobaiti, Adam Jenkins, and Kami Vaniea. 2021. A Case Study of Phishing Incident Response in an Educational Organisation. *J. ACM* 37, 4, Article 111 (August 2021), 32 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

## 1 Introduction

Keeping organisations secure requires effective procedures to handle reports of fraudulent emails aimed at deceiving employees into giving away valuable information. Such attacks are known as *phishing* and are often used to gain access to accounts and other information that is then used in more damaging attacks [97]. Protecting employees from such attacks is a key component of most large organisations' security plans, often including training employees on how to identify and report phishing as well as putting in place internal procedures to quickly respond to phishing reports.

Phishing is by far the most common and disruptive type of attack for UK organisations [22, 98] which can be partially seen in the amount of effort they put into managing it. Worldwide, 93% of organisations measure what phishing is costing them in terms of downtime, monetary losses,

Authors' addresses: Kholoud Althobaiti, University of Edinburgh, Edinburgh, United Kingdom, Taif University, Taif, Saudi Arabia, [k.althobaiti@sms.ed.ac.uk](mailto:k.althobaiti@sms.ed.ac.uk); Adam Jenkins, University of Edinburgh, Edinburgh, United Kingdom; Kami Vaniea, University of Edinburgh, Edinburgh, United Kingdom, [kvaniea@inf.ed.ac.uk](mailto:kvaniea@inf.ed.ac.uk).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

© 2018 Association for Computing Machinery.

0004-5411/2021/8-ART111 \$15.00

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

50 reputational damage, and time spent by their IT teams on remediation [79]. Organisations also  
51 spent considerable resource training end-users to identify phishing attacks, with 98% allocating  
52 more than 30 minutes to train each user and 61% training users at least every month [79]. The  
53 effort is showing results, with 63% of highly disruptive breaches being reported by staff as opposed  
54 to being found through automated monitoring [22]. The impact of on-the-job phishing training  
55 can even be seen in general user surveys where users report learning protection practices at work  
56 or learning from others who had training at work [77].

57 How organisations make use of phishing reports to update their defences is not currently well  
58 studied. Having a procedure in place to receive and act on phishing reports is an accepted industry  
59 best practice recommended by many authorities [7, 36, 97]. But very little is known about the  
60 practicalities and workflows staff use to accomplish these goals or how practical the various  
61 approaches are. What we do know is that effective phishing attacks can cause damage within  
62 minutes of making it to an inbox [97], so quickly responding to phishing reports is essential as it  
63 allows for rapid response and mitigation, which in turn limits the damage.

64 In this case study, we focus on answering how a large University in the UK handles phishing  
65 reporting and mitigation. Universities are an interesting organisation to study for several reasons.  
66 First, the education sector has the highest phishing clicking rate even compared to sectors such  
67 as finance and healthcare [97]. So Universities are a prime target for attack, and those attacks are  
68 currently successful. Second, Universities can be quite large, but their support staff are not funded  
69 or organised the same way a large financial institution might be. Third, the yearly turn over of  
70 students makes typical approaches such as training more challenging to do. Finally, Universities have  
71 valuable resources to protect. The most obvious examples are the personal details of students and  
72 staff, the content of ongoing research projects, and various intellectual properties [96]. Universities  
73 are also expected to protect access to contracted services such as the JSTOR digital library, with  
74 whom they have contracts promising to limit access to current staff and students.

75 Our investigation of the University began with shadowing staff working at the central Help  
76 Desk where phishing reports come in. We complemented these observations through contextual  
77 interviews with teams across the University that handle phishing-related issues.

78 We found that Help Desk staff become inundated by reports when large attacks occur and  
79 must balance their workload by prioritising security risks against potential impacts. For example,  
80 compromised staff or student accounts are potentially of greater threat than a single errant phishing  
81 email. Awareness of phishing incidents primarily come from end users' reports, however are  
82 not limited to this single source with alternative internal and external sources, requiring more  
83 sophisticated coordination and well-choreographed hand-offs between teams. This collaboration  
84 is seen as a distributed cognition across the teams which is essential for an effective response,  
85 enabling quick updates to automatic protections. However, lapses in current practice prevent teams  
86 from fully reflecting on incidents though practice flexibility also results in catching unexpected  
87 issues. End users who report phishing are typically given generic feedback which may not match  
88 their exact query. We also see that mitigation attempts can be hampered by the mitigating team  
89 having incomplete information, such as only looking at a single phishing example from a campaign,  
90 and therefore, missing subtle variations used by the attackers.

91 This case study is intended to form the basis of further studies into phishing report management  
92 processes. We believe that further research in this area will guide organisations to better inform  
93 their phishing incident plans. This case study highlights several challenges in handling phishing  
94 reports and the problems stakeholders face when managing phishing attacks. We also recommend  
95 that future research focus on augmenting phishing reporting systems using automation to help  
96 minimise staff time requirements while also making full use of every phishing report.

## 2 Phishing Protection Life Cycle

Here we explain how organisation phishing protections work in theory. The following description is similar to a high-level textbook description or the common narrative used at industry conferences. It also represents the goal state for many organisations, including the studied University.

Phishing management incorporates both proactive and reactive defence elements, which feed into each other. Proactive protections include activities like setting up firewalls, setting up phishing/spam detectors on incoming email, putting solutions in place to monitor and identify abnormal activity, and training end-users [75]. However, phishers are continuously trying to find new ways to bypass proactive protections, so organisations also need to have reactive processes to monitor for new attacks. These reactive approaches include monitoring automated alarms, providing avenues for users to report suspected phishing, and procedures for responding to identified attacks. Attacks identified via reactive methods are then mitigated as quickly as possible and used to update and improve proactive measures.

A typical example of reactive processes that feed back into proactive protections is *phishing campaigns*, where a phisher sends a large number of emails at an organisation, often with minor variations to each one to make them harder to detect and filter out. A phisher starts by sending one or more communications – most commonly emails – to organisation email addresses. If the organisation’s proactive measures are effective, the communications will be caught by a firewall or a mail relay filter and never make it to any inboxes. The Email filters are normally made up of a combination of heuristic rules and lists of known malicious patterns (blacklists). If the communication is not automatically identified as phishing, it will appear in organisation members’ inboxes.

If a phishing email bypasses all the automatic filters, users will see it and interact with it in different ways. Some users will open the emails and click on links (11%) or attempt to give away data (4%) [78]. If the proactive solutions are effective, users will be blocked from visiting the page by a firewall, their web browser, or blocked from sharing private data by internal software. If user training has been successful, some users will also identify phishing and report it, typically to an IT help desk or some other type of security operations centre. Hence, we began our research by studying a help desk.

Once a report is received, the help desk forwards it to one or more internal teams who use the content of the reported phishing email to immediately apply mitigation, such as adjusting and applying filters to mail inboxes to remove the offending phish and prevent any further interaction from users. The teams may also use log files to determine which, if any, users have already clicked on the malicious links and force a password reset. They also use the content to update proactive protections, such as updating the mail relay filter to identify and prevent similar future emails from breaching the defences again.

What should be clear from the above paragraphs is that phishing is a problem that impacts the organisation at all levels, with many teams needing to be involved in handling proactive and reactive phishing management. In this work, we focus on the phishing management’s reactive processes to understand the day-to-day process of handling phishing incidents.

## 3 Related Work

To situate our research, we consider relevant prior research that explores security incidents management and research that specifically focused on identifying, assessing, and responding to phishing incidents. We supplement this with previous work on general IT management and best practices.

### 3.1 Phishing as a security incident

The cybersecurity landscape can change rapidly as new vulnerabilities and mitigations arise, resulting in distinct working practices from other general IT incidents [28]. Security incident management, therefore, involves reporting, assessment, response, and learning from incidents to improve existing security practices [59]. Phishing incidents are not unique in this regard, as response procedures are developed and refined through reflective evaluation and feedback to adapt related procedures and best practice [25, 26, 30, 65] including the development of practices that consider the range of potential attack vectors, such as Social Media [105].

*3.1.1 Discovery and reporting.* Security incident management begins with the initial discovery of ongoing security issues, which are identified through internal (e.g. security monitoring mechanisms and employee reports) [27, 50, 63] and external sources (e.g. bug bounties or from another organisation) [50, 61, 63]. However, it can be challenging for these external groups to identify who to report to [50]. Internal security monitoring tools, although helpful at detecting security incidents [63], can overload IT staff with a large influx of reports arriving simultaneously for a single incident [103] or from false alarms [82]. These numerous generated reports take valuable staff resources because they must be sorted through to determine their accuracy and priority before action can be taken [82, 94].

Phishing is designed to trick both email servers and end-users into thinking that they are looking at a legitimate email, making discovery and accurate validation challenging. Security managers recommend concentrating efforts on the identification of phish through employee training and preventive technical measures [51]. Automatic preventive measures are designed with the goal of reducing the number of successful attacks by passively filtering incoming emails [16, 24, 33, 48, 89] and authenticating senders to block attempted forged emails [37].

While valuable, these solutions are not perfectly accurate, nor do they protect against more sophisticated phishing attacks, such as *lateral phishing* where the attack originates from a compromised account that is already verified and trusted [32]. Human awareness is therefore seen as an integral component within organisations' security management strategies [26], as it adds a layer of security while complementing the technological controls (e.g. filters). End-user training is considered a top priority within the industry, regularly being included in annual budgets for organisations [76], and additionally an area of focus for the academic community [23, 37]. Common approaches taken to staff training include web-based training materials [66], training courses [41], embedded training (simulated phishing attacks) [3, 13, 30, 43, 53, 54], cartoon-based training [53, 88], game-based training [5, 12, 85], real-time support [4, 99, 106], and phishing stories [100]. Despite the benefits of training it has been found to not be fully effective with between 15.57% [86] and 69%[47] of people still falling for simulated phishing tests [47, 65, 86], and employees and users are still expected to report phishing [36, 66].

Timely responses greatly aid organisations to react to and mitigate attacks, reducing the number of potential victims who would engage with the phishing communication and therefore reducing or avoiding potential organisational damage [51, 97, 103]. However, the number of phishing reports is still considered too low [97] as people usually only report phishing when they doubt its safety and need an expert's opinion [10, 57], know the spoofed sender, have a desire to protect other potential victims, or perceive the email to be particularly convincing and therefore dangerous [10]. Counter to this, research has found that users may not report due to a lack of awareness of legitimate reporting channels, concerns of mishandling, and perceived self efficacy [56]. To promote phishing reporting, prior work has looked at using staff feedback after training simulation to modify policies to better align with staff needs [47]. Phishing reporting can also be treated as a problem of Communal Knowledge Management, where employees report potential phishing to a website

197 that can be publicly accessed [45]. Praising a legitimate reporter and sharing with other staff has  
198 been investigated and shown to have a positive impact on encouraging staff phishing reports for a  
199 number of scenarios [30].

200  
201 *3.1.2 Incident assessment.* Assessing security incidents is non-trivial as it requires trained staff  
202 to review, communicate, and identify the causes in order to determine the relative importance of  
203 issues [103]. The general public struggles to accurately identify phishing resulting in numerous  
204 false-positive reports [38, 45] which must be verified and validated [50] for accuracy as well as  
205 determine how critical they are based on potential impacts. For example, the impact can be judged  
206 based on the number of affected users, the affected services, or the type of users affected [2, 63].  
207 Automatically prioritising sophisticated phishing emails helps the Incident Response Teams act on  
208 phishing that is more likely to harvest clicks [94].

209  
210 *3.1.3 Responding to incidents.* Responding to security incidents often involves coordinating with  
211 staff from many areas of an organisation, and with differing expertise [8, 82, 103]. Organisations are  
212 known to use established policies and procedures to help staff follow best practice when responding,  
213 which involves investigating the cause, escalating to the required teams while simultaneously  
214 documenting all actions taken [2, 63]. While these protocols are indeed impactful, they do not  
215 necessarily match the actions staff take when handling an incident. Staff responses can be influenced  
216 by their attitudes towards the applicable security policy and their interpretation of the policy within  
217 their working context [21]. Additionally, the usability of prescribed forensic tools can impact their  
218 abilities to follow the best practices [40, 63].

219 While there is minimal research detailing actions around handling phishing reports by experts,  
220 automated responses to phishing incidents have proved challenging. It directly depends on the  
221 accuracy of the initial report and; therefore, requires expert human validation [38]. While complex,  
222 looking at how to backtrack to the origin of a phishing attack and analyse it can help investigating  
223 social engineering crime [58].

224  
225 *3.1.4 Learning from incidents.* Learning from past security incidents is also challenging for or-  
226 ganisations, with some having no formal approach to gathering lessons or redistributing those  
227 lessons to staff [31, 84], which itself may take considerable time [55]. Organisation risk focusing  
228 solely on solutions to incidents without reviewing larger policies or organisational structures [2].  
229 Additionally, reviews may be biased when the focus is placed on rarer large-scale or severe incidents  
230 which obscures potential day-to-day lessons [2, 21] and results in overcompensation with security  
231 taking an overbearing role in incident management [21].

232 Best practices for learning primarily focus on the technical aspects and direct cause of inci-  
233 dents [93]; however, security policies may also be causes for learning. For example, Sasse and  
234 Brostoff [80] investigated the large number of incidents raised to an industrial organisation's help  
235 desk regarding password resets. The research was motivated by the high costs the organisation  
236 was incurring in help desk staff time. It concluded that modifying the current unusable password  
237 policy had the potential to reduce help desk staff time by 40%.

238 For phishing incidents, little is known about how to learn from them; however, qualitative and  
239 quantitative metrics are used to evaluate incident handling performance, with the latter being more  
240 dominant, such as response time, the number of tickets, and the number of incidents [51].

## 241 3.2 IT incident management

242 Incident management is a well-established space in Information Technology (IT) where service  
243 quality, users satisfaction, and system stability are examples of essential organisation requirements.  
244  
245



246 Many of the processes and policies implemented by organisations will be influenced by their  
247 chosen IT governance standard or frameworks such as Control of Business Objectives and Tech-  
248 nology (COBIT) [17] and Information Technology Infrastructure Library (ITIL) [42]. Frameworks  
249 similar to ITIL dictate the structure of their IT department, guiding how to organise teams, as well  
250 as how to handle communication and coordination between teams [72]. ITIL has been found to  
251 improve overall service quality [73], customer satisfaction [73], speed of incidents' responses [68],  
252 and the number of necessary escalations [68].

253 However, implementing such frameworks can be challenging [44] because of the natural tension  
254 between theory and practice. For example, frameworks rightly advocate solving the root problems of  
255 identified technical issues, but workarounds are much faster and easier for teams to implement [72].  
256 Frameworks also only offer high-level guidance, resulting in a variety of smaller implementation  
257 decisions between organisations [68, 91, 92]. With differences in implementation, IT infrastructure's  
258 performance is evaluated based on numerous indicators such as customer feedback, internal business  
259 processes, and the learning achieved [64].

