European Journal of Operational Research 216 (2012) 434–444

Contents lists available at SciVerse ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Innovative Applications of O.R.

# Information security trade-offs and optimal patching policies ☆

Christos Ioannidis [a,*], David Pym [b], Julian Williams [c]

[a] University of Bath, Department of Economics, Bath BA2 7AY, England, UK
[b] University of Aberdeen, School of Natural and Computing Sciences, King's College, Aberdeen AB24 3UE, Scotland, UK
[c] Business School, University of Aberdeen, King's College, Aberdeen AB24 3QY, Scotland, UK

ARTICLE INFO

*Article history:*
Received 15 April 2010
Accepted 26 May 2011
Available online 7 September 2011

*Keywords:*
Information security
Optimal policy
Risk reduction
Stochastic processes

ABSTRACT

We develop and simulate a basic mathematical model of the costly deployment of software patches in the presence of trade-offs between confidentiality and availability. The model incorporates representations of the key aspects of the system architecture, the managers' preferences, and the stochastic nature of the threat environment. Using the model, we compute the optimal frequencies for regular and irregular patching, for both networks and clients, for two example types of organization, military and financial. Such examples are characterized by their constellations of parameters. Military organizations, being relatively less cost-sensitive, tend to apply network patches upon their arrival. The relatively high cost of applying irregular client patches leads both types of organization to avoid deployment upon arrival.

## 1. Introduction

Software for computer networks, systems, and applications is typically subject to information security flaws, which, if exploited, may lead to substantial losses for the host organization. As vulnerabilities appear, software vendors periodically release patches in response. For large organizations, with tens or even hundreds of thousands of network devices, the deployment of patches is a costly exercise, impacting significantly on system availability, with consequences for properties of business processes, for credibility, and revenue. Failure to deploy a patch, however, risks exposing the host organization to exploitations of vulnerabilities.

The host organization's information security management team must make a judgement regarding the appropriate timing of the deployment of patches, in the light of the organization's policies. As in other areas of information security operations, decisions to deploy patches involve trade-offs between protecting the confidentiality of the system and maintaining its availability.

In recent years, there has been a good deal of research in the economics of information security. For example, Anderson et al. [1,3,2] have presented wide-ranging discussions of the issues, whilst Gordon and Loeb [21,22] have employed a microeconomic analysis of the costs and benefits of defences against given vulnerabilities.

Recent work by the present authors [25] has considered how to apply ideas and methods from utility theory and dynamic optimization to investment in information security. More specifically, by way of an illustration of the methodology, we have presented a dynamic model of trade-offs between confidentiality, availability, and investment in information security.[1] Our analysis has been motivated by those situations – including a detailed example in Beautement et al. [8], based on the use of USB memory sticks, as well as the work of Beres et al. [9,10] – in which the corruption of data (i.e., integrity) is a relatively minor issue. Here we intend confidentiality to refer to the system's state of protection against breaches of confidentiality, rather than the state of exposure of any particular data item. Similarly, we intend availability to refer to the system's readiness to supply its intended service. Our use of this example does not exclude the applicability of the methods we employ to situations in which integrity plays a major role. Such situations will be considered elsewhere. Moreover, in specific practical applications, it will typically be necessary to build richer models that incorporate more domain-specific details, such as the criticality of various information system components to business processes.

In this paper, we develop a model that is based on the confidentiality–availability trade-off model presented in [25], in which patch arrivals are interpreted as shocks to confidentiality and availability. We use this model to derive patching strategies in (large) organizations. We consider in detail the optimal timing of both client and network patching. As in [25], the purpose of this paper is to illustrate the application of modelling and reasoning

☆ This work has also been supported by the Cloud and Security Lab, HP Labs, Bristol.

* Corresponding author. Tel.: +44 1225383226.

*E-mail addresses:* c.ioannidis@bath.ac.uk (C. Ioannidis), d.j.pym@abdn.ac.uk (D. Pym), julian.williams@abdn.ac.uk (J. Williams).

[1] Modelling multiple trade-offs can be accommodated within the same methodology.

0377-2217/$ - see front matter © 2011 Elsevier B.V. All rights reserved.
doi:10.1016/j.ejor.2011.05.050

methods from utility theory and macroeconomics to questions in information security management that involve trade-offs between attributes of interest.

The remainder of the paper is organized as follows: in Section 2, we discuss related work; in Section 3, we develop our basic mathematical set-up, which draws upon methods from utility theory and dynamic optimization as deployed in economic and financial modelling, explaining how we model optimal responses in the presence of shocks to confidentiality and availability; in Section 4, we study an example of a model of the kind described in Section 3, to which to introduce a cost structure for implementing patches, and in which shocks to confidentiality and availability are given by the arrival (according to a Poisson process) of patches whose severity (or intensity) is drawn from a log-normal distribution; in Section 5, we describe the appropriate instance of the space of parameters employed in the model, and present and comment upon the results of our numerical simulations; finally, in Section 6, we explain our findings and their consequences for patching policies.

