Eliciting CM^3: Emergency Problem Management at Scandinavian Airline Systems

Mira Kajko-Mattsson

Department of Computer and Systems Sciences
Stockholm University and Royal Institute of Technology
Forum 100, SE 164 40 KISTA, Sweden
mira@dsv.su.se

Claus Nielsen, Per Winther, Brian Vang and Anne Petersen

Scandinavian Airline Systems (SAS Copenhagen)
SAS Dept. HB Commercial Systems, SAS Sales Systems & Support
Hedegaardsvej 88, 2300 Copenhagen, Denmark
claus.Nielsen@sas.dk, prodman@sas.dk, alp@sas.dk

Emergency software problems may present an immediate danger to public health, safety, general welfare or business. Hence, the organisations must be well prepared to handle them with the greatest expediency. Unfortunately, the software community has paid little attention to the emergency corrective maintenance. Today, we do not have any standard process models for handling emergency situations. In this paper, we outline an emergency corrective maintenance process model. The model is called CM^3: Emergency Problem Management. It is based on an industrial process model as defined at Scandinavian Airline Systems. In addition to the emergency process model, we present the status within the emergency process at Scandinavian Airline Systems, and describe the lessons learned.

ACM Classification: D.2 (Software Engineering), D.2.7 (Distribution, Maintenance, and Enhancement), D.2.9 (Management)

1. INTRODUCTION

Some unexpected software problems may present an immediate danger to people, environment, resources, general welfare or businesses (Bammidi *et al*, 1994). They often happen at inconvenient times or in inconvenient places. Many times, the harm they cause may render an organisation helpless, and at its worst, may lead to business closure (Andrews, 1994). Such problems are called emergency software problems. They are often unforeseen combinations of circumstances or unforeseen states that call for immediate actions to resume the system operation and to minimise its effects (Webster's, 1986).

The information about the emergency problems is often scarce, their causes are unclear and the actions to be taken are not always certain (Bammidi *et al*, 1994). Because life, environment, or business can be seriously threatened, rapid problem solution is critical.

The emergency problem solving is unusually challenging. One must act decisively, immediately, and with greatest expediency. Delaying action entails risk. For these reasons, it is pivotal to have a

Manuscript received: 22 December 2005
Communicating Editor: Hassan Reza

process and an appropriate organisational structure that is efficient and effective in supporting the management of emergency situations.

Despite its importance, emergency maintenance has received little attention within the software community. To the lead author of this paper's knowledge, there are no standard models whatsoever. Even the maintenance standards such as IEEE (1998) and ITIL (2003) do not provide any guidance. The only source of information can be found within the industry.

In this paper, we outline a process model of emergency corrective maintenance. Due to the fact that the model is part of a greater model *Corrective Maintenance Maturity Model (CM³)* (Kajko-Mattsson, 2001), we call it *CM³: Emergency Problem Management*. The model has been elicited at Scandinavian Airline Systems (SAS), the Nordic region's largest airline and travel group (SAS, 2005). Hence, this paper does not only provide an outline of an emergency process model but also reports on lessons learned at SAS.

In the remainder of this paper Section 2 outlines the process model and its constituents. Section 3 reports on how SAS implements their emergency process. Section 4 presents experience and lessons learned at SAS. Section 5 provides final remarks and gives suggestions for future work.

2. PROCESS MODEL STRUCTURE

CM³: Emergency Problem Management consists of six components. They are: (1) identification of the organisations/departments/teams involved in the process, (2) products to be managed by the emergency process, (3) roles involved in conducting the emergency activities, (4) point of contact through which one communicates all important emergency information, (5) emergency process and its phases, and (6) levels required for handling the emergency process.

Although most of these constituents are extant in any process model, they are extremely important within the emergency corrective maintenance. Inefficiencies in any of them may substantially affect the process results.

2.1 Identification of Organisations

Some software systems may be evolved and maintained by several organisations. They may also be integrated with many systems belonging to different organisations. Hence, there may be many organisations involved in an emergency process. Working effectively with all of them requires special care and co-ordination. For this reason, the very first step when defining an emergency process should be to identify all the organisations involved in the process.

2.2 Identification of Products

Not all products are critical to business. Therefore, as a next step, the organisations should identify the critical products to be encompassed by the emergency process. These products are usually safety-critical systems and business-critical systems.