260 The choice of tools used to manage IT incidents can have a direct impact on service quality and  
261 efficiency [20], which has resulted in research on the development of tailor-made software [20, 74]  
262 or the customisation of Off-the-shelf or outdated tools [18, 19, 87]. Still, tools alone cannot fix a  
263 broken process or workflow [18, 72]. Understanding the issues and finding an optimal workflow  
264 process before selecting the software can help organisations decide on the tools that best fit their  
265 needs.

266 The majority of IT incident handling begin with calls to the Help Desk and often take the shape  
267 of routine questions and issues that staff can answer confidently. However, around 10% of all  
268 calls require further research and escalation to the relevant teams [62, 83]. Research has found  
269 that end-user satisfaction is influenced by both the perceived quality of the solution [95] and  
270 their beliefs regarding the trustworthiness and level of expertise of staff resolving issues [14]. IT  
271 departments rely on several knowledge sources to alleviate pressure on staff and provide information  
272 regarding commonly reoccurring incidents. These knowledge sources can take numerous forms,  
273 including Internet repositories, cross-organisation shared knowledge base [104], Frequently Asked  
274 Questions [29], and peer advice [90].

275 With peer advice, for example, staff can seek their peers' help when they lack the needed expertise  
276 and the situational awareness due to the distribution of information based on the roles [60, 69], both  
277 justifying the necessity for hands-off between staff [83]. When issues regarding systems' stability  
278 occur, hands-off can involve numerous teams and staff [46] who can provide incident resolution at  
279 the right time and in the right context [83].

280

281

### 3.3 Distributed cognition

282 Distributed Cognition theory (DCog) [39] essentially describes the collaboration between mul-  
283 tiple agents as a single cognitive systems [11, 60]. This collaboration includes human-computer  
284 interactions; thus, DCog is well matched to understand the relationship between humans, tools  
285 and artefacts [34]. In this work, we use DCog as a lens for understanding coordination between  
286 embodied agents by analysing the interactions between the people, the problem, and the tools used  
287 both in planned and emergent cases [11, 71]. *Cues*, and *Norms* are two key features necessary for  
288 supporting work in a distributed context. A cue is defined as a signal that indicates to individuals  
289 the required actions and how to enact them. In contrast, norms are the procedures that ensure  
290 consistency between individuals' tasks [8, 11].

291 Similar to other ITSM findings [8, 9, 60, 102], our initial findings from observing the Help Desk  
292 showed that their work is highly distributed in nature. We therefore chose to use elements of DCog  
293 in our analysis and as a lens in the discussion to understand our results in a wider context.

294

We are not the first to observe that DCog is part of IT management. Individuals from various units of an organisation collaborate formally and informally to address IT issues, which are characterised by pattern recognition, hypothesis-generating, and testing for uncertain success [9]. For example, Maglio et al. combined distributed cognition framework with joint activity theory to understand a specific problem-solving instance in web-based administration [60]. Botta et al. further expand on this by applying distributed cognition within the context of security management and identifying its influence over organisational processes [8]. However, little work has been done on understanding distributed cognition when resolving phishing incidents by various IS teams.

#### 4 Participating Organisation

The University studied is an internationally recognised UK institution, which supports around 40,000 students and 15,000 academic and administrative staff members. It is distributed over multiple campuses inhabited by several academic schools, each with their own respective local IS management teams that manage their own resources. In total, there are around 1000 IT support staff.

The University's IT service management is guided by ITIL (Information Technology Infrastructure Library), a framework of best practices which is used by organisations worldwide and in diverse sectors and industries [1]. The University has adopted and adapted the ITIL framework for use by all IT services, creating a dedicated Quality Enhancement team to ensure compliance. Since the adoption of ITIL, the University has reported improvements such as the increased clarity of teams' roles and responsibilities, reductions in services outage, consistent logging of incidents, reduced running costs, and improved customer satisfaction. Improvements such as these are considered significant indicators of a successful implementation of ITIL [72]. Additionally, the University's IS achieved a Service Desk Certification<sup>1</sup>, which is an industry accreditation program specifically designed to certify service desk quality, indicating that the University's ITIL implementation and the workflow detailed within this case is comparative to other organisations using ITIL to inform their phishing practices.

#### 5 Methodology

The study data was collected from two sources: 1) ethnographic-style observations of the daily work of the Help Desk, and 2) interviews with other University teams. We followed our own University's ethics procedures and at all times ensured that participants were aware that participating in our research study was voluntary.

*Introductions and setup.* Before starting our research project we met with some of the stakeholders, namely the Chief Information Security Officer (CISO) for the University and the Help Desk manager. We explained our project goals and discussed possible project structures that would allow us to conduct the research in a minimally disruptive manner as well as provide insight that might be useful to the University. We also discussed the phishing-related issues they thought were most problematic for their teams. We used these insights as initial scoping for the semi-structured interviews discussed below.

*Observing the Help Desk.* The CISO and Help Desk manager confirmed that the Help Desk was the intended first point of contact for anyone reporting a phishing message. They are responsible for initial assessment of the report, escalating it if needed, and responding to the reporting user.

Given the Help Desk's central role, we started by observing their workflow. The Help Desk manager allocated a desk for the lead researcher so she could spend time with the team and

<sup>1</sup><https://www.servicedeskintstitute.com/service-desk-benchmarking/service-desk-certification/>

344 familiarise herself with their work practices. She started by shadowing Help Desk staff while they  
345 were doing their daily work. Initially, she only observed and asked about the full range of their  
346 normal work practices, then in the second week she started focusing more on phishing-related work  
347 practices. The normal mode of observation was to quietly observe the staff doing their work and  
348 then ask follow-on or clarifying questions when staff were free. The observations were contextual,  
349 bordering on ethnography.

350 The observations were done over two months, with the researcher taking notes and spending  
351 between one and two days a week observing. They conducted focused observation of six Help Desk  
352 staff with each observation lasting 4-5 hours. They also spent time at the provided desk observing  
353 the flow of the space and briefly observing different staff as interesting incidents arose. Observed  
354 staff included experienced staff, new staff, and an undergraduate computer science student doing  
355 work experience.

356 The researcher was also given limited access to the ticketing system used by the University to  
357 help teams track all types of issues within the University. She also attended the training sessions for  
358 using the ticketing system. The Help Desk uses the ticketing system to manage communications  
359 with users and other teams. They refer to all interactions with other groups as “calls” and track  
360 them through the ticketing system. Calls can be digital, but they can also be a phone call or someone  
361 walking into one of the physical Help Desks and asking a question, all of which are logged in the  
362 system. The lead researcher was able to use the ticketing system to better understand how phishing  
363 calls were handled and passed between teams as well as understanding communications between  
364 the Help Desk and end-users. Additionally, the lead researcher was given limited access to the  
365 Help Desk’s private knowledge base to better understand the observed practices. Throughout the  
366 research we took care to use these resources respectfully and quotes used from them in this work  
367 have been carefully redacted to protect staff and end-users.

368 As expected, phishing-related tasks are infrequent and tend to occur in clusters, such as when a  
369 single phishing campaign generates many calls in a short time period. Consequently, the researcher  
370 was only able to observe one live reaction to a phishing campaign. Instead, during breaks in work  
371 the researcher asked staff about their prior experiences with phishing calls. Because the ticketing  
372 system is normally open during work, it was easy for them to pull up prior phishing calls they had  
373 handled and discuss them.

374  
375 *Interviewing University teams.* Observing the Help Desk also gave insight into the work flows of  
376 the teams they work with, most of whom use the same ticket tracking system. To better understand  
377 the work practices of these other teams, we conducted interviews. Most of the interviews were  
378 contextual interviews [35] where during or after the interview the participant showed the lead  
379 researcher real phishing handling examples of how they do their work, the systems they use, and  
380 metrics from previous attacks. Interviews were mostly conducted in nearby meeting spaces to  
381 avoid disturbing other staff.

382 The Help Desk manager provided introductions to other University teams that deal with various  
383 aspects of phishing, even if their involvement was minimal. We were therefore able to interview one  
384 or two members from six teams, each of whom work on a range of phishing-related issues including:  
385 dealing with users, account resetting, desktop computer rebuilding, best practice management,  
386 email relay management, interface with Office365 email, and security. Unfortunately, we were not  
387 able to interview the team that manages the network and the virtual team focused on security.  
388 In total, we conducted about 25 hours of interview. All the interviewees were experienced staff  
389 who worked for the University for more than 4 years and most of them were their team’s manager  
390 or leader. The lead researcher also constructed a diagram of inter-team workflows from early  
391 observation and interview content, she then iteratively improved it by asking staff in following  
392



393 interviews about the accuracy of the identified inter-team interactions, procedures, and their  
394 phishing-related roles.

395 Team interviews started by explaining the project and the general goal of understanding how  
396 the University handles phishing reports. We then asked them to explain their team's mission in  
397 their own words and their general work practices. We then narrowed in on their phishing-related  
398 activities, including their interactions with other teams. The bulk of the interviews involved follow-  
399 up unstructured questions on issues they brought up or topics we were aware of from other team  
400 interviews or Help Desk observations.

401  
402 *Periodic review of findings.* During the Help Desk observations, the lead researcher setup some  
403 feedback sessions with staff where the researcher summarised their findings and ask for feedback  
404 or corrections on the observations. To help guide the research further, approximately twice a month  
405 the lead researcher would give a slide presentation detailing their latest interesting observations  
406 from the Help Desk shadowing as well as the interviews to our research lab who were encouraged  
407 to ask questions and comment. The presentations included the developing diagram of inter-team  
408 workflows as well as information flows within teams and between the Help Desk and end-users.  
409 The presentations were used to help the lead researcher process observations as well as identify  
410 areas that needed follow-up to understand. As these presentations happened regularly, the lab  
411 group was also able to provide needed external clarity. Notes of key points were taken during and  
412 after meetings.

413  
414 *Interview data analysis.* Interviews were audio recorded if the interviewees allowed it, if not, the  
415 researcher took detailed notes and completed a write-up immediately after finishing the interview.  
416 All audio recordings were transcribed by a researcher, with participant's personal information and  
417 team names being substituted for IDs both in the transcriptions and notes.

418 Two researchers reviewed all the transcripts and notes. Using open coding as they went, both  
419 researchers constructed their own independent codebooks focusing on the process of handling  
420 phishing. The researchers then met to discuss the process that had been observed. Through iterative  
421 coding and discussion sessions the researchers reached agreement on the workflow for handling  
422 phishing as well as the problems and friction points. Following each round of discussion, the two  
423 researchers provided feedback to the third researcher so as to guide reporting of results.

424 During the open coding, it became evident that phishing management at the University was  
425 an example of Distributed Cognition as the process of managing phishing clearly involved more  
426 than standard escalation of issues, and instead required multiple groups to communicate about  
427 their own unique perspectives of the incident and work together to manage it properly. DCog was  
428 therefore used to guide the analysis by putting more emphasis on the communications between  
429 teams, particularly, points where one team had access to data or resources not available to the other  
430 teams and how that information was being conveyed.

431 To ensure accuracy of the presented results, an early draft of this paper was shared with stake-  
432 holders and their comments were discussed and addressed.

## 433 6 Results

434  
435 Phishing is handled by multiple teams within the University, including: Security, Quality En-  
436 hancement, Help Desk, Mail Relay, Mail Exchange, Network, Incident Response (IRT), and Service  
437 Delivery teams. The level of involvement of each team is different; some teams routinely handle  
438 phishing-related issues while others are only involved in emergencies or other specific circum-  
439 stances. End-users are also an important part of this distributed process as they identify, report,  
440 and ask advice about phishing by contacting the Help Desk.

## 6.1 Phishing campaigns – managing the load

In this section we focus on how the Help Desk manages phishing calls. We observe that their largest problems involve: managing large numbers of reports coming in, deciding what reports need to be escalated to other teams, closing out the report calls efficiently, and using their own judgement on non-standard phishing reports.

The most common phishing attack handled by the Help Desk is *phishing campaigns* where the phisher sends emails to many recipients to increase the odds that one or more will interact with it. The emails are often visually similar but contain variations, such as putting the recipient's name in the email body (e.g. "Hello Alice,"), using slightly different body text or creating custom URL links for each recipient.

In most cases experienced by the Help Desk, a phishing campaign will use similar subject lines for all the emails. When reporting phishing, users often forward the email, causing the ticketing system to automatically adopt the subject of the phish as the subject of the ticket. If a Help Desk staff member finds a phishing email in their own inbox, they can confirm a campaign by comparing this email to those already reported in the ticketing system. It is one of the cues they use to verify phishing campaigns. Help Desk staff typically use the number of reports with similar-sounding subject lines reported in a short time frame as a signal of a campaign. Often, such sets of reports will happen in the morning due to users checking their email then.

[Help Desk] Morning usually is the peak time for us. We receive calls between 8 and 9 am because usually staff will come in the morning check their emails and report phishing.

After determining that they are looking at a phishing campaign, staff select one or more of the calls to escalate to the appropriate teams. The remaining similar-subject calls are either temporarily ignored by staff or grouped together into a single open call to reduce clutter in the ticketing queue. Later in the day, when new phishing reports have stopped coming in, a staff member will voluntarily go through and close all phishing calls at once.

The above workflow has naturally evolved as a way for Help Desk staff to manage phishing reports alongside their other service delivery tasks.

*6.1.1 The number of phishing reports can overwhelm.* The Help Desk's primary goal is to optimise the number of calls processed, either by closing or escalating them, ideally taking less than 15 minutes for most calls. A phishing campaign is problematic for the Help Desk because it generates a large number of phishing calls, each of which must theoretically be handled individually, taking time away from other calls. For 2019, they received phishing-related calls on at least 20 days out of every month, with call counts ranging between 2 and 170 per day. While that number may seem large, it only represents a small percentage of the University reporting phishing. If a theoretical phishing campaign were to target all University staff and students (about 58,000 people) and even 1% were to report it, that would be 580 calls, well above the normally observed number.

Most of the teams, including the Help Desk, agreed that having people report phishing was needed as it is the cue for identifying campaigns. However, they also recognised that the University did not have the resources to look through all the phishing reports. An Quality Enhancement staff member quote explained:

[Quality Enhancement] We want people to report [phishing emails] but we want to be able to manage the load of calls. The problem is we cannot manage them.

The Security team similarly recognised the problem of a lack of resources impacting the Help Desk's capacity for managing reports:

[Security] At the moment we don't have the resource to have someone look at them and triage them which is our main problem.

491 This overloading problem resulted in the Help Desk adapting their practices to fit the ticketing  
492 system's functionality, as we discuss later. Another tactic the Help Desk uses is to get help from  
493 other teams, such as the Mail Relay, to block still incoming campaign emails and remove existing  
494 phishing from peoples' inboxes. Doing so has positive security impacts, but more practically, it  
495 stops the flood of reports making it a strong immediate motivator to react fast.

