DOCTORAL THESIS

Title	**Organizational Learning From Information System-Related Incidents**
Presented by	MOHAMMAD HOSEIN REZAZADE MEHRIZI
Centre	ESADE BUSINESS SCHOOL
Research Unit	INSTITUTE FOR INNOVATION AND KNOWLEDGE MANAGEMENT
Department	DEPARTAMENT DE POLÍTICA D'EMPRESA, DIRECCIÓ DE RECURSOS HUMANS I SISTEMES D'INFORMACIÓ
Directed by	DAVIDE NICOLINI & JOAN RODON MODOL

ABSTRACT

This thesis examines how organizations try to leverage the experience of major information system-related incident (ISRI) to avoid the recurrence of failure and to reduce their impacts. Accordingly, the study utilizes a situated learning perspective and a practice view to respond to the question "*how do* IS *organizations learn from their internal, large Information system-related incidents?"*

Employing a multiple, inductive case study design, the study found that organizations adopt a wide range of practices during and after the incident handling process. This resulted in articulating five learning modes: 1) learning through incident handling, 2) post-incident reflection, 3) transversal learning, 4) outsourced learning, and 5) learning through material replacement. Although the first two learning modes are documented in other domains, the study shows how the characteristics of ISRIs affect the practices associated to these two learning modes. Further, the analysis focuses on the other three learning modes that seem to be typical of this sector.

Transversal learning refers to the fact that while some of the learning practices are focused on individual incidents, specific learning practices exist which take into account multiple adverse events. Outsourced learning indicates that capitalizing on the experience of an incident is often carried out through relying on specialized providers that handle incidents. Finally, the particular nature of the IS work processes and its material basis allow for learning through material replacement.

The thesis enriches our understanding of the processes whereby organizations learn from ISRIs, thus providing contributions to theoretical developments in knowledge and learning literature. More specifically, the results of the research challenge the established temporal view about *when* learning process takes place. While the existing literature on organizational learning suggests that learning takes place either during incident (through incident handling practices) or (right) after

(through post-incident reflection and learning), the current study suggests that the temporal pattern of learning process is not necessarily confined to this established dichotomy. The concept of transversal learning indicates the importance of looking at learning practices that take place in parallel and with a considerable temporal distance from incidents that occur in a proper moment of learning.

The thesis adds to knowledge literature by foregrounding the importance of the materiality regimes underlying the learning process. The idea of learning through material replacement shows that the modular, adaptable regime of materiality that dominates ISRIs can lead organizations to benefit from their incident experience, without necessarily knowing the causes of incidents and their potential solutions.

Finally, the thesis advances our understanding of the role of politics and governance of learning from incidents in the IS sector. It does so by highlighting the role of *neutralization* (versus normalization) and *dramatization* (versus rationalization) in the process of learning. The study found that the ignorance of influential actors about the technical aspects of the problem could be leveraged by learning agents to add to the political pressure needed to drive learning initiatives. The study also highlights the critical importance of organizational governance for the process of learning. The concept of outsourced learning underscores the importance of a learning governance system, which complements the two dominant learning governances in the literature –i.e., intra and inter-organizational learning. This shows that for capitalizing on the experiences of past incidents, organizations might avoid performing learning practices, especially when they are dealing with a wide range of changing technologies (learning abstinence), since the knowledge gained is expected to be obsolete or decrease in value through time.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

Chapter 1: Introduction

In facing critical incidents, it is not enough that organizations only recover from them. They also need to learn from them to avoid similar incidents in future and reduce their impacts. Incidents provide various learning opportunities for organizations as s*ources* (e.g., when organizations draw lessons from analyzing them) and *triggers* of learning (e.g., when they force organizations to take actions to avoid future incidents). However, incidents can pose serious challenges on learning process as part of the learning *context* (e.g., when the urgency, negative valance, and ambiguity of incident challenge learning practices). Recent studies have emphasized on the importance of looking at the *process* through which organizations capitalize on their incidents' experience. These studies call for a

deeper examination of the learning process with regards to the specific contextual factors related to various types of incidents. This has resulted in specialized studies in healthcare and safety incidents, industrial incidents, and natural disasters.

Literature has shown that information system-related incidents (ISRIs) also have their own specific characteristics that can influence learning process. However, the review of these studies shows that the focus has been mainly on the lessons that organizations (should) develop through formal post-incident analysis practices. Still we need to understand how organizational learning process unfolds in the specific context of ISRIs.

Based on a situated learning perspective and by adopting a practice view of organizational learning, the study tries to answer, "*How do organizations learn from their internal, large Information system-related incidents?"* Adopting a qualitative interpretative approach in four IT organizations (two information technology security and two supercomputer organizations), I focus on 15 major incidents and explore what the organizations did to learn from those incidents.

The results of the empirical study shows that the organizations adopt a wide range of practices *during* each incident handling process –i.e., handling incidents through ticketing system, doing Triage, and interactions with specialized providers- and *after* that –i.e., post-incident reflections and laboratory incident analysis. It also shows that the organizations adopt several practices –i.e., using Wiki systems, executing improvement projects, and performing upgrades- helping them to leverage the experience of *several* incidents.

The analysis of the learning practices resulted in articulating five learning modes: (1) learning through incident handling; (2) post-incident reflection; (3) transversal learning from incidents (4); outsourced learning; (5) learning through material replacement. This adds to the existing literature that mostly focuses on two traditional learning modes: learning through post-incident reflection and learning

through incident handling practices. Accordingly, the articulation of the new learning modes challenges the established temporal view about *when* learning process takes place. It also highlights the importance of materiality regimes underlying the learning process, advances our understanding of the politics of learning from incidents by highlighting the role of *neutralization* (versus normalization) and *dramatization* (versus rationalization), and shows the importance of considering various governance systems through which learning process takes place.

1.1 The importance of organizational learning (OL) from information system-related incidents (ISRIs)

Organizations increasingly face incidents related to their information systems, for instance due to security breaches and breakdowns of technological systems that disrupt their normal operations. Organizations are the key targets of such damages. For instance, according to the latest survey by Ponemon Institute, the average annual cost that cybercrimes posed on organizations like HP, in 2011, was around $5.9m, 56 % higher than the previous year (Ponemon Institute, 2013).

Of course, this is not confined to security incidents. Failures of information systems impose considerable financial, reputational, and privacy costs on organizations. For example, a simple breakdown of a banking system costs a fortune for a major bank like RBS that handles billions of transactions worldwide (Lex team, 2012). At the same time regulator authorities urge organizations ““to take reasonable care to establish and maintain such systems and controls as are appropriate to its business”, and to ensure business continuity” (Lex team, 2012).

Surprisingly enough, a large portion of these damages can be prevented easily. For example, in a report by The Verizon 2012 Data Breach Investigations, it was estimated that out of 855 data breach incidents to US secrete services and national security bodies, "97 per cent could have been avoided if simple security measures had been in place" (Taylor, 2012). This shows the importance of capitalizing on past ISRIs experience to avoid similar incidents in future or at least reducing their impacts. It also implies that the challenge often is not identifying *what* should be done; rather the difficulty is *how* to make sure that the experience of past incidents results in actual changes in practices that can prevent future incidents.

Various evidences indicate that learning attempts should not be merely confined to technological factors (Nelson, 2007). It is also critical that organizations take into consideration factors such as security policies, awareness programs, employees training, and motivational factors (Anderson & Moore, 2006). In other words, organizations need to learn from their experiences on both technical and organizational factors to become more immune and more prepared to deal with future incidents. This need heightens regarding the increasing number, heterogeneity, and severity of IT incidents (Ponemon, 2009; Ponemon, 2010).

1.2 The challenges of OL from ISRIs

Literature on organizational learning from incidents (Beck & Plowman, 2009; Carroll & Fahlbruch, 2011; March, Sproull, & Tamuz, 1991; March & Olsen, 1975; Rerup, 2009) has documented a wide range of barriers that organizations face when they engage in the learning *process* (D. Smith & Elliott, 2007). These include, for example, barriers due to cognitive limitations of organizations (Levitt & March,

1995), structural barriers and rigidities (Hannan & Freeman, 1984; Tripsas & Gavetti, 2000), and political challenges (Elliott, Smith, & McGuinnes, 2000; Smith & Elliott, 2007; Stern, 2002).

In general, past experience can result in improvements. This has been demonstrated in studies on learning curve in various sectors (Argote & Epple, 1990; Epple, Argote, & Devadas, 1991; Haunschild & Sullivan, 2002; Lapre, A. S. Mukherjee, & Wassenhove, 2000). However, recent studies on learning from incidents have shown that learning from experience can be different when it is through routine, normal experiences, compared with significant, unexpected incidents (Christianson, Farkas, Sutcliffe, & Weick, 2009; Lampel, Shamsie, & Shapira, 2009; Rerup, 2009; Starbuck, 2009). In the latter situation, there are strong elements of cognitive mismatch between established schemas and evidences, high level of shock that attracts attentions, and often low frequency that reduces the chance of multiple trial and error practices. Hence, the conclusion is that learning from major incidents is different from learning from daily, routine ones (Baumard & Starbuck, 2005; Lampel et al., 2009). The focus of this study is merely on learning from *major* incidents.

In addition, previous studies have argued that organizations learn differently from successful events versus negative incidents (Baumard & Starbuck, 2005). That is because the negative valence of incidents provides different psycho-political conditions for actors that can influence their collective reactions to incident experience (Cyert & March, 1963; Starbuck, 1983). This study only focuses on major *negative* incidents.

Literature suggests that organizations can learn from their major incidents by reflecting on incidents right after handling them, through a systematic analysis of the causes of incidents, drawing lessons on how future incidents can be prevented or managed effectively, and incorporating those lessons into organizational

routines and structures and culture (Carroll & Fahlbruch, 2011). Another different approach frames learning as a process that is situated in daily activities of organizational actors (Lave & Wagenr, 1991). In the context of learning from incidents, this view implies that learning emerges out of the engagement of actors in the very incident handling practices (Butler & Gray, 2006).

Building on these views, studies on learning *from incidents* have shown that the characteristics of incidents, such as their ambiguity, negative content, urgency, and revealing nature influence the learning process. For example, Haunschild & Sullivan (2002) show that organizations with a heterogeneous profile are more likely to learn from more complex incidents than simple events. In addition, some studies have documented evidences that the urgency of incidents, though might trigger the learning process, can also make organizations engage in a superficial, incomplete learning process, mostly to protect themselves against external pressure, rather than improving their defects that caused the incident (Elliott, 2009).

In addition to generalist literature, several studies have specialized in terms of specific incidents that inspire learning process. For instance, numerous studies on healthcare and safety incidents (Carroll & Edmondson, 2002; Nicolini, Powell, Conville, & Martinez-Solano, 2008; Rivard, Rosen, & Carroll, 2006), industrial incidents (Carroll & Hatakenaka, 2001; Carroll, Rudolph, & Hatakeneka, 2002; Carroll, 1995; Carroll & Fahlbruch, 2011), and natural disasters (Jasanoff, 1994; Parker, Stern, Paglia, & Brown, 2009) have articulated the challenges that organizations face in their learning process. The assumption of these studies is that the characteristics of incidents and the context in which they often occur influence the learning process. This has led to more specific and contextualized understanding of learning process and its nuances in each specific context.

Following this line of reasoning, the argument of this study is that ISRIs possess specific characteristics that have significant bearings on learning process. ISRIs are rooted in a wide range of interconnected technologies. This technical nature of ISRIs implies that diverse, specialized actors be involved in learning process (Loch, Carr, & Warkentin, 1992). These technologies also change frequently, which constantly faces organizations with emergent incidents (Egan, 2007). This requires further learning attempts by organizations to deal with new incidents and challenges on the applicability of past experiences in facing future incidents (Baumard & Starbuck, 2005).

However, the review of literature (chapter 2) shows that although information systems (IS) literature recognizes abovementioned characteristics, it does not contextualize learning process with regards to them. In fact, the specialization has revolved around *what* specific lessons organizations can learn from various categories of ISRIs. When it comes to learning *process*, the literature is quite generalist. The studies mostly refer to general learning processes such as post-incident reflection. Some empirical studies have questioned such a narrow, prescriptive view of learning from incidents by showing that organizations have limited chances to learn only through post-incident reflection (Gwillim, Dovey, & Wieder, 2005; Kasi, Keil, Mathiassen, & Pedersen, 2008).

What remains underexplored is how do the characteristics of ISRIs might affect the process of learning from them? Do organizations apply the same learning practices in the context of ISRIs, as they do in other types of incidents such as natural incident? Do ISRIs make organizations perform traditional learning practices such as post-incident reflection differently? Do organizations avoid some expected learning practices such as root-cause analysis that are well established in other contexts such as healthcare?

1.3 Research aims and expected contributions

The aim of this thesis is to understand the process through which organizations learn from their ISRIs. More specifically, I am interested in a contextual understanding of organizational learning process with regards to the characteristics of ISRIs. The goal is to openly explore the practices that organizations adopt to leverage their ISRIs experiences to enhance their capabilities for avoiding similar incidents in future and reducing their impacts.

The study seeks to articulate learning practices that are relevant in the context of ISRIs and examine how the characteristics of ISRIs might lead organizations to adopt them. It is expected that new learning modes be articulated that reflect the specific characteristics of ISRIs and their context. This, in turn, results in two specific contributions to the literature on organizational learning. First, it shows to what extent the traditional learning processes are relevant in the context of ISRIs. For instance, do organizations go through systematic root-cause analysis when they try to learn from their ISRIs? This helps specifying the boundaries of learning theories by examining whether they are relevant in the specific context of ISRIs or not.

Second, by capturing openly the practices that organizations apply for their learning purposes and by identifying new practices and articulating them into new learning modes, the study examines the relations between the learning practices and the contextual factors, especially the characteristics of ISRIs. Accordingly, the study tries to articulate insights about how the specific characteristics of ISRIs might qualify different learning practices, compared with other domains. This allows identifying some contextual factors (such as the temporality of technology and incidents, the specific underlying materiality regime that is present in ISRIs,

and the overall governance system through which ISRIs are handled) that warrant theorizing for understanding the learning process.

1.4 The structure of Thesis

The next chapter reviews literature to develop the conceptual framework for the study. I review literature to examine how previous studies have understood the process of learning from ISRIs. The result of this review suggests an overall framework that links the characteristics of ISRIs to OL process.

Chapter 3 describes the process of an empirical multiple case study design for understanding the process through which organizations learn from their ISRIs. After describing the research question, the overall research approach and design, the chapter reports data collection and data analysis processes. This is followed by describing the empirical settings of the study.

Chapters 4 and 5 describe the practices that the studied organizations adopted to avoid similar incidents in future or reduce their impacts. Chapter 4 focuses on practices related to a *single* incident (single-incident learning), both during and after incident handling. Chapter 5, reports practices that the organizations took in relation to *multiple* incidents (multi-incident learning).

Chapter 6 articulates five learning modes based on the practices observed in the organizations. In addition to two learning modes that are articulated in the literature (learning through incident handling practices, and post-incident reflection), the chapter focuses on analyzing three new learning modes that are less developed in the literature (transversal learning, outsourced learning, and learning through material replacement). I examine how the characteristics of ISRIs shape the way in which each learning mode unfolds.

In chapter 7, I discuss how the three new learning modes contribute to our understanding of 1) the temporal patterns of learning process, 2) the politics of learning from incidents 3) socio-materiality of learning process, 4) the governance of OL process, and 5) the organizational aspects of OL process.

Chapter 8 summarizes the findings and contributions of the research, highlights the limitations of research, suggests some lines for future studies, and comments on implications for practice.

Chapter 2: Background Literature

In this chapter, I first define organizational learning (OL) and its four dimensions (content, process, outcome, and context), information system-related incident (ISRI) and its characteristics, and four conceptual relations between OL and ISRIS (trigger, source, context, and trigger). Based on this conceptual framework, I review literature on information systems (IS), incident management, and general management to examine how previous works have studied OL from ISRIs. The analysis of the literature shows insufficient attention is given to OL *process* with regards to the characteristics of ISRIs. Accordingly, I develop an overall framework for addressing this gap.

2.1 Definitions

Understanding organizational learning (OL) from information system-related incidents (ISRIs) requires defining organizational learning, its various dimensions, and the way the study operationalizes this concept. Second, it is critical to define information system-related incidents and their characteristics. Third, it is important to elaborate what it means when we talk about OL *from* ISRIs. Box 2.1 summarizes the definitions used in the study.

Box 2.1: Definitions used in the study

Information System-related Incident (ISRI): "An unplanned interruption to an IT service or reduction in the quality of an IT service" (ITIL, 2012): 46).

Organizational Learning (OL): "A cognitive, discursive, and material process through which an organization and its members aim to expand their existing capabilities" (Nicolini et al., 2008; Nicolini, Mengis, & Swan, 2011).

OL "from" ISRIs: ISRIs can act as *sources* of drawing learning content, as *triggers* of learning process, as the *context* in which learning process takes place, and as the *targets* of learning when the aim is avoiding incidents or reducing their impacts (D. Smith & Elliott, 2007).

2.1.1 Organizational Learning (OL)

OL refers to the *cognitive, discursive, and material process through which an organization and its members aim to expand their existing capabilities* (Nicolini et al., 2008; Nicolini et al., 2011). Therefore, OL does not necessarily imply a rational, deliberate process. In fact, OL is often emergent (Nicolini, Gherardi, & Yanow, 2003; Orlikowski, 2002) and situated (Lave & Wagenr, 1991; Orlikowski, 1996) and it may not produce immediate and visible changes (Weick & Westley, 1996). This

definition is broad enough that allows capturing various related works in the literature.

Following the literature on OL, I distinguish between four dimensions: *content*, *process*, *context*, and *outcome* (Naot, Lipshitz, & Popper, 2004). The *content* of OL refers to the specific lessons and solutions that are suggested. It answers the question about *what* is (going to be) learned? (Naot et al., 2004). Learning content often refers to what is intended to be learned (in an ex-ante analysis), and can be of a technical, managerial, or cultural nature (Carroll 1998). For instance, organizations might think about defining a new organizational procedure in their production as a solution that is learned.

The practices undertaken by organizational members to identify and implement a learning solution are defined as learning *process* (Elkjaer, 2004). Investigating the learning process thus answers the question *how* learning takes place. A focus on the learning process brings to the fore concerns about what sort of activities are taken by organizational members, how these activities are structured (Daft, 1982; McKenney, Mason, & Copeland, 1997), what actors are involved (Cohen & Bacdayan, 1994), which roles and responsibilities are defined (Crossan, Lane, & White, 1999), and in what temporal pattern OL process unfolds.

Learning *context* refers to the factors and conditions that affect the learning process. More specifically, the learning context describes the structural (Lam, 2000), political (Blackler, 2000; Coopey & Burgoyne, 2000; Lawrence, Mauws, Dyck, & Kleysen, 2005), cultural (Blackler, 1995), legal (Mayer & Argyres, 2004), and material factors (Epple et al., 1991; Hargadon & Fanelli, 2002) either facilitating or hampering the learning process.

Finally, learning *outcome* refers to the actual product of the previous three dimensions. Learning outcomes can be intended (such as reducing the frequency

of incidents or their impacts), as well as unintended (such as creating new incidents or turbulences) (Epple et al., 1991).

Although these four aspects are intertwined, OL literature has shown that distinguishing between them can help us analyze and theorize on this phenomenon (Fiol & Lyles, 1985). This distinction also allows us to analyze the focus of various studies, how OL is framed in their theories, and avoids confusions due to mixing different OL dimensions.

2.1.2 Information system-related incidents (ISRIs)

Incidents, in general, are defined as *deviations from the expected or routine events that produce undesirable outcomes* (Haunschild & Sullivan, 2002; Mellahi, 2005; Van de Ven & Poole, 1995). These events contradict normal expectations, so, they entail an element of surprise and shock (Lyytinen, 1988; Weick & Roberts, 1993). Although all incidents have rather fuzzy temporal boundaries, they constitute a series of interrelated events that occur within a circumscribed and limited period of time (D. Smith & Toft, 2005; Turner, 1994a; Weitzel & Marchand, 1991). The identification of incidents is based on objective criteria such as the visibility and size of impact, and some subjectively and socially constructed criteria. As such, understanding both what an incident is and what events should be included in an incident account is partially subjective and context dependent (Elliott et al., 2000). Given the complexity of incidents and the subjectivity involved in understanding them, uncertainty is thus an inextricable aspect of all incidents (Carroll & Hatakenaka, 2001). A further critical characteristic of all incidents is their stressful nature. The negative impact of incidents and the need to recover from them quickly both imply great urgency and pressure (Borodzicz & Haperen, 2002; Stern, 2002). Hence, incidents often give rise to severe stressful conditions and external

pressures from stakeholders and public agents, such as the media (Moynihan, 2009).

Information system-related incident (ISRI) is broadly defined as "*an unplanned interruption to an IT service or reduction in the quality of an IT service"* (ITIL, 2012): 46). This definition is selected because it is close to empirical side that allows focusing on ISRIs. Second, the definition is inclusive enough to cover various types of ISRIs that this study concentrates on (security incidents and hardware failure).

The focus on organizational level implies that the negative impacts of ISRIs are defined with respect to the organization's goals and performance. In addition, the definition excludes planned interventions by IT service providers and any kind of interruptions and quality reduction event that is part of the standard operation of a service. From a process view of incidents (Carroll and Hatakenaka 2001), ISRIs are incidents in which information technology play a critical role in one or more incident stages, whether through their creation, incubation, or escalation (D. Smith & Elliott, 2007).

While ISRIs share several characteristics with all other types of incidents (e.g., they are perceived negatively, they have a limited duration, they are stress generating and give raise to the need for urgent remedial action), they also have some specific features that set them apart from incidents in other domains. First, unlike natural disasters or some other human-induced incidents, ISRIs have a strong technical dimension (Ang, Thong, & Yap, 1997). Accordingly, ISRIs are likely to derive from, impinge on, or being related to technical knowledge and professional know-how (Loch et al., 1992; McMullen, 2010).

Second, ISRIs are also transient as rapid technological changes in information technologies can easily pose novel sources of risk for organizations. This implies that the capabilities that organizations need to handle and avoid such incident might also change frequently.

Finally, ISRIs often occur at the interface between a wide range of actors such as suppliers of IT facilities and services (Van Eeten & Bauer, 2009), various users (Perrow, 1999; Turner, 1994b), and numerous potential attackers. These characteristics tend to be simultaneously present and interact in a complex way in most ISRIs (see Table 2.1).

Table 2.1: The Characteristics of ISRIs

Characteristics of ISRIs	**Description**	**References**
Technical nature	ISRIs rely on deep and specialized technical knowledge (IT).	(Ang et al., 1997; Loch et al., 1992; McMullen, 2010)
Transient and emergent	Rapid technological changes make ISRIs novel and give them emergent properties.	(Egan, 2007; Holmes & Poulymenakou, 1995; Strom, 1993)
Heterogeneous actors and institutions involved	Various agents (users, developers, supporters), from different institutional backgrounds, are involved in ISRIs.	(Perrow, 2007; Turner, 1994; Van Eeten & Bauer, 2009)

2.1.3 OL "from" ISRIs

While incident-handling process aims at controlling the negative impacts of the current incident and recovering from crisis, learning process aims at capitalizing on the experience of incident (handling) to avoid similar incidents in future or reduce their impact. This conceptual distinction, however, sometimes is not so explicit in reality, specifically when some practices can serve both incident handling and learning purposes.

Regarding the relation between OL and ISRIs, incidents can be considered as (1) a *source* for drawing learning content; (2) as a *trigger* to drive learning process; (3) as a *context* in which learning process takes place; (4) as the *target* when learning aims at avoiding them or reducing their impacts (Smith & Elliott, 2007).

Incidents as sources of learning

Incident can act as a source of providing insights about what should be learned to avoid similar incidents in future – i.e., learning content. In fact, the revealing nature of incident helps organizations see aspects that should be improved (Deverell, 2009; Smith & Elliott, 2007). Learning in this view involves various actions such as analyzing the patterns of incidents and their causes, comparing them with similar incidents, and simulating the models that are constructed out of the analysis of the incidents to test the effectiveness of the suggested solutions. In this view, frequent and rare incidents act differently in the learning process (Lampel et al., 2009; Rerup, 2009; Starbuck, 2009). Frequent incidents help actors see patterns among them (learning from communalities), detect the cause of incidents by comparing them (learning from variances), and experimenting the effectiveness of their learning solutions in the next similar incidents. However, rare incidents (March et al., 1991) can point to insights that organizations often ignore them. However, organizations might find it difficult to reach reliable solutions only by analyzing rare incidents (the basic problem of sampling and generalization).

Framing incident as a source of learning content, the validity, reliability and generalizability of the findings are key concerns. Therefore, the complexity (both detailed and dynamic complexity (Senge, 1990)) of incidents poses challenges in establishing relations between causes of incidents and their effects. Thus, organizations are prone to superstitious learning (March, Olsen, & Christensen, 1979) (wrong attribution between the incident and the causes of it), incomplete learning (lack of complete understanding of the whole cause-effect chain), over-generalization of lessons learned (to what extent the cause-effect map is generalizable to other similar incidents) (Hedberg, 1981), and the viability of the lessons learned (how long they can use the lessons that are learned from past

incidents without being concern about their obsolescence) (Baumard & Starbuck, 2005).

Incidents as triggers of learning

Incidents, especially when they are large and visible, can create serious needs for learning, justify learning process (Haunschild & Sullivan, 2002), destabilize the existing rigidities, and open learning windows (Keeler, 1993; Sabatier, 1993). The incident drives the learning process as it *draws attentions* (Rerup, 2009) towards improvements, *shows* the importance and need for learning (Elliott, 2009), *marshals* political and institutional pressure on actors to learn (Elliott et al., 2000), *justifies* the allocation of resources to learning process, and provides *excuses* for learning advocators to push for learning (Stern, 2002).

The triggering role of (frequent,) small incidents resides in their frequency, otherwise rare, small incidents might have little force to drive learning process as they can be easily dismissed (Baumard & Starbuck, 2005) and normalized (Turner, 1978). On the other hand, the triggering role of (rare,) large incidents are due to their visibilities and significant impacts that show the criticality of learning from them. In this view, organizations are dealing with challenges such as *filtering and ignoring* the signals of incidents, *late detection* of signals of incident, *normalizing* incidents (Turner, 1978; Turner, 1994a), and engaging in blame-game (D. Smith & Elliott, 2007), that all can diminish the triggering role of incident in the learning process.

Incidents as contexts of learning

Incident can play the role of context for learning process, especially when learning process takes place during or right after incident (Moynihan, 2009). Compared with normal conditions, incident situations involve more emotional excitement (Wang, 2008), pressure from stakeholders (Stern, 2002) to take urgent actions

(Borodzicz & Haperen, 2002), high level of uncertainty, and limited access to information (Carroll & Hatakenaka, 2001).

The most obvious case is when learning takes place during incident process (Moynihan, 2009). However, the contextual impacts of incidents can exist before and after its occurrence. For example, Y2K caused serious contextual pressure on organizations quite a few years before January 2000 (Perrow, 1999). Similarly, the tension and pressure caused by incidents can last for a long period afterwards.

Contextual influence of incidents on learning process can be positive or negative (Stern, 2002). The structural and political instability can provide organizations with flexibilities, which, in turn, can allow for learning that is more substantial. Cultural collapse and the disconfirmation of dominant views can also facilitate cognitive processes of learning (Stern, 2002). Incident can be a proper context in which experts across organizational boundaries can collaborate in learning process. Some information can be revealed mainly through dialogues and debates in such situations.

On the contrary, pressure and urgency caused by this context can result in defensive reactions of actors, leading to more rigidities and political struggles (Staw, Sanderlands, & Dutton, 1981), which, in turn, can block learning process. Limited access to information and emotional excitement can go hand in hand to reduce the performance of individual and collective decision making (Augilera, 1990). Due to political dynamics, some information about incident might be distorted or filtered, making organizations adopt immature solutions or apply superficial remedies instead of proper learning content. Finally, learning in such context might be so past-oriented or focused on the associated incident that organizations forget other learning opportunities.

Incidents as targets of learning

Finally, incidents can be targets for learning, when organizations attempt to avoid future similar incidents or reducing their impacts (Smith & Elliott, 2007). This view is more relevant in situations that organizations have to work under the paradigm of high-reliability organizations, which requires prioritizing preventive learning to avoid incidents before they occur (LaPorte, 1996). In this view, the focus is on potential negative impacts of the incident and avoiding them (learning outcomes). Detection of early warning signals that might indicate some future incidents, preventive actions (e.g. considering redundancies and contingency plans and …) (Rao, 2004), simulating future incidents to learn for them before they occur (Borodzicz & Haperen, 2002), and risk analysis are common learning practices.

2.2 A review of prior studies on OL from ISRIs

Based on the conceptual framework outlined before, this section reviews three bodies of literature that are directly related to learning from ISRIs:

1) The IS literature, especially works that focus on ISRIs and learning from them.

2) Studies in incident management literature that focus on OL from ISRIs.

3) Studies in managerial literature with specific attention to learning from incidents related to IT.

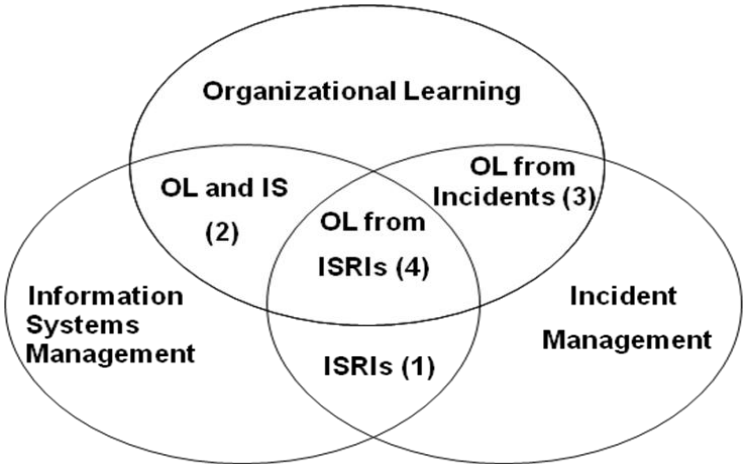

Figure 2.1: Map of the reviewed fields

As suggested by Figure 2.1, the focus of the review is on the study of learning from ISRIs (area Number 4). Accordingly, I *exclude* works that merely focus on ISRIs (area Number 1), general relations between OL and IS (area Number 2), and OL from other types of incidents rather than ISRIs (area Number 3).

2.2.1 Information systems literature

IS scholars and practitioners have been traditionally interested in ISRIs and how to handle and prevent them. Relevant insights on learning from ISRIs can be found in at least five IS domains: 1) IS application failure, 2) IS security and privacy incidents, 3) IS risks and early warning signals, 4) IS incident evaluation, 5) IS reliability and mindfulness. A summary of how learning from various ISRIs are discussed in the IS literature is presented in Table 2.2.

Table 2.2: Summary of insights about OL from ISRIs in the IS domains

IS domain	Main findings related to OL from ISRIs
IS application failure	- Learning from implementation failure is particularly difficult as many are not reported or are reported as quasi-successes. - Applications failures expose organizational weaknesses and tend to be concealed.
IS security and privacy incidents	- Security breach and privacy infringement are seen as increasing threats. - Prevention of security and privacy issues requires a combination of technical and managerial interventions such as better security infrastructures and applications; improved policies; and awareness programs. - Learning from security failures are often framed as a highly structured process of soliciting and disseminating information about threats and new policies.
IS risk & early warning signals	- Focus must be on improving resilience and 'designing out' vulnerabilities. - Major disasters always start with early signs that can be detected if the necessary processes are in place. - Early warning signals provide useful sources of information about wider risks and vulnerabilities.
IS (incident) evaluation	- Evaluation is a critical step for learning from ISRIs. - Ex-*post* analysis as a way of learning - Learning should be a stage at the end of the evaluation process. - Barriers to learning from evaluation of IS project include pursuit of formal compliance; clash of lessons learned with the existing body of knowledge; disincentives for learning; lack of resources and time for reflection; and political maneuvering.
Reliability and mindfulness	- Mindfulness both as a way of learning and preventing future mishaps - Learning as an embedded and emergent process - Importance of learning for maintaining reliability of operations - Focus on practical as well as abstract knowledge

IS Application Failures

Several works have focused on failures that take place during the operation of information systems which generate strong internal and external pressures to recover and learn from them, especially in case of major crisis and catastrophes (Holmes & Poulymenakou, 1995; Huang, Makoju, Newell, & Galliers, 2003; Muhren, Van Den Eede, & Van de Walle, 2007; Weitzel & Marchand, 1991; Westland, 2000). These failures can reveal technical or organizational weaknesses

and surface serious vulnerabilities, which constitute further learning opportunities (Faia-Correia, Patriotta, Brigham, & Corbett, 1999).

A number of scholars have analyzed these failures to extract lessons on their causes during the operation, as well as in previous stages such as IS development (Sarosa & Zowghi, 2005) and implementation (Orlikowski, 1992). Recommendations often point towards technical solutions such as upgrading the technology, reconfiguring it, and adopting new technologies. As for non-technical lessons, scholars advise organizations to create backup teams (Faia-Correia, Patriotta *et al.* 1999), improve communication between developers and users (Sarosa and Zowghi 2005), and align IT with business goals (Sarosa and Zowghi 2005).

IS Security and Privacy Incidents

A topic that is attracting increasing attention is the analysis of security and privacy incidents that take place during the operation of information systems (Culnan & Williams, 2009; Greenaway & Chan, 2005; Johnston & Warkentin, 2010; Kjaerland, 2006; Shedden, Ahmad, & Ruighaver, 2010; H. J. Smith, 1993; S. Smith, Winchester, Bunker, & Jamieson, 2010; Straub & Welke, 1998). Lessons, drawn from these works, point to three overall directions. First, authors suggest the need to invest in information security technologies and infrastructures (Cavusoglu, Mishra, & Raghunathan, 2005; Galbreth & Shor, 2010; Solms & Solms, 2004). Second, authors indicate a number of managerial interventions such as the adoption of information security management programs (Albrechtsen & Hovden, 2010; Solms & Solms, 2004) and improvements in policies, processes, and standards (S. Smith et al., 2010; Solms & Solms, 2004). A third approach is to raise responsiveness by establishing awareness programs (S. Smith et al., 2010; Solms & Solms, 2004), improving the knowledge of managers about attackers motivations (Willison and (Cremonini & Nizovtsev, 2009; Willison & Backhouse, 2006), and

motivating employees to comply with security policies (Bulgurcu, Cavusoglu, & Benbasat, 2010).

While these studies provided a wealth of suggestions for initiatives that need to be taken, marginal attentions are paid to the process whereby such suggestions might take place. Some scholars suggest that organizations should "disseminate information about security actions taken" (Straub and Welke 1998, p.460) or set up feedback loops that link the various stages of the incident response process (Muhren et al., 2007). What is common in all these models is that learning is mainly conceived as an activity to "extract useful information from incidents of all kinds and to use this information to improve organizational performance over time" (Cooke, 2003: 2). Such activity is often described as a sequence of formal stages such as documenting (Westland, 2000), standardizing, and sharing information on the incidents (Kjaerland, 2006).

This view is echoed in techno-policy documents such as ISO2700, where learning is defined as a step in the final stages of incident handling process, consisting of gathering and analyzing information from evaluation and handling security incidents "to identify recurring or high impact incidents" (ISO/IEC27002, 2005: 93). Recently, however, some authors have started to question such dominant view, noting that this view overlooks the importance of feedback timing, the need to facilitate double-loop learning, and taking into account informal learning processes (Shedden et al., 2010).

IS Risk and Early Warning Signals

Another stream of IS literature is shaped around a common finding that major incidents often start with some small errors or signals that could be detected by the organizations (Keil, 1995; Montealegre & Keil, 2000). Therefore, most of these studies identify, classify, and analyze common early warning signals of IS incidents (Ginzberg, 1981; Havelka, Rajkuma, & Serve, 2004; Kappelman, McKeeman, &

Zhang, 2006; Lyytinen & Robey, 1999), so that they can improve managers' ability to detect such signals early enough to take corrective actions and—in extreme cases—abandon faulty projects in their early stages (Lesca & Caron-Fasan, 2008). Alternatively, the lessons learned through risk assessment can feed new content to the awareness programs and improve tools and techniques for risk assessment (Straub & Welke, 1998).

Recently, authors have come to realize that drawing risk management lessons by identifying early warning signals may not be enough as "this does not ensure that actions are or will be taken" (Iversen, Mathiassen, & Nielsen, 2004: 412). Furthermore, applying lessons from past failures to future risks can be limiting and even misleading, since organizations might simply apply old remedies to new and poorly-understood problems. Thus, a small number of studies started to focus on the *process* through which organizations learn from IS risks. The view here is that the very process of detecting and analyzing risks is as critical as the content of the lesson learned. Such process should include a continuous activity of assessment, analysis, and planning for the IS risk. This way, the entire risk management process is considered as a learning process that runs alongside daily activities (Bandyopadhyay, Mykytyn, & Mykytyn, 1999; Spears & Barki, 2010).

IS (Incident) Evaluation

Part of the literature on IS *evaluation* (Nelson, 2005; Poulymenakou & Serafeimidis, 1997; Wilson & Howcroft, 2005), post-mortem analysis (evaluation) (Boddie, 1987; Kasi et al., 2008), and post-project evaluation (Pan, Pan, Newman, & Flynn, 2006) concerns analyzing ISRIs and trying to learn from them. For example, even in the post-project evaluation, one of the key themes is analyzing incidents (e.g., what went wrong) that took place during the project.

Unlike the foregoing domains, this line of research takes a process view. Most of these works approach evaluation as an ex-post activity that (ideally) takes place

right after an IS incident (Kanellis, Lycett, & Paul, 1999; Kasi et al., 2008) and in a formal way (Gwillim et al., 2005; Kumar, 1990). In many cases, although learning from ISRIs is part of this concern, these evaluations do not necessarily take place right after the incident (for example they might be postponed to the end of the project). Moreover, incident learning is often only mixed with the analysis of success events, risks, and other critical events.

These studies emphasize that when well-conducted, the evaluation of IS incidents results in lessons formulated in terms of explicit solutions that organizations should implement continuously after each failure (Poulymenakou & Serafeimidis, 1997). In this sense, the whole evaluation process is sometimes considered as a formal, staged learning process. Kasi, Keil *et al.* (2008) suggest, for example, a four-stage model of post-mortem analysis (identifying the underlying IS project, collecting data on it, analyzing data, and sharing and exploiting the resulting knowledge). Sometimes, however, learning is considered as *part* of an evaluation process (mostly the final stage). For example, Nelson (2005) models evaluation process in three stages "evaluating project performance, extracting lessons learned, and making recommendations for the future" (Nelson 2005: 361), where the last two stages pertain to learning.

The evaluation literature offers several insights on the barriers to learning from ISRIs. This often follows the finding that organizations rarely follow formal evaluation processes (Gwillim et al., 2005) and when they do so, they are often motivated by formal compliance (Kumar, 1990) or tactical reasons (such as justifying decisions and actions that have already been taken (Wilson & Howcroft, 2005)), rather than learning purposes (Kasi et al., 2008). Barriers to the actual exploitation of the results of post-mortem analysis (Gwillim et al., 2005; Kasi et al., 2008) include the clash of the lesson learned with the existing body of knowledge, disincentives for learning, inappropriate structural settings, lack of resources, lack

of time for reflecting on the past initiatives because of the necessity to start the next activity, and political issues (Kasi et al., 2008). The feeling is often that "evaluation will unearth problems better left undisturbed" (Kumar, 1990: 210), an attitude that clearly curbs learning (Kanellis et al., 1999). The time limitation can make learning agents too busy to reflect properly on the failures and learn from them (Kappelman et al., 2006).

Reliability and Mindfulness

Finally, a further small, but growing, body of IS literature is built on the concept of high-reliability organizations (LaPorte, 1996). The approach is important in that it offers a view of ISRIs as emergent and largely unforeseeable events that often cannot be prevented or pre-empted. This is because ISRIs are rooted in the complex, dynamic, fragile, componentized, socially situated, and hence, unreliable nature of information systems (Butler & Gray, 2006). The focus of these studies is thus not so much on the ex-ante forecast of risks or the ex-post evaluation of causes. Instead, they focus on the capacity of organizations to act as mindful and flexible systems during handling these incidents. By analyzing various types of ISRIs such as the unreliable operations of systems (Butler & Gray, 2006) and security problems (Wright & Marett, 2010), these studies show that mindfulness is a crucial element in handling ISRIs and learning from them in an emergent and situated mode (Butler & Gray, 2006).

Organizations that use information systems in their mission-critical activities should therefore organize for mindfulness, as they have no chance to let their members learn through trial and error or experimentation. Improving mindfulness is not only an important way to improve reliability; it can also enhance the capabilities of organizations to learn better from incidents. In this view, learning happens when individuals and collectives mindfully interact with incidents and try to cope with them. Hence, learning is not limited to the process of extracting

abstract lessons through a rational analysis of incidents. As Butler and Gray (2006) pointed out, mindfulness involves "focus on the present, attention to operational detail, willingness to consider alternative perspectives, and an interest in investigating and understanding failures" (p. 212). Acting mindfully gives organizations the ability to learn in a flexible, emergent way, as the nature of ISRIs incidents implies. Thus, the focus on acting mindfully brings to the fore the importance "learning in working" (Gherardi & Nicolini, 2000).

Summing up, the IS literature seems mostly focused on learning content by suggesting a wide range of specialized solutions to avoid various ISRIs. Regarding learning process, there seems to be a division of labor, with literature on early warning signals is focusing on learning *before* ISRIs and IS evaluation literature is concentrating on learning *after* ISRIs. In contrast, the literature on IS mindfulness and reliability contributes to understanding learning a process that happens *alongside* the incident process.

2.2.2 Incident Management Literature

Due to the importance and dominance of IS incidents, considerable attention has been given to this topic in the field of crisis and incident management. In particular, the worries about the Y2K bug and realization of its potential dramatic implications on almost every aspect of society was instrumental in drawing scholars' attention to ISRIs (Perrow, 1999). Although most of the hypothesized catastrophic consequences of Y2K did not materialize, the event triggered intense discussion of IS systems, their vulnerabilities, and the risks linked to IS activities. The debate was given new impetus with the advent of Internet and its associated security and privacy incidents. This trend has continued so far as the result of the

prevalence of a wide range of incidents linked to emergent information technologies such as social networks and cloud computing.

By and large, this body of literature has been informed by an analytical approach. Many of these studies on their own are in fact *post hoc* attempts to understand the risks and vulnerabilities of information systems with focus on incident prevention (LaPorte, 1996; Rochlin, LaPorte, & Roberts, 1987; Schulman, Roe, Van Eeten, & De Bruijne, 2004). Some studies concentrate on the vulnerabilities of critical infrastructures such as computer networks and the Internet (Fritzon, Ljungkvist, Boin, & Rhinard, 2007; Hills, 2005; LaPorte, 1996; Rochlin et al., 1987; Schulman et al., 2004; Van Eeten & Bauer, 2009; Wagenaar, 2009). The cases are usually of a supra-organizational nature. These works are also dominated by social (Perrow, 1999), national (Gorman, Schintler, Kulkarni, & Stough, 2004), political (Eriksson, 2001; Shin & Sung, 1995), and military (Demchak, 1999) concerns. The lessons drawn from the analysis are usually on sectoral, national and global levels, with the focus on governmental and public agencies as the main actors (de Bruijne & van Eeten, 2007). Such lessons, for example, include increasing the collaboration and interaction between public and private sectors for better incident analysis and avoidance; developing policies and programs at sectoral and national levels to enhance companies' preparedness; setting up knowledge-sharing programs between companies (Demchak, 1999).

Particular attention to ISRIs has been devoted by a specific sub-section of the incident management literature, namely the literature on healthcare management. Over the last decade, hospitals and community healthcare services have come to rely heavily on information system to carry out both clinical and non-clinical activities such as queuing, prescription management, and management of patients' information. Today, the most common modern medical devices rely greatly on information and communication technologies to operate

and this has attracted scholars' attention to the study of development and implementation failures of health information systems (Fauchart, 2006). Thus, in addition to ISRIs in medical devices, there have been other mishaps in the use of Health Information Systems (HIS) such as mismatches of information, wrong prescription of medicines (Y. Chen, Neil, Avery, Dewey, & Johnson, 2005), and breaches of privacy (Rangel & Friend, 1995).

While some of the studies in this field simply reflect general interest on OL and learning from incidents (Nikula, 1999; Storey & Buchanan, 2008), others have paid particular attention to ISRIs in healthcare and learning from them (Wallace, 2003). These studies suggest that organizations need to learn from these incidents by analyzing them. For instance, the analysis of previous cases can teach us ways of preventing identity theft (Amori, 2008) and reduce the failure of development and implementation projects by adopting a top-down plan to set a clear framework for interactions between users and implementers and remain flexible to properly react to uncertainties in the implementation process of HIS (Berg, 2001). Finally, some of these studies have taken into consideration the specificities of the healthcare sector as a large, institutionally fragmented, and geographically distributed entity (Southon, Sauer, & Dampney, 1999) and have discussed how these factors affect the way in which healthcare organizations learn from ISRIs.

2.2.3 In General Management Literature

Learning from incidents and failures has been a long-standing interest in the managerial literature on OL (Baumard & Starbuck, 2005; Carroll et al., 2002; Haunschild & Sullivan, 2002; Jasanoff, 1994; Kim & Miner, 2007; Shrivastava, 1988; Starbuck & Hedberg, 2001; Toft & Reynolds, 1992). Incidents are widely considered as valuable opportunities to shed light on the shortcomings of

everyday organizational and managerial processes. Although several of these studies are relevant to the understanding of ISRIs, only a few directly addressed this particular topic with attention to specific characteristics of ISRIs. For instance, Denrell (2003) studied IS incidents and found that a common structural challenge in learning process is the isolation between the learning actors that is, those who are supposed to learn from incidents versus those who are directly involved in handling incidents, mostly due to the traditional isolation of IS departments in large organizations from system users. In general, however, the managerial literature pays scant regard to ISRIs. Indeed, many of these studies are not primarily concerned about ISRIs, unless they use ISRIs as their empirical examples, without being concerned about the specific characteristics of ISRIs.

