

Risk assessment report

The SNOW project

Version 1.00

07.02.2008

Eva Henriksen, Johan Gustav Bellika,
Anders Baardsgaard (NHN), Johanna Nystad, Per Atle Bakkevoll, Monika Johansen

Norwegian Centre for Telemedicine

Summary

The project

The purpose of the Snow service for epidemiology is early detection of disease outbreak, fast and easy notification and management of disease outbreaks.

The risk assessment

The main part of this risk assessment was performed in October – November 2007. The following persons participated in the risk assessment:

- Risk assessment leader: Eva Henriksen.
- Project participants: Johan Gustav Bellika, Monika Johansen, Johanna Nystad, Per Atle Bakkevoll, Anders Baardsgaard (from NHN)

To analyse the security challenges of the Snow project, we performed a qualitative risk analysis of the information security aspects of the proposed architecture and the intended environment and use. The goal was to identify security threats to the institutions involved and to patient information confidentiality, and to find acceptable solutions to the threats. The threats were identified in two semi-structured brain storming sessions.

We performed the risk analysis by going through the five main steps described in the Australian and New Zealand standard for risk management [3].

1. Context identification
2. Threat identification
3. Analysis of the identified threats with respect to likelihood and consequence
4. Calculation of risk value for each threat as the product of consequence and likelihood, illustrated in a two-dimensional matrix
5. Proposal of risk-reduction treatment for all threats with a non-acceptable risk level

This methodology corresponds very well to the upcoming ISO standard 27005 for information security risk management [4].

Main conclusions

This first high level risk assessment of the Snow system identified no threats with an unacceptable *High* risk level. Only two threats have been given a *Medium* risk level, one of these (**c1**) is considered to be unacceptable. Threat **c1** concerns the possibility that the information produced by the Snow service is sensitive, i.e. not anonymous “enough”.

It is difficult to assess the information security risk of a system this early in its development process. Many threats can be identified as possible unwanted incidents, but it is impossible to foresee their risk level, in particular the likelihood for it to happen. The main result at this stage is therefore that we have been able to identify threats and possible unwanted incidents, and to foresee a *consequence* of these.

In the analysis we have focused particularly on threats with *severe* consequence. The argument for this is that with an increased likelihood these threats will easily get an unacceptable risk level. A tendency, based on consequence, is that all the confidentiality threats have severe consequence, while most of the availability threats have lower consequence. The integrity threats are distributed between severe and moderate consequence.

TABLE OF CONTENT

SUMMARY	2
<i>The project</i>	<i>2</i>
<i>The risk assessment.....</i>	<i>2</i>
<i>Main conclusions</i>	<i>2</i>
1 INTRODUCTION	4
2 CONTEXT IDENTIFICATION.....	5
2.1 DESCRIPTION OF SYSTEM AND SERVICES.....	5
2.2 SECURITY REQUIREMENTS	9
2.2.1 <i>Legal baseline</i>	<i>9</i>
2.2.2 <i>Requirements to the Snow service</i>	<i>10</i>
2.3 DEFINITION OF LIKELIHOOD, CONSEQUENCE AND RISK LEVELS	11
2.3.1 <i>Likelihood and consequence levels.....</i>	<i>11</i>
2.3.2 <i>Acceptance criteria</i>	<i>12</i>
2.3.3 <i>Risk levels.....</i>	<i>13</i>
3 THREAT IDENTIFICATION AND ANALYSIS OF RISK.....	14
3.1 THREATS WITH HIGH RISK LEVEL	15
3.2 THREATS WITH MEDIUM RISK LEVEL	15
3.3 THREATS WITH LOW RISK LEVEL.....	16
4 CONCLUSIONS AND RECOMMENDATIONS.....	18
4.1 RECOMMENDED RISK TREATMENT	18
4.1.1 <i>Treatment of threats with Medium risk.....</i>	<i>19</i>
4.1.2 <i>Treatment of severe threats with Low risk.....</i>	<i>20</i>
4.2 FURTHER INVESTIGATIONS NEEDED	22
4.3 CONCLUSIONS	23
REFERENCES	24
ABBREVIATIONS.....	25
ANNEX A THREAT TABLE	26
ANNEX B PLAN FOR IMPLEMENTATION OF SECURITY MEASURES	35
ANNEX C DESIGN ISSUES RELATED TO PRIVACY AND SECURITY.....	39
C.1 PRESERVE PRIVACY OF DATA FROM A SMALL DATA SET	39

1 Introduction

Risk analysis of information security is a basic requirement of ISO 27002 (formerly ISO 17799), internationally recognized as “the generic information security standard” [1]. Risk analysis is also required by national legislation as a vital part of an information security management system for any organisation. Risk analysis is performed with respect to the main information security aspects Confidentiality, Integrity and Availability. The risk acceptance criteria are defined by the information security policies of the affected organisation.

There are many methods and guidelines for how to perform risk analysis, but all of them include the central tasks of:

- identifying the threats or possible unwanted incidents
- analysing the impacts and probabilities of these threats
- evaluate risks with respect to the acceptance criteria

Our experience is based on the EU-funded research project CORAS from the fifth framework programme (FP5) of Information Society Technology (IST) [2] where a methodology for risk analysis was developed and tested on e-health systems. The methodology was based on the Australian and New Zealand standard for risk management (AS/NZS 4360/1999) [3], which clearly sets out the risk analysis process in five main steps:

1. Context identification; a description of the subject for analysis, i.e. the analysed system and its environment.
2. Threat identification; identify what could possibly happen.
3. Impact and probability analysis; a consideration of the consequences of the threats and the likelihood that these consequences may occur.
4. Risk evaluation; relating the resulting risk level with risk acceptance criteria.
5. Risk treatment; identification and assessment of treatment options.

Lately, the upcoming ISO standard 27005 for Risk Management in Information Security Management Systems [4] are built around a very similar approach.

2 Context identification

The Snow system is a peer-to-peer system that can be used for exchange of any kind of information between health institutions. In this risk assessment we only consider the Snow system as a decentralised disease surveillance system. The Snow system is in this context used to extract and distribute data about occurrences of communicable diseases from electronic health records systems (EHRs) within general practitioners' offices. We also want to exchange text-based messages using the Multi User Chat (MUC) service available in the Instant Messaging service which the Snow system is built on top of.

The subsystems and interfaces of the Snow service are illustrated in figures 1 and 2 below. The following subsystems have been subject for analysis in this risk assessment:

The client side:

- Text conferencing and Instant Messaging (IM) tool.
- The Snow Agent System Client is integrated with the IM client tool.

Together, these tools are used by the end users to receive textual information, and to initiate agent missions to collect, assemble and visualize disease surveillance data.

The server side components:

- The Mission Controller component (MC) which controls and manages agent missions on behalf of a client or Agent, by creating, monitoring, and terminating agent missions.
- The Agent Daemon component (AgD) which initiates the agent processes on the Snow Agent System servers, based on a Mission Specification from the requesting MC.
- The Agent application (disease surveillance specific) which performs local information retrieval and processing according to the mission specification.
- The Poller component (installed in the most secure zone) which is used to poll a server in the health network (less secure zone) for messages. The Poller is a technical solution to ensure that information is not sent from a less secure zone into a more secure zone.
- Post office (PO) component (installed in the health network) which is used to store and route messages on behalf of the servers located in the more secure zones. The PO could also have additional functionality like merging data for a region (processing of single results into aggregates).

Together the client and the server side components form the Snow Agent system.

2.1 Description of system and services

The system consists of clients used by end users and server(s) located in the health network, called post offices, in addition to the Snow Agent System (SAS) servers installed on the GP office's EHR server machines. The servers work together to provide the disease surveillance service. The servers are more or less identical:

- **General server features:** All servers provides the following features:
 - o Cache of latest results
 - o Snow institutional membership list
 - o User database
 - o Data needed to map geographical areas to data provider lists
 - o Geographical data to produce surveillance maps
 - o Demographic data from the local population used to produce disease statistics

- **Local level:** At this level data is extracted from the EHR systems. This is also where most of the targeted end users of the service are located. The end users use clients to request and view information provided by the set of servers. The Snow Agent System (SAS) server in each GP office is normally located on the same machine as the EHR server, but can also be located on a dedicated server. Specialised features:
 - Access patient information in the local EHR server
- **OPTIONAL: Regional (intermediate) level:** Dependent on the size of the surveillance network, a separate level for merging regional data may be needed. The Snow PO (Post Office) servers are used for this purpose. These servers merge data from all GP office Snow servers that is serviced by the post office. If an intermediate level is needed, data is forwarded to the top level which is described below. Specialised features:
 - Post office functionality (store and route messages)
 - Intermediate data merging
- **Top level:** Independent of surveillance network size there always exist an entity that merges all the surveillance data together and produce the final end result that is provided to the end user(s). This entity is called the “main agent”. The main agent is always the first process instantiated within a mission. Based on the specification of the epidemiological query the data provider list is constructed and disease statistics are requested from the local level by doing data extraction from the local EHR systems. Specialised features:
 - EHR data extraction request construction
 - Final data merging and end result production

Figure 1 shows the details of the Snow server at the *local* level. The Agent Daemon acts as a guardian for the local data, processor and storage. It constitutes the Ministry of Home affairs, if compared to a political system. This component is the entity responsible for initiating the processes (named Snow Agent in Figure 1) that extracts data from the EHR system. A Snow Agent process is created from a repository of trusted software that is stored locally and created during install and upgrade of the Snow server. Figure 1 also shows the local data cache that is used to provide fast response when multiple requests ask for the same data. The Mission Controller (MC) represents the Ministry of Foreign Affairs in our analogue to political systems. Its responsibility is to control computations performed on remote servers. It provides mission control service on behalf of local users and agents requested from remote servers. The MC always knows the whereabouts of agents belonging to the local users, but running on other sites. It may also provide migration or mission control service on request from locally running agents.

