

Incident Response Teams in IT Operations Centers

The T-TOCs model of team functionality

Judith Brown · Steven Greenspan · Robert Biddle

the date of receipt and acceptance should be inserted later

Abstract We studied the nature of incident response teams in 7 Operations Centers of varying size and types including service providers, a Security Operations Center, a Data Center, and two military training Operations Centers. All responded to incidents by forming teams. We asked: what is the context of incident response work?, how can we model incident response work?, and what are the implications for tool developers? Activity theory guided our research throughout. Using an ethnographic approach to data collection we shadowed 129 individuals for a total of 250 hours of observations, conducted 38 interviews, and facilitated 11 meetings with executives of Operations Centers. We produced rich descriptions of the work of operators and a model of incident team formation called the Tailor-made Teams in Operations Centers (T-TOCs). We position our results relative to other ethnographic studies and standards in the industry, showing how incident team formation has changed over time. Today's incident response team is ad hoc, i.e., tailor-made to the circumstances, and responsive to changing circumstances. Our model draws parallels between the incident response work of teams and human cognition. We conclude by pointing out that tools for tailor-made teams are in their infancy.

Keywords Ethnography of work · IT Operations Centers · Incident response teams

Judith Brown (corresponding author)
Computer Science & Cognitive Science
Tel.: +1-613-276-4821
E-mail: JudithM.Brown@carleton.ca
Carleton University

Steven Greenspan
CA Strategic Research Labs
E-mail: Steven.Greenspan@ca.com CA Technologies

Robert Biddle
Computer Science & Cognitive Science
E-mail: Robert.Biddle@carleton.ca Carleton University

1 Introduction

Information Technology (IT) Operations Centers are workplaces from which network and platform administrators, service desk staff, and security and service analysts work together to provide network and service continuity, to quickly diagnose and correct problems when they occur, and to help the customer or end user solve technological problems with their tools. While many organizations run their own generic Operations Centers, specialized IT Operations centers include Managed Service Providers (MSPs), Security Operations Centers (SOCs), Network Operations Centers (NOCs), and Command and Control Centers (C3s) that support IT systems for the military. Even Data Centers (DCs) can be seen as a special type of Operations Center, although the critical services they provide are a combination of software (running servers, and so on), hardware (computers, cables, and so on) and infrastructure (continuous power, and so on).

Operations Centers are especially concerned with maintaining IT services. An IT service is the “application of business and technical expertise to enable organizations in the creation, management and optimization of or access to information and business processes” [13]. A *service provider* is an organization or a part of an organization that manages and delivers an IT service or services to customers.

Services require support and monitoring, and when events develop into incidents multiple experts typically become involved. An *incident* is “any event that is not part of the standard operation of a service and causes, or may cause, an interruption to, or reduction in, the quality of that service” [23]. Incidents can originate from any point in the IT environment (software, hardware, or infrastructure). Some examples include: a user cannot receive email, a network circuit is down, a server is down, or an application is running slowly. The goal of those who respond to incidents is to restore normal operations, usually as quickly and as cost-ef-

fectively as possible. Incident response teams are an essential and necessary aspect of IT operations centers because teams provide dependability of services [32], resiliency in the face of new situations, [44] and because the variability and creativity of teams is “the most reliable antidote” to ‘cope with’ (i.e., control and manage) complexity [12]. Our focus is on incident response teams, their formation and structure. This work requires technical skills beyond those of the help desk operator or any individual specialist, and the work is often undertaken within moments of the incident occurring and under severe time constraints.

1.1 Activity Theory

To understand incident response work we drew on *activity theory* which is developed from Vygotsky’s cultural-historical psychology [33][39]. Activity theory is a theory of the human mind and how it develops. Activity theory has been used as a research framework in many domains [9], including the field of human-computer interaction [25]. In activity theory, activities exist to meet human needs.

We used the framework to identify activities and their elements [11]. *Individual activities* are mediated by tools. These include physical tools that act on the material world and psychological tools that are used to transform other people (e.g., persuasive arguments or models) and to direct one’s own actions (e.g., to do lists) *Group activities* are systems of inter-related ‘elements’ that adapt in response to any change in any one of the elements (including tools). While the concepts of individual activity and group activity are well integrated, in this research we focus on group activity.

Application of an activity theoretical framework directed our attention to certain elements of group activity, such as the people, their tools, tacit rules in the workplace, the makeup of groups or teams, and the way work is divided between members of a group. When the division of labor is well-structured, the group is called a team (e.g., an incident response team).

Activity theory also draws attention to the motive (or motives) for the activity (in this case, to fulfil the need for service continuity to communities of users), and the object or focus of the activity. In the case of an incident response team, the objective is transforming the incident into a resolved incident while maintaining or bolstering customer relations. [6] [11].

Activities may experience disturbances or disruptions to their flow. These may indicate ‘tensions’ which are unresolved conflicts that arise for multiple reasons, including conflicting motives, or conflicts between ‘human nature’ and the social or technological environment. The underlying forces that are in conflict, expressed as a tension, drive activities in multiple directions. Resolution of a tension can occur

through many means, one of which is the creation of new tools. Tensions in an activity are common.

Activity theory places a strong emphasis on the social. The *social world* is understood as “interactions with other real people, as well as interactions with the tools other real people have designed and left for others as part of their culture” [31]. In contrast to the social world, the mind is characterized by “basic capabilities such as attention, will, [and] intention” that “through human activity [i.e., social life], develop [into] what Vygotsky called the “higher psychological functions” [31].

Activity theory’s well-developed principles of *internalization* and *externalization* asserts that there is mutual influence between an individual’s internal life and social life. For example, learning in culturally-appropriate ways first happens on the social plane and then is internalized. Further, individuals may develop tools to aid internalized functions, a process called externalization. We will use these principles in two ways: 1) The transformation of tools/technologies (understood in the narrow sense) for incident management will be discussed. 2) We will show that incident response work in operations centers can be modeled as a system of coordinated externalized ‘higher psychological functions’, such as perception, decision-making, providing executive oversight, reflecting and anticipating. Specifically, we create a model detailing which externalized ‘higher psychological functions’ are relevant in incident response work at the team, i.e., group level. In doing so, we will draw an analogy between the work of incident teams and human cognition.

1.2 Research Focus and Questions

In this study we are particularly interested in how teams are structured and enabled by technology to act in dynamic, complex and uncertain environments. We present our field study of organic/natural incident response work at seven operations centers. We aim to shed light on understandings of service provisioning. We hope the model we will present depicting the structure of incident response teams will provide an aid to reasoning about the formation and nature of teams in OCs. In particular we aim to enable discussions about tool support. Our work is unusual in that it covers a great breadth; we look at incident response in typical operation centers, security operation centers, data centers and training environments. As well as being of interest to other researchers, our work is also of relevance to tool support developers, UX designers, and OC managers.

The specific questions that directed our study were:

1. What is the context in which operators form and participate in distributed incident response teams?
2. How can we model incident response work in a way that cuts across Operations Centers of many types?

3. What are the implications for tool developers?

2 Literature Review and Historical Analysis

In activity theory, activities are understood historically, i.e., the focus is on understanding how and why an activity changes over time, and the change or changes that would be fruitful in the future, given the objective of the activity. We therefore have taken this perspective when conducting our literature review. We begin by reviewing the position of key industry standards addressing incident response teams because these standards, which have been published for decades now, aim to capture best practices. Following this, we look at recent ethnographic research conducted in Operators Centers. Here we look for the development of team formation activity in Operation Centers since the release of standards.

In this section we review four standards that capture best practices for service providers.

ITIL & ISO/IEC 20000: ITIL for IT service management emerged in the 1980s in Britain [34]. Initially developed for the British government because of poor IT service quality, it has been accepted by industry and across the globe by large organizations. Newer versions were released in 2001, 2007 and most recently 2011. This standard has contributed to, and is aligned with, ISO/IEC 20000 standards for IT service management [5], also most recently published in 2011. ITIL is primarily process-focused, but of necessity, some roles have been identified. For example, ITIL v1 and v2 described a tiered system of incident escalation that is no longer actively supported in the 2011 standard [4], although the language of tiers (or levels) is still in use. In v2, tier 1, (i.e., ‘1st level support’, frontline, or ‘help-desk’) records, classifies and often resolves events (e.g., password resets). When tier 1 cannot provide support, an event is classified as an incident and is passed to tier 2 (internal technical experts) who may get help from tier 3 (software companies or hardware manufacturers). If tier 2 cannot resolve the incident, the incident is considered major and is passed to an incident manager who manages the incident by “dynamically establishing a team of [pre-designated] IT managers and technical experts” augmented by occasional specialists from other groups in the organization, e.g., applications analysts, service owners or technical specialists. This is a common, though not the sole, interpretation of the tiered escalation process.

CERTs: Security standards provide a different approach to incident management of serious security breaches. In 1998 the CERT Coordination Center of the Software Engineering Institute at Carnegie Mellon University created a standard for Computer Emergency Response Teams (CERTs) [40]. These guidelines were published in 1998 and updated

in 2003. They had a broad mandate covering IT and communications technology. The CERTs often serve regions as large as an entire country. Within CERTs (or within organizations) are Computer Security Incident Response Teams (CSIRTs) who respond to security incidents. CSIRT types include: a coordination team, a corporation team or a technical team [24]. Regardless of type, a CSIRT is a incident team that responds “like a fire department or an Emergency Response (ER) team” [40], meaning that particular trained responders from the CSIRT team are assembled at the time of an incident to respond to it.

ISO/IEC 27000: In 2004, and most recently 2013, ISO/IEC 27000 standards proposed Information Security Incident Response Teams (ISIRTs) [7] [21] [20] [19]. Technical reports TR-18044 (2004) and TR-27035 (2011) provide practical advice on incident response. The ISIRT team is similar to a CSIRT team. ISIRTs are found in medium to large organizations where “a planned approach is essential.” The members are skilled and trusted members of an organization who are responsible for incidents. There is a Point of Contact (PoC) person and the handbook states that “at times this team may be supplemented by external experts, for example, from a CSIRT or CERT” (TR-18044).

NIST: In 2007 and in 2012 NIST published a guide for security incident response teams [15]. In it they described three types of response teams: centralized, distributed, and coordinating, and declared that these may exist fully in-house, be partially outsourced, or be fully outsourced. A team manager pulls together a specific response team from diverse team members [17] [18] in the same way that CSIRT incident response teams are formed (like a fire department or ER team). The other teams in an organization that may be drawn in to an incident response are diverse and include: management, information assurance, IT support, the legal department, public affairs, and facilities management.

