Toward Viable Information Security Reporting Systems

F. Olav Sveen[1], J.M. Sarriegi[1], E. Rich[2], J.J. Gonzalez[3]

[1] Department of Industrial Management, Faculty of Technology, TECNUN, University of Navarra, Paseo de Manuel Lardizábal, 13, 20.018 Donostia-San Sebastián, Gipuzkoa, Spain
[2] Department of Information Technology Management, School of Business, University at Albany, BA 310, Albany, NY 12222 USA
[3] Agder University College, Faculty of engineering and science, Research Cell "Security and Quality and Organizations," Serviceboks 509, 4884 Grimstad, Norway (and Gjøvik University College, Norwegian Information Security laboratory, 2802 Gjøvik, Norway)
Email: fosveen@tecnun.es

Abstract

Reporting and resolution of information security incidents is the basis for continuous improvement of security through learning. Incidents have varying degrees of impact, financial risk and learning opportunity for the organization. This variability naturally leads to classification of information security incidents into low and high priority for review and action. However, this classification carries with it some insidious aspects. First, high priority incidents are more costly to mitigate and as a consequence also more "uncomfortable" to report. Reporters may face reprimands, ridicule, extra workload and various other recriminations. This favors reporting of low priority incident at the expense of important high priority incidents. Incentives tied to reporting, a common policy used to stimulate reporting, may reinforce the problem. In essence, reporters face incentives and disincentives based on effects on throughput but have limited knowledge of what is important or not to the organization's security. Second, if a highly successful incident reporting policy is developed, the organization may become victim of its own success, as a growing volume of reports put increasingly higher pressure on incident handling resources. Continuously hiring more personnel is unsustainable in the long run. Developing and continuously improving automated tools for incident response promises more leverage.

Keywords

Information Security, Reporting Systems, Security Management, Human Factors, Incidents.

1. Introduction

The oil industry on the Norwegian continental shelf is moving towards Integrated Operations, a new operating paradigm (Gonzalez et al., 2005). Previously isolated offshore platforms are now connected to shore by fiber optic cable, enabling new levels of connectivity for increased decision support and operational remote control. However, increased connectivity comes with increased security needs. Computer networks from platform to shore and the Internet that were once physically separated now only have logical barriers. Remote access may be exploited by computer attackers. Protection of these operations from computer attack is clearly important to achieving operational goals.

Safety in such an environment is highly dependent upon information and communication technology. Safety reporting systems have been mandated by the Norwegian government and have been in use for many years. The Norwegian Oil Industry Association has recently published Information Security Baseline Requirements for Process Control, Safety and Support ICT Systems (OLF, 2006), where information security reporting is recommended for member organizations. Owing to the connection between safety and security in the Integrated Operations regime, it is likely that the safety and security reporting systems will have many common features; the value of a shared perspective for safety and security has been recognized (Stoneburner, 2006), but has not been widely explored. We believe that such a conceptual link is overdue.

First, we outline some of the challenges facing these systems and review some of what is currently known about information security reporting. Second, we develop a conceptual System Dynamics simulation model of an organization's information security reporting system. System Dynamics is particularly well suited to complex, socio-technical systems. It views systems as governed by information and material delays, accumulations and feedback. Given the scarcity of material on information security reporting we adapt generic experiences from safety where necessary. We build upon previous modeling on safety reporting (Rich et al., 2006) and on computer security incident response teams (Wiik et al., 2005, Wiik, 2007).

2. Recent Incident Trends, Incident Classification and Reporting

In the past few years there has been a substantial increase in information security incidents. Data published by the CERT Coordination Center show a quasi-exponential increase in the amount of incidents. By 2003 they stopped reporting since the statistics no longer gave meaningful information in assessing the scope and impact of attacks (http://www.cert.org/stats/ cert_stats.html). Data for a typical CSIRT (Computer Security Incident Response Team), the DFN-CERT, 1999-2005 show that the trend of increasing incidents continues (Wiik et al., 2005, Wiik, 2007). Organizations today face a diverse range of threats with varying impacts. To combat incidents, organizations typically employ an incident handling team, either internal or external. For example, DFN-CERT (a non-profit company) is the incident handler for the much larger DFN (Deutsche Forschungsnetz, the German Research Network).

Increasing incident volume and other considerations (see below) force handling teams to prioritize incidents according to their risk. For example DFN-CERT uses nine categories of incident priority (Wiik, 2007), the highest prioritized being attacks on DFN's network infrastructure, as this threatens the whole network. Port scans on the other hand are less important and thus classified in a lower category. Other organizations may prioritize differently. For this paper it is not necessary to know exactly how handling organizations prioritize, only that there are some incidents that are considered more important than others. Hence we will in this paper restrict ourselves to two categories, high priority and low priority incidents. Although the perceived importance of an incident may vary depending on different agents, we here refer to important incidents as those who have a high impact and financial risk to the organization as a whole.

High priority incidents carry the greatest potential for harm to the organization; learning to mitigate them reduces future loss. A plausible assumption is therefore that high priority incidents carry the greatest potential for learning, i.e. to mitigate future incidents and fix current vulnerabilities. It is thus important that high priority incidents are reported and investigated.

