Preserving a balanced CSIRT constituency

Johannes Wiik
Jose J. Gonzalez
Security and Quality in Organizations research cell
Faculty of engineering and science
University of Agder
4898 Grimstad, Norway
johannes.wiik@gmail.com, jose.j.gonzalez@uia.no

Pål I. Davidsen
Institute for geography
University of Bergen
5020 Bergen, Norway
pal.davidsen@geog.uib.no

Klaus-Peter Kossakowski
Software Engineering Institute Europe
Carnegie-Mellon University
60 322 Frankfurt, Germany
kpk@sei.cmu.org

Abstract

Since their inception Computer Security Incident Response Teams (CSIRTs) have been afflicted by chronic problems concerning workload, QoS and sustaining their constituency. We have cooperated with one of the oldest CSIRTs to model the most challenging issues. Low- and high-priority incident response cause different problems. In companion papers we dealt with the impact of the exponential growth of low-priority incidents on the CSIRT workload and the effect of high-priority incident response on the CSIRT workload and QoS. Here, we focus on a severe consequence of instabilities in high-priority incident response: problems to retain the internal constituency – i.e., the customer base or community who by its funding enable the existence of the CSIRT. Such an external constituency – people and organizations outside the internal constituency – that are provided with limited services, is unavoidable and even desirable, since security incidents often involve sites outside the internal constituency. But our model shows that the instabilities in high-priority incident reporting create an imbalance that – if it persists – could threaten the very existence of the CSIRT. Our model suggests that a management strategy that reduces the turnover of the most frequent reporters is much better than any attempt to attract a higher number of frequent reporters.

Keywords

Information Security, Incident Response, Incident Management, CERT, CSIRT, Risk Management, System Dynamics

Introduction

This is the third of three papers dealing with chronic problems in a coordinating Computer Security Incidents Response Team (CSIRT). The first paper, entitled "Chronic workload problems in Computer Security Incident Response Teams", and the second paper, entitled "Persistent instabilities in the high-priority incident workload of Computer Security Incident Response Teams", have both been submitted to this conference.

For the sake of brevity we refer to the first paper for most details about CSIRTs in general, the specific coordinating CSIRT and the modeling process, including data mining, model verification, validation and policy testing procedures. In the remainder of this paragraph we only summarize the basic aspects.

The first CSIRT was established in 1988 and quickly new CSIRTs appeared. CSIRTs have been afflicted by chronic problems since the beginning. In 1994 a study concluded that the existing CSIRTs were insufficiently funded, understaffed, and overworked (Smith 1994, §3.8.1). A thorough recent CERT/CC report on the state of the practice of CSIRTs documented that the problems persisted in most of the now approximately two hundred external coordinating CSIRTs over the world (Killcrece et al. 2003). The workload in CSIRTs is overwhelming, implying a wide range of internal problems, such as insufficient funding (van Wyk and Forno 2001; Killcrece et al. 2003), lack of management support, shortage of trained incident handling staff, poorly defined mission and authority, and lack of coordination mechanisms (Killcrece et al. 2003, §3.11). The problems persist – at least, there is no published material that indicates a change for the better. Information available from existing teams indicate no change either.

Among the various types of CSIRTs, coordinating CSIRTs usually have the broadest scope and most diverse constituency among the CSIRT organizational models. A coordinating CSIRT is typically located in a single location, coordinating and facilitating the handling of incidents across a variety of organizations in dispersed locations. We have collaborated with one of the oldest coordinating CSIRTs, and for that matter one of the oldest CSIRTs, to shed light on the causes of the chronic problems and to test management policies. The project was conducted as a PhD thesis and the project team consisted of the PhD candidate, two supervisors from different universities offering a system dynamics programme and the manager of the coordinating CSIRT. The CSIRT staff provided access

to the necessary data (mental, written and numerical). For more details of this we refer to the first paper in this series and for an extensive description of the whole project we refer to the PhD thesis (Wiik 2007).

Owing to its restricted availability, the PhD thesis was allowed to disclose the name of the actual CSIRT and to use the real data series. In contrast, this paper and its companion two papers withhold the CSIRT name and use sanitized data – in accordance with the practice in information security.

A crucial finding was that low- and high-priority incidents cause major problems of their own. Low-priority incidents typically are "standard" attacks – such as port scans, spams, fake mails, etc. They are nuisances, but not serious. Their impact on the CSIRT workload comes from their sheer number and their growth rate, their incidence rate typically doubling in number every year (Killcrece et al. 2003§3.8.1). The first paper in this series, "Chronic workload problems in Computer Security Incident Response Teams", presented a system dynamics model of low-priority incident and discussed policies to contain the problem.