2.3 Analysis of literatures and what is missing: Insufficient attention to process and ISRIs contexts

As I have briefly shown in the previous section, several strands of IS literature have contributed to understanding how organizations learn from ISRIs. The topic has also attracted minor attentions from authors in incident management and general management literatures. In the present section, I focus on IS literature to identify existing patterns and gaps by mapping reviewed studies into the four OL dimensions –i.e., content, process, context, and outcomes (See Table 2.3).

Table 2.3: The mapping of IS literature unto OL dimensions

IS research domain	OL dimensions			
	Content	**Process**	**Context**	**Outcome**
IS application failure	- lessons for future development and implementation - Lessons on business operation continuity - Lessons on incident management process	- Learning as keeping records of errors and reporting them	- Time-pressure - The concern of reliability - The fear of publicity of the incident and external pressure	- Reliability of business process - Fewer incidents with less damage - Maintained organizational image - Faster recovery and business process continuity - Detecting and fixing the defects in organizational processes and structures
IS security and privacy incidents	- Improvement of security infrastructure - Improvement of security and privacy policies, procedures, and culture - Improvement of information security programs - Enhancing security awareness - Adapting incentive systems for both insiders and outsiders	- Learning as a feedback stage at the end of security planning	- Intentional and political concerns in creating incidents	- Fewer incidents with less damage - Maintained organizational image - Faster recovery and business process continuity - Detecting and fixing the defects in organizational processes and structures
IS risk & early warning signals	- Common IS risks and their causes - Improvements in risk identification, analysis and response capabilities - Improvements in other organizational aspects - List of key and common warning signals - Attention to the early and fast response and problem solution	- Learning as a stage at the end of the risk management process		- Preventing IS risks - Avoiding the escalation of small incidents to critical incidents
IS (incident) evaluation	- Suggestions and guidelines for more effective IS evaluation - Lessons on choosing appropriate criteria for evaluation	- Evaluation as a formal and staged learning process - Learning as a stage in IS evaluation vs. learning as an embedded activity in the whole IS evaluation process	- The overall political context of evaluation - Time pressure and lack of organizational slack for spending on IS evaluation	- More effective IS evaluation process - Deeper understanding of the problems and solutions related to the incidents
Reliability and mindfulness		- Learning as an emergent practice in incidents handling - The importance of early and proactive learning process - The importance of informal and daily learning practices - The importance of practical learning as well as cognitive learning	- motivational involvement as a necessary context - The emergent and unpredictable nature of ISRIs - Complexity of ISRIs and learning from them	- Being prepared and resilient in addition to preventing attempts - Maintaining reliability

As the second column in Table 2.3 shows, most of the reviewed studies focus on drawing lessons from the analysis of incidents. ISRIs are approached mostly as sources for extracting learning *contents* by identifying what organizations have learned or should learn to avoid such incidents in future. These lessons bear on understanding the characteristics of ISRIs, when they might occur, the main common causes of these incidents, their early warning signals, and how managers can predict, prevent, or successfully handle them. Although some of these lessons are drawn by organizations and practitioners, most of them are the results of analytical reflection by IS scholars.

The disproportionate focus of literature on learning content underscores the assumption that these learning contents are generalizable to similar cases and can be extended into future. However, this assumption is questionable given that organizational members interpret incidents in different ways (Lyytinen, 1988). Hence, organizations might perceive apparently similar incidents in such different ways that leaves no room to capitalize on earlier lessons learned by the company or by others. Furthermore, accelerating changes in IT and organizational environment goes hand-in-hand with the inherent ambiguity and uncertainty of incidents, leaving little room for such generalization. Overemphasis on drawing lessons from past can also lead to rigid and outdated incident response behavior.

The overemphasis on the study of the content of learning from ISRIs matches the description and promise of a number of *outcomes*, which should stem from the applications of the recommendations presented in the literature (last column in Table 2.3). These outcomes have been presented in terms of: avoiding future (similar) incidents, and repeated errors and mistakes (Eriksson, 2001; Fortune & Peters, 2005; Irani, Sharif, & Love, 2001; Kappelman et al., 2006; Kasi et al., 2008; Kjaerland, 2006; Shedden et al., 2010); reducing the likelihood of the future failures and incidents (Butler & Gray, 2006; Demchak, 1999; Eriksson, 2001;

Muhren et al., 2007; Raymond Caron, Jarvenpaa, & Stoddard, 1994; Salaway, 1987; Straub & Welke, 1998); mitigating the negative impacts of IS failures (Holmes & Poulymenakou, 1995; Straub & Welke, 1998); enhancing the organizational capabilities for better incident management (Shedden et al., 2010); and achieving a higher level of resilience (Butler & Gray, 2006).

Only a handful of studies indicate that learning practices can target prosaic and instrumental outcomes such as getting formal approvals (Kumar, 1990), obtaining some form of accreditation (S. Smith et al., 2010), enhancing market value (Bharadwaj, Keil, & Mähring, 2009), and improving external image (Gordon, Loeb, & Sohail, 2010). In addition, very few articles examine the negative and unintended outcomes of learning from ISRIs such as the tendency of managers to become hidebound and seriously skeptical of technological changes as a result of analyzing past incidents (McKenney et al., 1997).

When IS scholars have approached the study of learning *process*, they have done so in two main ways (third column in Table 2.3). First, many studies consider learning as a specific formal stage that often takes place at the beginning or at the end of IS activities. For instance, works in the IS risk management field recommend operational models with a learning stage at the end of the risk management process. This step should allow organizational members to reflect on their experience so that similar situations can be better managed in the future. A similar approach is also taken in literature on IS evaluation (Raymond Caron et al., 1994; Scott & Vessey, 2000). The main limitation of these studies is that they reduce learning to a rational process based on identifying and articulating lessons and applying them in future. As I mentioned earlier, this view is underpinned by the somewhat simplistic assumption that when the investigation is conducted in a rational way and the results are effectively disseminated, the lessons will be taken up by the organization and become part of its existing procedures. Therefore, the

mechanisms through which lessons learned are taken up and incorporated (or not) by the organization thus remain unexplored. This is particularly consequential as the OL literature is rife with examples of the so call “knowing-doing gap”, that is the incapacity of organizations to turn existing knowledge into concrete change and performance improvement (Pfeffer & Sutton, 2000).

While the abovementioned view is strongly aligned with the emphasis on the learning content and have a normative orientation (what the learning process should be), a second and different view conceives learning process as an embedded element that occurs *alongside* incident handling process, rather than framing it as a separate, formalized, and rational stage. This view matches approaches in organizational literature that consider learning as an inevitable feature of all work activities (learning-in-organizing: see (Gherardi, 1999; Gherardi & Nicolini, 2000)). A good example of this new, and still minority, view is found in works addressing the issue of IS reliability by promoting mindfulness. Butler and Gray (2006), for example, conceive the learning process as a series of emergent and daily practices that take place throughout incident management. In their view, the learning process is not reduced to specific and easily identifiable attempts to analyze and reflect on incidents. Instead, the learning process mainly consists of mindful practices of organizational members that take place during incident handling process. The benefit of this approach is that it can grasp the heterogeneity, imperfection, and emergent nature of learning process. The major limitation is that it does little to explain how local learning is absorbed by the entire organization as a whole.

As noted above, a further dimension is the *context* in which learning takes (or fails to take) place (fourth column in Table 2.3). Several cognitive, cultural, and structural factors affecting the capacity to learn from ISRIs have also been considered in the IS literature. The cognitive background of top managers and

employees, especially in IT (Straub & Welke, 1998), the way in which managers and experts frame IT as a threat (Eriksson, 2001), the experience of past failures (Raymond Caron et al., 1994; Robey, Boudreau, & Rose, 2000), the supportive climate for individual learning (Wastell, 1999), and the overall culture of mindfulness (Butler & Gray, 2006) have all been indicated as important factors affecting the capacity to learn from the experience of adverse events. Structural factors such as the quality of measures and feedback systems on the ISRIs, disincentives for learning (Lyytinen & Robey, 1999) and the lack of resources for formal learning processes (Iacovou & Dexter, 2005; Sarosa & Zowghi, 2005) have also been identified as important factors in determining whether organization will learn from incidents or not.

Apart from the general contextual factors, little attention has been paid to examining the characteristics of ISRIs in relation with learning *process*. In fact, recognizing specificities of ISRIs has been well reflected in learning *content* by providing specialized suggestions with regards to various types of ISRIs. Nevertheless, most of the studies in IS literature have not systematically examined whether and how the specific characteristics of ISRIs might influence the way in which organizations learn from their incidents.

We argue that the overemphasis on content and outcome runs against the recent insights from the broader literature on OL, which suggests that attention to *how* organizations learn from incidents is at least as important (if not more so) than focusing on *what* lessons can be drawn from specific incidents. This is so for at least two reasons. First, the peculiarities of each incident and organizational context often make it difficult to transfer past lessons to new cases. Accessing information on historical cases, while is important, does not tell the organization how this information can be translated into actionable changes. Second, unless organization establishes specific learning mechanisms, each incident risks being

addressed individually and some of the lesson that could be learned would go unnoticed. Thus, learning contents and lessons from previous incidents are useless unless organizations appropriate them and this leads to an expansion of their capacity to act. As I noted above, learning is not only about acquiring information and knowledge and requires also putting the right mechanisms in place.

2.4 Towards a framework for studying OL from ISRIs

As mentioned in the previous sections, the existing studies on organizational learning from ISRIs rarely pay close attention to how specific characteristics of ISRIs may affect learning process. The literature has a rather generalist flavor, the assumption being that general theories of organizational learning are immediately applicable to the study of learning from ISRIs. This contrasts with the results of studies in the managerial OL and incident management literature, where it has been shown that the contextual characteristics of incidents and the particular conditions in which specific organization operate work together to define the opportunities and constraints for learning. Therefore, the question that remains is that how does learning process from ISRIs look like. What is specific to learning from ISRIs? Are there any specific practices that exist in learning process from ISRIs? (And if there are specific aspects of learning process) what are the specific contextual factors that contribute to the way in which organizations learn from their ISRIs?

More specifically, although ISRIs share a series of characteristics with other types of incidents (Fitzgerald & Russo, 2005) they also possess specific characteristics which set them apart from incidents in other sectors such as healthcare, workplace safety, and the nuclear industry. Based on the foregoing discussion, the

study aims at providing a *contextualized* understanding of learning *process*, that allows both articulating what organizations do in their attempt to learn from ISRIS (learning process) and how these practices are shaped by the context of ISRIs.

Accordingly, I rely on the situated learning view of organizational learning (Handley, Clark, Fincham, & Sturdy, 2007; Lave & Wagenr, 1991) which "sees learning and knowing as processes which are integral to everyday practice in workplace, family, and other social settings" (Handley et al., 2007): 174). First, situated learning view focuses on the very social practices, which constitute the learning process. Compared with traditional cognitive views of learning that focus on the content of learning (Handley et al., 2007), situated learning brings to the fore the role of participants, their actions, and interactions. Second, another core assumption of situated learning perspective is that learning takes place in the very context of day-to-day practices. This is aligned with our intent to examine how *organizations* and their actors learn in their organizational context (not for example in the off-work training contexts). In addition, it fits with our aim to study learning process in the context of incidents where learning practices are influenced by the incident factors. In this way, situated learning view helps us consider context not as taken for granted factors, but as the factors that interact with the very learning practices (affect and are affected by them). This suits our aim in examining the relation between the contextual factors and learning process.

Situated learning view has been powerfully adopted in studies that have focused on learning process by focusing on social practices as the constitutional elements of this process. For studying learning practices, I rely on a specific *practice* view of organizational learning (Nicolini et al., 2003; Orlikowski, 2002; Silvia, 2001) that implies understanding OL process requires examining the actions are taken by various social actors, the interactions between them, the timing and temporal

pattern of those practices, the way these practices are organized in organizational settings, and the material aspects of such practices.

Accordingly, I suggest an overall theoretical framework for the study (see Figure 2.2). At the center of the framework, there is the learning process, which constitutes of the actions that organizations take to leverage incident experience to avoid similar incidents and manage them effectively in future.

The learning process is framed within a broader learning context. The framework distinguishes between three categories of contextual factors. First (from outside), learning process takes place in general organizational factors that are present regardless of having incident. Studies on organizational learning have shown the importance of several contextual factors such as psychological safety (Edmondson, 1999), political factors such as the distribution of power among learning actors (Gwillim et al., 2005; Lawrence et al., 2005), and structural features of the organization (Daft, 1982; Lam, 2000).

The second category of factors relates to general characteristics of incidents, such as the urgency, external pressure, negative content, complexity, and revealing nature of incidents (Baumard & Starbuck, 2005; Lave & Wagenr, 1991; Smith & Elliott, 2007). Considering these factors is critical to examine how learning process from incidents might be different from non-incident learning process such as learning by doing (Haunschild & Sullivan, 2002).

The third category of contextual factors pertains to the characteristics of ISRIs, including the technical nature, transient nature, and the heterogeneity of actors involved. The focus on these characteristics helps examining how the specific characteristics of ISRIs might contribute to the way in which organizations go through the learning process.

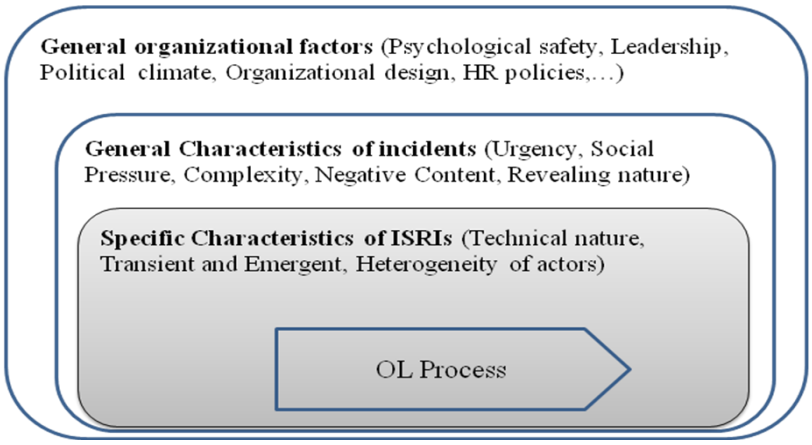

Figure 2.2: the overall theoretical framework for studying OL from ISRIs

The framework implies that empirical study should not only consider the learning process, but also the relations between learning process and the contextual factors. It also draws attention to different contextual factors, especially the characteristics of ISRIs. More specifically, by distinguishing between the general organizational factors and other contextual factors, the framework suggests examining how the observed learning process is affected by the context of incident (compared with general factors that exist in the absence of incidents as well). The distinction between general incident factors and the characteristics of ISRIs, the framework helps examining the contribution of specific characteristic of ISRIs to the learning process (compared with learning from any other type of incidents).

Chapter Summary: In this chapter, I defined organizational learning as a *cognitive, discursive, and material process through which an organization and its members aim to expand their existing capabilities* (Nicolini et al., 2008; Nicolini et al., 2011). Then, ISRIs were defined as "*an unplanned interruption to an IT service or reduction in the quality of an IT service"* (ITIL, 2012): 46). I discussed that the incident can be framed as the source, trigger, context, and outcome of learning.

Then, I reviewed how literature on information systems, incident management, and general management examined OL from ISRIs. The review showed that a major gap is understanding how the characteristics of ISRIs might act as part of the context in which OL process is shaped. Adopting a situated learning approach, I developed a framework that frames learning process as a collection of social practices that are embedded in various contextual factors. The characteristics of ISRIs are also framed as part of the context in which learning process takes place. The next chapter explains the research question and how it is empirically studied.

Chapter 3: Research Methodology and Settings

As indicated in the previous chapter, the research tries to understand how the process of learning from ISRIs unfolds, with regards to the characteristics of these incidents. This chapter operationalizes the research question and specifies the unit and level of analysis. It then describes the selected research design and data collection and analysis process. It is then followed by describing four empirical settings.

3.1 Research Question

The study aims at understanding OL process from ISRIs, with regards to the characteristics of ISRIs. Previous studies have shown that learning process from large incidents can potentially be different from learning process from small incidents (Baumard & Starbuck, 2005). Unlike small incidents, large incidents possess stronger negative content, pose more serious pressure on organizations, have more visibility, especially to a wider range of internal and external actors, potentially involve more diverse social actors and technological aspects (Deverell & Hansén, 2009), and their low frequency affects learning from them (Christianson et al., 2009; Rerup, 2009). This study merely focuses on *large* incidents because the characteristics of ISRIs are visible in such settings.

Literature has also indicated that organizations might follow different learning processes when they are learning from their own incidents (Baumard & Starbuck, 2005) versus when they try to learn from others' incidents (Kim & Miner, 2007). This can be, for example, due to the differences in terms of access to knowledge about incidents and incentives to learn from them (Haunschild & Sullivan, 2002; Kim & Miner, 2007). In fact, several socio-cognitive processes mediate the influence of external incidents on learning process. This study only focuses on *internal* incidents of organizations, because the characteristics of ISRIs that are happening inside the organizations can have more direct impact on the learning process, compared with situations that they take place in other companies.

This leads to the following research question:

How do *organizations learn from their internal, large Information system-related incidents?*

Regarding the dearth of studies on the process of OL with regards to the context of ISRIs, I will focus in particular on the learning *process*. Accordingly, the main unit

of analysis (Babbie, 2009) in this research is learning process, which has been defined as a collection of practices that organizations take to capitalize on the experience of their major ISRIs to avoid them or reduce their impacts in future. In other words, the operationalization of learning process is through learning practices. Learning process involves both individual and collective practices (Silvia, 2001). In examining learning practices, my attention will be focused on the *actions* taken, their associated *actors* and their *intentions*, the *situations* in which they have been taken, their timing, their involved *materiality*, and their (immediate) *consequences*.

The study focuses on two types of ISRIs that take place during the operation of information systems: (1) the failures of the operation of systems, such as hardware or software breakdowns; (2) security and privacy incidents. Both categories possess the concerned characteristics –i.e., technical nature, transient and emergent, and heterogeneity of involved actors and technologies.

To select large incidents, I only consider incidents that have been noticed by top managers of the organizations. This assures that incidents have been critical for organizational goals and performance, not simply being minor, local incidents.

3.2 Research Methodology

3.2.1 Overall Research design

Following the overall aim of this research, which is a contextual understanding of organizational learning process, regarding the nature of research questions ("How" questions), and considering the unit of analysis (learning *process*)

qualitative approach is selected (Denzin & Lincoln, 2005b; Eisenhardt, 1989; Lee, 1999). Qualitative approach fits the study because I am interested in deep understanding of a social phenomenon in its contextual setting (Eisenhardt, 1989). In addition, understanding the *complex* nature of learning process can be better attained through qualitative inquiry because it allows collecting and analyzing rich data about social phenomena (Miles, 1979). Through qualitative methods I can better grasp the role of contextual factors (Downey & Ireland, 1979; Yin, 2002) which is crucial for analyzing the influence of ISRIs as part of learning context. Finally, learning process is a dynamic phenomenon, which should be studied through methods capable to capture this dynamics over time (Langley, Smallman, Tsoukas, & Van de Ven, Andrew H., 2013).

Among different philosophical paradigms of qualitative research (Denzin & Lincoln, 2005a; Guba & Lincoln, 2005), an interpretative approach (Schuetz, 1953), which can be considered as a moderate epistemology between positivism and constructivism, fits the research framework. First, interpretative approach is capable to consider subjective aspects of human practices such as their meanings and purposes. Second, interpretative perspective not only captures objective aspects of social phenomena (e.g., the objective impacts of incident and the learning activities), but also admits the influence of subjective factors in shaping social phenomena (e.g., learners interpretations of incidents). Therefore, the outcomes of interpretative approach can be easily transformed into rich hypothetical theories that are deeply rooted in empirical data (Corbin & Strauss, 1990; Eisenhardt, 1989; Glaser & Strauss, 1980).

3.2.2 Multiple-embedded case study

Within the adopted methodological approach, case study (Eisenhardt, 1989; Yin, 2002) can be an appropriate research strategy. Case study approach aims at understanding a particular case deeply, with regards to its context. Case study helps examining the role of contextual factors on the learning process (Stake, 2005). Thus, it allows focusing on specific characteristics of our cases (Stake, 2005; Yin, 2002) to capture the specificities of learning practices in relation with ISRIs.

Case study also suits studying *processes* such as organizational learning (Stake, 2005). It is because the researcher can collect data on the very chain of events, actions, and interactions, with particular attention to their temporal aspects (their time of occurrence, their sequences, their duration, and their overall temporal patterns). Collecting rich data about all these aspects is critical to outline the details OL.

There are several options in designing case study (Yin, 2002). Case study can be exploratory (to openly explore the phenomenon) or explanatory (to look for the causes of a previously identified phenomenon). I rely on an *exploratory* approach since the aim is to openly examine how OL process unfolds in ISRIs context. This way, I am open to observe all sorts of actions, interactions, and actors that constitute the OL process.

I rely on a *multiple* case study design because it allows for exploring different types of ISRIs and enriches the contextual understanding of learning process through capturing the diversity of learning practices. I am not going to compare the selected cases (it is not a comparative case study). Instead, looking at different cases will help finding various patterns of OL process. In addition, multiple case study design reduces the risk of case attrition, which is specifically important for a sensitive topic such as learning from incidents. Through multiple case studies, I can

reach deeper understanding of the influence of ISRIs on learning process by paying attention to differences across cases (Eisenhardt, 1991; Yin, 2002).

Regarding the unit of analysis (learning process), and the level of analysis (organizational), the most aggregate level of the study is organizational. Hence, through an embedded-multiple case study design (Yin, 2002) I study multiple ISRIs and associated learning processes in each organization.

I analyze selected cases longitudinally to capture the dynamic nature of learning process (Leonard-Barton, 1990). It means that for each learning process, I will create the *story* of the learning process, consisting of the sequence of actions, interaction, events, and consequences. This longitudinal view helps me build a process data (Langley et al., 2013; Langley, 1999) through which I can make sense of the learning process.

However, the negative content of ISRIs might limit a real-time data collection. Therefore, a retrospective mode of inquiry is adopted to collect data (Leonard-Barton, 1990) around incidents that although were critical at the time of occurrence, they are not sensitive to be studied any more. More specifically, I focus on specific major incidents that happened in the last three years and I examine what the organizations did after the incident in order to avoid it in future or reduce its impact in future. This frees me from organizations sensitivity about current incidents. In addition, I can use the existing documents and reports that the organizations have developed during and after incident handling process.

3.3 Research Process

3.3.1 Case selection criteria

Regarding the research aim, purposeful sampling of cases is guided by two criteria: (1) their business should be IT intensive (any interruption in their IT services for several hours would damage the company's performance and be noticed by the top managers); (2) they must have experienced sever ISRIs in the last three years. The first criterion assures that in the selected cases, ISRIs are critical and deserve paying attention to them to be learned from. Therefore, the selected cases are more likely to show various practices related to learning from incidents, because such incidents are important for them to be prevented and controlled effectively. This way, I can increase the richness of cases in terms of observing the learning process.

The second criterion allows concentrating on specific incidents and examining what organizations did to learn from them. Three years is a tentative estimate of a proper period in which access to data is feasible and incidents are not so recent that makes organizations hesitated to reflect on. To confine the study at organizational level, the selected cases are in the same industry (IT services), and in the same national context (Spain) to get rid of the variability of factors at supra-organizational level.

The literature on ISRIs has points to two broad categories of ISRIs: the failures in the operations of information systems and IT security incidents. To increase the richness of empirical settings for observing various learning practices, two different categories of the companies are selected: two supercomputer centers that are mostly dealing with IS operation failures; and two IT security companies that mainly deal with handling security and privacy incidents.

3.3.2 Access

Due to the sensitivity of the topic (although my focus is on learning from incidents, not the incidents themselves), I started exploring a wide range of companies in various sectors. I got access to key actors in 10 organizations from which at least one initial interviews was conducted (out of more than 25 potential companies in financial, automotive, retail, utility, education, IT, petroleum, and manufacturing industries that showed some initial interests). In some cases, I had several interviews and visits before deciding to consider them in the analysis. Some cases stopped the process of data collection for political reasons (e.g. sensitivity of the topic, fear of media, and fear of revealing information) although in all cases I assured them about the confidentiality of all information and anonymity of results. This appeared when the access to the company was denied suddenly, and the given explanations were alien or unclear (e.g., "the top managers preferred to stop any further data collection").

Some other cases were dropped because the informants were too busy to be interviewed. Three of the cases were also excluded because after initial interviews, it appeared that the cases are not rich enough in terms of the types of incidents and learning process for further exploration. This process of filtering the cases allowed me to focus on each case with sufficient depth. I ended up with four organizations.

3.3.3 Data collection process

Data collection started in June 2011 and finished in October 2013 (16 months). Data was collected through interviews with informants, document analysis, and observations. Overall, 41 interviews were conducted with managers (at top and

middle levels), senior and junior experts, and employees who were involved in the process of handling incidents and learning from them. Interviews, by average, took 64 minutes. All interviews were semi-structured, based on a customized interview protocols according to the informants positions in the organizations, their roles in incident handling and learning process, their backgrounds, and their personal characteristics (as far as it was available before the meetings). Initial interviews with each informant were organized around three major sections: 1) understanding the role of the informant before, during, and after incident, 2) the story of the incident from his/her point of view, and 3) inquiring about what actions had been taken during and after incidents to avoid such incident in future or reduce its impacts. I asked specific questions about various aspects of learning process based on the interview protocol (see Appendix 1).

All the interviews were voice-recorded (with the consent from the interviewees before hand). I documented the reflections, personal notes, nonverbal messages, doubts and wonders within 24 hours after each interview (except for 4 interviews). This reflection helped customizing successive interviews, as well as arranging for follow-up interactions with the interviewees for complementary data collection. Interviews were transcribed verbatim. Around 20% of the interviews were conducted in Spanish. Ambiguities and doubts about some Spanish phrases were discussed with a native Spanish Speaker who is also an expert in management studies. Table 3.1 presents a summary of collected data in each case.

Table 3.1: The summary of cases and collected data

Organization	Description	Number of Interviews	Covered informants	Other sources of data
Security-Public	Small, public security agency of an autonomous community government	8	CEO; Top managers; Technical managers; Their provider; Technical experts	Documents of the creation and performance reports; Visits of their systems
Security-Private	Small, private company specialized in IT security services	11	CEO; Program manager; Project Managers; Functional managers; Experts	Documents; Observations of artifacts, Attending follow-up meetings; Visit of facilities; Observing them when they are working
Supercomputer -Large	Large, pubic supercomputer center	18	Department Manager; Group Managers; Experts	Website; Observations during interviews; Visits to the facilities,
Supercomputer -Small	Small, public supercomputer center	4	Department Managers; Middle managers; Project managers; Team managers; Senior experts	Documents; Websites; Movies
Total		41		

Another important source of data was the documents that all the organizations have provided about incident handling process, and their actions afterwards. Since all details of the incidents, actions applied to them, and internal projects are documented through various systems such as ticketing system, internal Wikis, and other information systems, I had the chance to get rich data about the selected cases. These documents specifically helped in several aspects of data collection. First, they provided background information about each company and the selected incidents, before conducting interviews. This saved many interviews that would otherwise be focused on getting background information. Second, it helped in developing questions that are more specific and focusing merely on issues that are not in the documents. Third, the documents were helpful for collecting more details about the issues discussed in the interviews. For instance, managers were

talking about several internal projects in their interviews. Then, the documents about the internal projects were consulted to glean more details. Fourth, the documents also allowed me to cross validate the information collected through interviews. In some cases, it resulted in identifying missing information in the interviews and some misunderstanding of documents that then was resolved through further interviews. Table 3.2 describes various types of documents, and their contribution to our data.

Table 3.2: The summary of documents analyzed for data collection

Document type	Contribution to data
Online documents (websites, public press)	information about organizations background, how they work, their internal structures, their external relations, their major changes over time, and their major incidents
Tickets in the ticketing system	Information about the incidents and actions applied to handle them
Wiki Articles	Information about lessons learned from incident analysis, the solutions for avoiding future incidents, and other potential sources of information to be consulted in case of similar incidents
Incident Analysis reports	Information about the causes of a specific incident, its potential solutions, and follow-up actions
Checklists	The details of activities in handling incidents, and (the changes in the checklists show) the changes made in the procedures due to the experience of incidents
Users manuals	Information about how users can work in such a way that do not face some incidents or be less affected by some incidents
Performance reports	Information about improving the quality of work as the result of experiencing some past incidents
Managerial reports	Decisions for some improvement projects and the decisions made for making some investments for technology improvement and organizational changes

Various online documents helped collecting information about the background of the organizations, their internal procedures and structures, their external relations, their major incidents, and their changes over time.

The content of tickets was a major source of information for getting detailed data about the selected incidents. The companies did not share all their tickets, but they did show the content of tickets related to several selected cases. This was helpful in getting detailed information about the incidents and the chain of events and practices during incident handling process.

Wiki articles provided information about the lessons that the organizations learned as the result of handling incidents and the solutions that they have proposed and implemented to prevent similar incidents or reduce their impacts. Sometimes, the wiki articles also showed some ideas that they have been thinking about for the sake of improving their work

Some companies have regularly created incident analysis reports for major incidents in which they were analyzing the specific incident, its causes, and the solutions they have provided, as well as the follow up actions that can help them avoid similar incidents or improve their work.

Some checklists about how to react to an incident and how to manage recovering from it revealed the specific actions that the organizations have learned as the result of experiencing past incidents. The changes in the checklists (through comparing different versions of them), showed such improvements.

The organizations have often developed and revised the users' manuals to help them work more effectively with their systems. The comparison between versions of each manual was helpful in pin pointing learning from past incidents.

Performance reports that the organizations have developed for their stakeholders have been also helpful in spotting changes in technologies, routines, and procedures. Some of the changes were done as the result of facing incidents.

Finally, several managerial reports that technical managers were developing helped identifying solutions that were proposed to managers and solutions that have been approved by the managers in order to avoid some incidents in future.

We also arranged for several on-site observations, in each case, to observe how the organizations work in their daily work, how they use different tools and artifacts, and what kind of documents they produce. These observations provided a detailed and contextual understanding of cases and raised several questions in terms of differences between what has been mentioned in the interviews and what was observed. This, in turn, resulted in further inquiry. Observations also allowed me to enter the communities of experts inside the company, building trust with them, and become able to grasp informal aspects of their work. Table 3.3 summarizes the various on-site observations done in each of the four companies and the focal points in the observations.

Table 3.3: The summary of observations done for data collection

Organization	On-site observations and the studied elements
Security-Public	• A whole day staying with various groups and department • Visiting their incident analysis laboratory • Observing their ticketing system and how they use it • Observing their wiki system and how they use it
Security-Private	• Observing their ticketing systems and how they use them • Navigating through their wiki system and observing how they work • Staying in their open seminar sessions as a silent attendant
Supercomputer -Large	• A tour to the main facilities, how it works, their technologies, and the related systems • Two day living with them in the company and chatting with experts and managers while they were working • Observing the hardware and software technologies they use and damaged technologies in some incidents • Observing how they work with the wiki systems in their incident handling process • Observing how they use the wiki system in their daily work
Supercomputer -Small	• A tour to the main facilities, how it works, their technologies, and the related systems

Finally, I had several informal interactions with informants through off-site meetings (mostly for taking coffee and lunch before or after interviews or in the transition time between two successive interviews). Several points have been mentioned during these interactions and some sensitive questions were asked and discussed then. These informal meetings were not voice-recorded. However, I kept detailed notes of the discussions in my reflection on the interviews right after each session. Some of the issues mentioned in these informal meetings allowed me to better interpret the assertions in the interviews. In some cases, it helped me formulate new questions and new lines of inquiry.

Several strategies were used to enhance the validity of empirical findings. First I stuck to critical incident interview approach, by focusing on indentifying critical events (e.g. an incident, a specific change, a key decision, and specific actions made by actors). I extracted the timeline of the actions, events, and consequences, following the sequence of critical incidents. I constantly asked about what has been really done, distinguishing them from what has been merely intended or decided to be done. I asked about the intentions and mentalities of actors in taking actions. I also asked about the personal interpretations of the situations, which then helped me in understanding and interpreting the opinions.

Second, I cross-validated main findings, ambiguous findings, and contradictory assertions by asking from at least two informants. In some cases, I kept asking questions from almost all informants. I also used complementary data sources (such as documents and observations) to cross validate the findings.

Third, the results of the each case study were presented to a selected group from the same organizations' managers to validate the findings. Throughout these sessions, I exposed ambiguous aspects of the case study, and potential lacking information. In addition, in each reflection session, I formulated questions about the observations that I had in other cases or are expected to be observed based on

the literature, but it was not observed in the first round of data collection. The meetings helped me fill gaps in the data and define complementary data collections in each case.

The data collection in each case continued until either I covered all the involved actors in the learning process or I reached to a point that further interviews showed little potential of revealing new insights. This was detected when repeated issues appeared in the last interviews and most of the findings were already mentioned in previous inquiries. Of course in Security-Public, the limited access to informants restricted collecting further data. In this case, the analysis focused on incidents and practices that I could collect sufficient data about them.

Data collection process in each case (organization) went through three steps (see Figure 3.1).

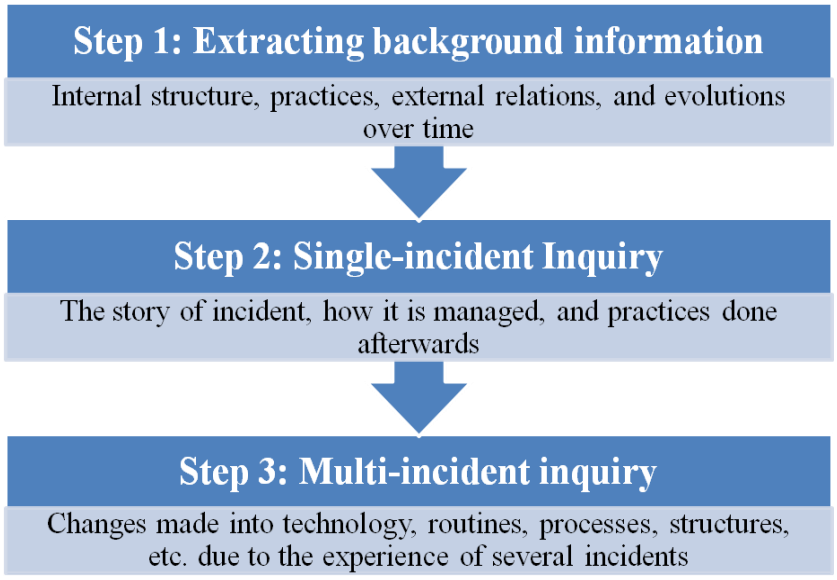

Figure 3.1: data collection process

Step 1: Extracting background information

In the first step, I tried to understand how each organization works. Data was collected about their internal structures, processes, services, procedures, human resources, financial aspects, technologies and artifacts, and their cultures. Data was also collected about the external relations of the organizations with their providers, clients, partners, and governmental and industrial institutions. This data provided background information of each case. At the end of this step, several major ISRIs were indentified in each organization that the organizations were willing to reflect on them. Overall, 15 major incidents were identified in the four organizations.

Step 2: Single-incident inquiry

Following the overall research questions, in the second stage, the inquiry focused on the identified incidents. In this stage, I concentrated on each of these incidents to examine what the organizations did after the incident to avoid it or reduce its impact in future.

The incident-based inquiry allowed for identifying practices that the organizations adopted as the result of facing that incident to reduce its likelihood or impact. Therefore, I could examine the relation between the incident and learning practices (the idea of learning *from* incidents). In addition, each major incident acted as a reference point to judge about the learning outcome and learning intentions of actors for each practice. More specifically, in order to be considered as a learning practice, either the adopted practice must have some impact in terms of reducing the likelihood of incident and its impact or at least learning actors have intended such learning impact.

In the incident-based inquiry, I focused on each incident and adopted an inclusive approach in collecting data about the incident itself, the incident management

process, and other practices that the organization adopted as the result of experiencing that incident afterward. This way, for each incident, I developed a rich narrative of all actions, interactions, actors, events, challenges, and contextual factors over time. These narratives included both the period during incident handling process (from the time that the incident is noticed by the organization to the time that it is controlled) and the time after incident handling process.

In these narratives, some specific practices were identified as learning initiatives that the organizations adopted as the result of each of the identified incidents to reduce the likelihood and impact of future incident. These actions, called *incident-specific learning accounts,* were specific to each incident. Upon identifying the learning accounts, I examined the process through which those learning accounts were identified and implemented. Table 3.4 summarizes the identified incidents in the four organizations and their associated learning accounts.

Table 3.4: Incident-specific learning accounts

Organization	Studied incidents	Incident-specific learning accounts
Security-Public	No specific incident was mentioned (due to the sensitivity of the organization); instead the organization referred to four incidents as typical incidents	Increasing the efficiency of delivering services Adopting more proactive approach in ERI[1] Expanding in-house ERI activities & improving the organizational level of ERI Adding technical facilities for mobile incidents Changes in the ticketing system (queues, labels, categories of incidents) The change in catalogue of services Clarifying the limits of the capabilities for delivering services The changes in the training programs Eliminating some internal processes or services that proved to be incapable
Security-Private	2 big Incidents in the web applications	Creating a CERT[2] for a client and moving the team to the site of client Creating an internal permanent project related to CERT of the client
Supercomputer-Large	Central Storage problem (moving to the new storage)	Training the users about the universal storage Creating and improving quasi-real-time mirroring process Improved general knowledge of managing mirroring process and other issues
	Spam filter server	Creating an application for reporting directly to the user about not delivered email
	HSM backup system	No specific action in this case, except documenting it in the ticketing system
	IP6 compatibility incident	Almost nothing, just some awareness about the possibility of similar incidents in future
	Temperature alarm and calibration	Adding new item to the regular maintenance
	Evaporating system	Making several technical changes to physical facilities
Supercomputer-Small	Cooling problem	Changing the mentality of the Parent University about the criticality of their infrastructure Taking the ownership of the infrastructure for the new site Defining and continuously improving the shut down process
	The storage problem (ZSF)	Developing index of files (where is where) for fast reporting of damaged files)
	The storage problem	Basically nothing special, expect documenting the incident in the ticketing system

[1] ERI refers to "Incident Response Team"

[2] CERT stands for "Computer Emergency Response Team"

Step 3: Multi-incident inquiry

Along with identifying incident-specific practices, I also identified practices that although were adopted for preventing and reducing the impact of future incidents, they were not necessarily related to one single incident. In fact, these practices were adopted as the result of experiencing several incidents. I named these practices as *multi-incident* learning practices.

I could identify these practices when I asked general questions such as "what have you done to avoid these kinds of incidents or to reduce their impacts?" and "what are the ways in which you assure that you would experience fewer of such incidents in future and with less damages?" These questions (that were not referring to any specific incident) helped me identify various practices that the organizations adopted for learning purposes. Therefore, in the third step, I adopted a change-based inquiry to extract the details of learning activities. As the result, I could identify 13 *multi-incident learning accounts* in the studied organizations (see Table 3.5 for details of the general learning accounts). These learning accounts, though were informed by several incidents, they were not necessarily taken after any single specific incident.

Table 3.5: Multi-incident learning accounts from ISRIs

Learning accounts	The description of learning process in relation with the related incidents
Security-Public	
Adopting more proactive approach in the incident response team	Over time, after experiencing a wide range of security incidents, the key experts in incident response team (ERI) got some free time during their annul strategic planning for the new year in 2012. They identified a pattern that last years, they have been acting more kind of a passive incident chaser. They decide to act more proactively, by not waiting until incidents happen. Instead, they should adopt a threat-detection approach that allows them to foresee future security attacks. They defined a project to implement this approach, as it required deploying network scanning tools, establishing procedures for monitoring social networks of suspicious groups, defining new associated queues in their ticketing system, and defining and negotiating associated roles and responsibilities inside their security service provider (in the form of a new type of service).
Expanding in-house incident response activities and improving the organizational level of incident response department	Several security incidents happened during 2011 that involved sensitive information of politicians. Following the overall outsourcing strategy, the organization delegated them to its provider. There were some critical moments that those sensitive users were quite concerned about the publicity of their incidents. Over last year, there has been a gradual accumulation of experience and interest inside the organization to revise this overall strategy and define criteria for identifying sensitive incidents that should be handled internally. This took place when the new CEO entered. He assigned his deputy to lead this initiative. As the result, a new (unusual) category of incidents "level 4" was defined in the ticketing system that corresponded to this type of incidents. Accordingly, the procedures and rules inside the organization and between the organization and the provider were defined. This change became more established by promoting the position of incident response team as one of the main divisions in the organization.
Change in the service model in relation with the provider: from body-shopping model to volumetrias	Security-Public has been working with its provider based on body-shopping model (paying based on the number of hours that provider works on the defined services). However, after 1.5 years working, and because both Security-Public and its provider learned how to deliver services more effectively and efficiently, the organization realizes that there are a lot of services that can be defined as a specific unit of service (volumetira) and be outsourced with a specific price. When the new CEO came to the company, the team of managers started reflecting on a wide range of incidents and "how" they were managed by the provider, their relations with the provider, and the way they were contracting with the provider. Learning from the pattern of various services was needed to help the Security-Public and its provider automate their tasks, reduce the prices of services, and identify what tasks cannot be automated. Accumulation of a wide range of incidents allowed for extracting the patterns and seeing the big picture of how they were managing their services, what services were lacking, and what sort of interactions they lacked in the provider's side. Therefore, they defined a project for changing the service model in the contract with their provider in 2014, when their previous contract will finish. This project involved defining categories of services, setting quality measures, negotiating prices, and defining new rules and procedures for interactions with provider.
Adding technical facilities for mobile incidents	A series of incidents revealed the lack of capacity in analyzing and handling mobile-based incidents. For the first few cases, Security-Public tried to rely on the existing capabilities and expand the use of available applications to solve the issues, but as far as it became a consistent and increasing need, it justified adopting various tools, specialized for mobile-incidents and learning how to handle them. This led to a series of actions for acquiring new equipments, defining procedures, establishing mechanisms, defining new categories of services and queues in the ticketing system, and defining roles and access rights.
Adding legal services as a new category of services to the service catalogue	After facing a series of incidents with various legal implications (criminal cases and fringing laws), Security-Public realized that a series of legal services should be defined and provided as part of the incident management process and as distinct preventive services. This took place when the legal expert who was leading the legal activities inside the organization spared some free time to reflect on various requests and incidents related to legal themes, identify a series of key legal services, define them, and put them into practice.
Periodically delegating more issues to the provider	Often new, critical incidents are handled internally by Security-Public. However, the organization has periodically identified a pack of incidents that can be outsourced to the provider. This happens when Security-Public gets enough experience and confidence about how to handle these incidents, find a kind of consistent pattern among them, and making sure that they do not involve any sensitive issue that should be kept in-house. However, Security-Public does this often when the load work decreases and the managers and experts of various divisions find time to articulate the pack, define some overall quality measures, and negotiate them with the provider's delegate.