All Standards Organizations: Across all of the standards capturing best practices, the concept of a fixed group of experts from which an incident team is formed is the most common. This is essentially CERT’s fire department metaphor. When an incident arises, some of the fire fighters are selected from the larger team to address the response. In IT incident response however, this team is also sometimes augmented by outside experts according to the standards. All standards also highlight the importance of communication; however, while some of the standards emphasize communication within an organization between team members (e.g. ITIL and NIST), others emphasize communication between teams or outside of the organization (e.g., CERT). Also, in ITIL there has been a strong emphasis on escalation of incidents, as reflected in the well-known tiered response mechanism. The concept of escalation is also found indirectly in other standards, e.g., captured by the idea of higher-level coordinating teams.

Standards Organizations		Corresponding Incident Response Team Name	
CERT	Computer Emergency Response Teams	CSIRT	Computer Security Incident Resonse Teams
IEC	International Electronic Technical Commission		
ISO	International Standards Organization	ISIRT	Information Security Incident Response Teams
ITIL	Information Technology Infrastructure Library		Escalation process or Incident-manager-led incident response team
NIST	National Institute of Standards and Technology		Team-Manager-led incident response team

Table 1 Standards organizations and incident response teams abbreviations and names

	OC1	OC2	OC3	OC4	OC5	OC6	OC7
Pseudonym	Finance-MSP	N&S-SP	UMA-SOC	Guard-SOC	NNO-DC	NNO-HD	Flexor-MSOC
Participants observed	34	5	18	30	6	15	20
Participants interviewed	10	n/a	5	8	2	4	9
Avg. yrs. of exp. of interviewees	12.8	n/a	1.6	2.8	15	2.8	9.0
Typical level of education	Diplomas + certifications	Bachelor's degree (tech.)	Masters in Elec. Eng.	Bachelors in Comp. Sci.	Varied	Bachelors degree (tech.)	Comp. Eng. Diploma + Networking

Table 2 Participants in our study

The next part of this review turns to research on operations centers for a scholarly perspective on the state of affairs with respect to incident teams.

2.1 Ethnographic Research on Incident Teams

We reviewed research taking an ethnographic approach to the study of incident response teams. Using publications in the last decade, we organize this review by different types of operating centers and then within each type, historically by date of publication. A number of these studies are of incident response teams within an organization that is not an operations center or does not contain an operations center. We included these because ethnographic studies of IT incident response teams in OCs are relatively rare, and because these are forms of incident response teams not included in our study.

SOCs: Early studies of SOC highlighted the relationship *between* security operation centers in different organizations. In 2007 Möller [30] reported on a coordinating CSIRT in an academic CERT in Germany and its relationship with local University CSIRTs. They stressed the enhanced need for collaboration between security teams that had arisen as a result of the recently established Grid Computing architecture being used to support research across Germany.

In 2009 Wiik et al. [41] [42] [43] studied the workload of operators over time in a coordinating CSIRT. The coordinating CSIRT was comprised of a 3-person fixed team that facilitated the handling of incidents across multiple client

CSIRTs in affiliated organizations. The researchers therefore had second-hand information of the form of affiliated CSIRTs. They described three types: 1) an ad hoc CSIRT where available personnel at the affiliated university handled the incidents; 2) a distributed, cross-faculty CSIRT; and 3) an internal centralized CSIRT consisting of a dedicated team within a university.

In 2011 Botta et al. [3] studied the relationship between security teams *within* a single organization. They conducted 35 interviews of IT professionals responsible for security in 16 organizations. In these organizations, incident teams were created as needed from loose networks of security people who knew each other. No standards were being applied, but there were established practices within the organizations.

In 2012 Ahmad et al. [1] studied two four-person security incident response teams in a large financial organization in which there was an Information Security department. One four-person dedicated team in this department was a high- and low-impact Incident Response Coordination Team, called the Network Incident Response Team (NIRT). In addition to this team, which had an operational focus, a four-person High-impact Incident Response Coordination Team (HIRCT) coordinated all incident responses, focusing on its structure and composition. The HIRCT had the potential to expand to as many as 12 people when they drew in others in the organization. The HIRCT managed the incident and liaised with the NIRT and others (including a service management team) as necessary. A major purpose of the HIRCT was to communicate in non-technical ways.

In 2014 Hove and Tårnes [17] and others [18] [36] studied incident teams in two organizations and one service pro-

vider. They demonstrated new complexities in incident teamwork introduced as a consequence of the use of IT service providers to the marketplace. They discussed main and second line teams at one organization, a team that works with an external service provider team at another organization, and a service provider that had a hierarchy of 3 types of teams: an incident response team (IRT), a Critical Incident Management team (CIM), and an Incident Management Board (IMB). They noted that a focus on preparedness and the importance of the non-technical aspects of information security had become stronger over time in all three organizations they studied, and that internal and external communication was the biggest challenge across all teams.

In 2014 Tøndel et al. [36] reviewed 15 research papers on the more general topic of IT security management and concluded that the nature of incident response was varied, that the essential aspect was communication and collaboration at any of the steps, that communication with the business side of an organization was very important, that tacit knowledge was relied upon heavily, that outsourcing often complicates collaboration for organizations, and that there is very little agreement about who (meaning which roles) should be engaged in incident response.

NOCs: In 2011 Norros et al. [32] conducted 20 interviews at one telecommunications company and provided two high-level models of work at Network Operations Centers (NOCs). One showed basic interactions between the NOC, their customers and the network. The other was a model of the ‘core task demands’ of operators. In their core task demands model they observed that the nature of collaboration was ‘constrained by’ the ‘dynamism of the object [the incident] and [the] fast action required [to address the incident]’ as well as the complexity of operator tools and the network itself. They concluded “preparedness for online problem-solving prompted by unexpected network failures is an integral part of [Communication Network Operations] CNO work” and identified communication and collaboration as challenges that need attention.

Service Management Operations Centers: In 2010 Cusick and Ma studied a team of service providers at Wolters Kluwer where a ‘tier 3’ team had been established. This fixed team consisted of one or more application support developers, a systems engineering representative, and a database management analyst.

In 2012 Greenspan, Brown, and Biddle [16] studied a service provider and discussed the increased need for group problem solving, support for multitasking, and tools for the management of heterogeneous, automated solutions. In 2013 Brown, Greenspan, and Biddle [6] provided a detailed description of a service management provider that had a team of incident response facilitators that coordinated incident response. They and Samaroo [35] also created a graphical formalism for depicting work, and used it to model tensions

between the activities of the operator of a large service management operations center.

Multiple types of Operations Centers: Metzger, Hommel and Reiser [29] compared major service and security incidents at an MSP and a SOC. They noted they were both managed with ticketing systems. With respect to security incident teams they noticed the role of the Security Incident Coordinator who may need to call others including sys admins, users, network admins, customer relations representatives, management, law enforcement, security experts, local infrastructure and service operators, and most importantly, the individual who reported the incident. They stressed the role of the Coordinator, not only in their communication task (internally and with other CSIRTs), but also their coordinating and delegating tasks. They suggested a blended incident response system. In that respect they are joined by Tøndel et al. [36] who also primarily studied security incident response teams, but noted similarities with service incident response teams.

All Centers studied: All in all, researchers exposed much more diversity in team formations than what was apparent in standards that capture best practices. Reality has moved beyond the fire-department metaphor to incident response.

Researchers have described situations where teams are very loose and dispersed [3], where internal teams partner with external teams (such as other security teams or service providers) [30], where there is a hierarchical organization of teams [1] [41] [42] [43] [17] [18] [36], where customer-focused teams work in parallel with operations-focused teams, and where the incident coordinator may be a team, such as a high-level, customer-focused crisis management team spanning an organization[1]; these represent new forms of work not captured by the standards.

3 Research Design

To understand incident management work in OCs, we needed to understand real work practices. We set up a series of field studies involving seven operations centers. Our goal was to study incident response work as it actually occurs.

Our study was designed to have a very high degree of ecological validity due to its exploratory nature. Three researchers participated in the data collection phase. We approached various OCs asking if they would be participants, particularly aiming for diversity in our sample. Within this requirement, we selected OCs opportunistically.

To collect data from each OC we designed a field study to observe operations center work directly because we knew that we would learn much more from observations than we could ever learn in interviews, where accounts of work can be idealized or overly simplified. From observations we wanted to know about implicit rules, the real pace of the work, dis-

turbances to the work, and so on, as these are fundamental elements of an activity as understood by activity theory [11] [38]. We used an observation technique called shadowing [28]. We aimed to observe at each center for 2-4 days. We planned to shadow participants who volunteered for the study for approximately 3 hours.

Observations were primary, but we also planned to interview individuals. Our aim was to develop a broader understanding of incident response work based on a full range of our participants' experiences, and not solely on the events that we happened to observe. We also wanted to develop a deeper understanding of incidents, i.e., we wanted to understand why we observed what we observed, and we wanted to understand the structure behind incidents (the roles of operators, the systems in place for responding to incidents, and so on). To this end, we designed a semi-structured interview, which would take approximately one hour to conduct. We used an interview template based on activity theory to structure and generate our interview questions [10].

Both the shadowing and the interviewing were conducted in the workplace environment. At each center we planned to observe first, then interview, which had the additional advantage of developing a common ground of knowledge, rapport, and trust between interviewer and interviewee. Within each OC we selected a diversity of participants to interview based on what we had learned while observing them.

It was very difficult to gain access to OCs; the major concerns were disruption to the work, skepticism about benefits for the OC, and the violation of their customer's privacy. We spent months building up relationships and establishing trust. Once we were permitted entry, we used direct observation augmented by photos where possible, interviews, and also normal interactions and meetings with a wide variety of operators, managers, and systems analysts working on actual incidents to develop our understanding. Initially we selected participants who worked *in* the OC, although as our study proceeded, we quickly realized that the walls of the OC were a misleading boundary since many individuals who are engaged in the work of deploying or maintaining services work in cubicles or offices elsewhere in the organization. The size of the OC was influenced by the need of the OC to conduct tours with prospective clients. OCs that conducted tours, had larger and flashier OCs and housed more operators and technicians. We therefore began to regularly include people in our study who performed OC functions, but who worked outside of the OC itself, e.g., system operators and OC executives. We did not go so far as to include remote participants including incident responders at remote locations (e.g. India) or vendor incident response participants, or incident responders from customer organizations.

To analyze our data we used Strauss and Corbin's version of Grounded Theory [8], with the goal of creating a model of the structure of incident response teams. This re-

quired us to transcribe interviews and integrate notes and photos in order to code our data. The qualitative analysis tool Atlas.ti facilitated this work.