In contrast to low-priority incidents, high-priority incidents are serious but relatively rare. Their impact on the CSIRT performance manifests itself through long-term instabilities in the workload and problems to retain a balanced constituency. The second paper in this series, "Persistent instabilities in the high-priority incident workload of Computer Security Incident Response Teams", presented a system dynamics model of high-priority incident response and its impact on the workload and discussed policies to contain the problem. This paper deals with another aspect of the high-priority incident response: difficulties in retaining a balanced constituency.

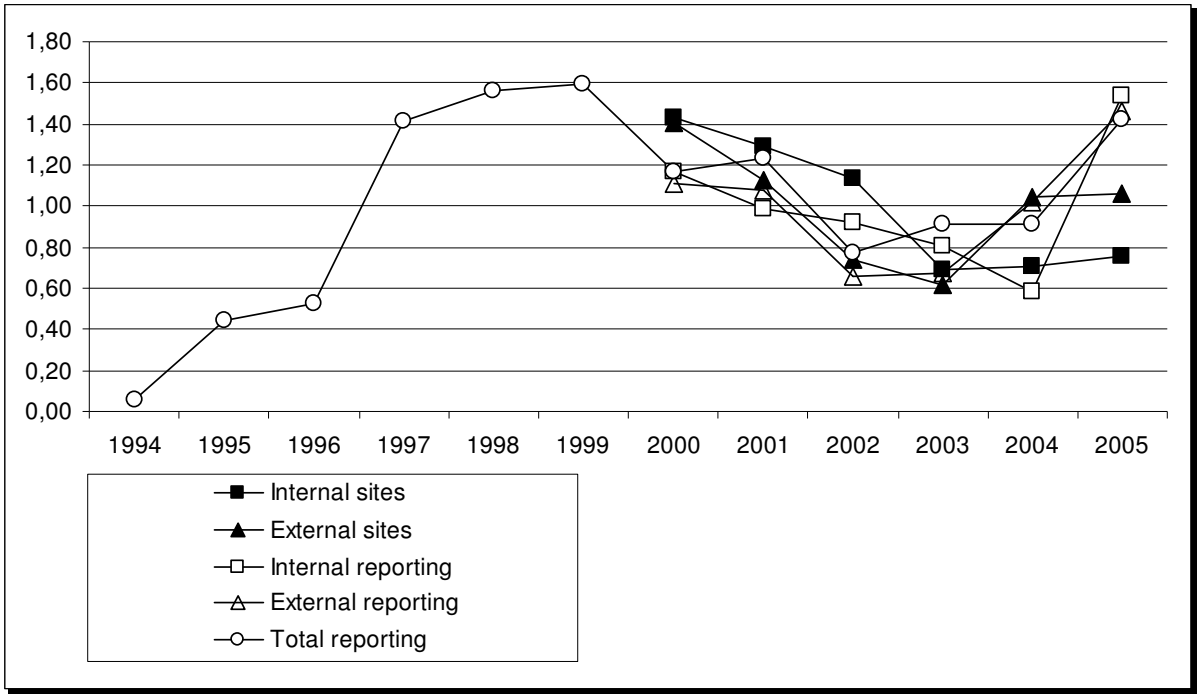

Figure 1 Historical variation relative to average for internal and external reporting sites as well as internal and external reporting in 2000-2005, and for total high-priority incident reporting 1994-2005

The retainment of a balance of internal and external sites

A CSIRT is dependent on the recognition of its constituency to remain effective (West-Brown et al. 2003, p. 17). A coordinating CSIRT faces a special challenge, since it has to deal with sites from both within and outside its constituency in order to coordinate the response to computer security incidents that impact sites in both groups (Killcrece et al. 2003, 2003; West-Brown et al. 2003). The global

nature of the Internet means that incidents can affect any site in the world. A coordinating CSIRT will be a focal point for receiving and responding to incidents that concern one or several sites within their constituency. Such an organization will warn internal as well as external sites involved in a potential incident, and help internal sites to analyze and mitigate the incident; and, viceversa, organizations will turn toward such coordinating CSIRTs to report attacks that are originating in its constituency. Thereby, from a CSIRT point of view, an incident is characterized by the information flow between the CSIRT and all the sites involved. To accomplish this task effectively, the CSIRT is dependent on receiving reports about such incidents, whether from inside or outside its constituency. A measure for the effectiveness of a CSIRT is the number of sites reporting incidents, and a high proportion of incidents reported from within the constituency is a good indicator for how well the CSIRT is recognized by its own constituency.

Our CSIRT case has experienced a large variation in reporting as well as in the number of reporting sites over time. Figure 1 provides a detailed view of the variation of five important variables over time. The variation in total reporting (the number of high-priority incidents reported) is shown relative to the average number of incident reports 1994-2005, the period across which such historical data are available. We only consider high-priority incidents, since most such incidents are reported manually by each site. We have ignored low-priority incidents, such as port scans and spam complaints, that are, for the most part, reported by automatic means, implying that reporting in this category is not primarily driven by the number of sites reporting (Wiik 2007). For the period 2000-2005, we also have data for the number of internal and external sites reporting as well as the amount of reporting from such sources. As before, the variation has been calculated relative to the average of these variables from 2000-2005.