Perceptions of the importance of an incident by the handling team will not always be shared by the staff affected by the problem. Some attacks are highly conspicuous and cannot go unreported. An example is denial of service attacks. Other important high priority incidents, such as successful social engineering attacks, may not be reported, even if the attack was successful and staff members recognize it afterwards. They may be compelled not to report because of embarrassment or fear of recrimination from management or colleagues. This may especially be a problem if the damage or potential for damage was considerable. Such high priority attacks may also bring with them an increased workload for the reporter who has to fill out forms and participate in investigations. In the face of economic performance pressures, reporting may be omitted.

Low priority attacks carry considerably less baggage. The damage from them is less and thus staff fear of recriminations should also be less. Reporting of such incidents can also to a large extent be automated (Wiik, 2007). For example firewalls may be set to automatically report port scans.

3. What we know about Information Security Reporting

Organizations that wish to be certified in the BS-7799/ISO-17799 standard are required to implement an information security reporting scheme (Calder and Watkins, 2005). Winkler (2005) strongly advises organizations to implement a security alert system. Schneier (2000) compares the state of security reporting to the success of air safety reporting systems and finds current practices in information security reporting lacking. Gonzalez (2005) views information security reporting as a quality improvement process that is essential to reduce incidents. Ernst and Young's Global Information Security Survey (2004) report that 56% of respondents have trained users to identify and report suspicious activities. So, we know that we should

have reporting systems for information security and that many organizations do have them. But, do we know if they actually work?

Wiant (2005) examined whether the presence of an official information security policy impacted incident reporting in American hospitals. He found that the presence of a written policy did not impact incident reporting. However, Wiant's study is too narrow to conclude that information security reporting does not work. There is arguably a lot more to incident reporting than just the presence of an official policy.

Wiik et al. (2004) studied the effectiveness of DFN-CERT, an external, coordinating CSIRT, and found that staff were overworked, as funding did not keep pace with growth in security incidents. The team was led into a capability trap. Working harder to cope with incidents stole resources away from development of time saving tools, leading over time to poorer incident response capabilities.

4. Safety and Information Security

In many cases safety and security is interrelated, as in eOperations. In such circumstances satisfactory safety relies on effective information security. Deliberate attacks or errors in ICT systems may cause serious accidents such as fires or explosions in production systems.

There are similarities between safety and information security reporting systems, both attempts to reduce risk by learning from incidents. We also find that similar factors affect the two. For example: Winkler (2005) outlines a series of social pressures that affect security incident reporting, e.g., bad relationship with superiors. Safety reporting is also subject to these kinds of social pressures (Johnson, 2003, Phimister et al., 2003). Furthermore, in safety, as in security, it is also common to sort incidents into high and low priority (Kjellén, 2000, Phimister et al., 2003). There are also some differences. In safety the incidents are usually unintended whereas in information security incidents are often caused by deliberate attackers.

5. Incident Reporting Causal Model

The causal structure for high and low priority incidents is essentially the same. The difference between high and low priority incidents lies in the resources assigned and the differing strengths of incentives and disincentives. To avoid repetition, the explanation of the causal model is limited to high priority incidents and the interaction between high priority and low priority incidents.