3.1 Operations Centers and Participants

Seven OCs participated in our study, as shown in Table 1. We sampled a variety of OCs. Two were traditional service providers, two provided IT services in a military context, one was a data center, another a help desk, and one was a managed security service provider. In the summary and rich descriptions below we sometimes refer to 'tiers' to simplify our descriptions and we use these terms as they were described in ITIL v2 (tier 1: help-desk; tier 2: internal experts; and tier 3: external experts).

Finance-MSP's objective: Finance-MSP is a very large service provider with many financial institutions as its customers. Almost 100 different platforms are being monitored and more are constantly being added.

Facilities: It is a large operations center that is an amalgamation of other operations centers which offers business transaction processing for many large financial institutions.

Staff: The operations center was capable of housing 50 people and most seats were occupied during the day. At night a skeleton staff of about 10 people kept the center functional.

N&S-SP's objective: N&S-SP provides services to an academic community including professors, staff and students.

Facilities: A traditional small OC, behind which there is a data center containing important computers for the university. There is a large area around the OC organized into cubicles. A help-desk is located in the library.

Staff: The OC has 3 people in it, but has dozens more systems administrators and network administrators working in cubicles around the NOC/SOC. The OC staff oversee the data center which is co-located, and all of the other computing infrastructure which is remote.

UMA-SOC and Guard-SOC's objective: UMA-SOC and Guard-SOC were a service provider with a combined SOC and NOC function established for training purposes. They provided services to a person playing the role of an end user in an imagined military unit. At both these sites we observed the final part of a real-life military simulation exercise that spanned a month or more. Teams built their own network and set up basic services such as email and word processing. Teams participating in the exercise defended their military unit against an external organization participating in the exercise as an attacker. The aim of the students was to keep their services up and running, so that their military end user was able to work. Their primary task was to be an effective SOC, i.e., to defend the network and services they built. When the exercise was completed the network was de-

commissioned.

Facilities: The students built their network and configured their computers with the mission software in their labs. Typically they had several co-located labs they could work in. The computers in this exercise were completely separate from the campus network.

Staff: Two or three professors were available for advice during the exercise, but did not perform any actions on the computers or network. The student numbers ranged from 20 to about 40 solved problems, formed and led incident response teams, and made changes as required to their system.

NNO-HD's objective: NNO-HD is a help desk for a health care community, including outpatients. The help desk provides support for more than 10 hospitals, and many more clinics. Tens of thousands of people in the hospital environment are clients. This includes, doctors, nurses and business people. Thousands of applications are supported by the OC. The help desk receives over 1000 calls per day from internal customers (e.g., doctors), insurance companies who need technical support, clients who request health information and technical support, and PIN resets for the public web portal.

Facilities: The OC room is divided into three parts, each serving a different part of the hospital community or a different major application. 'Tier 2' (primarily system admins) is located upstairs in the same building.

Staff: The help desk engages about 30-40 employees

NNO-DC's objective: NNO-DC's objective was to maintain and improve the IT infrastructure of the hospital.

Facilities: The data center is down the hall from NNO-HD's help desk. The operation center part had seating for about 10. NNO-DC is the data center for NNO-HD, which maintains 100s of servers, a robotic tape backup systems, electricity, backup generators, and coolers.

Staff: Approximately 8 operators worked in the data center.

Flexor-MSOC's objective: Flexor-MSOC is a moderately large, and rapidly expanding security operations center (SOC), that provides managed security services to clients. The SOC does monitoring for its own organization and for many other enterprises who are its clients.

Facilities: The center was located in a large secure area and housed approximately 50 people including monitoring staff, responding staff, and many others whose work was to set up secure software services for clients. The help desk for the SOC was located remotely, and we did not observe those operators.

Staff: There are about 5 tier 1 operators in the OC, who primarily do monitoring, several operations managers, and 15-20 tier 2 experts in the room, including incident responders who are technical liaisons with customers, deployment, and configuration.

Many more details about these operations centers are provided in Table 1. In total we shadowed 129 individuals totaling 250 hours of observations. The participants varied greatly with respect to their function ranging from tier 1 help desk operators, to systems experts, to developers, managers and executives. Of the 129 participants we shadowed, 38 (30%) were also interviewed. We also organized 11 larger meetings with executives and managers for briefings and discussions and we opportunistically participated in 3 additional tours of other Operations Centers. At two operations centers we had permission to take photos. We transcribed all of our interviews, meeting notes and observation notes and, where relevant, inserted photos into the appropriate points in the observation notes.

4 Results

4.1 Incident Response Workflow

To begin we provide rich descriptions of the structure and workflow of incident response teams at the various sites.

Finance-MSP: At Finance-MSP, the large service provider, there was a significant amount of attention to meeting service level agreements and the implementation of ITIL processes. The operators overall had significant amounts of experience, were primarily self-trained, and had certificates or diplomas or part of a degree.

Division of tasks At this center it was clear who worked at tier 1; this was the group of operators who worked in the front row, closest to the large overhead displays that were arranged prominently at the front of a large room that could accommodate more than 100 operators. Their main skills were people skills, but they had enough technological understanding to triage calls to other operators in other groups in the room. Six other groups in the room primarily operated at tier 2. They provided initial end-user terminal support, platform support, and network support. They (especially the NOC group) often relied on tier 3 hardware vendors and software companies, issuing them tickets as necessary and collaborating as appropriate. Another group of operators ran batch software and provided support for batch software clients in large financial organizations. A final group of operators who sat at the back overseeing the rest of the room functioned solely as incident facilitators and coordinators, and we referred to them as tier 1.5. The incident coordinators were supported by a group whose sole purpose was to contact individuals who needed to be on a call.

Workflow: When incidents escalated to the incident coordinators, the tier system was irrelevant; the incident coordinator assembled an appropriate team or teams to address the incident. Service level agreements effectively ensured an 'all hands on deck approach'.

Our goal is to work incidents relating to client impacting issues to get a client back up and working in under 45 minutes (Mary, team coordinator, Finance-MSP).

Throughout the duration of incident resolution, many of which were “challenging” (Mary) the tier of the operator was not important—more important was the capabilities of the operator or tech support person. Incident teams were broad at first and it was not always clear where the work should begin. Support staff for incident managers contacted people to join the bridge call for the incident. Mary, the team coordinator says,

The work of my support staff frees up my time to focus on asking questions to the teams: Are we looking at the server, are we looking at the database? (Mary)

Incident teams were sometimes structured into subteams.

At times we have had nine breakout sessions off of our main bridge call. (Mary)

Typically there were two teams, one that was customer-focused, and one that was resolution-focused. We were told of one unusual case where an incident team was comprised of 20 subteams. It was not uncommon for team members to be remote. Team members came and left the investigative team as the investigation unrolled; for example, they may have discovered their expertise was not useful (a common situation for the NOC operators who often had to demonstrate it was not a networking problem), because of ‘handovers’ at the end of a shift, or because of higher priority tasks drawing their attention away.

All the network people go to [*88], the commonly held breakout session [on our bridge calls], to determine if the network is an issue. (Mary)

Incidents could be resolved in less than 20 minutes or they could remain active for days.

Tensions: There were many times when there was nothing to do on a particular incident (e.g., when information was being dumped by a remote engineer) and so it was possible to ‘work multiple incidents’ (typically 2 or 3) simultaneously. Other background tasks that happened while incident response was ongoing included: monitoring, project work, tidying of one’s workspace, or learning.

Tools: Incident teams used a collection of tools to manage incidents including ticketing tools, chat windows and bridge calls. The work of the team occurred in a context where it was normal for an operator to be running six ticketing tools, six monitoring tools, have 25 chat windows open, and be participating on a couple of bridge calls.

N&S-SP: At N&S-SP we observed a traditional OC that ensures the continuity of services to an academic community. The role of the SOC side was primarily deployment of

security software, but detection and elimination of malware from the PCs of staff was also part of the service.

Division of tasks Because of the small size of this center, the operator’s role spanned tiers 1 and 2. On the non-security side of the operation center, the two operators were primarily responsible for monitoring (a tier 1 task), but they also had duties in the adjoining data center, such as the installation of some equipment. The tier 1 operator on the SOC side also monitored, but was primarily occupied with deploying security updates. He interacted with customers (individuals responsible for security within the various university units) across the university.

Workflow: Most service-related problems (such as the havoc that occurred when a squirrel partially chewed through a critical cable in a conduit) were resolved by moving services to backup systems and immediately alerting tier 2, the tech support group who were located in nearby cubicles, or alerting third party vendors under contract (tier 3), who would send over an expert or experts. When an incident occurred, one of the operators served as the incident manager, alerting relevant others, overseeing their activity (basically ensuring that there was movement towards incident resolution), and ensuring that the troubling indicators (e.g., red alarms) disappeared.

Tensions: The monitoring tools hid alarms from the operators or made alarms difficult to interpret.

Red alerts appeared in two parts of a networking diagram indicating two separate high-priority problems. Two vendors were called in. Eventually we discovered a squirrel had chewed through a cable cutting off one building from the network. The monitoring system showed when a problem was cleared, but did not help me understand the relationship between the two problems (the common source). (Pierce, alerts and system monitoring, N&S-SP)

Tools: Incidents were rare and email was the primary form of communication for the incident team. An array of large displays were being mounted at the front of the room to help with the monitoring of alarms.

UMA-SOC & Guard-SOC: At UMA-SOC and Guard-SOC we observed students who had set up a network and services and were defending it against an outside SOC who were attacking it as a competitive exercise; this was an intense learning experience for the students. For us it was a great opportunity to see highly trained operators responding to significant and intense attacks, something which is a rarity in SOC’s. The situation the students found themselves in, though unusual relative to established Service Centers, was not unlike what they might find in the field during a military operation. At UMA-SOC we observed graduate students of technology programs and at Guard-SOC we observed undergraduates, but in both cases the attacks were the same. At

these sites there was an equal amount of structure and chaos. They built and then defended a network. Services they provided to their military users included email, web services, an Internet browser, and a file directory.

Division of tasks Individuals were assigned multiple roles, but they moved fluidly into other roles based on the needs of the moment, directed by the individual who formed ad hoc teams around incidents. At UMA-SOC, none of the operators were solely responsible for continuity of service delivery as they also had duties deploying and configuring services; their roles were very broad. At UMA-SOC there was only one incident coordinator in the room at any given time. He assigned individuals to incidents based largely on the demands of the situation, their skills, and their availability; kept track of progress; and moved individuals from one task to another when necessary. The assignments were very fluid.