An investigation of the number of internal and external sites reporting to the CSIRT for 2000-2005 (Figure 2) reveals that the numbers of internal and external sites follow a similar pattern, but there are some differences:

1. According to the data the proportion of external reporting sites in 2000 was higher than that of internal sites. This indicates that the growth of external reporting sites in preceding years (1994-1999) was higher as well.
2. According to the data external sites are more volatile than internal sites: the number of reporters both declines and grows faster than in the case of internal sites. A closer study of the data material suggests that site turnover as well as attraction is generally higher externally than internally (Wiik 2007).

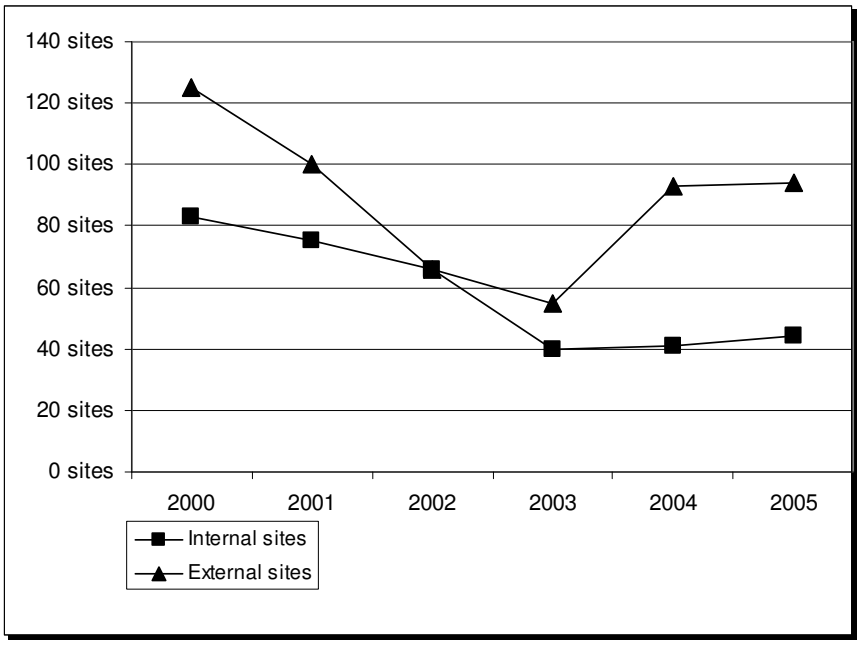

Figure 2 The number of sites varies with a similar pattern as the total number of incidents reported. However, the internal sites seems to vary more in absolute numbers, and an increasing gap in the number of reporting sites is emerging from 2003.

3. During 2003-2005 the number of external sites grows more than the corresponding number of internal sites.

Unless a CSIRT is used by its constituency directly (rather than indirectly), it will have a hard time to get funded. Hence, the increasing reliance on external reporting is a problem, since it indicates that the recognition of the CSIRT by its constituency is correspondingly weaker. It also means that external reporting fills up more of the incident response capacity. Furthermore, it is not a positive outlook for the CSIRT if this behavior indicating an increasing reliance on external reporting would continue in the future.

Based on the description and discussion about the problem behavior above, the following main questions will be addressed in this section:

1. What causes the changes to the relative proportions of internal and external reporting?
2. What is an expected long term development of the behavior shown in Figs. 1-2 (assuming no changes in CSIRT policies)?
3. What policies may preserve a more stable proportion of internal versus external reporting?