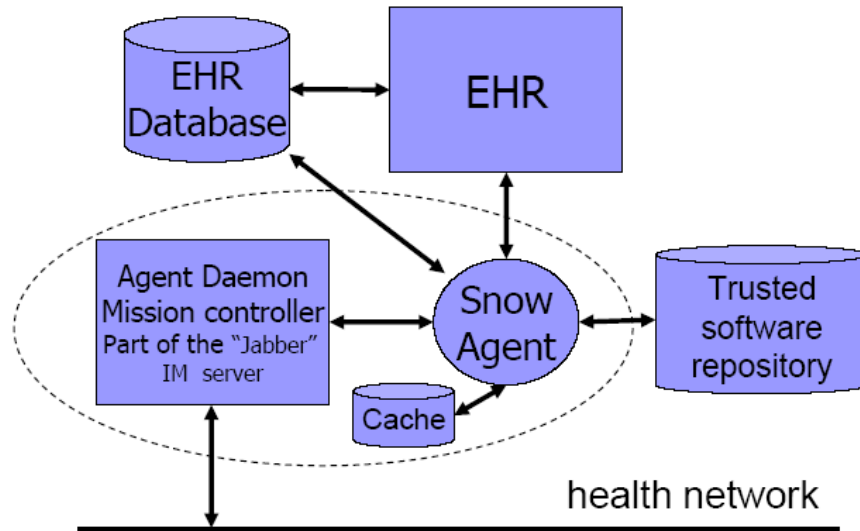


Figure 1: System overview, local level

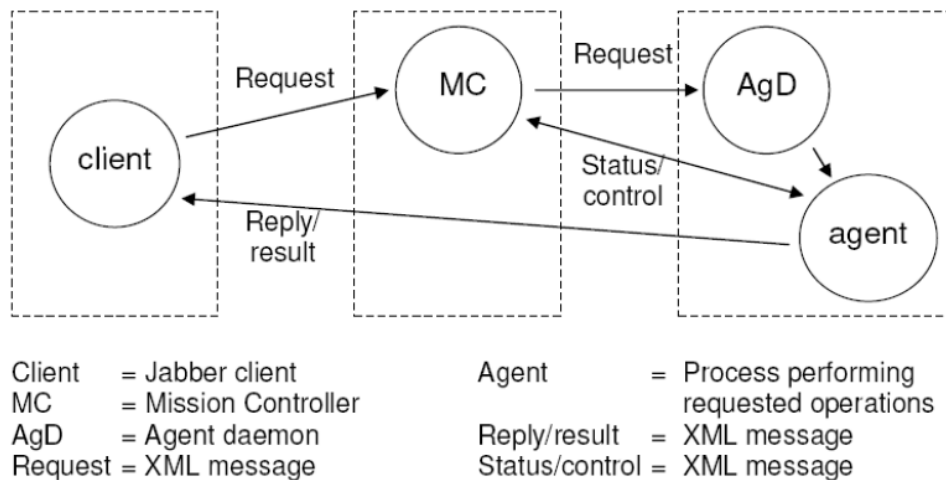


Figure 2: Process/agent initiation method [6]

In Figure 1 the Snow Agent is shown as accessing information in the EHR system or the EHR database. Figure 2 shows how and on what basis the agent is constructed. Normally the user sends a request to a Mission Controller (MC) in form of an XML message containing a “Mission specification” labelled “Request” in figure 2. The Mission Controller forwards this request to one or more remote agent daemons, depending of the type of service requested. If the Agent Daemon (AgD) accepts the request, a process (agent) is constructed from the local software repository shown in figure 1. After instantiation the agent performs the requested action, in this case to extract disease statistics from the local EHR system. To provide statistics, all local data are made anonymous and then transferred out of the local system for further processing as explained above. Both agents and clients can request mission control service from MC. If the word “client” is replaced by “Main Agent”, figure 2 shows how the main agent requests data from the EHRs which have data that is requested in the “mission specification” for the disease surveillance service.

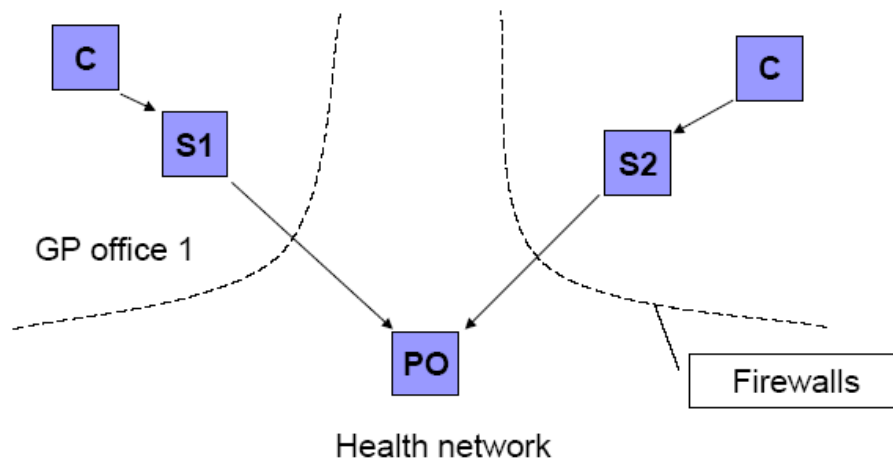


Figure 3: Architecture of the Snow system when Firewalls are used

Figure 3 shows the architecture of the system when firewalls are used to protect the GP office local area network for someone initiating a TCP connection into the software services inside the GP office. The “C” represents clients, The “S” represents Snow agent servers and the “PO” represents the “Post office” Snow agent server that enables communication between the participating institutions. Between the servers S1 and PO, and S2 and PO, we use the Poller and post office (PO) components mentioned above. The PO may run the “Main agent” described above and/or the intermediate regional PO process, if many servers are involved.

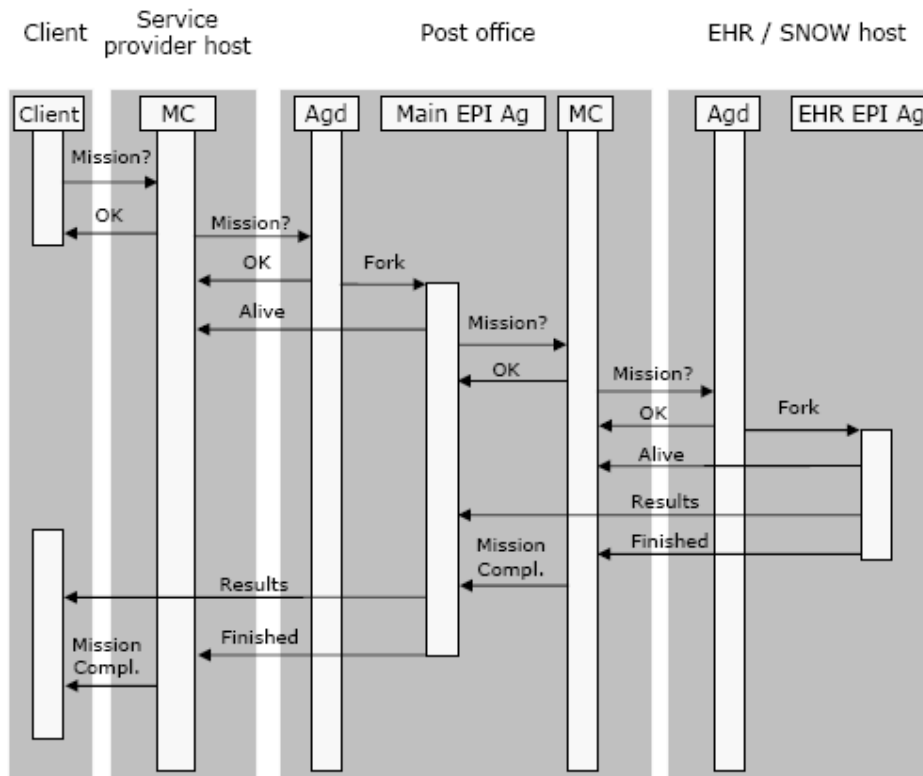


Figure 4: UML sequence diagram for a disease statistics extraction agent mission [5]

Figure 4 shows the sequence of actions performed by the involved actors to produce disease statistics. Time starts at the top and moves downwards, the grey boxes represents different machines where the rightmost represents the Snow agent server located on the same server as the EHR server in the GP office. The white vertical lines represent active processes. The diagram shows that “Main EPI Ag”, or the “Main agent”, is the first process instantiated. The Main agent requests processing of the “EHR EPI Ag”, i.e. the disease statistics extraction agent, which transmits the statistics to the Main agent for merging of data into the final result. The diagram does not show that many EHR EPI Ag processes may work simultaneously.

More information about the system and services are given in published papers [5] and [6] and in reports under www.telemet.no/opensource/snow.

2.2 Security requirements

Privacy is always an issue when patient data is involved. Protection of patient privacy is therefore an important issue to address for a disease surveillance system [6].

2.2.1 Legal baseline

Privacy requirements related to communication of sensitive patient-identifiable information establish the baseline for the information security needs. Confidentiality requirements originate from the professional secrecy and non-disclosure agreement imposed to all health-care workers. Requirements to electronic communication of patient information come from national legislation in European countries, which are also based on EU’s regulation on processing of personal data (Directive 95/46/EC) from 1995 [7]. At the lowest level these requirements become apparent through the security policies of the affected organisation.

According to Norwegian legislation, all health-related information concerning an identifiable person is considered sensitive information (Personal Data Act (Personopplysningsloven) §2) [8]. No one else than those who have a treatment relation to the person, should be able to access this person’s health information, unless the patient has given his or her consent.

Norwegian legislation requires risk assessment as part of an information security management system for any organisation. The legislation also defines information security to include the following aspects: Confidentiality, Integrity, Quality, and Availability [8, 9]. The risk assessment is performed with respect to these information security aspects.