496 Help Desk overload also impacts the University's ability to send simulated phishing messages  
497 to end-users as part of security measurement and training. Sending such fake phishing emails is  
498 currently an industry best practice to understand how well-trained staff are and if the procedures put  
499 in place are effective. However, when the University attempted such an exercise, they unintentionally  
500 overwhelmed the Help Desk with calls as the Mail Relay's normal work-management strategy of  
501 quickly blocking the incoming attack cannot be used on simulated phishing. The overload damaged  
502 the Help Desk's ability to do their daily tasks and their ability to provide customised feedback to  
503 reporting users, resulting in a loss of a potential user-training opportunity.

504 [Security] The problem is when people then report [the fake phishing email]; there is no  
505 inbuilt system on our email that says "Wait, this is a fake one, calm down". So with the issue  
506 in the Help Desk, Help Desk was completely swamped and that what happened after we ran  
507 the simulation the first time.

508  
509 As a result, the Security team temporarily stopped sending fake phishing emails and are working  
510 on finding ways to better conduct training and testing. Including agreeing to be part of the research  
511 project described in this paper as an effort to better understand how phishing reports flow through  
512 the organisation and where provisioning is needed.

513  
514 *6.1.2 Deciding what to escalate so as to not waste other teams' time.* The main point of the Help  
515 Desk is to decide what does and does not need escalation, so other teams only spend time on  
516 problems that require their expertise. Hence, one of the Help Desk key tasks is triage, where they  
517 sort through reports and identify which calls require escalation to the appropriate team.

518 In addition to wasting time, escalating unnecessary calls can also result in a polite rebuke from  
519 the other team. We observed several situations where a call was escalated, and the other team  
520 responded that they had already handled this one in the morning, or that necessary information  
521 was not present. Consequently, the Help Desk staff try only to escalate calls when necessary.

522 [Mail Relay] If they are asking us to block a specific email and they told us 10 times already,  
523 then we do not need to see it again ... but what we do not want is the Help Desk passing all  
524 the hundred calls to us just saying the same thing "This person received a phishing email".  
525 We only need to be told once.

526 Given the need for human validation of potential phish, there exists no formal approach for  
527 identifying the 'correct' cue and the applicable actions. Therefore, we detail the range of identified  
528 cues and the associated norms made available to Help Desk staff.

529  
530 *Is the reported phishing email from a University email address?* Compromised University accounts  
531 are a serious problem, so if a reported phishing email is from a University email address, the  
532 Help Desk will handle it differently. We further detail how compromised accounts are handled in  
533 Section 6.3.

534  
535 *Is this a campaign or a one-off attack?* A key criterion for the Help Desk is the number of reports.  
536 If there is only a small number of phishing reports, then requesting action to block incoming emails  
537 is likely a waste of other teams' time. The attack may be a one-off, or the user could just be confused.  
538 Neither case justifies escalation.

539

Looking at the phishing emails reported during May, we found 9 phishing campaigns with more than 4 phishing reports, 12 campaigns with 2-4 phishing reports, and 40 phishing reports of single phishing emails. Using the helpline strategy, only 21 calls would have been escalated in May.

*Do other teams already know about this phishing campaign?* If other teams are already aware of the problem, then escalating again is unnecessary. So before escalating a call, Help Desk staff check that it is not already being handled. The Help Desk maintains a separate list of phishing emails that are logged with Mail Relay. The list includes the subject lines of phishing calls that have already been escalated.

[Help Desk Knowledge base] Due to the current number of phishing attempts being made on our email service, we really need to monitor what we have sent to Mail Relay so we don't end up inundating them with calls.

Sometimes other teams become aware of a phishing campaign through other sources, such as getting the phishing email themselves or being notified by another University, detailed in Section 6.4. When that happens, they notify the Help Desk so that they do not escalate the calls again. However, the various sources of notifications resulted in missing them or delay the response process. In the below example, a call was escalated to the Mail Relay team, who then deescalated it because the IRT team had already created an earlier call for the same issue which the Help Desk would have been copied in on.

[Mail Relay deescalating a call] The Help Desk were passed a call about this from IRT earlier today.

*Is all the necessary data present?* When reporting, users often forward the problematic email or provide snippets of it in their communication. Nevertheless, the specialised teams such as Mail Relay need the header data from the original email, which is not included when the email is forwarded.

Fig. 1 shows two examples of why the original emails (.eml files) are needed. Both pictured emails are from the same campaign sent within minutes of each other to two Lecturers in the same University department. To a human, they are obviously part of the same scam, but they are quite different to a computer. The sender emails are different, the use of link text differs, the email text body itself is different, they have different HTML formatting, even the subject has variations. These example variations are consistent to those seen in prior work [57, 94]. Automatically identifying this scam from the email of a whole University is not possible using only the pictured information. However, the .eml files contain additional information such as a list of all the email servers the email passed through, their IP addresses, encryption signatures, and any checks performed by the servers.

The easiest way to get the header information is to have the user save and send the original email as an .eml file. Reporting tools, such as ones the University is currently pilot testing, will attach this file to the phishing report automatically, but at the time of observation, the only way to get it was for the user to provide it explicitly. To save time, the Help Desk has a pre-written response text, called a *Standard Solution* (SS)—detailed in Section 6.5, for asking the user to send the email as an .eml file, including instructions on how to do so. So if the .eml is missing, the staff managing the call will ask them for it using the SS. The user may then take some time to respond or not respond at all, delaying the report escalation.

*Escalating the call.* Once the Help Desk has all the needed data, they again check that the call has not yet been escalated, and if not, they escalate it. First, the Help Desk staff update the aforementioned list of escalated calls with their phishing incident's subject and then escalate the call to the Mail Relay, Mail Exchange, and IRT teams. Typically the escalation will be done on a single call so that all three teams can see comments made by the others. The purpose of the escalation to

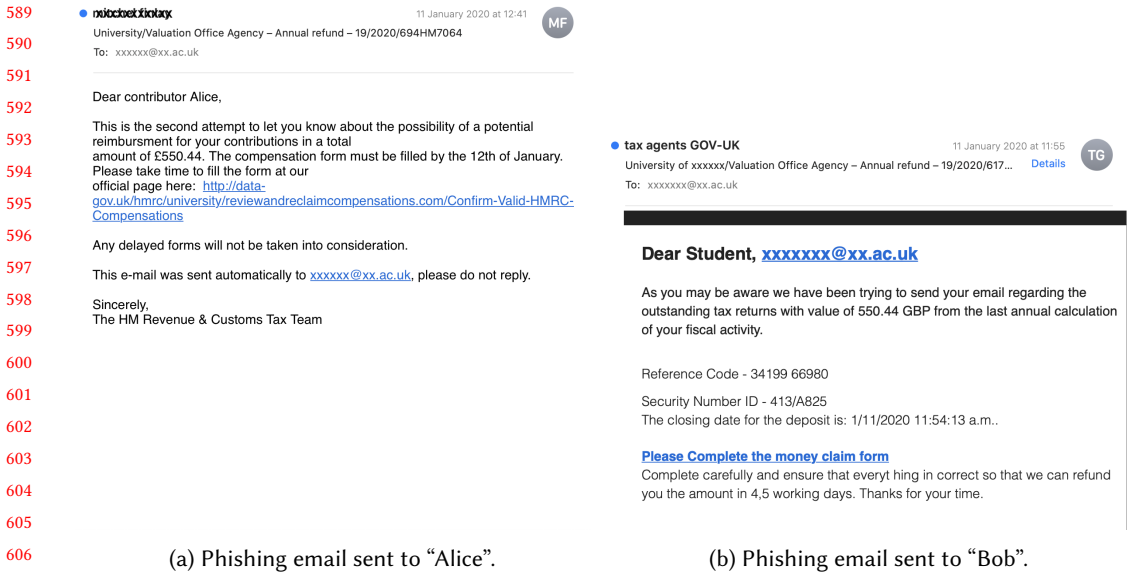


Fig. 1. Phishing emails sent to two University Lecturers in the same department on the same day. Mildly edited to make them anonymous.

the three teams is to ensure situational awareness across the ITSM department; however, only one team may work on the call depending on the situation.

6.1.3 *Closing calls is also time consuming.* Closing a phishing call involves providing feedback to the end-user and marking the call as resolved in the ticket tracking software.

[Help Desk] We would escalate one of the calls to the appropriate team to action in terms of blocking at the relays or removing the offending email from mailboxes. The rest of them we would just contact the user to say delete the email, change your passwords if you’ve clicked on a link etc.

To do so, a staff member has to look up the phishing SS in a long list, pick the appropriate one based on the user query in Table 1, select it, send it to the user, copy it to the call history, and then mark the call resolved. The actions are fairly simple, but when 100+ calls all require the exact same set of actions, it can get repetitive.

To combat this, the Help Desk staff will tag all incoming phishing calls with unclear subjects by adding “[Phishing]” to the subject and leave them open on the system. They will then wait about 30 minutes for any further calls or till the rate at which reports come in slows and then bulk close the calls via an internal communication system. They then select everything either tagged “[Phishing]”, or that is clearly part of the day’s phishing campaign, and close them out all at once, sending all the reporting users the same SS and marking all the calls as resolved as one bulk action; thereby, saving everyone quite a bit of call-closing time.

[Help Desk] Sometimes the volume gets beyond us, we are not going through this one by one and contact each user individually with SS. We are gonna just bulk and close these calls with the generic bulk closure message, it goes out with basically goes out and say if you want more information, look at the call in the service portal because sometimes we don’t have the capacity to chip through these.



638 *6.1.4 Help Desk staff are relied on to use their own judgement to identify non-standard problems.*  
639 While the Help Desk heavily relies on standard protocols, procedures, and processes they also  
640 encourage their staff to use their own judgement around how to handle each call and provide them  
641 with the tools to do so. While seemingly unremarkable, it is important to understand that staff are  
642 encouraged to handle calls within minutes and therefore tend to apply the heuristics described  
643 above to handle phishing calls quickly. Here we detail a couple of situations where staff noticed  
644 something odd and followed up in a non-standard way.

645 When shadowing a Help Desk staff member, we observed them take a call from a single phishing  
646 report which would typically be closed without escalation. The report complained that the forwarded  
647 phishing was impersonating their Department Head and asked the Help Desk to block the emails.  
648 The Help Desk staff member thought that the phishing email was particularly believable and  
649 therefore might constitute a high risk, but they were still uncertain if it was worth escalating. So  
650 they used a University provided message trace service on the email's from address and found that  
651 the phisher had sent the email to around a hundred users. So they immediately escalated it to the  
652 appropriate teams.

653 In a separate instance, the Help Desk received a large number of calls where users reported a  
654 legitimate email from within the University as phishing. The Help Desk team investigated the  
655 reasons for that misjudgement and found that an email sending service had been used that made  
656 the sent emails look like they might not actually be from the University.

657 [Help Desk] We have false positive occasionally. It has been assisted now. I am trying to think  
658 what is the bulk mail. XX mail is a commercial book mail relay that you can buy into that  
659 allows you to send lots and lots of emails to target audiences. Certain colleges, were using  
660 [it] for email campaigns and things like that. The problem with that is not coming from the  
661 University email address. It has The University name in some part of it but it is not xx.ac.uk  
662 and people were reporting that it is phishing but it is not. We are not having it anymore  
663 because changes were made to make the email that is sent to people more obvious that it  
664 is part of the University. It should have the University logo, University address, and looks  
665 professional.  
666

667 To address the problem, the University banned the use of email services that do not comply with  
668 best practice, such as providing URLs hosted under the University's domain.  
669

## 670 **6.2 Converting escalated calls into protections**

671 In this section we discuss how the Mail Relay and Mail Exchange teams handle escalated calls from  
672 the Help Desk along with other teams. Handling an escalated call most often involves finding a  
673 reliable way to identify the phish and then applying that method across several systems. Both of  
674 which may require hand-offs between several teams to find solutions and implement them. Also,  
675 some phish simply cannot be blocked or require several attempts to do so.  
676

677 A large component of handling an escalated phishing call is ensuring that University users are  
678 protected from the reported phishing attack. People have a wide range of skills at identifying  
679 phishing [49] with some quickly reporting the phish and other more vulnerable users clicking on  
680 malicious links. To protect these vulnerable users, teams can use reports to remove phish from  
681 inboxes as well as prevent new phishing from entering the University email system. Doing so  
682 primarily falls to the Mail Relay team who handle incoming and outgoing emails and the Mail  
683 Exchange team who handle stored emails as well as emails managed by Microsoft's Office365.  
684 However, other teams may be required if the situation dictates so. Phishing campaigns can vary  
685 dramatically in sophistication with the process being straight forward if all the phishing emails share  
686

687 a unique feature, such as all coming from the same email address. But other phishing campaigns  
688 can be more complex (e.g. Fig. 1).

689

690 *6.2.1 Finding a common reliable factor to protect users.* Many features are used to block phishing,  
691 but some of the most common are the from address, subject line, and any mail relays listed in  
692 the header. The teams use a combination of experience and educated guessing to select potential  
693 features and then run practice searches to see how effective it is.

694 While the main process is fairly basic, doing it reliably is not easily learned as it requires  
695 knowledge of email header's structure as well as an understanding of IP addresses and IP address  
696 ranges. Feature selection also has to be done carefully so blocking and deleting will not interrupt  
697 users' work. A simple example might be a phishing attack that closely mirrors a genuine Dropbox  
698 email. If the Mail Relay team chooses to block on features that are also shared by real Dropbox  
699 emails, they could inadvertently damage all users' ability to interact with Dropbox. Universities  
700 are also complex, with staff working on a wide variety of problems, using a wide variety of tools,  
701 and collaborating with a wide variety of people. So blocking whole IP ranges, domains, or tools  
702 cannot be done lightly.

703 The effect of blocking or deleting emails can also range from time consuming to impossible to  
704 reverse. The Mail Relay team can trap emails in quarantine and then release them for delivery if  
705 they were incorrectly identified as spam. However, when the Mail Exchange team deletes an email,  
706 it is permanently gone and cannot be recovered. As a result, the Mail Exchange team has internal  
707 procedures around email deletion, such as requiring the manager to approve all deletion commands.  
708 They also require specifying a date range, from address and subject line in the deletion commands  
709 to limit potential damage.

710 [Mail Exchange] Within Office 365, you can do what we call it content search. There is one  
711 various procedures to ensure we don't do it unless we have the approval to do it but yes we  
712 do have the ability to do that and any thing we do is fully auditable.

713

714 *6.2.2 Fixing requires knowledge, resource, and responsibility.* University teams are generally either  
715 centred around resources that need to be managed (e.g. Mail Relay) or around specialised knowledge  
716 that takes time to accumulate (e.g. Security). Knowledge regarding how to address a phishing  
717 incident is often located in many parts of an organisation. This situation is fairly common in IS  
718 support as different teams and members often gain in-depth understanding of the systems they work  
719 in and then collaborate with others as needed [46]. In terms of phishing, teams need knowledge,  
720 both in terms of making non-disruptive changes to running systems and making changes that  
721 are likely to have the desired security outcomes. Large computing systems are made of many  
722 interconnected components where a single change can have unexpected side effects. To balance  
723 this distribution of knowledge and ensure their changes meet the best practices, the University  
724 teams actively coordinate with the resource-based teams to make the actual changes to the system  
725 and also to ask for advice and sign-off from other teams, such as Security.