5.1 High Priority Incidents

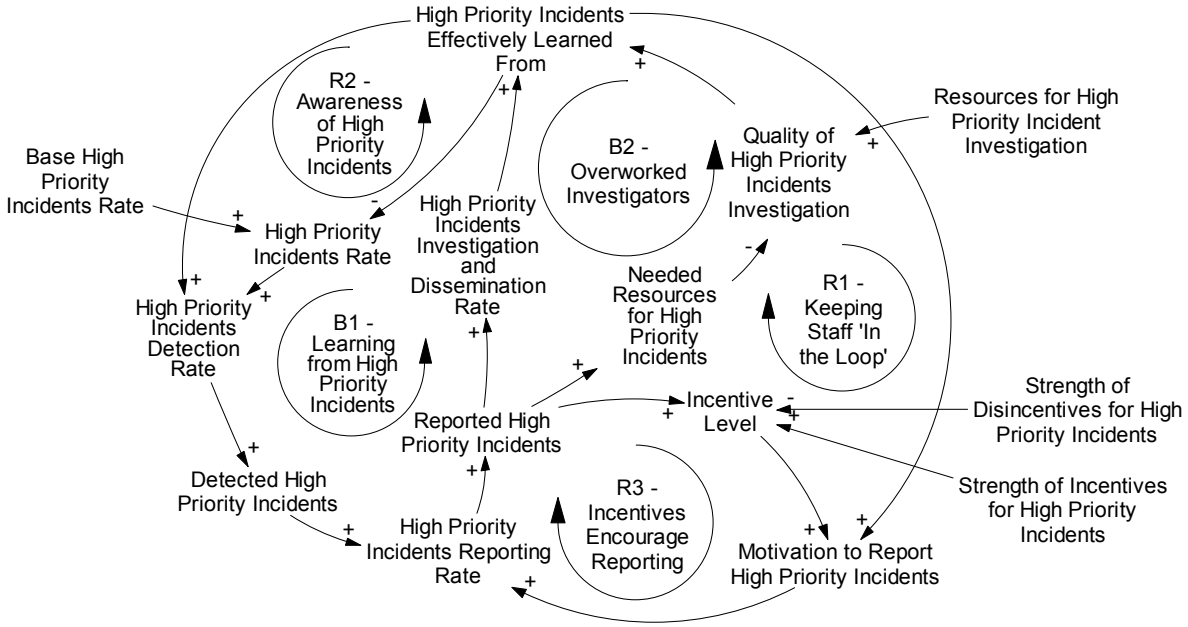

Figure 1: High Priority Incident Reporting Causal Structure

The sources of information security incidents are many. They may be software and hardware engineering errors, configuration errors or inadequate physical security which allows external attackers and malicious insiders to attack the system. Sometimes the source of a security incident may be simple mistakes. An example is the thousands of emails that are sent to wrong recipients every day. Some of those emails do contain sensitive information. To causally describe how incidents happen is beyond the scope of the model presented here. Our purpose is twofold: first, to describe how learning from incidents can prevent incident occurrence in the future, second, to describe some of the likely pitfalls an information security reporting system may run into. The source of high priority incidents is therefore modeled as an exogenous variable, *'Base High Priority Incident Rate'*.

The diagram above can be read as follows: The + and – signs at the arrow heads denote polarity. A causal link from A to B is positive if A adds to B, or if a change in A produces a change in B in the same direction. A causal link from A to B is negative if A subtracts from B, or if a change in A produces a change in B in the opposite direction (Sterman, 2000).

Reporting of incidents allows incidents to be investigated and learned from. This knowledge can be used to avoid such incidents in the future by putting into place technical and organizational countermeasures (Loop B1). We also assume that knowledge about previous incidents also improves the detection of future incidents (R2)

In System Dynamics terminology B1 is a balancing or goal seeking feedback loop. The loop attempts to balance the exogenous pressure of *'Base High Priority Incidents Rate'*. When *'High Priority Incidents Rate'* goes up more incidents are reported,

investigated and lessons disseminated. Ultimately, learning from those incidents reduces *'High Priority Incidents Rate'*.

R2 is a reinforcing feedback loop. Reinforcing loops work as either virtuous or vicious circles. If there is effective learning, more incidents will be detected, leading to more learning, a virtuous circle. Vice versa, if there is little effective learning, fewer incidents will be detected, which leads to less effective learning, a vicious circle.

Once detected, the incident can be reported. But detection does not imply that the incident will be reported. As previously mentioned, Winkler (2005) writes about social pressures that affect security incident reporting. We lack extensive evidence from studies of information security reporting systems, but we know from studies of safety reporting systems that there are many forces that reduce a person's willingness to report. One factor is that staff must see the usefulness of reporting. If not, they will be less likely to report incidents in the future (R1). It is therefore important that staff always receive feedback about what is happening with their report. Johnson (2003) termed this phenomenon "Keeping staff 'in the loop'." Low quality of incident investigations may lead staff to perceive reporting as less useful. Quality of investigations is described later in the paper.

A second but equally serious factor is the many forms of disincentives that may exist. Some may be punitive in nature. For instance, medical personnel often experience reprimands or other punitive measures if they make mistakes that endanger patient safety (Anderson and Webster, 2001). Lee and Weitzel (2005) describe how punitive culture in Taiwanese airlines causes pilots to avoid reporting of potentially dangerous near-miss situations in the air.

We do not consider outright punishment as the only disincentive present in incident reporting systems. In the face of economic pressures to produce, incident reporting may be seen as unnecessarily stealing time. For example the form to be filled out may be large and complicated (Nyssen et al., 2004), or the reporter may have to participate in lengthy investigations (Phimister et al., 2003).

In the model, recriminations for high priority incidents are assumed to be twice as strong as recriminations for low priority incidents. High priority incidents are by their nature more costly to the organization than low priority incidents. For example a configuration error that allows hackers to delete crucial information from a company server may cause reprimands or other forms of punishment for the technician who made the error. Whereas an incident such as a port scan most likely will not carry with it any form of punishment at all.

Just as there may be disincentives against reporting, organizations can also choose to reward incident reports, independent of their learning value. The literature indicates that incentives directly coupled to incident reporting do encourage reporting. (However, that rewarding zero-incident targets carries risk of underreporting (Kjellén, 2000).) We assume that incentives for low priority incident reports are more effective than incentives

for high priority incidents. The effect of incentives for high priority incidents have been modeled at half the strength of low priority incentives. Incentives and disincentives are represented in the model by the feedback loop R3.

To learn from an incident and avoid it in the future the incident's causes must be found (Johnson, 2003, Phimister et al., 2003). This implies that incidents have to be investigated, and thus, how that process is handled becomes of importance. The investigators must have the necessary competences (Phimister et al., 2003) and there must be enough time and people to do the job properly. The quality of an investigation has been modeled in a simplified manner as a function of the resources available and the workload. If the workload becomes higher than available resources, the investigative team will push investigations through faster at the expense of quality.