Personal identifiable information (PII) is any information that can identify a physical/natural person. The definition of personal information in the EU Directive [7] reads: “*Personal data shall mean any information relating to an identified or identifiable natural person (“data subject”); an identifiable person is one who can be identified, directly or indirectly, in particular by reference to an identification number or to one or more factors specific to his physical, physiological, mental, economic, cultural or social identity*”.

The “Article 29 group” is a “working party on the protection of individuals with regard to the processing of personal data”.¹ It has been formed with the regard to the EU Directive [7], in particular with regard to Article 29 and paragraph 1(b) of Article 30. (Hence the chosen short-name.) – In [11] the “Article 29 group” analyses further the concept of Personal Data. They state that “*a person may be identified directly by name or indirectly by a telephone number, a car registration number, a social security number, a passport number or by a combination of significant criteria which allows him to be recognized by narrowing down the group to which*

¹ http://ec.europa.eu/justice_home/fsj/privacy/workinggroup/index_en.htm

he belongs (age, occupation, place of residence, etc.)". This indicates that it also depends on the context of the particular situation which identifiers are sufficient to achieve identification.

On the other hand, if the information is *anonymous* it is not defined as personal identifiable information and *not* subject to the legislation given by the EU Directive [7] or the Norwegian Personal Data Act [8].

The Article 29 group defines anonymous data as any information relating to a natural person where the person can *not* be identified, neither by the data administrator nor by any other person, neither directly or indirectly. They state that a hypothetical possibility to single out the individual is not enough to consider the person as "identifiable": *If, taking into account "all the means likely reasonably to be used by the controller or any other person", that possibility does not exist or is negligible, the person should not be considered as "identifiable", and the information would not be considered as "personal data"*.

One example: The description "**gender: Male; age: 50-59; occupation: Bus driver**" is anonymous if we consider the whole population of Norway, or even Oslo. But if we consider the population of a small municipality of less than 1000 inhabitants, this could point directly to one specific person.

If the intention is to keep the information anonymous, the Article 29 group states: *If a criterion appears to lead to identification in a given category of persons, however large (i.e. only one doctor operates in a town of 6000 inhabitants), this "discriminating" criterion should be dropped altogether or other criteria be added to "dilute" the results on a given person.*

The Article 29 group concludes that the assessment of whether the information can be considered as anonymous, or the data identifies an individual, depends on the circumstances. A case-by-case analysis is therefore needed. Further, they state that: *This is particularly relevant in the case of statistical information, where despite the fact that the information may be presented as aggregated data, the original sample is not sufficiently large and other pieces of information may enable the identification of individuals.*

2.2.2 Requirements to the Snow service

The information handled by the Snow system is health-related information. Consequently, if the information is person identifiable (i.e. not anonymous) it will also be sensitive.

A core privacy principle in the design of the Snow system is to keep the sensitive person identifiable health data locally at each GP office and not transfer it to a central site for processing [5]. Information processing can be performed locally at the GP office, and only anonymous information should be communicated by Snow.

The question is whether we really manage to keep it anonymous. What about demographical and geographical information? One thing is to remove information that can identify a person directly or indirectly, such as name, telephone number, social security number, passport number, etc. But a person may also be recognized by narrowing down the group to which he belongs (age, place of residence, occupation, etc.).² For Snow the "narrowing down" is especially related to the size of the population in the selected geographical area (e.g. covering one postal code zone) and to the infrequency of the diagnosis (prevalence).

There are several worries that could be expressed by the GPs who are asked to be included in the Snow service [5]:

² Refer to the bus driver example in section 2.2.1

- Confidentiality and privacy: Is the transferred information really anonymous? Could it cause a deliberate or unintentional dump of patient data (from EHR) on the Internet?
- Availability of the local EHR: Can the new software interfere with the EHR system? Can the Snow agents contain malicious software (i.e. are they Trojan Horses)? Has Snow the capacity to bring down the EHR system in a GP's office due to programming errors? Could it take all computing resources away from the GP's local systems?
- Quality and representativeness: How widespread is the service? – The quality of the information is dependant on the correctness of the retrieved data and the coverage of the Snow service among the GPs.

2.3 Definition of likelihood, consequence and risk levels

We have chosen to use qualitative values for likelihood, consequence and risk levels.

2.3.1 Likelihood and consequence levels

We decided to use four levels for identification of likelihood and four levels for identification of consequence. The levels are defined in table 1 and table 2.

The likelihood levels can be described as frequency values or with respect to how easy it is for a person to exploit a threat. For some threats it is easier to think of the likelihood in the form of frequency or a probability value. This may often be the case for threats related to availability, e.g. caused by problems in sw or hw. For other threats it is easier to think of likelihood when related to ease of misuse or mistake, or to motivation for performing a malicious action. – For each threat or unwanted incident we choose the most appropriate column or the column that is easiest to use in order to estimate the likelihood for the threat.

Table 1: Definition of likelihood levels

Likelihood	Frequency	Ease of misuse and motivation
Very high	Very often, occurs more often than every 10 th connection, i.e. more frequently than 10 % of the time/cases.	Can be done without any knowledge about the system; or without any additional equipment being used; or it can be performed by wrong or careless usage.
High	Quite often. Occurs between 1 % and 10 % of the time/cases.	Can be done with minor knowledge about the system; or without any additional equipment being used; or it can be performed by wrong or careless usage.
Moderate	May happen. Occurs between 0.1 % and 1 % of the time/cases.	Normal knowledge about the system is sufficient; or normally available equipment can be used; or it can be performed deliberately.
Low	Rare. Occurs less than 0.1 % of the time/cases.	Detailed knowledge about the system is needed; or special equipment is needed; or it can only be performed deliberately and by help of internal personnel.

The consequence levels are described in terms of consequences for the patient (user) and consequences for the service or the service provider. In this case the service provider could be both the GP office (seen from the patient's viewpoint) and/or the project owner and the Snow service (seen from the GP's viewpoint).

For each threat or unwanted incident we choose the most appropriate description to estimate the consequence level for the threat.

Table 2: Definition of consequence levels³

Consequence:	
Small	<p><u>For the patient:</u> No impact on health; or negligible economic loss which can be restored; or small reduction of reputation in the short run.</p> <p><u>For the service provider:</u> No violation of law; or negligible economic loss which can be restored; or small reduction of reputation in the short run.</p>
Moderate	<p><u>For the patient:</u> No direct impact on health or a minor temporary impact; or economic loss which can be restored; or small reduction of reputation caused by revealing of less serious information (e.g. blood pressure level).</p> <p><u>For the service provider:</u> Offence, less serious violation of law which results in a warning or a command; or economic loss which can be restored; or reduction of reputation that may influence trust and respect.</p>
Severe	<p><u>For the patient:</u> Reduced health; or a large economic loss which cannot be restored; or serious loss of reputation caused by revealing of sensitive and offending information.</p> <p><u>For the service provider:</u> Violation of law which results in minor penalty or fine; or a large economic loss which cannot be restored; or serious loss of reputation that will influence trust and respect for a long time.</p>
Catastrophic	<p><u>For the patient:</u> Death or permanent reduction of health; or considerable economic loss which cannot be restored; or serious loss of reputation which permanently influences life, health, and economy.</p> <p><u>For the service provider:</u> Serious violation of law which results in penalty or fine; or considerable economic loss which cannot be restored; or serious loss of reputation which is devastating for trust and respect.</p>

2.3.2 Acceptance criteria

We use accept criteria to define the acceptable risk level for the service. We cannot expect to achieve a risk level equal to zero. Thus we have to define which level of risk we consider as acceptable for the service we are analysing. The accept criteria should be based on the security requirements for the service.

The Norwegian Health Personnel Act (Helsepersonelloven) states in chapter 5 the obligation to maintain secrecy with respect to health information a person has been acquainted with in his or her duty as health personnel [10].

The following acceptance criteria have been proposed for the Snow service:

It is not acceptable that⁴:

1. (C) – the likelihood is higher than **low** that unauthorised persons (i.e. anyone else than the patient, and those who have a treatment relation to the patient) get access to the patient's personal health data (i.e. to sensitive data). This is regardless of why, where, and how it happens. *(This means that in order to obtain unauthorised access to such data, detailed knowledge is needed about the technical system, or special equipment is needed, or it can only be performed by help of internal personnel.)*

³ These are the same four consequence levels as used by Helse Nord in their template for risk assessments.

⁴ The letter in parenthesis refers to the security aspects confidentiality (C), integrity (I), availability (A)

2. (A) – the likelihood is higher than **low** that the Snow service causes the local EHR system to be down for a period of time. *(This corresponds to up to 2.4 minutes of a 40 hours work week, or that it happens more infrequent than once for every 1000 Snow accesses.)*
3. (A) – the likelihood is higher than **low** that the Snow service causes data in the local EHR system to be destroyed. *(I.e. that it happens more infrequent than once for every 1000 accesses to the Snow service.)*
4. (A) – the likelihood is higher than **moderate** that the Snow service is unavailable for a period of time. *(This corresponds to up to 24 minutes of a 40 hours work week, or that it happens not more than once for every 100 Snow accesses.)*
5. (I) – the likelihood is higher than **low** that the Snow service causes information in the local EHR system to be modified. *(I.e. that it happens more infrequent than once for every 1000 accesses to the Snow service.)*
6. (I) – the likelihood is higher than **low** that information in the Snow system (request, results) are being modified. *(I.e. more infrequent than once for every 1000 accesses to the Snow service.)*

2.3.3 Risk levels

We have decided to use three distinct levels for risk: *Low, Medium, and High*. Our risk level definitions are presented in table 3.

The risk value for each threat is calculated as the product of consequence and likelihood values, illustrated in a two-dimensional matrix (figure 5). The shading of the matrix visualizes the different risk levels. Based on the acceptance criteria, the risk level *High* is decided to be unacceptable. Any threat obtaining this risk level must be treated in order to have its risk reduced to an acceptable level.