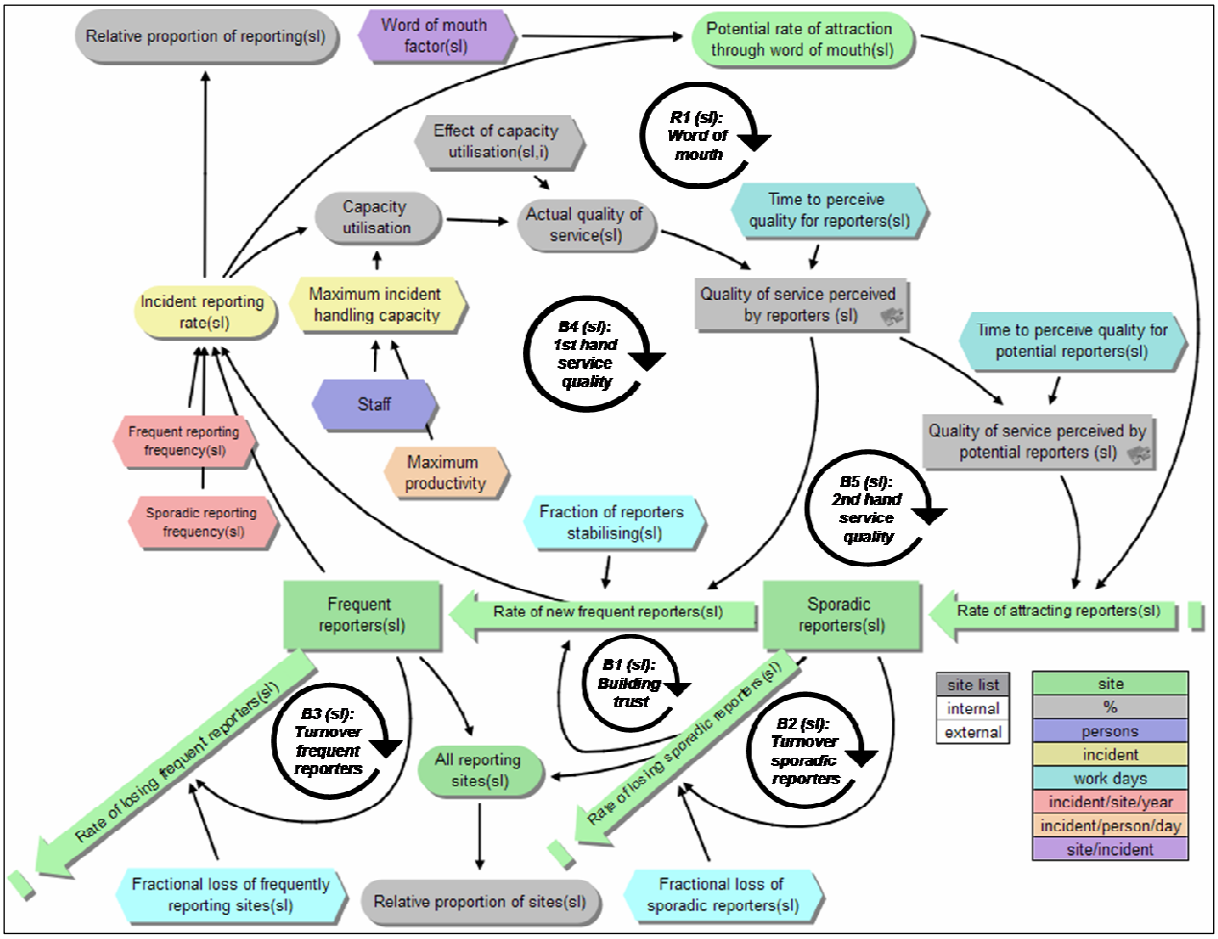

Figure 3 The handling capacity represents an internal limit to the growth of the CSIRTs workload and this forms several balancing feedback loops that may counteract growth of sites by slowing down the rate of attraction of new sites, B5 (sl) and the rate of new frequent reporters through B4 (sl)

To answer this questions we use an enhanced version (Figure 3) of the model employed in the second paper in this series (Wiik et al. 2009b). As for the two companion papers, the model has been created with the new object-oriented system dynamics modeling tool Smia, developed by Dynaplan. The figure shows the basic stock-and-flow structure for sites reporting to the CSIRT. Arrays are identified through the abbreviation "(sl)" after the variable name. The term "sl" is short for 'site list,' containing

two elements for internal and external sites respectively. Variables without such a term are scalars. Each array element is governed by a feedback structure resembling the scalar model of the second paper in this series (Wiik et al. 2009b). See the appendix for the model equations.

We run the simulation results from 1993 to 2015, i.e., ten years beyond the historical time frame.

The results from the base case are shown in the four graphs on Figure 4. The ensuing description of the behavior has been organized in time periods so that we may discuss the governing feedback structure for each period. Hence, the shift from one period to the next also indicates a relative change in dominating structure.