In addition to this, the organisations should specify the severity of the problems encountered in these systems. Severity measures the effect of disruption and harm caused by a problem. It is used for indicating how seriously a software problem is perceived; the higher the severity, the bigger the harm.

Underestimating the impact of severity of the problem could be significantly detrimental to the organization. Therefore, the organisations should define a pertinent scale for recording the severity and determine which severity levels should be covered by the emergency process. They should also provide guidelines for how to assign severity values to software problems. If the emergency process covers several severity levels, then they should specify in what way the management of the emergency problems differs.

2.3 Designation of Roles

Various roles are involved in the emergency process. We group them into two classes: permanent and temporary roles. By permanent roles, we mean the roles exclusively dedicated to manage the emergency situations. They are: *Emergency Administrator*, *Emergency Process Manager*, *Task Force Leader* and *Task Force Team*. By temporary roles, we mean the roles temporarily involved in the emergency process. They are *Support Personnel*, *System Users*, *System Managers*, and other roles which either are responsible for the problematic system or are users of the system.

Irrespective of whether the roles are permanent or temporary, they should all have responsibilities assigned to them and they should make commitments to fulfil those responsibilities. Their job descriptions should reflect their responsibilities and assignments. Deviations from these responsibilities can seriously complicate the emergency process. To promote expedient problem resolution, a line of authority between the key personnel must be established and decision hierarchies should be identified and defined. Still, they should not be too rigid as to prevent the free flow of ideas. Below, we list and describe these roles. We map them on the process phases and operational levels described in Sections 2.5-2.6, respectively. The roles are the following:

- *Emergency Administrator:* This is a permanent role usually consisting of several people. It is the focal point of contact during the entire emergency process. One of the persons possessing this role is always on emergency duty. Her task is to accept and control all the emergency problem reports and take appropriate measures. As depicted in the upper part of Figure 1, this role is the owner of the problem in the first phase of the emergency process, the *Alert Level 1 – Normal Operation* phase.
- *Emergency Manager:* This is a permanent role responsible for keeping the descriptions of the emergency procedures updated and for performing periodic reviews together with the *Task Force Leader*. As depicted in the upper part of Figure 1, this role is also the owner of the problem in the second emergency phase, the *Alert Level 2 – Increased Attention* phase.
- *Task Force Leader:* This is a permanent role responsible for managing the resolution of the emergency problems. She leads the *Emergency Task Force Team* and is the owner of the problem in the third emergency phase, the *Alert Level 3 – Emergency Situation* phase (see the upper part of Figure 1).
- *Task Force Team:* This role consists of two groups of roles:
  - *Permanent Task Force Team Roles* consisting of key roles and management responsible for vital collaborating areas within the organisations involved. The permanent team is responsible for all the emergency situations.
  - *Temporary Task Force Team Roles* consisting of *System Managers*, *Support Personnel*, *Programmers*, and other roles vital for resolving the emergency problem. The *Task Force Team* is responsible for the overall problem resolution, co-ordination of the emergency activities and tracking of the problem resolution. The team is the owner of the problem during the *Alert Level 3 – Emergency Situation* phase, and stays so until the system has been stabilised. The group ensures that appropriate actions are taken by all the parties involved in order to re-establish the normal operation with a minimum delay.
- *Support Personnel:* The support personnel on *Support Line 1* and *2* to whom the customers report on the encountered problems (Kajko-Mattsson, 2003). It is this role who mainly reports on all the emergency problems to the *Emergency Administrator*.
- *System Manager:* The role responsible for the system in which the problem was encountered.
- *Programmer:* The role responsible for changing the affected code. Usually, this role belongs to the team managed by the *System Manager*.

2.4 Point of Contact

An emergency problem may be encountered in various ways by various roles such as end users, system managers, external organisations, or other. Each serious problem should be immediately reported to the relevant group which constitutes a focal point of contact. Such a point of contact is the *Emergency Administrator* role. One should also specify the administrator's availability, both within and outside office hours.

2.5 Process Phases

It is important to define an emergency process. The process spells out how the organisation responds to emergencies. As outlined in the upper part of Figure 1, the emergency process consists of five phases. The first three phases however, are the explicit emergency phases, during which one resolves the emergency problems. We call them *Emergency Phases* and they are *Alert Level 1 – Normal Operation*, *Alert Level 2 – Increased Attention*, and *Alert Level 3 – Emergency Situation*. Regarding the last two phases, the so-called *Post Emergency* phases, they are conducted after the problem has been resolved. They are *Emergency Closure* and *Emergency Follow Up*.