Table 3: Definition of risk levels

Risk level:	
Low	Acceptable risk. The service can be used with the identified threats, but the threats must be observed to discover changes that could increase the risk level.
Medium	The risk can be acceptable for this service, but for each threat the development of the risk must be monitored on a regular basis, with a following consideration whether necessary measures have to be implemented.
High	Not acceptable risk. Can not start using the service before risk reducing treatment has been implemented.

Figure 5: Risk matrix showing the defined risk levels

Consequence:	Small	Moderate	Severe	Catastrophic
Likelihood:				
Low	Low	Low	Low	Medium
Moderate	Low	Medium	Medium	High
High	Low	Medium	High	High
Very high	Medium	High	High	High

3 Threat identification and analysis of risk

Approximately 30 threats and unwanted incidents have been identified. The threats are listed in the threat table in Annex A. For each possible threat we wanted to evaluate its impact or consequence and the likelihood that it would occur. Threats were given qualitative values for consequence and likelihood, according to definitions in tables 1 and 2.

Many of the threats have, however, been difficult to analyse with respect to consequence and in particular with respect to likelihood. It is problematic in the design phase to imagine the *likelihood* for possible unwanted incidents to happen in a system that has not yet been fully implemented. However, we have more than ten years of research activity on this concept.

At this early stage we do not know what the user interface will look like. Typically, we have not been able to analyse the likelihood for threats related to software development (software errors/bugs, software functionality, wrong usage), and it is also difficult to evaluate quality threats resulting from limited use and coverage (too few users). It is much easier to foresee *consequences* of these threats and possible unwanted incidents.

The risk value for each threat⁵ is calculated as the product of consequence and likelihood values. The unique ID of the threat is written into the corresponding cell of the matrix, as shown in figure 6.

Figure 6: Risk matrix for the Snow service

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				

The placing of the identified threats in the risk matrix shows a particular tendency: Most of the threats have been analysed to have a Low risk level. It is mainly the *likelihood* for these unwanted incidents that has been evaluated to be Low. First of all this is related to the problems, mentioned above, of analysing a system which is in its early design phase.

Ten of the threats have, however, been analysed to have Severe *consequence*. If the likelihood for these threats increases, their risk level will soon be unacceptably high. This is for instance the case for all the threats related to confidentiality (c1-c5). For these threats we can argue that the low likelihood is based on two important design assumptions:

⁵ The following threats from the table in Annex A have not been given a risk value and are therefore not included in the matrix: **g1, a8a, a8b, a9, q1a, q1b, q2, q3**

- The information that is retrieved from the EHR is (meant to be) anonymous
- End-to-end encryption is imposed between Snow nodes.

The matrix also shows that while all the confidentiality threats are analysed to have Severe consequence, most of the threats to availability (and integrity) are analysed to have a lower consequence. A reason for this is that while confidentiality breaches are violation of law, the consequences of availability breaches in this case are more related to the trust and reputation of the Snow service. (This could of course in some cases be devastating for Snow.) It is also a tendency that threats to availability and integrity of the local EHR system are considered more serious than similar threats to availability and integrity of the Snow system.

Some of the identified threats are discussed in more detail in the following subsections.

3.1 Threats with High risk level

In this analysis none of the threats have got an unacceptable *high* risk level.

3.2 Threats with Medium risk level

Only two threats have got a *medium* risk level. Even if these in principle could be acceptable risks, each of them should be investigated separately to see if they can cause additional problems. In this case the conclusion is that one of them (c1) is unacceptable, while the other (a3b) can be accepted.

c1 – Sensitive (i.e. person identifiable) information is extracted from the EHR by the agents, and communicated in the Snow system.

The likelihood for this threat is uncertain, but is analysed to be *higher than low*. According to acceptance criterion 1 in section 2.3.2 this threat is therefore unacceptable: “It is not acceptable that the likelihood is higher than *low* for unauthorised persons to get access to sensitive data.” – It is difficult to suggest a likelihood for this to happen. Originally, we have said that this must be further investigated. But we predict that if this is not especially handled by the functionality of the system, the likelihood will be more than Low (i.e. at least Moderate).

The consequence is analysed to be *severe* for the system/service⁶. 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 that will influence trust and respect for the Snow system/service for a long time (maybe forever).

The legal baseline and definitions of person identifiable information (PII) and anonymous information is discussed in section 2.2.1 above.

When considering whether the information is anonymous, one must also take into account the possibility of design and programming errors in the development of this functionality of the Snow service.

a3b – Increased load on the local systems at the GP office, and correspondingly decreased responsiveness, because of features in the Snow system.

Examples of such features in the Snow system could be that too many missions (requests) and corresponding agents are executing simultaneously. It could, for instance, happen during outbreaks that many GPs issue similar requests at the same time. Load problems could be

⁶ Even if the consequence is not Severe, but only Moderate, the Risk level will still be Medium.

caused by missing limitations/restrictions imposed in the system, or by errors/bugs or wrong configuration.

This threat is analysed to have *moderate consequence*, or less. It depends on how often and how long they experience problems with decreased responsiveness. The result of this threat is mainly annoyance for the user (GP) and reduced reputation for the Snow system. On the other hand, we have indicated a *high likelihood* for this threat, merely to point at the importance of taking care of the load problem during design and implementation.

If the increased load does not cause the local EHR system to be completely down for a period of time, this is considered an acceptable risk (acceptance criterion 2 in section 2.3.2).

In connection to this threat, we refer to paper [6] which documents the scalability of the Snow system and concludes that the responsiveness of the Snow system is minimally affected when the number of Snow participants grows.

3.3 Threats with Low risk level

The remaining 20 threats have a *low* risk level. These are therefore acceptable risks, but one should occasionally keep an eye on them to see if they can cause new problems. Some of the risks could for instance change due to modifications of the service.

It is particularly important to observe the ten low-risk-threats which have been analysed to have *severe* consequence. If the likelihood for these threats increases, their risk level will soon be unacceptably high. This could happen if the communication is done without encryption and the information transferred is not anonymous. The low risk threats with *severe* consequence are therefore discussed separately here.⁷

g2 – False or bogus software modules can be installed on the Snow servers or in the GP's local systems.

Such false modules must be considered malicious software (malware). They can do all sorts of harm to the confidentiality, integrity, and availability of the information and the service, and lead to a lot of other threats. For example, a fake/bogus client could send a request to a corresponding fake/bogus agent which extracts patient id information from the EHR.

If this happens it will be devastating for the trust and reputation of the Snow system/service. The likelihood, however, has been set to *low* because this is foreseen to be taken care of in the development of the system, by several means. For instance, the access rights to the EHR database must be limited with respect to which information that can be extracted. The malicious software must then, in addition, be able to modify or overrule these access rights.

c2a – Sensitive information from the GP's EHR is revealed to unauthorised persons by false or bogus agents that are able to extract sensitive information from the local EHR.

This threat is directly related to threat g2 above. If it is not possible to introduce such fake/bogus software modules into the Snow system, this threat disappears more or less.

c2b – Sensitive information from the GP's EHR is revealed to unauthorised persons because errors/bugs in the Snow software makes it possible to extract sensitive information from the local EHR.

The likelihood for this to happen is probably lower than the likelihood for threat c2a above. The Snow system's access rights to the EHR database will be limited, and a programming

⁷ Low-risk-threats with consequence analysed to be lower than *severe*, are not discussed here.

error would not be able to violate these rights. – On the other hand, there could also be an error in the setting (configuration) of the access rights to the EHR database.

c3 – Sensitive information is exposed during transfer because of wiretapping, unauthorised persons “listening in” to the communication.

The likelihood for this risk is foreseen to be (very) low because end-to-end encryption will be imposed on the communication. And this adds to the fact that the information extracted and transferred is intended to be anonymous and thus not sensitive. (See the discussion of threat c1 in section 3.2 above.)

c4 – The GP intentionally performs a copy-paste operation from the EHR into a message which is submitted to the JID for a receiver

It can be discussed whether this threat should be analysed in our context at all, because it is heavily connected to the GP’s own ethics and law obedience. If the GP wants to distribute such information he will have several means to do so, e.g. e-mail. The only reason to include this as a possible threat is that the Snow system gives the doctors a new and easy-to-use tool for communication with colleagues.

It is very difficult to give any likelihood for this threat, but the consequence if this happens is considered to be *severe*.

c5 – Unintentional delivery of information from GP, caused by an unintentional copy-paste, or by sending a message to a wrong receiver address (JID)

Also for this threat it is very difficult to anticipate a likelihood. An incident like this is related to the possibility of wrong use of the system, and thus to usability aspects of the user interface of the Snow service. – Is it too easy to place sensitive information into a message? Is it too easy to send a message to a wrong JID? For instance, if the “disease prevention doctor” wants to send (multicast) a message about a possible epidemiological outbreak to all GPs in his area, is it then possible that he, by incidence, also includes the JID to another receiver?

a1a, a1b – The Snow system crashes the local EHR server, which results in either disk crash with destroyed data or the EHR system being unavailable for a period of time.

These incidents could be caused by malicious DoS-attack utilizing weaknesses in the Snow system, or by other errors/bugs in Snow. In the case of disk crash, we must assume that the GP offices have established their own information security management system which also includes verified routines for backup and restore of data.

These incidents are of course serious for the GP office that loses their EHR system for a while, but the consequence is even worse for the trust and reputation of the Snow service.

i1a, i1b – The Snow system causes modification of data/information and relations in the local EHR system, which results in wrong patient treatment.

This could be caused by false/fake/bogus software modules doing this type of harm maliciously (see threat g2 above), or it could be caused by errors/bugs in the Snow system. But the Snow modules do not need write access to the EHR database, so either the malicious software must also modify the access rights to the database, or the configuration of the 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 therefore considered being minimal.