Learning accounts	The description of learning process in relation with the related incidents
Security-Private	
Changing the service model (from corrective to preventive mode)	For a specific client, Security-Private starts working as incident handler. After one year, Security-Private not only handled the incidents, but could detect a wide range of vulnerabilities in client's system that could cause serious damages. Using this experience, Security-Private leveraged its good performance and suggested changing the overall service model with the provider to act for preventing incident by working closely with the application providers of the client to make sure that they commit security principles during their development projects and when they try to implement their application.
Creating a network scanning tool out of the solutions that were temporarily developed for the incidents of clients	In several tough incidents happened to various clients, Security-Private had to examine a wide range of suspicious network nodes that were distributed globally over the internet. It was needed to do so for identifying the scope of the incident, the possible sources of attacks, and potential future threats. Security-Private has done several local actions in different projects. Some were like developing short software codes (scripts) to run such a scanning, some others were mostly doing some manual inspections. When similar case emerged in a project with a big client, Security-Private had enough resources (and freedom) to develop a tool that helped scanning the nodes automatically. Although this tool was developed in that specific project, Security-Private was seeing it as its own tool that can be used as a technology in other projects. However, there were some debates between Security-Private and the client about the property right of the tool.
Supercomputer-Large	
Developing script to automatically handle the problems of quotas in the case of central storage movement	The need for moving data from one of the main HPCs to the new central storage and the temporal situation for keeping both original and copy files, caused some serious problems when users wanted copy large files. This could cause data loss, and wasting a lot of time and processing resources. After a couple of initial incidents, Supercomputer-Large started a project to develop a script that not only solves similar problems in future, but also could incorporate several other improvements that had been appeared out of the experience of some other incidents before.
Creating an application for reporting directly to the user about not delivered email	Supercomputer-Large has been constantly faced with problems related to its Email-servers. Although most of these problems could not be prevented (because they required a totally new application or some of them are quite normal), the big problem from the view of users was that they were not informed about the incidents (e.g. not delivering their emails). This issue was handled manually, case by case, when the users were noticing this through other ways (for example the colleague calling them to ask why the email is not sent). The System-Administrator department defined and internal project to develop a tool to be installed on the email-server application that can provide detailed report to the users about incidents. This project was scheduled as one of the internal projects, and once the other more urgent projects were finished, the two experts in this domain started working on it.
Improving the messages to the users about the failures of backup (HSM) system	Backup systems for storing large files in long term (HSM) that are partly mechanical often are prone to damages and failures. Any damage in the system can corrupt data, and might cause in permanent data loss. Although not so often, but Supercomputer-Large had suffered from various cases in which the backup system failed during storing data, but the user (who was copying data in the backup system) was not able to detect the incident during the copying process because the message that the system gives to the user is quite general ("input-output error") without indicating what is the error, and where the copying process is failed. This led the expert of HSM with the head of System-Administrator group defined an internal project to improve this situation. They first went through the provider of the system, and the provider lunched an internal project to deal with this issue. But soon after, it appeared that the provider is not capable enough and because this is just a very peculiar case, this issue was not a priority for the provider. Therefore Supercomputer-Large started working on the tool internally, and finally after testing various methods, could develop a script that could provide specific message to the user of the system about why the system has failed, and which specific part of the system is damaged.
Supercomputer-Small	
Defining and continuously improving the shut down process	Break of electricity is common (happening once every two months, more or less). This is mostly unpredictable, and depends on the host university's infrastructure. Supercomputer-Small can only manage the shut-down process to avoid burning its equipments and avoid losing data and jobs running in the systems. The shut-down process should be done in less than 5 minutes. There are numerous systems, applications, and network nodes that have complex relations. The sequence of switching off each equipment is very critical because it can cause the loss of data and job (imagine that you shut down the core processes of a distributed process first, before you close the peripheral processes). The difficulty is that the relations between the systems change

Learning accounts	The description of learning process in relation with the related incidents
	almost quickly as new equipments, applications and new configurations of the system appear quite frequently. For the first cases, Supercomputer-Small handled the shut-down process manually. However, the accumulated experience and the increasingly level of difficulty, and experiencing several incidents during this process led Supercomputer-Small to develop a script that tries to automatically shut-down the systems. This script needs to be updated continuously as the result of other changes in the systems.
Developing index of files for fast reporting of damaged files	Working as part of a bigger network of supercomputers, Supercomputer-Small can retrieve the lost data from copies in other centers. Previous incidents that caused data loss were always handled in this way. Although retrieving lost data requires time and energy (for finding the copies, reprocessing them, and transferring them), it is often not so difficult. However, in a recent case, one of the main file systems of the Supercomputer-Small was damaged during the copying process which could cause potentially several peta-bytes of data loss (a "nightmare"!). This time, although more than 90 percent of data could be retrieved from the damaged file system, it took several days to do so. During this time, Supercomputer-Large had to report the total (possible) data loss for being recuperated from other sources, because it was not clear which parts of the overall file system were damaged. This required a huge amount of time and energy for recovering the potentially lost data. This experience, and using the experience of other cases led Supercomputer-Small to define an internal project to dynamically store the physical address of various parts of data files, so that, when similar incident happens, they can easily determine which parts of the data is damaged and needs quick recovery from other sources.
Taking the ownership of the infrastructure for the new site	Supercomputer-Small had suffered from numerous infrastructural incidents (electricity breaks, cooling system damages) mostly because the host university that owns the infrastructure is not considering their infrastructure as a critical onse. The managers of Supercomputer-Small have tried to change the mentality of university managers to convince them that their infrastructure is critical. However there are various formal and physical limitations to deal with the infrastructure of Supercomputer-Small differently. This led Supercomputer-Small to define an overall strategy to set up its own computer room, with the ownership of the infrastructure. This resulted in a series of projects to equip the new room, based on direct contracts with providers. In doing so, they considering most of the experiences of past infrastructural incidents in designing the new room.

3.3.4 Data Analysis process

Data analysis went through five steps (see Figure 3.2).

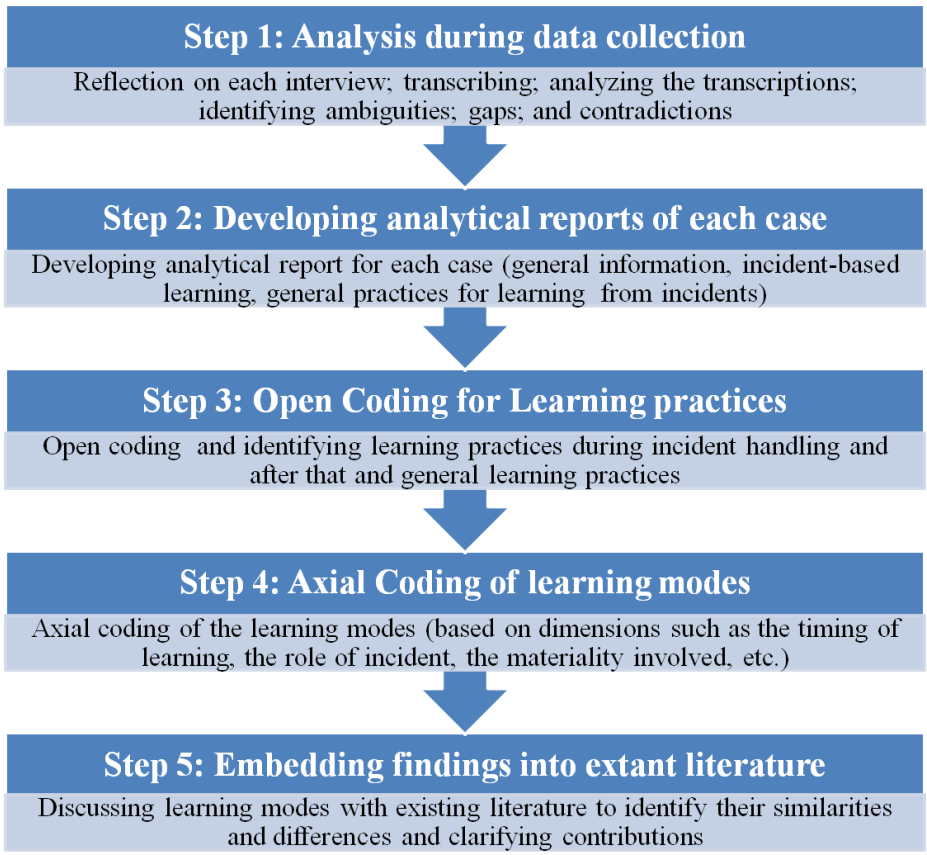

Figure 3.2: data analysis process

Step 1: Analysis during data collection

During data collection process, the interviews were voice recorded and transcribed verbatim. Spanish interviews were then translated into English and checked for potential misunderstandings by consulting a native Spanish speaker. Parallel with data collection, I ran the initial analysis of each interview (sometimes several interviews together) to 1) check for ambiguous themes (which resulted in follow-up inquiry and customizing successive interviews), 2) detect contradictory findings (which became part of the further inquiry), and 3) identify lacking information regarding learning process (which were translated into specific questions and points of inquiry in the following data collection activities). This

analysis was done often within 48 hours after each interview and in some cases within one week.

Step 2: Developing Analytical reports

Upon completing data collection for each organization, all collected data (i.e. interviews, observations, documents, and reflections) were integrated to compose a complete and detailed analytical report. This resulted in four analytical reports for the four organizations. The attempt was made to include all the details. Hence, no summarizing or rewording was done on the original data. The analytical reports were created with specific attention to different sources of data (e.g., whether they are from interviews or from official documents, or from informal documents, or are personal reflections of the researcher) by using different colors and footnotes to specify data sources. Small sentences or phrases were used as sub-sections' titles. The reports were then checked to make sure they were complete and clear. In some cases, this reflection led to some additional data collection.

Each story had three main sections. First, I described how the organization works, including its external relations, internal structure, processes, routines and systems used in their work. This information helped me understand the context of each case. Second, I focused on specific major incidents in each organization to describe what and how the organizations learned from these incidents (listed in Table 3.5, above). This section covered both actions during and after incident handling that the organizations took to prevent similar incidents in future or reduce their impacts. These actions were specific to one incident (*single-incident learning accounts*). Third, I articulated the other practices that organizations took as the result of experiencing several incidents (*multi-incident learning accounts*).

Step 3: Open Coding

The third round of analysis was done based on each of the four analytical reports. In this step, I first openly coded data (Corbin & Strauss, 1990). Following the overall principle of open coding- i.e., the process of "breaking down the data into distinct units of meaning" (Goulding, 2002): 76) - I analyzed the analytical reports line by line to identify meaningful themes regarding the overall research question. More specifically, the coding process was sensitized by considering the actors who were involved, the actions they made, their intentions, the consequences of their actions in terms of avoiding or managing future incidents, the characteristics of ISRIs, the relations between the practices and the incidents, and other contextual factors.

In addition to coding, numerous memos were developed in order to reflect on concepts that emerged out of the analysis, potential relations between the concepts, and the role of some contextual factors in the learning process. I used ATLAS.ti (version 7.0.82) to organize better the data, codes, and memos. This step resulted in 212 codes and 229 memos.

The initial list of codes was then classified around more abstract concepts that they emerged throughout the reflection on the codes. These overall concepts were iteratively refined, to be clearly defined. Accordingly, some codes were clarified, some codes were merged, and some codes were broken down into more specific sub-codes.

As shown in Table 3.6, the codes are classified around nine major categories: (1) Learning practices ("OL practices"); (2) Cognitive aspects of learning practices ("Cognitive"); (3) Structural aspects of learning practices ("Structure"); (4) Politics aspects of learning practices ("Politics"); (5) ISRIs Characteristics ("Incident"); (6) Relations between ISRIs and learning practices ("Relation"); (7) Contextual factors ("Context"); (8) Temporal aspects of learning process ("Time"); (9) Objects in the

learning process ("Objects"). The second and third columns in Table 3.6 provide the definition of each category and the list of associated codes.

Table 3.6: Categories and associated definitions and open codes

Category	Definition	Related Open Codes
Organizational Learning Practices	Practices related to learning from incidents, including the actions, inactions, intentions, interactions, and actors	OL-Actors OL-Agenda of learning OL-Alternative Learning paths OL-Apprenticeship OL-Automatization OL-Barrier OL-Change in identity and self image OL-Collaborative OL-Competing learning items OL-Cost OL-Cross-incident analysis OL-Detached learning process OL-Digitalization vs. codification OL-Documentation for formalization OL-Eliminating obsolete OL-Embedding into artifacts and systems OL-Experimentation OL-Formal / Formality OL-Fundamental vs. remedy OL-improvisation OL-Inaction OL-Incentive / motivation OL-Incident Analysis OL-Incident reporting system / post incident reporting / post hoc analysis / reflection on the incident OL-Intention OL-Incident Simulation OL-Learning from others´ experience OL-Learning accounts / learning solutions / learning cases / learning content OL-Learning Base OL-Learning Conflicts OL-Learning trajectory / complementary OL-Open sessions follow up sessions OL-Overlearning OL-path avoiding OL-Path dependency OL-Perception / framing incidents and incident handlers police vs. doctor OL-Reducing the possibilities for complaining OL-Roles and Responsibilities OL-Routine learning / learning curve / learning by repeating and doing OL-Slack OL-Specialization of OL based on Cognitive boundaries OL-Standardization OL-Training OL-Trigger OL-Trying not to learn: Grey zone of learning OL-Turnaround OL-Vulnerability Analysis

Category	Definition	Related Open Codes
Cognitive	The way that actors understand incidents and their understanding change over time	Cognitive-accumulation of K from learning Cognitive-Abstract Knowledge Cognitive-Attention / Attention Management Cognitive-Awareness Cognitive-codified Knowledge Cognitive-Contextual Knowledge Cognitive-Detailed vs. transversal Knowledge Cognitive-Embedded and Embodied and embrained Cognitive-Expert vs. novice Cognitive-Heuristic Knowledge and learning Cognitive-Holistic Knowledge Cognitive-Know-how Cognitive-Knowledge Characteristics Cognitive-Knowledge Codification Cognitive-Knowledge Gap Cognitive-Knowledge Integration Cognitive-Knowledge lack Cognitive-Knowledge Localized Cognitive-Knowledge loss Cognitive-Knowledge Obsolete Cognitive-Knowledge Ownership Cognitive-Knowledge replication Cognitive-Knowledge slack Cognitive-Knowledge Specialization Cognitive-Knowledge Structure Cognitive-Knowledge Superficial Cognitive-Knowledge Superficial vs. Deep Cognitive-Knowledge Validation Cognitive-Know-who Cognitive-Lessons Learned Cognitive-Reliance on others´ Knowledge Cognitive-Specific Knowledge / Atomic Knowledge / Pieces of Knowledge Cognitive-Subjective Knowledge Cognitive-Tacit Knowledge Cognitive-Tentative Learning Tentative Knowledge Cognitive-Trick Cognitive-Uncertainty
Structure	Aspects that are related to the stable patterns of social actions, such as routines, procedures, processes, roles, incentives, responsibilities, etc.	Structure- Formal affiliations Structure-Dominant Logic Structure-Informality-flexibility in structure Structure-Legal learning Structure-Motivational gap Structure-Need for Learning Structure-Resource limitation Structure-Responsibility of L Structure-Rigidity Structure-Rules and regulations that facilitate learning process Structure-Stabilizing / institutionalizing / fixing / establishing Structure-Structural Gap
Politics	Issues related to the relative power of actors and the dynamics of their power	Politics- Detour Politics- Passing the ball to other´s field Politics-Control and surveillance Politics-Informal and personal relations with key actors Politics-Justifying and Legitimizing Politics-Key Actors Politics-Political support Politics-Prestige Politics-Removing the possibilities for claiming Politics-Scarce Resources Politics-Sensitive decisions / sensitive issues Politics-Status and Legitimacy

Category	Definition	Related Open Codes
Incident	The labels, types, and characteristics of ISRIs	Incident- Types and categories Incident-abnormal Incident-Characteristics Incident-complexity Incident-concern about incident Incident-Criticality Incident-Desirable incidents that show market is expanding Incident-Dramatizing Incident-from legacy systems Incident-Heterogeneity-actors Incident-Heterogeneity-Knowledge Incident-Heterogeneity-suppliers Incident-Heterogeneity-Technology Incident-Incidents as a result of learning Incident-Intentional Incident-Intentional incidents and smart changes in the way they are formulated Incident-Label Incident-Mistake Incident-Novelty Incident-Pain Incident-Pattern of incident Incident-Potential Incident-Proactive detection Incident-process Incident-Rare incidents Incident-The potential of incident for learning changes Incident-Visibility
Relation	The relation between ISRIs and learning practices	Relation- the link between learning and incident Relation-Learning for Relation-learning from Relation-Learning Opportunities
Context	Factors that are stable during the learning process and affect learning process	Context-Free services Context-HR incentive system Context-lack of ownership Context-Legal Context-Low Demand Context-luck Context-Open Source and cost free learning
Time	Temporal aspects of OL process	Time-Absolute Time-Amount of time / as a resource Time-Coincidence Time-Concentration Time-Fast change of technology Time-Frequency Time-limitation Time-Moment of learning Time-Sequence Time-Slack Time-Speed / pace of OL Time-Temporal distance/gap Time-temporal Resonance Time-temporality vs. durability
Object	Any physically observable element related to learning practices	Object-Article- An executable entity Object-Decommissioning Object-Interoperable Object-Learning space Object-Material gap Object-Materiality and Artifact Object-Redundancy and slack Object-ticket and ticketing system Object-Ticket vs. Knowledge article Object-Ticket-Interaction Object-wiki Object-wiki-incident specific articles Object-Wiki-K articles

Step 4: Axial coding of learning practices

Upon the development of open codes, in the next step I focused on learning practices (as first-order codes). I Followed the axial coding procedure – i.e., "moving to a higher level of abstraction and is achieved by specifying relationships and delineating a core category or construct around which the other concepts revolve" (Goulding, 2002:78)- to make sense of learning practices and develop patterns among the learning practices. In doing so, I concentrated on the open codes related to "OL practices". I abstracted them into categories of practices. The codes related to the relation between practices and incidents ("relation") were helpful to come up with two major categories of learning practices (second-order codes): *single-incident* learning practices (described in Chapter 4) and *multi-incident* learning practices (described in Chapter 5). The single incident learning practices were then classified into two sub-categories depending on the temporal aspects of learning practices ("time") which are: learning *during* incident handling process and learning practices after incident handling. Figure 3.3 summarizes the learning practices.

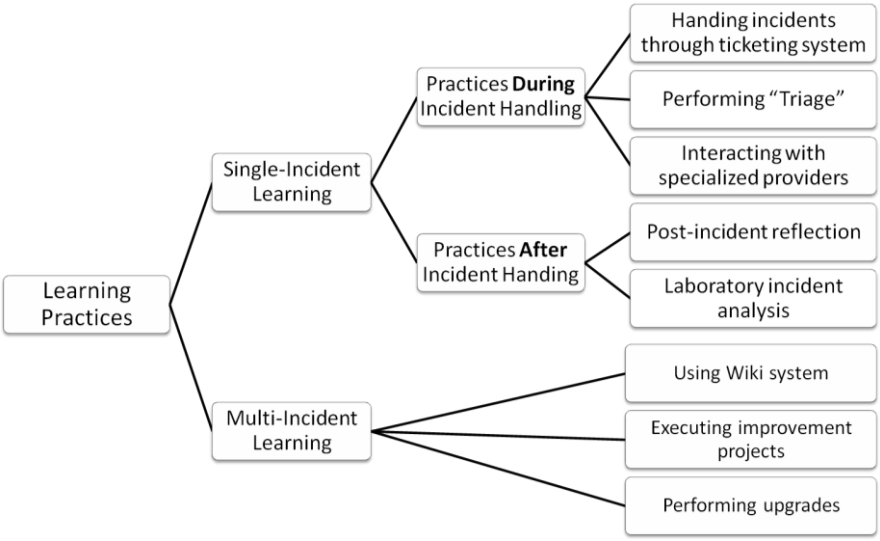

Figure 3.3: Learning practices and their categories

I described each learning practice by defining what are the action that various actors took in performing that practice, the interactions and challenges in doing so, the temporal aspects of those practices, the cognitive, structural, and political aspects of those practices, the relations between the practices and past incidents, the conditions in which the practices were observed, and the role of objects in the practice.

As an important part of this step, I examined how the characteristics of ISRIs were affecting the emergence and evolution of learning practices. These findings were discussed with several scholars in the domain of organizational learning to make sure they are relevant and clear enough. Several comments from the scholars (such as what were the conditions, what contextual factors were present, how the findings were different in different cases) resulted in several iterations with data.

Step 5: Embedding findings into extant literature and articulating learning modes

In the last step of analysis, the articulated learning modes from ISRIs were embedded into extant literature to ground the findings in the previous studies, identify similarities and differences, and explain them. Comparing the learning practices identified in step 5, with learning practices in the literature, five patterns of learning practices were identified that are, called *learning modes*: 1) learning through incident handling practices, 2) learning through post-incident reflection, 3) transversal learning from incidents, 4) outsourced learning, and 5) learning through material replacement (described in Chapter 7). The relation between the learning practices and the five learning modes is depicted in Figure 3.4.

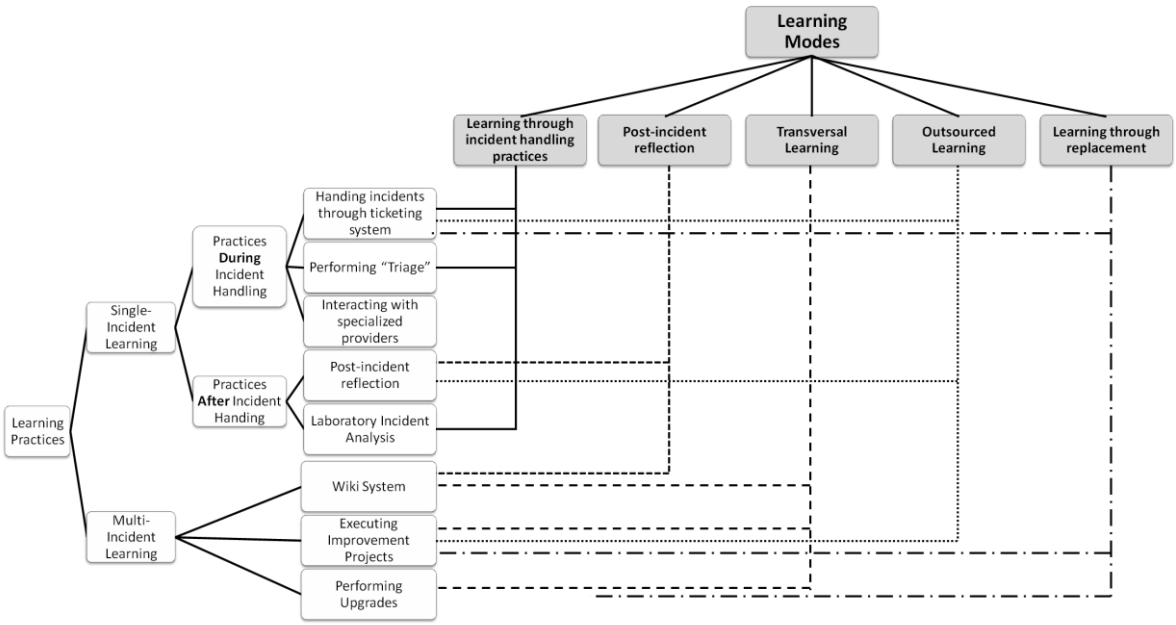

Figure 3.4: Learning modes and their relations with learning practices

As illustrated in Figure 3.4, several learning practices are contributing to various learning modes. For instance, the outsourced learning mode is abstracted from the pattern of outsourcing incident-handling practices, the pattern of outsourcing some post-incident reflections, and the pattern of outsourcing some improvement projects. Moving from describing the learning modes to explaining their emergence in different situations, I could reflect on the role of ISRIs characteristics in shaping learning process (Chapter 5).

3.4 Ethical issues

The study followed the ethical research principles, namely autonomy, informed consent, privacy and confidentiality, beneficence, and justice (Minnesota, 2003; Orb, Eisenhauer, & Wynaden, 2001). In line with autonomy I did not force any company to take part or remain in the research. All companies were voluntarily

participating in the study. As mentioned before, several companies were dropped from the research process due to their hesitation in giving some information and fear of potential harms. In these cases, the data collection was immediately stopped when the company expressed its wonders.

As informed consent principle implies, all the informants (both managers and experts) were completely informed about the aim of the study, the supporting institution, the details of research process, the final outcomes (e.g., thesis, papers, and reports), and the way that the results might be published. This was done firstly at the beginning of studying each company, and was applied to each informant. In case of interviewing an informant several times, each time this process was repeated. I informed the interviewees before the interview sessions, so that they could decide whether to take part in the study or not.

Pursuant to privacy and confidentiality, all the companies and informants are anonymous. Even in cases that some companies consented to publish their names, I did not do so. In addition, I offered each informant and the managers of each company to read and check the interview transcriptions and other research outcomes before being used in my data analysis. This was to allow them to check for privacy and confidentiality concerns. As for two security companies, I signed non-disclosure agreements with specific terms and policies that the companies suggested. Before each interview, I asked the interviewee about the possibility of voice recording the meeting. In a few occasions, the informants asked me to switch-off the recorder when they were providing some information, which I committed to that. Finally, I did not ask about any private issue or about any information that might threaten or create potential harm to the companies and the informants.

Regarding the beneficence principle, I offered the companies a final practical report about the findings that might be helpful for improving their internal

practices. I also offered them a workshop on learning models that these organizations could benefit from in establishing new learning practices. I also openly offered the companies any other academic outcome that they might find relevant for their practices.

Finally, I tried to minimize the unnecessary meetings with informants by consulting available documents. I did not force any temporal limitation on the meetings and other interactions with the informants.

3.5 Challenges and limitations

There were three main challenges in the research process. First, although the research focused on learning process, its relation to incidents made it sensitive for the organizations to reflect on. This is reflected in the large number of case attritions (from around 25 cases to eventually four ones). I tried to reduce this challenge by being open to the target industries, while still focusing on cases that are IT intensive. Fortunately, the eventual cases that constitute the empirical data were committed to providing required information.

Second, the sensitivity of topic, namely inquiry about incidents of the organizations was limiting the chance of collecting data prospectively. Therefore, the study focused on retrospective data collection to inquire about learning from incidents that were no longer sensitive to the organizations at the time of the research. Several methods helped increasing the richness of data and assuring the validity of the findings. First, the selected incidents were in the last 3 years, to make sure that still the actors have a fresh record of what happened in those cases. Second, I focused on major incidents about which the organizations have more detailed records such as tickets, reports, meeting notes, and wiki articles.

Third, some selected incidents were dropped in the data collection process and then in data analysis due to insufficient details.

The third challenge was related to capturing learning practices. Since the concept of learning has been used and discussed among managers and experts, this could affect the findings, if I would have used this term in my inquiry. Thus, the data collection focused on inquiring about *what did the organizations do* as the result of the experience of major incidents. This way, I could capture a wide range of practices that then were analyzed based on the learning definition. Although I captured many other practices than learning practices (e.g., routine incident handling practices), it avoided filtering out nuances of learning process and focusing on what some social actors might subjectively consider as learning.

3.6 Research settings

The empirical study is based on four organizations. Two companies are responsible for handling security incidents, one of them is public and the other is private, referred to as Security-Public and Security-Private[3] respectively. The other two companies are international supercomputer centers. Both are public and hosted in local universities. The large one is called Supercomputer-Large and the small one is named Supercomputer-Small. This setting allowed me to limit the variation on the type of incidents and working context that could influence learning process. As for the two security companies, the focus has been only on the security incidents. The supercomputer organizations did not face any serious security incident, but

[3] For the sake of confidentiality, we use pseudo names.

suffered from hardware and software incidents. The following sections describe the background of the organizations, how they work (e.g., internal processes and routines), and what types of incidents they have faced.

3.6.1 Security-Public

Security-Public is a public foundation dedicated to handling security incidents, delivering security services to public organizations, fostering the development of strong and competitive information security industry, and increasing the resilience and preparedness of the various communities (businesses, citizens, and public organizations) by providing preventive security services. Facing various security incidents and the fear of potential negative impacts of security incidents in future made the government of a Spanish autonomous region defined an IT security plan in 2009 and established Security-Public as the main body that is responsible for implementing the plan in the same year. The government owns 51% of its share and the rest is held by other nonprofit organizations. Its parents include the shareholders and five other organizations that represent various stakeholders such as industries and universities.

Security-Public has a main contractor for delivering most of its services. The main contractor has sub-contracted many of its duties to a specialized IT security company (we call it the provider). On the daily basis, Security-Public is dealing directly with the sub-contractor. The provider was a key actor in the process of defining the IT security plan and designing and creating Security-Public since 2007 (two years before the creation of Security-Public). The provider was disconnected from the process during 2008 when local government was doing internal negotiations for approving Security-Public and allocating resources. Later on, the provider did not take part in the first bid for outsourcing services by Security-

Public. However, as the winning company (the main contractor) was not able to deliver services at the expected quality level, the provider came back to the process as the sub-contractor, in late 2009.

In 2010, the new general director of Security-Public was appointed. He called for a change in the services model with the provider to move from body-shopping model to service-based contracts. In the body-shopping model, which was in place since the creation of Security-Public, the basis of payment to the provider was the number of hours that each level of experts from the provider would spend on delivering services. However, budget limitation, linked with increasing demand for services, went hand in hand with the capability of the provider to improve the quality and efficiency of delving various services, making Security-Public suggest the new service model in which the payment is based on the number of services that are delivered in each category of services. This process has been going on with various challenges due to defining the categories of services, setting prices for each category, and negotiating the processes.

Internal structure, processes, routines, and systems at Security-Public

Security-Public has a rather flat structure. It has outsourced the delivery of its services to specialized IT security companies. There are six middle and top managers affiliated to Security-Public and there are around 21 experts from the provider that work for Security-Public.

Security-Public provides two main categories of services: proactive and reactive. The proactive services include vulnerability analysis, generating notices about security threats, awareness and training programs, and promoting the development of IT security industry. There are 8 persons involved in the provision of proactive services. Reactive services relate to handling incidents that are relayed to Security-Public or Security-Public detects them. These services are delivered based on the overall incident handling process articulated by Carnegie

Mellon model. In this area, there are three levels of experts. Level 1 includes employees from the provider who are receiving incidents, log them, and introduce them to the ticketing system. In level 2, the experts (mostly from its provider) deal with more sophisticated issues that are escalated from level 1 to them. Level 3 experts (one person from Security-Public and two persons from the provider) are responsible for the overall incident handling process, defining frameworks for interventions, doing forensic analysis, and providing reports and recommendations.

The organizational structure is composed of a general director, who is appointed by the government. There are two staff mangers, i.e. administration and general managers, who support the director in managing internal activities. There are five internal divisions: (1) Operations, responsible for delivering proactive services and managing the internal operations of Security-Public; (2) Incident Handling, responsible for handling incidents and doing security intelligence for detecting incidents proactively; (3) Legal Services, which is responsible for delivering legal services to internal employees and external clients; (4) Promotion, which is responsible for providing and delivering contents related to information security for various communities; (5) Communication, that handles the interactions between Security-Public and other organizations.

The incident handling process starts when a new incident is reported by a client or detected by experts inside Security-Public or inside the provider. There are cases that the external bodies such as general security agencies or major IT companies (e.g., Google and Facebook) send alerts to Security-Public related to incidents. Incidents are sent to specialized queues when they are sent through specific email addresses. The incidents might also arrive through the call center that is operating 24/7 by first-level incident handlers. Upon the arrival of an incident, a ticket is opened in the ticketing system. The ticketing system is the main tool in the

incident handling process. First level experts log information about the incident, assess it, and if needed, they escalate it to level two to be handled. All incidents go through an initial phase called Triage, in which three managers from inside Security-Public examine the incident, assess its criticality, and define the process of handling it. In cases that the incident is sensitive, they might decide that the incident be handled internally, without sending it to the provider. The legal and criminal consequences and considerations are also discussed in Triage.

Major incidents at Security-Public

Security-Public is responsible for handling and preventing a wide range of security incidents in the local government. These incidents, for example, include hacking incidents, privacy attacks, data breaches, and identity theft. Being part of the local government, Security-Public considers these incidents as its own incidents.

3.6.2 Security-Private

In 1995, a project of establishing a new Computer Emergency Response Team (CERT) was the turning point for creating Security-Private as a start-up within the local university that was running the project. One of the senior managers of this project started exploring opportunities to establish Security-Private. In 1999, when there was still no formal company, two key experts, who later on became the managers of Security-Private, were running several IT security projects in banking sector and some others. In 2007, Security-Private was formally established as a private company specialized in delivering IT security services. This allowed the managers to attract financial resources, define their own catalogue of services, and equip themselves with required resources.

Since then, one of the two managers has served as the CEO of the company, and the other handles the projects (as projects manager) and manages internal

processes. The company experienced a major growth in 2007 and 2008. The company has been able to maintain its relations with most of its original clients, as well as establishing relations with new clients. The company has pursued quality strategy, by providing a rather specialized domain of security services at a premium quality. Security-Private has been working with clients from financial sector, public administration, education and universities, and manufacturing. In 2011, 30% of its revenue came from international projects and the plan is to reach 50% of revenue from international projects.

Internal structures, processes, routines, and systems at Security-Private

The company has a matrix structure, mostly shaped around functional departments; each specialized in delivering a specific service based on specific domains of expertise, such as incident response team, creating CERT, and vulnerability analysis. Security-Private provides a wide range of IT security services such as 1) handling their clients' incidents based on Carnegie Mellon Standard, 2) vulnerability analysis, and 3) designing and creating CERT teams.

Security-Private has 25 full-time employees. In the domain of incident handling, the employees are organized into three levels of seniority. In other domains of activities, the experts are not so formally classified. Instead, they are categorized into junior and senior. The employees have a fixed working schedule (40 hours per week), with almost fixed salary. A small part of their salary depends on their engagement in activities such as documenting their experience, and their performance, which is annually evaluated quantitatively and qualitatively. Extra-work hours are not paid formally, but the friendly and informal context of Security-Private allows compensating it through vacation leaves and recognitions (such as increasing the salary next year or promoting to a higher level). Most of the employees are young (average around 25 years old), specialized in IT security.

Incident handling process is based on the stages of the Carnegie Mellon Standard, consisting of 1) Identification, 2) Registration, 3) Modification, 4) Containing, 5) Resolving the incident (recuperation), 6) Analyzing the incident (technical), 7) Documenting and reporting, 8) Escalating (optional, just if other incidents would be created as the result) and 9) Follow-up. Normally Security-Private performs all these stages on behalf of its clients. However, in case of sensitive incidents, the client might take over some of the stages, especially the initial stages (e.g., Triage) and latter stages that require interactions with the final user and handling potential legal and criminal impacts.

As often required by its clients, Security-Private is working with the ticketing systems of its clients, as well as other complementary systems (e.g., Wiki, intranet, and databases) that each provider imposes. In parallel, Security-Private is operating its own system, which basically handles its internal events and incidents, as well as incidents of clients that do not force their own systems. In addition, Security-Private is using an internal Wiki system in which all information and documents and experiences of Security-Private projects and its members is store. This repository includes all sorts of issues such as formal documents, delivery items, internal procedures, and knowledge articles.

Major incidents at Security-Public

Security-Private is dealing with a wide range of security incidents for their own clients and their own systems. Security-Private has contracts with its clients through which the internal team of Security-Private are located inside client's site permanently and act as incident handling team. These teams perceive incidents of their clients as their own incidents.

3.6.3 Supercomputer-Large

Supercomputer-Large is a public international research center dedicated to supporting basic science researches that require analyzing huge amounts of data. Supercomputer-Large is running eight High Performance Computers (HPCs). Supercomputer-Large was born through a project in 2003 for launching one of the most powerful supercomputers in Europe at that time. In 2004, Supercomputer-Large was formally established with partners from the local university, local government and national government. Still this supercomputer, which was once upgraded in 2006, is running a main supercomputer at Supercomputer-Large. Over time, Supercomputer-Large has been able to add several other supercomputers.

Supercomputer-Large is more than a supercomputer center. It is a large, public, international research center, dedicated to supporting basic science researches that require analyzing huge amounts of data using supercomputers. Supercomputer-Large is part of a public University.

Supercomputer-Large has two categories of user (researchers) groups: internal and external groups. Internal users, are structured around “computer science” (178 employees), “earth science” (43 employees), and “life science” (86 employees). External users are researchers who do not belong to the University. Around 30% of the computing capacity of the main supercomputer is dedicated to the internal groups and the rest is reserved for external applications. Each project (application for a specific amount of processing hours) should submit its application to a scientific committee. The committee evaluates the application based on its scientific relevance and importance. This assessment takes place every 4 months. The committee ranks the applications and evaluates the amount of processing hours that should be granted to each project. At the beginning of each period, the users receive their account information and instruction.

Supercomputer-Large is responsible for assuring that each project can utilize the granted total amount of processing hours in the period of 4 months.

Internal structure, processes, routines, and systems at Supercomputer-Large

In parallel with research groups, Supercomputer-Large has a big operations department, which is responsible for operating the supercomputers, maintaining them, assuring the service continuity, handling the interactions of users with the systems, and resolving their problems. Operations department is also responsible for handling any incident that might happen to the supercomputers and its related infrastructure and application. Operations department is structured into three main groups: (1) Facility-Management, which is responsible for electricity, cooling systems, and other basic office infrastructures (2 persons); (2) System-Administration , which is responsible for all processing and storage facilities, network, basic applications for running the systems, website and other web servers, as well as all basic services such as email (11 persons); (3) Users-Support; which is responsible for interactions with the users in terms of managing their access to the systems, helping them running their jobs on the systems, getting their data properly stored, and resolving any issues that users might face in their work with the systems (6 persons).

The focus of this study has been on the operations department as this is the main body related to IS incidents. In fact, my inquiry showed that final users and top managers often do not enter into the process of handling IS incidents. Experts who work in the operations department are mostly from computer engineering domain, and some others are specialized in electrical engineering. They have a fixed salary with fixed amount of working hours (40 hours per week). The heads of System-Administration and Facility-Management groups should be available (through mobile phone) on a 24/7 basis to assure timely response to critical incidents. In spite of specialization of tasks and responsibilities, all the experts

from operations department are close friends and there are a lot of informal interactions and with a friendly climate in the department.

The operations department deploys an overall ticketing system for handling its daily works. Each group and sub-groups have their own separated queues in the ticketing system. Requests from users or other internal groups arrive to each specialized team. To facilitate the process for users, there is a general email for each of the three groups to which the users can send their requests. Once this arrives, a first level expert takes the ticket. If the ticket is easy enough to be solved at first level, it will be done. Otherwise, it will be escalated to level 2 experts who are specialized in different themes. The ticketing system for each group has a very specialized categorization of queues, which in most of the cases is not visible for the final user. Often one or a few experts are handling the incidents related to each queue. There are tickets that move across various queues as they deal with several specialized aspects of the system. In addition, Supercomputer-Large has recently forced users to provide a report every two weeks. In this report, the users can present their problems, concerns and requests. This way, Supercomputer-Large can detect the problems of the users on a regular basis and avoid complains from users at the end of the period of access.

There are, however, differences between groups in terms of handling the incidents and daily activities. In Users-Support group that receives numerous requests, experts take turn for handling incidents. This way, each of the three experts at level 2 is responsible for handling all the tickets of the group in that specific week, so that he has 2 weeks with no responsibility about tickets. This allows experts to concentrate on their internal projects. System-Administration group is highly specialized (6 sub-groups). Although specialized providers provide most of the basic facilities (hardware), the group has been taking the responsibility of handling the applications (often, open source). In some cases, because the technology is

very advance and unique, the group has not been able to find appropriate provider to solve their problem. Hence, they often start developing new tools by themselves. In Facility-Management group, there are many facilities operating, which are all outsourced. The group has various contracts with service providers for the maintenance of systems. In case of incidents, these service providers often are those who come and solve the issue. Even fundamental changes are often formulated in terms of changing the hardware, scheduled repairs, and periodic or specific upgrades.

In addition to handling requests and incidents detected by the operations department, the groups spent almost more than half of their time on their internal projects. These projects are long-term changes for upgrading systems, expanding some capacities and basic repairing. Most of these internal projects are around developing some tools and applications (often open-source software).

Major incidents at Supercomputer-Large

Supercomputer-Large has experienced a wide range of major failures to its critical systems such as damages to its storage and processing systems. As a result of these incidents the jobs that are running on supercomputers might be lost and some damages to the systems (e.g., hardware damage) might happen. However, Supercomputer-Large has not experienced any serious security incident.

3.6.4 Supercomputer-Small

Supercomputer-Small was founded in 2003 and maintained through a collaboration agreement among the local university and the local government. Supercomputer-Small is a data center dedicated to scientific-data processing by supporting scientific groups working in projects, which require strong computing resources for the analysis of massive, distributed data. Since the beginning,

Supercomputer-Small has been collaborating closely with the LHC Computing Grid (LCG) as a Tier-1 centre for WLCG project at CERN. Around 80% of this capacity (both storage and processing) is dedicated to WLCG project.

Supercomputer-Small has been pursuing to acquire other research projects from other domains such as astronomy, life science, and medicine, though most of them have been minor pilot projects. Each project often has its own funding. This funding sometimes is in the form of bringing new machines (processors or storage) to Supercomputer-Small.

Supercomputer-Small interacts with the host university (as the owner of the infrastructure), with tier-0 of WLCG project at CERN, with other tier-1 sites of WLCG, with its own tier-2 sites, and its providers and vendors of systems and technologies (that are directly connected with Supercomputer-Small).

Internal structure, processes, routines, and systems at Supercomputer-Small

Supercomputer-Small is organized into five Areas: Administration, Infrastructures, Services, LHC, and Projects Area. Around 30 experts work at Supercomputer-Small. Their backgrounds are mostly in computer, IT, and physics.

There are two computer rooms at Supercomputer-Small. The traditional one is bigger, with older machines, and its facilities (electricity, network, and cooling) are owned by the host university. As for the second one, Supercomputer-Small owns the infrastructure and operates them through directly outsourcing to providers. This gives Supercomputer-Small more control and autonomy for handling the infrastructure. One of the challenges of Supercomputer-Small with the host university is how to convince the university to consider Supercomputer-Small's facility as a critical system, not something like other departments.

Energy is one of the most critical concerns of Supercomputer-Small. The cost of energy (for cooling systems and for operating machines) is the main concern for

the host university. In addition, reliable supply of electricity has been constantly a challenge and has been challenged with several incidents. This concern justifies the upgrade of the machines (both processors and storages) quite frequently (almost every 2 years), although most of them are still working properly in the time of decommissioning. However, operating them is not cost effecting with regard to the high energy costs. This concern, however, has facilitated justifying the need for frequent upgrading, making Supercomputer-Small working on the latest technologies, with a constant trend of expanding its capacity.

Supercomputer-Small is using numerous technologies from a wide range of providers. As an overall strategy, Supercomputer-Small has avoided being locked into a specific vendor and technology. Supercomputer-Small has an informal and friendly climate. Most of the activities are organized through daily and informal interactions. Relations with the host university and other suppliers are often at technical levels, without the involvement of top managers. Each week, one person from the experts and middle managers at Supercomputer-Small takes the responsibility of “manager on duty”. This person is responsible for being available on a 24/7 basis, and is the one who is responsible for handling incidents and interactions with other groups and providers. Supercomputer-Small uses a ticketing system to organize its tasks, requests from the users, and incidents. Even big incidents are handled through the ticketing system. In addition, Supercomputer-Small has its internal Wiki system.

Major incidents at Supercomputer-Small

Supercomputer-Small has experienced very similar incidents as Supercomputer-Large did. In some cases, major damages to hardware have resulted in partial damages in stored data, as well as losses in running jobs. The former effect is often considered more serious than the latter.

Chapter summary: The chapter articulated the specific research question that is "how does the process of organizational learning from ISRIs unfold, with regards to the characteristics of ISRIs"? It described that a multiple case study design is selected for looking at the learning process, including what organizational actors do to capitalize on their incident experience. The chapter then described the data collection process in each case, which took place in three steps: understanding the background of the organization, incident based inquiry, and change based inquiry. The chapter described the process of analyzing the data, which went through four steps: analyzing during data collection, developing analytical reports for each case, open coding, axial coding for learning practices, and embedding findings into the extant literature and articulating learning modes. Ethical considerations and main challenges in the research process were then reported. Finally, the four case settings (two security companies and two supercomputer centers) were described. The next two chapters describe what the studied organizations did to leverage their incident experience.

Chapter 4: Findings: Single-incident Learning

In this chapter, I focus on single incidents to examine what the organizations did in relation with each of their major incidents to benefit from the experience of that incident for future (the next chapter concentrates on actions related to multiple incidents). Examining how organizations benefit from their experience of a given incident surfaced a series of practices that organizations adopted *during* and *after* handling major incidents (Figure 4.1). The chapter first describes the practices that the organizations took during incident handling process. These practices are classified around three overall themes that emerged inductively from the empirical study: “handling incidents through ticketing systems”, “performing Triage”, and “interacting with specialized providers”. In the second part, I describe

the practices that the organizations adopted after handling each incident. These practices are articulated under “post-incident reflection” and “laboratory incident analysis” practices.

The chapter describes these practices by explaining the way these practices were taken, in which situations, by which actors, and with what temporal patterns. An illustrative example of each practice is presented in boxes. In addition, different patterns of the practices across organizations are described.

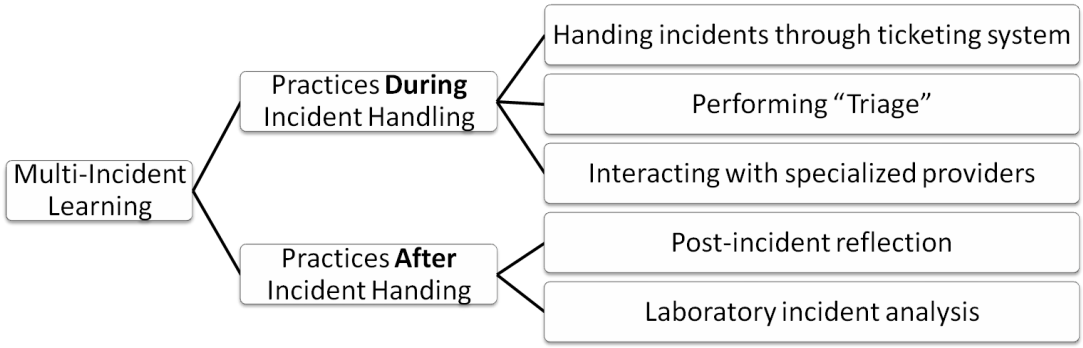

Figure 4.1: single-incident learning practices

4.1 Practices during incident handling process

The studied organizations were actively engaged in three sorts of practices during incident handling process, namely from the time that the organizations faced an incident to the time that they could control the impacts of the incident and return to the normal situation. Table 4.1 summarizes these practices and how each practice is differently adopted by the studied organizations.

Table 4.1: Learning practices during incident handling process

Practice	Description	Specificities of the organizations			
		Security-Public	**Security-Private**	**Supercomputer-Large**	**Supercomputer-Small**
Handling incidents through ticketing systems	Handling incidents itself provides intimate experience about various types of incidents, how they work and how they can be prevented or managed in future. In addition, using ticketing system facilitates documenting the events, experiences, ideas, and reflections during incident handling process.	- Incident handling process is outsourced to a specialized provider who operates Security-Public's ticketing system. - Politically sensitive incidents and their tickets are internally managed. - Triage is only done by internal experts. - Ticketing system is supposed to keep the record of all incident handling experience, although it is done by the provider	- Operating several ticketing systems for various clients creates the challenge of not being able to systematically move tickets and documents across different ticketing systems. - The limited access to ticketing system due to confidentiality and sensitivity for the provider	- High level of specialization in queues and incident handling process (at System-Administration group) - Rotations for handling incidents among senior experts (at Users-Support group)	- Ticketing system as a repository that everything about incidents should be logged there - Ticketing system is also used for tracking follow-up actions
Interactions with specialized providers	The organizations rely on the experiences of their specialized providers during incident handling process. This way, they try to establish relations with their providers to allow the providers leverage their experiences and keep their relations with the providers to help them avoid similar incidents and manage them better in future cases.	- Asking the provider to document the experience of incident handling - Changing the contract with the provider from body-shopping to volumetrias	- Not Applicable	- In some cases trying to collaborate with the provider during incident handling process, while in some other cases trying not to be involved and delegate all the process to the providers	- Mostly relying on the specialized providers - A highly diversified collection of providers
Performing Triage	Before starting to handle incident, senior experts examine the scope of incident, previous related incidents and experiences, and set an overall framework for handling it.	- Done internally (unlike other stages of incident handling process that are outsourced) - Doing it formally for all incidents	- Doing it rather informally	- Doing it informally and occasionally only for novel incidents	- Doing it informally and occasionally only for novel incidents

4.1.1 Handling incidents through ticketing system

The organizations handle all incidents using ticketing systems. Large incidents are handled, sometimes, through several tickets. Ticket is an artifact in the ticketing system that is created when a specific incident, request, or task arrives (See Figure 4.2 bellow). For doing that, the user of the ticketing system (often first-level experts), opens a new form in which the user first introduces the name of the ticket, which is a rather short name (e.g., around 5 words) to show what the incident is about. Then, the user enters a description of what is the incident about (around one paragraph). The expert writes this description based on the information gained from end-users, from observations of the incident, and from further investigations done at the time of receiving incident (e.g., searching in weblogs to see if any similar virus is reported recently). The time that the user enters these two items and creates the ticket, the system automatically records the "issue time" of the ticket (date, hour, minute, and second). Upon the creation of the ticket, it is considered as an "open ticket", meaning that it still should be handled. From this point on, the ticket becomes the core element in the system upon which all incident-handling activities are focused. Experts, who work on the ticket, often do not use "incident" or "failure" terms in their interactions. Instead, they use "ticket" to refer to the incident.

