Medical Informatics in a United and Healthy Europe
K.-P. Adlassnig et al. (Eds.)
IOS Press, 2009

doi:10.3233/978-1-60750-044-5-710

Threats to Information Security of Real-Time Disease Surveillance Systems

Eva HENRIKSEN [a,1], Monika A. JOHANSEN [a], Anders BAARDSGAARD [b], Johan G. BELLIKA [a,c]
[a] *Norwegian Centre for Telemedicine, University Hospital of North Norway, Tromsø, Norway*
[b] *The Norwegian Health Network, Tromsø, Norway*
[c] *Department of Computer Science, University of Tromsø, Norway*

Abstract. This paper presents the main results from a qualitative risk assessment of information security aspects for a new real-time disease surveillance approach in general, and for the Snow surveillance system in particular. All possible security threats and acceptable solutions, and the implications these solutions had to the design of the system, were discussed. Approximately 30 threats were identified. None of these got an unacceptable high risk level originally, but two got medium risk level, of which one was concluded to be unacceptable after further investigation. Of the remaining low risk threats, some have severe consequence, thus requiring particular assessment. Since it is very important to identify and solve all security threats before real-time solutions can be used in a wide scale, additional investigations are needed.

Keywords. disease surveillance, risk assessment, information security, privacy

1. Introduction

The idea of real-time disease surveillance becomes more and more relevant these days as current systems struggle with access to timely data. It is in this regard that the Norwegian Centre for Telemedicine (NST) and partners are investigating a real-time peer-to-peer disease surveillance solution. As a first phase in the development of the surveillance solution, the "Snow Agent System" (Snow) was created, to automatically extract anonymous data from a variety of health service providers' sources in a defined geographic area [1, 2], including the general practitioners' (GPs) electronic health record (EHR) systems, hospital systems and laboratories. Data shall be extracted in real-time from operational systems with maximum security requirements, and be processed and presented in real-time. A detected outbreak can be communicated by multicast or single messages, using instant messaging technology. This totally new surveillance approach introduces new security threats. In the future, most surveillance systems are expected to be based on similar real-time solutions. It is therefore important to identify the new security threats, and investigate how to solve these challenges. This paper presents the main results from a qualitative risk assessment of information security aspects of this new real-time disease surveillance approach.

[1] Corresponding Author: Eva Henriksen, Norwegian Centre for Telemedicine, University Hospital of North Norway, N-9038 Tromsø, Norway; E-mail: eva.henriksen@telemed.no.

2. Methods

Security challenges were analysed with regard to the intended environment and use in Snow as the case. Possible security threats and acceptable solutions were discussed, together with implications these solutions had to the design of the system. New design implications resulted in identification of new security threats. This process was accomplished in weekly semi-structured brainstorming sessions over a period of two months. The working group included the risk assessment leader, the chief engineer from the Norwegian Health Network's infrastructure department, the architect behind the Snow concept, and three project participants with more general technical background. Medical personnel were consulted when needed, to sort out questions related to the operational EHR systems at the GP offices.

NST has developed the risk assessment methodology in use, based on the Australian and New Zealand standard for risk management [3]. This methodology corresponds very well to the new ISO standard 27005 for information security risk management [4]. The risk assessment process consists of five main steps: 1) Context identification; a description of the analysed system and its environment, 2) Threat identification; identify what could possibly happen, 3) Impact and probability analysis; a consideration of consequences of the threats and the likelihood that these may occur, 4) Calculation of risk value for each threat as product of consequence and likelihood, and 5) Proposal of risk-reduction treatment for all threats with non-acceptable risk level.

As part of the preparations for the risk assessment, definitions of qualitative values for **likelihood** *(Low, Moderate, High, Very high)* and **consequence** *(Small, Moderate, Severe, Catastrophic)* were discussed and agreed. Acceptance criteria were also discussed. **Risk** values *(Low, Medium, High)* as product of likelihood and consequence, were defined, and illustrated in a two-dimensional risk matrix. Threats obtaining a high risk level are not acceptable, threats with medium risk level have to be further investigated and low risk threats may all be acceptable. (Definition details can be found at *www.telemed.no/opensource/snow,* as an Appendix under the title of this paper.)