4 Conclusions and recommendations

4.1 Recommended risk treatment

There are basically four different approaches to handle a risk [3, 4]:

1. **Accept** the risk, in accordance with the organisation's security policy. This approach is usually applied for the risks with an acceptable risk level. *It is worth remembering that accepting the risk does not mean accepting the unwanted incident indicated by the threat.*
2. **Reduce** the risk to an acceptable level. Since the risk is a product of likelihood and consequence, this means to reduce the likelihood, the consequence, or both. It is often difficult to reduce the consequence of a threat, so the focus should first of all be on reduction of the likelihood.
3. **Avoid** the risk, i.e. try not to be exposed to the risk, not do the things that could lead to the risk. (In our case this could mean not installing the Snow service.)
4. **Transfer** the risk to a third party (e.g. an insurance company)

In this analysis we will mainly stick to strategies 1 and 2 above and recommend security measures that can *reduce risks* to an acceptable level. Risk reduction should be subject to a cost/benefit analysis, and if cost effective, risks should be reduced based on the ALARP principle (As Low As Reasonable Possible).

There are also examples of using *risk avoidance* (3) as a strategy, and *risk transfer* (4) could possibly be considered for threats to the local EHR systems

The table in Annex B lists the threats with risk level Medium (from section 3.2) and the threats with Severe consequence (from section 3.3), together with security measures proposed for treatment of these threats. Treatments of each of these threats are discussed in more detail in sub-sections 4.1.1 and 4.1.2.

Some of the other threats with Low risk level will also benefit from treatment proposed to threats with higher risk. Table 4 summarizes the proposed treatments for risk reduction, and lists all the threats that would have their risk reduced by implementing these treatments.⁸ Different treatment options are grouped under a few main headings, indicating the need for routines to be defined, design decisions and configuration, encryption, quality assurance, and user training. But even if treatment options are grouped in this way, they are also depending on each other, and risk treatments for a certain threat will therefore appear in more than one group. – For instance; access rights and restrictions to the EHR database have to be defined by policies/routines *and* it has to be configured. Another example is encryption issues where routines for key administration also have to be defined.

Table 4: Proposed risk reduction treatment and affected threats

Treatment; security measures	Affected threats
Policies, routines and procedures to be defined for: <ul style="list-style-type: none"> - Installation, upload, and upgrade of Snow sw at GP offices and PO - Access rights and restrictions to EHR database - Administration of encryption keys - Security measures and protection at GP office and PO - Restriction of results from geographic areas with limited population 	g2 c2a, c2b, i1a, i1b g2, c2a, c3 c3, a1a, a2, a3a, i1a c1

⁸ In this table threats with severe consequence are written in bold face print, while the rest of the threats are written in normal print.

Treatment; security measures	Affected threats
Design and configuration decisions: <ul style="list-style-type: none"> - Access rights/restrictions to EHR database - Database access method (EHR database) - Where to process and aggregate sensitive information - Limit the possibility to paste (sensitive) information into messages - How to handle GP offices (Snow servers) that do not respond - Strategies for limitation of load - Filters to restrict communication - "Karma" – black-listing of clients 	c2a, c2b, i1a, i1b c2a, c2b, i1a, i1b c1 c4, c5 a6a a2, a3a, a3b , a5 a4 a1a, a3b , a4
PKI and encryption: <ul style="list-style-type: none"> - Digital signature on installed/uploaded sw modules - End-to-end encryption of transferred data 	g2, c2a , i3b c3 , i2
Quality assurance and test procedures	c2b, a1b, a3b, i1b , a8a, a8b
User education and training: <ul style="list-style-type: none"> - Information on legal aspects and risks; awareness - Good (self-evident) user interface with on-line help - Simple user manuals - User training/education 	c4, c5 , i3a c5 , a9 c5 , a9 c5 , a9

For a couple of threats, *risk avoidance* could be an alternative to risk reduction: The threats related to disclosure of sensitive information in messages (threats **c4** and **c5**) could be avoided by not implementing the functionality of sending messages.

Risk transfer could be an alternative treatment for a couple of other threats, i.e. pay an insurance to cover for a possible economical loss as a result of an unwanted incident. This could be a solution for threats causing the GPs EHR system to be unavailable for a period of time (**a1a** and **a1b**) and for threats to the integrity of information in the EHR database (**i1a** and **i1b**).

4.1.1 Treatment of threats with Medium risk

This section discusses the threats from section 3.2.

c1 – Sensitive (i.e. person identifiable) information is extracted from the EHR by the agents, and communicated in the Snow system.

Anonymisation is imposed, meaning that obviously identifiable information like name, full address, and personal number⁹ are not retrieved from the EHR. It is an open question whether gender, age, or age group should be extracted. But some data could still be sensitive because of small geographic area (municipality, postal code area) and rare diagnosis. Routines must be defined for restrictions with respect to rare diagnoses, limited geographical areas, gender, age, etc. Finally, the doctors should be given the possibility to define/configure which diseases they will (not) provide information about, by defining their own "disease profiles".

A part of this is to decide where the sensitive information is to be processed, controlled and assembled. Sensitive information could be transferred from GP office to PO if end-to-end

⁹ i.e. fødselsnummer

encryption is imposed. But the question is whether the result could be made public if there is only one case in a small population. – Which size should the population have to allow the result to e.g. be shown on a map? (*Treatment of this threat is further discussed in Annex C.*)

a3b – Increased load on the local systems at the GP office, and correspondingly decreased responsiveness, because of features in the Snow system.

Black-listing of clients who send (too) many requests (“Karma”).

Quality assurance and extensive test procedures are needed to avoid programming errors, bugs, wrong configurations, and missing limitations/restrictions. In particular, extensive testing of the Snow system together with a running EHR system should be performed.

There are many ways to restrict this type of load:

- Processing of extensive requests (for a wide time period) could be restricted to be performed only at night time.
- A maximum number of agents per Agent Daemon (configurable). A new agent will not start until there are available resources.
- During outbreaks it is foreseen that many equal requests are submitted at the same time. For this purpose, a cache could be kept in PO for “fresh” results.
- Timeout on unsuccessful actions to avoid deadlock or infinite loop; the corresponding agents are killed.

4.1.2 Treatment of severe threats with Low risk

This section discusses the threats from section 3.3.

g2 – False or bogus software modules can be installed on the Snow servers or in the GP’s local systems.

Routines must be made for upload, installation, and upgrade of the Snow software to the servers at the GP office (and to the PO). A central part of the solution should be digital signature on the uploaded software. – One could compare this to how the EHR system vendors install and update their software. This is part of the general information security strategy at the GP offices.

The information (spec luggage) sent to the SAS server for a specific mission must not interfere with the SAS software repository. There should be dedicated ports for input/output of such data to/from the server.

c2a – Sensitive information from the GP’s EHR is revealed to unauthorised persons by false or bogus agents that are able to extract sensitive information from the local EHR.

This threat is related to threat g2 above: If it is possible to introduce fake/**bogus** software modules, then they can do all sorts of harm. Access rights to the EHR database must be restricted with respect to which information it is possible to extract. Then the fake/malicious software must also modify the access rights to the EHR database (the configuration).

Different alternatives for database access should be considered. Among these are the use of web services (which is a language independent solution), and a combination of stored procedures and view.

c2b – Sensitive information from the GP’s EHR is revealed to unauthorised persons because errors/bugs in the Snow software makes it possible to extract sensitive information from the local EHR.

Quality assurance and extensive test procedures must be performed in order to avoid programming errors. This also includes the configuration of database access.

Access rights to EHR database must be restricted (see c2a above), including the right to modify the configuration. EHR access is limited to the specific application, e.g. epidemiology. A programming error cannot violate that. An agent can not extract more information than what has been defined as accessible.

c3 – Sensitive information is exposed during transfer because of wiretapping, unauthorised persons “listening in” to the communication.

Data transferred between Snow server at GP office and PO will be encrypted end-to-end, by use of PKI solution and session keys. Routines for key administration must be defined.

Data will be decrypted in PO, for further processing. This means that the PO must be placed in secured/trusted environments, with protection like firewalls and virus control. This is assumed to be the case if PO is located in the Norwegian health net. (See also discussion of threat c1 in section 3.2 above.)

c4 – The GP *intentionally* performs a copy-paste operation from the EHR into a message which is submitted to the JID for a receiver.

There is not much to do if the doctor really intends to do this. But measures discussed for c5 below could also reduce the risk for this threat.

c5 – Unintentional delivery of information from GP, caused by an unintentional copy-paste, or by sending a message to a wrong receiver address (JID)

Information should be given to the users (GPs) about legal aspects and the possible risks (awareness). Education and training of users is necessary to avoid user mistakes, combined with good, obvious user interfaces and instructing user manuals.

A more strict measure is to remove the possibility to paste data from EHR into messages. Or, if this functionality is really necessary, have a pop-up dialogue box asking the user for a confirmation of the paste operation, with default set to “no”, and where the text to be pasted is shown in the dialogue box.

A strategy for risk avoidance is to *not* open up for the possibility to send messages, at least not from the ordinary GPs. This could, however, be a useful functionality for the “disease prevention doctor”.

a1a, a1b – The Snow system crashes the local EHR server, which results in either a disk crash with destroyed data or the EHR system being unavailable for a period of time.

These incidents could be caused by malicious DoS-attack utilizing weaknesses in the Snow system, or by other errors/bugs in Snow. The local system at the GP office is assumed to have an up-to-date information security management system with firewall, virus control, and regular security patching. The information security management system shall also have routines for backup and restore. This should minimize the damage in case of disk crash.

It should also be investigated whether it is possible to retrieve lost messages again from the communication service, i.e. the latest messages, sent after the last backup.

Black-listing should be imposed of clients who send (too) many requests (“Karma”).

Bugs that steal processing time will cause problems for the Snow service but not for the local EHR. Area reserved for each process is normally protected from the area of other processes. We must assume that these types of weaknesses in MS Windows is detected and corrected.

In any case, quality assurance and extensive test procedures is always needed to avoid programming errors.

A different and additional approach could be *risk transfer*: Consider the possibility to pay an insurance to cover for possible economical losses of the GP if the EHR system is unavailable for a period of time.

i1a, i1b – The Snow system causes modification of data/information and relations in the local EHR system, which results in wrong patient treatment.