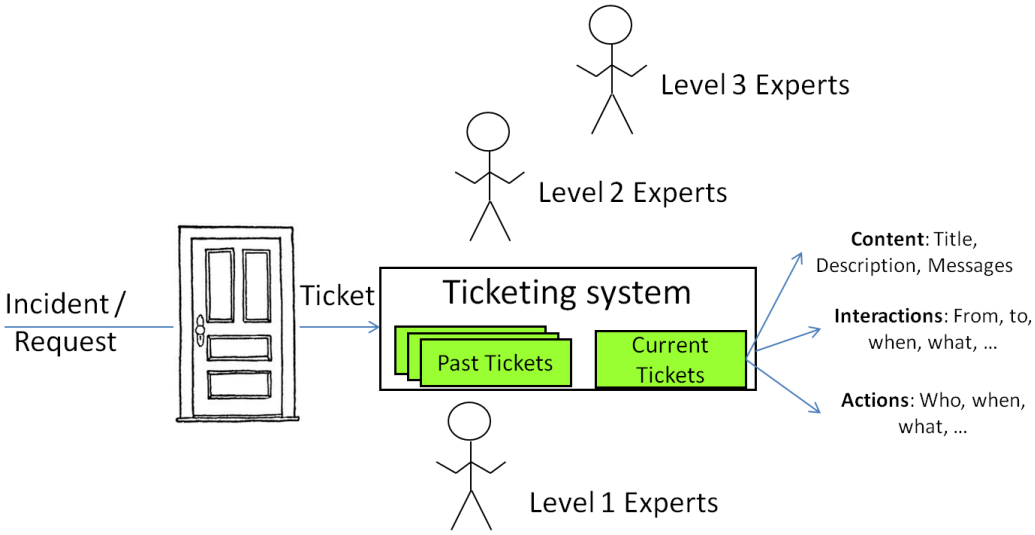

Figure 4.2: The process of handling incidents through ticketing system

In the ticketing systems, there are categories in which open tickets are stored. These categories are called “queues”. The queues are often specialized based on departments, groups, and sub-groups. For instance, at Security-Public, each of the three departments –i.e., operations, incident handling, and legal- has its own queues. In many cases, the queues are further specialized in terms of technical systems, incident types, and importance levels. For example, in the area of incident handling at Security-Public, there are several queues, each related to different levels of importance (e.g., level 1, level 2, level 3, and level 4). At Supercomputer-Large, for instance, in the System-Administration group, each specialized sub-group has its own queue. Sometimes the distinction between queues relates to different ways in which incidents arrive. For instance, some tickets are automatically generated at Security-Public from on-line services that some companies such as Google provide. Finally, the specialization of queues sometimes is done to facilitate the access control of various users. For instance, at Security-Public, a queue is only for Triage to which the provider of Security-Public cannot access.

At any time, each ticket has only one owner who is responsible for taking care of the ticket. When the ticket arrives, it should be taken (owned) by a user (the expert who handles the incident). During incident handling process, a ticket might move from one owner to another owner.

Each user of the system has its own account in the ticketing system. This means that when the user enters the system, the system records all his activities done on the tickets. From the time that a ticket is opened, all the transactions that are done on it are automatically stored in the log of the ticket. This includes who, when, did what action on the ticket. These actions, for instance, can be picking the ticket for handling it, observing it, moving it to another user, and editing the content of the ticket (for example revising the description of the ticket).

Almost all actions that experts perform for handling incident either are done through or at least are reflected in the related tickets. For instance, all the tools used for handling incidents, all trials and errors, all tests, all failed actions, and all personal reflections are logged in the ticketing system. Part of this is done automatically. For instance, when an expert writes back to the user to ask him to run a specific file to solve his problem, this message, the file that is attached, and the details of the interaction (e.g., when and who sent the message) are automatically stored in the ticket.

Other users, even if they are not the owners of the ticket, can see all the tickets in the queues to which they have access. Tickets can be seen and be commented on by other actors who are not its owners. For instance, in the Security-Private, when new tickets arrive, the head of the incident-handling department who is not often involved in the details of handling incidents receives emails related to those tickets. He sometimes opens some of the tickets just "*to see what is going on*" [the head of incident handling department].

Experts who handle a ticket write short notes describing their actions on the ticket. For instance, in a ticket, the user reports an incident related to batch files (the files that should be run on the supercomputers). The incident arrives to Users-Support group. They guess that the problem can be solved if the user uses a different process of uploading the job. They send the message back. The user tries the new method, but it does not work again. Then the user sends back the ticket and the experts in Users-Support group realize that the problem is more complex and System-Administration group should handle it. Therefore, they move the ticket to the queue of the System-Administration.

The ticket at each time has one of these states, depending on its position in the incident handling process: (1) “not contained”, from the time that it is created until the time that one expert owns it; (2) “Open”, during the period that the ticket is under handling by experts; (3) “Pending”, when the real activities for handling incident are finished, but still it is not formally approved as finished by the person who has the authority to close tickets; (4) “closed”, once the ticket is approved by the person who has the authority to formally close the case.

Apart from these normal states, some innovative states are also defined for specific cases. For instance, at Security-Public’s ticketing system, there is a state for tickets that is added by its provider, called “pending closed”. This corresponds to tickets that from the view of the provider, everything that should be done for its handling is finished, and they are almost sure that Security-Public’s delegate will approve it (based on their experience), but he has not had time to do so. This category is important because at the end of each month, the provider can report the handled incidents (by summing the “pending closed” and “closed” tickets) for billing purposes.

The system allows that similar tickets be grouped. Large tickets can be split into smaller ones. Tickets can have sub-tickets (“child tickets”). In Security-Public and

Security-Private, this feature is used for defining sub-tickets that different actors should handle independently. For example, once a sever security attack arrives, it might be considered as a single ticket. Then, during the process, it might appear that the technical department should control the attack and recover the lost data and the legal department should start filing legal cases and start evaluating the legal consequences. Then, the ticket would be split into two tickets for each department. At the end of the process of handling incident, the two tickets will be integrated.

Tickets, although sometimes might go through several departments and several actors to work on them, even though they are often individual tasks, related to a single person. Since large, challenging tickets are specialized, a senior expert often handles them individually. This means that most of tickets just have one single owner throughout their life. For example, in Supercomputer-Large, more than 95% of the tickets are single owners. In fact, tickets that are handled by several users are either exceptional or are those that were mistakenly arrived to the wrong person, and then moved to the next person. In fact, tickets are attached to individuals ("*was this your ticket?*", as a common way of talking about tickets). The experts either handle the ticket alone, or ask the related providers to handle the incident. For example, at Supercomputer-Large, the experts at System-Administration group each have a narrow domain of specialization. The person responsible for network, for instance, handles less than 10 incidents per month. These serious incidents are quite specialized. Since the specialized systems are often outsourced to specialized providers, these internal experts work with the suppliers to handle the incident.

The system allows searching in the tickets' titles and contents. It also provides various reporting facilities. For instance, the users can find all the incidents in the last month related to batch files. The search and reporting facilities of ticketing

system is used for finding a specific ticket that can help in handling a current ticket. It often happens when the expert who is handling the ticket remembers that he had handled a similar ticket in past. They use this feature when they do not remember by memory, which often happens to incidents that repeat occasionally. Previous similar tickets are useful because they show the full history of all actions taken in previous cases, including both things that worked and things that did not work. This helps experts not repeat the same unsuccessful actions.

Each expert can search in the tickets that are handled by other experts to see if there has been similar incidents in past. However, this is done quite rarely, mostly because first, they are not sure whether other experts have handled similar cases or not. Second, they are not sure which specific keywords other experts have used when they were writing the content of similar tickets. Third, they are not sure if what others have written can help them. It is because when an expert puts comments on the ticket, it is mostly for his own use, to remember what he did in case of future similar incidents. Fourth, other ways, such as just asking loudly in the room from colleagues would be much easier and more preferable to see if others know something about this incident. Finally, because the search capabilities of the system is not as strong as the web-search engines, it is much easier and more effective that instead of searching in the ticketing system, they just “Google it”.

There are differences between organizations and their departments in terms of using the ticketing system. For instance, in Security-Public, the system is shared with the provider. The provider has access to almost all queues, except some of them that are for very politically sensitive incidents. In fact, the ticketing system at Security-Public is forced to the provider, as the provider has to do everything related to incident handling on Security-Public’s ticketing system. This is because Security-Public wanted from the very beginning, to have all knowledge and

experience about the incident handling be stored inside it, to reduce the risk of losing that knowledge in case of switching to another provider. In terms of daily work, the provider's experts do most of the actions in the ticketing system.

Security-Private uses several ticketing systems from various clients, as well as an internal ticketing system for its own. These systems are disconnected. Therefore, if a ticket is generated in one system, it cannot be moved or replicated in another system. This separation is mostly because of the confidentiality of the incidents that each client has.

In Supercomputer-Large, the ticketing system is specialized in the different departments. In Users-Support group, there is only one queue for all requests from users. It is to make it easier for users to interact with a single contact point. In System-Administration group, there are very specialized queues, each related to a specific aspect of the technology or a specific technical system. Box 4.1 illustrates the process of handling a recent incident that interrupted the email system at Supercomputer-Large.

Box 4.1: Handling Email-servers incident at Supercomputer-Large

In a recent, rare incident, two research institutes that were receiving the services from Supercomputer-Large could not receive any email from Supercomputer-Large. This incident caused serious problems because they could not get their jobs run and examine the progress of their jobs on the supercomputer. The incident was detected when a user in one of the institutes called the Users-Support group. At Users-Support groups, the Email-servers expert, who was responsible for incident handling during that week, opened a ticket related to the incident. He documented the information about the incident, including when and how it happened, who were the users that were affected, and what were the specific problems created.

Realizing that the problem might be related to the email service, the expert at Users-Support group moved the ticket to the expert at System-Administration group that is responsible for incidents related to Email-servers.

Upon receiving the ticket, this expert read the ticket's description and he doubted if there was any problem with the email-server applications. He ran some tests on the system and checked if the messages could be delivered properly to other clients. Then, it appeared that the problem was not from the Email-servers. He thought that it might be a problem at the network level. Therefore, he passed the incident to the expert who is specialist in network incidents.

The expert at System-Administration group who deals with network incidents received the ticket and read the description that his two colleagues had written in the ticket. He started testing the network connections with the two clients, using several specialized network probes. He then ran tests on the suspicious network nodes. He realized that the problem must be outside the Supercomputer-Large network. In fact, Supercomputer-Large is connected to these clients through a dedicated network that is operated by a special network provider. He then opened a ticket in the ticketing system of the network provider to pass the incident to them. Doing so, he used the same description in the internal ticket and added his own part. He then descried that the tests he had done showed that the problem was outside Supercomputer-Large network. In parallel, he called the provider on phone, explaining the incident.

After a couple of hours, the experts in the network provider resolved the problem. It turned out that it was a problem with the new configuration that they recently applied for migrating to IP6. The provider did not explain so much details of the incident, as it was not related to Supercomputer-Large's network.

Figure 4.3 illustrates the flow of the ticket during incident handling. Table 4.2 summarizes the interactions inside Supercomputer-Large and between Supercomputer-Large and its network provider during incident handling.

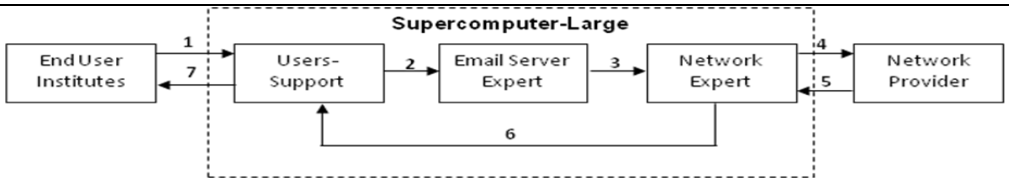

Figure 4.3: The flow of actions and interaction in handling IP6 incident

Table 4.2: The steps of handling incident and related ticket activities

Step	Duration	Interactions through ticket
1	Several minutes	- Opening the ticket by User's supports - Documenting the detail of the incident into the ticket - Recognizing that the incident cannot be handled by Users-Support group and should be moved to System-Administration group
2	Around half an hour	- Users-support moves the ticket to the expert in the email servers - They also talk face to face with the expert to inform him about sending the incident - The email-server expert runs several tests and realizes that the problem might be from the network level - These actions and tests and the results are documented into the ticket by the mail server expert
3	Around one hour	- Email server expert sends the ticket to his colleague who is responsible for network - He also talks with the network expert, face to face, to let him know about this incident and They discuss the issue and he explains the story - The network expert runs several tests to find the source of the problem - It appears that the problem is in the external network - The network expert documents these points in the ticket
4	Several hours	- The network expert opens a ticket in the ticketing system of the network provider (by sending an email to their ticketing system) - The network expert calls the provider and explains the story and informs them about the problem - The provider works on the problem and detects the problem and solves it
5	Several minutes	- The provider calls back to the network expert and explains the problem, its causes, and says that it is solved - The ticket between the provider and the network expert is closed
6	Several minutes	- The network expert describes the story of the problem into the ticket briefly ("the problem was because of the configuration of the IP6 in routers of the network provider") - The network expert sends the ticket back to Users-Support and let them know that the issue is solved
7	Several minutes	- The Users-Support informs the end user institutes that the issue is solved - The end user institutes confirms that their problem is solved - The ticket is closed by Users-Support group with the note that "the problem is solved"

Finally, in Supercomputer-Small, the ticketing system is considered as “the only must”, stated the head of technical department. The organization wants to have a repository of everything that happens in incidents. Therefore, from the very initial minutes of the incident, everything is managed through the ticketing system. Supercomputer-Small uses the ticketing system also for tracking the progress of its follow-up actions. In fact, there are two types of tickets: those that are related to incidents and tickets that represent a specific internal project. Although the organizations use ticketing system differently, it is a key working tool that helps them use their past experience during handling incidents and store their experience for future incidents.

4.1.2 Performing Triage

A second practice that Security-Public and Security-Private used during incident handling process is Triage. These two organizations use Carnegie Mellon methodology, which is a detailed, staged incident handling process. In this process, after receiving and containing the incident, there is a stage called “Triage” in which incident handlers hold a meeting to discuss the incident, evaluate its scope, its criticality, and its potential impacts, set the overall framework for handling it, define roles and responsibilities for managing it, and make sure that legal and criminal considerations are taken into account. All incidents must go through Triage. The same ticketing infrastructure that the organizations use for incident handling is helping the Triage stage. Among all other queues in the ticketing system, there is a specific queue, for doing Triage where the incidents that are waiting to be Triaged are stored. The access to this queue is critical and only a few senior experts have access to it. Box 4.2 illustrates a typical Triage at Security-Public.

Box 4.2: A typical Triage at Security-Public

When an incident enters any of the queues in the ticketing system, the first-level experts in provider make sure that its information is complete and clear. This includes, basic information about what is the incident about, when and how it happened, and all the initial information gained in the time of creating the ticket. Then, these experts move the ticket to a queue called “Triage”. In cases that the provider recognizes that the incident is of high priority, uses colored flags to indicate its priority (red for the most critical incidents).

There are four second level experts inside Security-Public who are members of Triage team: (1) The head of incident response team (ERI), who is a senior IT security expert; (2) The internal executive administrator who is managing all internal activities and coordinates interactions among different departments (as the “right hand” of the CEO); (3) The legal expert who is a lawyer with 10 years of experience in IT law; (4) The representative from local police who is expert in IT crime.

The head of ERI is responsible for Triage. He reads the incidents in the queue, before the meeting. For minor cases, he might decide that no formal Triage is needed. Therefore, he passes those tickets to their associated queues (depending on the type of incident). From that point on, the provider that has access to these queues starts handling the incidents. There are other incidents that the ERI manager put some short notes inside their tickets. This helps the provider in handling them. These notes can be suggestions about how the incident should be handled, what specific cares should be taken, what potential incidents might result during or after this incident, and sometimes some references to similar incidents that were handled before.

For novel and serious incidents, the ERI manager reads the details of information that is written in the ticket. He sometimes searches in weblogs, forums, and daily

notices that are related to similar incidents. In some cases, they go back and look at some specific incidents that they handled before. It is mostly in cases that they remember some specific past incident and they want to remember a specific point, for example the hacking methods used or the specific solution that did work last time. However, no systematic analysis of past tickets is often done.

The four experts meet normally every two days physically where they examine the tickets in the Triage queue. In urgent situations, they hold extra meetings when an urgent, critical incident arrives. It is common that the meeting takes place on a telephone conference session, or by using the dedicated network to which these four people have access.

The ERI manager prepares a short list of novel, critical incidents for the meeting. He manages the meeting by introducing incidents one by one. He starts from incidents that are more serious. This can be related to the potential damage of the incident, as well as the urgency of handling it. He opens the ticketing system in his laptop and goes straightforward to the Triage queue that is already cleaned before the meeting by removing the unimportant tickets. He starts from the most important ticket that from his point of view needs the opinion and decision of other Triage team members. Other team members also have access to the tickets in the Triage queue and look at it during the discussion.

ERI manager opens the ticket, which might have some of his notes that he had prepared during his search before the meeting. He describes briefly the ticket to the other members using his own words. He outlines the technical aspects of the incident, when and how it happened. Knowing the whole incident management process then the ERI manager asks specific questions that require discussions with other colleagues responsible for organizational, legal, and criminal aspects.

They discuss the incident among themselves. They discuss the potential threats that the incident can bring about, possible related incidents that might not be so

visible in the first place, the incidents that might occur as the result of this incident or during handling it, the legal considerations that should be taken into considerations when the experts are handling the incident, and potential criminal implications of the incident. Finally, they define the steps through which the incident should be handled and identify which expert should take which part of the job. More specifically, technical tasks (e.g., controlling the virus) is assigned to ERI manager, and then, he break them down into sub-tickets to be handled by the provider's experts. Legal tasks are delegated to legal expert. The police expert handles the tasks related to criminal activities (e.g., filing a record in police office).

The ERI manager then puts notes in the ticket that reminds him and other team members about what each of the four persons are supposed to do. If needed, new specific sub-tickets are defined under the overall ticket that they need to be handled by each of the four members. Sometimes they conclude that there is some significant lacking information about the incident that needs to be gathered. This can imply that they need to redo Triage once they got the picture clearer.

Figure 4.4 depicts various aspects of Triage in the incident handling process.

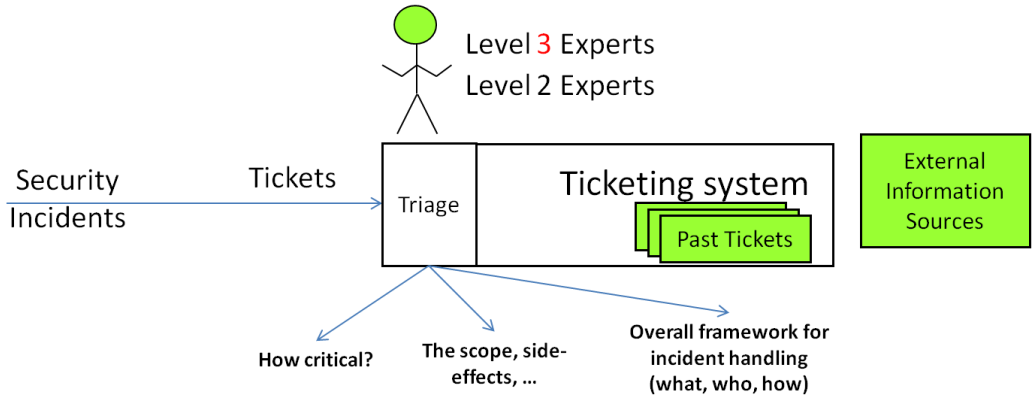

Figure 4.4: Triage in the incident handling process

In Security-Public, the composition of actors involved in Triage has changed over time. At the beginning, it was done by the technical expert who was the head of

incident handling group and another middle manager who was the head of operations. Later, the legal expert was added to the team, as the legal consequences of incidents and incident handling became more and more observable. Recently, a representative from local police is also added to Triage team, due to facing incidents that have criminal implications.

Although all incidents should go through Triage, small or familiar incidents do not take a long time to be analyzed in Triage. It might be a matter of minutes to decide that these incidents should be assigned to the provider to handle them. However, for complex, novel incidents, Triage might take a long time (sometimes several long meetings). These incidents might require Triage to be repeated several times. This happened in cases that new issues, for example related to legal consequences of identity theft, appeared during incident handling process. Sometimes one incident might be part of a chain of interrelated incidents (e.g., a series of security attacks). Then, it is critical that the possible relations between the detected incident and other incidents be examined. This task, that is mostly done through Triage is called *correlating,* which means various pieces of information related to the current incident and other incidents are put together to see the whole picture of the incident. Correlating, as described by a senior expert in Security-Private, does not necessarily confine to Triage meetings. Rather, it is done throughout the whole incident handling process. He said,

"you constantly do correlating when you are handling a specific incident. It is a constant thinking about what might be related to this incident, searching and finding clues, and connecting them to each other" [a senior expert at Security-Private].

As mentioned before, one of the aims of Triage is to make sure that incident handling does not yield in further damages. However, the fact that all incidents should go through Triage makes this stage sometimes a bottleneck in the incident

handling process. This is particularly the case in Security-Public where most of incidents are delegated to the provider to be handled, after Triage. The provider first receives the incidents, files them, and sends them to Security-Public for Triage. In many cases, the incidents are quite simple and routine. However, officially, the provider should wait until it receives the incidents back from Security-Public. Over time, the provider has learned which incidents are not sensitive. Thus, the provider benefits from the trust that Security-Public has on it, does an internal assessment (they call it "auto-Triage"), and starts the process of handling them if they are not sensitive.

At Supercomputer-Large and Supercomputer-Small, Triage is mostly done informally and occasionally for very novel incidents. Even in such cases, it is not a specific stage in the incident handling process. It is more embedded with other incident handling activities such as discussing the incident among the experts.

To conclude, Triage as a formal step in incident handling process provides an opportunity for the organizations to reflect on the incident, before rushing into its handling. This way, the can draw on their past experience early on in the incident handling process.

4.1.3 Interactions with specialized providers

The third category of practices during incident handling process pertains to the interactions with specialized suppliers. All the organizations (except Security-Private) were relying on their specialized providers in handling incidents, since they outsourced a wide range of specialized maintenance services to their technology vendors and specialized services providers. Handling incident is part of contracts with their providers. This includes both preventive revisions to detect the potential incidents and resolve them before they result in severe damage, and

corrective actions to control the incidents and fix the damages. Therefore, most of incident handling activities take place in the site of their providers. Box 4.3 illustrates this process in a recent incident at Supercomputer-Small.

Box 4.3: ZSF File-system damage at Supercomputer-Small

One of the critical responsibilities of Supercomputer-Small is to store processed data coming from all parts of WLCG project. The data is stored in huge file-systems (several peta-bytes) on storage systems. In a recent upgrade, after making several changes to the software and hardware, the whole system was switched on. However, when new data arrived, the storage system generated errors, because the new data was written over the existing data. This damaged around 250 tera-byte data on the disk. This huge amount of data was likely to be lost. As expressed by the head of LHC at Supercomputer-Small, "*the nightmare is data loss*" [the head of LHC at Supercomputer-Small].

The head of facility manager quickly switched off the system and opened the ticket in the system. Then, she contacted the service provider of the damaged storage system. Since the system was a rather old-fashion technology, it took around two days to find an expert of the system and contact him. The provider's expert got access through Internet to the system and started retrieving damaged data. Around 90% of data then was retrieved.

During this period, experts of Supercomputer-Small were only giving access to the provider's expert, providing him with some information about the history of the system, and finally checking the recovered data. As mentioned by the facility manager at Supercomputer-Small, "*we were here on-line and waiting*" [the Facility Manager at Supercomputer-Small]. She added, "*These technologies are very special. Even for big providers, it takes a while to assign a proper expert to us*" [the Facility Manager at Supercomputer-Small].

Further analysis showed that the interactions with specialized suppliers for incident handling might take place in different ways. In some situations, the organizations fully outsource incident-handling process to their providers, while in some other occasions they collaborate with the providers to the extent they make sure that incident handling process is taking place properly. All the organizations, except security-public, have outsourced their incident handling processes to various specialized suppliers. Security-Public has been mostly working with a single security supplier that handles the whole range of its incidents, mostly due to its overall outsourcing strategy. Finally, there are some circumstances in which incident handling process is done internally by the organizations. The next paragraphs describe these ways of interacting with suppliers and illustrate them.

Fully outsourced incident handling to numerous providers

In some situations, the internal experts are detached from the details of incident handling process. It is because the organizations are operating a wide range of technologies, all outsourced to specialized suppliers and providers. In addition, in cases such as Supercomputer-Large, the contract forces that all changes in the hardware should be done by the provider through scheduled upgrades. For instance, Supercomputer-Large works with more than 30 specialized service providers for different technologies (e.g., processing systems, storage systems, backup systems, simulation systems, network infrastructures, servers, operating systems, and data processing applications). Each department and specialized group is working with its own specialized providers. For instance, the backup expert in Supercomputer-Large is working with 5 major service providers that are only specialized in backup systems. As the head of Facility-Management group at Supercomputer-Large mentioned about the process of handling a recent hardware damage, the organizations did not involve in the process, since it was delegated to the specialized supplier:

"I don't know, they just come, fix the system and they go… for me, it is important that I have the guarantee and they have to solve the issue…" [The head of Facility-Management group at Supercomputer-Large].

The excerpt suggests that in such situations the expert becomes detached from the incident handling process and its details. While the practice ensures a swift resolution of the problem, how this is done becomes a black box. There are, of course, good reasons for adopting such procedure. For example, organizations operate a wide range of technologies, and these tend to be outsourced to specialized suppliers and providers. However, as I will see later, this arrangement may have unintended effects on learning process.

At Supercomputer-Large, there are two groups that are mostly dealing with their providers for handling incidents: Facility-Management and System-Administration. In the Facility-Management group, the model of working with the providers is quite outsourced. In case of incidents, the internal team (one junior expert and one senior expert) just make some initial actions to control the incident. This often involves switching off the systems that might be damaged. Since the systems are modularly designed, they can easily detect which part of the system has the problem. A wide range of alarms and sensors help them spot exactly the faulty part. If needed, they disconnect the faulty part and they call the related service provider. The service providers should arrive in 48 hours. As described by the head of the group,

"From that point on, I do not do anything. They just come and do their job, fix it and go" [the head of Facility-Management group at Supercomputer-Large].

This is partly because the process of solving the problem relies on replacing the damaged part with a new one. The group often does not get any specific feedback from the provider on why the incident happened and how it can be prevented. It is mostly because

"We cannot do anything to the system. It is designed. It is fixed. We cannot prevent these incidents. They happen sometimes. The hardware might break down. We just make sure that we have contract with the providers" [the head of Facility-Management group at Supercomputer-Large].

This is very similar to the situation at Supercomputer-Small. As the head of infrastructure group pointed out, they almost rely on whatever their, specialized providers do:

"It is easy for us. We just have spare parts. We replace the iron! Then, the providers come and take the damaged part and examine what was the problem" [the head of infrastructure group at Supercomputer-Small].

Outsourcing to single provider by Security-Public

Although Security-Public is mostly relying on one specialized provider, the overall strategy is based on outsourcing handling almost all incidents. This has allowed Security-Public to constantly improve its capability in terms of handling more incidents, in shorter period, and more efficiently. For the first 3 years, Security-Public had been contracting with the provider based on body-shopping model in which the payment to the provider was based on the number of hours that experts from different levels spend on handling their incidents. Since in that period, many incidents were quite new to the provider, the provider constantly improved delivering its services as gaining more experience.

In 2011, Security-Public changed the service model in its contract with the provider from body-shopping to "volumetrias". In the volumetrias model, the provider is paid based on the number of incident handling services in each category of services. Thanks to the experience of the first three years, Security-Public and the provider could define categories of incident handling services and define a price for each category based on its difficulty. Within an overall annual

budget, Security-Public asks its providers to be prepared to provide an expected amount of services in each category (for example 30 lost data recovery). These services are common in most of the incidents. One service can be enough for handling a small incident, such as a specific virus infection. However, for big incidents, such as big security attack, often a range of services should be integrated to handle a big incident. This way, Security-Public makes its provider improve the services to be done more efficiently, so that the provider can gain a higher margin.

Obviously, there are services that are difficult to be standardized like this. This is a key debate among managers and senior experts inside Security-Public about to what extent the new service model can be applied. They fear that applying this services model for a series of highly sensitive and non-standard services would make them compromise their quality. As one of the middle managers at Security-Public mentioned,

"The analysis of the vulnerabilities or response remotely … are more "art", these ones are more complicated" [middle managers at Security-Public]

This shows that the expert thinks some incident handling services cannot be defined clearly and be outsourced at a specific price. In these cases, Security-Public tries to give more flexibility to the provider in terms of the amount of resources available to handle the incidents. Sometimes, some contingent, small projects are defined for handling these incidents, instead of breaking them into standard services. For example, when a series of serious security attacks were launched on the government servers in parallel with an international conference, Security-Public defined an extra project (beyond the normal contract) to manage the incidents during this conference. In parallel, the senior experts inside Security-Public engaged in the process of handling novel incidents. As a rather stable

pattern in the last four years, non-standard services have become standardized by the provider and then, are added to new contracts.

Collaborating with providers

There are also cases that the internal experts enter incident-handling process to collaborate with the supplier. For instance, in Security-Public, when a novel incident happened, the experts from Security-Public collaborated actively with their provider, instead of fully delegating it to their provider. As a senior incident handler at Security-Public mentioned,

"In novel cases, I act as a member of the provider team. I take part in their discussions. We do tests on the faulty system …" [A senior incident handler at Security-Public].

As internal experts actively engage in incident handling, they get a practical sense of what is going on in this kind of incidents. They can, thus, determine to what extent these incidents can be delegated to the provider, what specific measures should be considered for assessing the quality of incident handling by the provider, and what are the risks of outsourcing their incident handling process.

Experts in the System-Administration group at Supercomputer-Large also engage with their specialized providers in incident handling process. These experts are highly specialized in terms of different technologies and systems. In many incidents, they work closely with providers' experts to identify the problems, find the solutions, and implement them.

Exceptions for outsourcing incident handling

In several situations, the organizations handled their incidents internally. First, Security-Private does not outsource its incident handling process, and handles the incidents internally. With the specialized experts, Security-Private handles the incidents of its clients and its own incidents internally. At Security-Public,

politically sensitive incidents are also handled internally, and the provider does not have access to these incidents. These incidents are not necessarily complex, but they relate to sensitive authorities or sensitive data. Another exception is Triage that is done internally by Security-Public.

To sum up, several practices during incident handling process help organizations benefit from their experiences of incidents to avoid them and manage them better in future. Although these practices were primarily taken for incident handling sake, they also helped the organizations benefit from their ISRIs experiences. The very actions that experts did to handle an incident (e.g., restarting a broken system) provided them with a detailed, intimate understanding of incidents, and practical capabilities and skills for handling similar incidents in future. In addition, the past experiences about other incidents are integrated and discussed through Triage collectively to define how the current incidents can be managed effectively. Finally, in their interactions with their specialized providers, the organizations were constantly adding to their mutual experience and repertoire of practices that they can use in handling other incidents. This adds to their capabilities to handle incidents internally, as well as getting their specialized providers handle their incidents more effectively.

4.2 Post-incident learning practices

The studied organizations adopted several practices *after* handling incidents that helped them examine the incidents and identify ways in which similar incidents can be prevented or, at least, be managed better in future. These practices are classified under “post-incident reflection” and “laboratory incidents analysis”.

Table 4.3 summarizes the two practices and the way they were differently adopted by the studied organizations.

4.2.1 Post-incident reflection

Reflecting on handled incidents, discussing them, and making sense of the experience of incident handling is a common practice that takes place through informal talks, periodic meetings, open seminars, and incident reporting systems. The current research was also another opportunity for the organizations to reflect on some of the already managed incidents during the interviews.

Table 4.3: Practices after incident handling process

Practice	Description	Specificities of the organizations			
		Security-Public	**Security-Private**	**Supercomputer-Large**	**Supercomputer-Small**
Post-incident reflection	Senior experts reflect on major incidents after handling them, through various informal and formal interactions, to discuss their experiences and ideas.	- Mostly in relation with the experts inside their provider	- Running "open seminars" across specialized teams	- Mostly through informal talks	- Adopting a formal incident reporting system
Laboratory incident analysis	Novel incidents are analyzed by specialized technologies to understand the causes of incident and how it can damage systems	- Done internally by third-level experts - The results are formally documented	- Done mostly by using the technologies inside the clients - In some minor situations are done in a trial way	Done for reproducing the incidents for the provider	- Done for reproducing the incidents for the provider

Informal reflection on incidents

The fact that experts and technical managers discuss incidents in their daily work informally reduces the need for specific extra meetings. In these talks, For instance, a senior expert at Supercomputer-Large mentioned that

"Well, we talk about the incidents [that] we handle all the time. My manager is sitting just beside me. We constantly talk about the issues" [a senior expert at Supercomputer-Large].

This indicates that reflection on past incidents is done informally as part of the daily routines.

Talking about incidents is also a common part of the daily breakfast and lunchtime in all the companies. Box 4.4 illustrates this point at Security-Public.

Box 4.4: Discussing incidents during lunch time

The operative manager of Security-Public nocks the door of the room in which ERI manager is meeting with some external visitors and invites them to have lunch together in a restaurant, located beside the building. He knows the visitors and says

"I am hungry. But take your time. I will wait for you. Let me know when you finish with your talk then we go for lunch together" [The operative manager of Security-Public].

After around 30 minutes the ERI manager and the visitors stop the meeting, go to the office of the operative manager, and call him for going outside for lunch. In Spanish culture, lunchtime takes almost one hour. There are four people at the table, two managers from Security-Public and two visitors. They order lunch and wait for the first plate. The conversation starts from the recent political changes in the local government and, after several minutes, it moves to discussing various recent incidents. The ERI manager talks about specifically a recent major incident

that was grabbing the attentions of a wide range of local organizations. He explains a bit the technical aspects of the incident, and then quickly moves to other similar incidents that he had recently come across when reading a weblog today. [The first plate arrives and conversation interrupts for seconds].

The ERI manager continues his speech, while the operative manager seems to be unaware of some information that he was saying about other similar incidents. At a point in the conversation, the ERI manager and operative manager appear to disagree about one of their actions that they took in the last incident. They both turn to the visitors, who are sitting in front of them on the other side of the table, and say:

"*Well, you know, we are still growing. We have to improve a lot our ways of working and handling incidents. These kinds of discussions are what we have every day, especially now that we are thinking about restructuring our service model*" [The operative manager and ERI manager of Security-Public].

The friendly atmosphere and close personal ties facilitates this process. Experts often bring their *interesting* stories to the discussion, to say something that is interesting for their colleagues. Normal incidents often are not talked about. The experts are not necessarily thinking systematically about any particular lesson that might help any specific audience. As one expert at Supercomputer-Large mentioned:

"Sometimes I just say an interesting point that I just learned from an incident loudly to my two other colleagues who are sitting just in front of me… it might help them " [an expert at Supercomputer-Large].

This shows that little preparation is done for such conversation and is mainly derived by the desire to bring out some unusual point in the discussions.

In addition, the specialization of experts makes it irrelevant that experts talk about the details of incidents because those technical details would be alien to their specialized domain, unless, the detailed discussion about incidents would be confined to a very limited number of experts. For instance, at Supercomputer-Large, network incidents are discussed only among the expert dedicated to network administration and his group manager. Furthermore, the fact that most of incidents are handled by providers reduces the need for discussing them internally.

Periodic meetings

Another common time for discussing major, handled incidents is during periodic (e.g., weekly or monthly) meetings. For example, in Security-Public, post-incident reflections are mostly organized through weekly meetings. At management level, the heads of departments who are senior experts in their own field meet every Tuesday morning. As part of this meeting, they talk about the big incidents that were handled in the last week. The discussion is mainly for informing other managers about incidents. So that, in case they found some other similar incidents, they could benefit from the experience of past incidents.

Similar weekly meetings between the head of projects, project managers and functional managers take place at Security-Private. According to the informants, together with daily activities and urgent cases, participants often also discuss any important incident, which occurred in the past week.

This discussion is often very short because normally other colleagues already know about the big incidents and because the details of the experience are irrelevant to them as they are specialized in other types of incidents.

In Supercomputer-Small, Most of the reflection on the internal incidents happens through weekly meetings that are called “hand-over meetings”. Each week, one of

the senior experts at Supercomputer-Small is responsible for handling incidents. This person, called Management on Duty (MOD), is the first one who is informed about the incident. He leads incident-handling process, by taking the initial control actions (e.g., disconnecting the faulty part or shutting down the whole system), calling the provider, and taking further recovery actions. At the end of each week, the current MOD, the next week's MOD, the head of technical department, and the head of LHC department meet and talk about what happened in the last week and plan for the next week. As the head of LHC department said,

"We talk [about] these incidents …. Manager on Duty has a list with all the incidents that has happened in that week…. And sometimes the follow-up actions … are also discussed" [the head of LHC department at Supercomputer-Small].

By embedding post-incident reflection into periodic meetings, the organizations frame such reflection as part of their ordinary work. This way, even major incidents are discussed as normal events, beside other tasks. Second, the organizations pick and discuss the incidents in the background of their current activities. They pick those aspects of the incidents that are relevant to their tasks at hand. For instance, in meetings about getting prepared for upgrading the main supercomputer at Supercomputer-Large, whenever the members of the meeting (senior experts from three units) saw it relevant, they commented on some specific incidents that they had experienced. This gives more practical relevance to the incidents that are discussed. Hence, the discussion about incidents in the meetings is often grounded in the current workflow, which makes the organizations connect their incident-reflection to specific practical decisions related to their current tasks. As another example, when deciding about which specific incidents should not be outsourced to the provider, the head of operations department at Security-Public bring out a recent incident that showed the difficulty of handling such type of incidents.

Most of the conversations about incidents in the meetings are shaped around *what* happened to give a hint to others to be aware of that, and *how* they can act on it in other incidents or make some other improvements. There are very few detailed technical conversations in these periodic meetings, because normally senior experts are assigned to their own domain of specialization. For instance, in Supercomputer-Large, a recent backup system damage was discussed to the level that the head of Users-Support group becomes alert about the potential failures that might take place to users who are trying to store huge amount of data in a short time.

Formal, systematic post-incident reflection

In addition to informal talks and discussions during periodic working meetings, the organizations also adopt practices through which they formally analyze their major incidents and extract lessons. An illustrative example of systematic analysis of incidents is incident-presorting system as Supercomputer-Small. At CERN, there is a public, web-based database of critical incidents that all tier 1st sites should report their major incidents there. Supercomputer-Small, as the tier 1st site of CERN, is asked from the very beginning (2003) to report its major incidents by developing an analytical report soon after incident is controlled. The reports are directly uploaded on the website that is accessible to other peer sites.

The expected audiences of the reports are technical experts in other tier-1 sites. As noted by the experts and managers at Supercomputer-Small, there have been some cases that experts from other sites have contacted Supercomputer-Small to inquire about further details on their reports, showing that the reports had been read by some colleagues.

The report is developed for “critical” incidents, meaning incidents that are important for all tier-1 sites and are important enough to be reported. CERN does not provide any definition or criteria of “major incident”. However, for

Supercomputer-Small, it is easy to identify which incident is critical enough to be reported. The head of LHC at Supercomputer-Small, who is in charge of developing such reports, elaborates this point by referring to the impact of incident on the main services as the criterion for judging about the criticality of incidents.

"We have many incidents and not for all of them we file an incident report. … there is a decision … is this incident sever enough to file in the services incident…. it is not such a huge decision. It is quite obvious sometimes. … . it depends very much on the time of the outage and what services were affected. … If the main services are affected, for a reasonable amount of time, … it is pretty obvious" [the head of LHC department at Supercomputer-Small].

The reports have a rather fixed structure (see Figure 4.5 as a typical structure of incident reports), *suggested* by CERN. It includes 1) description of the incident, 2) the impacts of incident, 3) timeline of events during the incident, 4) the analysis of cause, and 5) the follow-up actions, and 6) a final summary. The follow-up section is related to actions that should be taken to avoid similar incidents in future.

[Title of the incident]

Description

[What happened, when, where, the situation]

Impact

[The damages, the interruptions in the service, any data loss, potential impacts in future]

Time line of the incident

[The timing of events and the interventions]

Analysis

[The details of technical causes of the incident and the technical solutions provided, both those that did work and those that failed to be effective]

Follow up actions

[Other actions that should be done after controlling the situation, including both short-term actions and long-term ones]

Summary

[A summary of the incident and making sure that all actions are properly considered]

Figure 4.5: The structure of incident reports at Supercomputer-Small

Within this overall structure, there is a lot of room for adopting different styles and providing different levels of details. Over time, the reports have changed in terms of the content, mostly towards more specific and concise material, and just providing very crucial content. As the head of LHC department at Supercomputer-Small, said:

"I think there is an evolution … because of sort of experience and so on. … and the format … it is a suggestion top down but also very light level. Essentially the headers of the sections. There should be an analysis, there should be a timeline, there should be a follow-up actions.. more or less. But there is no guidance imposing in the level of details or the content" [the head of LHC department at Supercomputer-Small].

The head of LHC starts developing the report the day after the incident. However, there is no formal time limitation for that. As he mentioned

"We try to do it as soon as possible. But … we don't have a timeline. Sometimes … we can do … it depends on the work we have …" [the head of LHC department at Supercomputer-Small].

He starts the process by contacting the experts who were present in the incident handling process, especially the person in charge of incidents in that week (called "manager on duty"). He talks to them face to face, and sometimes on phone and asks about the story of the incident, what happened, actions taken, the current situation, the perception of the experts about the causes, and the actions that are needed to be taken afterwards. He takes notes of the information by trying to fill out the template: He describes this process:

"I am the one starting the template and just saying … OK, you did these things, please provide us with the explanation. Technical explanation that we can put

here. … just to make sure that everything which was relevant get into here" [the head of LHC department at Supercomputer-Small].

The report is developed through a shared document (Google Docs) to which all the experts involved in the incident have access. Most of the time, the head of LHC is the one who is writing the material. In some cases, an expert might just adjust a small detail or just give a small feedback. If any specific content is needed to be developed by an expert, the head of LHC specifically contacts that expert. At the end of the process, a final message is sent to all collaborators to get their final approval of the report. Normally, other experts do not respond. As the head of LHC mentioned "silence implies confirmation". However, it has happened that experts could not correct some mistakes in the report, just because they were too busy.

"It is true that probably writing something which is rotten and just the experts are too busy to read and will never find it… it could happen. Because we don't formalize …" [the head of LHC department at Supercomputer-Small].

Another important source of information is tickets related to the incident. However, the content of tickets that are filled by experts is often quite technical. The head of LHC has a PhD in Physics. He has more than 10 years of experience working with supercomputers. However, he does not have such a detailed technical knowledge, as computer engineers do. Thus, an important job for him in writing the report is to make sense of very technical issues that the engineers have put in the tickets. He does this by trying to abstract the technical content into more general issues and asking the experts to tell the technical stories in a way that he can understand. He explains this process:

"Dealing with this, I go and talk to the guy … and ask what hell is this. Because in the ticket … what most of the information that you find are like commands. Like Linux commands. So, I just ask … sometimes I just know what they are doing here is

starting a services or reconfiguring this module" [the head of LHC department at Supercomputer-Small].

Due to the details in the tickets, he also needs to filter out a lot of *technical noises* when he is analyzing the technical content related to the incident:

"I don't need to know the details. But … I need to know … what they mean, at the higher level. … I make sure that the relevant things are written here. It is a bit deciding what is important and what is not. Because … easily it can be quite a lot of noise. Technical noise. And … so, the cleaning up of that noise and putting in a human readable sentence is something … I am doing as the editor of this… " [the head of LHC department at Supercomputer-Small].

The last section of each report is devoted to "follow-up actions". These are actions that Supercomputer-Small identifies as the result of the analysis of the incident, to fix the damages of the incident, avoid it, and reduce its damage in future. For instance, in the case of cooling incidents, they considered "improving the shut-down process" as a follow-up action that helped them handle the process of switching off the whole systems in case of cooling incidents. Some follow-up actions are simply asking for further detailed technical examination (mostly by the provider) to identify the causes of incidents. These follow-up actions are not so systematically derived.