Sporadic emphasis and management fear of liability may hinder success in an incident reporting system (Phimister et al., 2003). In the model, management commitment is partially represented by incentives, disincentives and resources for investigation. Management also decides policy. We will see later that different policies can have widely different long term effects.

5.2 Interaction between High and Low Priority Incidents

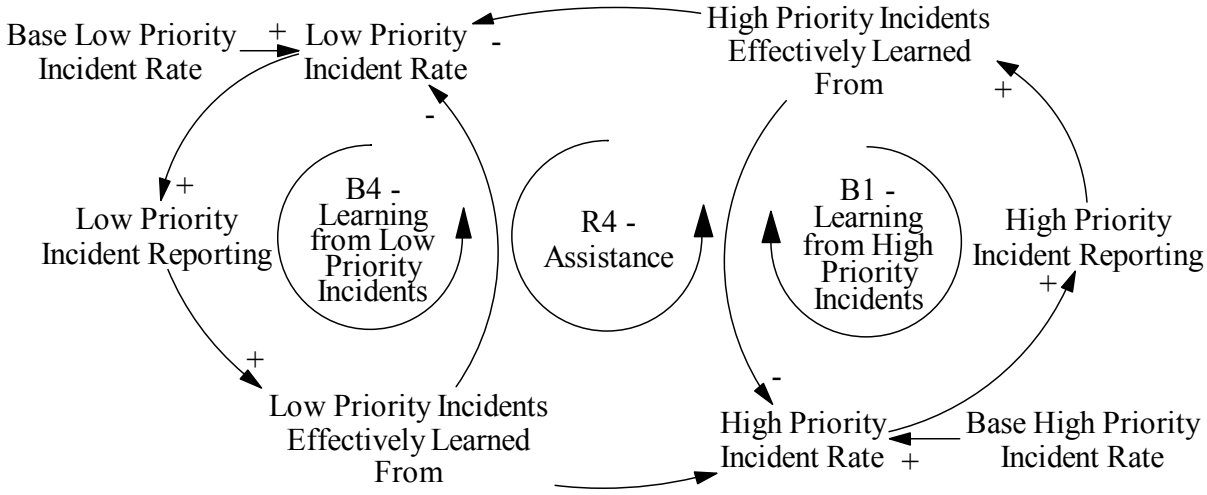

Figure 2: Interaction between high and low priority incident learning

We will now turn to describing the interactions between high and low priority incidents. We assume that lessons learned from high priority incidents will allow an organization to reduce not only high priority incidents but also low priority incidents (and vice versa). However, a crucial assumption in the model is that more can be learned from high priority incidents than from low priority incidents. In particular, learning from high priority incidents is more effective at assisting in the reduction of low priority incidents than learning from low priority incidents is in assisting with the reduction of high priority incidents (R4). Learning effects in the model have been modeled using power law learning curves (Zangwill and Kantor, 1998). For every

doubling of *'Low Priority Incidents Effectively Learned From'* there is a 5% reduction in *'High Priority Incident Rate'*. A doubling of *'High Priority Incidents Effectively Learned From'* reduces *'Low Priority Incident Rate'* by 15%.

Causal loop diagrams, as seen in the previous paragraphs, are useful for describing a system's feedback structure. However, they do not say anything about the relative strength of the feedback loops, or in other words, the system's behavior over time. To investigate behavior over time we next turn to simulation.

6. Simulation Runs

6.1 Assumptions and scenarios

Although DFN-CERT represents an instance where incident handling has been outsourced, we believe that an organization with internal incident handlers would face much the same challenges in terms of the development of high and low priority incidents. We have therefore modeled *'Base High Priority Incident Rate'* and *'Base Low Priority Incident Rate'* to correspond with the trends shown in the published material on DFN-CERT (Wiik et al., 2004). This is also in agreement with the statistics that CERT has published up to 2003. Figure 3: **Exogenous Base Incident Rates** has two scales to improve readability.