The local system at the GP office is assumed to have an up-to-date information security management system with firewall, virus control, and regular security patching.

As for c2a above, access restrictions must be imposed on the EHR database, by the EHR system vendors. The Snow modules do not need to have write access to the EHR database. Alternatives for database access must be decided (Web Services or a combination of stored procedures and view).

As for c2b above, quality assurance and extensive test procedures must be imposed, to avoid programming errors.

An additional approach could be *risk transfer*: Consider the possibility to pay an insurance to cover for possible economical losses of the GP if damages to the EHR system cause wrong patient treatment.

4.2 Further investigations needed

Not all threats could be analysed in this first risk assessment, mainly because of uncertainties in the architecture and design of the Snow system. At such an early stage, threats have been difficult to analyse with respect to consequence and in particular with respect to likelihood.

Most of the threats have got a Low risk level, mainly because the *likelihood* for these unwanted incidents has been evaluated to be Low. As mentioned above, this is related to the problems of analysing a system which is in its early design phase.

For the following threats we have not at all been able to give a value for likelihood and/or consequence:

g1 – Use of open source software, and/or publishing the Snow software as open source.

This problem is foreseen to be addressed in a separate study, at a general level.

a8a, a8b – Snow service unavailable as a result of software errors, at the GP office and PO, respectively. It is difficult to predict the likelihood for this, but in any case it is important to impose quality assurance routines and extensive test procedures.

a9 – Snow service unavailable as a result of wrong use (wrong user actions) at the GP office. Also for this it is difficult to predict the likelihood; data concerning usage will be needed. In any case it is important to have an intuitive user interface, good user manuals, and training and education of users.

Quality threats **q1-q4** – These threats are mainly related to the coverage of the results: Is the number of participating GPs large enough? Is the resulting information representative of the real situation? There is also a basic question whether these quality threats concerns the information security at all. One argument for keeping them in the risk assessment for information security is that

A second and more thorough risk assessment is needed when the system design is stable and the implementation has been tested to a certain extent. In addition to testing the quality and functionality of the Snow system, tests should be performed related to performance, throughput, and system load. It is particularly important to test the system load when Snow is running together with the EHR system. Paper [6] presents a scalability testing that has been done, but

in that case only the performance of Snow was evaluated, not the performance of an EHR system running with Snow.

Before the next risk assessment there should also be clarifications concerning the expected coverage of the Snow service (number of GP offices participating, geographical area covered), and the responsibilities regarding operations, maintenance and support of the Snow service should be decided.

4.3 Conclusions

In general, it is difficult to assess the information security risk of system that has not been developed yet, or not even designed. Many threats can be identified as possible unwanted incidents, but it is impossible to foresee their risk level, the consequence, and in particular the likelihood for it to happen.

The main result at this stage is therefore that we have been able to imagine (identify) threats and possible unwanted incidents, and to foresee *consequences* of these.

In this first risk assessment of the Snow service no threats with an unacceptable *High* risk level have been identified.

Only two threats have been given a *Medium* risk level, one of these (**c1**) is considered to be unacceptable. Threat **c1** concerns the possibility that the information produced by the Snow service is sensitive, i.e. not anonymous “enough”.

Among the threats with *Low* risk level we have focused especially on the threats with *severe* consequence. The argument for this is that with an increased likelihood these threats will easily get an unacceptable risk level. A tendency, based on consequence, is that all the confidentiality threats have severe consequence, while most of the availability threats have lower consequence. The integrity threats are distributed between severe and moderate consequence.

Measures for *risk reduction* have been proposed in section 4.1. Among these treatment options we want to mention the following from table 4:

- Definition of policies, routines and procedures for specific areas, like:
 - installation, upload, and upgrade of Snow software modules at GP offices and PO
 - access rights and restrictions to the EHR database
 - administration of encryption keys
 - security measures at GP office and PO (firewall, virus protection, backup, etc)
 - how to restrict results from areas with limited population so that it is still anonymous.
- Design and configuration decisions for e.g.:
 - access to EHR database (access rights and methods)
 - filters to restrict communication and limit system load
- Encryption and PKI – to protect confidentiality and integrity
- Quality assurance and test procedures
- Training and education of users

In section 4.1 we have also mentioned *risk avoidance* and *risk transfer* as alternatives to risk reduction for a few of the risks. It is up to the project owners to decide whether these are realistic approaches.

References

- [1] ISO/IEC 17799 (renumbered into ISO/IEC 27002) Information technology – Security Techniques – Code of Practice for Information Security Management
<http://www.iso27001security.com/html/27002.html> (last visited 31 January 2008)
- [2] The CORAS project. <http://coras.sourceforge.net/> (last visited 31 January 2008)
- [3] Risk Management. Standards Association of Australia. AS/NZS 4360:1999
- [4] ISO/IEC 27005 Information technology – Security Techniques – Information Security Risk Management (currently a DIS – Draft International Standard).
<http://www.iso27001security.com/html/27005.html> (last visited 31 January 2008)
- [5] Bellika, J.G., et al. Propagation of program control: A tool for distributed disease surveillance. *Int. J. Med. Inform.*, 76 (4), April 2007, 313-329
- [6] Bellika, J.G., et al. Properties of a federated epidemiology query system. *Int. J. Med. Inform.*, 76 (9), Sept 2007, 664-676
- [7] Directive 95/46/EC on the protection of individuals with regard to the processing of personal data and on the free movement of such data. European Parliament and Council of the European Union, 24 Oct 1995.
<http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CELEX:31995L0046:EN:HTML>
(last visited 31 January 2008)
- [8] Act of 14 April 2000 no. 31 relating to the processing of personal data [Personal Data Act] <http://www.ub.uio.no/ujur/ulovdata/lov-20000414-031-eng.pdf>
LOV-2000-04-14-31 – Lov 14. april 2000 nr. 31 om behandling av personopplysninger (Personopplysningsloven).
<http://www.lovdata.no/all/hl-20000414-031.html>
- [9] Act on personal health data filing systems and the processing of personal health data [Personal Health Data Filing System Act]
<http://www.ub.uio.no/ujur/ulovdata/lov-20010518-024-eng.pdf>
LOV-2001-05-18-24 – Lov 18. mai 2001 nr. 24 om helseregistre og behandling av helseopplysninger (Helseregisterloven).
<http://www.lovdata.no/all/hl-20010518-024.html>
- [10] Act of 2nd July 1999, no 64 relating to health personnel etc. [The Health Personnel Act] http://www.regjeringen.no/en/dep/hod/Documents/lover_regler/reglement/2002/Act-of-2-July-1999-No-64-relating-to-Health-Personnel-etc.html?id=107079
LOV-1999-07-02-64 – Lov 2. juli 1999 nr. 64 om helsepersonell m.v. (Helsepersonelloven).
<http://www.lovdata.no/all/hl-19990702-064.html>
- [11] Article 29 group: 01248/07/EN WP 136 - Opinion 4/2007 on the concept of personal data, 20.juni 2007
http://ec.europa.eu/justice_home/fsj/privacy/docs/wpdocs/2007/wp136_en.pdf (last visited 31 January 2008)
- [12] ISO/IEC 27799 Health Informatics – Security Management in Health using ISO/IEC 17799 (→ using ISO/IEC 27002) (currently a draft)

Abbreviations

AgD	Agent Daemon (in Snow)
ALARP	As Low As Reasonable Possible
DoS	Denial of Service
EHR	Electronic Health Record
EU	European Union
GP	General Practitioner (doctor)
HW	Hardware
ID	Identifier / Identification
IM	Instant Messaging
JID	Jabber IDentifier
MC	Mission Controller
MUC	Multi-User Chat
PII	Personal Identifiable Information
PKI	Public Key Infrastructure
PO	Post Office (server in Snow)
QA	Quality Assurance
SAS	Snow Agent System
SW	Software
TCP	Transmission Control Protocol
UML	Unified Modelling Language
XML	eXtensible Mark-up Language

Annex A Threat table

The following table indicates likelihood, consequence and resulting risk level for each identified threat. (The values actually used in the analysis are indicated in **bold**.)

General comments, valuable for system design, have been gathered at the end of (below) this table.

ID	Threat / Unwanted incident	Cause	Likelihood	Consequence	Risk level	Comments and descriptions of implemented security measures
General threats						
g1	Open source software - use of open source modules - making open source sw		??	??	??	To be investigated
g2	False/fake sw modules can be installed on SAS server or in the GP office's systems - can do "all sorts of harm", to confidentiality, integrity, availability ex.: a false/fake client that sends a request to a false/fake agent that extracts patient id etc from EHR...		Low	Severe for the service, the trust and reputation of Snow	Low	Routines for upload, installation and/or upgrade of sw to SAS server at GP offices. - Should be compared to how the EHR system vendors install, upgrade and update their sw. This is part of the general information security strategy at the GP offices, like firewalls, virus protection, security updates, etc. The information (spec luggage) sent to the SAS server for a specific mission must not interfere with the SAS sw repository. Solutions? Dig. signature on installed modules

ID	Threat / Unwanted incident	Cause	Likelihood	Consequence	Risk level	Comments and descriptions of implemented security measures
Confidentiality <i>For all these threats: consequence = severe, because it is a violation of law if sensitive information falls into hands of unauthorised persons.</i>						
c1	The information extracted from the EHRs by the agents is sensitive, i.e. person identifiable information (not anonymous)		To be further investigated - but >Low if not taken care of in the design, e.g. Moderate?	Severe? - it depends on who sees this information	Medium?	Routines for restrictions of rare diagnoses, limited geographical areas, gender? Anonymisation is imposed, but it is a question whether the transferred information is sensitive or not. – We assume some data could be sensitive due to small geographical area (municipality, postal code area) and rare diagnosis. Where is the information controlled and assembled? Locally at GP or in PO?
c2a	Worst case: Sensitive information from EHR falls into hands of unauthorised persons	False/fake Agent Daemon (or Agent) that extracts sensitive information from EHR	Low? see g2 above	Severe - serious loss of reputation for service provider and patient	Low	Access to EHR must be limited wrt. which information is possible to extract. The malicious sw must then also modify the access rights to the EHR DB (config.?) Alternatives for database access: - Web Services (language independent), or a combination of stored procedures and view?
c2b		SW bugs: Erroneous Agent Daemon (or Agent) that extracts sensitive information from EHR	Very Low Can not extract more information than what has been defined as accessible.	Severe	Low	Programming errors: Testing! EHR access limited to the specific application, e.g. epidemiology. A programming error cannot violate that. (Is this a sort of configuration? What if the configuration is wrong – i.e. configuration error? How is the configuration set? Changed? → Routines needed...)