726 Access to resources, such as systems, services, and channels, is also often restricted to ensure that  
727 people with insufficient knowledge do not make well meaning but potentially very problematic  
728 changes. For example, it would be much faster and more efficient if the Help Desk could directly  
729 go from identifying an accurate phishing report to applying an appropriate rule to the mail relay  
730 to stop new phish (and reports) from coming in. However, as the Mail Relay team notes below,  
731 knowledge of how to modify the relay safely is not quickly learned.

732

733 [Mail Relay] It is so complicated only experts can understand [filter rule construction], I don't  
734 think the Help Desk have access to the black list.

735

736 Consequently, the Help Desk cannot edit themselves and instead raise a call and wait for the Mail  
737 Relay team to take action.

738 University systems also frequently impact each other, requiring teams to coordinate or ask  
739 each other for help addressing an issue. For example, a Mail Relay team member recalled a prior  
740 phishing campaign when so many emails were coming in the same time, causing the relay server  
741 to experience overload. So despite applying mail filtering rules, the phishing emails were still  
742 negatively impacting users' ability to get emails. To solve the problem, they reached out to the  
743 Network team to temporarily block all incoming emails from the phisher's IP address range. Doing  
744 so reduced the load on the mail relays and resulted in a successful blocking of the malicious emails  
745 with minimal impact on University services.

746 [Mail Relay] We do communicate with Network team in active attacks from a specific IP.  
747 Network team will block the traffic to save resources.  
748

749 In situations requiring specialised knowledge, teams may require consultation, guidance, or sign-off  
750 from other teams before enacting a change. For example, phishers will sometimes send emails from  
751 an IP address range to make blocking the email harder. IP address allocation is a constantly shifting  
752 feature on the Internet as addresses are assigned and reassigned to different Internet Address  
753 Registrars. Some ranges are unallocated and can be safely blocked, while other ranges might be  
754 used by valid users. The Mail Relay team's knowledge does not necessarily include an in-depth  
755 understanding of the current IP Address allocation issues, so when choosing to block IP ranges  
756 they may consult with teams like Security or Network who have this knowledge.

757 [Mail Relay] We consult Security with policy and security. For example, if we have an attack  
758 from a country, we will escalate it to them for decision ... We can't block all the emails from  
759 Nigeria. It will damage our reputation when someone from Nigeria tries to contact us. We  
760 should just balance the cons and pros of each rule. Many emails will be trapped by blacklists.  
761

762 While some teams play a direct role in the day-to-day responsibilities of handling phishing, others  
763 monitor the ongoing incidents and processes used. For example, the Security team observes the  
764 internal processes and advises on change to better promote security practices. They do this by  
765 tracking ongoing security incidents handled by the IRT through monitoring of the IRT's resources:  
766

767 [Security] We have access to the IRT's inbox. So we can see what is coming in and out. So, we  
768 know what is happening and we can actually see the [ticket system] calls that are assigned to  
769 IRT's inbox. So that how we know what is going on.

770 *6.2.3 Some messages cannot be blocked or removed from inboxes.* Sometimes, even with the help of  
771 multiple teams, phishing emails cannot be reliably blocked or deleted. During the data collection,  
772 a large phishing campaign hit the University. It was sent from different email addresses and had  
773 unique subject lines, often involving recipients' names. The Mail Relay applied several blocks on  
774 the relays, but the number of from addresses involved made it challenging to fully block or remove  
775 them from users' inboxes. Instead, the Quality Enhancement manager relied on the users as the last  
776 point of defence and sent an email on to all University members explaining the email characteristics  
777 and asking them to delete any email that matched.  
778

779 [Quality Enhancement email] ... the University is receiving a very significant number of  
780 phishing attempts which have the recipient's name in the subject field and some text in the  
781 body of the message ... If you have received a message like this can you please delete it without  
782 clicking on any links within. If you have received a message like this and clicked on the link,  
783 please contact [Help Desk] for advice.  
784

785 6.2.4 *Email blocking does not always work on the first try.* Normally, the Mail Relay team will use  
786 an escalated call to construct a blocking rule that effectively blocks all emails of that type. But  
787 sometimes the blocking rules miss a portion of the campaign.

788 In the following case, a campaign had happened, the Mail Relay team had been escalated to, and  
789 end-users had been sent the SS. The user questioned the block's effectiveness and sent a follow-up  
790 call in response to the first phishing report to ensure that a block was correctly applied on the mail  
791 relays. However, they were not satisfied with the response and sent in yet another call.

792 [End-user] Got another three e-mails through from the xxxx@gmail.com address, so that the  
793 block you have put on does not seem to be effective. Could you check what has gone wrong?

794 In response to the initial call, the Mail Relay team had applied a block based on the phisher's email  
795 address, but the phisher was sneaky and bypassed the block. A Help Desk staff correctly identified  
796 that the above email was an indication that the Mail Relay team's solution was not working and  
797 escalated the call despite it already being on the list of previously escalated calls. Using the new  
798 information, the Mail Relay member consulted the staff within the for a solution; thus, Mail Relay  
799 team expanded their rules to look for the problematic email address in multiple header locations.  
800 Also, given the sneaky attacker behaviour they also chose to block all email traffic originating from  
801 the IP range the attacker was using.

802 [Help Desk escalation] Hi Mail Relay, The block on the below email put in place in call  
803 xxxx2241 appear not have worked. Can you advise on what else can be done to try and stop  
804 these emails getting through?

805 [Mail Relay member 1] Any idea why my entry in the access file didn't work?

806 [Mail Relay member 2] The simple way of blocking an address relies on the return path -  
807 which is what the mail system sees - being the same as what is in the from header. That is  
808 usually the case, but in this case it was not. They are being rather more sneaky than usual. I  
809 have put in a different sort of block which should block anything with "xxxxxxx@gmail.com"  
810 in the From: header. I also identified four network [IP] blocks from which these spams have  
811 recently originated (though there may be others I have not found) and blocked all traffic from  
812 them.  
813

### 814 6.3 Compromised accounts – a serious reputation and workload problem

815 Accounts can be compromised as a result of a successful phishing attack. Once the victim provides  
816 their credentials on a fake web page, the attacker can use them to gain access to the victim's  
817 accounts. In this section, we discuss University practices of reported compromised accounts. More  
818 specifically, we detail the cases in which they hand-off the calls to IRT team or rely on users to  
819 revoke the attacker's access to their account. IS teams have strong security awareness with regards  
820 to compromised accounts; thus, reacting on them is independent from reacting to other phishing  
821 incidents.  
822

823 Compromised accounts were found to be mutually understood by all teams, and were unani-  
824 mously identified as definite security concerns, given that multiple services can be accessed using  
825 the University's single authentication system. Among other actions, they can use the accounts to  
826 gain access to send out more phishing emails. Similar to findings from previous research [32], staff  
827 from Mail Relay, Help Desk, and Mail Exchange believed that this is the most common scenario for  
828 compromised accounts in University.

829 [Mail Relay] Every time an account is compromised, effectively theft of personal information.

830 Blocking phishing emails is done internally by the spam filter on the mail relays managed by the  
831 University's Mail Relay team and also by the default spam filter on Office365 Exchange Server  
832 managed by Microsoft. Phishing emails from compromised accounts are treated more seriously  
833

834 than other phishing emails since they do not go through the relay's spam filter and they get a lower  
835 spam score from the Exchange filter because they are from the same domain (internal users), which  
836 is considered trustworthy.

837 [Mail Relay] Messages sent internally from one student to another student do not go through  
838 any of that relay's scoring. It is all set within Microsoft Office365.

839 [Security] Add scores to the emails is the things that Office 365 and Microsoft do. They have  
840 what is called spam ratings and they basically just go: "this looks like spam and this does not",  
841 based on the huge number of criteria. We use the scores to some extent to block things that  
842 have basically the highest score because always entirely they will be.  
843

844 Compromised accounts are sometimes used to send out spear and loosely targeted phishing emails  
845 aimed at specific groups of individuals. Both of these attacks are hard to detect by end users [49].  
846 We observed several cases of these attacks while shadowing the Help Desk staff. In one case, the  
847 attacker compromised the victims' account and used it to send several emails to other University  
848 members hoping that one of the recipients knows the victim and replies to the attacker.

849 [Help Desk] The user here reports [Spear Phishing]. So the user was aware of the phishing  
850 type. The email said only "Are you available?". If the user knows the person who sent the  
851 email, they will communicate with the attacker and maybe transfer money to them. Mail  
852 Relay told us that the sender accounts are compromised.

853 Compromised accounts also pose a reputation problem. They can be used to send a malicious  
854 email to other organisations with the University's name associated with it. Considering that  
855 University is moving toward applying the cryptographic signature of the University's mail relay,  
856 other organisations use these signatures in their spam filters, so email coming from a University  
857 compromised account is more likely to make it through other organisations' spam filters and harm  
858 reputation when their teams have to handle the messages.

859 [University public Knowledge base] Phishing campaigns lead to compromised accounts and  
860 as it stands currently, risks the integrity of the University's reputation, data and ultimately,  
861 University business.

862 Compromised accounts are one of the most critical impacts of phishing campaigns as they can  
863 quickly snowball in scale as internal emails can reach more users who can potentially become  
864 compromised and phish other users.

865 Despite their importance, fixing compromised accounts is time-consuming and hard to do at  
866 scale as each account must have its password reset individually.

867 [Mail Relay on behalf of IRT] The problem we have is that often in the time between the  
868 initial compromise and the sending of spam and then us trying to delete that phishing from  
869 other users mailboxes, even if that is 10 to 15 minute, we get dozens of people who then get  
870 compromised. We then have an ongoing problem of phishing going around the University,  
871 moving from one person to another person to another person, and we can't keep deleting  
872 them all. So ideally, what we want to do is stop getting them in the first place, which is a  
873 really difficult thing to do.  
874

875 *6.3.1 Reported phishing from new compromised accounts should be escalated.* For every phishing  
876 incident, the Help Desk staff look at the sender of the phishing email; if it comes from an internal  
877 user, such as a student, it is considered a compromised account and requires more immediate  
878 actions beyond those discussed in Section 6.1.

879 The escalation of compromised accounts is similar to the escalation of phishing campaigns  
880 in that only new cases should be escalated; however, the Help Desk should escalate one call for  
881 every phishing campaign whereas all distinct compromised account reports should be escalated.  
882



883 To ensure the account has not already worked on, the Help Desk staff normally look at the list of  
884 blocked compromised accounts provided by the IRT team and only escalate the call to the IRT if the  
885 account is not on that list. Unlike the list of phishing emails logged with the Mail Relay, the known  
886 compromised account list is maintained directly by the IRT team. Unlike phishing escalation, Help  
887 Desk staff are not expected to verify if the account is compromised before escalating it.

888 [End-user] Just to make you aware, please see above most likely a phishing email that has  
889 been sent from a student account.

890 [Help Desk] Hi IRT, I can't see this account on our list of compromised accounts.

891 [IRT] I see no evidence that the account is compromised.

892 A Mail Relay staff on behalf of the IRT team was positive on hand-off calls that ask if an account is  
893 compromised or not.

894 [Mail Relay] If the Help Desk are asking us to look at a specific email from different account  
895 saying is this account compromised, we will look at each individual one of those.

896 Reporting compromised accounts is necessary since IRT staff can use it to trace other compromised  
897 accounts. The example below illustrates how the IRT staff used the report to reset the compromised  
898 account, escalate it to Mail Exchange to remove the emails, and also lock other user accounts  
899 proactively because they were observed sending the phishing emails from the same location and at  
900 the exact time.

901 [End-user] I received the attached email over the weekend, I think it might have gone to other  
902 staff. It seems to be a phishing attempt and comes from what purports to be a student's email  
903 address. I did not click the link.

904 [Help Desk] Hi IRT, see attached email and compromised account.

905 [IRT] Passwords have been reset. Call also sent to Mail Exchange to delete the emails. User1,  
906 user2, user3 are also sending phish around the same time from the same place. Passwords  
907 have been reset.

908  
909 **6.3.2 Helping users who self-identify potential compromise.** Some users contact the Help Desk  
910 looking for advice because they clicked on a phishing email, resulting in a compromised account.

911 [End-user] I opened a dodgy email and now my email has been hacked and is sending spam  
912 all over, what should I do?

913 In this case, the Help Desk asks the user to change their password, and no escalation is required  
914 because in most cases changing the password will lock the attacker out. Time-wise, this is the  
915 best-case scenario since the user can solve the problem and do not lose any access to their account.  
916 Supposing a compromised account is identified by someone other than the owner, then any calls  
917 have to be escalated to the IRT team who resets the user's accounts with a temporary password  
918 and asks the user to change later, which takes more time for everyone.

919 This scenario aligns with Help Desk strategy of shifting the responsibility to users to reduce  
920 future calls when the user re-encounters the same issue.

921 [Help Desk] We also rely on "shift left". Moving everything to the users. E.g. knowledge on  
922 the University website so when users contact the Help Desk, they will receive a link if users  
923 can help themselves.

## 924 925 926 **6.4 Information flow patterns are not in fixed orders or directions**

927 While our focus started with the Help Desk, awareness about phishing campaigns and compromised  
928 accounts may actually originate from various sources. In this section we detail how teams become  
929 aware of an ongoing campaign by discussing the notification procedures within teams or from  
930 outside the ticketing system. Teams may become aware of incidents by receiving notification from a  
931

number of internal and external channels, including the team's own monitoring systems, a UK-wide educational network called Janet, and the contracted organisation that handles out-of-hours Help Desk calls. Regardless of the source used to identify an ongoing phishing incident, each team works to ensure that all relevant teams are aware of new attacks so as to maintain organisational situational awareness. These cues result in responsive action which may not be reported to helpline staff, but will result in communication between service teams as they coordinate operations for protection.

6.4.1 *Phishing awareness can originate from within teams.* Standard monitoring practices by teams can identify suspicious activity within their systems or through their own personal email inboxes. Once aware, teams proactively adjust their systems to protect users from phishing attacks. For example, Mail Exchange can notice phishing incidents if a compromised user exceeded the limit on sending emails.

[Help Desk] Generally, Network look to see when there is strange activity on the network because they have the network logs. So they have some massive traffic ... So what you got here about quota limit, they see some of the other behaviour around that kind of concept that things like massive downloading from users who wouldn't normally do that and then they will feed it in there and say OK here is the thread and start investigating.

Similar to the Help Desk, the team creates a call based on the observed problem and escalate it to other relevant teams to mitigate their own systems. Maintaining vigilance for anomalies allows teams to react immediately.

6.4.2 *External partners can inform teams.* Apart of the University IS teams, awareness of phishing incidents might originate from external partners. For example, many Universities in the UK gain access to the Internet through the Janet Network, which is a dedicated network infrastructure for the "UK research and education community"<sup>2</sup>. If a Janet admin identifies a serious attack, they may choose to block URLs associated with it UK-wide and notify Security or Network teams directly.