Workflow: The formation of ad hoc teams happened a lot because the attacks were frequent and priorities changed rapidly.

Anytime an incident was reported or observed, then I was involved. Most of the time it was multitasked. So I could be in the middle of doing some malware analysis and then an incident would happen. It was reactionary. I'd have to stop what I was doing, deal with the incident, talk to everyone who needed to be involved and get them going, get them all on the same page, ... then I could get back to what I was originally doing. (Jack, team coordinator, UMA-SOC)

Tensions: At UMA-SOC, and to a lesser extent, at Guard-SOC, a contributing factor to the chaos was the near absence of even primitive tools for situation awareness within the room. At both sites the students called out to one another, passed notes, or wrote very minimal amounts of information on a whiteboard, such as keywords to highlight ongoing incidents, to keep track of the activity in the room. At both sites there was no history of actions taken, no ticket-writing, and lots of uncertainty across the room with respect to the status of various incidents. Because situation awareness was very poor there was some duplication of work resulting in frustration.

Tools: The tools to organize incident response were minimal (e.g., no chat, no ticketing tool) and would have been completely ineffectual had they not been co-located in a single large area. Face-to-face communication was their primary tool for collaboration.

NNO-DC: At NNO-DC we observed a team of operators running a Data Center where the goal was to ensure continuous delivery of hardware systems within a hospital environment. Although Data Centers are not often compared with Service Centers, the work had many similarities with traditional service centers that provide continuous delivery of software services. The data center was very forward look-

ing; it was constantly being re-built and adapted to meet future demands.

Division of tasks: All operators had similar roles and were more academically qualified than the operators in a service center. One operator functioned as an operational manager and incident coordinator. The operators communicated largely via face-to-face or email. The task of deploying and configuring hardware (e.g., servers) and DC infrastructure (cables, generators and the like) occupied approximately half of their time. The other half was spent monitoring or performing essential background tasks like tape backups of patient data. The incident manager organized the overall activity of the DC; a ticketing tool was used to track incidents.

Tensions: These operators were much more mobile than other operators in this study, but they were not equipped with mobile devices. When they were in the part of the data center that housed important infrastructure equipment and hospital servers, they were away from their workstation and disconnected from monitoring tools and their fellow operators. It was not unusual for operators to work off site at one of the hospital campuses. When we visited, two were responding to remote incidents.

Workflow: Mostly operators worked on non-incident impacting tasks, like checking servers, changing cables, or receiving new equipment at the arrival dock. Alerts (such as rooms overheating or power outages) mainly arrived through software systems and were displayed on large boards at the front of the room. Incidents were infrequent. For this reason they regularly ran drills (e.g., switching over to another generator even though it wasn't necessary) that involved most of the operators. The inability of these operators to respond to a crisis (such as a power failure) could impact the entire hospital and result in loss of lives.

The data center has gone black from times in its entire history. Once it happened accidentally. Another time a loose washer caused a blackout. On these 5 occasions, the generator didn't always kick in, due to maintenance issues so now, once a month we run through drills so everyone knows what to do and everyone works together effectively. (Jarvis, data center supervisor, NNO-DC)

During these exercises the incident teams relied heavily on checklists to organize their work.

Tools: Email was relied upon heavily to communicate with IT staff. A ticketing tool was used to track any incidents. Specialized monitoring tools were used to produce large displays of all of the critical infrastructure systems at the front of the room.

NNO-HD: NNO-HD was a traditional help desk in a large room subdivided into three parts and holding roughly 50 people. The distinguishing characteristic of this help desk was its commitment to quality. Any incident that impacted

the hospital's commitment to quality care was investigated across the hospital. The hospital was also strongly committed to continuous improvement through reflective practices and collaboration across different parts of the organization. A practice that began during a critical incident (standup meetings in the morning with managers of departments) became a regular practice to ensure that issues that crossed organizational boundaries within the hospital were addressed before they impacted the hospital's work.

Division of tasks: All the operators were both tier 1 and tier 2 operators. There were three groups in the room supporting three different types of services.

Workflow: Each operator processed a continuous stream of requests on a daily basis working from a queue of calls. In the afternoon when things were quieter they performed background tasks, such as tuning their system, learning, or checking requests for new accounts to systems. Difficult issues were quickly forwarded to systems operators located several floors above. John was one of those analysts.

There are three back-end systems we are responsible for. ... People tend to send in tickets in a variety of ways, mostly through the ticketing tool, but people also just contact me directly through IM or email. Those can be anything, from being locked out of an account, to highly complex problems where data is not making it to its endpoint. ... You have to dig in and investigate in the code. ... Priorities are clinical systems that can affect patient care, or things that are very important to know like patients coming in, or changing rooms. You just have to prioritize from experience. There are 8 other people on my team and they're always helpful. Sometimes I might not be looking at a problem in the right way—I've read something incorrectly, it can be as simple as that. Sometimes one of the team members is the developer for the code, and that can be very helpful. I can be working on 20 different projects at once and it can be confusing.

Tensions: When incidents were escalated to operations/systems engineers, the systems engineers complained that the ticketing tools, which were the primary tool of help desk operators, were not designed with system engineers in mind, e.g., ownership/responsibility stayed with the person (or machine) that generated the ticket and did not transfer to the systems engineer. Systems engineers bounced tickets between themselves and consulted with others until the engineer or engineers able to address the issue was engaged. It was sometimes difficult to find the correct person to address the issue described in the ticket.

Tools: Email and chat were the primary means that the systems engineers used to communicate with each other. There was no identified team coordinator, except when a systems

engineer assumed that responsibility out of a concern for resolving the incident. One person simply took charge if that seemed necessary.

Flexor-MSOC: Flexor-MSOC was a commercial security operation center that offered security services to its many clients. The room was organized to complete deployment and configuration work on security architectures, to ensure continuity of security services for the operation center's clients, and to make an impression on new clients who toured through frequently.

Division of tasks: Several new operators were responsible for monitoring and replying to routine incidents. Senior analysts oversaw the work of the junior analysts, and liaised with customers on complex incidents or projects. Architects of security services (also known as solutions engineers) worked in another room and there was a significant disconnect between them and the operators who deployed, configured and maintained the services they designed. Incident response (in the case of a security issue) was handled by the operation centers' clients who were informed of incidents via tickets or emails, as per their customer agreements.

Workflow: Tool-related incidents (e.g., errors from logging tools) were handled within the room by the engineer or engineers who were available and most equipped to respond it. Typically, the response team was small and assembled rapidly. Managers served as coordinators, but were only involved if required by the circumstances. There was a war room nearby where incident teams could work, but the war room was rarely used. The coordination between those who deployed & configured, and those who monitored was good (accomplished via email), and usually false alarms that were common and attributable to system changes were correctly ignored and not incorrectly identified as an incident. At this site the client relationship was very important.

We built a tool to manage the relationship with our clients. Clients have different requirements for escalation of problems, different preferences, and the end person at the client site we interact with could be a business person, an admin person or a techie. It's hard to get the language right when communicating to such a wide variety of people. We try not to have special per client processes, but at the same time we do not allow cut and paste responses (Bert, SOC manager, Flexor).

Tension: In the case of some incidents, a broader team response to a client's incident would have been useful, e.g., there was a missed opportunity to inform architects that their solutions were challenging to deploy and maintain. For example, because incident response teamwork in the room was often face-to-face it was not easy to include others outside the room (such as the solutions architects).

Tools: When communication was to remote clients, the op-

erators used the method of communication preferred by the client.

Trends across all sites: The number of incidents in an operators' day varied significantly by operation center. Finance-MSP operators typically experienced multiple incidents in a day. N&S-SP might experience one a month. Across all OCs incidents were triggered in a number of ways: 1) Alerting tools and 2) Direct system monitoring, 3) A customer request, 4) An event of an important customer may be escalated to incident status, or 4) And finally, the launch of a new project may introduce significant changes in normal activity and trigger an incident.

Overall we observed that the tier terminology was inconsistently applied within and across sites, and not a prominent part of the work we observed; when a significant incident arose, tiers were unimportant, but incident teams were not. Rather than process dictating how teams are assembled, the needs of the incident determined the makeup of the team. There were a great variety of incident responders and often the most technical individuals (e.g., engineers and developers) were not a part of the team, but many non-technical people were. For example, communication experts or executives liked to be involved when important clients were experiencing difficulties.

We found that incident response teams were commonly structured and composed of multiple teams. One common scenario was a technical team for problem resolution and a customer-focused team to liaise with the customer.

We observed that the role of the team coordinator was very important as they provided guidance around the team formation, facilitated team decision-making, and served in a coordinating capacity.

We found that the way that incident teams were formed and the way they worked was not well supported by existing incident response systems such as bridge calls, chats and ticketing tools. This was especially the case because of the multi-tasking nature of the work. Operators worked multiple incidents at a time on occasion and also had background tasks such as monitoring alerting tools, monitoring platforms, working on special projects, and learning. This made for very complex arrangements of windows on their displays (typically 3 or 4 of these) and the specific bridge calls, chats, and relevant tickets and other artifacts associated with an incident were not conveniently grouped.

4.2 The T-TOCs model: Incident Response Functions

In the descriptions in the previous section, we stressed the workflow around incident response and its ad hoc nature. In this section we present a general model that applies to all the diverse sites described in the previous section and captures and explains the functions performed by ad hoc team mem-

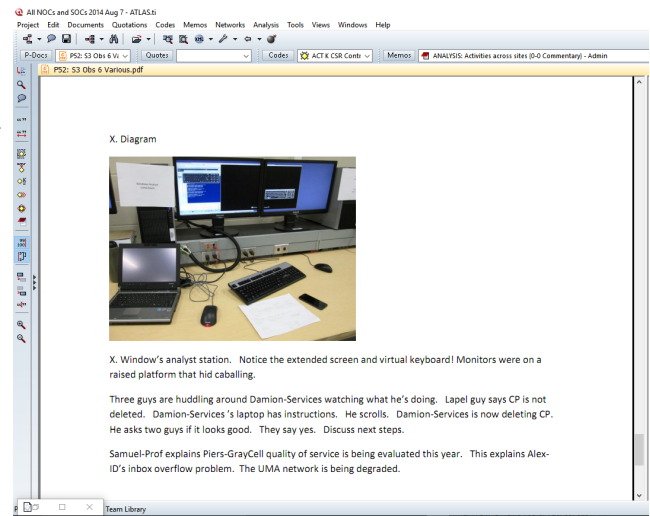


Fig. 1 Screen shot of observational notes in Atlas.ti. At UMA-SOC we were able to take photos so these are embedded in our notes. Depicted is a typical operator workstation. The text describes how three operators huddled around one operators' workstation make and check a change to their system.