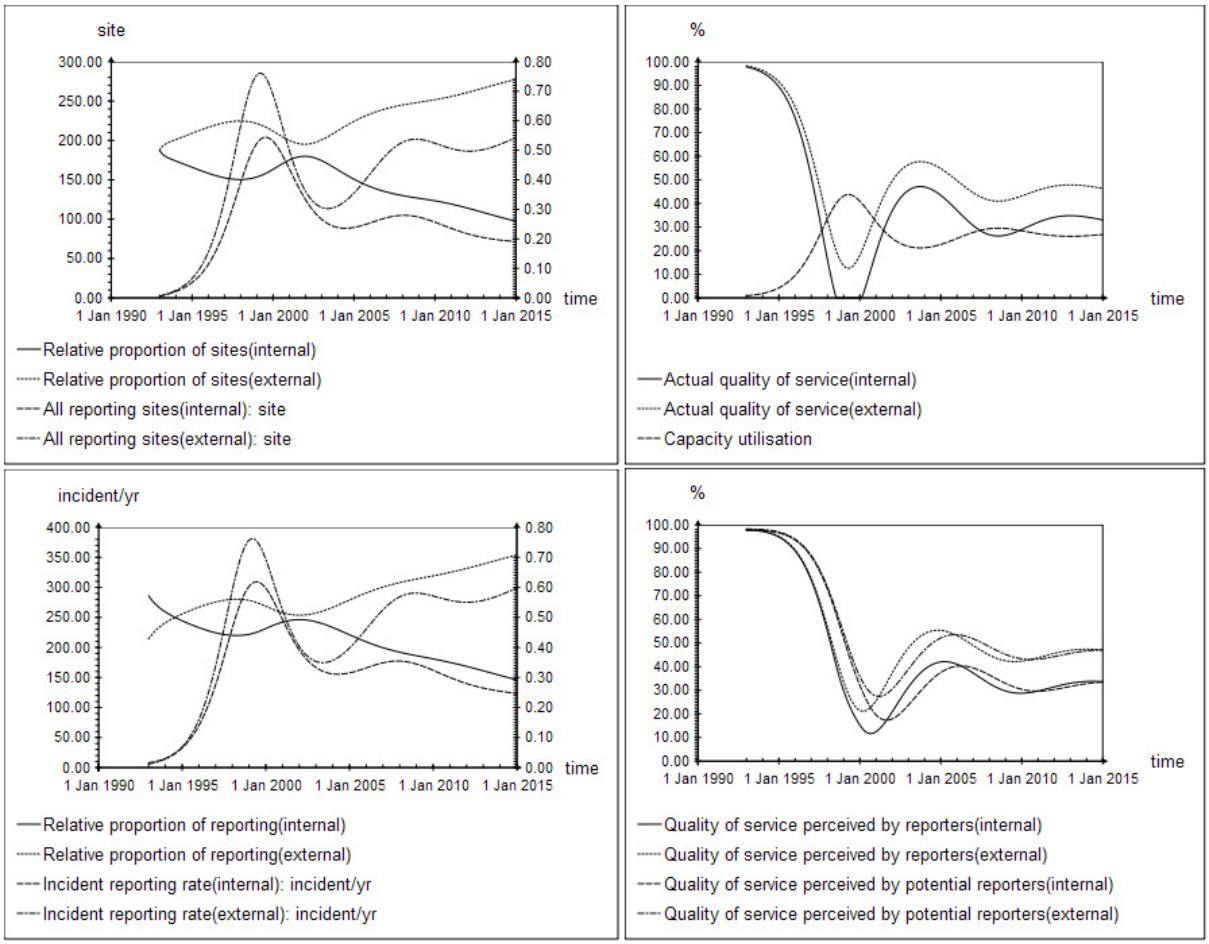

Figure 4 The base case scenario shows the behavior of key variables from 1993 to 2015, using the historical policies identified in the case

During 1993-1997, the number of sites, the rate of reporting and, thereby, the capacity utilization increase exponentially. The quality of service declines exponentially as the capacity utilization increases. The governing feedback generating the behavior in this period is "R1 (sl): Word of mouth". As the number of external sites grows more rapidly than the number of internal sites, we may assume that the external reinforcing loop is stronger than the internal loop. The main reason for this is the external community networks (closed forums/mailing lists, conferences, participation projects, common interests) that amplify word-of-mouth. This means that the relative proportion of external reporting sites and, consequently, external incident reporting rates increases as well.

During 1997-1999, the behavior patterns change from exponential growth (and decline) to goal-seeking growth (and decline). Several balancing feedback processes start to govern the behavior. First, the quality of service loops B4(sl) and B5(sl) become increasingly important and weaken the reinforcing feedback processes R1(sl) governing the behavior in the previous period. The lower

sensitivity of external sites to the quality of service means that the relative proportion of external sites and their reporting continues to increase in this period. Owing to the delay in perceiving the quality of service, the actual rock bottom in late 1999 is not noticed until more than one year later, indicating that capacity utilization has been severely overstretched: The number of reporting sites and, thereby, the incident workload had overshot any useful service that the CSIRT was able to provide with its limited workforce.

During 2000-2002, the reaction to the overshoot is a decline in the number of reporting sites. This behavior is governed by a combination of low perceived quality of services that weakens word-of–mouth, while site turnover B2 & B4 remains high. The decline in the number of sites is higher externally than internally, since external turnover is stronger. Furthermore, less sites are turned into stable reporters. As the turnover is higher among sporadic reporters (B2) than frequent reporters (B4), both internally and externally, the decline in the number of sites is even more prominent than it otherwise would have been. As the decline in sites and reporting rates continue, the pressure on the workforce decreases and the actual quality of service improves. The higher turnover among external sites means that the relative proportion of external reporting sites and the resulting reporting rates is brought closer to that of the internal proportion. This process continues until the quality of service again improves.