During the brainstorming sessions, identified threats were documented in a threat table. Threats were afterwards grouped with respect to information security aspects. Each threat was given a unique identifier, starting with a character indicating the group it belongs to: confidentiality (c), integrity (i), availability (a), or quality (q). In addition, some threats were grouped as general (g). Within each group consecutive numbering was used. Threats were then analysed with respect to likelihood and consequence and placed in the corresponding cell of the risk matrix, Table 1.

3. Results

Nearly 30 threats or unwanted incidents were identified in the risk assessment. As shown in the risk matrix (Table 1), none of these threats got an unacceptable *high* risk level. Two threats got *medium* risk level, these were investigated separately to see if they could be accepted or not. In our case the conclusion was that one of them (c1) was *unacceptable*, while the other (a3b) could be accepted. Ten of the *low* risk threats were analysed to have *severe* consequence. These threats are in principle acceptable, but because of their severe consequence they should be observed to see if their likelihood, and thus their risk level, increases. The two threats with medium risk are discussed in the following, together with the ten low risk threats with severe consequence.

Table 1. Risk matrix

Consequence / Likelihood	Small	Moderate	Severe	Catastrophic
Low	a7a	a2, a3a, a4, a5, a6b, a7b, i2, i3a, i3b	g2, c2a, c2b, c3, c4, c5, a1a, a1b, i1a, i1b	
Moderate	a6a		c1	
High		a3b		
Very high				

c1: Sensitive (person identifiable) information is extracted from the EHR and presented by the surveillance system. The legal baseline is that person identifiable health information is sensitive, and it is therefore important to ensure the anonymity of the information. The challenge is to avoid that information *indirectly* identifies a person. If this is not especially handled by the system, the likelihood is assumed to be more than *low* (at least *moderate*). According to our acceptance criteria, this threat is then unacceptable: "It is not acceptable that likelihood is higher than *low* for unauthorised persons to get access to sensitive data." The *consequence* is analysed to be *severe*: Revealing this kind of sensitive information is a violation of law which could result in penalty or fine, and it would cause a serious loss of reputation which would influence trust and respect for the surveillance system for a long time, maybe forever.

a3b: Increased load on the local systems at the GP office, and correspondingly decreased responsiveness, caused by features in the surveillance system. An example could be that too many requests and corresponding processes are executed simultaneously. For instance during outbreaks, many GPs would issue requests at the same time. Load problems could be caused by missing restrictions imposed in the system, or by software bugs or wrong configuration. This threat is analysed to have *moderate consequence*, or less. It depends on how often and how long these problems with decreased responsiveness are experienced. The result of this threat is mainly annoyance for the user (the GP) and reduced reputation for the surveillance system. On the other hand, a *high likelihood* was indicated for this threat, merely to point at the importance of taking care of the load problem during design and implementation. This is considered an acceptable risk if the increased load does not cause the local EHR system to be completely down for a period of time. Scalability testing of the Snow system concludes that the responsiveness of the system is minimally affected when the number of Snow participants grows [1].

The remaining threats have a *low* risk level. These are basically acceptable risks, but should be monitored to see if they can cause new problems, for instance due to modifications of the service. It is particularly important to observe low-risk threats with *severe* consequence. If the likelihood for these threats increases, their risk level will soon be unacceptably high. The low risk threats with *severe* consequence are therefore discussed specifically here.

g2: Fake software modules can be installed on the surveillance system's servers or in the GP's local systems. This is malicious software which can do all sorts of harm to confidentiality, integrity, and availability of the information and the service,

and lead to a lot of other threats which would be devastating for trust and reputation of the surveillance system. The likelihood was set to *low* because this is planned to be handled in the system development. The access rights to the EHR database will be limited with respect to which information can be extracted, so the malicious software must then be able to modify or overrule these access rights.

c2a and c2b: Sensitive information from the GP's EHR is revealed to unauthorised persons by fake processes which are able to extract sensitive information from the EHR (c2a), or because errors in the surveillance software make it possible to extract sensitive information from the EHR (c2b). Threat c2a is directly related to threat g2 above: If it is not possible to introduce such fake software modules into the surveillance system, this threat disappears more or less. In the Snow system the access rights to the EHR database will be limited and a programming error would not be able to violate these rights (c2b). On the other hand, there could also be an error in the configuration of the access rights to the EHR database, so this threat should be carefully monitored.

c3: Sensitive information is exposed during transfer because of wiretapping, unauthorised persons "listening in" to the communication. The likelihood for this threat is foreseen to be low because end-to-end encryption will be imposed on the communication. This adds to the fact that the information extracted and transferred is intended to be anonymous and thus not sensitive. (See discussion of threat c1 above.)