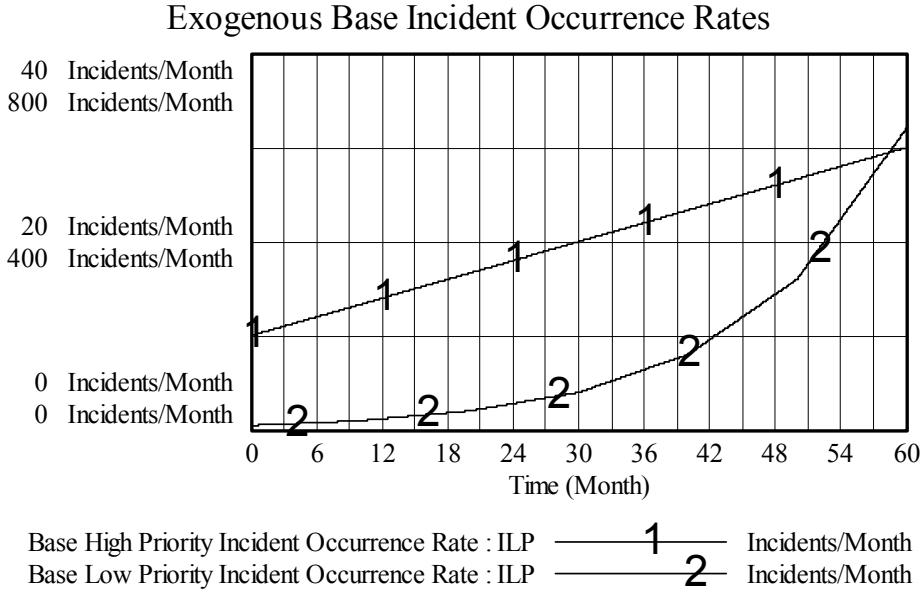

Figure 3: Exogenous Base Incident Rates (with different scales for high and low priority incidents)

We assume that high priority incidents take twice as long to investigate as low priority incidents. This is a conservative assumption, as low priority attacks will likely be well known and can therefore quickly be resolved. Capacity is 2.4 and 3.2 incidents / month respectively for high and low priority incidents (60-40% split in resources). Surplus resources in one category are fed into the other. Disincentives, if present, are assumed

to be stronger than incentives. The incident reporting system is introduced at time zero, with no prior reporting system in existence.

We ran a series of experiments to determine different policies' impact on our conceptual system. Different combinations of disincentives, incentives and limited resources were run.
Figure 4: **Table of Policies** shows an overview of the policy experiments.

Scenario	***Low Priority Incentives***	***High Priority Incentives***	***Low Priority Disincentives***	***High Priority Disincentives***	***Limited Resources***
1: ILP	X				
2: ILP LR	X				X
3: ILP DIHP LR	X			X	X
4: IHP DILP LR		X	X		X

Figure 4: Table of Policies

6.2 Incentives and Disincentives under Unlimited Resources

Scenario **1: ILP** assumes that incentives are only effective for low priority incidents. Unlimited incident handling resources lead to an improvement in incident rates compared to base incident occurrence rates. High priority incidents are stabilized with only a slight increase at the end of the simulation. Low priority incidents, although still growing significantly, are about 100 at the end of the simulation period, much less than the maximum base rate of 650.

6.3 Effect of Limited Resources

When limited resources are added to the **ILP** scenario a different behavior emerges (scenario **2: ILP LR**). High priority incident rate initially improves, but as increasing low priority reports put higher strains on incident handling teams, the system runs out of resources. Initial gains are reversed.

3: ILP DIHP LR adds strong recriminative culture around reporting of high priority incidents. Predictably, the simulation shows little gain in high priority incident rate and low priority incident rate. The situation is significantly worse compared to **2: ILP LR**. There are two effects that cause this result. First, there is an initial decrease in the rate of high priority incidents. But high priority incident reporting drops when staff experience disincentives and, as a result, learning slows down. Second, the absence of disincentives and presence of incentives triggers a flood of low priority incident reports – to the detriment of reporting of high priority incidents. The system has surplus resources for high priority incidents and these resources are fed into low

priority incidents. However, the surplus resources are not enough to compensate for the increase in reported low priority incidents. The quality of investigation of low priority incidents falls, causing a further drop in learning from incidents.