[Quality Enhancement] We have for example a huge issue last year where a lot of different Janet organisations were being attacked and Janet is the education network. What happened is gradually different universities started to get attacked but some universities put proactive steps because they were able to share that information to reach a solution.

Janet also informs the Network team about outgoing suspicious communication, so the network team become aware of compromised accounts or machines.

[Help Desk] They are all interlinked because you could end up getting referral from Janet on the basis that someone clicked on a phishing email and gave out some details which is compromised their accounts and it has downloaded some sort of bot that then start trying to connect to all over the world throw your traffic out and trying to contact other hosts which are known to be malicious.

Another external source of information is NorMAN. NorMAN is an Out-Of-Hours service to support University members 24/7. After the work hours, the Help Desk directs all the calls to the NorMAN support desk. NorMAN cannot solve all the calls because they do not have access to the ticket system, so at the end of their work, they provide a report about resolved and not resolved calls. Then, the Help Desk integrate all the calls to the call system and resolves all the calls which could not be handled by NorMAN.

[Help Desk] NorMAN will try to resolve some of the issues but not all of them because they don't have access to [ticket tracking] services.

<sup>2</sup><https://www.jisc.ac.uk/janet>

981 In the case of phishing attacks, NorMAN staff reply to users without escalating the calls to other  
982 teams. Information about any phishing campaigns is included in the NorMAN support report for  
983 the Help Desk so they can prepare for the potentially large number of early morning calls and  
984 escalate them to the other teams. During shadowing, a first-line staff retrospectively told us how  
985 they learned about an ongoing phishing campaign from the manager relaying a NorMAN report.  
986

## 987 **6.5 Providing feedback and guidance to end users**

988

989 The above sections highlight the importance of user generated phishing reports in detecting and  
990 managing phishing attacks. In this section we look at interactions with users, particularly the use  
991 of pre-written standard solutions and impacts of bulk closures on users. Standard solutions may  
992 save staff time, but potentially at a risk to user satisfaction which is seen as a necessary component  
993 of service quality as well as likely influencing users willingness to report future phishing the  
994 encounter.

995 Users contact the Help Desk to ask about guidance around phishing. For example, a user may be  
996 uncertain if an email is real or not and is asking for guidance about if they should respond. They  
997 may have already taken some action, such as reporting the phishing to their bank, and are asking  
998 if that was the correct course of action. Or the phishing email may be threatening them in some  
999 way, such as threatening to shut down their account, so they want reassurances that their account  
1000 will not be deactivated. Some users even engaged with the phishing before realising that it was  
1001 fraudulent and are seeking guidance about how to best manage the situation.

1002 Help Desk staff endeavour to respond to every call in a professional way that promotes user  
1003 satisfaction and resolves any service problems. However, the number of calls and time limitations  
1004 can impact communicating with users.  
1005

1006 *6.5.1 Standard solution (SS) design.* While the range of queries can be broad, most users are looking  
1007 for only a small set of guidance, such as reassurance that the email is indeed phishing, that they  
1008 have done the right thing by contacting the Help Desk, that they should delete the email, how  
1009 to ensure that their device is malware free, and how to protect their account if they did interact.  
1010 The high overlap in needed guidance is a perfect match for a Standard Solution (SS) where several  
1011 sets of guidance can be written by a qualified person and then re-used. Having an available SS  
1012 is also helpful for Help Desk staff who may not themselves be experts in phishing and may feel  
1013 uncomfortable providing self-written guidance to others. SS also allow them to quickly provide  
1014 consistent professional guidance with detailed steps.

1015 The benefits of the SSs were acknowledged by Help Desk members we spoke to, with all reporting  
1016 that they reply to all phishing calls using them. SS were also valued for the time they saved first-line  
1017 staff. Instead of working 5 minutes on every phishing call and writing a reply to every phishing  
1018 email they receive, first-line staff can quickly use an SS for common phishing calls.  
1019

1020 *SS content design.* Phishing-related SSs were developed by the Help Desk based on the most  
1021 common reasons for contacting the Help Desk and focus on: 1) thanking users who reported  
1022 phishing emails, 2) confirming an email is phishing, 3) confirming an email is not phishing, 4)  
1023 helping users who clicked on links, and 5) informing users about phishing simulations. Contents of  
1024 phishing-related SS are detailed in Table 1.

1025 The wording used in the SS were carefully chosen through a collaboration between the Help Desk  
1026 and Security teams to ensure that it both matched what users were asking and that the responses  
1027 were technically accurate.  
1028  
1029

User's query	Summary of SS content
Ask if an email is phishing	Yes it is, staff report informing email team, and provide mitigation steps in case the user clicked on links or opened attachments. No it is not, staff manually write the reason(s).
Report phishing	Thank them, staff report informing email team, and provide steps if clicked on links or opened attachments.
Report phishing (simulated attack)	Thank them, inform that it is simulated, educate user, no action required.
Clicked on a link	Ask them to delete the email immediately and follow provided mitigation steps.

Table 1. Phishing Standard Solutions used by the Help Desk when responding to phishing-related calls.

[Security] And we have occasionally gone to [the Help Desk], the wording you are means the people around being inclined to not report any more. Why you are doing it this way. So, there is some back and forth.

*Consistent messaging.* SSs are also helpful because they provide consistent messaging to users which is generally considered an effective approach in public safety communication [81]. The Help Desk manager was a proponent of consistent and accurate communications because they felt it would enable them to develop a trusting relationship with users. Managing relationships with users also partially motivated the creation of multiple SSs addressing different common requests, because they felt that the template variations would signal that the Help Desk is listening to each user and providing custom feedback.

[Help Desk] We focus on the consistency in answering users. It is important to build up the relationship over the years.

[Help Desk] Standard solution is used for specific processes such as phishing. When a user forget their username, we give them technical process.

*6.5.2 Bulk closing calls has costs.* As discussed earlier in Section 6.1, we observed that Help Desk staff typically read the first few phishing calls in a campaign and then assume that the remaining reports with similar subject lines have similar content and should therefore all receive the same SS. They do so to efficiently use valuable first-line staff time. However, bulk closures can result in sounding tone-deaf as well as lead to missing important information.

Some users may be giving or looking for information beyond reporting or inquiring if the email is phishing. The bulk closures can cause these calls to be missed. For example, one of the campaigns was a cyberbully type of phishing attack, where the attacker threatened to share victims' unpleasant secrets with everyone in their address book. After the bulk closure of the emails, a user replied to the SS trying to get an answer to their specific question:

[End-user reply to SS] I asked whether I should be informing police - given than this was an attempt to extort money from me by threats? It is not simply a phishing email.

Bulk closures can also result in missing valuable information from reporting users. For example, the user below noticed that there were different senders in a two week-long campaign.

[End-user] Back in February fake emails of the "are you on campus?" type were being sent in the name of Miguel, with address miguelxxxx@gmail.com. I reported this.

I just got another one, this time from a different address, miguelyyyy@gmail.com.

1079 The user makes an important point about how the phisher is using multiple addresses, which may  
1080 be helpful to the Mail Relay and Mail Exchange teams. Unfortunately, emails like this one are at  
1081 risk of being missed during a campaign because of bulk closure.

1082

## 1083 7 Discussion

1084

### 1085 7.1 Phishing management is a distributed cognition process

1086 Prior research has found that information service management is, by its very nature, a distributed  
1087 cognitive (DCog) process [8, 9, 102] that requires hand-offs between distinct teams [60]. In a DCog  
1088 system the cognitive elements of a problem are solved not by a single person but by offloading them  
1089 onto an environment made up of a combination of people and technology. Because the technical  
1090 systems of a large organisation, like a university, are managed by multiple teams who must pass  
1091 problems between them, they often evidence DCog in their problem solving approaches.

1092 However, that does not necessarily mean that all processes in service management involve DCog  
1093 as many tasks may only require a single person or small team. In this work we have shown that  
1094 phishing management, at least in this case, evidences many of the elements of DCog. Phishing  
1095 touches on many distributed elements of the University's IS operations, people from different  
1096 teams collaborate to pool their knowledge, unique system perspectives, and technical ability to  
1097 solve phishing problems which occur on a nearly daily basis. To do so effectively, they also make  
1098 heavily use of tools like ticketing systems, knowledge bases, jointly maintained lists, and standard  
1099 solutions, all of which allow them to communicate and share knowledge between members to  
1100 jointly manage phishing. Ticketing systems allow for quick communication that is archived and  
1101 visible to all involved teams, enabling joint problem solving. Tools like knowledge bases and jointly  
1102 maintained lists allow for storing of knowledge for use by other teams or future team members.  
1103 For example, the maintenance of lists of known compromised accounts by various teams so the  
1104 Help Desk can accurately respond to users. Standard solutions are also an interesting inter-team  
1105 example where the Security team worked with the Help Desk to construct user-facing feedback  
1106 about phishing that is accurate which the Help Desk regularly uses in their communications.

1107 Our work suggests that phishing management might be a DCog system for organisations that  
1108 divide their IS operations up into different teams and have those teams coordinate through ticketing  
1109 system like tools. The observation has wider implications because it allows for a deeper and more  
1110 nuanced understanding of the workflows than is necessarily presented in this case. Break downs in  
1111 DCog systems are evidenced partially through challenges that the people involved experience. Such  
1112 as having a fragmented common ground where each individual has their own set of knowledge that  
1113 does not necessarily overlap well with the knowledge of others, requiring not only information  
1114 sharing but also facilitating collaborative operations. In DCog, information also tends to pass  
1115 through networks of people either directly, or more as we observed, through the passing of a ticket  
1116 that progressively builds up contextual cues. This dialogue between teams is needed because no  
1117 one individual has full understanding of the system or the problem, necessitating the information  
1118 sharing.

1119 By thinking of the problem as a DCog issue, several possible avenues of future work and  
1120 interpretation open up, some of which we expand on in the following sections. The most obvious  
1121 space for future work is to look at the role of ticketing systems in supporting the DCog activities  
1122 around phishing management. These tools were heavily used by staff, but didn't provide an easy  
1123 way to do things like give everyone access to the full set of phishing reports, or allow for easy  
1124 highlighting of evidence. Other tools, like the lists of known compromised accounts, were kept  
1125 separate from the ticketing system making it easier for some teams to miss information, leading  
1126 inter-team friction.

1127



1128 Phishing is also an interesting DCOg issue because of its frequency. Other work on DCOg with  
1129 system administrators [46] talks about how they regularly have to solve a wide variety of issues.  
1130 Phishing is different in that it is a nearly daily occurrence and while the details of each attack  
1131 vary, the processes needed to address it are far more stable. Which may mean that the types of  
1132 support needed here can be more easily built into tools than the standard IT management DCOg  
1133 interactions.

## 1134 **7.2 The University wants reporting, but can it handle all the reports?**

1136 Everyone across the security community agrees that more phishing reporting is better [22, 79, 98]  
1137 which is a view that was also shared by University staff as well as the University CISO. It is easy to  
1138 see where this view comes from since staff reporting is the leading way to learn about security  
1139 breaches [22]. Humans also have the nice feature that they are not deterministic, which makes  
1140 their actions harder for attackers to predict and it also means that any sufficiently large campaign  
1141 is likely to include several staff who are skilled at phishing detection, or just overly observant that  
1142 day. Reporting phishing remains one of the best ways IS staff have of finding the attacks that are  
1143 getting through the automated protections, so of course getting more such reports is important.

1144 The problem though is scale. In order to be effective and identify all phishing that is getting  
1145 through, all staff need to report any phishing they see. In theory this is a good practice, but it has  
1146 the side effect of producing a very large number of reports which must be processed by someone.  
1147 The need to manually process the reports is expensive in staff time. Currently, each report must be  
1148 processed individually, even if they are bundled together, that still requires a human to individually  
1149 select each report and then group them. This problem is one that is likely shared by all organisations  
1150 that manage their phishing reporting through ticket tracking systems. With such a manual process,  
1151 an obvious future work area is to look at potential automation or human support systems.

1152 Unlike other service disruption type issues that help desks get reports about, phishing reports  
1153 have a relatively high level of internal similarity. While phishing emails are designed to hide well  
1154 in inboxes, they stand out as a group in the ticketing system due to the similar subject lines, report  
1155 times, senders, and URLs. The University Help Desk already makes use of many of these features  
1156 to manually group such emails. However, important information can also be overlooked due to  
1157 the number of incoming reports. Phishing emails often have some amount of variance to avoid  
1158 detection. A ticketing system doesn't have the tools necessary to identify which reports are the most  
1159 important to look at because they exhibit different features. Reporting users may also be pointing  
1160 out important information that staff do not have time to look at resulting in a lost opportunity  
1161 to learn and improve technical defences [31]. For example, pointing out that the emails are still  
1162 coming in through the filter.

1163 We believe that it is impossible to fully automate the phishing report processing and incident  
1164 management. Instead we recommend looking into ways to use automation to better support users  
1165 in managing increasingly complex situations [6]. The use of automation can make the Help Desk's  
1166 work more efficient while following the best practices [15, 101]. Help Desk staff could use support  
1167 to better leverage their limited time. Their judgement is needed to both decide if the reported  
1168 phishing is real [38], and to handle more specific questions users might have. One solution for  
1169 the large number of reports is to automatically cluster a reports of similar phishing emails as a  
1170 campaign (e.g. emails in Figure 1) and then based on those reports, the system would auto generate  
1171 a detailed report about that campaign including the number of reports received and a list of reported  
1172 compromised accounts. This potential solution can assist the Help Desk so they only need to label  
1173 the first batch of reported phish and also decide if the escalation is necessary.

1174 There is also some room to look at the effectiveness of AI conversational chatbots to look at the  
1175 incoming user questions and automatically assign SS based on the question text. Chatbot use has  
1176

1177 been explored in SOC's to analyse and convey system alerts to security analysts [70] but have not  
1178 yet been used to generate auto response to security incidents reported by users.

1179

### 1180 **7.3 Ticketing tools support distributed cognition but have room for improvement**

1181 We observed the critical role played by the ticketing system in facilitating much of the inter-team  
1182 collaboration. In this situation, the ticketing system itself can be considered an embodied agent  
1183 with direct impact on the success of the distributed process [11] as it stores the phishing evidence,  
1184 and it is how the different teams record what they are doing and communicate. It is also how the  
1185 IRT team monitors what is going on so that they can maintain situational awareness and interject  
1186 information when needed.

1187 While inter-team collaboration seems reasonably well supported by the ticket tracking tool  
1188 they use, there are still some potential areas for improvement and future research. In the case of  
1189 phishing, teams need to see all the incoming reports associated with a campaign to accomplish  
1190 two tasks: 1) tailor their mitigation to cover all the reported phishing variations, and 2) check if  
1191 the mitigation has indeed stopped the attack. The second point also includes a need to see future  
1192 incoming reports.

1193 Ticketing systems are usually built around the standard information service model where a  
1194 problem might happen, such as a power failure. While a service disruption results in many reports,  
1195 they are basically identical in terms of information so only one report ever needs to be escalated.  
1196 Phishing though has some differences because often users each receive a slightly different phishing  
1197 email. So there is a need to be able to group reports together in a way that is visible to other teams.