They are often based on examining the tickets related to the incident, the experience of handling it, the discussions during and after incident handling, and some ideas in the weekly meetings. However, the way in which the follow-up actions are articulated is mostly the product of reading the whole story of the incident (once the incident report is composed) and then thinking about those follow-up actions. This point is mentioned by the head of LHC when he said about how he develops the follow-up actions in the reports

"Typically… sometimes, they don't think about it and it is only at the end that … we read the complete story. It is a good way to detect the useful follow-up actions. So, I would say many times it comes at the end. … There are some which are not that obvious … that once everything is quiet and you read everything from A to Z, you say… why all this mess happens in the first place. And you say yes… because we have this typical monitoring that is not good enough. So, it is afterwards you decide let's … to make sure that it will not happen again, let's write it here" [the head of LHC department at Supercomputer-Small].

There is no systematic way that the follow-up actions written in the reports be implemented. Depending on their priority and available time and resources, some of these follow-up actions result in some internal improvement projects. For these internal projects, a ticket is opened for tracking its progress.

Another way of performing post-incident analysis is through open seminars that Security-Private has been formally implemented. Box 4.5 describes a typical open seminar at Security-Private. The story shows how these seminars have evolved over time from regular incident-specific seminars for the all the technical experts to rather occasional seminars with more homogenous audience.

Box 4.5: "Open Seminars" at Security-Private

From the very beginning, Security-Private designed two-week seminars among all technical experts in which one group or an expert from a group would present their work that they have done recently, reflect on it by examining what went right and what went wrong, and present the lessons they learned. As an important part of the seminar, experts were supposed to discuss the major security incidents that they handled for their clients or faced internally. The presentation had to be finished in 45 minutes. The meetings were scheduled often on Friday afternoons, when experts have already closed their weekly works and are relaxed from their operational duties. After the presentation, the slides of the presentation were

uploaded in a publicly accessed space in the internal Wiki. The meetings were not obligatory, but the invitation was sent to all the project teams.

The original idea of the seminar came from the CEO of Security-Private, as he put out:

"Every two weeks, the team inside Security-Private have their open sessions in each project that they reflect on what they have learned from incidents handled and they share their experience. In some cases, the incident is so complex that we do not explain the lessons learned to the client. But learning from incidents help us to suggest new services to the market" [the CEO of Security-Private].

Security-Private could hold these seminars for almost the first year of its operation. In practice, most of these open seminars were gradually scheduled at the end of each project. This way, the experts could reflect on a whole project, in which, they often faced various incidents. Then, they could pick some critical and novel incidents to discuss in open seminars.

"And at the end of the project the manager thinks about what elements of the gained knowledge should be distributed and shared with other people… When each project finishes, we do this seminar, talking about 1) what has been done in this project, 2) what has been done well, 3) what has not been done well, 4) what has been learned and where is it situated in the intranet" [the projects manager at Security-Private].

After around one year, the open seminars took place in a monthly manner. It was mostly because experts were too busy to have seminars and prepare for that every two weeks. Another challenge was that often the contents of seminars were too specialized, making it irrelevant for experts from other groups.

Gradually, these open seminars took place irregularly, mostly determined by the load work of the current projects. However, the CEO, as the proponent of the

seminar, has been insisting on having the seminars regularly. In practice, the open seminars of many finished projects just took place long time after, when the load work of experts diminished. In the last year, except for a few seminars, most of the seminars were postponed to summer time, when the clients are on vacation, thus, Security-Private has more time to dedicate a whole week for holding a series of open seminars, all together.

Reflections during the interviews!

Finally, organizations reflected on their critical incidents also when I interviewed them. Even in case of several critical incidents, they had never had the chance to reflect on it in a specific meeting. In several occasions, the managers and experts mentioned

"emm… well… interesting, Now you are asking me, it is my first time to think about this …",

For instance, in a recent incident at Supercomputer-Large, the thermometer of the main supercomputing room alarmed as the temperature went above the threshold. This is a very dangerous event since high temperature reduces significantly the processing power and can easily damage the hardware. They made quick actions to reduce the temperature by increasing the cooling power. However, it was very surprising that they could not feel any higher temperature or any event that might increase the temperature. Two days later, the technical expert from the provider came and started measuring the temperature using his own thermometer. Surprisingly, he found that the temperature is much lower than the threshold. In comparing with the thermometers in the room, it appeared that the problem was that these thermometers deviated from their calibration, so, they showed a higher temperature than it really was. Around four months later,

when the research team interviewed the head of Facility-Management group, he mentioned that they have never thought about what they could learn from this incident.

"Maybe nothing. Because the providers do all the job" [the head of facility manager group at Supercomputer-Large].

Thinking a bit more, just a few seconds after this point, he interrupted and said

"So, we know that when the next time they come we have one more item in the checklist. I can say we learned that we can put a point on the checklist to check the calibration of the thermometers" [the head of facility manager group at Supercomputer-Large].

In other similar cases, inside Supercomputer-Large, I found instances that the experts did not formalize their learning from specific incidents. They said, "It is quite obvious", which does not need to be done in such a formal way. In some cases, they were the only person in the whole organization dealing with the incident, so that, it would not be needed to do any formal assessment of what the expert perfectly has in his mind.

4.2.2 Laboratory incident analysis

In addition to reflections on incidents after handling them, security companies also performed laboratory analysis on novel security incidents to figure out how they are working and how they can be prevented in future. Box 4.6 illustrates a typical example of laboratory analysis at Security-Public. The laboratory analysis focuses on *novel* incidents that have some technically unknown elements for the organizations. The aim of laboratory analysis is to understand how the novel incidents work and create damages to the systems. It is mostly a technical process that only senior experts who have been trained on that can perform. It requires

specialized technologies and instruments. The results of such analysis are often translated into some specific new methods, tools, and databases of incidents, from which incident handlers can act on similar future incidents.

Box 4.6: A typical example of laboratory analysis at Security-Public

Take the example of a new virus that has affected a hard disk. During incident handling process, the hard disk is first disconnected from the other devices to avoid distributing the virus to other systems. The system resumed working by replacing the infected hard disk and retrieving data from a backup. After handling incident, the infected hard disk is sent to laboratory, with special care to make sure that the hard disk is well isolated and the data on it is frozen. At laboratory, there are special technologies, such as isolated hosts and specific test applications that are adopted for laboratory analysis purposes. These technologies are rather expensive and they constitute an important part of the investment by Security-Public. These technologies are also specialized to different types of incidents. For example, some technologies are specialized for analyzing incidents in mobile devices. Even when the provider of Security-Public needs to do some laboratory analysis, it is done by these technologies that are inside Security-Public.

At security-Public, only the head of ERI department and two high-level experts (they are third-level experts who are specifically trained for incident analysis) perform laboratory analysis. They take the infected hard disk and connect it to the analysis technologies. The software tries to re-activate the virus in a protected virtual environment to examine its behavior in various situations. Several experiments can help experts identify how the virus works and how it might affect systems. This sometimes requires several trial and errors to run the virus under different conditions to detect its behavior. For instance, the virus might be activated (to reproduce the incident) in a controlled network and its behavior be monitored to understand how it can damage systems. They might also analyze the

log of the affected system to track the behavior of the virus. There are always many uncertainties about how the virus might behave in new situations.

Only novel incidents that their behaviors are unknown for the experts are sent to laboratory analysis. For example, in 2011, Security-Public sent only 100 incidents to the laboratory analysis (out of around 500 managed incidents). As the head of operations department mentioned:

"Some incidents are highly complex or technically sophisticated. Especially if the incident is abnormal, we analyze it more seriously. The time that we spend for the analysis of an incident depends on the complexity and severity of the incident" [the head of operations department at Security-Public].

The results of the analysis are often documented in specific articles inside Wiki system. As the result of this analysis, experts might detect some defects and vulnerabilities in the existing systems, which, in turn, can result in new tickets for fixing them. The insights that experts gain from this analysis help them be aware of new sources of threat and control them in case of future events. The head of ERI at Security-Public referred to this point by saying:

"They learn why such incident happened, what should we do to avoid it in future. What should we do if it happens again…. Then they make recommendations for security issues. For example, they say that this incident happened because this proxy was not well set up. Or this firewall is not well tuned" [the head of ERI at Security-Public].

At Supercomputer-Large and Supercomputer-Small, laboratory analysis of incidents is not done as a specific formal task. Instead, in case of some novel incidents, they try to *reproduce* the incident in order to make sure that the incident is not due to random events. In addition, reproducing the incident is

needed for opening the case for the provider and showing to the providers that the problem is related to systems that are maintained by those specific providers.

To summarize, the organizations adopted a wide range of formal and informal practices to reflect on and analyze critical incidents once they are handled. The results of these analyses can help organizations be more prepared for similar incidents in future or take some actions to avoid them. Overall, Security-Public and Security-Private have more systematic ways for post-incident learning, while these activities at Supercomputer-Large and Supercomputer-Small are done more informally, as embedded parts in daily practices, on a less regular basis. The exception is when some formal structures are imposed from outside (the case of incident reporting at Supercomputer-Small). Most of these practices, although are designed to be done collectively, in practice, take place among a limited number of experts in the same domain or simply by single experts individually. Only senior experts (level 2 and 3) are involved in these practices.

Chapter Summary

The chapter narrated the practices that the organizations adopted in relation with each major ISRI that allowed them benefit from their incident experience to avoid and manage future incidents. Two categories of practices were identified: practices *during* incident handling, including handling incidents based on their ticketing system, performing Triage, and interacting with specialized suppliers, as well as practices *after* incident handling process, such as post-incident reflection and laboratory analysis. However, there are practices that the organizations adopted in relation with *multiple* incidents that are described in the next chapter.

Chapter 5: Findings: Multi-incident Learning

The previous chapter described practices that the organizations adopted during or after handling each major incident to avoid similar incidents in future or reduce their impacts. The organizations have also adopted other practices for learning purposes that they are related to *several* incidents (multi-incident learning). In fact, these practices allow organizations to capitalize on the experience of several incidents that have some similarities and relations. This chapter narrates these practices that are classified under 1) using Wiki system, 2) internal improvement projects, and 3) upgrading systems. Table 5.1 summarizes the three practices and the differences among the four organizations in terms of adopting these practices.

Table 5.1: Multi-incident learning practices

Practice	Description	Specificities of the organizations			
		Security-Public	**Security-Private**	**Supercomputer-Large**	**Supercomputer-Small**
Using Wiki System	The organizations try to document their incident handling experiences and improvement ideas in their internal Wiki system and share it among their experts.	- Both internal experts and provider's experts develop Wiki content - The content is specialized based on the departments - Lack of time for documenting experiences by senior experts - Gradual nature of writing articles and developing them over time (semi-empty articles)	- Wiki as a repository of all information that might be needed in handling incidents in future	- It is only used for documenting the tools and changes that are developed internally	- Only using for documenting the tools and changes that are developed internally
Executing improvement projects	The organizations run internal projects through which they improve some aspects of their technologies or their processes, often by developing an open-source application or script.	- The projects are defined by internal senior experts and outsourced as commercial services to the providers	- Done as part of the services for the clients - In few cases they have been organized as specific internal projects	- Improvement projects are the main tasks of senior experts - Relying on open source applications	- Limitedly done, only for short-term changes - Mostly avoided since the specialized providers automatically solve the issue in the next upgrades
Performing Upgrades	Regular upgrades of systems and parts of them help the organizations avoid past incidents to occur in future.	- Mostly done by the provider	- Done for some open-source applications	- Easily done for small parts and parts that already have their own financial resources, but delayed for major upgrades	- Done frequently to a wide range of technologies

5.1 Using Wiki system

The first category of practices pertains to articulating the experience of incidents into Wiki system. All the studied organizations use internal Wiki systems

(sometimes, called knowledge-based systems) to document and store their experiences of handling incidents and other useful information that can help them avoid similar incidents or handle them better in future. This tool though, is used differently in the organizations, helps them keep some tracks of critical incidents by documenting its experience and ideas for improvements.

Security-Public has an internal Wiki system, called knowledge-based system, to document and maintain the knowledge gained from handled incidents. It has been one of the key ideas of Security-Public founders that all the knowledge and information should remain in Security-Public systems, even though the provider handles them. For this reason, one of the duties of the provider is to introduce knowledge and important issues into the Wiki system. Security-Public urged its provider to document all experience, tools, knowledge, information, methods, and lessons that they develop and gain through handling its incidents, into the Wiki system. This way, they want to make sure that

"All the information of the cases is in the knowledge database" [e.g., the head of LHC at Supercomputer-Small].

In parallel, the experts of Security-Public continuously compose articles related to various services and incidents that they face. These articles are only used by the internal experts and the experts from the provider, so that external users do not have any access to them.

The content of the Wiki is structured around specialized departments (e.g., incident response, operations, and legal services). This way, each department has their own space for creating sub-categories of articles. However, there is no imposed structure on the content inside each department. Normally, senior experts who are also the heads of the departments identify an important topic as they are handling various incidents and develop an article about that topic. For instance, a recent article about "how to protect your Smartphone when

downloading apps" was created by the head of ERI department as he detected several serious incidents on this issue. This way, most of the articles are not specific to any single incident. Rather, they focus on a specific solution.

All experts inside Security-Public have access to all the content in the Wiki system. The provider only has access to articles in ERI departments. Normally the articles that the provider writes are specific and technical (e.g., how to isolate a Mac-Hard disk that is affected by some virus). Using the Wiki system, other experts who have the access, can revise the article (removing parts, adding some content, and putting comments on the article).

Except for very few cases, these articles are written individually. The owner (writer) of the article is known. However, quite rarely an expert revises the articles that are written by others. Interestingly, articles are mostly read by those who write them to remember specific ideas. The exception is when the expert who is dealing with a specific domain is leaving the company or is on vacation. In these situations, other closely specialized colleagues read articles to figure out how to perform some specific technical actions.

The legal expert is almost an exception who tries to read the technical articles of other technical groups. The most active group in terms of developing articles is the incident response team as they are constantly facing new themes and issues related to handling incidents. In addition, the head of incident response team takes Wiki as an opportunity to document his experience in handling incidents so that the provider can easily access it. This, in turn, facilitates the process of outsourcing incident handling to the provider in future.

Since Security-Public is in its initial years of development, the experts have many ideas in their mind, but they are waiting to spare some free time to document them into knowledge articles. Therefore, experts often have a list of articles in their mind or in their personal notes. They mostly enter Wiki system to write

something, rather than reading the existing articles. The main bottleneck seems to be time limitations. They have the system open, in parallel with their daily work of handling incidents. They constantly take notes and write articles as a progressive task. There is no specific plan, unless for the year 2012, it is planned that the system be fully organized, which means that

"… everybody knows who, how, and what is wiring the other people" [the head of operations department at Security-Public].

Box 5.1 describes the process of developing and using a typical knowledge article at Security-Public.

Box 5.1: the process of developing a typical knowledge article at Security-Public

The legal expert at Security-Public recently wrote an article about the legal considerations of handling security incidents in smart phones. He decided to write this article after receiving a series of security incidents related to smart phones. He started by opening an article, with a name and just a few bullet points, after the experience of a few initial incidents. At that time, still he did not have such a clear mind about the exact content of the article. As new incidents related to smart phones happened and new legal aspects of such incidents were more discussed in legal communities, he gradually got more specific ideas. Opening a rather empty article helped him to collect information and ideas about it, and in case other peers had some ideas on it, they could share with him. He took notes of ideas through successive incidents, until a time that he felt that

"I have enough in my mind to put into the article" [the legal expert at Security-Public].

Then, he started writing the article, when he could spare a few relaxed times. He mentioned that he often postpones writing articles until a time that he feels that he has got to a rather *"stable knowledge"* about it. This is important especially in

the IT sector as many technical issues change quite rapidly. Another reason is that the content of knowledge often changes dramatically when the expert experienced the second, third, and next incidents of a kind. The legal expert described the nature of legal knowledge of IT security, as it is highly subjective, rapidly changing, and a matter of holistic interpretation.

"On the legal part, that's quite curious,… I mean the legal knowledge is different, … is quite abstract…. We normally say that for some particular situation, we have seen that the court has decided this or that…. And then it is quite probable that if you do this, this will happen. You cannot be certain that this is going to be every time, … tomorrow there is another decision, from the court that is different. …. " [the legal expert at Security-Public].

This led the expert to avoid hasty reflection. Therefore, he postponed writing the article until he reached a rather stable and holistic view ("a big picture" as he mentioned) to be reliable enough to be put in an article:

"That's why I am also quite cautious with introducing some kind of particular information…" [the legal expert at Security-Public].

The uncertain nature of legal knowledge adds to this situation, as even seasoned experts cannot be sure about their judgments. This made the expert consider the different interpretations and changes, which might not fit with the content of their article. He stressed this point saying:

"For my experience in this sector, from 10 years ago, we have worked on absolutely grey areas, … …. You need to understand what means Internet, and what is security, … and the data protection part, … the legal procedures part, the contracts part. It is like trying to build up a theory from understanding how the legal part works" [The legal expert at Security-Public].

At Security-Private, the Wiki system is used as a repository for all information that is worth storing, in addition to the ticketing system. Security-Private does not allow its employees to use any other data storage rather than the Wiki. All information, files, and documents, about both technical and non-technical issues must be stored in the Wiki system. This includes everything such as lessons learned from a specific incident, an idea for improvement related to a specific incident, some new organizational procedures that can improve the work, as well as more general content such as the formal procedure of interactions with the clients in reporting incidents.

Each user has his own personal space for storing and even sharing information that is related to the person, his role in the organization, and the involved projects. This part is mostly consulted by its owner except in situations that some specific invitations is sent to other colleagues to see some content there. This part is more related to the personal notes, tasks, concerns, ideas, and personal documents. The rest of the content is structured around projects. Projects have an overall fixed structure which includes: (1) “Management” of the project that includes all documents and information about how to manage the project, its methodology, schedule, costs, managerial reports, and so forth; (2) “Work”, including all documents that are used and produced during the project, such as technical reports, technical methodologies, and technical tools; (3) “Deliverables”, which includes all the documents and outcomes that are delivered to the client as the deliverables of the project. Only those who are participating in each project have access to these three areas.

In addition to personal and project spaces, there are some general spaces for documenting ideas for improvements, ideas for new services, and best practices that are related to various projects. Very few articles are related to a specific incident. Only a specific part of the Wiki is devoted to open seminars (mentioned

in previous chapter) for uploading the slides of all open seminars. This space is shared and all groups have access to it. Another category is devoted to documenting all the tools, methods, and innovative solutions for handling incidents that Security-Private develops them when delivering security services to its clients. The CEO, project managers, and department managers have always thought about their Wiki as *"a place that everything we need for our work should be put there*".

Although the Wiki provides opportunities for collaborative writing of articles and commenting on other's articles, this rarely happens, as most of the articles are written individually. For example, just in a very special case, the head of incident management team created an article and one of his experts commented a few words on it. However, this feature is used for developing reports for the clients. For example, in a recent project, they have to create a manual that provides security instructions for citizens when using Internet. This document was created through an article in the Wiki system. The project manager created an initial template (with the overall bullets that shows the structure of the report). Then, it was shared with two other experts and the head of the group. They divided work and each expert was responsible for writing a specific section. There have been some interactions on the document, when an expert asked another colleague to comment on a specific section. Therefore, this collaborative feature is used to support developing articles in an individual way.

At Supercomputer-Large, the Wiki system is considered as a complementary documentation system. The content of the Wiki is updated only when a new tool or new procedure is put in place (e.g., a new script that detects network failures) to describe how the tool works. Figure 5.1 shows an example of Wiki entries at Supercomputer-Large that describes how jobs should be executed on a specific

machine (please note the specific information is not presented for the sake of confidentiality).

How to run jobs on Machine XYZ

Introduction

[The general characteristics of the machine, its processing capabilities, its architecture, normal algorithms that can be used, and the limitations]

The steps for running jobs on the machine

[Explaining the steps for uploading, executing and examining the jobs on the machine, with some examples of instructions, and the codes that are needed for running the job and comments on how to use them and adapt them for different cases]

Practical implications

[The possible problems that might happen in getting jobs run on the system, an example of each and solutions, and some tricks for solving potential problems (e.g., put a space character before the first line juts to avoid some strange crazy errors in compiling the job)]

Figure 5.1: An example of a Wiki entry at Supercomputer-Large

This is mostly to help the same expert who has developed the tool, remember the technical details when he is trying to use or change the tool in future (personal memory). In some cases, the articles help other experts in situations that the principle expert is absent or leaves the company.

Almost all articles are about a specific technical solution that has been developed for a specific problem. This way, articles often have an introduction of what is the problem, then it follows by the solution, often accompanied by the source code of the software developed (the script), and finally, is completed by some practical considerations about how to use the tool. For instance, an article is about a recent solution that allows laptops remain connected to different Wi-Fi zones without being disconnected when the laptop is moving from one zone to another zone. The same person who developed the tools is responsible for documenting it in the Wiki system.

All the experts from the same group have access to all articles. However, in practice, just in exceptional cases that the owner of the article is not able to handle the case, other colleagues might consult others' articles. The articles are often written after the development of the solution. It sometimes takes a long time (even more than six months) for the experts to find a relaxed time to document the tool. The updates in the tools can make the experts revise their related articles.

Supercomputer-Small uses its Wiki system very similarly to Supercomputer-large. It is used for documenting specific information that can help in future. It can be about a specific solution, how to deal with a specific type of incident, or simply documenting a specific tool or procedure. In fact, Wiki is used as a *complement* of ticketing system. This is described by a middle manager at Supercomputer-Small,

"I would say it is for documents which are less dynamic than the ticket. The ticket is a flow of information. And when it is like procedures… for instance … systematic procedures …. So this is like a document which is … les dynamic and this is in our Wiki…" [a middle manager at Supercomputer-Small]

Similar to Supercomputer-Large, the articles in the Wiki are mostly consulted by those who have written them, except when the author is absent and some urgent needs arise. Even the authors of articles consult their own few articles only when they want to remember some specific technical information.

Wiki system, at Supercomputer-Small, is not so critical, as it is used only for specific contents. In addition, experts use their own ways, such as an Excel file, to document some information more easily. The articles in the Wiki system are not incident-specific. They are related to overall themes or several incidents. Incident specific content is documented in the ticketing system. Distinguishing between what should be stored in the Wiki system, versus what should be put in the

ticketing system is reflected in the following point said by the head of LHC department:

"The Wiki perhaps the information is more transversally … ticket is the first level documentation" [the head of LHC department at Supercomputer-Small].

Unlike the security organizations, both supercomputers have a thin content in their Wikis (very few articles). However, the security organizations have tended to document all information that they think might be needed in future.

To conclude, the organizations use internal Wiki systems to help them document their experience in handling incident and ways of improvement for future incidents. However, much of the content is not specifically related to any incident. Moreover, the content is often individually created and often is used by the same person who has developed it. In fact, although the content is shared, it is highly specialized and is used mostly as a memory support for individual experts. In addition, there are variations among the organizations in terms of how to deploy Wiki system and how to incorporate it into their daily work. For security organizations, it is a crucial, active tool, while for supercomputer organizations it is a support tool that is consulted contingently.

5.2 Executing Improvement projects

The second sort of practices that the organizations adopted in relation to several incidents is executing improvement projects. The studied organizations have constantly run improvement projects to avoid incidents in future and reduce their impacts. In fact, defining internal projects and creating new tools as the result of experiencing (repeated) incidents is a routine practice in the organizations. These projects often focus on addressing a specific change or making a specific

improvement, mostly by developing a specific technical solution such as a new software application, and a script that improves the functionality of the existing systems. The organizations often work on open-source applications that allow them to make changes that solve their specific problems. In a few cases these solutions requires some hardware accompaniments. These projects are done by senior experts (from level two or three), and are guided by middle managers (managers with strong technical background).

At Security-Public, the internal projects are mostly outsourced to the provider. In fact, the provider does two things for Security-Public: handling the incident, and executing long term projects. For example, a recent improvement projects at Security-Public is designing, acquiring, and implementing the whole set of technologies and methods for handling mobile security incidents. Some of these projects are suggested by the departments' heads who were observing some needs as the result of facing several incidents. For example, the provider was asked to develop a database of vulnerabilities that Security-Public has been using for detecting security vulnerabilities. This project was defined because the head of operations department at Security-Public realized that the traditional vulnerability database does not include some recent vulnerabilities that they were detecting in some incidents or through their laboratory analysis. Sometimes, the provider suggests some similar projects, during handling incidents. For instance, the provider suggested a specific Wiki system to Security-Public that could help them document their experience more effectively than their traditional file-system.

Each project is defined as a specific service and should be delivered within a specific budget limit. Normally the senior experts inside Security-Public define the project, delegate it to the provider, lead it through continuous contacts with the provider team, and finally test and approve it. For instance, for developing the database of vulnerabilities, the head of operations department defined a clear

project, with the specifications of the database, and with a clear schedule. He was helping the provider team during the project in designing the structure of the database and implementing it. In the case of legal services, the legal expert from Security-Public directly contacts some legal companies (that are different from the technical provider) for outsourcing specific projects.

Security-Private also runs internal projects for improving its work because of experiencing various incidents. However, most of these activities are defined as part of the service that they deliver to their customers. It is because Security-Private does not have enough extra resources to spend on developing internal projects. As the head of projects at Security-Private mentioned,

"habitually, these things have been in the form of small projects… small development,… by small, I mean those that require like 10 working days … 15 days of working … and very good experts can do things very good in a very short time… but these hours normally is charged as part of the work for the client….. all these things are paid by the client" [the head of projects at Security-Private].

This way, Security-Private can improve its work and do not need to spare extra resources for the improvement projects. In a few cases, however, there have been some small internal projects funded by Security-Private. The projects that cannot be framed as part of the work for the client are very carefully discussed between the experts who suggest the project and the head of the projects. Once the idea is initially assessed and has been proved to be a profitable proposal, it would be suggested to the CEO. He is also the head of "business development" department. The proposal should have a clear estimation of benefits for the company. Upon the approval of the CEO, it would be assigned to an expert or some experts to be implemented. The costs of these projects are kept under control.

Recently, in some cases, Security-Private faced challenges with some of its clients about the property rights of the tools that they had developed for delivering their

services. For instance, during delivering some web security services, Security-Private team could develop a script that can quickly search critical nodes of a network and find security threats. This was done as the result of an urgent need of a client. Although the times that experts have spend on this project were considered as part of the service to the client, the developed tool per se was not a deliverable to the client. Security-Private is trying to clarify the property rights of the tool with the client.

At Supercomputer-Large, internal projects are central to the daily work. This is particularly true for System-Administration and Users-Supports groups, which constitute the main technical groups inside Supercomputer-Large. There, more than 70% of the time of experts is devoted to their improvement projects. In fact, for the experts (often from level 2), incident handling is an *inferior task*, which distracts them from their "main, important" tasks, which are developing their internal projects. Therefore, the experts have implemented a turning system in which each expert is dealing with the incidents for one week, and the other weeks, he is free to work on his own projects. Box 5.2 explains a recent improvement projects at Supercomputer-Large.

Box 5.2: Central Storage Mirroring Project at Supercomputer-Large

Supercomputer-Large started a project to install a new central data-storage (see Figure 5.2) that would provide services to various high performance computers (HPCs). In parallel, one of the main HPCs (HPC #1 in Figure 5.2) had to be replaced with a new one (HPC #2). Therefore, the data stored in HPC #1 had to be moved to the central storage because the new machine (HPC #2) does not have any internal storage. However, meanwhile, Supercomputer-Large was not able to make HPC #1 work with the central storage because it required a lot of modifications and changes in the operating system and related applications. In addition, it was because still the central storage was not reliable enough to be connected to the

main HPC. This created a transitory situation in which Supercomputer-Large had to keep the data in inside HPC #1, and at the same time, start migrating parts of the less active data in HPC #1 to the new central storage.

This transition period took more than one year, even though it was supposed to last only a few months. It was mostly because of the delay in buying HPC #2. Meanwhile, some users who were intensively using HPC #1 and some other machines faced several problems. The first critical problem came from a user who had processed his job, and the huge amount of produced data had to be quickly stored (almost without delay). Otherwise, all the millions of processing-hours used for producing the results would be lost.

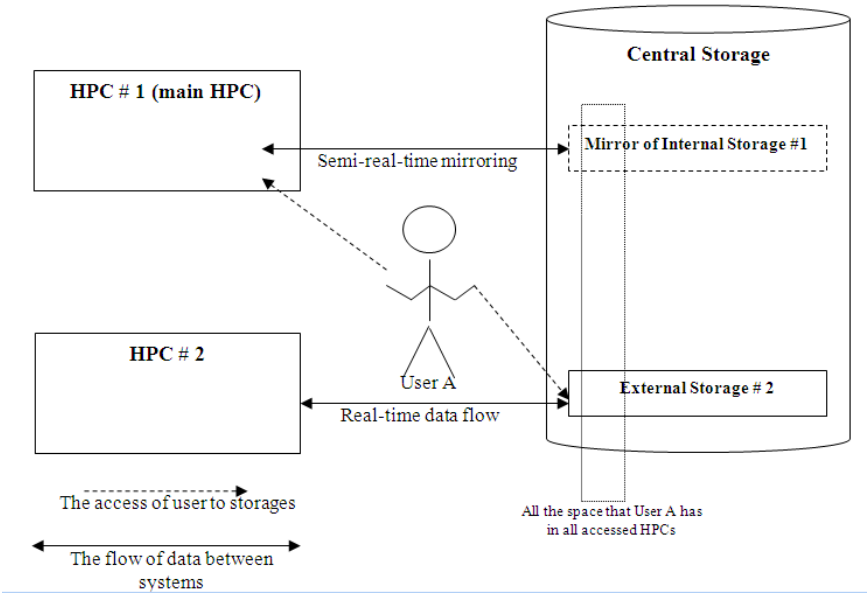

Figure 5.2: Central storage, related HPCs and the user's space

The problem was that the user was receiving a general error "input-output error" when trying to copy the results to its storage space (quota). The user was scary about losing the whole results. Then the user reported an error and the ticket arrived at the Users-Support group. Since the Users-Support group does not handle basic infrastructure, they forwarded the ticket to the System-

Administration group. Due to the urgency of the case, experts at System-Administration extend the storage space of the user temporarily to allow him to store his results. The Users-Support group got back to the user and informed the user about the resolution of the problem.

The experts at System-Administration group who were responsible for handling the quotas and the transition to the central storage started thinking about whether this way of solving the issue was appropriate or not. This point was discussed, in the days after, between the experts and their manager who was aware of this critical incident. They did not reach a concrete solution because there were other more urgent tasks to be done.

After one week, the second incident came from another user who received similar error, with the same problem. This made the experts at System-Administration group think more seriously about some other solutions. The third similar incident arrived the following week. Again, the same process toke place for resolving the incident. However, this time, the System-Administration group expressed its concerns to the Users-Support group about turning this solution into a routine. It would be problematic, because they had to extend constantly the spaces available to the users, while there were quite limited spaces available for each user.

The experts in the System-Administration group shared their concerns with their manager. This point remained in the mind of the manager until the next weekly meeting among the groups' heads, where the heads of the Users-Support group and System-Administration group discussed the issue. Users-Support group preferred keeping this temporal remedy solution as it was the easiest way to deal with the users, and keeping them happy. However, System-Administration group showed objections and requested that clear limits be defined for the extension of spaces. The argument of the System-Administration group was that this is

something that we have to “teach them” that they do not have unlimited space, and these incidents are good excuses for making them learn that.

Finally, they came to an agreement that in similar cases, they only extend the space by 15%, and Users-Support group should inform the users that they had to start deleting their legacy files that they did not need any more to create free space. However, there was a discussion that critical and sensitive users (e.g. the heads of the research departments) are exceptional and should be treated with more flexibility.

The senior expert at System-Administration group who was mostly taking care of this issue suggested to System-Administration group head that “*why not to create a tool (a specific code) that solves several issues together*”. He mentions that they have had the problem of unclear messages to the users. That is, instead of giving a specific error message that shows the specific problem (e.g., “the space is full”), it gives an unclear basic error. He added, “*We can automate the 15% extension of the space by this tool. We also know that there have been some imperfections in our mirroring process (copying data from HPC to the central storage for having a provisional copy)*”. “*Now that we are creating a tool*”, he added, “*let us create a script that solves all these three issues together*”, he continued.

The manager agreed about this internal project and started working with the expert to develop, test, and install the script on the system. Once the tool was finalized, the expert who was working on it created an article in the Wiki page to describe what this is tool for, how it should be used, and practical points about how to use it. He also copied the source code of the tool (script) in the article to be accessible for the other colleagues.

The tool worked properly, especially when the fourth incident arrived. However, the fourth incident was slightly different. This time, the user was supposed to have enough space in the main storage. However, the system was saying that the space

is not enough. It was because the user was using several HPCs in parallel. In these cases, there is one single overall storage space. Therefore, if the user exceeds its storage limit in one of the HPCs, this would deduce from the space available in other HPCs. This point although was solved by adding the 15% of space, opened a new issue among Users-Support group and System-Administration group that they should "*change the mindset of users about local spaces for each HPC" to "a universal space for all HPCs*". This resulted in a series of training activities for conveying this point to the users.

Table 5.2 summarizes the events happened during each incident in the second column (mostly captured through ticketing system), the actions that Supercomputer-Large took related to developing the new tool (improvement project) in the third column, and the formal documentations of the project and tool inside their Wiki System, in the fourth column.

Table 5.2: Central storage problem at Supercomputer-Large

Time (tentative)	Ticketing system	Learning practices	Documentation in the Wiki knowledge base system
Day 1st	Ticket 1 (day 1): - User cannot copy its results in its disk space - Actions: o Checking with System-Administration group o System-Administration explains the issue that is because of the mirroring o Discussion between System-Administration and Users-Support on how to solve it o Final decision: let's handle case by case and increase the space for urgent cases - Answer to the user: your problem is solved. Now you can copy it		
Some days after		"I started thinking that should we keep solving this issue in this way"? "I started talking with my boss, when we were taking coffee" "he already knew about this issue" We said, let's see.	
Day 7	Ticket 2 - The second User: cannot copy its results on the disk space - Actions o Sending request to System-Administration to increase the space of the user o System-Administration increases the space o System-Administration: this cannot be done each time in this way... let's make it a bit automatic o Starting to think about the overall idea of doing that automatically - Answer to the user: your problem is solved. Now you can copy.		
The days after		I was thinking that this way cannot be maintained. The problems will be solved when the new storage system arrives, but we don't know when it will be done. Maybe we can think about some mid-term solutions, something like a script that handles this automatically	
Day 15	Ticket 3 - The third user: I cannot copy my files on my disk space - Actions o Sending request for extending the space to System-Administration o But explaining to the Users-Support that this should not become a routine! The basic solution is that we make users delete their legacy files, and get them know that they have quite limited space - Answer: This happened because you have exceeded your space limit. We already extended your space, now you can copy it. But next time, you need to first free space by removing your legacy files		

Time (tentative)	Ticketing system	Learning practices	Documentation in the Wiki knowledge base system
A couple of days after		I said to people in the Users-Support that this way is not working. Let's decide about a specific percentage of increasing the space as something fixed.	
In the next weekly meeting		The two heads of the two groups discuss the issue. They agree about extending 15% of spaces of users that face this problem. But at the same time taking care of the sensitive cases	
A few days after the meeting		I started writing a script that does several things together 1) instead of generating an input-output error, it gives a clear message that what is the problem, 2) it automatically increase the space of the user, 3) it also improves the mirroring process from old storages to the new storages in a way that users who do not face this problem would not notice the movement.	- Documenting the script (explaining what is it for, how it should be executed, and some practical considerations) - The code of the script (the source code is copied inside the article)
After around 1 months	Ticket 4 (around 1 month later) - The fourth user: I cannot copy my files. Also I cannot see why all my space is occupied. I have only 10 G files. So I must have enough space… - Actions: o Sending request to System-Administration o System-Administration gets back to Users-Support saying that it is because the user has files in other HPS and they are now centralized to its central space. That's why he cannot copy. o Negotiation between System-Administration and users support for solving this specific case o System-Administration extends temporarily the space, so that the use can copy it - Answer: this happened because the files that you have in other machines are occupying your space. We have integrated them all into your central space. Next time, make sure that you have examined all your files, and remove unnecessary files before you start copying a new file. Be sure that you consider all the files that you have in all machines together.	The script was working quite well in this case.	
Several days after		I tried to do some improvements to the script. I do it once in a while. It is now almost stable. Works fine.	A minor revision of the script in the article (just copy and pasting the new code inside the article)

* Quotations are by the senior expert responsible for the improvement project

Like the example in Box 5.2, senior experts often suggest the projects as they face several specialized problems. Each project aims at improving a specific problem by providing a technical solution. The experts often provide their ideas about improvement projects in their daily talks and weekly meetings with their group managers. Group managers are always hands-on in the incident handling process, as well as their internal projects. These technical managers know almost everything at a very detailed technical level. However, they often just lead the activities and intervene when they feel that things might go wrong, if they do not help their experts.

The heads of groups provide some insights about the priority of the projects. Each senior expert, who is specialized in a specific technical domain, has its own list of internal projects to develop. The expert goes through the list based on the priorities. Often experts are working on two or three projects at the same time. Most of the projects only involved one expert, except a few projects in which two experts collaborate. The result of each project is first tested by asking other colleagues to use it internally. Then, the tool is implemented in the whole Supercomputer-Large. Finally, the expert who has developed the tool is responsible for documenting the tool in the Wiki system. This includes defining what is the problem of interest, what is the solution (describing the tool and how it works), and providing practical guide for using the tool.

In some cases, Supercomputer-Large preferred to perform improvement projects internally, although, the support contracts with the provider could allow outsourcing them. For example, in a recent case, Supercomputer-Large faced a problem in one of its backup systems. They had to modify the operating system of the machine in a way that shows exactly where the process of copying files fails. The expert in storage started working with the provider company. After several months, it turned out that the company is not capable to solve the problem. It was

because the provider had other priorities rather than this specific problem. Hence, Supercomputer-Large's expert started solving the problem on his own.

At Supercomputer-Small, internal projects also exist, although they are not as frequent as Supercomputer-Large. Internal projects are limited to cases in which experts can solve a problem by writing a short script. Changes that are more fundamental are outsourced to their providers. Being part of the whole CERN network, Supercomputer-Small often suggests more basic improvements to be done centrally at CERN. Finally, the fact that Supercomputer-Small upgrades its systems frequently reduces the need for developing internal projects because most of the expected changes would be automatically addressed in the next version of technology.

5.3 Performing Upgrades

The third category of practices that the studied organizations were actively doing in relation with several incidents was performing upgrades. Upgrading problematic and outdated technologies is a common practice that the organizations do to avoid past incidents in future and reduce their likelihood. This practice is more visible in Supercomputer-Large and Supercomputer-Small, as they operate a wide range of different systems. The systems are often very specialized, modularly designed, and are supported by specialized providers.

For instance, at the Supercomputer-Large, there are several storage systems from different technology, supported by various providers. Normally, a specific expert handles systems with similar functionality (e.g., all backup systems) internally. Inside each system, there are modularly distinct parts. Therefore, replacing one part would not require changes in the other working parts of the system. Box 5.3

illustrates a typical replacement case at Supercomputers-Large. The example shows how the accumulation of several incidents makes Supercomputer-Large replace its faulty system to avoid similar incidents in future.

Box 5.3: Upgrading the Backup System at Supercomputer-Large

Supercomputer-Large has a backup system for long-term storage of huge data. This system is composed of an electro-mechanical robot (controller) that links supercomputers to the backup hard disks. When a user wants to store big data for several months or years, it should move the data from temporal storages in the supercomputers to the hard disks. Since the robot is electro-mechanical and has not been updated since 2006, there have been several major failures in the backup system. These incidents have occurred in different steps of copying data, including in reading a temporal copy of data to intermediary memory, moving the data to the disks, verifying the stored data, and removing the temporal copy. For example, in a serious incident, the controller failed to pick the proper disk. Then, it started overwriting data on a disk. Although the damaged data could be recovered from the backups, it took a lot of time and delayed the copying process.

In the last years, the speed at which data are entering the storage system has increased dramatically. At the time of the study, the prediction was that in a few months, the storage system would collapse since it cannot cope with the pace at which data should be stored.

The solution was to upgrade the backup system with a new one in which the memory size and the pace of storing data would be enhanced. The expert who is in charge of the storage system made the analysis of the problems of the existing system using all past incidents and did the prediction of the failure of the system due to slow data storage process. He submitted his proposal for upgrade to the head of System-Administration group. He approved the upgrade and sent the

request to the head of operations department for allocating budget and organizing the bid process.

At Supercomputer-Large, upgrading has done quite frequently in the isolated systems such as the cooling system, power transformers, network, and some peripheral systems such as security alarms. However, the core processing system, including all the processing tracks, internal storages, all wiring and network, and associated infrastructure, required to be upgraded in a whole project. This was almost the most complex and challenging project at Supercomputer-Large. Since the establishment of the first supercomputer in 2003, there has been one upgrade in the supercomputer in 2006.

The project requires changing almost everything related to the main supercomputer. Since the new technology is used in the new version, all the 250 applications that are used for processing the jobs have to be revised again to be adapted to the new technology. This requires testing the applications on the new system, fixing them, in case they have some incompatibilities, updating their users' guides, and installing them on the new system. The following upgrade was supposed to be done in 2008. However, regarding the huge cost of upgrading, the project was delayed to be done in 2010, and was postponed again until 2012.

However, the four year delay in the update created a long period of transition in which, the experts had to run internal improvement projects to fix the problems of the existing system in short term.

Regardless of the benefits of upgrades, large costs are sometimes very complicated and almost impossible to be justified at Supercomputer-Large, given the public nature of the organization and the bureaucratic process of budget allocation. The current national crisis has heightened this problem. As the result, I could find many internal improvement projects that would not be needed if upgrades could be done as quickly as expected. An example is developing a script

that allowed parallel usage of backup system for coping with the high speed of data generation that cannot be stored by the existing, slow backup systems. This issue would be solved when, next year, Supercomputer-Large buys the new storage system and replaces the existing one. However, the urgent need of users does not allow waiting until that time. This point was stressed by an expert in backup systems who had to develop a tool for faster storage of huge amount of data, before installing the new storage systems:

"Although next year we might upgrade this backup system, but I don't know when, I cannot wait for that. It depends on many things, such as the money from the ministry, the government, the politics. But what I have to do is to backup this data now. The new version of this backup system has a lot of capacity and will solve most of the problems, but there are many other uncertainties" [backup expert at Supercomputer-Large].

Compared with upgrades, which often require big investments, internal improvement projects do not often need any other resources except the time that internal experts spend on them. Given that the overall strategy of Supercomputer-Large is to use open source applications, there is often no need to buy any software license. Thus, in case of problems, always there is a discussion about to what extent the efforts should be channeled towards developing internal projects, versus pushing towards faster upgrades.

Finally, in the upgrade projects, the provider does most of the activities. The internal team needs to decide about when to do upgrade, negotiate for the budget, contact the provider, arrange for the upgrade, and help the provider's team make the changes.

Regarding the large number of systems and providers, especially in the Facility-Management group, even in cases that the problem might be simple and can be solved through some internal efforts, it is preferred to be done through upgrades.

This is first because the internal team does not have enough time to work on those improvements. Second, the reliability of new systems is often higher than the repaired systems. Third, the newly upgraded systems have new guarantee contracts, which assures enough support in case of incidents.

Supercomputer-Small has had much more frequent upgrades in its systems, compared with Supercomputer-Large. This is a simple and rather routine task:

*"For us, it is easy. We bring out the damaged track, and insert a new one from inventory" [*the head of infrastructure department at Supercomputer-Small].

The process of repairing the faulty systems is then a complex job that is done by the specialized providers. Most of upgrades involve changes in the hardware as well as some related software. In fact, many of the local improvements are automatically addressed in the next version of the technology, removing the need for short-term remedies or developing internal projects to fix them.

Supercomputer-Small has also executed frequent upgrades in its core facilities. Almost all processing and storage systems have been upgraded every 3 to 4 years. This is mostly due to energy consumption concern that is one of the most critical factors from the view of the host university. In many cases, the systems that are working properly are decommissioned because the cost of energy consumption would surpass the cost of upgrading them. In addition, Supercomputer-Small can benefit from higher performance of new systems. These points are illustrated in the following quote from the head of LHC department at Supercomputer-Small:

"We retire them every four years. … At some points, essentially after four years, it is better to … unplug all stuff. The new technologies are smaller and … in the storage, the kilowatt per tera-byte follows like an exponentially growing and then after 3 or 4 years, typically 4 years, it is so much better than … a 4-year machine, you are essentially paying everything for electricity.… the optimal decision in terms

of spending money is to replace the old machine with more tense machine" [the head of LHC at Supercomputer-Small].

Supercomputer-Small has not been dependent on a specific supplier. There is a wide range of technologies, even for a specific system such as storage. This has helped Supercomputer-Small benefit from the competition among its providers. In the supercomputer sector, the upgrade of core facilities is a huge contract. Therefore, the providers are willing to consider their clients' needs to increase their chance in winning the next tender. In addition, using a wide range of technologies has made Supercomputer-Small capable enough to operate various technologies and systems and make them work with each other. Therefore, Supercomputer-Small has learned to manage the interfaces between various systems when each of the systems is upgrading independently. The downside, however, is that this adds to the complexity that Supercomputer-Small has to manage in each case of upgrading.

The rapid and frequent upgrading has reduced the need for many internal projects at Supercomputer-Small. There were cases that internal experts developed some temporary solutions, just until the next release of the system. A middle technical manager at Supercomputer-Small explained this point:

"For example, now, we have a hardware that has some problems and it is in the production and we know that we should have a new hardware in one or two months. … then we don't take follow-up actions … we could spend a lot of time there trying to tune the rate or whatever. But we don't [do] that. We know that we can do it in the infinite manpower scenario (laughing). We would learn something about it. But in two months it would be unplugged" [a middle technical manager at Supercomputer-Small].