Graph for High Priority Incident Occurrence Rate

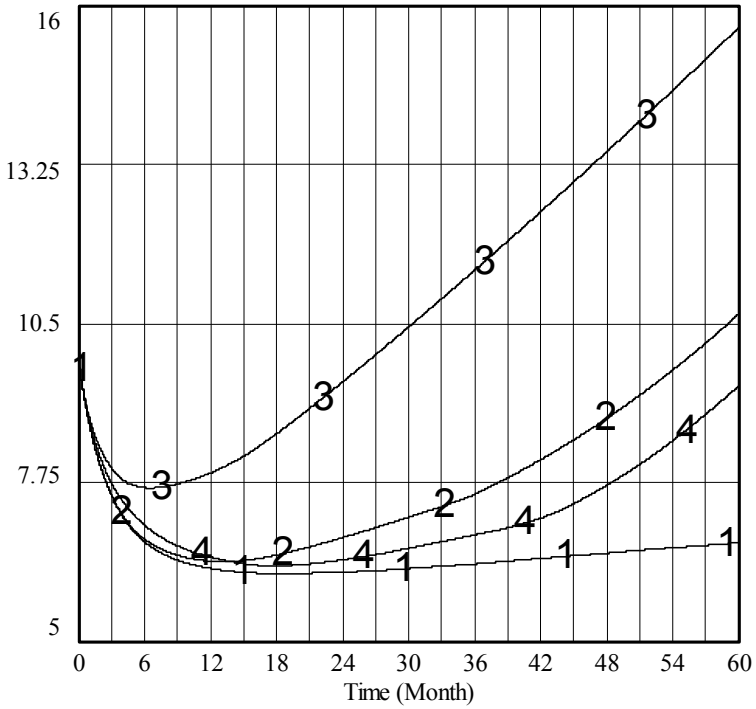

Graph for Low Priority Incident Occurrence Rate

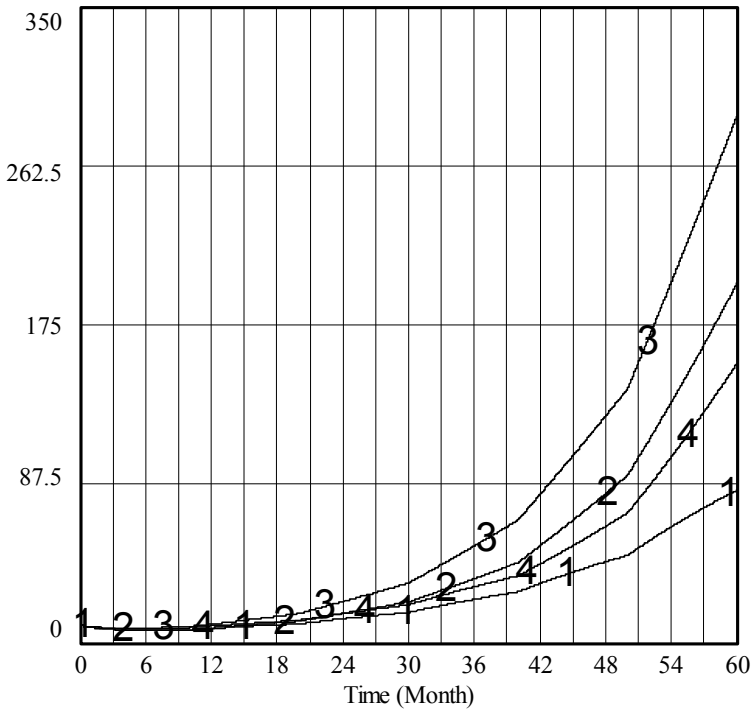

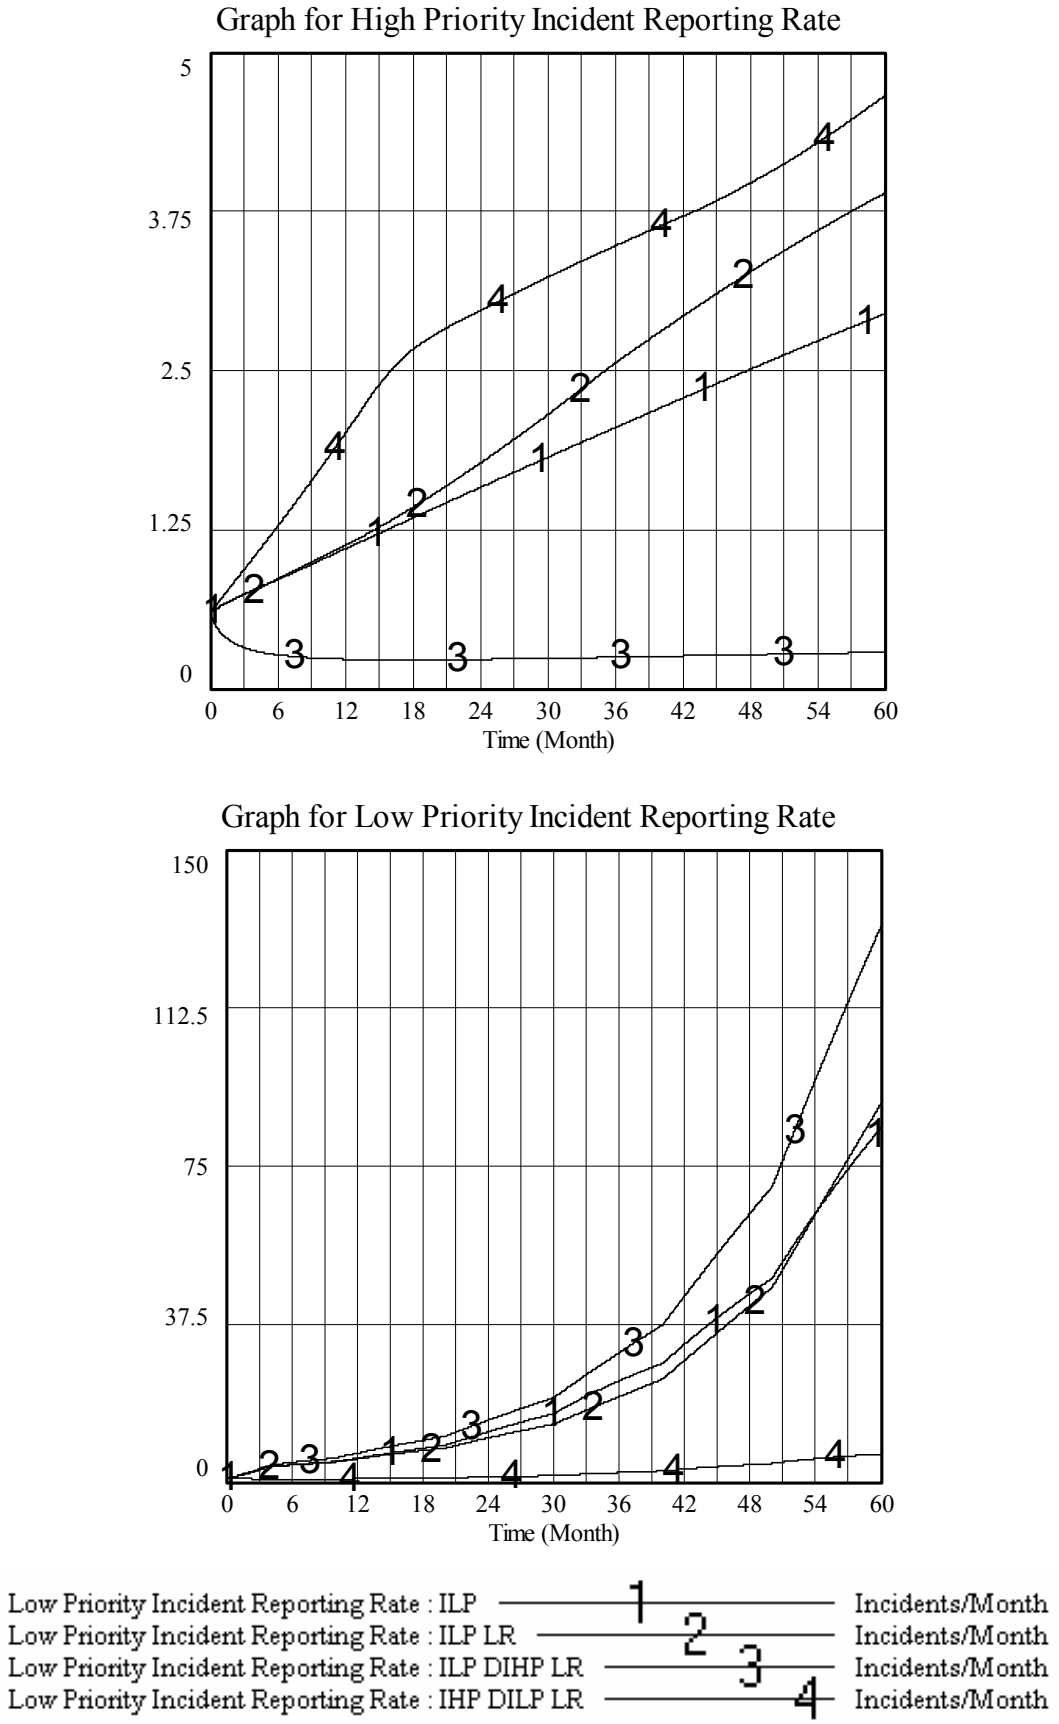

Figure 5: Low and High Priority Incident Occurrence and Reporting Rates

Discouraging low priority incident reports (**4: IHP DILP LR**) improves the situation significantly. The compounded effects of low priority incentives and high priority disincentives are removed. Although the initial gains are slightly lower than in **2: ILP LR**, the occurrence of both kinds of incidents drops below that of **2: ILP LR**. As low priority incident reports are discouraged, the system's resources last longer. However, this highly successful policy leads to a significant increase in high priority incident reports, straining the system's resources. High priority reports are also more labor intensive. The system eventually reaches the resource bottle neck. The problem is only postponed.

6.4 Victim of Own Success

In both the **ILP LR** and **IHP DILP LR** scenarios the system becomes a victim of its own success. The growth in incident reports overwhelms the investigative resources available. The first solution that springs to mind is to hire more people. However, it is unlikely that such a policy would be economically viable, given the exponential growth in low priority incidents. More leverage could be obtained by developing tools to handle low priority incidents. Their high frequency and relative low sophistication make them good candidates for automatic procedures.

7. Observations and Future Work

The introduction of critical and complex ICT infrastructure into critical infrastructure, such as oil and gas production, elevates the need to manage computer incidents to the best practice in safety management. This paper begins an examination of the challenges of incident management by looking at how incident reporting policies and incident handling efforts interact to produce successful or unsuccessful outcomes. We draw heavily from the experience of industrial safety reporting systems to structure the analysis.

Successful security incident handling requires effective reporting, mitigation, and learning. Within the organization, however, these activities are not always seen as beneficial to those tasked with reporting. The relative impact of reporting incentives and disincentives will affect their frequency, reliability, and learning value. When incident reporting is discouraged because of fear of recriminations or pressure to appear secure, the decision to notify the incident handlers rests in the hands of the reporters. These staff may have only a limited understanding of the effects of the incident on the organization's security. In contrast, a high volume of low-impact reports, stimulated by incentives, pushes the evaluation for review and mitigation onto the incident handling team. If the incident handling team becomes overwhelmed with reports of limited value, their effectiveness will drop, reducing their ability to identify areas for operational change and improvements for security.