1198 Currently the Help Desk spends a great deal of time collecting together similar tickets manually,  
1199 they also spend time curating tickets to make sure that the ticket they escalate has all the information  
1200 later teams may need while the suggested solution supports the Help Desk goal to assess the impact of  
1201 the campaign. An obvious improvement for the ticketing system would be to leverage the clustering  
1202 idea to flag a representative number of reports with the features necessary for capturing variations  
1203 within distinct campaigns. Such a system would save Help Desk time to find the variations, but  
1204 it would also allow them to escalate in a way where the most useful reports are highlighted to  
1205 other teams while also giving them visibility of the other reports, indicating the scale of problem.  
1206 Flagging such information would provide teams with the contextual cues necessary for facilitating  
1207 dialogues with other teams when further knowledge and resources are needed, aiding collaboration  
1208 across the distinct systems. As a future research direction, it would also be interesting to investigate  
1209 the contextual cues necessary for a range of incidents reported to help desk staff, and attempting  
1210 to address limitations in current systems where break-downs in DCog prevent suitable incident  
1211 response, not just regarding security incidents.

1212

### 1213 **7.4 Best practice, is it helping?**

1214 The University made great efforts to align their structure and practices with industry best practice.  
1215 They leaned heavily on a framework called ITIL, which provides best practices for IS service  
1216 management designed to help organisations align practices with business needs. ITIL defines IS  
1217 services from the customers' point of view to satisfy their needs and to bring value to them without  
1218 ownership of risks [52]. It provides guidelines rather than rules as it determines "what should be  
1219 done" as opposed to "how it should be done". Therefore, the implementation of ITIL is different  
1220 between organisations.

1221 ITIL identifies sets of activities, called processes, that respond to a specific trigger to accomplish  
1222 specific objectives. The workflow discussed in this paper imply two main processes for handling  
1223 phishing attacks, namely "Incident Management" and "Problem Management". Incidents are defined  
1224 as an unplanned event to the service, such as the daily calls to the Help Desk (i.e. queue for printer  
1225

1226 help, password resets etc.) [52]. In our case, each phishing report, or request for guidance is an  
1227 example of an incident. The same call can potentially evolve into a “problem” once the Help Desk  
1228 receives several different calls regarding that specific attack, which in our case represents the  
1229 individual phishing campaign. The calls are now considered more critical than the incidents as the  
1230 impact to the system is greater than one off attacks; and thus, as discussed in Section 6.1.2, should  
1231 be treated differently. However, given the difficulty in identifying repeating incidents [44] and the  
1232 need for a quick reaction, current practices may allow for damage to occur when waiting for more  
1233 phishing calls before escalation. Another method for triggering a problem should be considered  
1234 when dealing with phishing to minimise the impact of incidents that cannot be prevented.

1235 The studied University made an effort to develop an environment that encourages phishing  
1236 reporting in line with guidance [67]. Much of this effort was dictated by ITIL, such as providing  
1237 details for how to contact the Help Desk on their website, allowing users to contact them via several  
1238 potential communication channels, providing customised feedback to users, and aiming to quickly  
1239 resolve users’ queries. The Security team also ran campaigns to educate staff and students about  
1240 phishing and encourage them to report it. However, the number of phishing reports the Help Desk  
1241 receives is relatively low considering the size of the University. Similar to previous research [10], we  
1242 noticed some of the reports come to the Help Desk because reporters wanted help in determining  
1243 the safety of an email, they already suspected that the sender was spoofed or that the email looked  
1244 too sophisticated.

1245 While prior research attempted to explore the effectiveness of ITIL in general IS operations [68,  
1246 73], our observations imply that the ITIL framework might not fully fit the workflow of phishing  
1247 handling. Further research is needed to understand how organisations handle phishing while  
1248 adhering to the ITIL framework and what barriers might arise from using such a framework.

## 1249 8 Limitations

1251 The case study may be suggestive of the situation of organisations but generalising the results  
1252 requires further research. The case study looked at an academic institution that likely differs  
1253 from other sectors. Universities also have a wide range of internal structures, so while this case is  
1254 interesting and instructive, other Universities likely have different structures and may be impacted  
1255 by things like their size and how centralised IT services are. However, we argue that many aspects  
1256 of this case have similarities with other organisations; for example, using ticketing systems is quite  
1257 common across sectors. We therefore believe that many of our high-level findings may be useful in  
1258 future work around how to better support how IT handles phishing reports.

1259 Both interviews and observations were used to collect data. While observations allow the  
1260 researcher to observe work practices directly, interviews with participants are complimentary,  
1261 gaining retrospective accounts of events that have happened across a wide time frame and validating  
1262 observations made. That said, retrospective interviews are known to be somewhat biased towards  
1263 memorable events such as particularly impactful phishing campaigns, which may have caused us  
1264 to over-sample these events. Interviews also suffer from social desirability bias where participants  
1265 may provide a version that does not fully reflect reality. To partially counter this issue, we asked  
1266 every team about what they thought the other teams do and detail communication between them.  
1267 We also attempted to provide validation of interviews through analysis of calls in the ticketing  
1268 system. Due to limited access to the system and the use of other communication channels, we were  
1269 not able to see all interactions between teams.

## 1270 9 Conclusions

1272 We explored the process of handling phishing incidents in a large University using a combination  
1273 of interviews and observations. The University uses industry best practices aligned with ITIL to  
1274

efficiently react to and prioritise incidents based on their potential impact. One observation is that large phishing campaigns can result in many reports which overwhelm Help Desk staff, making it challenging for them to respond individually to each report. We also find that the Help Desk operates as a kind of report triage, shielding third line staff, such as those that manage the email relays, from being inundated by reports that may not have the data they need to take action. Similar to earlier works [46], we also find that communication among staff in different teams is a key aspect of coordinating phishing attack mitigation. We also believe that managing phishing reports is an example of distributed cognition, where the different teams work together through the ticketing system to coordinate solving a multi-system problem. We believe that although it is impossible to fully automate phishing response, there is potential to better support IT staff through improved tools that allow them to handle the scale and complexity of phishing attacks.

## References

- [1] June 10 and Joe Hertvik. 2020. Who Uses ITIL in 2020? <https://www.bmc.com/blogs/who-uses-til/>. Accessed June. 2019.
- [2] Atif Ahmad, Justin Hadgkiss, and Anthonie B. Ruighaver. 2012. Incident response teams - Challenges in supporting the organisational security function. *Computer Security* 31, 5 (2012), 643–652. <https://doi.org/10.1016/j.cose.2012.04.001>
- [3] Abdullah M. Alnajim and Malcolm Munro. 2009. An Anti-Phishing Approach that Uses Training Intervention for Phishing Websites Detection. In *Sixth International Conference on Information Technology: New Generations, ITNG*. IEEE Computer Society, Las Vegas, Nevada, USA, 405–410. <https://doi.org/10.1109/itng.2009.109>
- [4] Kholoud Althobaiti, Kami Vaniea, and Serena Zheng. 2018. Faheem: Explaining URLs to people using a Slack bot. In *2018 Symposium on Digital Behaviour Intervention for Cyber Security (AISB)*. University of Liverpool, Liverpool, UK, 1–8. <http://aisb2018.csc.liv.ac.uk/PROCEEDINGS%20AISB2018/Digital%20Behaviour%20Interventions%20for%20CyberSecurity%20-%20AISB2018.pdf>
- [5] Nalin Asanka Gamagedara Arachchilage, Steve Love, and Konstantin Beznosov. 2016. Phishing threat avoidance behaviour: An empirical investigation. *Computers in Human Behavior* 60 (2016), 185–197. <https://doi.org/10.1016/j.chb.2016.02.065>
- [6] Rob Barrett, Paul P. Maglio, Eser Kandogan, and John H. Bailey. 2005. Usable autonomic computing systems: The system administrators' perspective. *Advanced Engineering Informatics* 19, 3 (2005), 213–221. <https://doi.org/10.1016/j.aei.2004.11.001>
- [7] SparkCMS by Baunfire.com. 2020. Report Phishing. <https://education.apwg.org/report-cybercrime/>. Accessed 2019.
- [8] David Botta, Kasia Muldner, Kirstie Hawkey, and Konstantin Beznosov. 2011. Toward understanding distributed cognition in IT security management: the role of cues and norms. *Cognition, Technology & Work* 13, 2 (2011), 121–134. <https://doi.org/10.1007/s10111-010-0159-y>
- [9] David Botta, Rodrigo Werlinger, André Gagné, Konstantin Beznosov, Lee Iverson, Sidney S. Fels, and Brian D. Fisher. 2007. Towards understanding IT security professionals and their tools. In *Proceedings of the 3rd Symposium on US Able Privacy and Security, SOUPS*, Vol. 229. ACM, Pittsburgh, Pennsylvania, USA, 100–111. <https://doi.org/10.1145/1280680.1280693>
- [10] Pavlo Burda, Luca Allodi, and Nicol Zannone. 2020. Don't Forget the Human : a Crowdsourced Approach to Automate Response and Containment Against Spear Phishing Attacks. In *European Symposium on Security and Privacy*. IEEE, Virtual Conference, 6.
- [11] J.S Busby. 2001. Error and distributed cognition in design. *Design Studies* 22, 3 (2001), 233–254. [https://doi.org/10.1016/S0142-694X\(00\)00028-4](https://doi.org/10.1016/S0142-694X(00)00028-4)
- [12] Gamze Canova, Melanie Volkamer, Clemens Bergmann, and Benjamin Reinheimer. 2015. NoPhish App Evaluation: Lab and Retention Study. In *Internet Society, 8 February 2015 (Usec '15, Vol. 453)*. The Internet Society, San Diego, CA, USA, 1–10. <http://dx.doi.org/10.14722/usec.2015.23009>
- [13] Deanna D. Caputo, Shari Lawrence Pfleeger, Jesse D. Freeman, and M. Eric Johnson. 2014. Going Spear Phishing: Exploring Embedded Training and Awareness. *IEEE Security & Privacy* 12, 1 (2014), 28–38. <https://doi.org/10.1109/msp.2013.106>
- [14] Christopher L. Carr, Patrick J. Bateman, and Saral J. Navlakha. 2008. They Call for Help, But Don't Always Listen: The Development of the User-Help Desk Knowledge Application Model. In *Learning from the past & charting the future of the discipline. 14th Americas Conference on Information Systems, AMCIS*. Association for Information Systems, Toronto, Ontario, Canada, 387. <http://aisel.aisnet.org/amcis2008/387>

- 1324 [15] Sonia Chiasson, PC van Oorschot, and Robert Biddle. 2007. Even experts deserve USAbLe security: Design guidelines  
 1325 for security management systems. In *SOUPS Workshop on USAbLe IT Security Management (USM)*. Citeseer, Pittsburgh,  
 1326 PA, USA, 1–4.
- 1327 [16] Asaf Cidon, Lior Gavish, Itay Bleier, Nadia Korshun, Marco Schweighauser, and Alexey Tsitkin. 2019. High Precision  
 1328 Detection of Business Email Compromise. In *28th USENIX Security Symposium, USENIX Security*. USENIX Association,  
 1329 Santa Clara, CA, USA, 1291–1307. <https://www.usenix.org/conference/usenixsecurity19/presentation/cidon>
- 1330 [17] Cobit 2020. COBIT: Control Objectives for Information Technologies. <https://www.isaca.org/resources/cobit>. Accessed  
 1331 10 Oct. 2020.
- 1332 [18] Matthew J. Conlon. 2007. Overhaul your helpdesk ticketing system. In *Proceedings of the 35th Annual ACM SIGUCCS  
 1333 Conference on User Services*. ACM, Orlando, Florida, USA, 37–40. <https://doi.org/10.1145/1294046.1294056>
- 1334 [19] Rachael Cottam, Jeff Goff, and Peter Nguyen. 2012. Extending the centralized helpdesk functionality to improve  
 1335 decentralized support. In *ACM SIGUCCS Annual Conference, SIGUCCS '12*. ACM, Memphis, TN, USA, 153–156.  
 1336 <https://doi.org/10.1145/2382456.2382493>
- 1337 [20] Lynne M. Coventry and T. B. Kane. 1993. The automation of helpdesks. In *Proceedings of the 1st International Workshop  
 1338 on Intelligent User Interfaces, IUI*. ACM, Orlando, Florida, USA, 219–222. <https://doi.org/10.1145/169891.169991>
- 1339 [21] Albese Demjaha, Tristan Caulfield, M. Angela Sasse, and David J. Pym. 2019. 2 Fast 2 Secure: A Case Study of  
 1340 Post-Breach Security Changes. In *European Symposium on Security and Privacy Workshops, EuroS&P Workshops*. IEEE,  
 1341 Stockholm, Sweden, 192–201. <https://doi.org/10.1109/EuroSPW.2019.00028>
- 1342 [22] Media & Sport Department for Digital, Culture. 2020. *Official Statistics Cyber Security Breaches Survey 2020– Chapter  
 1343 5: Incidence and impact of breaches or attacks*. Technical Report. National Cyber Security Centre. Also available  
 1344 as [https://www.gov.uk/government/publications/cyber-security-breaches-survey-2020/cyber-security-breaches-](https://www.gov.uk/government/publications/cyber-security-breaches-survey-2020/cyber-security-breaches-survey-2020)  
 1345 [survey-2020](https://www.gov.uk/government/publications/cyber-security-breaches-survey-2020/cyber-security-breaches-survey-2020). Accessed Jan. 2021.
- 1346 [23] Rachna Dhamija, J. D. Tygar, and Marti A. Hearst. 2006. Why phishing works. In *Proceedings of the Conference on Human  
 1347 Factors in Computing Systems, CHI*. ACM, Montréal, Québec, Canada, 581–590. <https://doi.org/10.1145/1124772.1124861>
- 1348 [24] Sevtap Duman, Kubra Kalkan-Cakmakci, Manuel Egele, William K. Robertson, and Engin Kirda. 2016. EmailProfiler:  
 1349 Spearphishing Filtering with Header and Stylometric Features of Emails. In *40th Annual Computer Software and  
 1350 Applications Conference, COMPSAC*. IEEE Computer Society, tlanda, GA, USA, 408–416. [https://doi.org/10.1109/  
 1351 compmac.2016.105](https://doi.org/10.1109/compsac.2016.105)
- 1352 [25] Edwin Donald Fraunstein and Rossouw von Solms. 2009. Phishing: How an Organization can Protect Itself. In  
 1353 *Information Security South Africa Conference 2009, School of Tourism & Hospitality, Proceedings ISSA2009*. ISSA, Pretoria,  
 1354 South Africa, University of Johannesburg, Johannesburg, South Africa, 253–268. [http://icsa.cs.up.ac.za/issa/2009/  
 1355 Proceedings/Full/12%5FPaper.pdf](http://icsa.cs.up.ac.za/issa/2009/Proceedings/Full/12%5FPaper.pdf)
- 1356 [26] Edwin Donald Fraunstein and Rossouw von Solms. 2013. An Enterprise Anti-phishing Framework. In *Information  
 1357 Assurance and Security Education and Training - 8th IFIP WG 11.8 World Conference on Information Security Education,  
 1358 WISE 8, Proceedings (IFIP Advances in Information and Communication Technology, Vol. 406)*. Springer, Auckland, New  
 1359 Zealand, 196–203. [https://doi.org/10.1007/978-3-642-39377-8\\_22](https://doi.org/10.1007/978-3-642-39377-8_22)
- 1360 [27] George Grispos, William Bradley Glisson, David Bourrie, Tim Storer, and Stacy Miller. 2017. Security Incident Recog-  
 1361 nition and Reporting (SIRR): An Industrial Perspective. In *23rd Americas Conference on Information Systems, AMCIS*.  
 1362 Association for Information Systems, Boston, MA, USA, 1–10. [http://aisel.aisnet.org/amcis2017/InformationSystems/  
 1363 Presentations/15](http://aisel.aisnet.org/amcis2017/InformationSystems/Presentations/15)
- 1364 [28] Eben M. Haber and Eser Kandogan. 2007. Security administrators: A breed apart. In *SOUPS Workshop on USAbLe IT  
 1365 Security Management (USM)*. Citeseer, Pittsburgh, PA, USA, 3–6.
- 1366 [29] Christine Halverson, Thomas Erickson, and Mark S. Ackerman. 2004. Behind the help desk: evolution of a knowledge  
 1367 management system in a large organization. In *Proceedings of the ACM Conference on Computer Supported Cooperative  
 1368 Work, CSCW*. ACM, Chicago, Illinois, USA, 304–313. <https://doi.org/10.1145/1031607.1031657>
- 1369 [30] Rand Abu Hammour, Yousef Al Gharaibeh, Malik Qasaimeh, and Raad S. Al-Qassas. 2018. The status of information  
 1370 security systems in banking sector from social engineering perspective. In *Proceedings of the Second International  
 1371 Conference on Data Science, E-Learning and Information Systems, DATA*. ACM, Dubai, UAE, 14:1–14:7. [https://doi.org/  
 1372 10.1145/3368691.3368705](https://doi.org/10.1145/3368691.3368705)
- 1373 [31] Ying He and Chris Johnson. 2017. Challenges of information security incident learning: An industrial case study in a  
 1374 Chinese healthcare organization. *Informatics for Health and Social Care* 42, 4 (2017), 393–408. [https://doi.org/10.1080/  
 1375 17538157.2016.1255629](https://doi.org/10.1080/17538157.2016.1255629) arXiv:<https://doi.org/10.1080/17538157.2016.1255629> Pmid: 28068150.
- 1376 [32] Grant Ho, Asaf Cidon, Lior Gavish, Marco Schweighauser, Vern Paxson, Stefan Savage, Geoffrey M. Voelker, and  
 1377 David A. Wagner. 2019. Detecting and Characterizing Lateral Phishing at Scale. In *28th USENIX Security Symposium*.  
 1378 USENIX Association, Santa Clara, CA, USA, 1273–1290. [https://www.usenix.org/conference/usenixsecurity19/  
 1379 presentation/ho](https://www.usenix.org/conference/usenixsecurity19/presentation/ho)