Due to the urgency of the situation, each process phase should be assigned a predetermined time window with specific goals. This window should provide enough time for completion of work, but not so much as to allow assignments to linger. Below, we briefly describe these phases.

- *Alert Level 1 – Normal Operation:* As soon as the *Support Personnel* gets a report on a serious problem, they are obliged to escalate it to the *Emergency Administrator*. Within a predetermined amount of time, for instance 10-30[1] minutes of the problem existence, the problem is managed by the *Emergency Administrator*, who is the focal point of contact towards all the problem submitters, customers, and all the parties concerned during the whole emergency process. She assesses the emergency situation, monitors user reactions and evaluates the impact and severity of the problem. It is the *Emergency Administrator* who initiates the emergency process.
- *Alert Level 2 – Increased Attention:* After some predetermined period of time, if the problem still causes trouble, it gets escalated to the next alert level, called *Alert Level 2 – Increased*

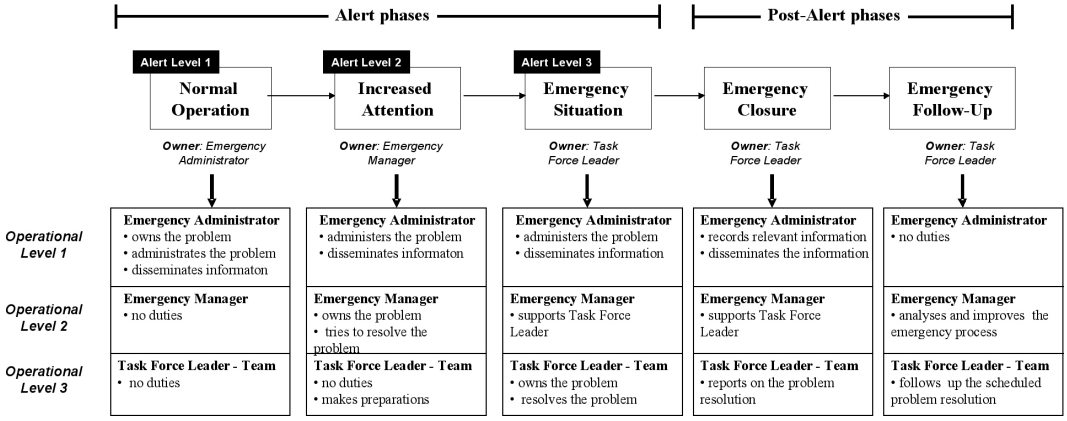

Figure 1: CM³ emergency process and operational levels

[1] This time limit varies across the organisations and products.

Attention. At this level, it is the *Emergency Manager* who becomes the owner of the problem. During this phase, the *Emergency Manager* and one or several *System Managers* evaluate a possible workaround for the problem.

- *Alert Level 3 – Emergency Situation:* After some predetermined period of time of the problem existence, the problem gets escalated to the *Alert Level 3 – Emergency Situation* phase. This phase ends when the problem gets resolved. It may take 20 minutes or it may take several days. In this phase, the *Task Force Leader* is in charge of the emergency situation. Her first action is to establish the *Task Force Team*. After the team gets assembled, it commonly makes an effort to solve the emergency problem.
- *Post-Emergency – Emergency Closure:* After the problem gets resolved, it steps into the *Post Emergency – Emergency Closure* phase. This task is mainly executed by the *Task Force Leader*. In this phase, she writes a report on the problem and distributes it to all the parties concerned.
- *Post-Emergency – Emergency Follow-Up:* During this phase, the *Task Force Leader* together with the *Emergency Manager* investigates the problem with the purpose of finding root causes underlying the problem. These causes provide an important feedback to process improvement.

2.6 Operational Levels

As depicted in the lower part of Figure 1, the whole emergency process is conducted on three operational levels. They are: (1) *Operational Level 1* managed by the *Emergency Administrator*, 2) *Operational Level 2* managed by the *Emergency Manager*, and (3) *Operational Level 3* managed by the *Task Force Leader* and *Task Force Team*.