Where then shall the problem detection burden be placed? An effective security policy in a setting of limited resources encourages reporting of high priority incidents and discourages reporting of low priority incidents. The ability to differentiate high

priority incidents from low can be cultivated in reporting teams, but such policies may not integrate with other social and economic demands. If this differentiating ability is not cultivated, and unimportant incidents flood the reporting process, incident handling teams must triage, focusing resources on the important and time-sensitive problems that are presented to them, or become overwhelmed and lose their effectiveness as an agent for learning and future protection. On the other hand, discouraging reporting reduces the effects of reporting on staff and on the reporting teams, but may mask hidden vulnerabilities. A collective view of the tradeoff between security and operational costs is needed to ensure appropriate management. This view may be stimulated by common risk analysis and goal-setting, though transfer of lessons among organizational units takes time to reach convergence (Martinez-Moyano et al., 2007).

While the parallels between safety reporting and computer incident reporting are clear, there are critical differences that must be considered in future work. Safety systems strive for high reliability (Cooke and Rohleder, 2006), but they do not face exponential growth in low priority incidents that have been observed in computer security incidents. From the incident reporter's perspective, a safety incident may have very visible and immediate risk of personal injury, where a computer incident's effects may be far removed from the immediate worksite. Finally, the continued growth of sophisticated and innovative attacks on computer systems, driven in part by the speed of change in the ICT environment, creates new opportunities for failure with each generation of technology. It may well not be possible to provide staff the knowledge needed to keep up with these changes, increasing our dependence on technology to separate high priority from low priority problems.

References

Anderson, D.J. & Webster, C.S. (2001) A system approach to the reduction of medication error on the hospital ward. *Journal of Advanced Nursing,* 35**,** 34-41.

Calder, A. & Watkins, S. (2005) *IT Governance,* London and Philadelphia, Kogan Page.

Cooke, D.L. & Rohleder, T.R. (2006) Learning from incidents: From normal accidents to high reliability. *System Dynamics Review,* 22.

EYGM (2004) Global Information Security Survey. Ernst & Young.

Gonzalez, J.J. (2005) *Towards a Cyber Security Reporting System - A Quality Improvement Process,* Berlin Heidelberg, Springer Verlag.

Gonzalez, J.J., Qian, Y., Sveen, F.O. & Rich, E. (2005) Helping Prevent Information Security Risks in the Transition to Integrated Operations. *Telektronikk,* 101**,** 29-37.

Johnson, C. (2003) *Failure in Safety Critical Systems: A Handbook of Incident and Accident Reporting*, Glasgow University Press.

Kjellen,U. (2000) *Prevention of Accidents Through Experience Feedback,* London and New York, Taylor & Francis.

Lee, P.I. & Weitzel, T.R. (2005) Air carrier safety and culture: An investigation of Taiwan's adaptation to western incident reporting programs. *Journal of Air Transportation,* 10.

Martinez-Moyano, I.J., Rich, E., Conrad, S., Andersen, D.F. & Stewart, T.R. (2007) A Behavioral Theory of Insider-Threat Risks: A System Dynamics Approach. *ACM Transactions on Modeling and Computer Simulation,* (forthcoming).

Nyssen, A. S., Aunac, S., Faymonville, M. E. & Lutte, I. (2004) Reporting systems in healthcare from a case-by-case experience to a general framework: An example in anaesthesia. *European Journal of Anaesthesiology***,** 757-765.

OLF (2006) OLF Guideline No. 104:
Information Security Baseline Requirements for Process Control, Safety and Support ICT Systems. Norwegian Oil Industry Association.

Phimister, J.R., Oktem,m U., Kleindorfer, P.R. & Kunruether, H. (2003) Near-miss incident management in the chemical process industry. *Risk Analysis,* 23**,** 445-459.

Rich, E., Sveen, F. O. & Jager, M. (2006) Overcoming Organizational Challenges to Secure Knowledge Management. *Secure Knowledge Management Workshop.* New York, US.

Schneier, B. (2000) *Secrets and Lies*, Wiley Computer Publishing.

Sterman, J. D. (2000) *Business Dynamics: Systems Thinking and Modeling for a Complex World*, Irwin McGraw-Hill.

Stoneburner, G. (2006) Toward a unified security/safety model. *IEEE Computer***,** 96-97.

Wiant, T. L. (2005) Information Security Policy's Impact on Reporting Security Incidents. *Computers & Security,* 24**,** 448-459.

Wiik, J. (2007) Dynamics of incident response effectiveness – A system dynamics approach. Bergen, University of Bergen.

Wiik, J., Gonzalez, J. J. & Kossakowski, K.-P. (2005) Limits to effectiveness of Computer Security Incident Response Teams (CSIRTs). *Twenty Third International Conference of the System Dynamics Society.* Boston, MA, The System Dynamics Society.

Winkler, I. (2005) *Spies Among Us*.

Zangwill, W. I. & Kantor, P. B. (1998) Towards a Theory of Continuous Improvement and the Learning Curve. *Management Science,* 44**,** 910-920.