ID	Threat / Unwanted incident	Cause	Likelihood	Consequence	Risk level	Comments and descriptions of implemented security measures
c3	Sensitive information is exposed during transfer	Wiretapping, unauthorised persons "listening in" to the communication <i>(Other ways of doing this?)</i>	Very Low - because data will be encrypted during transfer	Severe	Low	End-to-end encryption during transfer, use of PKI solutions and session keys (TBD). Data is decrypted in PO, for processing. This means that PO must be secured/trusted.
c4	Intentional, but unauthorised, delivery of information	The GP does an intentional copy-paste operation from the EHR into a Jabber message and sends it to the JID for the receiver	Low? But this is a new tool that makes it easier to send info	Severe	Low	This is the same threat as for other services (e.g mail) available for the GP. - Could remove the possibility to paste data from EHR into messages? - Or a pop-up dialogue box asking the user for a confirmation of the paste operation, with default = "no" – and showing the text to be pasted.
c5	Unintentional delivery of information	Unintended copy-paste, or sending to the wrong JID	Low? - remains to be seen...	Severe	Low	(As above.)

ID	Threat / Unwanted incident	Cause	Likelihood	Consequence	Risk level	Comments and descriptions of implemented security measures
Availability						
a1a	<p>Worst case: The SAS system crashes the local EHR server (or the database)</p> <p>This is a wide problem, the result could be different:</p> <ul style="list-style-type: none"> a) Disk crash with destroyed data b) System/service is down for a period of time 	DoS-attack (intentional)	Low	Severe? due to loss of reputation and trust	Low	<p>“Karma” – ”black listing” of clients who sends (too) many requests.</p> <p>If disk crash: The GP office will/shall have an information security management system with necessary routines for e.g. backup. – Could it also be possible to have routines for retrieval of lost messages from the communication service? (Recent messages, after the last backup.)</p>
a1b		Erroneous sw, bugs e.g. that steal processing time and resources, e.g. never free memory...	Low Must assume that quality assurance and test procedures are good enough	Severe? due to loss of reputation and trust	Low	<p>Testing, test procedures, input verification.</p> <p>Bugs that steal processing time will cause problems for Snow but not for the local EHR. Area reserved for each process is normally protected from other’s area. Must assume that these kinds of weaknesses in Windows is detected and corrected...</p>
a2	DoS-attack crashes the Snow system	Malicious action from “internal” (<i>i.e. internally in the Snow service...</i>) (One way to do this is to repeat a request for data from a very long time period.)	Low	Moderate - maybe worse if the GPs are being dependant of Snow’s functionality	Low	<p>Motivation? E.g. to prevent a successful service or a successful result of the project... <i>How will this be discovered from the outside (from support)? PO performs a status control by frequent polling and logs and reports the (missing) responses.</i></p> <p>Post office will have the first pressure if the attack comes from the outside. PO can be more powerful than the computers at the GP offices. PO must control the amount of requests that are being forwarded, e.g. one at the time? (<i>Or queuing at the GP server?</i>)</p>

ID	Threat / Unwanted incident	Cause	Likelihood	Consequence	Risk level	Comments and descriptions of implemented security measures
a3a	Increased load on local system that significantly decreases the responsiveness	DoS-attack (intentional)	Low	Moderate - or small? Depends on how long the problem lasts	Low	See above, a1a and a2. A DoS-attack does not necessarily crash the system; but it may overload it.
a3b		SW features, e.g. too many missions/requests and corresponding agents executing simultaneously.	High - if not taken care of in sw development	Moderate - or small? Depends on how long the problem lasts	Medium	Errors, bugs, wrong configuration, etc. Missing limitations/restrictions. Many ways to restrict this type of load. – Could for instance restrict extensive requests (for a wide time period) to be performed only at night time. Max number of agents per Agent Daemon (configurable). A new agent will not start until there are available resources. E.g. during outbreaks: Many equal requests submitted at the same time. Cache in PO for “fresh” results. (<i>What if the requests are different?</i>) Testing: What <i>is</i> the system load? Deadlock or infinite loop: Timeout on unsuccessful actions, the corresponding agents are killed
a4	Unauthorised persons communicates with local Jabber server → increased load on system	JID for the Agent and Agent Daemon is known externally (sort of DoS-attack)	Low - must be client at an “approved” server	Moderate	Low	The JID is available only inside the “closed” system, i.e. inside the health net. No <i>relevant</i> communication from Internet into the health net. (Only well-defined application proxy- and relé-based services are let through the firewall.) Filters (in PO) to limit how widely the JID should be known. Filter to restrict who could contact these JIDs. (Karma... see a1a above)

ID	Threat / Unwanted incident	Cause	Likelihood	Consequence	Risk level	Comments and descriptions of implemented security measures
a5	Undue consumption of storage space in local system	Storage of old data/information	<p>Low</p> <p>- AgD will monitor the need for disk space and compare with the available disk space, before a new agent is launched</p>	<p>Moderate</p> <p>The disk will fill up quicker, and some applications could stop working</p>	Low	<p>Cache routines</p> <ul style="list-style-type: none"> - should old caches be stored (historic data)? Could instead generate results again for time periods in the past. Thus also avoiding extra storage of possible "sensitive" data. - historic data for research purposes? <p>Restrict extensive requests (e.g. annual statistics) to be performed only at night time and cache the results in PO.</p> <p><i>If the date & time field is not indexed, there is no difference in workload to make a request for one year than it is to request for one week.</i></p> <p>("No more data than it would be from debug-logging...")</p>
a6a	<p>SAS service unavailable due to other technical problems:</p> <ul style="list-style-type: none"> - Network problems 	- at GP office	<p>Moderate</p> <p>(the health net guarantees a network connection of 98,5 % for the GP offices, between 8am and 8pm)</p>	<p>Small?</p> <p>Depends on the situation.</p> <p>The coverage, completeness or "sensitivity" of the result gets worse, but the client can try again and get a better result.</p>	Low	<p>Has a certain consequence for reputation, trust in the service.</p> <p>Snow would assume that all GP offices are connected.</p> <p>Approx. 80% of GP offices in the pilot region are "always connected" (20% still ISDN).</p> <p>Standard SLA for Wanda-connected GP offices guarantees 99.5% availability on a monthly basis.</p> <p>Could skip the GP offices that do not respond, and/or do a retry on them later on in the round.</p> <p>The coverage or completeness ("sensitivity") of the result should be stated, i.e. the amount of GP offices that are included.</p>

ID	Threat / Unwanted incident	Cause	Likelihood	Consequence	Risk level	Comments and descriptions of implemented security measures
a6b		- at PO (in health net)	Low - less likely that the PO is not connected?	Moderate The PO <i>and</i> all the GPs in its area will be out of reach	Low	Has a certain consequence for reputation, trust in the service.
a7a	SAS service unavailable due to other technical problems: - HW failures	- at GP office	Low? Would assume that the GP offices have support contracts to deal with hw problems quickly	Small? - for the Snow service - it is not caused by Snow...	Low	
a7b		- at PO	(Very) low Less likely?	Moderate The PO <i>and</i> all the GPs in its area will be out of reach	Low	Has a certain consequence for reputation, trust in the service.
a8a	SAS service unavailable due to other technical problems: - SW errors	- at GP office	Difficult to predict...	Moderate? Or severe? Such errors will most likely exist in all GP installations	??	Has a certain consequence for reputation, trust in the service. Extensive testing and QA during the development process.
a8b		- at PO	Difficult to predict...	Moderate The PO <i>and</i> all the GPs in its area will be out of reach	??	

ID	Threat / Unwanted incident	Cause	Likelihood	Consequence	Risk level	Comments and descriptions of implemented security measures
a9	SAS service unavailable due to other technical problems: - Erroneous user actions	- at GP office	Difficult to predict... will need data about usage	Small Will hit only one GP office?	??	Has a certain consequence for reputation, trust in the service. Training and education. – It must not be too easy to change or damage the system by erroneous usage.
Integrity						
i1a	Worst case: The SAS service modifies data/information and relations in the EHR system, which causes modification to patient treatment.	False/fake sw modules (i.e. malicious sw)	Low?? (What should be the motivation?)	Severe	Low	The Snow modules do not need to have write access to the EHR database. This threat depends on which access rights that the EHR provider (here: Profdoc) give to Snow.
i1b		SW bugs/errors in the Snow system	Low? - difficult to predict.	Severe	Low	
i2	Data (Snow results) is being modified during transfer	- deliberately, by intruders in the network	Very Low	Moderate	Low	Encryption procedures – to be investigated <i>Is it more serious if it is the <u>requests</u> that are being modified?</i>
i3a	Production of false results from Snow	Wrong/fake info inserted in the EHR - deliberately (done by the GP?)	Low	Moderate	Low	Could cause panic... It is not possible for an outsider to join the system as a new GP office. It has to be managed in the PO etc, it has to be registered. Means that there is an admin job to be done for the PO and Main. - It is not possible for a fake GP to be connected to NHN and registered in HER. Before a GP is connected to the Health Network, the 9-digit business ID is verified against the national register in Brønnøysund, and a check is also performed against the Fastlege directory. - The client receives the result and interprets and presents it (in a map) → the client must verify the origin of the result (the sender)
i3b		False/fake clients, users or agents - the client can for instance receive the fake/false result as a chat message?	Low? (What should be the motivation?)	Moderate	Low	

ID	Threat / Unwanted incident	Cause	Likelihood	Consequence	Risk level	Comments and descriptions of implemented security measures
Data quality Must consider whether these threats really concerns information security... For further study – evaluation.						
q1a	Data from a limited set of GPs	“All” the GP offices (EHRs) are not online	??	??	??	EHR must be online 24/7 ? see above a6a
q1b		Too few GPs participates in the system because they do not trust/believe in the service	??	??	??	Motivation... Indicate the amount of GP offices participating, from system logs. Could also be done before the request: - Indicate the amount that is configured to “no”/“yes” for optional. Then the requester could determine if it is useful to launch the mission or not...
q2	Too old data	“All” the GP offices (EHRs) are not online	??	??	??	Data freshness: GP must be aware of how old/fresh the resulting information is. Cache routines to be defined. EHR must be online 24/7 ? See above a6a
q3	Data correctness is low	?? Misinterpretation of info, e.g. when performing free text search	??	??	??	GPs report to be afraid of getting “too many false positive”. Depends on use of correct code for diagnosis and symptom and that the GP updates the code in the EHR when it is confirmed from e.g. lab results.
Other threats						

Annex B Plan for implementation of security measures

Proposed measures for threats with risk level *Medium* and threats with severe consequence are listed in the following table. The table should be used in the follow-up of the risk treatment.