- 1373 [33] Grant Ho, Aashish Sharma, Mobin Javed, Vern Paxson, and David A. Wagner. 2017. Detecting Credential Spearphishing  
1374 in Enterprise Settings. In *26th USENIX Security Symposium, USENIX Security*. USENIX Association, Vancouver, BC,  
1375 Canada, 469–485. <https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/ho>
- 1376 [34] James D. Hollan, Edwin Hutchins, and David Kirsh. 2000. Distributed cognition: toward a new foundation for  
1377 human-computer interaction research. *ACM Transactions on Computer-Human Interaction* 7, 2 (2000), 174–196.  
<https://doi.org/10.1145/353485.353487>
- 1378 [35] Karen Holtzblatt and Sandra Jones. 1995. Conducting and Analyzing a Contextual Interview (Excerpt). In *Readings in*  
1379 *Human-Computer Interaction*, RONALD M. BAECKER, JONATHAN GRUDIN, WILLIAM A.S. BUXTON, and SAUL  
1380 GREENBERG (Eds.). Morgan Kaufmann, San Francisco, CA, USA, 241–253. [https://doi.org/10.1016/B978-0-08-051574-](https://doi.org/10.1016/B978-0-08-051574-8.50028-5)  
1381 [8.50028-5](https://doi.org/10.1016/B978-0-08-051574-8.50028-5)
- 1382 [36] Homeland Phishing Reporting 2020. Report Phishing Sites. <https://www.us-cert.gov/report-phishing>. Accessed May,  
1383 2020.
- 1384 [37] Jason Hong. 2012. The state of phishing attacks. *Commun. ACM* 55, 1 (2012), 74–81. [https://doi.org/10.1145/2063176.](https://doi.org/10.1145/2063176.2063197)  
1385 [2063197](https://doi.org/10.1145/2063176.2063197)
- 1386 [38] Martin Husák and Jakub Cegan. 2014. PhiGARo: Automatic Phishing Detection and Incident Response Framework.  
1387 In *Ninth International Conference on Availability, Reliability and Security, ARES*. IEEE Computer Society, Fribourg,  
1388 Switzerland, 295–302. <https://doi.org/10.1109/ares.2014.46>
- 1389 [39] Edwin Hutchins. 1991. The social organization of distributed cognition. In *Perspectives on socially shared cognition*.  
1390 American Psychological Association, 283–307. <https://doi.org/10.1037/10096-012>
- 1391 [40] Suhaila Ismail, Arniyati Ahmad, and Mohd Afizi Mohd Shukran. 2011. New method of forensic computing in a small  
1392 organization. *Australian Journal of Basic and Applied Sciences* 5, 9 (2011), 2019e25.
- 1393 [41] IT governance training 2020. Phishing Staff Awareness E-Learning Course. [https://www.itgovernance.co.uk/shop/](https://www.itgovernance.co.uk/shop/product/phishing-staff-awareness-e-learning-course)  
1394 [product/phishing-staff-awareness-e-learning-course](https://www.itgovernance.co.uk/shop/product/phishing-staff-awareness-e-learning-course). Accessed May, 2020.
- 1395 [42] Itil 2021. ITIL - IT service management. <https://www.axelos.com/best-practice-solutions/itil>. Accessed 10 Oct. 2020.
- 1396 [43] K. Jansson and Rossouw von Solms. 2013. Phishing for phishing awareness. *Behaviour and Information Technology*  
1397 32, 6 (2013), 584–593. <https://doi.org/10.1080/0144929x.2011.632650>
- 1398 [44] M. Jäntti. 2012. Examining Challenges in IT Service Desk System and Processes: A Case Study. In *The Seventh*  
1399 *International Conference on Systems(ICONs)*. 105–108.
- 1400 [45] Matthew L. Jensen, Alexandra Durcikova, and Ryan T. Wright. January 4-7, 2017. Combating Phishing Attacks: A  
1401 Knowledge Management Approach. In *50th Hawaii International Conference on System Sciences, HICSS*. ScholarSpace /  
1402 AIS Electronic Library (AISeL), Hilton Waikoloa Village, Hawaii, USA, 1–10. <http://hdl.handle.net/10125/41681>
- 1403 [46] Eser Kandogan, Paul P. Maglio, Eben M. Haber, and John Bailey. 2012. *Taming information technology: Lessons from*  
1404 *studies of system administrators*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780195374124.001.0001>
- 1405 [47] Wayne D. Kearney and Hennie A. Kruger. 2013. Phishing and Organisational Learning. In *Security and Privacy*  
1406 *Protection in Information Processing Systems - 28th IFIP*, Vol. 405. Springer, Auckland, New Zealand, 379–390. [https://doi.org/10.1007/978-3-642-39218-4\\_28](https://doi.org/10.1007/978-3-642-39218-4_28)
- 1407 [48] Mahmoud Khonji, Youssef Iraqi, and Andrew Jones. 2011. Mitigation of spear phishing attacks: A Content-based  
1408 Authorship Identification framework. In *6th International Conference for Internet Technology and Secured Transactions,*  
1409 *ICITST*. IEEE, Abu Dhabi, UAE, 416–421. <http://ieeexplore.ieee.org/document/6148475/>
- 1410 [49] Sabina Kleitman, Marvin K. H. Law, and Judy Kay. 2018. It’s the deceiver and the receiver: Individual differences in  
1411 phishing susceptibility and false positives with item profiling. *Plos One* 13, 10 (10 2018), 1–29. <https://doi.org/10.1371/journal.pone.0205089>
- 1412 [50] Erka Koivunen. 2010. "Why Wasn't I Notified?": Information Security Incident Reporting Demystified. In *Information*  
1413 *Security Technology for Applications - 15th Nordic Conference on Secure IT Systems, NordSec, Revised Selected Papers*  
1414 *(Lecture Notes in Computer Science, Vol. 7127)*. Springer, Espoo, Finland, 55–70. [https://doi.org/10.1007/978-3-642-](https://doi.org/10.1007/978-3-642-27937-9_5)  
1415 [27937-9\\_5](https://doi.org/10.1007/978-3-642-27937-9_5)
- 1416 [51] Faris Bugra Kokulu, Ananta Soneji, Tiffany Bao, Yan Shoshitaishvili, Ziming Zhao, Adam Doupe, and Gail-Joon Ahn.  
1417 2019. Matched and Mismatched SOCs: A Qualitative Study on Security Operations Center Issues. In *Proceedings of*  
1418 *the 2019 ACM SIGSAC Conference on Computer and Communications Security, CCS*. ACM, London, UK, 1955–1970.  
1419 <https://doi.org/10.1145/3319535.3354239>
- 1420 [52] Abhinav Krishna Kaiser. 2017. *Become ITIL Foundation Certified in 7 Days*. Vol. 1st edition. Apress. 280 pages.  
1421 <https://doi.org/10.1007/978-1-4842-2164-8>
- [53] Ponnurangam Kumaraguru, Yong Rhee, Alessandro Acquisti, Lorrie Faith Cranor, Jason I. Hong, and Elizabeth Nunge. 2007. Protecting people from phishing: the design and evaluation of an embedded training email system. In *Proceedings of the 2007 Conference on Human Factors in Computing Systems, CHI*. ACM, San Jose, California, USA, 905–914. <https://doi.org/10.1145/1240624.1240760>



- 1422 [54] Ponnuram Kumaraguru, Steve Sheng, Alessandro Acquisti, Lorrie Faith Cranor, and Jason I. Hong. 2010. Teaching  
1423 Johnny not to fall for phish. *ACM Transactions on Internet Technology* 10, 2 (2010), 7:1–7:31. [https://doi.org/10.1145/](https://doi.org/10.1145/1754393.1754396)  
1424 1754393.1754396
- 1425 [55] Marshall A Kuypers, Thomas Maillart, and Elisabeth Paté-Cornell. 2016. An empirical analysis of cyber security  
1426 incidents at a large organization. *Department of Management Science and Engineering, Stanford University, School of*  
1427 *Information, UC Berkeley*, [http://fsi.stanford.edu/sites/default/files/kuypersweis\\_v7.pdf](http://fsi.stanford.edu/sites/default/files/kuypersweis_v7.pdf), accessed July 30 (2016), 1–22.
- 1428 [56] Youngsun Kwak, Seyoung Lee, Amanda Damiano, and Arun Vishwanath. 2020. Why do users not report spear  
1429 phishing emails? *Telematics Informatics* 48 (2020), 101343. <https://doi.org/10.1016/j.tele.2020.101343>
- 1430 [57] E.E.H. Lastdrager, Pieter H. Hartel, and Marianne Junger. 2015. Apate: Anti-Phishing Analysing and Triaging  
1431 Environment (Poster). In *36th IEEE Symposium on Security and Privacy*. IEEE Computer Society, United States, 2.
- 1432 [58] Christina Lekati. 2018. Complexities in Investigating Cases of Social Engineering: How Reverse Engineering and  
1433 Profiling can Assist in the Collection of Evidence. In *11th International Conference on IT Security Incident Management*  
1434 *& IT Forensics (IMF)*. IEEE, Hamburg, Germany, 107–109. <https://doi.org/10.1109/imf.2018.00015>
- 1435 [59] IsecT Ltd. 2013. ISO/IEC 27002:2013 - Information technology - Security techniques - Code of practice for information  
1436 security controls (second edition). <https://www.iso27001security.com/html/27002.html>. Accessed May. 2020.
- 1437 [60] Paul P. Maglio, Eser Kandogan, and Eben Haber. 2003. Distributed Cognition and Joint Activity in Collaborative  
1438 Problem Solving. In *Annual Meeting of the Cognitive Science Society*, Vol. 25. ACM, 758–763. [https://doi.org/10.1007/978-](https://doi.org/10.1007/978-1-84628-901-9_6)  
1439 1-84628-901-9\_6
- 1440 [61] Thomas Maillart, Mingyi Zhao, Jens Grossklags, and John Chuang. 2017. Given enough eyeballs, all bugs are shallow?  
1441 Revisiting Eric Raymond with bug bounty programs. *J. Cybersecur.* 3, 2 (2017), 81–90. [https://doi.org/10.1093/cybsec/](https://doi.org/10.1093/cybsec/tyx008)  
1442 tyx008
- 1443 [62] Lena Mamykina and Catherine G. Wolf. 2000. Evolution of Contact Point: a case study of a help desk and its users. In  
1444 *CSCW Proceeding on the ACM Conference on Computer Supported Cooperative Work*. ACM, Philadelphia, PA, USA,  
1445 41–48. <https://doi.org/10.1145/358916.358950>
- 1446 [63] Stefan Metzger, Wolfgang Hommel, and Helmut Reiser. 2011. Integrated Security Incident Management - Concepts  
1447 and Real-World Experiences. In *Sixth International Conference on IT Security Incident Management and IT Forensics,*  
1448 *IMF*. IEEE Computer Society, Stuttgart, Germany, 107–121. <https://doi.org/10.1109/imf.2011.15>
- 1449 [64] Norshidah Mohamed and Jasber Kaur A. P. Gian Singh. 2012. A Conceptual Framework for Information Technology  
1450 Governance Effectiveness in Private Organizations. *Information Management and Computer Security* 20, 2 (2012),  
1451 88–106. <https://doi.org/10.1108/09685221211235616>
- 1452 [65] Katelin A. Moul. 2019. Avoid Phishing Traps. In *ACM SIGUCCS Annual Conference, SIGUCCS*. ACM, New Orleans, LA,  
1453 USA, 199–208. <https://doi.org/10.1145/3347709.3347774>
- 1454 [66] NCSC 2018. Phishing attacks: dealing with suspicious emails and messages. <http://bit.ly/3fTwQpC>. Accessed Aug.  
1455 2019.
- 1456 [67] NCSC guidance 2018. Phishing attacks: defending your organisation. <https://www.ncsc.gov.uk/guidance/phishing>.  
1457 Accessed Feb. 2019.
- 1458 [68] V R Palilingan and J R Batmetan. 2018. Incident Management in Academic Information System using ITIL Framework.  
1459 *IOP Conference Series: Materials Science and Engineering* 306 (Feb. 2018), 012110. [https://doi.org/10.1088/1757-](https://doi.org/10.1088/1757-899x/306/1/012110)  
1460 899x/306/1/012110
- 1461 [69] Sharoda A. Paul and Madhu C. Reddy. 2010. Understanding together: sensemaking in collaborative information  
1462 seeking. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work, CSCW*. ACM, Savannah,  
1463 Georgia, USA, 321–330. <https://doi.org/10.1145/1718918.1718976>
- 1464 [70] Vihanga Heshan Perera, Amila Nuwan Senarathne, and Lakmal Rupasinghe. 2019. Intelligent SOC Chatbot for  
1465 Security Operation Center. In *International Conference on Advancements in Computing (ICAC)*. 340–345. <https://doi.org/10.1109/ICAC49085.2019.9103388>
- 1466 [71] Mark Perry. 2003. CHAPTER 8 - Distributed Cognition. In *HCI Models, Theories, and Frameworks*. Morgan Kaufmann,  
1467 San Francisco, 193–223. <https://doi.org/10.1016/B978-155860808-5/50008-3>
- 1468 [72] Carol Pollard and Aileen Cater-Steel. 2009. Justifications, Strategies, and Critical Success Factors in Successful ITIL  
1469 Implementations in U.S. and Australian Companies: An Exploratory Study. *Information Systems Management* 26, 2  
1470 (2009), 164–175. <https://doi.org/10.1080/10580530902797540>
- [73] BC Potgieter, JH Botha, and C Lew. 2005. Evidence that use of the ITIL framework is effective. In *18th Annual conference of the national advisory committee on computing qualifications*. Citeseer, Citeseer, Tauranga, NZ, 423–427.
- [74] Robert Prince, Jianwen Su, Hong Tang, and Yonggang Zhao. 1999. The design of an interactive online help desk in the Alexandria Digital Library. In *Proceedings of the international joint conference on Work activities coordination and collaboration*. ACM, San Francisco, California, USA, 217–226. <https://doi.org/10.1145/295665.295692>
- [75] Swapan Purkait. 2012. Phishing counter measures and their effectiveness - literature review. *Information Management & Computer Security* 20, 5 (2012), 382–420. <https://doi.org/10.1108/09685221211286548>