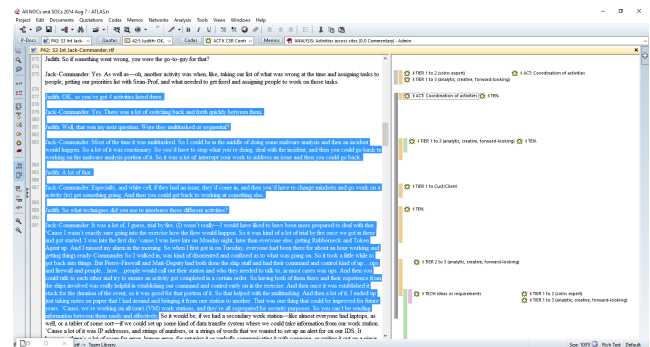


Fig. 2 Screen shot of Transcript in Atlas.ti. A portion of the transcript with Jack-Commander, team coordinator at UMA-SOC. To the right of the transcript a vertical tagged line indicates a quote in the transcript has been identified and coded. There are many codes in this segment of the transcript. One is highlighted. It describes how Jack coordinates the incident response work of others in the SOC.

bers. We also discuss the structure of teams we observed across OCs, emphasizing new forms of work.

The model we produced of team functionality was generated through application of Strauss and Corbin's grounded theory method [8], which we applied to our transcribed observation notes and interview transcripts. See Figure 1. Incidents were our unit of analysis; we knew of these first hand because of our observations of incident teams and indirectly through interviews with incident team members. Like Strauss and Corbin, we declare our analytic frame when applying the grounded theory method. While their's was action theory, ours is activity theory—we used the concept of externalized higher psychological functions from activity theory as an analytic lens. With this analytical lens we developed a set of externalized higher psychological functions involved

in incident response work, which were then grouped into categories. The functions are most clearly expressed as categories in the model we produced. Functions were not readily available from knowledge of the operator's title, which in most cases, were more clearly tied to their compensation rather than their role.

We next report on the steps that led to the creation of our model. We worked together to produce this model, verifying each other's judgments as we went.

The functions we identified emerged as a result of applying the first step of grounded theory, i.e., open coding, to our data. Each time we learned of a function that an operator performed, we coded that. For example, we learned that "learning new platforms and practices" was something that all operators worked on every day. Here's an example where Alice, a technical person for networking issues, explains this function that was triggered by a new release of software.

We received a new console 4 or 5 months ago and we weren't notified that it was going to be integrated with other software products. We all had to learn how to sign into it. We had a huge bridge call because we use that application a lot and we had not been told how to use it. I learned how to sign into it and use it, and then I had to do screen prints and a PowerPoint presentation which I sent to my coworkers so they could learn the same way I did.

We coded similar quotes from all other operators and then, using Atlas.ti to find patterns in our data, we were therefore able to conclude that this was a common function. See Figure 2.

Over all we discovered that team members performed diverse functions. Some of these were common to all the participants. Two of these common functions were continuous: 'Maintaining awareness of the situation' and 'Communicating with end users/clients/vendors or partners inside or outside the organization'. Other common functions that all participants engaged in were background tasks that were performed when time allowed. This included 'Mentoring others on technical matters,' 'Learning about technical matters,' 'Creating reports,' 'Advancing special projects,' and 'Tuning and tinkering with tools.' We do not discuss these important functions further, because of our focus on incident response, but we do note that these functions are fundamental to work in Operations Centers. They capture what operators do when they are not working on an incident, but also capture what is inter-leaved with incident response work.

The other functions that we observed operators perform, or that we became aware of, are listed in Figure 3. This table shows the counts of the numbers of participants and their primary function/role and was an important outcome of the grounded theory process. These functions emerged from the coding process. Sixteen such functions were coded. The first

Functions	Finance-MSP	N&S-SP	UMA-SOC	Guard-SOC	NNO-DC	NNO-HD	Flexor-MSOC
End User			1				
Vendor Rep	2		1				
Ops Manager		1			1	2	3
Team Coordinator	9		1	3			
Front Line Help						6	
Alerts-Sys Monitor	14	3	6	11	4	1	4
Customer Rels							2
Infrastructure			1	6	1	2	4
Technical Support	4			8			
Ops Engineer	3		4				
Trouble-Shooter			2				3
Strategic Manager	2	1	2	2	1	1	
Total Individuals	34	5	18	30	7	12	17

Fig. 3 Count of Primary functions of participants by Site.

Groups	Functions	Finance-MSP	N&S-SP	UMA-SOC	Guard-SOC	NNO-DC	NNO-HD	Flexor-MSOC
Ext Environment	End User			1				
	Customer Rep							
	Vendor Rep	2		1				
	Partner Rep							
Coordination	Ops Manager		1			1	2	3
	Team Coordinator	9		1	3			
Initial Response	Front Line Help						6	
	Alerts-Sys Monitor	14	3	6	11	4	1	4
Ext Communication	Tech Liaison							
	Customer Rels							2
Acting	Infrastructure			1	6	1	2	4
	Technical Support	4			8			
	Ops Engineer	3		4				
	Trouble-Shooter			2				3
Creating	Tool Crafter							
	Architect							
	Engineer							
Reflecting	Strategic Manager	2	1	2	2	1	1	
	Business Manager							

Fig. 4 Count of Primary functions of participants by Site, with the additional of function groups.

on the list is 'Operations' Manager' and the last is 'Business manager.' These functions were very similar to identifiable roles within the operations center and in a final step of our process we tweaked all our code names so that they resembled role names because there appeared to be connection between the two, and this switch made for a more readable and useful model. Long code names like 'Maintaining the operations environment,' which captured what the operator

did (i.e. a function), were replaced by short code names like ‘Ops Engineer.’

The second step of grounded theory is axial coding. In this step higher-level codes, i.e., categories, emerged. These are shown in Figure 4 column 1, and include ‘Coordination’, ‘Initial response’, ‘External communication’, ‘Acting’, ‘Creating’, and ‘Reflecting’. This table shows clearly how the categories are mapped to the functions. For example, the category of ‘Acting’ includes the functions ‘Infrastructure support’, ‘Technical support’, ‘Ops engineer’ and ‘Trouble-shooter’ (a highly skilled person with very broad expertise and relational knowledge).

Selective coding identified the functions associated with incident response work and those that were not. The functions regularly performed by members on incident response teams are Teams Coordinator, Technical Support, Infrastructure Support, Ops Engineering and Technical Liaison with Customers. Functions that are often performed by members on incident response teams are Operational Manager, Customer Relationship Manager, Architect, ‘Development, Services & Solutions Engineer,’ Strategic Manager for the OC, Customer Rep, Vendor Rep, and Partner Rep. Functions rarely performed by members on incident response work were Front Line Help, Alerts & System Monitors, and Tool Crafter.

In the final step of grounded theory we pulled everything together and generated the model shown in Figure 5, which we called the T-TOCs model, standing for Tailor-made teams in Operations Centers. The term tailor-made was carefully chosen to indicate that while the teams are ad hoc, i.e. tailor-made in the moment, the teams are just right, or more than adequate, for the incident, and all necessary functions for investigating the incident are present. The name is also intended to suggest fluidity as new functions are drawn in as required (new operators may be called in), and unnecessary functions are dropped (operators ‘drop off’ the incident).

Working a problem from scratch, . . . there’s no script.
The team has to go off of what we know. (Alice, tech support for networking issues, Finance-MSP)

There are four main parts to the model: 1) The Products, Services, Solutions or Networks maintained by the OC, 2) The External Environment, represented by External Stakeholders including End Users, Customer Representatives, Vendor Representatives and Partner Representatives. 3) The Operations Center and 4) The tools that enable the work of the incident response teams. The products, services, solutions and networks maintained by the operators are carefully positioned to show that these technologies may physically exist in the External Environment, within the Operations Center, or a combination of both. Access to software to change settings or restart it, is generally accomplished using remote software, though occasionally DC operators directly

	Primary	Secondary	Tertiary	Quaternary
Finance-MSP	34	21	6	
N&S-SP	5	3	1	
UMA-SOC	18	9	1	
Guard-SOC	30	23	9	
NNO-DC	7	6	3	1
NNO-HD	12	5		
Flexor-MSOC	17	14	7	5

Fig. 6 Count of Number of functions performed by an individual by Site. Most operators perform two functions, and some more.

act on computing devices. The model is a reminder of the rich ethnographic information summarized in the previous section about the ad hoc nature of teams, and the makeup and structure of teams.

When considering the model, it is important to be aware that the functions did not map to individuals on a one-to-one basis. Figure 6 shows that in most centers individuals fulfilled, on average, 2 functions, one of which was their primary function. It also shows that at least six individuals in the study performed 4 functions. This could obviously be stressful at times.

Sometimes it was difficult, but it was just because I had so many roles going on at the same time. (Jack, team coordinator, customer relationship manager, and trouble-shooter, UMA-SOC)

We also discovered that certain pairs of roles were commonly performed by the same individual Figure 7. 14 individuals whose primary role was Alerts-System Monitoring were also involved in providing Infrastructure Support, 11 were involved in Technical Support, and 10 were also Team coordinators. Conversely, 12 individuals whose primary role was Technical Support were also involved in Alerts-System Monitoring.