ID	Threats, unwanted incidents	Security measures	Responsible	Deadline	Status
g2	False/fake software modules can be installed on the Snow servers or in the GP's local systems	Routines for upload, installation, and upgrade of Snow sw to servers at the GP office (and to the PO) Digital signature on the uploaded sw. Spec luggage to the SAS server for a specific mission must not interfere with the SAS sw repository. Dedicated ports for input/output of such data to/from the server.			
c1	Sensitive information is extracted from the EHR by the agents, and communicated in the Snow system	Routines for restrictions of rare diagnoses, limited geographical areas, gender, etc. E.g. decide where the sensitive information is to be controlled and assembled: Locally at GP or in PO?			
c2a	Sensitive information from the GP's EHR is revealed to unauthorised persons by false/fake agents that are able to extract sensitive information from the local EHR	Access restrictions imposed on the EHR database. – This must be handled by the EHR system vendors. This also includes the possibilities to modify access rights (configuration). Decide on alternatives for database access: - Web Services (language independent) - combination of stored procedures and view			

ID	Threats, unwanted incidents	Security measures	Responsible	Deadline	Status
c2b	Sensitive information from the GP's EHR is revealed to unauthorised persons because errors/bugs in the Snow software makes it possible to extract sensitive information from the local EHR	Same as for c2a above. Plus: Quality assurance and extensive test procedures to avoid programming errors.			
c3	Sensitive information is exposed during transfer because of wire-tapping, unauthorised persons "listening in" to the communication	End-to-end encryption of transferred data between Snow server at GP and PO, using PKI solution and session keys. Routines for key administration have to be defined. Protection of PO, like firewall and virus protection. Data is decrypted in PO, for processing, so PO must be placed in a secured/trusted environment. This is assumed to be in place if PO is located in the Norwegian health net.			
c4	The GP <i>intentionally</i> performs a copy-paste operation from the EHR into a message which is submitted to the JID for a receiver	Information to users about legal aspects and possible risks (awareness). Education and training of users. Good user interfaces, user manual.			
c5	Unintentional delivery of information from GP, caused by an unintentional copy-paste, or by sending a message to a wrong receiver address (JID)	Remove the possibility to paste data from EHR into messages. – Or have a pop-up dialogue box asking the user for a confirmation of the paste operation, with default = "no" – and showing the text to be pasted. A more strict approach to these threats is <i>risk avoidance</i> : Do not implement the possibility to send messages.			

ID	Threats, unwanted incidents	Security measures	Responsible	Deadline	Status
a1a	A DoS attack through the Snow system crashes the local EHR server, which results in either a disk crash with destroyed data or the EHR system to be unavailable for a period of time	<p>The local system at the GP office is assumed to have an up-to-date information security management system with firewall, virus control, and regular security patching.</p> <p>The information security management system shall also have routines for backup and restore.</p>			
a1b	The local EHR server crashes because of errors/bugs in the Snow system, which results in either a disk crash with destroyed data or the EHR system to be unavailable for a period of time	<p>This should minimize the damage in case of disk crash.</p> <p>Black-listing of clients who sends (too) many requests ("Karma").</p> <p>Quality assurance and extensive test procedures to avoid programming errors.</p> <p><i>Risk transfer</i> is a different (and additional?) approach: Pay insurance to cover for possible economical losses of the GP if the EHR system is unavailable for a period of time.</p>			
a3b	Increased load on the local systems at the GP office, and correspondingly decreased responsiveness, because of features in the Snow system	<p>Black-listing of clients who sends (too) many requests ("Karma").</p> <p>Quality assurance and extensive test procedures to avoid programming errors, bugs, wrong configurations, and missing limitations/restrictions.</p> <p>Ways to reduce load:</p> <ul style="list-style-type: none"> - Extensive requests (for a wide time period) to be processed only at night time. - Max number of simultaneous agents per Agent Daemon (configurable). - Cache the "fresh" results in PO, for reuse. - Timeout on unsuccessful actions, and kill the corresponding agents. 			

ID	Threats, unwanted incidents	Security measures	Responsible	Deadline	Status
i1a	Malicious sw introduced through the Snow system causes modification of data/information and relations in the local EHR system, which results in wrong patient treatment	<p>The local system at the GP office is assumed to have an up-to-date information security management system with firewall, virus control, and regular security patching.</p> <p>As for c2a and a2b above:</p>			
i1b	Errors/bugs in the Snow system causes modification of data/information and relations in the local EHR system, which results in wrong patient treatment	<p>Access restrictions must be imposed on the EHR database, by the EHR system vendors. The Snow modules do not need to have write access to the EHR database.</p> <p>Alternatives for database access must be decided (Web Services or a combination of stored procedures and view).</p> <p>Quality assurance and extensive test procedures must be imposed, to avoid programming errors.</p> <p>Additionally, consider to pay insurance to cover for possible economical losses of the GP if the EHR system is modified so that it causes wrong patient treatment (<i>risk transfer</i>).</p>			

Annex C Design issues related to privacy and security

C.1 Preserve privacy of data from a small data set

The discussion in this section is related to threat **c1** – “Sensitive (i.e. person identifiable) information is extracted from the EHR by the agents, and communicated in the Snow system.”

This is the scenario: In a search for diagnosis/disease XY (e.g. HIV...) the agent of a local GP server returns the result “Number of cases in area 1234 is: 1”. And everyone knows that in area 1234 there are only 345 inhabitants. If in addition the gender and/or age are given, the number to choose from is much smaller. And if the number of cases changes at the same time as newcomers arrive to the area, this can be related.

The result to be presented at the end is the *aggregated* information for e.g. the whole city or county, but before it is aggregated it is transferred as a single result. → How can we ensure that the receiver of the single information (or a “hacker” who eavesdrop the communication) is not able to “guess” the identity of this unfortunate person?

The system must not allow presentation of information which is person identifiable in any ways. This can be avoided by necessary consideration taken during design of the information collection functionality and/or the result presentation functionality.

Collection of information:

There are two possible ways to collect the information from all the Snow “members”: **Jump** or **Spread**. Figure C1 illustrates the two methods. Because of the communication delay the Jump method is assumed to take longer time, while the Spread method gives a more instant view of the situation.




Figure C1a: Jump





Figure C1b: Spread



Using the Spread method, each node (GP Snow server) will return his single result; the result can be identified at the initiating node, or by an eavesdropper. Using the Jump method, each node will see the aggregated result from the preceding nodes. So still the second node will see the result of the first node... One way to prevent this is to initiate the request with a random “salt” value: The initiator starts the round with a random value, the next node adds his result to this value, and so on, and finally the initiator subtracts the “salt” value from the aggregated

value he receives at the end. A random “salt” value can also be used in the Spread case, this could prevent any eavesdropper from knowing the single result (... unless he also eavesdrop the request and reads the “salt” value...). But the initiator will still be able to identify the result from each single node.

Encryption can further protect eavesdropping. By use of asymmetric encryption (PKI), the information is encrypted with the public key of the receiver (which is found in his certificate). For the Jump method, this would be the public key of next node. For the Spread method this will be the public key of the initiator. It depends on who should perform the aggregation, and therefore needs to decrypt the information.

By use of XML encryption one can choose to encrypt only parts of the message. The same is the case for digital signing parts of the XML message.

For further study:

Here is another problem that should be discussed, but is too big to be solved in this project:

Independent of Jump or Spread method – how can we be sure that the same person/case is not counted twice – without communicating the person’s ID? The same case can be registered in an (active) EHR both at the person’s GP, at the emergency unit, and at the laboratory.

Gustav sketched a special pseudonymisation solution: Each person id (e.g. “personnr”) is mapped to a *set* of pseudonyms, as indicated in the table below. (*The size of n needs to be discussed – do we really need one per site, or can we reuse pseudonyms, e.g. choose among ten different?*)

Person ID	Pseudonym 1	Pseudonym 2	Pseudonym 3	...	Pseudonym n
A	PA1	PA2	PA3	...	PAn
B	PB1	PB2	PB3	...	PBn

Each site gets *one* of these pseudonyms for the person. When one node returns his result of a request, he could include the pseudonym(s) for the individuals. For example:

Node 1 says: - Number of cases: 1 - Pseudonyms: PA2

Node 2 says: - Number of cases: 2 - Pseudonyms: PA6, PB3

Node 3 says: - Number of cases: 1 - Pseudonyms: PA1

The initiator (PO) says: - Total number of cases: 2 (because he knows that PA1, PA2 and PA6 refer to the same person.)