- 1471 [76] PwC. 2015. *Managing cyber risks in an interconnected world: Key findings from the global state of information security*.  
1472 Technical Report. The Global State of Information Security. Also available as [https://www.pwc.com/gx/en/consulting-  
1474 services/information-security-survey/assets/the-global-state-of-information-security-survey-2015.pdf](https://www.pwc.com/gx/en/consulting-<br/>1473 services/information-security-survey/assets/the-global-state-of-information-security-survey-2015.pdf). Accessed  
1475 Aug. 2020.
- 1476 [77] Elissa M. Redmiles, Sean Kross, and Michelle L. Mazurek. 2016. How I Learned to be Secure: a Census-Representative  
1477 Survey of Security Advice Sources and Behavior. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and  
1478 Communications Security, October 24-28, 2016*. ACM, Vienna, Austria, 666–677. <https://doi.org/10.1145/2976749.2978307>
- 1479 [78] ProofPoint Phishing Report. 2019. *State Of The Phish*. Technical Report 1. Proofpoint, Inc. 22 pages. <https://doi.org/10.1038/sj.jp.7211019> Also available as <https://bit.ly/2O1n18O>. Accessed May. 2019.
- 1480 [79] ProofPoint Phishing Report. 2020. *State Of The Phish— An in-depth look at user awareness, vulnerability and resilience*.  
1481 Technical Report 1. Proofpoint, Inc. 48 pages. Also available as [https://www.proofpoint.com/sites/default/files/gtd-  
1483 pfpt-us-tr-state-of-the-phish-2020.pdf](https://www.proofpoint.com/sites/default/files/gtd-<br/>1482 pfpt-us-tr-state-of-the-phish-2020.pdf). Accessed Jan. 2021.
- 1484 [80] M. A. Sasse, S. Brostoff, and D. Weirich. 2001. Transforming the 'Weakest Link' – a Human/Computer Interaction  
1485 Approach to USable and Effective Security. *BT Technology Journal* 19, 3 (Jul. 2001), 122–131. [https://doi.org/10.1023/a:  
1487 1011902718709](https://doi.org/10.1023/a:<br/>1486 1011902718709)
- 1488 [81] Annie Saunders. 2004. Online solutions: looking to the future of knowledgeBase management. In *Proceedings of the  
1489 32nd Annual ACM SIGUCCS Conference on User Services*. ACM, Baltimore, MD, USA, 194–197. [https://doi.org/10.1145/  
1491 1027802.1027848](https://doi.org/10.1145/<br/>1490 1027802.1027848)
- 1492 [82] Ankit Shah, Rajesh Ganesan, Sushil Jajodia, and Hasan Cam. 2019. Understanding Tradeoffs Between Throughput,  
1493 Quality, and Cost of Alert Analysis in a CSOC. *IEEE Transactions on Information Forensics and Security* 14, 5 (2019),  
1494 1155–1170. <https://doi.org/10.1109/tifs.2018.2871744>
- 1495 [83] Nikhil Sharma. 2008. Sensemaking handoff: When and how?. In *People Transforming Information - Information  
1496 Transforming People - Proceedings of the 71st ASIS&T Annual Meeting, ASIST (Proceedings of the Association for  
1497 Information Science and Technology, Vol. 45)*. Wiley, Columbus, OH, USA, 1–12. [https://doi.org/10.1002/meet.2008.  
1499 1450450234](https://doi.org/10.1002/meet.2008.<br/>1498 1450450234)
- 1500 [84] Piya Shedden, Atif Ahmad, and A B. Ruighaver. 2010. Organisational Learning and Incident Response: Promoting  
1501 Effective Learning Through The Incident Response Process. In *Proceedings of the 8th Australian Information Security  
1502 Mangement Conference*. Edith Cowan University, Perth, Australia, 131–142. <https://doi.org/10.4225/75/57b6771734788>
- 1503 [85] Steve Sheng, Bryant Magnien, Ponnurangam Kumaraguru, Alessandro Acquisti, Lorrie Faith Cranor, Jason I. Hong,  
1504 and Elizabeth Nunge. 2007. Anti-Phishing Phil: the design and evaluation of a game that teaches people not to fall  
1505 for phish. In *Proceedings of the 3rd Symposium on USable Privacy and Security, SOUPS, Vol. 229*. ACM, Pittsburgh,  
1506 Pennsylvania, USA, 88–99. <https://doi.org/10.1145/1280680.1280692>
- 1507 [86] Hossein Siadati, Sean Palka, Avi Siegel, and Damon McCoy. 2017. Measuring the Effectiveness of Embedded Phishing  
1508 Exercises. In *10th USENIX Workshop on Cyber Security Experimentation and Test, CSET 2017, August 14, 2017*. USENIX  
1509 Association, Vancouver, BC, Canada, 8. [https://www.usenix.org/conference/cset17/workshop-program/presentation/  
1511 siadatii](https://www.usenix.org/conference/cset17/workshop-program/presentation/<br/>1510 siadatii)
- 1512 [87] Christian J. Sinnett and Tammy Barr. 2004. Building a champagne helpdesk on a beer budget. In *Proceedings of the  
1513 32nd Annual ACM SIGUCCS Conference on User Services*. ACM, Baltimore, MD, USA, 351–353. [https://doi.org/10.1145/  
1515 1027802.1027884](https://doi.org/10.1145/<br/>1514 1027802.1027884)
- 1516 [88] Sukamol Srikwan and Markus Jakobsson. 2008. Using Cartoons to Teach Internet Security. *Cryptologia* 32, 2 (2008),  
1517 137–154. <https://doi.org/10.1080/01611190701743724>
- 1518 [89] Gianluca Stringhini and Olivier Thonnard. 2015. That Ain't You: Blocking Spearphishing Through Behavioral  
1519 Modelling. In *Detection of Intrusions and Malware, and Vulnerability Assessment - 12th International Conference, DIMVA,  
1520 Proceedings (Lecture Notes in Computer Science, Vol. 9148)*. Springer, Milan, Italy, 78–97. [https://doi.org/10.1007/978-3-  
1522 319-20550-2\\_5](https://doi.org/10.1007/978-3-<br/>1521 319-20550-2_5)
- 1523 [90] Tracy Ann Sykes. 2015. Support Structures and Their Impacts on Employee Outcomes: A Longitudinal Field Study of  
1524 an Enterprise System Implementation. *MIS Quarterly* 39, 2 (2015), 437–495. [http://misq.org/support-structures-and-  
1526 their-impacts-on-employee-outcomes-a-longitudinal-field-study-of-an-enterprise-system-implementation.html](http://misq.org/support-structures-and-<br/>1525 their-impacts-on-employee-outcomes-a-longitudinal-field-study-of-an-enterprise-system-implementation.html)
- 1527 [91] Xiaojun Tang and Yuki Todo. 2013. A Study of Service Desk Setup in Implementing IT Service Management in  
1528 Enterprises. *Technology and Investment* 4, 3 (2013), 190–196. <https://doi.org/10.4236/ti.2013.43022>
- 1529 [92] Abtin Refahi Farjadi Tehrani and Faras Zuheir Mustafa Mohamed. 2011. A CBR-based Approach to ITIL-based  
1530 Service Desk. *Journal of Emerging Trends in Computing and Information Sciences* 2, 10 (2011), 476–484. <http://www.doaj.org/doi/jfunc=fulltext{&}aId=868208>
- 1531 [93] Inger Anne Tøndel, Maria B. Line, and Martin Gilje Jaatun. 2014. Information security incident management: Current  
1532 practice as reported in the literature. *Computers and Security* 45 (2014), 42–57. [https://doi.org/10.1016/j.cose.2014.05.  
1534 003](https://doi.org/10.1016/j.cose.2014.05.<br/>1533 003)

- 1520 [94] Amber van der Heijden and Luca Allodi. 2019. Cognitive Triaging of Phishing Attacks. In *28th USENIX Security*  
1521 *Symposium, USENIX Security*. USENIX Association, Santa Clara, CA, USA, 1309–1326. [https://www.usenix.org/  
1522 conference/usenixsecurity19/presentation/van-der-heijden](https://www.usenix.org/conference/usenixsecurity19/presentation/van-der-heijden)
- 1523 [95] LEX S. VAN VELSEN, MICHAËL F. STEEHOUDER, and MENNO D. T. DE JONG. 2007. Evaluation of User Support:  
1524 Factors That Affect User Satisfaction With Helpdesks and Helplines. *IEEE Transactions on Professional Communication*  
1525 50, 3 (2007), 219–231. Issue 3. <https://doi.org/10.1109/tpc.2007.902660>
- 1526 [96] Verizon. 2018. *2018 Data Breach Investigations Report*. Technical Report. Verizon Trademark Services LLC. Also  
1527 available as <https://vz.to/2Rzk8ZW>. Accessed Aug. 2019.
- 1528 [97] Verizon. 2019. *2019 DataEnterprise Phishing Resiliency and Defense Report Breach Investigations Report*. Technical  
1529 Report. Verizon Trademark Services LLC. Also available as [https://enterprise.verizon.com/resources/reports/2019-  
1530 data-breach-investigations-report.pdf](https://enterprise.verizon.com/resources/reports/2019-data-breach-investigations-report.pdf). Accessed Jun. 2020.
- 1531 [98] Verizon. 2020. *2020 Data Breach Investigations Report*. Technical Report. Verizon Trademark Services LLC. Also  
1532 available as <https://enterprise.verizon.com/resources/reports/2020-data-breach-investigations-report.pdf>. Accessed  
1533 Jun. 2020.
- 1534 [99] Melanie Volkamer, Karen Renaud, Benjamin Reinheimer, and Alexandra Kunz. 2017. User experiences of TORPEDO:  
1535 T0oltip-poweRed Phishing Email DetectiOn. *Computer Security* 71 (2017), 100–113. [https://doi.org/10.1016/j.cose.  
1536 2017.02.004](https://doi.org/10.1016/j.cose.2017.02.004)
- 1537 [100] Rick Wash and Molly M. Cooper. 2018. Who Provides Phishing Training?: Facts, Stories, and People Like Me. In  
1538 *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, CHI 2018, April 21-26, 2018*. ACM,  
1539 Montreal, QC, Canada, 492. <https://doi.org/10.1145/3173574.3174066>
- 1540 [101] Rodrigo Werlinger, Kirstie Hawkey, and Konstantin Beznosov. 2008. Security practitioners in context: their activities  
1541 and interactions. In *Extended Abstracts Proceedings of the Conference on Human Factors in Computing Systems, CHI*.  
1542 ACM, Florence, Italy, 3789–3794. <https://doi.org/10.1145/1358628.1358931>
- 1543 [102] Rodrigo Werlinger, Kirstie Hawkey, and Konstantin Beznosov. 2009. An integrated view of human, organizational,  
1544 and technological challenges of IT security management. *Information Management & Computer Security* 17, 1 (2009),  
1545 4–19. <https://doi.org/10.1108/09685220910944722>
- 1546 [103] Rodrigo Werlinger, Kasia Muldner, Kirstie Hawkey, and Konstantin Beznosov. 2010. Preparation, detection, and  
1547 analysis: the diagnostic work of IT security incident response. *Inf. Manag. Comput. Secur.* 18, 1 (2010), 26–42.  
1548 <https://doi.org/10.1108/09685221011035241>
- 1549 [104] Kevin F. White, Wayne G. Lutters, and Anita Komlodi. 2008. Towards virtualizing the helpdesk: assessing the  
1550 relevance of knowledge across distance. In *Proceedings of the 2nd ACM Symposium on Computer Human Interaction*  
1551 *for Management of Information Technology, CHIMIT*. ACM, San Diego, California, USA, 3. [https://doi.org/10.1145/  
1552 1477973.1477977](https://doi.org/10.1145/1477973.1477977)
- 1553 [105] Heidi Wilcox and Maumita Bhattacharya. 2016. A framework to mitigate social engineering through social media  
1554 within the enterprise. In *2016 11th Conference on Industrial Electronics and Applications (ICIEA)*. IEEE, Hefei, China,  
1555 1039–1044. <https://doi.org/10.1109/iciea.2016.7603735>
- 1556 [106] Weining Yang, Aiping Xiong, Jing Chen, Robert W. Proctor, and Ninghui Li. 2017. Use of Phishing Training to Improve  
1557 Security Warning Compliance: Evidence from a Field Experiment. In *Proceedings of the Hot Topics in Science of Security:  
1558 Symposium and Bootcamp, HoTSoS*. ACM, Hanover, MD, USA, 52–61. <https://doi.org/10.1145/3055305.3055310>