## 2. Related work

Beres et al. [9] provide a process model (written in the Demos 2 K modelling language [16], now superseded for our purposes by the Gnosis modelling language [13,19]) of vulnerability management policies in a large organization, and explore the effectiveness of both standard (or regular) patch-management and emergency (or irregular) escalation-based policies. In designing their model, which is concerned with patching clients, they examine the decision making process followed by the security operations managers of several large organizations, together with the different mitigation and patching measures that might be selected. They also identify external threat environment events that influence the type of mitigations that are deployed and the time at which they are deployed. They focus on examining the 'risk exposure window', defined as the time from public vulnerability disclosure to when an organization believes the risk is mitigated, as a measure of the effectiveness of these processes. By designing a model of these processes and running stochastic simulations, they examine the effectiveness of security operations processes and protection mechanisms based on external environment events.

In [23], it is postulated that both attackers and defenders behave strategically, whilst Beres et al. [9] seeks, treating attackers exogenously, to enable the decision-makers in IT security to predict the outcome of investment decisions or changes in policy in advance of putting them into effect. Their results show the impact of increasing the effectiveness of early mitigations and of speeding up patch deployment on reducing the risk exposure window.

The importance of timely patching in networks in the presence of externalities has been addressed by August and Tunca [6], in which they develop a set of incentive structures for users to implement effective patch management when their actions impact upon the welfare of other users. They show that software vendors can offer rewards to encourage timely patching when vulnerabilities occur in both proprietary software and freeware and, given the differential costs of patching to users, conclude ([6, p. 1718]) that 'a "one-size-fits-all" approach is unlikely to be an immediate remedy'.

The timing of vulnerability disclosures by vendors is modelled formally by Arora et al. [5], where it is shown that, with no regulation, the vendor releases a patch less frequently than is socially optimal.

The relationship between the release of patches by vendors and their implementation has been studied recently by Cavusoglu et al. [12]. They classify patching cycles into time-driven and event-driven. They show that social loss is minimized when vendor releases are synchronized with the time-driven cycles of the system operator. Their analysis is done in the context of single vendor and a single system operator. When such synchronization cannot be achieved because it is costly, the imposition of liability of the vendor for delayed release cannot achieve socially optimal disclosures.

When system operators employ a variety of applications, patch arrivals to the system operator will appear as random events, without apparent periodicity. In this paper, we capture patch arrivals as a Poisson process, and we decompose patching implementation into time-driven and event-driven incidents.

The 'Vulnerability Timeline', reproduced from Beres et al. [9]in Fig. 1, is a reference point for many studies of patching policies.

The timeline provides a detailed description of the sequence of events from the discovery of a vulnerability to the deployment of a patch. Similar accounts of such a timeline have been given by other authors; for example, Arbaugh et al. [4] and Schneier [34]. Of particular interest is Frei et al. [17], which illustrates the distributions and frequencies of vulnerabilities, using data from several large databases. The vulnerability arrival timeline given by Frei et al. is augmented in Beres et al. [9]. Arora et al. [5] calculate the socially optimal time interval between discovery and disclosure, $T_0$. August and Tunca [6] calculate, in the presence of externalities, the optimal period Patch Available to Patch Deployed, $T_3$–$T_5$, when vendors offer incentives to the system operator. Cavusoglu et al. [12], calculate the socially optimal window of exposure an decompose the patching process into time- and event-driven incidents.

Beattie et al. [7] explore the factors affecting the best time to apply patches so that organizations minimize disruptions caused by defective patches. Their results indicate that patching during the period of 10–30 days after first patch release date is the optimal time for minimizing the disruption caused by defective patches. The adoption of a real options methodology for determining choices of the appropriate technology in the presence of multiple sources of uncertainty and market entry has been addressed by Bobtcheff and Villeneuve [11] and Pennings and Lint [32]. In a similar vein, Gordon et al. [20] offer a framework around which decisions to delay the implementation of patches are integrated into a financial model that exploits deferment.

From the arguments discussed above, it is apparent that the timing of patch deployment matters because their deployment

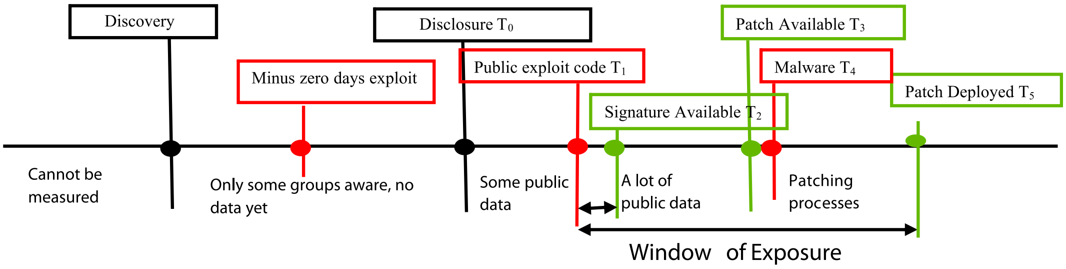

Fig. 1. Vulnerability timeline: the sequencing of events in this timeline is not fixed; the aim is to illustrate the various stages in the vulnerability life cycle (Beres et al. [9]).