However, the same long time lags that caused an overshoot in the workload lead this time to an undershoot. Even though the quality of service improves for the better, first externally and then internally, it takes time before this has an impact on the word of mouth process generating renewed growth in 2003-2005, similar to the behavior experienced during the period 1993-1997. There are, though, some differences compared to the situation in 1993-1997. First, the number of external sites is slightly higher than the number of internal sites. This means that the word of mouth process is relatively stronger externally than internally at this point. Furthermore, due to the lower sensitivity to quality and a shorter perception times externally the growth process starts earlier externally than internally. Hence, the external reporting sites will tend to fill up more of the capacity before the internal sites starts to grow again. This behavior pattern is indeed present in the historical data on Figure 2.

For the remaining time 2005-2015, the oscillatory behavior pattern continues with damped oscillations. However, due to the stronger growth processes externally, the relative reliance on external reporting increases at the expense of internal reporting. Even though this evens out a little in downward cycles, it is not enough to compensate for the preceding build up of external reporting sites during growth. This is a serious problem since it means that the CSIRT is gradually less recognized by its own constituency even though all incidents reported do involve internal sites.

Alternative scenarios

We comment two aspects of the historical behavior pattern studied above.

1. The gradual shift towards increasing reliance on external reporting over time is not desirable. This does not mean that the CSIRT would like to obtain a 50-50 split. Rather, the CSIRT would like to find a stable mix between the two and at least preserve a stable number of internal sites to remain well-recognized by its constituency.
2. The oscillations are undesirable as a CSIRT would prefer to maintain a stable quality of service such that as many sites as possible are helped with a sufficient amount of effort to provide a useful service. The oscillations mean that the CSIRT overshoots and undershoots such a target. As shown in the second paper of this series (Wiik et al. 2009b), the main path to stabilization is finding ways to reduce perception times. For now, however, we are primarily concerned with the changing proportions of internal versus external reporting – the point being that we do not want to amplify the oscillations while trying to find a remedy for the drifting proportions in reporting.

We test two scenarios where we attempt to stabilize the relative proportion of internal and external sites over time while, at the same time, avoiding an amplification of the oscillatory behavior.

1. Increase word of mouth internally
2. Reduce site turnover internally

These two options are used to change the relative strength between internal and external reporting. The main difference seemed to be the sensitivity to parameter adjustment, as the following two scenarios will illustrate.

In the first scenario we increase the internal word of mouth factor by 25%. An interpretation of this measure in the real world would be to increase the marketing effort of the organisation in such a way that word of mouth between sites get stronger. In the second scenario we reduce the turnover of frequently reporting internal sites by 25%. An interpretation of this secenario in the real world would be to maintain stronger relations with existing frequent users of the service to assure continued reporting.

For simplicity we only show the relative proportion of internal and external reporting sites as well as the sum of reporting sites as an indication of side effects on the undesirable oscillatory behavior (p. 7). Figure 5 shows the results from these two scenarios compared to the base case.