c4: The GP *intentionally* performs a copy-paste operation from the EHR into a message which is submitted to a receiver. Whether this threat should be analysed at all in our context can be discussed, because this is closely connected to the GP's own ethics and law obedience. If a GP wants to distribute such information there are several means to do so, e.g., e-mail. On the other hand, Snow gives the GPs a new and easy-to-use tool for communication with colleagues, even multicasting. It is very difficult to estimate likelihood for this threat, but if it happens the consequence is severe.

c5: Delivery of information from GP, caused by an *unintentional* copy-paste, or by sending a message to a wrong receiver address. Also for this threat it is difficult to anticipate likelihood. It relates to the possibility of wrong use of the system, and thus to usability aspects of the surveillance service's user interface: The system must be designed so that it is not too easy to place sensitive information into a message or to send a message to wrong address. For instance, if the Municipality Disease Prevention Doctor shall send (multicast) a message about a possible epidemiological outbreak to all GPs in his area, he must not, by accident, be able to also include another receiver.

a1a, a1b: The surveillance system crashes the local EHR server, resulting in a disk crash and destroyed data, or the EHR system being unavailable for a period of time. This could be caused by different errors in the surveillance system or by malicious software exploiting weaknesses in the system and, for instance, causing a denial-of-service attack. In case of disk crash, it is assumed that the GP offices have established the required information security management system which also comprises verified routines for backup and restore of data. These incidents are, of course, serious for the GP office that loses access to the EHR system for a while, but the consequence is even worse for trust and reputation of the surveillance service.

i1a, i1b: The surveillance system causes modification of data and relations in the local EHR system, resulting in wrong patient treatment. This could be caused by fake software modules doing this type of harm (see g2 above), or it could be caused by software errors in the surveillance system. For the Snow case, the modules do not

have write access to the EHR database, so either the malicious software must also modify the access rights to the database, or configuration of access rights must be wrong. The motivation for someone to intentionally modify information in the EHR is considered to be very small, and the likelihood that someone will use Snow to do that is considered minimal.

4. Discussion

None of the 30 identified threats got an unacceptable *high* risk level, but two got *medium* risk level, of which one (c1) was concluded to be unacceptable after further investigation. *Low* risk threats with *severe* consequence are in principle acceptable, but they should be carefully observed since increased likelihood for these could cause their risk level to be unacceptably high.

Many of the identified threats were difficult to analyse with respect to consequence, but particularly with respect to likelihood for unwanted incidents to happen in a system that has not yet been implemented; or nearly not designed at the time of the risk assessment. The likelihood for threats related to software development (software errors, software functionality, wrong usage) has not been possible to analyse, and it has also been difficult to evaluate quality threats resulting from limited use and coverage (too few users). It is much easier to foresee *consequences* of these threats. While using a preliminary design as basis for risk assessment will make the assessment results more uncertain, including risk assessment as part of the early design process will, on the other hand, contribute to a more secure system being developed.

Since it is important to identify and solve all security threats before real-time solutions can be used in a wide scale, additional risk assessments are needed in the design of similar systems and in the development and testing of the deployment version of the Snow system. – More details and future work will be presented on the project web-site, *www.telemed.no/opensource/snow.*

5. Conclusion

The present risk assessment has contributed to the identification of threats to the Snow system in particular, but more important, to future real-time disease surveillance systems in general. The assessment has also contributed to the development of the traditional design process by using the results from the risk assessment as system requirements and input to the design of the overall disease surveillance system.

References

[1] Bellika, J.G., Sue, H., Bird, L. et al. (2007) Properties of a federated epidemiology query system. *International Journal of Medical Informatics* 76(9):664–676.

[2] Johansen, M.A., Scholl, J., Hasvold, P. et al. (2008) "Garbage in, garbage out" – Extracting disease surveillance data from EPR systems in primary care. In *Proceedings of the ACM 2008 Conference on Computer Supported Cooperative Work (CSCW'08)*, ACM, New York, 525–534.

[3] Standards Association of Australia (1999) *Risk Management*. *AS/NZS 4360:1999*.

[4] International Standard ISO/IEC 27005 (2008) *Information Technology – Security Techniques – Information Security Risk Management*. First edition 2008-06-15.