Being part of the whole CERN network, there are central software providers that continuously receive the change requests and release their new versions of

applications almost twice a year. Supercomputer-Small takes part in monthly meetings with the experts from other tier 1st sites, to discuss the new change requests with application providers. A recent example illustrates this point. In various incidents, some parts of the stored data were prone to be damaged. Apart from recovering the damaged data, which often takes several days or weeks, Supercomputer-Small needed to know very quickly (in less than one hour) which files are damaged and report it to CERN. This way, the whole network of CERN can quickly see if there is any backup of the damaged file. There is an overall application that managers the data storage in the whole CERN network, called d'cache. This need, once urgently detected in an incident inside Supercomputer-Small, made the internal team develop quickly a simple script in short-run to locate where the damaged files are located. In parallel, this point, which has been also reported by other sites, was added to the change requests in the d'cache system to be done by the specialized provider:

Sometimes during the upgrading process, new incidents occur. This is mostly due to the increasing complexity of systems that are connected. Thus, recently, Supercomputer-Small has recognized the need for developing a systematic process for upgrades, in which specific steps should be taken for discontinuing the existing system, as well as installing the new one. The process is complex, because the newly installed system is prone to some initial incidents due to inappropriate installation or configuration. Thus, the existing systems are often kept for a while, as a backup, to allow rolling back the upgrade in case of incidents in the new systems.

Chapter summary: The chapter described various practices that the organizations adopted in relation with multiple incidents. These practices are adding to the practices that were specifically taken for each specific incident (previous chapter).

The multi-incident practices were described under three categories: using Wiki systems for documenting the experience of several incidents and using it as a support for future incident handling activities, executing and implementing improvement projects that address some defects that have been relent in several incidents, and making upgrades in the faulty systems that has created or can potentially create future incidents. The next chapter focuses on analyzing the learning practices that have been taken during and after each incident, as well as the practices taken in relation with several incidents.

Chapter 6: Analysis: Beyond Traditional Learning Modes

Chapter 4 and 5 described various practices that helped the organizations capitalize on their experience of ISRIs to avoid them in future or manage them more effectively. In chapter 4, I described the practices that were adopted in relation with single incidents, either during incident handling process (handling incidents through ticketing system, doing Triage, and interactions with specialized suppliers) or after handling the incident (post-incident reflection and incident laboratory analysis). Chapter 5 described various practices that the organizations

adopted in relation to several incidents (using Wiki system, executing improvement projects, and upgrading the systems and technologies).

This chapter articulates these practices into five patterns that I call them learning *modes*: (1) Learning through incident handling practice; (2) Learning through post-incident reflection; (3) Transversal Learning from Incidents; (4) Outsourced Learning; (5) Learning through material replacement. Each learning mode shows a specific pattern of learning practices that is based on the abstraction from one or several categories of learning practices.

More specifically, learning through incident handling practices refers to the process through which the very engagement in incident handling provides detailed experience of the incident, its causes, and potential ways for preventing it. This learning mode is manifested in handling incidents through ticketing systems, performing Triage, and doing laboratory analysis when the experts engage with incidents and work on them.

Learning through post-incident reflection is a learning mode in which the actors reflect on their incident experience and draw lessons that can help them avoid future incidents or manage them effectively. Learning through reflection pertains to post-incident analysis practices and developing some Wiki articles that are specific to single incidents.

Transversal learning mode is a pattern in which the organizations run an improvement project that straddles several incidents, but often takes place long after them. Transversal learning mode is observed in internal improvement projects, developing most of Wiki articles, and upgrades.

Outsourced learning mode describes the pattern in which the specialized providers perform the very learning practices, while the focal organization is still benefiting from the learning outcomes. Outsourced learning is evident when the

organizations outsource their incident handling processes, post-incident reviews, and their improvement projects.

Finally, learning through material replacement refers to a common pattern in the organizations that they tried to avoid future incidents or reduce their impacts by replacing the faulty part of their systems. This learning mode can be observed when the organizations replaced their IT artifacts during incident handling process, through their improvement projects, and in their scheduled upgrades.

The first two modes are well documented in the literature. However, I highlight some nuances that show how the characteristics of ISRIs affected these two modes in the studied organizations (sections 6.1 and 6.2). The rest of the chapter (section 6.3, 6.4, and 6.5) articulates the other three learning modes. I also analyze how the characteristics of ISRIs contributed to the emergence of these modes and the way they unfolded.

6.1 Learning through incident handling practices

A first learning mode, emerging from this study, rotates around the actual incident handling practices. This learning mode means that *the engagement of individuals and collectives in the incident handling process provides them with experiences and capabilities that help them in handling future incidents.* The experience gained through handling incidents is important for handling future incidents. A senior expert at Security-Private mentioned:

"Of course, you get a lot experience of security threats, the way that hackers implement their strategies, various sources of vulnerabilities just as you handle these incidents" [A senior expert at Security-Private].

This learning mode manifests in three specific practices: handling tickets, doing Triage, and laboratory analysis of incidents. In this section, I comment on five aspects of this learning mode that shows how the context of ISRIs can contribute to it. First, I show how the organizations neutralized incidents when it entered into the ticketing system. Second, I comment on Triage as an integrative learning step at the beginning of incident handling process. Third, I comment on the high level of specialization that predominates incident handling process and how the escalation of incidents and selective involvement of actors influence learning opportunities throughout incident handling process. Fourth, I analyze how ticket is serving as an individual, shared memory during incident handling process. Fifth, I explain how the organizations were relying on the ticketing system as an automated, complete memory of incident.

6.1.1 Neutralizing incidents through ticketing system

Incidents have negative content that might reduce the willingness of experts to dig into them for learning from them (Baumard & Starbuck, 2005; Stern, 2002). However, in the studied cases, this negative content was hardly observable. It was rather surprising that in the interviews, the managers and experts reacted to the question "what kind of incidents you have here", by saying

"We don't call them 'incidents'. We call them 'tickets'" [e.g., the head of Users-Support groups at Supercomputer-Large]

Further exploration showed that this is relevant for major incidents as well. A senior expert at Security-Private commented on how the incident is named "ticket" and is perceived as a neutral, urgent task, after entering the organization:

"From the time that they enter our ticketing system, they are "tickets"" [A senior expert at Security-Private].

Tickets then become *important tasks* that should be done urgently. This way, incidents become *neutralized* in the incident handling process. It means that the negative content of Incident –i.e., "what was wrong", "who did the wrong task", "whose fault is this"- is removed by being transformed into a ticket that a specialized experts has to handle it. However, tickets still maintain the urgency and importance of the incident –i.e., "why it is critical to act on it", "how important is to control it quickly and avoid its occurrence in future".

Tickets play an important role in this process. First, it changes the language that the organizations use instead of incident. It provides a neutral substitute for notions such as "problem", "incident", and "failure". Often, those who have been affected by the incident (e.g., the users who could not get access to supercomputers) express these negative terms. However, inside the organizations, the notion of ticket replaces all these words. In addition, ticket, when it is open, has the connotation of "*must be handled urgently".* In fact, tickets are different from normal tasks (e.g., preparing a report for a manager), because they represent urgent needs.

The literature has shown that organizational actors are likely to ignore or overlook the incidents that have negative connotations (normalizing phenomenon (Perrow, 1999)). However, the ticket is acting in an opposite way. The moment that a ticket is created for an incident, it gives the incident an objective presence ("an open ticket") that demands attention. Open tickets scream that "please admit me; I am an urgent task that has not been handled yet". The tickets cannot be closed unless the head of the department or the one who is responsible for customer relations approves it. This setting reduces the chance that incidents be filtered out when they enter the organization or be overlooked beneath the daily activities.

Neutralizing allows experts to talk freely about incidents and interact openly during and after handling them. In addition, the attribution of tickets to an expert (e.g., “this is your ticket for handling”) does not foist any negative connotation on the expert to be considered as guilty or so. Instead, it can even be a badge of recognition, since the expert is considered capable enough to handle the incident.

6.1.2 Integrative learning through Triage

Triage is a specific stage in incident handling process that contributes to learning through incident handling practices. Triage involves practices that have unique learning aspects. First, learning during Triage has an *integrative* orientation, because the attentions and efforts are directed towards the current incident at hand. In other words, the actors draw upon their past experiences that can help them in managing the specific underlying incident. Past experiences and newly gathered information are all integrated to set a proper framework for handling the current incident. At the same time they take into account what happened in the past in terms of any similar incident. They think about possible links between the current incident and the previous ones. They consider if any experience in the past can help understanding and managing better the incident. In doing so, they put together various knowledge elements coming from their own experience and other sources.

By the same token, Triage highlights the importance of having a *holistic* view of the incident, its scope, its potential impacts, and the patterns over similar incidents. This is reflected in the aim of Triage, which is setting the overall framework for handling the incident.

Triage also brings to the fore the importance of *heuristic learning*, meaning that actors need to rely on their tentative and intuitive knowledge of the existing case,

as still little information is available about the incident in this stage (compared with after-incident reflection, when more stable and detailed information about the incident is available).

Triage makes the incident handling process different from what is normally perceived; that is actors rush into the incident to control it without having time to reflect on it (the firefighting metaphor). However, Triage, as a formal stage, is an *institutionalized reflection time in the midst of disaster*. As a middle manager in Security-Private mentioned, it is a time that they put aside the pressure of incident handling process, and they say

"OK, wait a minute, and let's think about it, before we take any action …" [a middle manager in Security-Private].

This helps the experts become free of the pressure and urgency of incident context that can prevent them from deep, comprehensive understanding of the incident.

Compared with learning that individuals get through their practical engagement in incident handling, Triage is a *collective* process. As illustrated in Box 4.2, Triage is done among senior experts who are responsible for different aspects of the business (e.g., different technologies or different business activities). The judgments about the importance and scope of incidents are the key themes of Triage that is discussed from the view of the senior experts. In addition, taking decisions about how incident-handling process should be structured, who is responsible for what, and what are the potential side effects, are all collective decisions that these experts make together. This is because various opinions from technical, organizational, and legal points of view are discussed for evaluating the criticality of the incident. In addition, setting the overall framework for handling incident requires that actors from various domains work together. The fact that

actors from different domains are involved in this stage provides many opportunities for learning about other domains.

6.1.3 Specialization, escalation, and selective involvement

In learning through incident handling practices, exerts learn through their engagements with various incidents. Particularly for junior experts, this practical engagement is critical because it complements what they had learned about IT incidents through formal trainings (e.g., university courses), and make them apply their knowledge in the real situation of incident handling process, where a lot of time pressure, ambiguity, and sensitivity in terms of potential mistakes exist.

Senior experts with ample experience in incident handling mostly learn by dealing with *novel* incidents that allow them to learn new themes. As the head of incident handling department at Security-Public mentioned,

"In case of new types of security incidents I personally take part. I try to work, in these cases with our provider, because these cases might have sensitive issues that should be handled with specific care. At the same time, I can learn new things" [the head of incident handling department at Security-Public].

Several mechanisms helped senior experts detect and involve merely in novel incidents. First, according to the *escalation* process in the incident handling process, once an incident arrives at the organizations, it is taken by first-level operators. If the incident is not routine (in the list of normal incidents with simple responses), it is escalated through ticketing system to the second-level experts.

In specific cases that incident is complex or has to do with sensitive decisions that are beyond the knowledge and responsibility of second-level experts, it is escalated to third-level experts. This situation provides opportunities for junior

experts who have just promoted from level 1 to level 2 to learn through engaging with specialized incidents.

In addition, the senior experts are also "*aware of all incidents and what is going on during handling them*" [the head of System Administration Group at Supercomputer-Large] through ticketing system and daily interactions. However, they only intervened when they feel that "*something might go wrong*", mentioned by a senior expert at Supercomputer-Large. In fact, organizations saw the engagement of their junior experts in various incidents as a key mechanism through which they "*develop experts who are capable to handle the very specific types of incidents*" [the head of Users-Support group at Supercomputer-Large]. This, in turn, adds to the overall capacity of the organizations in dealing with similar incidents in future. It also filters out routine incidents that can yield little learning at the second and third levels of experts. A senior expert at Supercomputer-Large mentioned that

"Each day, we might receive tens of incidents. I do not go through all of them. Just those that are critical and technical enough that cannot be handled by our helpdesk team, is sent to me" [a senior expert at Supercomputer-Large].

This selective intervention, not only assures appropriate incident handling, provides junior experts with learning opportunities through their interactions with senior experts. For instance, they can learn what is often called "*tricks*". Tricks are specific solutions or techniques that help solving puzzling problems, in a rather strange way. For instance, the Email-server expert at Supercomputer-Large faced a situation that some emails were strangely considered as Spam. He checked them and resent them from different senders to different recipients to find the problem. It turned out that the problem was not related to the sender, or the receiver. Even the problem was not about the subject of email or its attachments. The problem was very strange, since there was no reason that such emails would be considered

as Spam. Accidentally, the expert changed the first character of the email content. He put a space character at the beginning of the text and the problem was solved. He did several tests and realized that the problem was because the email text started with “LM”. He did not know why that happened and how putting a space at the beginning solves the issue. However, he was referring to this as an example of many other tricks that just happen in their work and cannot be learned from reading formal books. He reflected on this point saying

“I do not know why it works (laughing). But it works. It is funny. That happens. You just put a space and you see the problem is solved” [email-server expert at Supercomputer-Large]

These tricks, though are very specific and sometimes just relevant in very special situations, are critical because junior experts cannot learn them through other mechanisms (such as formal training or reading available information).

Second, and in addition to this formal escalation process, senior experts detect novel incidents as they are constantly checking out the in-flow of incidents. This selective attention is often backed by the experience of senior experts that provides them with clues for detecting novel incidents. In doing so, experts often look for *counter-intuitive* and *strange* incidents that look unfamiliar, with regards to their experience of past incidents. The head of System-Administration group at Supercomputer-Large said

“… normally I am on top of the tickets. So, I know more or less the issues that are happening and normally these people are reporting me … problem. That we have to solve it and we have to find alternative solutions” [The head of System-Administration group at Supercomputer-Large].

This shows that the intuition of the senior expert is helping him to detect strange incidents.

Third, the *specialization* of experts in specific categories of incidents also works as a mechanism for selecting *relevant* incidents to engage with. Within the domain of IT incidents, all organizations (except Supercomputer-Small) had a detailed specialization at the second-level experts. This specialization, in some cases, was based on historical preferences. For example, in Supercomputer-Large, in Users-Supports group, four second-level experts are specialized in handling incidents related to different applications. Once expert "A" handles the first incident related to a specific application, then the next incidents related to this application will be assigned to the same person to leverage his experience. Similarly, in System-Administration group at Supercomputer-Large, the team of senior experts is specialized in handling incidents related to network, security, backup, CPUs, and servers. Although in most of the cases experts can handle incidents outside their domain of expertise, they intentionally avoid that because it is not efficient.

"It is not efficient to do so. Other colleagues can handle these network incidents, maybe if they spend like 10 hours to understand it, but for me it is a matter of 2 hours" [network expert at Supercomputer-Large].

Over time, this gradually created clear domains of specialization among the four experts. This phenomenon is known as transactive memory system (Brandon & Hollingshead, 2004; Liang, Moreland, & Argote, 1995) that allows some actors rely on the knowledge of other actors without needing to learn about others' domains.

The specialization of incident handling prevents the organizations from spending further resources to re-learn what some specific experts have already learned. Instead, experts could spare time to learn about incidents that are really new to the organization. Regarding the fact that even in large organizations often one expert is assigned to a specialized domain, most of the learning that they had through their engagement in handling incidents is individual.

6.1.4 Ticket as an individual, shared memory

During incident handling process, tickets keep the record of all the actions (including both effective and failed actions) that have been taken for handling incident. This way, experts can learn from the actions applied to the current incident, when they are still handling the same incident. This is important when a group of experts (often first-level helpdesk) fail to handle the incident, so that, they move the ticket to more specialized experts (often specialized experts at level 2 and 3). In fact, the ticket is acting as a *temporal, live memory* that connects different actors. For instance, the ticket that was created for the Email-server incident at Supercomputer-Large (please see box 4.1), first, was filled with the comments of the users support experts. The same comments were then consulted by the expert responsible for the Email-servers, while still the ticket was open. This expert changed part of the comment of the Users-Support to correct the description of the ticket in a way that shows it is not due to the Email-servers. Then until the point that the network expert was resolving the incidents, the ticket was used as a memory of what has been done so far.

In addition, ticket acts as a historically contextual memory of handled incidents because it reflects the context in which the incident took place, the very actions that were applied on it, and the consequences of those actions. The chronological sequence of the ticket content acts as a story that reminds the situation in which incident handling process took place.

Tickets related to major incidents are often quite specialized in narrow domains. Therefore, although tickets are shared, they mostly act as *occasional, individual memories.* It means that each expert refers to his own past tickets occasionally to refresh his/her mind about a specific piece of information. The very attempt of putting detailed descriptions of actions and solutions facilitates retrieving required

information in future similar cases. However, this is often done with the assumption that the same person who handled it will consult the ticket.

6.1.5 Ticketing system as an automated, complete memory

Ticketing system provides an infrastructure that facilitates learning during incident handling process. The way tickets are designed and, thus accumulatively store information during incident handling (e.g., with the automatic logs), helps organizations keep a detailed, retrievable record of incidents that does not require extra efforts. In fact, experts do not need to stop their very incident handling practices to document some information into tickets. Unlike Wiki, actions on the ticket are perceived as the very incident handling practices, not as extra-work burdens. Although this might make tickets less structured and articulated than Wiki articles, it is aligned with the context of incident handling process in which actors are of short of time.

Another important characteristic of ticketing system is that it keeps a rather *complete* record of the incident handling process. This completeness has been quite important for the organizations as almost managers in all the organizations mentioned this point as one of their rationales for using ticketing system. For instance, the head of LHC at Supercomputer-Small mentioned that

"What is important is that we have all *[emphasis in the original voice] information of an incident in its ticket"* [the head of LHC at Supercomputer-Small].

It is striking that, this completeness is rarely exploited in practice. For instance, organizations rarely performed comprehensive search and analysis on the information stored in their ticketing systems. It seems that this completeness is mainly concerned for the sake of *peace of mind.* In other words, a complete memory of incidents serves the managers psychologically, rather than cognitively.

This point is reflected in the answer of the same manager at Supercomputer-Small in responding to the question about how they benefit from this completeness:

"In case that one day we need some specific information, we are sure that we have everything in the ticketing system" [the head of LHC at Supercomputer-Small].

In fact, being complete, is a subjective judgment that the managers of the organizations pursue. There is no specific measure that shows whether the content of tickets is complete or not. Even there is no reference point in the organizations against which the completeness can be assessed. In other words, it reflects more the *fear of lacking important information,* rather than a specific need that drives their zeal to store everything possible.

To sum up, various characteristics of ISRIs affect the way organizations learning through handling incidents. The technical nature of incidents and their high level of specialization make this process be predominately individual. The critical role that ticketing system plays in this process also allows neutralizing incidents into tickets, which in turn allows for storing a detailed, historical log about each incident. In addition, the specialization of incidents due to their technical natures, joint with the escalation process, helps allocating learning opportunities to learners, without compromising the quality of delivering incident handling services (due to selective intervention of senior experts).

6.2 Learning through post-incident reflection

A second group of practices, common to all the organizations, rotated around the efforts to reflect on and analyze major incidents after handling them. The studied

organizations performed various practices through which they *analyzed and reflected on handled incidents to draw lessons for future incidents and improvements*.

This learning mode is clearly reflected in both informal and formal post-incident reflections such as incident reporting and open seminars. In addition, some Wiki articles that were specifically developed for single incidents also manifest this mode.

This learning mode is the dominant view in the literature about how organizations learn from their incidents. However, three nuances in the way that organizations adopted this learning mode unravel some of the characteristics of ISRIs. First, the organizations intentionally avoided doing post-incident reflections in various situations. Second, although post-incident reflection is supposed to be done collectively, it has been conducted individually or in small teams. Third, in spite of the original intention of post-incident reflection for extracting the causes of incidents (know-why), the focus of organizations has been mostly on identifying practical solutions (know-how). In fact, the organizations were skipping the root-cause analysis of this learning mode. I will go through these three nuances and discuss how they can be related to the contextual factors of ISRIs.

6.2.1 Learning abstinence

The studied organizations did not perform post-incident reflection and analysis on all major incidents. They were quite selective in doing so. Even in organizations like Security-Private with a systematic post-incident analysis stage, the experts were careful in deciding whether to analyze an incident or not. The experts were mentioning that not all incidents are worth being analyzed. One major reason was

that technologies change rapidly. Therefore, the next incidents will be different. Hence, lessons from past incidents are not applicable in case of future incidents.

In this sense, learning from ISRIs presents unique challenges that are different from other industries. In fact, the period through which the lessons learned from past incidents can be used is too short that does not justify spending time and effort by the organizations. Sometimes, even the learning process (with its delays and challenges) can take longer than the period that such learning might be relevant. This point was clear for instance, in the case of backup systems at Supercomputer-Large that the organization was not entering into the analysis of its incidents since the technology would change very soon. The new backup systems are based on a fundamentally different architecture and design, that even thinking about the causes of past incidents could mislead experts about how the new system works.

As the second reason, which was more relevant to novel incidents, the experts were worried about drawing immature lessons, before they see a rather stable pattern. Even in situations that the technology is not changing fast, the experts resisted systematic analysis of incidents causes until a good number of incidents are accumulating. They had to see some overall patterns over time to see what aspect of the technology might be the source of incident, and consequently find the appropriate intervention point. Their intuition was indicating that too early analysis might make them draw incomplete lessons and erroneous conclusions. In addition, they were worried that this might lead to suggesting and sustaining some remedies that prevent them from finding solutions that are more fundamental.

6.2.2 Collectively designed, but individually performed learning

Although post-incident reflection and analysis were designed to be done collectively, senior experts often conducted them individually. For instance, mainly the head of LHC department did incident reporting at Supercomputer-Small. Similarly, single, specialized experts who handled the incident mostly did post-incident documentations and analysis at Supercomputer-Large. This is because major incidents are often quite specialized. Therefore, the expert who is specialized in that domain handles them. In addition, the fact that IT systems are modular makes incidents be technically separated. For instance, although a network incident at Supercomputer-Large affected Email-servers, the lessons drawn from the former incident was quite unrelated to the latter. This was manifested in the case narrated in Box 4.1 in which the Email-servers expert mentioned that

"I do not know exactly what was the issue. I did not ask about it, to be honest (laughing). It was not related to Email-servers. It was something about network, I guess" [Email-servers expert at Supercomputer-Large].

Likewise, open seminars at Security-Private, that were designed to be done collectively, faced a problem that a senior expert at the company referred to as the problem of *"being boring and unrelated to the rest of the team*".

As we see, over time, collectively designed practices such as open seminars and post-incident reflections gradually changed into a rather individual task. Even their formality (e.g., all projects must document their lessons) did not help in performing them collectively. It has been observed in most of the organizations that a specific expert is almost doing everything that is needed for creating such document. Even in cases such as Supercomputer-Small, formal post-incident

reports are developed with the centrality of a principle expert who occasionally consulted other experts for specific information.

The core issue is the *relevance* of the analysis to the organizational actors. The fact that most of the incidents that are important and novel enough to be worth being analyzed, are so specialized that either one or just very few experts relate to them. Even in cases such as Security-Private that the organizations have prepared all sorts of motivational and resources for holding such analysis sessions collectively, the lack of relevance resulted in rather useless participation of expert.

6.2.3 Shifting the attention from "know-why" to "know-how"

Post-incident analyses were mostly focused on practical implications, useful tools and methods, and, specific tricks that can help in future. For instance, the Wiki articles at Supercomputer-Large that were developed as reflections on specific incidents were mostly focused on what they call it "how-to-do" content. These articles were documenting a specific tool or technique that can solve a specific problem. As illustrated in the previous chapters, the content of wiki articles focuses on technical know-how, such as step-by-step instructions for running a script and modifying it.

Similarly, the incident reports at Supercomputer-Small turned into shorter reports that focus less on the detailed causes of the incident. The part related to the causes of incidents shrank over time, giving more centrality to the description of the incident and the actions done to handle it and follow up actions that should be taken afterwards.

There were several reasons for this pattern. First, the experts mentioned that the audiences of such documents are either themselves or experts in their narrow domain of expertise who "all know what is behind". Therefore, for the experts it

was not making much sense to document the causes of incidents that they have already handled. As mentioned by a senior expert at Supercomputer-Large, there are two scenarios. Either the expert has handled the incident. Therefore, he already knows everything about the incident. Then, he just needs to put some practical points that in the article that in case he forgot something, he can consult the article. Or, the expert is not involved in the incident and the incident is handled by the provider. In that situation, the expert does not need to know the causes of the incident. Perhaps, he only needs to know a little about how to control the incident in the first place and hand over the incident to the provider.

Second, the fact that technical aspects behind incidents change frequently (e.g., Internet browsers change, so that the ways in which hackers enter computers does) reduces the need for detailed documentation of their causes. In fact, it is not efficient in many situations that the experts dig into the incident and spend a considerable time to detect and document incident causes, when those causes are not going to be relevant in future incidents.

To summarize, rapid changes in information technologies made it unnecessary that organizations perform post-incident reflection. Furthermore, the organizations sometimes strived to avoid learning practices. Even in cases that post-incident reflection was conducted, it has been done mainly by individual experts or among limited number of experts who are related to the specialized incident. Finally, in documenting the post-incident analysis, the efforts has been mostly focused on documenting know-how, with little attention to extracting and codifying know-why. This is because either the experts already know and do not need to document incident causes for themselves or the technology changes so quickly that makes such knowledge irrelevant for the next incidents.

6.3 Transversal learning from ISRIs

The previous two groups of practices (learning through handling incidents and learning through reflection) resemble processes that have been documented by prior studies as typical learning modes that take place in various contexts. However, this study foregrounds a third leaning mode and associated practices that is typical of the IS (and other technologically dense) environments. This has to do with learning not from individual, but rather several incidents, that is called transversal learning (TL). Transversal learning (TL) refers *to the process through which the organization tries to solve a specific problem or make a specific improvement based on the experience of several related incidents*.

Improvement projects, which almost all organizations were running for making local improvements in their systems, are clear examples of TL. For instance, when Supercomputer-Large launched the central storage mirroring project (see Box 5.2 in chapter 5), the organization tried to address three categories of past incidents.

Other examples of transversal learning can also be found in some of the wiki articles at Security-Public. For instance, the incident response team created an article about considerations that should be taken for isolating the affected hard disk of Apple Laptops. This article, was developed when the experts faced several incidents in which Apple Laptops were affected and they got the experience of some special considerations related to isolating their hard-disk.

There are several characteristics that all together define TL mode: (1) Multiple relation with past incidents; (2) Focusing on specific solution (thematic nature), (3) Asymmetric temporal relation with incidents; (4) Being based on a temporal, purposeful cognitive base; (5) Being based on accumulated political pressure; (6) Being materialized through socio-technical bases. Table 6.1 summarizes the characteristics of TL.

Table 6.1: The characteristics of transversal learning

Characteristics	Description
Multi-incident relation	The learning process straddles multiple incidents
Thematic with specific purpose	TL is focused on a specific solution, which often relates to a theme that straddles a wide range of incidents
Asymmetric temporal relation with incidents	The timings of incidents and learning process are not synchronized; learning talks place either in a periodic or in ad-hoc manner; the temporal distance between incident and learning process is significant
Loose and vague relation with incidents	The relation between the learning process and incidents is loose and vague; it is hard to pin point a specific and clear relation between TL and related incidents
Articulated around socio-material elements	In TL, there are socio-material elements that integrate various learning practices and incorporate learning solutions into specific socio-material elements.

6.3.1 Multiple relation with past incidents

TL, as exemplified in these two cases, is informed by the experience of *several incidents*. In fact, the relation between the incident and learning process is not one-to-one. The learning process is often based on several incidents that share some similarities. Those similarities are selected by learning agents depending on the purpose of learning (learning theme). For instance, numerous incidents might be selected to be analyzed, by focusing on their similarities in terms of the legal issues that they raise, while on another occasion, based on their technical similarities.

The past incidents have been discussed in a rather vague and abstract way in the learning sessions to disguise the relations with specific incidents. This rather ambiguous reference to incidents helped the organizations avoid political tensions in the learning process. In fact, organizational actors were intentionally *de-politicizing* learning process by blurring the relations between learning process and the related incidents.

Although TL occurs in parallel with incident handling process, it often requires a different mood of action in which actors free themselves from the pressure of incident handling, to see “the big picture”, take a more “long-term” perspective, and to see “whether it is worth doing” or not. In some companies, the experts were taking turn for handling incidents, so that they could spare “relaxed time” for spending on learning projects.

6.3.2 Focusing on specific solution (thematic nature)

TL is articulated around a specific solution (a *theme*). For instance, adding a legal service for assessing the legal implications of security incidents at Security-Public was informed by the experience of various past incidents with significant legal implications. The focal purpose of TL determines which past incidents are relevant to be drawn upon in the learning process, and more specifically, which aspects of them should be taken into consideration. In addition, this central theme determines relevant learning practices that are needed to be conducted in the learning process.

6.3.3 Asymmetric temporal relation with incidents

Another characteristic of TL mode is its timing. TL is often not triggered by any of the related incidents. TL often takes place long after the incidents are handled, when the political climate of incident calmed down. Sometimes, attempts were made to add more *temporal distance* between the incidents and the learning process to make it hard to pin point a specific relation between learning practices and the related incidents.

This temporal distance was particularly important in cases that technology was changing quickly, so that the organizations could bypass temporal changes to reach to a stable pattern of incidents for defining their learning initiatives. For instance, Security-Public handled the first sever mobile security incident by using its general tools. Once a stable pattern of mobile security incidents was recognized, the managers and experts realized that specialized incident handling tools are needed for effective management of mobile incidents. This led them to acquire new set of technologies for handling mobile security incident.

Instead, TL takes place when a proper *moment of learning* appears. The organizations transversally learned, for example, when they were upgrading their systems and technologies, in periodic events, (such as renewing contracts), when they were designing new systems, during scheduled interventions for doing overhaul or preventive changes, when the daily load work plummeted (e.g. during summer time), when required resources (financial, technological, or human) became available, and simply when key influential actors were present to be involved in the process. For instance, Security-Public has used its annual planning as a critical time to revise its service catalogue by adding new services related to new incidents detected in the last period. It also helped Security-Public make changes in the processes through which the services are delivered.

6.3.4 Being based on a temporal, purposeful cognitive base

TL mode benefits from the experiences that organizations accrue during incident handling process and in their reflections after each incident. However, it is mostly based on the knowledge that is specifically developed and integrated for the purpose of the learning project. In fact, the actors focus on a specific targeted solution (e.g., the script that creates alarms before the failure of storage system).

They often start developing their knowledge exclusively for creating the solution. This, for instance, can be done by assisting some specific practical workshops (e.g., knowing how the storage system might fail), collecting some new information (e.g., searching in the web to see how other similar organizations have solved this problem), and conducting some specific experiments (e.g., reproducing the failure of the storage system). This results in a body of knowledge that has been developed purposefully for the underlying learning project.

In doing so, the organizations also needed to draw upon their experiences from past, related incidents. For each transversal learning solution, often, there have been several related incidents. These incidents are sometimes very different in terms of their origin, impacts, and handling process. However, each adds to the local knowledge base that is needed for developing the intended solution. For instance, in the central storage case in Supercomputer-Large (see Box 5.2), the experience of three different types of incidents were used: failure in storing data files, failure in accessing to the existing files, and failure in copying data with sufficient pace.

Unlike in post-incident reflection, in creating the local, temporal knowledge base, the organizations were focusing on a central theme (the solution) and this allowed them to identify what part of their experience was relevant. This temporal judgment helped them *filter out the noises* and focus on what they were seeing relevant for the sake of the current learning project. Unlike learning through incident handling and post-incident reflection, the relevance of past incidents derives from their potential contribution to complete the local knowledge base that is needed for developing the solution, not necessarily due to their novelty per se.

In developing the local cognitive base, the organizations adopted an *integrative* mode of knowing in which the search for new knowledge was similar to

completing a puzzle. Therefore, it was different from the type of discussions that are mostly present in post-incident reflection sessions that experts brainstorm on what ideas can be drawn *from* the incident. Instead, the question, in TL, is what is needed to be known *for* developing the solution.

6.3.5 Being based on accumulated political pressure

Transversal learning mode is also based on mechanisms that try to accumulate political pressure to drive and lead the learning process. Since TL takes place long after the related incidents, the sense of urgency and pressure that exists during and right after incidents fades away. Instead, TL mode is based on the accumulation of pressure that has been accrued out of several related incidents. In running TL process, the organizations were building on the past incidents by reopening them and upholding the needs for change that those incidents were showing. This way, they could create a temporal political pressure that could drive the learning project.

Furthermore, the experts who were leading TL process were playing on the knowledge lack that some influential actors such as top managers or external customers were suffering from. Hence, they could *dramatize* further by drawing upon a wide range of past incidents, as well as potential future damages. This way, a series of incidents were used to construct a discourse that could increase the political pressure to a level that drives the learning process. The pain that actors could reconstruct and envision based on past experiences and potential future incidents is an important aspect of incidents that contributes to dramatizing.

6.3.6 Being materialized through socio-technical bases

TL mode is solution oriented. Therefore, it is often materialized around the development and refinement of a specific socio-material element such as a specific hardware, a piece of software, or simply a specific technical standard and an organizational procedure. This is the objective facet of TL mode. The socio-material elements sometimes are created through the learning process (e.g., developing a script for increasing the pace of mirroring data) and, sometimes, are used as the basis upon which learning practices are applied (e.g., adding a script to the backup operating system). This way, socio-material elements act as a basis for integrating learning practices. At the same time, they incorporate part of the knowledge and experience that has been used for developing them. They also institutionalize the learning solution into existing organizational practices.

To summarize, TL mode refers to a pattern of learning that is based on several incidents, is focused on a specific intended solution, and often takes place with considerable temporal distance from the related incidents. As shown in lower part of Figure 6.1, bellow, TL mode is based on the experience of several incidents. In fact, the experience that individuals and collectives gain through incident handling process and through post-incident reflection adds to the organizational memory. However, for the specific intended learning solution, a local cognitive base is created out of the overall organizational memory (see the middle part of Figure 6.1). In doing so, the organizations relied on various mechanisms such as filtering the noises, which means selecting parts of organizational memory that are related to the underlying solution. Benefiting from several incidents' experience, the organizations could identify some patterns across various incidents and gain deeper understanding of some problems that existed in several incidents. This experience helped the organizations gain a holistic picture about the incident, its cause, and the potential solutions. In addition, various knowledge domains that

are related to the past incidents are often integrated to build the local knowledge base for the development of the solution.

In addition, TL is based on the accumulation of a local political pressure to drive the learning process (right side of Figure 6.1). The organizations rely on various mechanisms such as dramatizing on several past incidents to add to the perceived urgency and need for learning. In addition, since TL is related to several incidents, sometimes the organizations disguised the specific relation with particular incidents to avoid some political interference such as blame-game that can be created by reflecting on specific incidents. In addition, the support of some key actors also triggered some learning projects.

TL is also based on creating a socio-material bases upon which the intended solution is developed. This socio-material base sometimes has a scaffolding role (e.g., a temporal model or prototype of the final solution) for the creation of the ultimate solution, while in other occasions might become part of the final solution (e.g., an initial versions of a script). Creating a socio-material base requires mobilizing various resources, deploying material elements, and developing some organizational structures (such as a new routine or a project) upon which the learning solution is built.

The final solution in the TL process is then the product of these three pillars (cognitive, political, and socio-material) that can be realized in an appropriate learning moment.

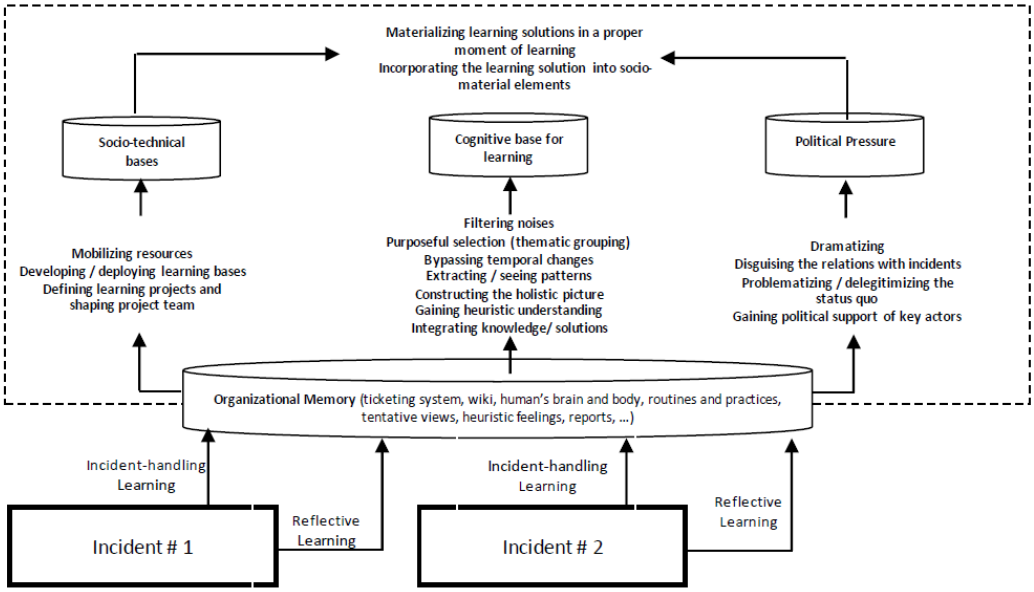

Figure 6.1: Transversal learning and its relation with learning through incident handling and post-incident reflection

Some recent theoretical developments in the literature on learning from incidents emphasize on the importance of inter-incident learning (Moynihan, 2009). Some patterns about the cause of incidents can be detected when organizations reflect on several incidents over time. These studies mainly focus on the cognitive aspects of learning from multi-incident analysis.

Although TL relies cognitively on multiple incidents (as various sources of information), it also benefits from other structural and political mechanisms that are based on a series of past incidents. As discussed above, multiple incidents accumulate enough political pressure to justify and lead the learning process, the fact that might hardly be achieved through single incident learning. Most of the focused incidents in this study were significant incidents. Therefore, the empirical findings suggest that even single, large incidents might be incapable to mobilize resources for learning process, especially when learning process requires significant amount of resources and political barriers of learning are sever.

TL also provides a temporal pattern for learning process that might suit well learning from *significant, repetitive* incidents (sometimes known as middle-size incidents (Baumard & Starbuck, 2005)). These incidents are significant enough that require specific learning initiatives (e.g., internal improvement projects). They are also repetitive enough (even though not so frequent) that it justifies a transversal learning mode be taken to formulate solutions based on cross-incident analysis.

6.4 Outsourced learning from ISRIs

The fourth learning mode refers to the pattern observed in situations that the organizations outsourced their learning practices to their specialized suppliers and providers, which served their learning purposes. Analyzing various learning practices in the organizations showed a rather prevalent pattern that the organizations rely on and benefit from the learning practices that take place inside their providers. They benefit from this learning, by maintaining their relations with their providers for handling similar incidents in future. In fact, the organizations *relied on* what and how the providers have learned from incidents.

In other words, outsourced learning mode implies that not only the incident handling process, but also the very learning process takes place inside the provider. This pattern is evident in organizations such as Security-Public that outsourced developing post-incident reports and knowledge articles and running improvement projects.

The organizations even avoided some internal learning practices to pursue their outsourced learning strategy. For instance, the junior expert in Facility-Management group at Supercomputer-Large has had the chance to take part in

some training workshops to learn handling some critical incidents internally. He, however, avoided that temptation, since

"The strategy is that we outsource all these activities" [the junior expert in Facility-Management group at Supercomputer-Large].

The heterogeneity of specialized technologies and their associated incidents also contributed to this strategy. For instance, running a wide range of technologies, Supercomputer-Small decided, from the very beginning, that a single internal facility manager be responsible for dealing with a wide range of providers, and making them improve their incident handling services.

6.4.1 Learning from others, through suppliers

Outsourced learning not only helped the organizations leverage the learning capabilities of their providers, but also could learn from other organizations' incidents through their providers. In fact, the organizations needed to learn from the experience of other similar organizations. However, they had to do this on a wide range of specialized domains. Outsourced learning provided the organizations with the opportunity to benefit from the learning that their specialized suppliers accrue from handling incidents for various clients. The organizations often did not have any entry into the specific knowledge and practices that the providers have learned from other clients.

In addition, rapid technological changes pave the way for adopting outsourced learning mode. This way, the specialized providers can learn and develop new technologies, while, the studied organizations could hardly do so on a wide range of specialized domains.

Nevertheless, the downside is when the incident is occurring for the first time or in a peculiar way that the experience of other clients is irrelevant or even misleading. This happened in some cases in which Supercomputer-Large tried to rely on the provider, while it proved that the case is different from what the provider had experienced in other clients. The situation became worse when the provider tried to use the same solution that had developed for apparently similar incidents of other clients. As a result, Supercomputer-Large acted on its own experience and developed the solution internally through a trial and error process.

6.4.2 Contracts for learning

The relation between the organizations and their suppliers is important in outsourced learning mode. The challenge is how this relation be developed to make the providers learn as quickly and effectively as possible. As exemplified in the case of Security-Public, the organization was benefiting from the learning that its provider had gained through the first two years of engagement with various incidents. However, moving from body-shopping to volumetrias contract led the provider invest in changing tools, methodologies, procedures, and even structures that could allow for handling incidents more efficiently. In fact, paying based on the resolved incidents, rather than number of working hours, encouraged the provider to find ways in which the same quality of services could be delivered at a lower price. These two types of contracts (body-shopping versus volumetrias) illustrate how learning can be embedded in the duties of the provider. What was seen from the view of the organizations as learning from incidents was considered as the core daily business of the providers.

On the other hand, the provider had the fear that too much learning on the services could result in substantial improvement in their efficiencies, which, in

turn, might make Security-Public reduce the service prices in the following year. There were different opinions inside the provider about the impact of too much learning for delivering services. For instance, the CEO of the provider was not concerned about too much learning since he was seeing a trend of fast market expansion in security sector, as he said

"We are not worried about this. The IT security sector is expanding very quickly. It is changing very quickly as well. We need to learn and this allows us to open new opportunities for delivering new services" [the CEO of Security-Public's provider].

This contrasted with the projects manager's opinion that he saw their substantive improvements a tricky decision that should be made concerning the expectations that can create inside the clients:

"We have improved a lot our services. Now we can deliver them much more efficiently. However, the [clients] should not feel that this justifies decreasing the prices. We invested a lot to reach to this point" [the project manager of Security-Public's provider].

To summarize, outsourced learning mode helps the organizations push the learning practices outside their organizational boundaries, while still reaping their benefits for future incidents. This pattern seems quite relevant in ISRIs due to technological specialization, the heterogeneity of these technologies, and rapid technological changes. The organizations pursued this learning mode that helped them avoid entering into a wide range of specialized domains. Instead, they could black box the learning process and incorporate it into the services that they received from their providers. The way that the organizations arranged their relations with the providers was also important for motivating providers to learn effectively and in a timely manner.

6.5 (not) Learning through material replacement

The fifth learning mode refers to a common pattern observed in situations in which the studied organizations avoided similar incidents and reduced their impacts by simply *replacing* parts of their systems. This practice was observable during incident handling process and when the organizations chose to replace damaged parts of their systems. It was also evident in various internal improvement projects that resulted in replacing a specific part. In addition, the organizations were running periodic upgrades through which they were replacing hardware and software elements that could bring about future incidents.

Although this does not fit with the traditional understanding of learning, it has been a common pattern in the studied organizations for capitalizing on the experience of past incidents to avoid them in future or reduce their impacts. This learning mode surfaces some of the assumptions about how organizational learning process takes place in organizations.

The fact that technologies changed quickly made the organizations pursue this strategy actively. This way, they could easily benefit from the learning that the specialized suppliers have gained in developing new versions of technology. In fact, the organizations tried not to engage in the technical aspects and confine their learning scope to just knowing how to replace the damaged parts and manage their suppliers in a way that the experience of past incidents yield future improvements. The modularity and malleability of IT systems and the access right of the organizations for changing their systems are factors affecting learning through material replacement.

6.5.1 Modularity versus malleability

Learning through material replacement becomes more feasible when the underlying technologies become more modular; that is replacing one part of the system does not require changes in other parts. The modularity of systems supported this learning mode, since the organizations could replace one part of their systems, without being forced to replace the other related parts. For instance, Supercomputer-Small is running a very heterogeneous collection of processing and storage systems from different technologies and brands. The fact that all systems can work with each other through an open-source platform helps them replace one part easily.

This learning mode is more relevant for hardware components that are closed, meaning that the organizations cannot make changes into them. In fact, changes in hardware module are often too complex and specialized that only very specialized suppliers are able to do. However, in case of software components that are more open to changes (especially when they are open source), the organizations were relying on improvement projects that aimed at changing and manipulating them (in TL mode). As the case of central storage at Supercomputer shows (see Box 5.2), sometimes the original preference of the organizations was replacement. However, the budget or other practical limitations made them embark on learning how to make changes into the objects. This shows two different ways of capitalizing on the experience of past incidents: learning by knowing their causes and making related solutions versus trying not to get into them and only benefiting from the learning of the suppliers.

6.5.2 Access right for learning

The access right and the arrangements with the providers can challenge outsourced learning mode. For example, Supercomputer-Large has been dependent on a primary provider for its main supercomputer. In spite of system's modularity, Supercomputer-Large has not been able to replace some parts, since the provider has been pushing for an *overall upgrade* in the whole system. This delayed upgrading process because it required a big investment.

This point becomes clearer when we compare Supercomputer-Large with Supercomputer-Small. Supercomputer-Large is based on processing systems that making changes in them are mainly in the control of the provider. However, Supercomputer-Small has built its own systems based on an open-source strategy; that is any technology and brand is welcomed as far as it can work with other systems and technologies through the open-source platform. Therefore, they can make changes to the systems independent of the providers. In the former case, learning through material replacement took place less frequently and more fundamentally than the latter.