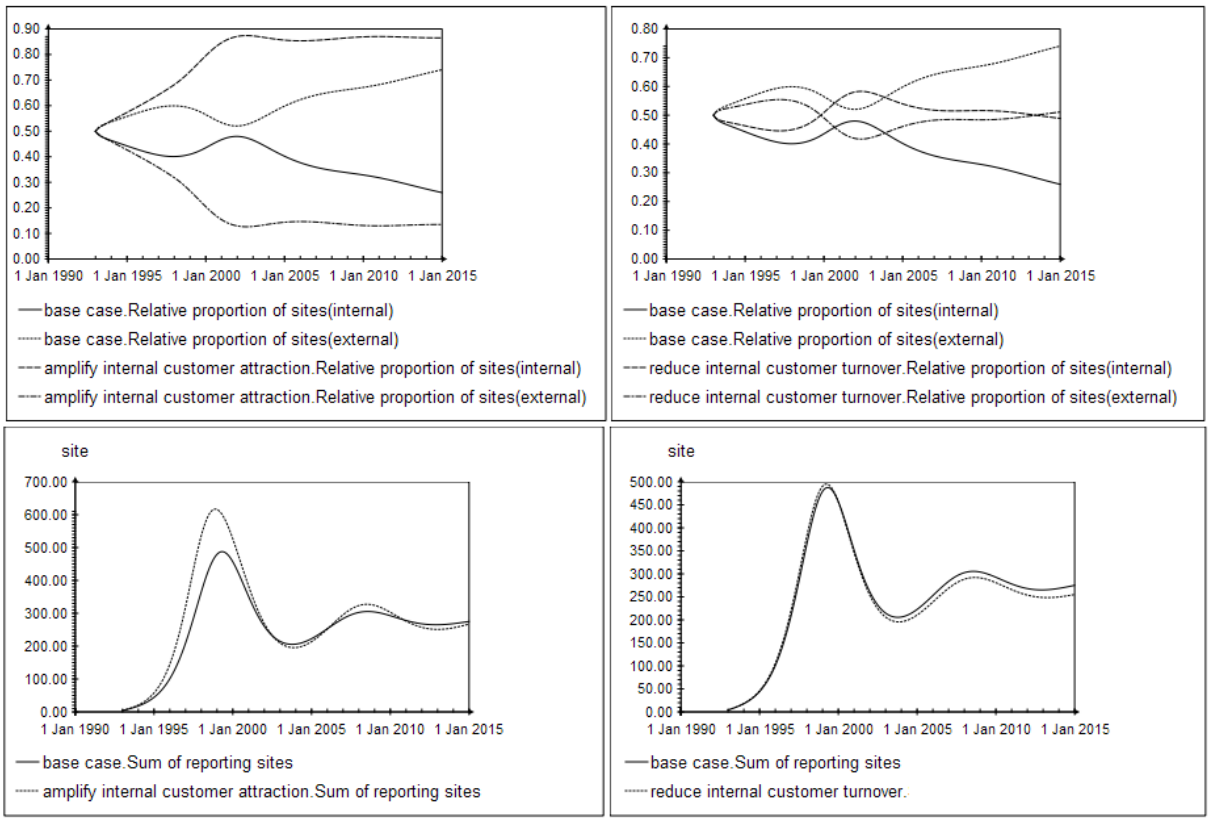

Figure 5 Simulation results comparing the base case to amplified attraction (left column) and preserving reporting sites (right column)

Both scenarios show that the relative proportion of internal reporting may be increased over time by any of the two measures. The increase of internal attraction seems to have a significant impact – the fraction of internal reporting sites rising close to 90% at the end of the simulation. The effect of preserving frequent reporting sites internally also has an impact, though less marked, so that the share of reporting sites varies from 40 to 60% percent during the simulation.

Amplifying internal customer attraction will tend to amplify the oscillatory behavior pattern of the overall number of sites reporting. The number of sites drives reporting. Hence, higher peaks of reporting sites will lead to peaks in reporting and subsequent overloads as well. As previously mentioned, this is an undesired effect. For the second scenario, where we reduced internal customer turnover, we observe a similar, but much smaller effect relative to the base case.

Discussion

The system dynamics model of our CSIRT case provides insights into the main mechanisms behind oscillations in high-priority incident reporting over time and the undesired gradual change towards an increased reliance on external reporting sites.

It is probably not possible to establish a 50-50 split of internal and external reporting over time. However, the model indicates two intervention points that can be utilized to preserve or even increase the proportion of internal constituent sites. Both interventions have similar effects, but the model seems much more sensitive to the policy of amplifying internal attraction than to the alternative of reducing turnover for frequent reporters. Since the side-effects seem to be larger for the most sensitive intervention point, reducing internal customer turnover seems to be the best approach for preserving a larger pool of internal reporters. Also, this last policy is probably easier to implement, and preserving the pool of internal frequent reporters is probably less costly than amplifying the attraction of new sites (Wiik 2007). Actually it is also less of a burden for the automation which will be needed to integrate new reports, if such interaction should be automated as the first paper in this series suggests (Wiik et al. 2009a). Hence, automation is best utilized for long-term and stable reporting processes which are based on internal frequent reporters.

References

Killcrece, Georgia, Klaus-Peter Kossakowski, Robin Ruefle, and Mark Zajicek. 2003. *Organizational Model for Computer Security Incident Response Teams*. Pittsburgh, PA, USA.

———. 2003. State of the Practice of Computer Security Incident Response Teams (CSIRTs). Pittsburgh, PA, USA: CMU/SEI.

Smith, Danny. 1994. Forming an Incident Response Team. Paper read at FIRST Annual Conference, at University of Queensland, Brisbane, Australia.

van Wyk, Kenneth R., and Richard Forno. 2001. *Incident Response*. Sebastopol, CA, USA: O'Reilly and Associates Inc.

West-Brown, Moria J., Don Stikvoort, Klaus-Peter Kossakowski, Georgia Killcrece, Robin Ruefle, and Mark Zajicek. 2003. *Handbook of Computer Security Incident Response Teams (CSIRTs)*. 2nd ed. Pittsburgh, PA, USA: CMU/SEI.

Wiik, Johannes. 2007. Dynamics of Incident Response Effectiveness -- A System Dynamics Approach, University of Bergen, Bergen.