5 Discussion

5.1 The gap between standards and incident response work

There are few descriptions of incident response work because operations centers are inaccessible to most researchers, since a considerable amount of trust must be established for studies in these environments. There are also significant barriers to access for tool developers. Traditionally in software

T-TOC: The 'Tailor-made Teams in OCs' model

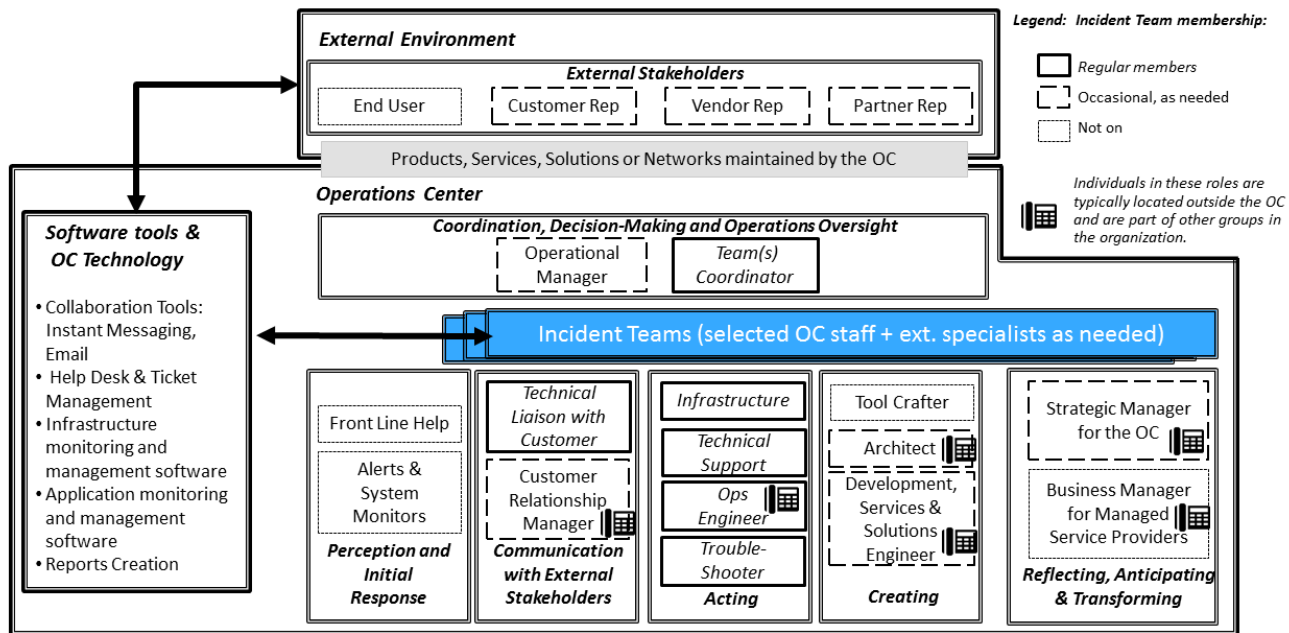


Fig. 5 Tailor-made teams in OCs: The T-TOCs model depicts the functions of operators as various socio-cognitive capabilities. Response to an incident is modeled as a distributed response across a system (the operation center) that involves multiple functions simultaneously and is coordinated by an executive function (typically an incident manager). The model depicts the reality that multiple incident teams may be formed to respond to a single incident.




Fig. 7 Off-diagonal entries indicate pairs of functions that were commonly performed by the same individual, e.g., 12 individuals who performed the function Technical Support also took on the function Alerts System Monitor.

companies it is difficult for designers to access end users, but it is much more difficult when the customers are other

businesses and end users work in restricted environments. Designers face barriers with the sales department and customer representatives within their own companies because of the need to 'control the relationship with the customer,' but they also face barriers in operations centers where there are delicate relationships with suppliers, potential disruption to time-critical work, and restricted access to information to be considered. This type of research fills an important gap in understanding.

Throughout this research we have used concepts from activity theory (group activity, motivation, object of an activity, social world, tools, division of labor, tensions and higher psychological functions) to direct and organize our work. Our paper has been an indepth study of the development of the activity of incident response over time. We have looked at the development of this activity through our review of the standards, the literature, and our own current understandings achieved through ethnography. We have shown the standards community We of the fire department metaphor to describe team formation for incident response. Incident teams are formed from a core group of trained responders, with the exception that incident response teams may also draw in other experts from outside the core group of responders as required.

However, we have also shown there is a large gap between standards, which represent best practices, and actual incident response work. Recent ethnographic research in the domain has captured some of these differences. Our literature review showed that there were many ways that incident response teams have deviated from the norm (see the final paragraph of the literature review). We see these deviations in a positive way as new forms of work.

5.2 Incident response team formation

In our own research we observed that incident response work is an important issue and many factors must be balanced in determining the best way to respond to an incident. This includes allocation of scarce resources, costs, impact on the customer relationship, the opportunity to learn, and so on. Team formation is intertwined with all of these factors.

5.2.1 *Modelling incident response work*

In our review of standards we have shown that the standards community has provided classification systems for incident response teams. For example, CERT standards describe incident response teams as coordinating, corporate and technical classifications. As another example, the NIST standard classifies response teams by type and then by the degree of distribution). The ethnographic research, especially recent research has provided rich descriptions of incident response teams and their varied nature. We present a model that stresses how incident response teams function in response to an incident. Rather than being a workflow model, our model shows the potential parts/players of the incident response team and the functions they perform. It puts forward the idea that the incident response team is a social-cognitive system by applying a cognitive metaphor of the functions that need to be performed (e.g., perceiving, decision-making, reflecting, and so on). Our model stresses that there are regular members of incident response work (a coordinator, people who act to address the issue, and others who communicate with customers) and then a large number of others who may join an incident response team on an as-needed basis. It also stresses the central role of team coordinators. Our model is grounded in ethnographic research, with significant explanatory power.

Our modelling work builds on Norros et al.'s work in particular [32] who have also taken a modelling approach in their description of one telephone operator's operations center. Our model fits in very well with the Norros et al.'s models. Their first model of the environment matches with ours closely—both models identifying three large parts: the operation center, the external world, and the technology being maintained. Comparing their second model with ours, we note that they are complementary because our model is

team-focused, and theirs is focused on the individual operator. Both models identify important issues related to incident response.

We next review some of the trends that we observed and described throughout this paper, referring back to our model as we go. We speculate on the reasons behind the changes we observed in our discussion.

Functional escalation: One aspect of incident response work is the escalation of incidents. In some centers, most incidents that cannot be handled by an operator are escalated to a single technical person whom the operator believes will be able to resolve the issue. Malega [27] points out that escalation can be hierarchical (as in ITIL's tiers) or functional. In our study we observed a move away from hierarchical escalation to functional escalation. With this type of escalation scheme there is a significantly reduced emphasis on an individual's tier. Instead individuals with different specializations are brought in to the team or dropped off as required, which is a more efficient and effective form of work. Our model does not incorporate the idea of tiers, but does rely heavily on the notion of function.

Team work: When incident teams are formed the trend is towards forming teams with members who exhibit a breadth of skills rather than escalating incidents to an individual in a higher level tier. It appears that operation centers are acting cautiously and building broader teams from the start, to explore many possibilities, knowing that it is possible for team members to withdraw easily. Escalation (even functional escalation) is not an approach that is necessarily effective, because a wrongly classified incident leads to an inappropriate escalation, which leads to wasted time and possibly negative consequences for customers or end users, or large fees when service organizations fail to meet their service level agreements. As Malega [27] has also pointed out, classification of incidents can be very hard. We observed this especially at NNO-DC. When the help desk operator did not classify an incident appropriately, the incident was bounced around between systems/ops engineers. Standards argue for tier 1 to classify incidents, but this creates tensions between help desk operators and systems/ops engineers. At Finance-MSP we observed that operators classified incidents by the service that was affected, but if the incident was not immediately solvable within a short period of time (say 15 minutes), the incident was assumed to require a team approach. Other Operations Centers like UMA-SOC and Guard-SOC quickly formed broad team as well. In both cases, the approach taken was that team members that were not required could quickly drop off the team. At Finance-MSP operators with oversight of network issues were regularly asked to investigate if the network was an issue, and regularly reported that it was not, then dropped off the incident team. This is not like the fire department metaphor, where surplus fire fighters at a fire cannot easily return to the fire department.

At Flexor-MSOC, incidents were immediately reported to customers by junior analysts, with senior analysts providing oversight. If the incident was a failure of a security tool, usually the individual responsible for the tool investigated, but others, like managers, who acted as the team coordinator, also joined the team early on to provide oversight. We observed that incident response teams were formed very early, and that the team approach seemed to be preferred to a tiered response. Our model clearly shows the potential for a broad response to an incident. An incident response team may be composed of individuals from inside and outside the organization, and may include individuals with diverse skills including technical skills, coordination skills or communication skills.

The team coordinator: Another trend that we observed across all sites was the single point of contact for the customer. Escalation of the management of the incident is a common practice described by Malega [27]. However, escalation of the management of an incident can be very confusing from the point of view of the customer, a view confirmed by Jäntti et al.'s study of service desks [22]. At Finance-MSP, rather than escalate the management of the incident, a single incident coordinator assumed responsibility for the incident until their shift changed, at which point there would be a handover, but no change in the way the customer communicated with the incident manager, i.e., the new team coordinator simply took over from the previous team coordinator on the bridge call, and used notes and logs to catch up on the action. Further, as the incident progressed, escalating levels of management and individuals with responsibility for customer relationships would join the call, but not coordinate it. Great care was also taken to ensure a single point of contact for the customer at N&S-SP, UMA-SOC and Guard-SOC. In our model we highlight the team coordinator's role as central to coordination, decision-making and oversight.

The customer relationship: In all operations centers the customer relationship is important, but in managed service providers, the customer relationship is particularly important. In larger centers like Finance-MSP, sometimes separate teams were established for managing the customer relationship (for important customers these were established immediately). These customer teams were as important as the technical teams addressing issues. In our model we emphasize the importance of this function by stressing the importance of communication with external stakeholders.

Outside vendors and incident response teams: We also observed the practice of engaging outside vendors or other service providers early as a new form of work. This was a clear indication of the increasing complexity of service provisioning and the dependencies involved. It is no longer possible to solve many incidents with a team assembled from a group of incident response experts and the occasional help of outside experts, as standards suggest. In-

creasingly, teams of teams are being assembled; and in many cases, some of these teams are outside of the organization. Distributed incident response teams are much more the norm than co-located teams. As evidence of this, at Flexor-MSOC and Finance-MSP, war rooms are not being used. In the review of standards we saw a simple form of team-to-team interaction in the role of coordinating CSIRT teams in CERTs, but some of the relationships with teams that we observed were more complex and ongoing. For example, at Finance-MSP we saw the network operators engaging in constant communication with vendors of switches and terminals as the good relationship between the operations center and the vendor was beneficial in both directions. For the operations center it helped greatly when the response to an incident required vendor involvement. The evidence of complex relationships was also clearly evident at Flexor-MSOC where the operators had developed a specialized tool to help them manage customer relationships. The tool recorded the customer's communication preferences under varied circumstances and the preferences were followed precisely by the operators. In the case of a managed security service provider, the customer is typically responsible for responding to incidents; collaboration with the service provider is one part of that mix. For each customer, the role required of the service provider is unique and a specialized tool was the only way that Flexor-MSOC could get their role right in each case. Our model includes the role of these important participants in an incident response by including them as actors in the 'external environment' of the operations center.