To summarize, learning through material replacement is similar to outsourced learning in the sense that the organizations did not enter the details of incidents and their causes. In fact, in both cases, the organizations skipped the process of knowing about the causes of incidents, and tried to learn from them from distance. In other words, the learning process becomes a black box, either by being outsourced to the provider or by being embedded in new artifacts. In both cases, although the organizations have learned literally (because they have done something that helped them avoid similar incidents in future or reduce their impacts), they have minimized their cognitive and practical involvements in the learning process. This is contrary to the normal conceptualization of learning

process that assumes learners develop new knowledge and practices for avoiding future incidents and reducing their impacts.

Chapter summary: This chapter articulated five learning modes by extracting patterns observed in learning practices that the organizations adopted in relations with single or multiple incidents. Table 6.1 summarizes the five learning modes, their definitions, their associated learning practices, and the characteristics of ISRIs that are relevant to each mode. As discussed in each learning mode, various characteristics of ISRIs qualified the adoption of each mode and the way it unfolded in the studied organizations. In learning through incident handling process, the specialization and technical nature of incidents helped the organizations neutralize incidents, yet not normalizing them. In doing post-incident reflection, the technical nature and specialization of ISRIs led to the dominance of individual learning practices, in spite of their original collective design. In addition, rapid technological changes results in learning abstinence that prevented the organizations to embark on analyzing incidents before they accumulate insights on a series of them. Finally, the outsourced nature of IS activities and heterogeneity of ISRIs makes the organization focus on “know-how”, rather than “know-why” in their post-incident reflections.

Rapid technological changes and the high level of specialization created a situation in which the organizations learned from multiple incidents for making specific solutions, based on a temporal cognitive base. The technical nature of ISRIs paved the way for technical experts to dramatize on the past, related incidents to increase the political pressure that is needed for moving the learning process forward.

The modularity of ISRIs and the outsourcing trend of numerous technologies that the organizations operate went hand in hand with the high level of specialization

of ISRIs to lead the organizations to rely on outsourced learning mode. The organizations were opting to follow outsourced learning mode since the technologies were changing so quickly that justified pushing learning practices towards their specialized providers.

The same characteristics resulted in the dominance of learning through material replacement. This way, the organizations were avoiding engagement in specialized practices and understanding technical ISRIs that will be automatically solved when the new artifacts are adopted.

In general, most of the learning practices related to five learning modes were done individually, and informally. Sharing the experiences and reapplying it in other groups happened only to a limited extent because of limited *relevance* of such experience and action to others. This was mostly due to the high level of specialization. For instance, learning from network incidents was irrelevant to other groups (e.g., email server) since the technology, processes, and types of incidents were different. Related to the high level of specialization that made learning process be predominantly individual, it also made learning process be rather informal since it did not make any sense that the organizations add much formality for just one or a couple of people. In spite of such individual and informal nature of learning practices, the organizations incorporated the experiences of incidents into their organizational processes and technological elements (for example through improvement projects). However, this incorporation is done often in an individual, informal manner. Next chapter discusses the insights from learning modes in relation with the existing literature on organizational learning from incidents.

Table 6.2: Five learning modes from ISRIs

Learning mode	Definition	Related practices	Related ISRIs characteristics
Learning through incident handling practice	The engagement of individuals or teams in incident handling process provides them with experience for handling future incidents.	- Handling incidents through ticketing system - Doing Triage - Laboratory analysis	- Specialized and technical nature
Learning through post-incident reflection	Analyzing the incident (when, how, why it happened) after handling it and drawing lessons for future improvements	- Post-incidents reflections - Some Wiki articles	- Rapid technological changes - Specialization of technology
Transversal Learning from Incidents	A process through which the organization tries to solve a specific problem or make a specific improvement based on the experience of several related incidents.	- Internal Improvement projects - Wiki articles - Upgrades	- Negative valance of incidents - Rapid technological changes
Outsourced Learning	The organization relies on and benefit from the learning practices that take place inside its providers.	- Outsourcing incident handling process - Outsourcing post-incident analysis and reflection - Outsourcing improvement projects	- Modularity - Heterogeneity of technologies - Specialization of technologies - Rapid technological changes
Learning through material replacement	By replacing part of the technology, the organization avoids past incidents or reduces their damages in future.	- Replacements during incident handling and Improvement projects - Upgrades	- Modularity - Heterogeneity of technologies - Specialization of technologies - Rapid technological changes

Chapter 7: Discussion

The aim of this research is to understand the process of learning from ISRIs, with regards to the characteristics of these incidents (Figure 2.2 in chapter 2). By adopting a situated learning view, I focused on learning process in relation with contextual factors, particularly the characteristics of ISRIs. Relying on a practice view, I examined the actions (and inactions) that the organizations performed in relation with single incidents (chapter 4) and in relation with multiple incidents (chapter 5) to capitalize on the experience of past incidents. The previous chapter articulated five learning modes that emerged from the analysis of various learning practices observed in the studied organizations (See Table 6.2 in chapter 6).

Therefore, the first contribution of the study is articulating the differences that ISRIs context make for the known learning modes. As discussed in the previous chapter, learning through incident handling mostly revolves around managing tickets that help organizations objectify incidents and reducing their negative connotations (neutralizing). The very technical nature of ISRIs is reflected in the prevalence of individual learning practices, with limited engagement of other peers. In other words, organizational learning is very much dependent on specialized learning that single experts accrue in their incident handling practices. However, the exception is Triage, which is done collectively across a wide range of domains and responsibilities. During Triage, the organizations integrate their experiences of past and current incidents into decisions and solutions related to the incident at hand.

Analyzing how the organizations reflected on their handled incidents showed that unlike the original intentions, this process often takes place either individually or among a few experts. This is mostly due to the high level of specialization of technologies that stems from the technical nature of ISRIs. In addition, rapid changes in IT prevented the organizations from reflecting on incidents and drawing lessons for future incidents, because those lessons would be irrelevant and even misleading when new technologies (with new ISRIs) are in place. In fact, the effort sometimes should be made to avoid such immature and ineffective incident analysis (learning abstinence). Finally, the technical nature of ISRIs and the high level of specialization made the organizations focus on documenting "know-how" (how they can perform if such incident happens in future), rather than going through the causes of the incident ("know-why").

The second contribution of the study is articulating three new learning modes that seem to be prominent for helping the organizations benefit from their past ISRIs experiences. Rapid changes in technology led the organizations to define learning

solutions through the analysis of multiple incidents (transversal learning). This way, they could address fundamental changes that were relevant in a wide range of incidents. Furthermore, the technical nature of incidents allowed them to leverage the unfamiliarity of non-expert stakeholders to dramatize on the past incidents. This, in turn, added to the political pressure needed for moving the learning process forward.

Prior studies have shown the importance of learning from individual or few incidents (March et al., 1991). In addition, recent studies have stressed the importance of learning from rare events (Christianson et al., 2009; Lampel et al., 2009). Contemporary theoretical developments in the literature on learning from incidents emphasize on the importance of inter-incident learning (Moynihan, 2009). For instance, a series of incidents over time can point to a more fundamental defect (Fauchart, 2006) than what can be extracted from a single incident. Sometimes, the heterogeneity of various incidents acts as a source for drawing lessons that might not be drawn from similar incidents (Haunschild & Sullivan, 2002). These studies have mostly focused on the advantages of having multiple sources of information to increase the validity and reliability of learning *content*, compared with the knowledge that can be drawn through analyzing rare events (Lampel et al., 2009).

However, the articulation of political aspects of transversal learning (TL) process shows that incidents play other roles, rather than only providing sources of insights and information. In particular, the technical nature of ISRIs facilitates dramatizing on past incidents that might not necessarily be used in drawing insights for the development of the final solution. In addition, TL, as observed in the cases, is solution-oriented rather than incident-oriented. In other words, the question is not what we can learn from a series of similar or different incidents (the focus of multiple incident learning in the literature). Instead, the question is,

for the specific concerned solution, what aspects of past incidents can help, both cognitively and politically. In fact, the efforts are directed towards finding a specific solution. The relatedness of past incidents is, then, judged based on their potential contributions in providing insights for developing the intended solution, no matter how different or similar they are.

The results of the study also show that the organizations actively pushed the boundaries of their learning process to their providers (outsourced learning). This is because the organizations operate a wide range of specialized technologies that change rapidly. This makes the organizations unable to learn about all the rapid technological changes. In addition, outsourced learning helps them benefit from the learning that their specialized suppliers accrue from handling other clients' incidents. The attempt, then, focuses on defining a relation with the suppliers that assures reaping learning outcomes, without being involved in them.

Finally, my analysis shows that the organizations actively replaced their systems and their parts as a common practice to avoid similar incidents in future. In fact, the organizations preferred not to learn about the causes of incidents and the solutions that can prevent them. All this understanding for them either is useless, since the technology is changing very fast, or is almost impossible to attain due to the heterogeneity of their specialized technologies. On the other hand, the modularity of information systems helps them make rapid replacements in specific parts of the system or major upgrades to their overall systems.

In other words, in both outsourced learning and learning through material replacement, the technical nature and the rapid changes of ISRIs make the organizations learn from past incidents, without knowing about them. This is different from the dominant perspective of literature that envisions knowing about the causes of incidents (either through engagement with the incident or by analyzing it a posteriori) as the first step in the learning process.

In the remaining of this chapter, I examine how the articulation of these learning modes can advance our understanding of learning process from ISRIs. These three new modes, not only show that organizations go beyond the traditional expectation of literature in their attempt to learn from their ISRIs, but also show how learning from ISRIs might be different from learning in other contexts. The claim is that understanding OL from ISRIs requires rethinking (1) the temporal pattern of organizational learning process, (2) the uncertain and emergent nature of learning, (3) materiality of knowledge and learning, (4 and 5) the politics of learning, (6) learning governance, and (7) the organizational aspects of learning process.

More specifically, the articulation of transversal learning mode questions the established view that learning from incidents either takes place during or after incidents. Second, the uncertainty and changing nature of ISRIs makes the organizations rely on empty knowledge articles as tools that not only expose their knowledge lack, but also materialize their learning needs and inspire a transversal learning on critical themes.

Third, the modularity and malleability of IT artifacts, linked with their rapid changes, influence the material aspects of learning and knowing, different from situations in which the artifacts are stable and tightly coupled.

Fourth and fifth, the technical nature of ISRIs and the dominance of ticketing system in the process of incident management have paved the way for the emergence of some political mechanisms in the learning process, such as neutralizing and dramatizing.

Sixth, the widespread outsourcing of IT services has shown the importance of considering new governance systems in analyzing how learning process is organized; beyond the traditional intra versus inter organizational learning dichotomy.

Finally, the individual pattern of learning practices poses questions on the organizational aspects of learning process. This, in turn, makes us rethink the appropriate unit of analysis for studying OL process, when we focus on situations that are characterized by high-level of specialization and extensive outsourcing of wide range of modular technologies.

The discussion of these aspects is grounded in the existing literature, which, in turn, allows extracting new theoretical insights.

7.1 Temporal pattern of OL process

Studies on organizational learning from incidents have shown that learning process can take place *during* incident handling process, known as intra-incident learning (Moynihan, 2009), and *(right) after* handling incidents, known as post-incident learning (Boddie, 1987; Kasi et al., 2008). The review of IS literature (chapter 2) shows that the dominant focus has been on post-incident reflection and analysis, through practices such as post-mortem analysis and post-incident evaluation. These studies suggest that involved actors should run learning process quickly after incident handling process to avoid losing important information and sufficient attention.

A growing body of literature also considers learning during incident handling by focusing on learning through working and practicing (so called “knowing” (Nicolini et al., 2003; Orlikowski, 2002)) and mindfulness (Butler & Gray, 2006; Weick & Roberts, 1993). In these views, learning takes place *while* organizations handle incidents. Therefore, there is no temporal separation between OL and incident handling process. This distinction has been insightful as it shows the different

opportunities and challenges that organizations might face during and after incident handling, when they leverage their incident experiences.

This study identified transversal learning (TL) as a learning mode that manifests a different timing of OL. Instead of focusing on a single incident, which is the case in learning through incident handling and in post-incident reflection, TL concentrates on multiple incidents that cognitively and politically contribute to the intended learning solution. Unlike the previous two learning modes, in which OL process is triggered by a specific incident, TL is not triggered necessarily by any specific incident. In other words, transversal learning mode does not necessarily take place during or after any (major) incident. It occurs *in parallel* with incidents.

Transversal learning takes place when a proper *moment of learning* appears. Organizations learn transversally, for example, when they are upgrading their systems and technologies, in periodic (mostly seasonal and annual) events such as renewing contracts when they are designing new systems, during scheduled interventions for doing overhaul or preventive changes, when the daily load work drops (e.g. during summer time), when required resources (financial, technological, or human) becomes available, and simply when key influential actors are present.

Moreover, TL often takes place *long after* a series of incidents. This *temporal distance* helps organizations build a holistic, stable cognitive base about the solution, mount sufficient political pressure around the concerned solution, and develop required socio-material elements as the basis for materializing the learning initiative. This temporal pattern of learning seems to be particularly relevant in ISRIs since rapid technological changes require that organizations avoid hasty learning actions by developing a stable understanding of incidents. In addition, this temporal distance helps them bypass the negative content of

incidents (e.g., by disguising the relations with incidents) that often poses barriers in learning during and right after a specific major incident.

This challenges the recommendation of literature that organizations should engage in learning process right-after incidents to prevent information loss and capitalize on the existing political pressure due to fresh wounds (Elliott, 2009). TL offers an alternative learning mode that takes place in an appropriate moment, when required cognitive, political and socio-material elements are already constructed.

7.2 Learning by seeding through semi-empty articles

Normally, it is expected that Wiki articles document the experience of an incident (in post-incident reflection learning), and codifying lessons that have been drawn from a series of incident (in transversal learning mode). However, we observed that the articles in the Wiki system are sometimes rather *empty*. The experts, who feel that some themes are important for learning, create an empty article with a tentative, broad title (for example “legal consequences of exposing organizational passwords”). This often happens when such learning need arises due to facing one or several incidents. However, the socio-technical environment (e.g., the laws, the capabilities of technologies, the patterns of users’ behavior) is still too uncertain to allow the experts to draw any specific lesson. They, at most, might put a few bullet points as the potential issues related to the article that should be considered as *points of attention*. For example, one bullet point in the abovementioned empty article was “check if the formal employee contracts have any point about this issue.”

The empty article objectifies and materializes the learning need. The empty article is there, in the Wiki system and is observed by other experts. It generates questions about a specific topic, and makes other experts think about it and search for it. It draws attentions towards an overall theme that is not known, but is important to be known. In other words, the empty article articulates the knowledge lack. Because Wiki articles are shared, it publicizes such need that can make the community of experts search for it. In this the sense, empty article acts as a hole that intakes learning efforts. This way, empty articles contribute to the need for learning among experts, rather than supplying some knowledge to them.

The prevalence of empty articles is partly because technologies change quickly. This has been one of the rationales of the experts for creating empty articles that could help them avoid hasty conclusions about issues that are still changing and are uncertain (of course, another reason for creating semi-empty articles is the time limitations that experts face).

Accordingly, (semi)empty articles act as seeds that plant learning practices. They are created as starting points around which new learning practices are taken and new knowledge is developed. They grow, as they become complete and are developed into complete articles. In other words, they act as holes, that different learning agents throw their learning efforts into them. Of course, the shape of the hole – i.e., the words selected for defining the empty article and the few initial bullet points- affect what kind of new knowledge is thrown into them.

Empty articles are also helpful for the sake of integrating various ideas and opinions. They act as discussion meetings in which an open question is going to be discussed. Empty articles say loudly that please let me know if you have ideas about this theme. This way, each expert knows what other peers do not know, but they need to know. This is different from the idea of transactive memory (Brandon & Hollingshead, 2004), which relies on what actors know about other actor's

knowledge. Empty articles, instead, let others know what other actors do not know. This can help in avoiding naïve reliance on others in domains.

Empty articles are different from open questions that an expert can pose through a wiki system. From social perspective, the creation a semi-empty article implies some level of expertise. Unlike questions, empty articles articulate learning need, without putting the creator of the article in a lower status in front of other experts. In fact, when an expert creates an empty article it means that he is the most knowledgeable person about this theme, who still does not know much about it.

Finally, empty articles can contribute to the learning attitude among experts. When senior experts who are supposed to know a lot in their domains develop several empty articles, it signals to junior experts that there would be no shame in doing so. Empty articles institutionalize admitting knowledge lacks, especially because those wiki articles have an objective presence (more than just keeping the questions in the mind).

7.3 Regimes of materiality and learning process

The pattern of OL process in the studied organizations shows that modifying, upgrading, and replacing a specific part of information technology is a common practice. As mentioned in the previous chapters, most of the improvement projects are defined based on developing a local script that can solve a problem and avoid some incidents. Similarly, the organizations have constantly considered upgrades and replacements as a critical strategy to avoid incidents in future.

This resulted in the prevalence of a new learning mode, which is learning through material replacement. This learning mode contrasts with the conventional learning

modes in which the organizations need to develop an understanding of the faulty technology, the causes of incidents, and ways of preventing them. Nevertheless, learning through material replacement addresses learning aims –i.e., avoiding similar incidents in future and reducing their impacts – while the organizations do not engage in, and even avoid, understanding the incident and its causes. In doing so, they only replace the system or part of that with a new element that automatically reduces the likelihood of past incidents.

This pattern results from the modularity of information systems that allows organizations to make local modifications without changing other parts or the basic design. Modularity refers to the systems' design that is composed of components that can be independently changed, without requiring changes in other related components. These components are interacting through interfaces when they are operating as parts of the whole system. In fact, a modular system lends itself to be decomposed into parts in such a way that changes in one part do not require changes in the other parts. Regardless of the agency of social actors, the very design of the system determines the extent to which the system is integral versus modular (Kallinikos, 2012).

The fact that organizations are running a wide range of specialized technologies also made them perform such changes frequently by relying on their specialized providers. In most of the cases, the organizations just need a thin understanding for identifying which part needs to be replaced and a minimum effort to get a specialized provider to do so. In this way, learning through material replacement and outsourced learning can complement each other when the replacement is mostly done through the supplier. Then, the supplier analyzes the faulty system to understand why such incident happened.

In addition, the high pace of technological changes requires that organizations introduced quick modifications and changes in their technologies, while some

other parts (e.g., some basic technologies) are still unchanged. Nevertheless, the digital aspects of information technology, linked with the overall open-source strategy, helps in situations that revising or developing a local software element (known as scripts) can postpone or remove the need for making changes in some stable parts of the system.

This pattern relates to the discussion on materiality in the learning process (Leonardi & Barley, 2008; Leonardi, 2010; Nicolini et al., 2011; Orlikowski, 2010). Previous studies have shown that material objects can play different *roles* in the learning process, including transferring, translating, and transforming knowledge across specialized communities (boundary object) (Carlile, 2002), acting as an object to be known and investigated (epistemic object) (Miettinen & Virkkunen, 2005), serving as an element upon which social actors perform their practices (practice object) (Engeström, 2000), and being taken-for-granted material conditions within which social coordination takes place (Nicolini et al., 2011).

Our empirical findings add two specific nuances to this discussion. First, the abovementioned roles assume that the object itself remains almost unchanged, although it is embedded in various social processes such as knowing and practicing. However, the learning during incident handling and in transversal learning mode often requires changes in material objects (e.g., upgrading a faulty technology). In normal situations of operation, material objects are considered as stable and unchanged, while they are evolving during the design period. During incident handling process, the doors of materiality open as the faulty systems are not under their normal operations. This is the time that organizational actors can introduce changes to the material objects, as well as doing several experimentations on it. This *window of material opportunity* becomes selectively (e.g., only some specific parts of technology) open again in transversal learning mode and when doing upgrades and replacements (e.g., through scheduled

interventions). In fact, these two learning modes (learning during incident handling and transversal learning) are two critical moments concerning material learning. Although previous studies have recognized the importance of opening social windows that appear during and right after incidents (e.g., “policy window” (Stern, 2002)), little attention has been paid to material window that often opens during incident handling process and during transversal learning mode.

More specifically, when the windows of materiality are open, the learners can *experiment on the material object* and *introduce changes* into it which both are important for the sake of learning from incidents. This is different from the epistemic and practice roles of objects when social actors try to investigate and work with the object, respectively. In both cases, the object remains unchanged. However, the technology (faulty system) is changing and revising constantly during incident handling process and then during other learning modes.

Furthermore, even when the policy windows are open after the control of incidents (e.g., through post-incident analysis), the organizations might not be able to perform their analysis of the incident until the time of scheduled interventions that they can experiment on their systems. In that periods, they can reproduce the incident (in a controlled environment, and only in a specific part of the whole system) to understand how the incident works and how it can be prevented.

The characteristics of ISRIs, in particular, and information systems in general, allow for this learning mode. The fact that IT systems are modular provides a situation in which local, quick replacements and modifications become feasible. The digital aspects of information systems, and in particular the open source applications, also add to this feasibility since the organizations can adapt some digital aspects – e.g., a specific module of the application- even when changes in the hardware and the overall system architecture is not feasible.

This pattern can extend the discussion on the role of material objects in the learning process by pointing to a specific *regime of materiality* that dominates ISRIs. This regime of materiality, as contrasted with materiality regime in situations such as healthcare system, paves the way for making more chances to their material conditions *as* they are working.

The second nuance relates to the dominance of replacement and upgrades as a common way of capitalizing on past ISRIs experience. Following my overall definition of learning from incidents, practices taken for replacing material objects are considered as learning practices. However, these practices do not require that the central organization go through cognitive levels of learning. This is because the organizations are operating a wide range of specialized technologies that makes it difficult for them to gain sufficient knowledge about all of them. Moreover, their outsourcing strategy also makes them abstain from learning about their technical incidents, even in situations that they can do. This is partly because these technologies change so quickly that the knowledge about them becomes obsolete.

This pattern modifies the dominant models of organizational learning that assume that at least part of the learning process is based on developing an understanding of the incident and its causes. For example, Kim (1993) argues that organizational learning starts with the understanding of experts about the problem, and then it becomes shared into collective schemas, which, in turn, result in creating changes at organizational level. Learning through material replacement, however, relies on a very limited cognitive involvement of organizational actors (just at the level of detecting the problem), and mostly relies on changes in material components. This can question an implicit assumption that organizational actors need to develop the knowledge of incidents (known as "cognitive aspects of learning" (Fiol & Lyles, 1985)) to make changes in their behaviors (known as behavioral learning). In fact,

learning through material replacement is a complete learning process in which cognitive learning is at minimum.

7.4 From normalizing to neutralizing

OL studies have shown that organizations are prone to ignore or filter out the signals that do not fit their expectations (Hedberg, 1981; March & Olsen, 1975). Further studies have shown that even organizational actors tend to overlook, suppress, and distort negative signals that can threaten their power (Lawrence et al., 2005). Similarly, Brown & Starkey (2000) have argued that there are "identity defense mechanisms", such as denial, rationalization, idealization, fantasy, and symbolization that prevent organizations from learning.

In case of incidents, the chances are high that organizations adopt strong defensive mechanisms such as normalizing incidents (Turner, 1976) into unimportant signals. This can lead to ignoring the signs of incidents and misinterpreting incidents, so that organizations might fail to tap learning opportunities, especially when still windows of opportunities for making changes are open.

However, what we have seen in the case of *neutralizing* process is that organizations reconstruct incidents in a way that they do not question the established self-concept. Neutralizing is different from normalizing because organizations do not overlook or ignore the incident. Instead, they admit it attentively (e.g., by entering it into their formal incident handling process, through ticketing system). What they do, however, is that they re-create the incident as a as "ticket" which implies "an important task". Thus, the incident loses its identity-threatening connotation, while still keeps its inherent criticality (urgency).

Neutralizing has to do with the discourse that revolves around incident handling process. More specifically, the language that the studied organizations had established to talk about incidents was concentrated around neutral concepts such as “tickets” and “tasks”.

In addition, objectifying incident through ticket and handling it through ticketing system facilitates neutralizing. Similarly, the centrality of technology (e.g., the faulty hardware) seems to divert the attention from people to objects. Hence, the discussions and efforts focus on “what” (e.g., what is broken and which part of the system is faulty) rather than “who” (e.g., who was guilty, who did so).

Neutralizing also reflects the fundamental assumption about technology in IT departments. For them, even in case of sever incidents, the blames would often be diverted towards technology. This is the technology that is assumed as “something that always fails”, and “never, it is perfect”. This is well illustrated in the technical discourse of programmers who believed that “any code has its own bugs”. This facilitates neutralizing since the incident is considered inevitable. Instead of interpreting incident as the misbehavior or failure of social actors in complying with their responsibilities, the focus is diverted towards acting quickly and effectively to solve it and making improvements that reduce the chance of similar incidents in future.

This perspective is dominant even in the studied organizations that they were handling very critical infrastructures such as supercomputers and national-wide information systems. Such organizations are classified in the literature as high-reliability organizations (La Porte & Consolini, 1991; Roberts, 1989; Roberts, 1990), in the sense that they have little room to compromise their service quality. However, the picture inside the studied organizations, especially in the time of major incidents, is different from the prediction of this literature. The studied organizations have tolerated a wide range of incidents, although they maintained

their service quality through other mechanisms such as technological redundancies. For them, incidents were inevitable. More seriously, their default assumption was that technology fails. This way, they could maintain their reliability through continues learning, with little fear about working through and reflecting on past incidents. This has been facilitated through neutralization.

Neutralizing suggests a different strategy than what Brown & Starkey (2000) suggest in their psychodynamic analysis of identity-defense mechanisms. They suggest that for organizations to learn, they need to foster a culture and mechanisms that make their members critically reflect on their collective identity. In doing so, the organizations need to work on the collective wisdom, which allows for controlling and counteracting identity-conformation mechanisms. They suggest that organizations can foster organizational learning by focusing on mechanisms that "promote attitudes of wisdom", through mechanisms such as "critical self-reflexivity" and "identity-focused dialogues". Their suggestion is that organizations should learn how "to promote critical reflection upon organizational identity" (Brown & Starkey, 2000: 103) through "understanding and the mitigation of those ego defenses that tend toward a regressive retreat" (Brown & Starkey, 2000: 103).

However, in the case of neutralizing, such identity-threatening aspects of incidents are removed from the very beginning. Instead of acting *against* the self-defense mechanisms, organizations *reframe* and *reconstruct* incidents in a way that do not contradict their identities, so that they could tap the learning opportunities. The reconstruction of incident into a neutral socio-material element (e.g., a ticket) helps circumventing such identity-defense mechanisms, which, in turn, removes the need for counteractions in order to foster learning process.

7.5 From rationalization to dramatization

The significant negative impacts of major incidents can both create *political pressure* for organizations to embark on learning process (Stern, 2002) and pose challenges in the learning process if *blame-game* (Pearson & Mitroff, 1993; Stead & Smallman, 1999) and "protecting vested interest" (Elliott & McGuinness, 2002) become prevalent. Previous studies mostly consider the political role of incident during and right after incident handling process, when the flames of incident and the memories of pains are still in place.

Even when incidents are too big to be normalized, the organizations might sacrifice their learning opportunities during and right after incidents by rationalizing negative incidents (e.g., attributing them to uncontrollable, external factors) (Baumard & Starbuck, 2005). Therefore, organizations might fail to see incident as a learning opportunity, resulting in no incentive for adopting learning practices. For instance, Baumard & Starbuck (2005) analyzed several large negative incidents in a telecommunication company and realized that the managers do not consider the incidents as learning opportunities due to the attribution of them to uncontrollable causes.

Although the current study focused on major negative incidents, the studied organizations showed little tendency to rationalize their incidents. It first appeared in the willingness of the organizations to openly talk about their incidents. Even when the organizations were asked to pick incidents that they are more willing to reflect on, they were quite open to pick any incident that the researchers would suggest. Second, it was reflected in the self-critical language of most of the managers, without being embarrassed.

As described in the previous chapter, during incident handling process and even right after that, the political pressure was reduced when incidents became

"tickets" at the beginning of the incident handling process. In fact, they became neutralized as tickets, rather than being perceived as negative incidents. The technical levels, in my cases, were mostly considering even big incidents as *another task* or simply a *request* that should be solved.

However, the findings showed that, during transversal learning, the organizations were even *playing on past incidents* and intentionally rejuvenating their pains to mount political pressure that is needed for driving learning projects. More specifically, transversal learning mode surfaced an important political dynamic: *dramatizing,* in which senior experts who were already aware of the importance of some learning initiatives *reconstructed* and *added to* the political pressure around that specific initiative. The targets of dramatizing were managers and the clients who had little technical knowledge. Dramatizing could create a feeling of necessity among these actors who had to decide about learning initiatives (managers) and actually perform them (clients). Senior experts who were dramatizing were drawing upon several incidents that could be interpreted as somehow relevant to the learning initiative. This way, dramatizers were revitalizing the pain and envisioning potential damages. Lack of technical knowledge made targets of dramatizing believe in the frightening stories narrated by senior experts, although they were not aware of their technical contents.

Dramatizing resonates with the insights in change management literature about creating urgency for change. More specifically, change management literature has shown the importance of enhancing the readiness for change by persuading actors (Armenakis, Harris, & Mossholder, 1993; Kotter, 1995; Mento, Jones, & Dirndorfer, 2002). In the case of dramatizing, the past, related incidents are selected by technical experts (not managers of the organization) and have been reinterpreted in a way that show the importance of the underlying specific change project (mostly in the transversal learning mode).

While the literature on incident learning is predicting that the organizations tend to rationalize, rather than dramatize on their past incidents, the insights from this study point to the contrary. This might be because the existing literature has focused on two learning modes (learning during incident handling and post-incident reflection), while, dramatizing is more likely to appear in the transversal learning mode in which the multiple, unclear relations with past incidents reduce the chance of blame-game. In addition, neutralizing incidents facilitates dramatizing on past incidents, since those incidents have not been perceived as threats to individual and collective identities. Furthermore, the technical nature of ISRIs contributes to dramatizing since IT experts can act on the unfamiliarity of other learning actors (e.g., managers and users) with the technical aspects of ISRIs.

7.6 The governance of OL process

The studied organizations were capitalizing on the experience of ISRIs by relying on the learning practices that took place inside their providers (outsourced learning). This way, the organizations tried to cope with the heterogeneity of technologies and their related incidents and the rapid technological changes, by relying on distributed learning practices located outside their organizations. The modularity of their systems also facilitated this process. As a result, the specialized providers could learn and apply their learning to the associated technology, without affecting other parts of the system.

Outsourcing IT services is a well-documented pattern in IT industry (Cha, Pingry, & Thatcher, 2008; Levina & Vaast, 2008; Rai, Maruping, & Venkatesh, 2009; Ramasubbu, Mithas, & Kemerer, 2008). In this line, various studies have examined how a central organization can learn through interactions with its suppliers (Cha et

al., 2008; Ramasubbu et al., 2008). For instance, Cha et al. (2008) suggest a model in which the more organizations involve in knowledge transfer from their providers, the more long-term benefits (e.g., reduction in production cost) would be gained and the less would be the rate of internal knowledge depreciation. The results of their simulations indicate that organizations should actively engage in transferring knowledge from external companies into their internal memory to avoid being locked in long-term ineffective relations.

However, our findings show that the organizations were adopting a different strategy. They were intentionally *striving not to learn*, and rely on the learning inside their suppliers. This way they could benefit from the competition among their providers in successive contracts. What they learned was how to develop and improve their contracts to enhance their benefits in such relations. For instance, moving from body-shopping to volumetrias, Security-Public motivated its provider to make its services more efficiently. Similarly, Supercomputer-Small tried to create a competition among its suppliers by frequent, open bids, in which defining the requirements for new systems was informed by past incidents. Instead of trying to learn technical aspects from the provider, the organizations learned how to contract them (Mayer & Argyres, 2004).

In fact, the organizations created a specific governance system in which they could accrue the benefits of learning that takes place inside their specialized providers, without being dependent on internal learning practices. This learning governance is different from the dominant view that considers that learning process occurs inside the organizational hierarchies. It also contrasts with the idea of collaborative learning that takes place when organizations learn through their interactions with other organizations (e.g., see literature on interorgnizational learning (Ingram, 2002; Kraatz, 1998; Larsson, Bengtsson, Henriksson, & Sparks, 1998; Scott, 2000) and vicarious learning (Denrell, 2003; Kim & Miner, 2007)).

Outsourced learning is close to the findings of literature on outsourced R&D that argues organizations tend to outsource their knowledge-intensive activities especially when they establish a long-term relation with suppliers that have proved to be capable for improving their services and providing knowledge-intensive services (Maskell, Pedersen, Petersen, & Dick-Nielsen, 2007). In their study on Danish international companies, Maskell et al. (2007) show that in the first step, the companies outsourced their activities to reduce their costs. In the next step, the managers broadened their perspective about the advantages of outsourcing knowledge-intensive activities. Hence, they pursued outsourcing more knowledge intensive activities to their providers to seek new knowledge that those providers can bring about. Recent studies on R&D outsourcing and cross-organizational knowledge collaboration have argued that organizations reduce their knowledge collaboration to reduce their risk of "unintended leakage of valuable" (Martinez-Noya, Garcia-Canal, & Guillen, 2013: 68) knowledge and technology to their suppliers.

However, our analysis shows that the studied organizations did not consider knowledge drain to their suppliers as a threat. Instead, they were willing to push the process of learning outside their boundaries to avoid investing in quickly changing domains of expertise, while making their specialized suppliers learn more quickly and effectively (through multiple clients) than what they might have been able to do internally. In fact, outsourcing, for them, is not another path for seeing knowledge about their incidents and ways of avoiding them. Instead, it is a way to *avoid* acquiring knowledge about the incidents and technologies that are heterogeneous, specialized, and rapidly changing, while still benefiting from the results of learning practices done by the providers.

7.7 Where is *organizational* learning?

One of the tenets of organizational learning literature is that it is not enough for the organizations that only their individuals learn. The reason is that organizations need to translate individual learning into organizational level in order to expand and replicate learned practices. Therefore, they do not reinvent the wheel in parallel places inside the organization. They also need to institutionalize such learning through organizational routines, collective practices, and material objects; so that organizations do not lose their learned capabilities if some of their members leave the organization (Easterby-Smith, 1997; Kim, 1993). In other words, learning journey is incomplete, if it only confines to local, individual learning practices.

The current study has focused on organizational learning, by examining the practices that both individuals and collectives performed in order to enhance organizational capabilities for avoiding future incidents and reducing their impacts. In doing so, the study explored practices, procedures, structures, and mechanisms that helped the studied organizations transform their individual learning into organizational capabilities. However, the empirical findings showed that many learning practices are done individually or, at most, among a few actors. For instance, a specialized expert in a narrow domain of technology often performs most of incident handling practices, post-incident reflections, improvement projects, and even managing outsourced learning. Even practices like open-seminars that were designed to be done collectively, gradually turned into individual tasks of documenting lessons or ceremonial meetings among a limited number of experts.

The analysis showed that such individual pattern is mostly due to the high-level of specialization that stems from the technical and modular nature of IT systems.

This creates situations that only one or a few experts are related to learning practices. In fact, the involvement of other actors in learning practices is often irrelevant. Therefore, the individual pattern of learning practices is not an indication of organizations' failure in progressing towards organizational learning. Instead, it reflects the very specialized, and predominantly outsourced nature IS activities.

The prevalence of outsourcing IT services, and particularly outsourced learning practices, contributes to this pattern. Most of the interactions in learning process pertain to interactions between (single) experts and the associated specialized provider. In fact, the *relevant domain of learning process* is not the whole organization; rather it is the single expert plus the specialized provider.

Such individual learning pattern, however, involves the risk of losing learned capabilities in case the specialized expert leaves the organization. However, the organizations were relying on their providers through their outsourced learning mode to alleviate such a risk. For them, maintaining relations with specialized providers who are actively learning from incidents of their clients has been as important as keeping their experts. This is because the strategy has been not engaging in learning about so many technical domains. In addition, learning through material replacement contributes to this pattern since the internal experts can simply learn how to replace faulty parts (which is not a specialized task), when some domain experts are absent.

The fact that organizational learning is pursued through individual learning practices seems contradictory with the overall intuition of OL literature that suggests learning process is incomplete if it is confined to the individual level. The findings of this study suggest that such intuition should be revised by changing the unit of analysis from the whole organization to a system of actors that involves (few) internal experts and the specialized providers. In other words, the

meaningful unit of analysis for examining whether an organization has learned from its ISRIs or not should include the specialized providers. Thus, the traditional boundaries of organizations do not make much sense for defining and studying *organizational* learning, since outsourced learning is a dominant learning mode.

Chapter summary: The chapter discussed how the articulation of various learning modes advances our understanding of OL process in relation with the characteristics of ISRIs. I discussed how insights from various learning modes help us rethink the temporal aspects of OL process, the role of materiality in the learning process, the politics of learning process, the governance system through which learning process is organized, and the organizational aspects of learning. Table 7.1 summarizes the findings of the empirical study, the related learning modes, the involved characteristics of ISRIs, and implications for research and theory.

More specifically, transversal learning mode questions the dominant view that learning process takes place either during or right after incidents. It broadens our perspective by showing important learning practices that take place long after several related incidents. Learning through material replacement points to the importance of open windows of materiality, besides open windows of policies. The findings also show how the characteristics of ISRIs, such as their technical nature, allow organizations neutralize those incidents, while avoiding their normalization, during incident handling process and right after that. In addition, the articulation of transversal learning mode allows us to observe and explain how organizations dramatize on their past incidents, while avoiding the rationalization of incidents. Outsourced learning mode unravels the tendency of organizations to avoid learning practices related to specialized, rapidly changing technologies.

Finally, our analysis shows that the technical and specialized nature of heterogeneous ISRIs leads to the prevalence of individual learning practices, even when organizational learning outcomes are intended. This pattern, thus, invites rethinking the appropriate unit of analysis from the whole organization to a specialized system of actors across the organization and its specialized providers.

Table 7.1: Summary of discussion points and implications for research

Empirical finding	Related learning mode	Contributing ISRI characteristics	Implications for research
Learning process might take place with a significant temporal distance with associated incidents, in a proper learning moment	Transversal learning	Rapid changes in technology	We need to think beyond the established dichotomy that implies learning either takes place during or (right) after incident handling process.
Semi-empty articles help seeding learning efforts and planting knowledge about emergent themes	Transversal learning; Post-incident reflection	Changes in technology	It is important to pay attention to the emergent and evolutionary nature of knowing, with regards to the social aspects of learning communities; attentions should also be paid to the demand for learning.
Modularity and malleability of IT systems and open windows of materiality pave the way for various material learning practices	Transversal learning; learning through material replacement	Specialized nature of ISRIs and their underlying technologies; modularity of information technologies	Understanding the materiality of learning also requires looking at periods and processes through with objects are changed and modified; Different regimes of materiality can influence learning practices.
Neutralizing incidents through tickets and material objects helps avoiding political tensions and normalizing them	Learning through incident handling	Technical nature of ISRIs	Mechanisms such as neutralizing can prevent identity-threatening mechanisms from the very beginning.
Dramatizing on a series of past incidents helps mounting political pressure that is needed for driving learning process	Transversal learning	Technical nature of ISRIs	Although during and right after incident handling process organizations might rationalize their incidents, dramatizing can help in rejuvenating the pains of incidents long after.
Outsourced learning mode shows the importance of considering the governance system through which learning practices are pushed outside the boundaries of organizations	Outsourced learning	Heterogeneity of ISRIs; Technical nature of ISRIs; Rapid changes in ISRIs	Outsourced learning questions the established dichotomy of learning governance –i.e., intra versus inter-organizational learning.
The prevalence of individual learning practices in the process of learning from ISRIs shows limited relevant actors in the learning process	Learning through incident handling; Post-incident reflection; Transversal learning; outsourced learning	Technical nature of ISRIs; Specialization of ISRIs; Outsourced IT services	The proper unit of analysis for judging about the organizational level of learning is not the whole organization; instead, it is a system of actors that include the few internal experts and the specialized providers.

Chapter 8: Conclusions

This chapter summarizes the findings of the study and their implications for theory and research. It follows by commenting on the limitations and suggestions for future research. Finally, I comment on some potential implications for management practice.

8.1 Summary of findings

The research aimed at understanding "*how does organizational learning process from major, internal ISRIs unfold?*" Through a qualitative multiple case study, I examined the practices that four IT organizations adopted to leverage the experience of their major ISRIs. The results of the empirical study showed that the organizations adopted a wide range of practices *during* incident handling process – i.e., handling incidents through ticketing system, interactions with specialized providers, and doing Triage- and *after* that –i.e., post-incident reflections and laboratory incident analysis. It also showed that the organizations adopted several practices –i.e., using Wiki systems, executing improvement projects, and performing upgrades- that allowed them to leverage the experience of *several* incidents through specific improvement initiatives.

Analyzing these learning practices resulted in articulating five learning modes: (1) learning through incident handling process; (2) post-incident reflection; (3) transversal learning (TL); (4) outsourced learning; (5) learning through material replacement. The articulation of these learning modes helped examining how the characteristics of ISRIs can influence learning process.

As for the first two learning modes, the study showed that the technical nature of ISRIs implies that the organizations mostly act individually or in small, specialized groups when they are learning from their major ISRIs. The specialization of technology, linked with its modularity, makes it unnecessary for many actors to take part in most of learning practices. The fact that ISRIs change frequently, due to technology changes, makes the organizations be cautious in performing the expected learning practices, such as post-incident reflection, and even further, attempt to avoid such practices (learning abstinence). In addition, the emergent and changing nature of ISRIs might require organizations to learn in an

evolutionary way, across various incidents over time (learning by seeding), rather than relying on definitive lesson drawing approaches.

In addition, the characteristics of ISRIs such as their technical and changing nature can lead organizations to adopt other learning modes rather than learning through incident handling process and post-incident reflection. More specifically, the findings of this study showed that transversal learning from a series of incidents is a very common pattern that allows organizations to learn fundamentally from various incidents over time. In addition, by relying on outsourced learning and learning through material replacement, the organizations can still benefit from their incidents' experience, while they do not engage in knowing the causes of incidents and how they can fix them. In other words, the organizations can bypass the cognitive process that is often assumed in traditional learning theories.

8.2 Implications for theory and research

The analysis of learning practices shows that the characteristics of ISRIs –i.e., technical nature, heterogeneity of technologies and actors involved, the modularity of the technology, and rapid technological changes- create a situation in which organizations not only engage in traditional learning modes (learning through incident handling and post-incident reflection), but also actively embark on other learning modes such as transversal learning, outsourced learning, and learning through material replacement.

The findings suggest that learning from ISRIs might not be necessarily triggered by incidents. They can take place with a temporal distance, and in parallel, with a series of incidents, in a proper moment of learning. This indicates that we need to look beyond the established dichotomy in the literature that suggests learning

practices take place either during or (right) after major incidents. This becomes particularly important in cases like ISRIs that the underlying technologies change rapidly. Therefore, organizations need to construct stable patterns across incidents by bypassing technical noises for formulating fundamental learning initiatives. They can also mount enough political pressure by leveraging (e.g., through dramatizing on) a wide range of related incidents.

The findings also show that the underlying *material regime* can influence the learning practices. More specifically, the findings indicate that ISRIs rely on a modular, changing material regime, which allows for adopting local, fast technical improvements, through replacements and upgrades. This pattern can contrast with integrated, stable material regimes in which learning from major incidents would be implemented through basic design changes.

The articulation of transversal learning mode, vis-à-vis the two traditional learning modes also reveals how blame-game (that often occurs due to the proximity to the incident time) can be reduced by introducing temporal distance between learning practices and the incidents, and through disguising the relations between incidents and the learning practices. In addition, transversal learning surfaces the role of *dramatizing* as a political dynamic through which learning agents can add to the political pressure for driving learning initiatives, by leveraging the ignorance of influential actors.

Finally, the results show that in situations like ISRIs that organizations deal with a wide range of heterogeneous, and rapidly changing technologies, the learning process might be outsourced to specialized providers. Thus, organizations might not engage in internal learning process and through interactions with external partners. However, this learning governance requires that contractual relations between organizations and their providers be set in a way that allows them to accrue the results of fast, specialized learning efforts performed by providers.

8.3 Limitations and future studies

The study is an initial step for understanding how learning from technical incidents occurs. There are several limitations that call for future studies. First, the findings are limited to the specific empirical setting, namely IT organizations. The fact that the studied organizations were specialized in IT domain might have resulted in the emergence of several patterns. For instance, it is possible that the studied organizations embarked on various improvement projects internally due to their knowledge about IT domain. Therefore, the organizations that are not in IT industry (e.g., pharmaceutical companies), might lack internal IT capabilities to involve in such learning practices from ISRIs. This way, they might rely more on outsourced learning mode and learning through material replacement.

Second, being the first study, the attention has been mostly paid to *exploring* the patterns of learning practices (learning modes) and examining how the characteristics of ISRIs might influence them. Thus, capturing the *heterogeneity* and *contingency* of learning from ISRIs has been beyond the scope of this study. Thus, further studies are needed to examine the differences between various types of organizations (e.g., large versus small), in different industries (e.g., IT intensive versus non-IT-intensive), in different contexts (e.g., public versus private), and based on different organizational cultures. Previous studies have shown that these factors, among many can influence the way in which organizations learn from their incidents.

Third, and in this line, it is important to compare between learning from two main categories of ISRIs: 1) system failures and 2) security and privacy incidents. IS literature has shown some basic differences between these two categories of incidents. For instance, security incidents (P. Chen, Kataria, & Krishnan, 2011; Dhillon & Backhouse, 2001; Eriksson, 2001; Zafar & Clark, 2009) are often

intentionally created by attackers. This way, security incidents are more dynamic and emergent than system failures, since the attackers might change their strategies as they foresee how their targets (e.g., the affected organizations) might learn from those incidents.