Wiik, Johannes, Jose J Gonzalez, Pål I Davidsen, and Klaus-Peter Kossakowski. 2009a. Chronic workload problems in CSIRTs. Paper read at Twenty Seventh International Conference of the System Dynamics Society July, at Albuquerque, NM, USA.

———. 2009b. Persistent instabilities in the high-priority incident workload Computer Security Incident Response Teams. Paper read at Twenty Seventh International Conference of the System Dynamics Society July, at Albuquerque, NM, USA.

Appendix

```
vendor dynaplan
product smia
version 4
language enGB
def {
    submodel model {
        var 'Actual quality of service' = {sl=='site list'|'lookup linear'('Capacity
            utilisation','Effect of capacity utilisation'(sl,*),false)} as %
        var 'All reporting sites' = 'Frequent reporters'+'Sporadic reporters'
        var 'Capacity utilisation' = sum('Incident reporting rate')/'Maximum incident handling
            capacity' as %
        var 'Effect of capacity utilisation' = {sl='site list',i=0.0 to 0.5 step 0.1 |
        _ 1.0 ,3/4,2/4,1/4,0.0,0.0,
        _ 1.0 ,4/5,3/5,2/5,1/5,0.0}
        var 'Fraction of reporters stabilising' = {sl='site list'|15 ,15} as %/yr
        var 'Fractional loss of frequently reporting sites' = {sl='site list'|40,50} as %/yr
        var 'Fractional loss of sporadic reporters' = {sl='site list' |50,90} as %/yr
        var 'Frequent reporters' = stock {sl=='site list'|1 site} inflow 'Rate of new frequent
            reporters' outflow 'Rate of losing frequent reporters'
        var 'Frequent reporting frequency' = {sl='site list'|7,5 }as incident/site/yr
        var 'Incident reporting rate' = ('Frequent reporting frequency'*'Frequent
            reporters'+'Sporadic reporting frequency'*'Sporadic reporters')
        var 'Maximum incident handling capacity' = Staff*'Maximum productivity' as
            incident/yr
        var 'Maximum productivity' = 5 incidents/persons/'work days'
        var 'Potential rate of attraction through word of mouth' = 'Incident reporting rate'*'Word
            of mouth factor'
        var 'Quality of service perceived by potential reporters' = sd.'delay information'('Quality of
            service perceived by reporters','Time to perceive quality for potential reporters')
        var 'Quality of service perceived by reporters' = sd.'delay information'('Actual quality of
            service','Time to perceive quality for reporters')
        var 'Rate of attracting reporters' = flow 'Potential rate of attraction through word of
            mouth'*'Quality of service perceived by potential reporters'
        var 'Rate of losing frequent reporters' = flow 'Frequent reporters'*'Fractional loss of
            frequently reporting sites'
        var 'Rate of losing sporadic reporters' = flow 'Fractional loss of sporadic
            reporters'*'Sporadic reporters'
        var 'Rate of new frequent reporters' = flow 'Sporadic reporters'*'Fraction of reporters
            stabilising' *'Quality of service perceived by reporters'
        var 'Relative proportion of reporting' = 'Incident reporting rate'/sum('Incident reporting
            rate')
        var 'Relative proportion of sites' = 'All reporting sites'/sum('All reporting sites')
        var 'Sporadic reporters' = stock {sl=='site list'|1 site} outflow 'Rate of new frequent
            reporters' outflow 'Rate of losing sporadic reporters' inflow 'Rate of attracting reporters'
        var 'Sporadic reporting frequency' = {sl='site list'|1,1} as incident/site/yr
        var 'Sum of reporting sites' = sum('All reporting sites')
        var 'Time to perceive quality for potential reporters' = {sl='site list'|12,12 } as mth
        var 'Time to perceive quality for reporters' = {sl='site list'|18,12} as mth
```

```
        var 'Total reporting rate' = sum('Incident reporting rate')
        var 'Word of mouth factor' = {sl='site list'|1 ,1.3 } as site/incident
        unit 'incident/person/day' = incident/persons/dy
        unit 'incident/site/year' = incident/site/yr
        unit 'site/incident' = site/incident
        var Staff = 1.5 persons
        submodel globals {
            var 'exchange rate'&exchrate = {r ≤ text|"EUR" =>1}
            var 'game step' = 'time step'
            var 'report step' = 1 yr
            type 'role list' = [ ]
            var 'start future' = 'start time'
            type horizon = date(1993,Jan) to date(2015,Jan) step 1 mth
            submodel sd {
                component 'delay information' = template original
                component 'delay information' {
                    var 'change in ouput' = flow (input-output)/'delay time'
                    var 'delay time' = in:2 'time step'
                    var 'initial output' = optional in:3 input
                    var input = in:1 0.0
                    var output = return stock 'initial output' inflow 'change in ouput'
                }
            }
        }
    }
}
```