The role of non-technical experts: We also saw a definite trend towards the inclusion of experts with non-technical skills as the norm on incidents (e.g., communication experts, business managers), which again is not something captured strongly by the standards. The resolution of incidents is very important from a business perspective. Incident resolution is not just about technical issues; it is also about building and strengthening the customer relationship. All of the operations centers, whether they were service providers or not, were very concerned about their relationship with their customers. Our model includes two important categories of non-technical experts. One category includes individuals with expertise in communicating with external stakeholders. The second category includes individuals with expertise in reflecting on, anticipating and transforming the incident response activity itself.

Multi-tasking: A final trend that we observed was intensive multi-tasking by operators. It was not unusual for operators to be working on several incidents at once. This was possible because there are often periods of inactivity for operators who are working on an incident when they are waiting for actions or results from others, such as the upgrade of software or the collection of information for a subsequent forensics investigation. Individuals with a more

technical bent were particularly in high demand and could have as many as 20 ‘projects’ open at one time, some of which would be classified as incidents. We observed that the pressure to multi-task comes from the desire of operations center managers to reduce their costs as much as possible. Operators also multi-tasked with other non-incident activities. These are listed in section 4.2 and are: ‘mentoring others on technical matters’, ‘learning about technical matters’, ‘creating reports’, advancing special projects,’ and tuning and tinker with tools’. Operators also performed multiple roles, as we showed in our results section; Also, monitoring was often multi-tasked with other tasks.

Tailor-made teams: Overall, we have observed that incident response teamwork is very unlike planned teamwork for longterm IT projects, where team members are allocated to the project at the beginning and stick with it to the end (a duration which transpires over months or possibly years). Tailor-made teams are also comprised of qualified experts who must communicate, coordinate, and collaborate, but these teams are unplanned in the sense that they must be assembled in response to an incident that is occurring and must imminently be resolved in the next moments, hours or days. Further, unlike planned teams, the participation of team members in unplanned teams is fluid, with individuals arriving and departing at any time for a variety of reasons. The need for tools to support unplanned teams, i.e., the type of team we observed in this study, and which is emerging as a new form of work in operations center, is in its infancy. We did not see good tool support for this type of work. In our model we emphasize this aspect by stressing that incident response teams are tailor-made.

5.3 Tool support for Tailor-made teams

In previous papers we have described the environment of the operators [6] and usability problems operators encounter with tools [16]. We described the operator’s workspace as a highly complex technological setup requiring 2 or 3 large displays, as many as half a dozen alerting tools, 20 to 40 simultaneous chat sessions and sophisticated phones that allow an operator to participate on several calls at once. We observed that none of the tools the operators use were specifically designed to support tailor-made, unplanned teamwork.

In 2015 Alsabbagh and Kowalski [2] addressed the issue of tool support for incident response in SOC’s. They concluded that their literature review and their own experience indicated a lack of customizable incident response tools that facilitate communication and elaboration within organizations during incident management. They also called for holistic approaches to understanding incident management that considered both the technical and social aspects.

An ideal tool would enable situation awareness, decision-making, coordinating, ... essentially all of the functions

that are commonly required on teams as reflected in our model. Ticketing systems do go a long way towards easing these issues but are not a complete solution (we saw the stress amongst operators at N&S-SP and UMA-SOC where no ticketing tools were used), but we also observed that ticketing systems had really been designed for tier 1 operators and not for teams at NNO-HD.

We would like to see designers designing explicitly for the physical environment of operators. This would include creating software experiences for multiple-display environments: ensuring dialog boxes appear on the correct display, dashboards potentially spanning multiple displays, etc. Tools that were more closely integrated with display and Window management would help operators manage their windows more effectively. There is also a potential for developing mobile applications for mobile operators, such as the data center operators who monitored and serviced physical equipment either in the clean room of the data center or out in the field.

Team management tools to support work on incidents are in their infancy and would require significant design effort. Many current tools are not suitable; all of the tools we have reviewed for managing teams (not reported in this paper) were too heavyweight, and we saw none of them in use in any of the operations centers we visited. A team management tool for unplanned teamwork would have to address how work is actually assigned. In incident response work, tasks may be assigned to teams, particular roles or individuals. Either teams or individuals can generate or pick up tasks, and tasks can be dropped at any time if new information comes to light.

To enable the work itself we imagine some useful techniques. One would be the design of modifiable digital templates that could be converted into checklists. Checklists were in use at several sites. Checklists are a very powerful tool for specialists under stress [14], and could prove useful in operations centers. Much of the work of operators is highly repetitive, but with variations for the situation. Modifiable team composition templates would also be very useful for setting up teams quickly, especially if they were designed so that team coordinators and team members could reach individuals through function names in the model rather than through their individual names.

Secondly, to support incident response work, it is important that operators have an environment where they can rapidly switch between multiple tasks. Even though an operator is “working an incident”, it is not always the case that they are actively working on that incident 100% of the time. There are many instances when ‘field work’ (dumping memory, reconfiguring servers or software, waiting for repair parts, waiting for the results of a query, waiting to find an expert, ...) holds up the work and at these times operators switch to other tasks. Therefore it is important to have

tools that support both planned and unplanned work because the operators' work is a mixture of both. Operators need the ability to quickly switch between incident work and other current tasks like monitoring, report writing, or learning and have their entire workspace rearranged to support their new primary task, while still allowing them to keep an eye on background tasks. Essentially what is required is support for restoring context when context switching between incidents or between incidents and other types of work.

Thirdly, unplanned teamwork support tools would have to be designed to support fluid team memberships because operators do not necessarily stay with the incident until it ends. The work may be handed over at the end of shift, higher priority work may come along, or their particular skills are no longer needed. It would also have to be designed to enable collaboration with 3rd parties such as vendors or coordinating ISIRTs. Further, because the team composition can become complex, there needs to be a mechanism to enable situation awareness across incident teams that are comprised of multiple sub-teams. This need not be very elaborate, but would help each person or each team know their place in the work, and therefore know how they could best contribute. The tool should also enable and manage the communication channels associated with an incident. It is not unusual for an incident to be supported by multiple bridge calls and a dozen chat sessions. It can be easy to lose track of which channel of communication is associated with which incident when an operator is working on multiple incidents.

Finally, tools are needed to enable collaborative decision-making. In incident response work, the people with the most expertise are not necessarily the team coordinators, and therefore it makes sense that the discussion that enables the team coordinator to make a decision be spread across the team.

5.4 Limitations of the findings

Our study did not include a case where technical teams in OCs worked with CERTs. This form of incident response work is therefore not apparent in our model. To remedy this we could add a function within the external box labelled 'external OC'.

We also did not study incident response teams that were not a part of an operations center within large companies as did Botta et al. [3]. This form of incident response teams appears to be very loose and does not necessarily involve a team coordinator. The model does however capture this type of team of experts scattered across various departments who are able to act when incidents arise. Other more recent research on incident response teams within organizations shows significantly more structure to the response and the type of teams as in Hove and Tärnes [17] and others [18]

[36]. The T-TOCs model does not clearly show internal and external teams working together, or hierarchies of teams, as we did not witness the former and saw only flat team structures. Incorporating the structure of teams of incident teams into our model is future work.

Another form of team we did not encounter was described by Ahmad et al. [1]. They described a coordination team and a technical team, which is not suggested by our model. However, this may simply be a semantic matter, as we did see multiple teams being coordinated by a team coordinator. Very often one team was customer facing (as in Ahmad et al.'s coordination team) and the other technology-facing (as in Ahmad et al.'s technical team).

We did not sample evenly across functions. For example, Figure 6 shows that our knowledge of the function of the work of front line help is limited to only one of the OCs. Our knowledge of some functions is primary, while other functions such as Technical Liaison, Customer Rep, Partner Rep, Architect and Business Manager is based on second hand accounts obtained through interviews. We also tend to know more about the work of operators in the room than those outside of it. We do not have first hand information about those on teams who are outside of the organization (e.g. vendor representatives).

5.5 Comparison to other research

In previous papers we have detailed the multi-tasking nature of operators [35], tensions between activities within the OC [6], and usability problems with the operator's tools [16]. In this paper we have focused on the nature of incident teams.

Our work is consistent with the advice of standards, but standards take either a prescriptive approach or a classification-approach when prescribing incident response teams. In comparison, our work produced rich descriptions of case studies and the T-TOCs model. Our work is much like other ethnographic studies, but differs with respect to the number of sites involved in the study.

We contribute to the relatively small body of research on tools in operations centers by proposing tools for supporting ad hoc unplanned work. Other work has, for example, explored the use of timelines [26] or the failure of knowledge management systems [37]. We advocate for an ethnographic approach to designing tools. Our impression is that designers have been too influenced by standards and guidelines and also have not had enough direct contact with their end users.

6 Conclusions

Using an activity-theoretical perspective, we studied incident response in an IT operations center. Building on the work of a number of researchers and Norros et al.'s general

models of operator work in particular [32], we focused on incident response and the teams that are formed to respond to these incidents. We asked: what is the context of incident response work?, how can we model incident response work?, and what are the implications for tool developers?

Standards for service management processes and security incident response processes recommend that organizations establish a fixed group of trained experts to respond to incidents, as in the fire department metaphor. Responders are selected from this group to respond to an incident and experts may augment these teams.

Reviewing the work of other researchers who studied operations centers, we found that team composition has increased in complexity over time. The idea of a team of experts, selected from a fixed group who function as responders with the help of occasional others is breaking down. Potential augmentation of the team with a very wide range of experts is becoming the norm. Further, while incident teams of the past may have been co-located, it is now common for some team members to be co-located while others (sometimes many others) are not co-located. Through our review we found that today's response teams often engage in complex interactions with other response teams to resolve incidents. These other teams may be within or outside the organization. We also found that incident response teams regularly included non-technical participants, such as high-level managers and customer representatives. In some situations the structures of teams could be quite complex and could be teams of teams. Our new findings build on that of other researchers.

We studied 7 OCs, and saw many of the patterns reported in the literature. We found that incident response teams were formed very early after the appearance of incident, and that the team approach seemed to be preferred to escalating an incident from one individual to the next. We noted there is a tendency for the team membership to be over-estimated early on with team members dropping off when their skills are not required. Other experts may be brought onto the team as required, and generally they appear early. We also observed that incident response teams have increasingly complex structures. Teams are commonly comprised of sub-teams, and it's not unusual for sub-teams to be external to the organization. Highly distributed incident response teams are very normal.

We observed a strong focus on customers. Teams worked hard to ensure there was a single point of contact for customers, and that the team liaised with the customer in the way the customer preferred. In larger OCs, and especially managed service providers, separate teams existed to manage the customer relationship and these were as important as the technical teams. In addition we found managers will sometimes add themselves to an incident management team to track an incident when a high-valued customer is involved.