Although our empirical setting allowed for studied both types of incidents, we did not intend to make such a specific comparison between the two categories of incidents. It was because the aim of the study was to openly explore the learning practices with regard to the general characteristics of ISRIs. However, in our analysis, we could capture the emergent nature of security incidents when organizations were considering the changing nature of incidents in their learning abstinence and in adopting transversal learning. Future studies might explore such differences through comparative studies. A proper design can be focusing on organizations that suffer from both types of incidents at the same time, then, comparing the process of learning between the two categories of incidents. Such comparative study can help identifying conditions in which some learning modes are more relevant than other ones.

Fourth, the study intended to examine organizational learning practices, namely practices that individuals and collectives adopt to avoid recurrence of incidents and reduce their impacts, from the view of the organizations. The study openly explored various learning practices and specifically sought to see how individual learning is then transformed into collective practices and is institutionalized into organizational elements. However, the technical nature of specialized technologies revealed a strong individual pattern across various learning practices. Hence, it showed that the *relevant* learning actors inside the organizations are limited to single or few experts, though they are interacting with other experts inside their providers. Although this pattern was captured through outsourced learning mode, the locus of study was on the central organization. Further studies

can look at this mode from an industry-level perspective. More specifically, it is interesting to know how a division of labor is shaped among IT companies and their clients in terms of learning practices, how IT companies learn across their clients, and how IT organizations act as specialized learning centers for other industries.

8.4 Implications for practice

8.4.1 Beyond traditional learning modes

As for managerial implications, the findings of this study show that there are ample learning opportunities for organizations beyond the two traditional modes. More specifically, learning efforts of the organizations do not confine to their practices during and right after each incident. The organizations can also engage in learning process when they learn transversally from a series of incidents, to materialize a specific solution. Furthermore, the findings suggest that transversal learning can provide more chances for organizations to focus on specific solutions, and engage in a process through which they build a local understanding of the concerned solution and put it into practice by developing political pressure and available socio-technical elements. This way, organizations can bridge their knowing-doing gap. Some learning initiatives that might not be identified and be feasible during and right after each incident, might be more justifiable and feasible when organizations reflect on a series of incidents and integrate political pressure and mobilize resources to put them into practice.

This suggestion becomes important particularly when we notice that the practitioner-based literature focuses on learning from incident as a stage that should be done right after incident handling, through the analysis of the causes of incident and drawing lessons. This suggestion, though might be useful in some situations that the incident is stable and repetitive, can be misleading if it is considered as the only or the best learning practice. Particularly, our empirical findings showed that in cases like ISRIs that incidents are changing, an overemphasis on post-incident analysis can make organizations waste their resources on drawing lessons that would become obsolete quickly. In addition, organizations might run the risk of applying obsolete lessons to future incidents.

8.4.2 Embracing new learning governances

In line with the previous implication, the study shows that learning process does not necessarily confine to *acquiring more knowledge* about the incidents and their causes, through either internal practices or external interactions. In fact, the very context of ISRIs showed that organizations might find it even more effective if they rely on learning practices that take place inside their providers. Furthermore, the organizations might intentionally avoid internalizing learning activities due to the technical nature of incidents that changes quickly. The dominant trend of outsourcing IT incident management (Jopson, 2013) resonates with the tendency among both IT and non-IT companies to outsource their learning practices as well.

However, this can bring about risks in situations that the relation between the organization and its provider does not create enough incentives for deep, timely learning by the providers. In this line, it is critical for managers to set up their contractual relations in a way that supports outsourced learning mode. In fact,

learning how to contract would be an agenda for managers when they realize that internalizing learning practices is not a viable solution anymore.

8.4.3 Neutralizing incidents through ticketing system

The empirical findings show that ticketing system can help organizations neutralize their incidents by transforming them into tickets. This way, organizations can prevent blame-game and other political and social tensions that can negatively influence the learning process. Although this practice seems to be prevalent in IT industry, organizations in other industries can adopt the same approach.

Ticketing system is different from some traditional documentation and filing systems that passively store information about the incident. Ticket can be used as a live, impersonal entity that objectifies incidents. Therefore, actors act on ticket, as an important task, rather than a sign of failure or defect. In addition, tickets reduce the efforts needed for documenting information during incident handling process. This is particularly important since organizational actors are often very busy during incident handling process. Hence, documentation of incident during incident handling through traditional systems is often incomplete and requires extra efforts.

8.4.4 Incident mining

The study showed that ticketing system is an established working tool in incident handling process. The organizations use it not only to organize their incident handling process, but also to keep a detailed, contextual history of their incident experiences. They even put efforts to make sure that *all* possibly useful information about incident handling process is stored in tickets.

Surprisingly, the study found that the organizations rarely went back to this rich source of information and analyzed them. Although they have explicitly announced their need for identifying patterns of incidents by analyzing a large number of accumulated tickets, they do not perform such analysis. Most of incident reporting activities are confined to incident counting for the sake of measuring the amount of hours spent on tickets of one type of another.

There seems to be several reasons for this situation. First, the managers and experts are often too busy to spare time for doing such transversal analysis on the tickets from various specialized domains. Second, for doing such analysis, they need to develop and implement analytical tools that help them search for meaningful patterns. Third, they need to spend a good amount of time in analyzing qualitatively the tickets, their differences and their similarities, in addition to plane overall patterns.

This shows an important opportunity for providing incident-mining services to the companies. Incident mining aims at analyzing both quantitatively and qualitatively the tickets of the organizations over a long period to identify meaningful patterns of incidents, analyze their causes, and suggest potential solutions. Incident mining provides the opportunity of identifying and analyzing patterns of incidents that are not easily detectable through reflections on single or a few incidents (like trying to see jungle beyond the daily observation of trees). Incident mining is not simply a large-scale data analysis. It is specialized in terms of incident elements, namely chains of events and impacts, the potential underlying causes, the effect of various solutions, and the patterns of secondary incidents that might appear subsequently.

Appendix 1: Interview protocol

All interviews were based on a customized interview protocol that was derived from the overall case study protocol. The general interview protocol that was used for customization is presented bellow.

Interview Protocol / Report

Date	Time	Place	Org. / Department		Case
Interviewee(s):					
Name & Family		**Position / Role**	**Tel Number**	**Email**	**Mobile**
Interviewer:					

1- Check points before the interview

- Thanks for having time.
- Introducing myself
- Why this study
- Who is the beneficiary
- History of work
- Why you?
- Confidentiality
- Recording?
- Your introduction
    - What is your specialty?
    - How many years and months in present position? _____years _____months
    - How many years experience with this company? _____years _____months
    - Your position
    - Your functions / tasks

2- Initiating the meeting

a. Introducing the topic and project (3 minutes presentation)
    i. Academic study on how organizations learn from incidents that happen to there is
    ii. We are not interested in the incident itself, but mainly in **how organizations** learned from it
    iii. Learning = any change in order to reduce the possibility or impact of such incident or other incidents in future
    iv. All data will be confidential, anonymous, and fully checked by you and the company's managers
        1. As we arranged, we record the session, only to be able to analyze the data accurately and focus better in the session on the questions
        2. All these records will be destroyed after the project and will be kept meanwhile fully confidential
        3. But at any point if you preferred that we pause recording, please just let us know

b. Identifying/Specifying the case

i. We are now focusing on the case [name of the case]
ii. We are interested in what **really happened** in this case, regardless of the formal procedures of the project
iii. We have studied the documents that you sent us about this case
1. [Saying a bit about the case to clarify what is exactly the case that are focusing on]

c. Introducing yourself in relation to this case
i. Before starting to ask our questions, do you want to say a bit about
1. What is your background?
2. What is your position and role in the company?
3. What is/was your role in this case?
ii. Before asking the questions, do you want to mention any general point?

3- Extracting the overall story of incident management

a. Exploring the **narrative story**
i. (let them talk freely)
ii. **Start**
1. When it started?
2. How did it start?
3. Why the company decided to do so?
iii. **Process**
1. What happened afterwards
iv. **End**
1. What is the situation now?
2. What are now in place in this regards?

4- The **learning process** (overall story)

a. What did the company after the incident
b. What did it afterwards?
c. Which kind of changes took place

d. **Mapping on the Timeline**

Events:

- Incidents
- (Dis)Approvals
- Decisions
- Bold actions
- Unexpected happenings
- Oppositions
- Changes

Actions

- (to emphasize more the intentional actions)
- By managers
- By employees
- By consultants
- By IT-experts

Actors (who were active (affecting / be affected) actors)

- Internal
  - Managers
  - Employees (users)
  - IT experts
  - Elites (key but informally important)
  - Departments
  - Other roles that might be defined during the project / case
- External

- Suppliers
- Customers
- Other actors

Artifacts
- Information systems
- Products
- Other technologies
- Facilities
- Layouts / locations / positions

Discourse
- Formal assertions
- Informal assertions
- Hearings (gossips)
- Wonders / doubts

Coalitions
- Main positions around the case
- Proponents (in favor of the change)
- Opponents (against the change)
- Reasons / rationales /
- Interactions

Conflicts
- Major cases of conflict
- Who were involved
- How it happened

Interventions
- What was not expectable to be done
- Someone who was not expectable to intervene

Inactions
- What was supposed to be done, but never happened
- What were decided and planned, but not executed
- What suddenly terminated or stopped

e. Dimensions of change in learning process
   i. **Systems / Technologies**
      1. What did they do to the systems?
      2. How did they change the existing technology
      3. Did they change anything about the
         a. Software?
         b. Hardware?
   ii. **Procedures / processes**
      1. Did any change happened to the processes?
   iii. **People**
      1. Who were involved?
      2. Who **were not** involved?
      3. Who did what?
      4. Any changes in people? (hiring, firing)?
      5. Any changes to the HR activities / policies?
   iv. **Structures**
      1. Any departmental changes?
      2. Creating any new team / committee / …?
      3. Merging?
      4. Removing?
   v. **Training**
      1. Any new training activities?

2. Internal vs. external?

vi. **R&D**
1. Any research project?
2. Any new exploration?

vii. **Communication**
1. How did they communicate about this issue?
2. During the incident?
3. After the incident?
4. External vs. internal?
5. To the top managers?
6. To the technical experts?
7. To those who **were not** involved?

viii. **Physical changes**
1. Any change in locations? (location of servers …)
2. Any change in hardware?
3. Any change in the arrangements and settings?
4. Any change in the design?

f. Any other changes / activities?
1.
2.

5- Specific Questions

- **[to be added here]**

6- Complementary resources

a. Other informants
i. ……..
ii. ……..
iii. ……..

b. Reports
i. ……..
ii. ……..
iii. ……..

c. Websites
i. ……..
ii. ……..
iii. ……..

d. Memos
i. ……..
ii. ……..
iii. ……..

e. Films
i. ……..
ii. ……..
iii. ……..

f. News
i. ……..
ii. ……..
iii. ……..

g. Other sources
i. ……..
ii. ……..
iii. ……..

7- To do List

a. …..……..
b. …..……..
c. …..……..

8- Post-Interview Actions

- Thanks for your time
- How to contact you
  - Email
  - Mobile
  - Fixed
- Time of next meeting
- How to follow up issues
  - directly
  - Secretary
  - colleague
- Thanks again,

9- Transcript of the Interview

10- Analysis after the session

a. Impressions
b. Contradictions
c. Disagreements (with others)
d. Supports for assertions

11- Appendixes

12- Links to Support Documents

References

Albrechtsen, E., & Hovden, J. 2010. Improving information security awareness and behaviour through dialogue, participation and collective reflection. an intervention study. ***Computers & Security***, 29(4): 432-445.

Amori, G. 2008. Preventing and responding to medical identity theft. ***Journal of Healthcare Risk Management***, 28(2): 33-42.

Anderson, R., & Moore, T. 2006. The economics of information security. ***Science***, 314(5799): 610-613.

Ang, K., Thong, J. Y. L., & Yap, C. 1997. ***IT implementation through the lens of organizational learning: A case study of insuror***. Atlanta, Georgia, USA: Association for Information Systems.

Argote, L., & Epple, D. 1990. Learning curves in manufacturing. ***Science***, 247(4945): 920-924.

Armenakis, A. A., Harris, S. G., & Mossholder, K. W. 1993. Creating readiness for organizational change. ***Human Relations***, 46(6): 681-681.

Bandyopadhyay, K., Mykytyn, P. P., & Mykytyn, K. 1999. A framework for integrated risk management in information technology. ***Management Decision***, 37(5): 437-445.

Baumard, P., & Starbuck, W. H. 2005. Learning from failures: Why it may not happen. ***Long Range Planning***, 38(3): 281-298.

Beck, T. E., & Plowman, D. A. 2009. Experiencing rare and unusual events richly: The role of middle managers in animating and guiding organizational interpretation. ***Organization Science***, 20(5): 909-924.

Berg, M. 2001. Implementing information systems in health care organizations: Myths and challenges. ***International Journal of Medical Informatics***, 64(2): 143-156.

Bharadwaj, A., Keil, M., & Mähring, M. 2009. Effects of information technology failures on the market value of firms. ***The Journal of Strategic Information Systems***, 18(2): 66-79.

Blackler, F. 1995. Knowledge, knowledge work and organizations: An overview and interpretation. ***Organization Studies***, 16(6): 1021-1046.

Blackler, F. 2000. Power, mastery and organizational learning. ***Journal of Management Studies***, 37(6): 833–851.

Boddie, J. 1987. The project post-mortem. ***Computerworld***, 7(December).

Borodzicz, E. P., & Haperen, K. V. 2002. Individual and group learning in crisis simulations. ***Journal of Contingencies & Crisis Management***, 10(3): 139-147.

Brandon, D. P., & Hollingshead, A. B. 2004. Transactive memory systems in organizations: Matching tasks, expertise, and people. ***Organization Science***, 15(6): 633-644.

Brown, A. D., & Starkey, K. 2000. Organizational identity and learning: A psychodynamic perspective. ***Academy of Management. The Academy of Management Review***, 25(1): 102-120.

Bulgurcu, B., Cavusoglu, H., & Benbasat, I. 2010. Information security policy compliance: An empirical study of rationality-based beliefs and information security awareness. ***MIS Quarterly***, 34(3): 523-A7.

Butler, B. S., & Gray, P. H. 2006. Reliability, mindfulness, and information systems. ***MIS Quarterly***, 30(2): 211-224.

Carlile, P. R. 2002. A pragmatic view of knowledge and boundaries: Boundary objects in new product development. ***Organization Science***, 13(4): 442–455.

Carroll, J. S., & Edmondson, A. C. 2002. Leading organisational learning in health care. ***Quality and Safety in Health Care***, 11(1): 51-56.

Carroll, J. S., & Hatakenaka, S. 2001. Driving organizational change in the midst of crisis. ***MIT Sloan Management Review***, 42(3): 70-79.

Carroll, J. S., Rudolph, J. W., & Hatakeneka, S. 2002. Learning from experience in high-hazard organizations. ***Research In Organizational Behaviour***, 24: 87–137.

Carroll, J. S. 1995. Incident reviews in high-hazard industries: Sense making and learning under ambiguity and accountability. ***Organization & Environment***, 9(2): 175-197.

Carroll, J. S., & Fahlbruch, B. 2011. "The gift of failure: New approaches to analyzing and learning from events and near-misses." honoring the contributions of bernhard wilpert. ***Safety Science***, 49(1): 1-4.

Cavusoglu, H., Mishra, B., & Raghunathan, S. 2005. The value of intrusion detection systems in information technology security architecture. ***Information Systems Research***, 16(1): 28–46.

Cha, H. S., Pingry, D. E., & Thatcher, M. E. 2008. Managing the knowledge supply chain: An organizational learning model of information technology offshore outsourcing. ***MIS Quarterly***, 32(2): 281-306.

Chen, P., Kataria, G., & Krishnan, R. 2011. Correlated failures, diversification, and information security risk management. ***MIS Quarterly***, 35(2): 387-422.

Chen, Y., Neil, K. E., Avery, A. J., Dewey, M. E., & Johnson, C. 2005. Prescribing errors and other problems reported by community pharmacists. ***Therapeutics and Clinical Risk Management***, 1(4): 333–342.

Christianson, M. K., Farkas, M. T., Sutcliffe, K. M., & Weick, K. E. 2009. Learning through rare events: Significant interruptions at the baltimore & ohio railroad museum. ***Organization Science***, 20(5): 846-860.

Cohen, M. D., & Bacdayan, P. 1994. Organizational routines are stored as procedural memory: Evidence from a laboratory study. ***Organization Science***, 5(4): 554-568.

Cooke, D. L. 2003. ***Learning from incidents***. Paper presented at In Proceedings of the 21st International Conference of the System Dynamics Society, New York.

Coopey, J., & Burgoyne, J. 2000. Politics and organizational learning. ***Journal of Management Studies***, 37(6): 870–85.

Corbin, J. M., & Strauss, A. 1990. Grounded theory research: Procedures, canons, and evaluative criteria. ***Qualitative Sociology***, 13(1): 3-21.

Cremonini, M., & Nizovtsev, D. 2009. Risks and benefits of signaling information system characteristics to strategic attackers. ***Journal of Management Information Systems***, 26(3): 241-274.

Crossan, M. M., Lane, H. W., & White, R. E. 1999. An organizational learning framework: From intuition to institution. ***Academy of Management Review***, 24(3): 522-537.

Culnan, M. J., & Williams, C. C. 2009. How ethics can enhance organizational privacy: Lessons from the CHOICEPOINT and TJX data breaches. ***MIS Quarterly***, 33(4): 673-687.

Cyert, R. M., & March, J. G. 1963. ***A behavioural theory of the firm***. New Jersey: Prentice Hall.

Daft, R. L. 1982. Bureaucratic versus non-bureaucratic structure and the process of innovation and change. In S. B. Bacharach (Ed.), ***Research in the sociology of organizations, 1***, pp. 129-166. Greenwich: JAI Press.

de Bruijne, M., & van Eeten, M. 2007. Systems that should have failed: Critical infrastructure protection in an institutionally fragmented environment. ***Journal of Contingencies & Crisis Management***, 15(1): 18-29.

Demchak, C. C. 1999. 'New security' in cyberspace: Emerging intersection between military and civilian contingencies. ***Journal of Contingencies & Crisis Management***, 7(4): 181-198.

Denrell, J. 2003. Vicarious learning, undersampling of failure, and the myths of management. ***Organization Science***, 14(3): 227–243.

Deverell, E., & Hansén, D. 2009. Learning from crises and major accidents: From post-crisis fantasy documents to actual learning in the heat of crisis. ***Journal of Contingencies & Crisis Management***, 17(3): 143-145.

Dhillon, G., & Backhouse, J. 2001. Current directions in IS security research: Towards socio-organisational perspectives. ***Information Systems Journal***, 11(2).

Easterby-Smith, M. 1997. Disciplines of organizational learning: Contributions and critiques. ***Human Relations***, 50(9): 1805-1113.

Edmondson, A. 1999. Psychological safety and learning behavior in work teams. ***Administrative Science Quarterly***, 44(2): 350-383.

Egan, M. J. 2007. Anticipating future vulnerability: Defining characteristics of increasingly critical infrastructure-like systems. ***Journal of Contingencies & Crisis Management***, 15(1): 4-17.

Elkjaer, B. 2004. Organizational learning the 'Third way'. ***Management Learning***, 35(4): 419-434.

Elliott, D., Smith, D., & McGuinnes, M. 2000. Exploring the failure to learn: Crises and the barriers to learning. ***Review of Business***, 21(3/4): 17-24.

Elliott, D. 2009. The failure of organizational learning from crisis – A matter of life and death? ***Journal of Contingencies & Crisis Management***, 17(3): 157-168.

Engeström, Y. 2000. Activity theory as a framework for analysing and redesigning work. ***Ergonomics***, 43(7): 960-974.

Epple, D., Argote, L., & Devadas, R. 1991. Organizational learning curves: A method for investigating intra-plant transfer of knowledge acquired through learning by doing. ***Organization Science***, 2(1): 58–70.

Eriksson, J. 2001. Cyberplagues, IT and security: Threat politics in the information age. ***Journal of Contingency and Crisis Management***, 9(4): 211–222.

Faia-Correia, M., Patriotta, G., Brigham, M., & Corbett, J. M. 1999. ***Organizational back ups: Reconfiguring technology in a telemediated environment***. Paper presented at 32nd Hawaii International Conference on System Sciences, Maui, Hawaii.

Fauchart, E. 2006. Moral hazard and the role of users in learning from accidents. ***Journal of Contingencies & Crisis Management***, 14(2): 97-106.

Fiol, C. M., & Lyles, M. A. 1985. Organizational learning. ***Academy of Management Review***, 10(4): 803-813.

Fitzgerald, G., & Russo, N. L. 2005. The turnaround of the london ambulance service computer-aided despatch system (LASCAD). ***European Journal of Information Systems***, 14(3): 244–257.

Fortune, J., & Peters, G. 2005. ***Information systems - achieving success by avoiding failure***. New York: John Wiley & Sons, Ltd.

Fritzon, Å, Ljungkvist, K., Boin, A., & Rhinard, M. 2007. Protecting europe's critical infrastructures: Problems and prospects. ***Journal of Contingencies & Crisis Management***, 15(1): 30-41.

Galbreth, M. R., & Shor, M. 2010. The impact of malicious agents on the enterprise software industry. ***MIS Quarterly***, 34(3): 595-A10.

Gherardi, S. 1999. Learning as problem-driven or learning in the face of mystery? ***Organization Studies***, 20(1): 101-123.

Gherardi, S., & Nicolini, D. 2000. The organizational learning of safety in communities of practice. ***Journal of Management Inquiry***, 9(7): 7-18.

Ginzberg, M. 1981. Early diagnosis of MIS implementation failure: Promising results and unanswered questions. ***Management Science***, 27(4): 459–478.

Gordon, L. A., Loeb, M. P., & Sohail, T. 2010. Market value of voluntary disclosures concerning information security. ***MIS Quarterly***, 34(3): 567-594.

Gorman, S. P., Schintler, L., Kulkarni, R., & Stough, R. 2004. The revenge of distance: Vulnerability analysis of critical information infrastructure. ***Journal of Contingencies & Crisis Management***, 12(2): 48-63.

Goulding, C. 2002. ***Grounded theory A practical guide for management, business and market researchers***. London: SAGE Publications.

Greenaway, K. E., & Chan, Y. E. 2005. Theoretical explanations of firms' information privacy behaviors. ***Journal of the Association for Information Systems***, 6(6): 171-198.

Gwillim, D., Dovey, K., & Wieder, B. 2005. The politics of post-implementation reviews. ***Information Systems Journal***, 15(4): 307-319.

Handley, K., Clark, T., Fincham, R., & Sturdy, A. 2007. Researching situated learning participation, identity and practices in Client—Consultant relationships. ***Management Learning***, 38(2): 173-191.

Hannan, M. T., & Freeman, J. 1984. Structural inertia and organizational change. ***American Sociological Review***, 49(2): 149-164.

Hargadon, A., & Fanelli, A. 2002. Action and possibility: Reconciling dual perspectives of knowledge in organizations. ***Organization Science***, 13(3, Knowledge, Knowing, and Organizations): 290-302.

Haunschild, P. R., & Sullivan, B. N. 2002. Learning from complexity: Effects of prior accidents and incidents on airlines' learning. ***Administrative Science Quarterly***, 47(4): 609-643.

Havelka, D., Rajkuma, T., & Serve, P. 2004. ***Early indicators of troubled IS development projects***. Paper presented at Proceedings of Americas Conference on Information Systems, New York.

Hedberg, B. 1981. How organizations learn and unlearn. In P. C. Nystrom & W. H. Starbuck (Ed.), ***Handbook of organizational design, volume 1***: 3-27. New York: Oxford University Press.

Hills, A. 2005. Insidious environments: Creeping dependencies and urban vulnerabilities. ***Journal of Contingencies & Crisis Management***, 13(1): 12-20.

Holmes, A., & Poulymenakou, A. 1995. ***Towards a conceptual framework for investigating IS failure***. Paper presented at Proceedings of the Third European Conference on Information Systems, Athens, Greece.

Huang, J., Makoju, E., Newell, S., & Galliers, R. D. 2003. Opportunities to learn from 'failure' with electronic commerce: A case study of electronic banking. ***Journal of Information Technology***, 18(1): 17-26.

Iacovou, C. L., & Dexter, A. S. 2005. Surviving IT project cancellations. ***Commun. ACM***, 48(4): 83-86.

Ingram, P. 2002. Interorganizational learning. In J. A. C. Baum & T. J. Rowley (Ed.), ***Companion to organizations***: 642–663. Malden, MA: Blackwell.

Irani, Z., Sharif, A. M., & Love, P. E. D. 2001. Transforming failure into success through organisational learning: An analysis of a manufacturing information system. ***European Journal of Information Systems***, 10(1): 55-66.

ISO/IEC27002. 2005. ***Information technology security techniques code of practice for information security management***ISO/IEC.

ITIL. 2012. ***Best management practice portfolio: Common glossary of terms and definitions***ITIL.

Iversen, J. H., Mathiassen, L., & Nielsen, P. A. 2004. Managing risk in software process improvement: An action research approach. ***MIS Quarterly***, 28(3): 395-433.

Jasanoff, S. 1994. ***Learning from disaster***. Philadelphia: University of Pennsylvania Press.

Johnston, A. C., & Warkentin, M. E. 2010. Fear appeals and information security behaviors: An empirical study. ***MIS Quarterly***, 34(3): 549-568.

Jopson, B. 2013. ***Subcontractors are chink in cyber armour***. New York: Financial Times.

Kallinikos, J. 2012. Form, function, and matter: Crossing the border of materiality. In P. M. Leonardi, B. A. Nardi, & J. Kallinikos (Ed.), ***Materiality and organizing: Social interaction in a technological world***: 67-87. Oxford: Oxford University Press.

Kanellis, P., Lycett, M., & Paul, R. J. 1999. Evaluating business information systems fit: From concept to practical application. ***European Journal of Information Systems***, 8(1): 65-76.

Kappelman, L. A., McKeeman, R., & Zhang, L. 2006. Early warning signs of it project failure: The dominant dozen. ***Information Systems Management***, 23(4): 31-36.

Kasi, V., Keil, M., Mathiassen, L., & Pedersen, K. 2008. The post mortem paradox: A delphi study of IT specialist perceptions. ***European Journal of Information Systems***, 17(1): 62-78.

Keil, M. 1995. Pulling the plug: Software project management and the problem of project escalation. ***MIS Quarterly***, 19(4): 421-447.

Kim, D. H. 1993. The link between individual and organizational learning. ***Sloan management review***, 35(1): 37-50.

Kim, J. I. Y., & Miner, A. S. 2007. Vicarious learning from the failures and near-failures of others: Evidence from the U.S. commercial banking industry. ***Academy of Management Journal***, 50(3): 687–714.

Kjaerland, M. 2006. A taxonomy and comparison of computer security incidents from the commercial and government sectors. ***Computers & Security***, 25(7): 522-538.

Kotter, J. P. 1995. Leading change: Why transformation efforts fail. ***Harvard Business Review***, 73(2): 59-67.

Kraatz, M. S. 1998. Learning by association? interorganizational networks and adaptation to environmental change. ***Academy of Management Journal***, 41(6): 621–643.

Kumar, K. 1990. Post implementation evaluation of computer-based information systems: Current practices. ***Communications of ACM***, 33(2): 203-212.

La Porte, T. R., & Consolini, P. 1991. Working in practice but not in theory: Theoretical challenges of `High reliability' organizations. ***Journal of Public Administration Research and Theory***, 1(1): 19-47.

Lam, A. 2000. Tacit knowledge, organizational learning and societal institutions: An integrated framework. ***Organization Studies***, 21(3): 487-513.

Lampel, J., Shamsie, J., & Shapira, Z. 2009. Experiencing the improbable: Rare events and organizational learning. ***Organization Science***, 20(5): 835-845.

Langley, A., Smallman, C., Tsoukas, H., & Van de Ven, Andrew H. 2013. Process studies of change in organization and management: Unveiling temporality, activity, and flow. ***Academy of Management Journal***, 56(1): 1-13.

Langley, A. 1999. Strategies for theorizing from process data. ***Academy of Management Review***, 24(4): 691-710.

LaPorte, T. R. 1996. High reliability organizations: Unlikely, demanding and at risk. ***Journal of Contingencies and Crisis Management***, 4(2): 60–71.

Lapre, M. A., A. S. Mukherjee, & Wassenhove, L. V. 2000. Behind the learning curve: Linking learning activities to waste reduction. ***Management Science***, 46(No.): 597-611.

Larsson, R., Bengtsson, L., Henriksson, K., & Sparks, J. 1998. The interorganizational learning dilemma: Collective knowledge development in strategic alliances. ***Organization Science***, 9(3): 285-305.

Lave, D., & Wagenr, E. 1991. ***Situated learning: Legitimate peripheral participation***. Cambridge University Press, New York, NY.: .

Lawrence, T. B., Mauws, M. K., Dyck, B., & Kleysen, R. F. 2005. The politics of organizational learning: Integrating power into the 4I framework. ***Academy of Management Review***, 30(1): 180–191.

Leonardi, P. M. 2010. Digital materiality? how artifacts without matter, matter. ***First Monday***, 15(6-7).

Leonardi, P. M., & Barley, S. R. 2008. Materiality and change: Challenges to building better theory about technology and organizing. ***Information and Organization***, 18(3): 159-176.

Lesca, N., & Caron-Fasan, M. 2008. Strategic scanning project failure and abandonment factors: Lessons learned. ***European Journal of Information Systems***, 17(4): 371–386.

Levina, N., & Vaast, E. 2008. Innovating or doing as told? status differences and overlapping boundaries in offshore collaboration. ***MIS Quarterly***, 32(2): 307-332.

Levitt, B., & March, J. G. 1995. Barnard and the intelligence of learning. In O. E. Williamson (Ed.), ***Organization theory : From chester barnard to the***

present and beyond (expanded edition ed.). New York: Oxford University Press.

Lex team. 2012. ***RBS: Systems down***Financial Times.

Liang, D. W., Moreland, R., & Argote, L. 1995. Group versus individual training and group performance: The mediating role of transactive memory. ***Personality and Social Psychology Bulletin***, 21(4): 384-393.

Loch, K. D., Carr, H. H., & Warkentin, M. E. 1992. Threats to information systems: Today's reality, yesterday's understanding. ***MIS Quarterly***, 16(2): 173-186.

Lyytinen, K., & Robey, D. 1999. Learning failure in information systems development. ***Information Systems Journal***, 9(2): 85-101.

Lyytinen, K. 1988. Expectation failure concept and systems analysts' view of information system failures: Results of an exploratory study. ***Information & Management***, 14(1): 45-56.

March, J. G., Sproull, L. S., & Tamuz, M. 1991. Learning from samples of one or fewer. ***Organization Science***, 2(1): 1–13.

March, J. G., & Olsen, J. P. 1975. The uncertainty of the past: Organizational learning under ambiguity. ***European Journal of Political Research***, 3(2): 147-171.

March, J. G., Olsen, J. P., & Christensen, S. 1979. ***Ambiguity and choice in organizations***Universitetsforlaget.

Martinez-Noya, A., Garcia-Canal, E., & Guillen, M. F. 2013. R&D outsourcing and the effectiveness of intangible investments: Is proprietary core knowledge walking out of the door? ***Journal of Management Studies***, 50(1): 67-91.

Maskell, P., Pedersen, T., Petersen, B., & Dick-Nielsen, J. 2007. Learning paths to offshore outsourcing: From cost reduction to knowledge seeking. ***Industry and Innovation***, 14(3): 239.

Mayer, K. J., & Argyres, N. S. 2004. Learning to contract: Evidence from the personal computer industry. ***Organization Science***, 15(4): 394-410.

McKenney, J. L., Mason, R. O., & Copeland, D. G. 1997. Bank of america: The crest and trough of technological leadership. ***MIS Quarterly***, 21(3): 321-353.

McMullen, A. 2010. ***Don't leave IT to the techies***.

Mellahi, K. 2005. The dynamics of boards of directors in failing organizations. ***Long Range Planning***, 38(3): 261-279.

Mento, A. J., Jones, R. M., & Dirndorfer, W. 2002. A change management process: Grounded in both theory and practice. ***Journal of Change Management***, 3(1): 45.

Miettinen, R., & Virkkunen, J. 2005. Epistemic objects, artefacts and organizational change. ***Organization***, 12(3): 437–456.

Minnesota. 2003. ***A guide to research ethics***UNIVERSITY OF MINNESOTA, CENTER FOR BIOETHICS.

Montealegre, R., & Keil, M. 2000. De-escalating information technology projects: Lessons from the denver international airport. ***MIS Quarterly***, 24(3): 417-447.

Moynihan, D. P. 2009. From intercrisis to intracrisis learning. ***Journal of Contingencies and Crisis Management***, 17(3): 189-198.

Muhren, W., Van Den Eede, G., & Van de Walle, B. 2007. ***Organizational learning for the incident management process: Lessons from high reliability organizations***. Paper presented at In Proceedings of the

Fifteenth European Conference on Information Systems, Österle H, Schelp J, Winter R eds.

Naot, Y. B., Lipshitz, R., & Popper, M. 2004. Discerning the quality of organizational learning. ***Management Learning***, 35(4): 451-472.

Nelson, R. R. 2005. Project retrospectives: Evaluating project success, failure, and everything in between. ***MIS Quarterly Executive***, 4(3): 361–372.

Nelson, R. R. 2007. IT project management: Infamous failures, classic mistakes, and best practices. ***MIS Quarterly Executive***, 9(2): 67-78.

Nicolini, D., Gherardi, S., & Yanow, D. 2003. Introduction: Toward a practice-based view of knowing and learning in organizations. In D. Nicolini, S. Gherardi, & D. Yanow (Ed.), ***Knowing in organizations: A practice-based approach***: 3-31Armonk, NY: ME Sharpe.

Nicolini, D., Mengis, J., & Swan, J. 2011. Understanding the role of objects in cross-disciplinary collaboration. ***Organization Science***.

Nicolini, D., Powell, J., Conville, P., & Martinez-Solano, L. 2008. Managing knowledge in the healthcare sector. A review. ***International Journal of Management Reviews***, 10(3): 245-263.

Nikula, R. E. 1999. Organisational learning within health care organisations. ***International Journal of Medical Informatics***, 56(1): 61-66.

Orb, A., Eisenhauer, L., & Wynaden, D. 2001. Ethics in qualitative research. ***Journal of Nursing Scholarship***, 33(1): 93-96.

Orlikowski, W. J. 1992. ***Learning from notes: Organizational issues in groupware implementation***. Paper presented at Proceedings of the 1992 ACM conference on Computer-supported cooperative work, Toronto, Ontario, Canada.

Orlikowski, W. J. 1996. Improvising organizational transformation over time: A situated change perspective. ***Information Systems Research***, 7(1): 63-92.

Orlikowski, W. J. 2002. Knowing in practice: Enacting a collective capability in distributed organizing. ***Organization Science***, 13(3): 249–273.

Orlikowski, W. J. 2010. The sociomateriality of organisational life: Considering technology in management research. ***Cambridge Journal of Economics***, 34(1): 125-141.

Pan, S. L., Pan, G. S. C., Newman, M., & Flynn, D. 2006. Escalation and de-escalation of commitment to information systems projects: Insights from a

project evaluation model. ***European Journal of Operational Research***, 173(3): 1139-1160.

Parker, C. F., Stern, E. K., Paglia, E., & Brown, C. 2009. Preventable catastrophe? the hurricane katrina disaster revisited. ***Journal of Contingencies and Crisis Management***, 17(4): 206-220.

Pearson, C. M., & Mitroff, I. I. 1993. From crisis prone to crisis prepared: A framework for crisis management. ***Academy of Management Executive***, 7(1): 48-59.

Perrow, C. 1999. ***Normal accidents: Living with high risk technologies***. New Jersey: Princeton University Press.

Pfeffer, J., & Sutton, R. 2000. ***The knowing-doing gap: How smart companies turn knowledge into action***. Boston: Harvard Business School Press.

Ponemon Institute. 2013. ***The post breach boom***Ponemon Institute LLC.

Ponemon, L. 2009. ***The state of privacy & data security compliance sponsored by sophos***Independently conducted by Ponemon Institute LLC.

Ponemon, L. 2010. ***Security in the trenches comparative study of IT practitioners and executives in the U.S. federal government***Sponsored by CA Independently conducted by Ponemon Institute LLC.

Poulymenakou, A., & Serafeimidis, V. 1997. The role of evaluation in dealing with information systems failure: Conceptual explorations. ***Failure and Lessons Learned in Information Technology Management***, 1: 167-177.

Rai, A., Maruping, L. M., & Venkatesh, V. 2009. Offshore information systems project success: The role of social embeddedness and cultural characteristics. ***MIS Quarterly***, 33(3): 617-A7.

Ramasubbu, N., Mithas, S., & Kemerer, C. F. 2008. Work dispersion, process-based learning, and offshore software development performance. ***MIS Quarterly***, 32(2): 437-458.

Rangel, J. L., & Friend, G. N. 1995. Confidentiality of health care information in the computer age: A litigator's perspective. ***Journal of Healthcare Risk Management***, 15(3): 2-10.

Rao, H. 2004. Institutional activism in the early american automobile industry. ***Journal of Business Venturing***, 19: 359-384.

Raymond Caron, J., Jarvenpaa, S. L., & Stoddard, D. B. 1994. Business reengineering at CIGNA corporation: Experiences and lessons learned from the first five years. ***MIS Quarterly***, 18(3): 233-250.

Rerup, C. 2009. Attentional triangulation: Learning from unexpected rare crises. ***Organization Science***, 20(5): 876-893.

Rivard, P. E., Rosen, A. K., & Carroll, J. S. 2006. Enhancing patient safety through organizational learning: Are patient safety indicators a step in the right direction? ***Health Service Research***, 41(4): 1633–1653.

Roberts, K. H. 1989. New challenges in organizational research: High reliability organizations. ***Organization & Environment***, 3(2): 111-125.

Roberts, K. H. 1990. Some characteristics of one type of high reliability organization. ***Organization Science***, 1(2): 160-176.

Robey, D., Boudreau, M., & Rose, G. M. 2000. Information technology and organizational learning: A review and assessment of research. ***Accounting, Management and Information Technologies***, 10(2): 125-155.

Rochlin, G. I., LaPorte, T. R., & Roberts, K. H. 1987. The self-designing high-reliability organization. ***Naval War College Review***, 40: 76–90.

Salaway, G. 1987. An organizational learning approach to information systems development. ***MIS Quarterly***, 11(2): 245-264.

Sarosa, S., & Zowghi, D. 2005. ***Recover from information system failure: An indonesian case study***. Paper presented at EMCIS2005, The American University Cairo, Cairo, Egypt.

Schuetz, A. 1953. Common-sense and scientific interpretation of human action. ***Philosophy and Phenomenological Research***, 14(1): 1-38.

Schulman, P., Roe, E., Van Eeten, M., & De Bruijne, M. 2004. High reliability and the management of critical infrastructures. ***Journal of Contingencies & Crisis Management***, 12(1): 14-28.

Scott, J. E. 2000. Facilitating interorganizational learning with information technology. ***J. Manage. Inf. Syst.***, 17(2): 81-113.

Scott, J. E., & Vessey, I. 2000. Implementing enterprise resource planning systems: The role of learning from failure. ***Information Systems Frontiers***, 2(2): 213-232.

Senge, P. M. 1990. ***Fifth discipline. the art& practice of the learning organization***. New York: Doubleday.

Shedden, P., Ahmad, A., & Ruighaver, A. B. 2010. ***Organisational learning and incident response: Promoting effective learning through the incident response process***. Paper presented at Proceedings of the 8th Australian Information Security Management Conference, Edith Cowan University, Perth Western Australia.

Shin, R. W., & Sung, D. 1995. Disaster recovery plans for computer system failures: An empirical study of the state of preparedness of american county government. ***Journal of Contingencies & Crisis Management***, 3(2): 91-102.

Shrivastava, P. 1988. Industrial crisis management: Learning from organizational failures. ***Journal of Management Studies***, 25(4): 283-284.

Silvia, G. 2001. From organizational learning to practice-based knowing. ***Human relations***, 54(1): 131-139.

Smith, D., & Elliott, D. 2007. Exploring the barriers to learning from crisis: Organizational learning and crisis. ***Management Learning***, 38(5): 519-538.

Smith, D., & Toft, B. 2005. Towards an organization with a memory: Exploring the organizational generation of adverse events in health care. ***Health Services Management Research***, 18(2): 124-140.

Smith, H. J. 1993. Privacy policies and practices: Inside the organizational maze. ***Communications of the ACM***, 36(12): 105-122.

Smith, S., Winchester, D. W., Bunker, D., & Jamieson, R. 2010. Circuits of power: A study of mandated compliance to an information systems security de jure standard in a government organization. ***MIS Quarterly***, 34(3): 463-486.

Solms, R. v., & Solms, B. v. 2004. From policies to culture. ***Computers & Security***, 23(4): 275-279.

Southon, G., Sauer, C., & Dampney, K. 1999. Lessons from A failed information systems initiatives: Issues for complex organisations. ***International Journal of Medical Informatics***, 55(1): 33-46.

Spears, J. L., & Barki, H. 2010. User participation in information systems security risk management. ***MIS Quarterly***, 34(3): 503-527.

Starbuck, W. H. 1983. Organizations as action generators. ***American Sociological Review***, 48: 91-102.

Starbuck, W. H., & Hedberg, B. L. T. 2001. How organizations learn from success and failure. In M. Dierkes, A. B. Antal, J. Child, & I. Nonaka (Ed.),

Handbook of organizational learning and knowledge: 327-350. Oxford: Oxford University Press.

Starbuck, W. H. 2009. Cognitive reactions to rare events: Perceptions, uncertainty, and learning. ***OrganizationScience***, 20(5): 925–937.

Stead, E., & Smallman, C. 1999. Understanding business failure: Learning and un-learning from industrial crises. ***Journal of Contingencies and Crisis Management***, 7(1): 1-18.

Stern, E. 2002. Crisis and learning: A conceptual balance sheet. ***Journal of Contingencies and Crisis Management***, 5(2): 69-86.

Storey, J., & Buchanan, D. 2008. Healthcare governance and organizational barriers to learning from mistakes. ***Journal of Health Organization and Management***, 22(6): 642-651.

Straub, D. W., & Welke, R. J. 1998. Coping with systems risk: Security planning models for management decision making. ***MIS Quarterly***, 22(4): 441-469.

Strom, J. 1993. Executives pushing IT panic button. ***IT Magazine***, 25(11): November 1993.

Taylor, P. 2012. ***Warnings of 'war' serve to focus minds***. May 31, 2012 3:06 pm: Fiancial Times.

Toft, B., & Reynolds, S. 1992. ***Learning from disasters***. London: Butterworth.

Tripsas, M., & Gavetti, G. 2000. Capabilities, cognition, and inertia: Evidence from digital imaging. ***Strategic Management Journal***, 21(10/11, Special Issue: The Evolution of Firm Capabilities): 1147-1161.

Turner, B. A. 1976. The organizational and interorganizational development of disasters. ***Administrative Science Quarterly***, 21(3): 378-397.

Turner, B. A. 1978. ***Man-made disasters***. London: Wykeham.

Turner, B. A. 1994a. Causes of disaster: Sloppy management. ***British Journal of Management***, 5(3): 215-219.

Turner, B. A. 1994b. Software and contingency: The text and vocabulary of system failure? ***Journal of Contingencies and Crisis Management***, 2(1): 31-38.

Van de Ven, A. H., & Poole, M. S. 1995. Explaining development and change in organizations. ***The Academy of Management review***, 20(3): 510-540.

Van Eeten, M., & Bauer, J. M. 2009. Emerging threats to internet security: Incentives, externalities and policy implications. ***Journal of Contingencies & Crisis Management***, 17(4): 221-232.

Wagenaar, P. 2009. Power and security in the information age: Investigating the role of the state in cyberspace, myriam dunn cavelty, victor mauer and sai felicia krishna-hensel (eds). ***Journal of Contingencies & Crisis Management***, 17(2): 142-142.

Wallace, C. 2003. Health care IT system failures: Risks, prevention and recovery. ***Journal of Healthcare Risk Management***, 23(4): 9-15.

Wastell, D. G. 1999. Learning dysfunctions in information systems development: Overcoming the social defenses with transitional objects. ***MIS Quarterly***, 23(4): 581-600.

Weick, K. E., & Westley, F. 1996. Organizational learning: Affirming an oxymoron. In S. R. Clegg, C. Hardy, & W. R. Nord (Ed.), ***Handbook of organization studies***: 440-458. Thousand Oaks, CA: Sage Publications.

Weick, K. E., & Roberts, K. H. 1993. Collective mind in organizations: Heedful interrelating on flight decks. ***Administrative Science Quarterly***, 38(September 1993): 357-381.

Weitzel, J. R., & Marchand, D. A. 1991. The US stock market crash of 1987: The role of information system malfunctions. In F. W. Horton & D. Lewis (Ed.), ***Great information disasters: Twelve prime examples of how information mismanagement led to human misery, political misfortune, and business failure***. London: Aslib.

Westland, J. C. 2000. Research report: Modeling the incidence of postrelease errors in software. ***Information Systems Research***, 11(3): 320-324.

Willison, R., & Backhouse, J. 2006. Opportunities for computer crime: Considering systems risk from a criminological perspective. ***European Journal of Information Systems***, 15(4): 403–414.

Wilson, M., & Howcroft, D. 2005. Power, politics and persuasion in IS evaluation: A focus on relevant social groups. ***The Journal of Strategic Information Systems***, 14(1): 17-43.

Wright, R. T., & Marett, K. 2010. The influence of experiential and dispositional factors in phishing: An empirical investigation of the deceived. ***Journal of Management Information Systems***, 27(1): 273-303.

Yin, R. K. 2002. ***Case study research: Design and methods***. Newbury Park: SAGE Publications.

Zafar, H., & Clark, J. G. 2009. Current state of information security research in IS. ***Communications of the Association for Information Systems***, 2009(24): 571-596.