There is also a strong emphasis on the role of team coordinators (usually a single person, but sometimes a team).

The role of vendors on incident teams is also changing. Outside vendors are engaged early in the incident response process. Relationships of operators with vendors are also prioritized; good operators maintain close contact with their vendors.

And finally, although our focus is on incident response we noticed that operators were broadly interested in incidents being handled at any point in time, whether or not they were on the incident team, because of the potential for inter-connectivity between incidents and the help the knowledge may be in understanding future incidents.

We found that the way that incident teams were formed and the way they worked was not well supported by existing incident response systems such as call systems, chats and ticketing tools. We identified unplanned teamwork as a type of work that is unsupported by specific tools. This is especially apparent when considering the number of tools that exist for planned work.

We suggest areas in which new tools are needed to reduce the stress experienced by operators due to cognitive overload, to reduce avoidable errors, and to reduce the time to resolve incidents. Our suggestions included modifiable checklists that become to-do lists for ensuring task completion as an aid for specialists under pressure. Modifiable templates for team composition that are based on a customized version of the T-TOCs model, and the ability to reach out to individuals who can perform necessary functions, through the function that is required, rather than through the person's name. Secondly, to support multi-tasking we suggest the development of tools that allow operators to quickly switch between incident work and other current tasks like monitoring. Such a tool would support the operators to focus on their current task while maintaining a minimal awareness of changes in other tasks. Thirdly, while we observed ticketing tools support responsibility and resolution of an incident, we suggest tools that provide support for incident response teamwork. This would require a new sort of project management tool for unplanned ad hoc work that would allow very quick team setups, fluid team membership, reporting, support for situation awareness and support for managing the many channels of communication within a sub-team and between sub-teams. Finally, within the context of the new project management tool we see a need for simple tools for enabling collaborative decision-making in the context of incident response work.

We produced a model depicting the composition of incident response teams. Our model could be used to assess and then tune incident response teams by enabling discussions on the necessary functions for a particular incident. Such a discussion would be a way of implementing a sort of functional escalation scheme. If used in the centers, it could

also aid with communication. Currently operators communicate with other individuals engaged in the incident response, but as the response teams grow larger and more complex, it may be easier to communicate through functions e.g., report changes in the environment to the Customer Relationship Manager, rather than a specific individual by name e.g., Brian Smith, which is the current method. The model also has potential as a visualization for conveying who is on the call, a seemingly simple question that was frequently not one an operator could answer easily. In this way the model could serve as a basic coverage tool for an incident.

Our findings build on and are not in opposition to work on standards such as ITIL. What this research adds is additional information about the lived experiences of modern operators (many of whom adhere to standards) and how their work can be supported.

7 Acknowledgments

Thank you to the Natural Sciences and Engineering Research Council of Canada (NSERC), ISSNet and SurfNET for funding; special thanks to Mitacs and CA Technologies who funded this work with a Mitacs Elevate grant. Thank you also to our study participants and the Operations Centers that allowed us access to their sites and employees.

References

1. Ahmad A, Hadgkiss J, Ruighaver AB (2012) Incident response teams—challenges in supporting the organisational security function. *Computers & Security* 31(5):643–652
2. AlSabbagh B, Kowalski S (2015) Security from a systems thinking perspective-applying soft systems methodology to the analysis of an information security incident. In: *Proceedings of the 58th Annual Meeting of the ISSS-2014 United States*
3. Botta D, Muldner K, Hawkey K, Beznosov K (2011) Toward understanding distributed cognition in IT security management: The role of cues and norms. *Cognition, Technology & Work* 13(2):121–134
4. Boylan, Don (2014) ITILtopia: The Tyranny of Tiers. <http://itiltopia.com/?p=458>
5. Brewster E, Griffiths R, Lawes A, Sansbury J (2012) IT service management: a guide for ITIL foundation exam candidates. BCS, The Chartered Institute for IT
6. Brown JM, Greenspan SL, Biddle RL (2013) Complex activities in an operations center: A case study and model for engineering interaction. In: *Proceedings of the 5th ACM SIGCHI symposium on Engineering interactive computing systems*, ACM, pp 265–274
7. Calder A (2013) ISO27001/ISO27002: A pocket guide. IT Governance Publishing
8. Corbin J, Strauss A (2014) Basics of qualitative research: Techniques and procedures for developing grounded theory. Sage publications
9. Daniels H (2008) Vygotsky and research. Routledge
10. Duignan M, Noble J, Biddle R (2006) Activity theory for design from checklist to interview. In: *Human Work Interaction Design: Designing for Human Work*, Springer, pp 1–25
11. Engstrom Y (2000) Activity theory as a framework for analyzing and redesigning work. *Ergonomics* 43(7):960–974
12. Flach JM (2012) Complexity: learning to muddle through. *Cognition, Technology & Work* 14(3):187–197
13. Gartner (2014) IT Glossary. <http://www.gartner.com/it-glossary/it-services>
14. Gawande A, Lloyd JB (2010) The checklist manifesto: How to get things right, vol 200. Metropolitan Books New York
15. Grance T, Kent K, Kim B (2012) NIST special publication 800-61r2: Computer security incident handling guide. <http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-61r2.pdf>
16. Greenspan S, Brown J, Biddle R (2012) The Human in the Center: Agile decision-making in complex operations and command center. CA Labs Research p 12
17. Hove C, Tårnes M (2013) Information security incident management: An empirical study of current practice. Master's thesis, Norwegian University of Science and Technology
18. Hove C, Tårnes M, Line M, Bernsmed K (2014) Information Security Incident Management: Identified Practice in Large Organizations. In: *IT Security Incident Management IT Forensics (IMF), 2014 Eighth International Conference on*, pp 27–46, DOI 10.1109/IMF.2014.9
19. Humphreys E (2011) Information security management system standards. *Datenschutz und Datensicherheit-DuD* 35(1):7–11
20. ISO/IEC (2013) Information technology – security techniques – code of practice for information security controls. <http://www.iso27001security.com/html/27002.html>
21. ISO/IEC (2013) Information technology – security techniques – information security management systems – requirements. http://www.iso.org/iso/home/store/catalogue_ics/catalogue_detail_ics.htm?csnumber=54534
22. Jäntti M, Cater-Steel A, Shrestha A (2012) Towards an improved it service desk system and processes: a case

- study. *International Journal on Advances in Systems and Measurements* 5(3 & 4):203–215
23. Kapella V (2003) A framework for incident and problem management. *International Network Services* whitepaper
 24. Killcrece G, Kossakowski KP, Ruefle R, Zajicek M (2003) Organizational models for computer security incident response teams (csirts). Tech. rep., DTIC Document
 25. Kuutti K (1996) Activity theory as a potential framework for human-computer interaction research. In: Nardi B (ed) *Context and Consciousness*, MIT Press, Cambridge, MA, chap 2, pp 17–44
 26. MacEachren AM, Jaiswal A, Robinson AC, Pezanowski S, Savelyev A, Mitra P, Zhang X, Blanford J (2011) Senseplace2: Geotwitter analytics support for situational awareness. In: *Visual Analytics Science and Technology (VAST)*, 2011 IEEE Conference on, IEEE, pp 181–190
 27. Malega P (2014) Escalation management as the necessary form of incident management process. *Journal of Emerging Trends in Computing and Information Sciences* 5(6):641–646
 28. McDonald S (2005) Studying actions in context: a qualitative shadowing method for organizational research. *Qualitative research* 5(4):455–473
 29. Metzger S, Hommel W, Reiser H (2011) Integrated security incident management—concepts and real-world experiences. In: *IT Security Incident Management and IT Forensics (IMF)*, 2011 Sixth International Conference on, IEEE, pp 107–121
 30. Möller K (2007) Setting up a Grid-CERT: experiences of an academic CSIRT. *Campus-Wide Information Systems* 24(4):260–270
 31. Nardi BA (1998) Concepts of cognition and consciousness: Four voices. *ACM SIGDOC Asterisk Journal of Computer Documentation* 22(1):31–48
 32. Norros L, Norros I, Liinasuo M, Seppänen K (2013) Impact of human operators on communication network dependability. *Cognition, technology & work* 15(4):363–372
 33. Roth WM, Lee YJ (2007) Vygotsky's neglected legacy: Cultural-historical activity theory. *Review of educational research* 77(2):186–232
 34. Sallé M (2004) IT service management and IT governance: review, comparative analysis and their impact on utility computing. *Hewlett-Packard Company* pp 8–17
 35. Samaroo R, Brown JM, Biddle R, Greenspan S (2013) The day-in-the-life scenario: A technique for capturing user experience in complex work environments. In: *2013 10th IEEE International Conference and Expo on Emerging Technologies for a Smarter World (CEWIT)*, pp 1–7
 36. Tøndel A, Line MB, Jaatun MG (2014) Information security incident management: Current practice as reported in the literature. *Computers & Security* 45:42–57
 37. Trusson CR, Doherty NF, Hislop D (2014) Knowledge sharing using it service management tools: conflicting discourses and incompatible practices. *Information Systems Journal* 24(4):347–371
 38. Turner P, Turner S (2001) A web of contradictions. *Interacting with Computers* 14(1):1–14
 39. Vygotsky L (1934) *Thinking and speech. The collected works of LS Vygotsky, vol. 1.* New York, NY: Plenum
 40. West-Brown MJ, Stikvoort D, Kossakowski KP, Killcrece G, Ruefle R (2003) Handbook for computer security incident response teams CSIRTs. Tech. rep., DTIC Document
 41. Wiik J, Gonzalez JJ, Davidsen PI, Kossakowski KP (2009) Chronic workload problems in CSIRTs. In: *Twenty Seventh International Conference of the System Dynamics Society* July, at Albuquerque, NM, USA
 42. Wiik J, Gonzalez JJ, Davidsen PI, Kossakowski KP (2009) Persistent instabilities in the high-priority incident workload of CSIRTs. In: *Twenty Seventh International Conference of the System Dynamics Society*
 43. Wiik J, Gonzalez JJ, Davidsen PI, Kossakowski KP (2009) Preserving a balanced CSIRT constituency. In: *Twenty Seventh International Conference of the System Dynamics Society* July, at Albuquerque, NM, USA
 44. Zieba S, Polet P, Vanderhaegen F, Debernard S (2010) Principles of adjustable autonomy: a framework for resilient human-machine cooperation. *Cognition, Technology & Work* 12(3):193–203