The designation of these levels is very important. The process execution is strongly dependent not only on the emergency phase the process is in but also on the operational level. Each group of roles has clearly defined responsibilities for each phase and operational level. The decision making within the process is strongly dependent on the phase and the operational level. Below, we describe each operational level, map it on the process phases and list the responsibilities of the roles involved.

2.6.1. Operational Level 1

The *Operational Level 1* is mainly conducted by the *Emergency Administrator*. It is involved in four process phases, the three emergency phases and one post-emergency phase – *Emergency Closure*. The *Emergency Administrator* provides the first point of contact towards the users during the whole emergency process. Below, we describe her involvement within each emergency phase.

Alert Level 1 – Normal Operation

When a problem of a high severity is encountered, it is reported to the *Emergency Administrator*. The submitter may be any role such as *Support Personnel*, *System Manager*, *Customers* reporting on the problem via the *Support Personnel*, or other.

The reported problem should be first investigated and verified as serious. If it is of high severity, then the *Support Personnel* should immediately contact the *Emergency Administrator* and hand over the emergency handling.

The *Emergency Administrator* should confirm the problem, establish an internal emergency log, record relevant information in it, and distribute it to all parties concerned. The information should basically describe the problem, specify the time of its occurrence, its cause, impact, and other relevant data. Due to the fact that the problem is very serious, the *Emergency Administrator* must also determine a point in time when more information will be provided about the problem.

Alert Level 2 – Increased Attention and Alert Level 3 – Emergency Situation
During the *Alert Level 2 – Increased Attention* and *Alert Level 3 – Emergency Situation* phases, the *Emergency Administrator* continues administrating the problem, and informing all the parties concerned. However, at the *Alert Level 2*, she should escalate the problem to the *Operational Level 2* and pass over its ownership to the *Emergency Manager*.

Post-Emergency – Emergency Closure
When the emergency problem has been resolved, the *Emergency Administrator* records all the known information about the problem and informs all the parties concerned about its resolution.

2.6.2. Operational Level 2

The *Operational Level 2* is conducted by mainly two roles: *Emergency Manager* and *System Manager(s)*. The *Emergency Manager* has many responsibilities. One of them is to support the *Emergency Administrator* in all emergency situations. She also coordinates the workarounds received from the *System Manager(s)*. The responsibility of the *System Manager(s)*, on the other hand, is (1) to manage the problem resolution in their respective system parts, (2) to be available to the *Emergency Administrator* and the *Emergency Manager* and (3) to provide them with the necessary information about the system and its status.

Alert Level 1 – Normal Operation
At the *Normal Operation* phase, neither *Emergency Manager* nor *Task Force Leader* has any duties.

Alert Level 2 – Increased Attention
During the *Increased Attention* phase, the *Emergency Manager* becomes the owner of the problem. She should establish contact with the *Emergency Administrator* and continuously inform her about the problem status. She should also involve the *System Manager(s)* responsible for the affected systems or system parts in creating problem solutions or workarounds and distribute information to the relevant management.

Alert Level 3 – Emergency Situation
During the *Alert Level 3 – Emergency Situation* phase, the *Emergency Manager* should escalate the problem to the *Operational Level 3*, and thereby hand over the problem ownership to the *Task Force Leader*. She together with the *System Manager(s)* should however continue to support both the *Task Force Team* and the *Emergency Administrator*.

2.6.3 Operational Level 3

The *Operational Level 3* is mainly conducted by the *Task Force Leader* and *Task Force Team*. The number of participants in the team varies depending on the type and scope of the problem. If, for instance, three systems are involved, then it automatically implies that three *System Managers* and their teams are involved.

Alert Level 1 – Normal Operation and Alert Level 2 – Increased Attention
During the *Alert Level 1 – Normal Operation* phase, the *Task Force Leader* does not have any responsibilities. During the *Increased Attention* phase, she evaluates the situation in order to get prepared for the emergency situation.

Alert Level 3 – Emergency Situation

During the *Alert Level 3 – Emergency Situation* phase, the *Task Force Leader* establishes a *Task Force Team* and ensures that the team is in place. Afterwards, the course of actions varies depending on the problem. However, the *Task Force Leader* acts as a focal point of entry for all the management contacts, ensures that all parties concerned are informed, leads the *Task Force Team*, co-ordinates the emergency activities, initiates activities leading to the reduction of user impact, makes sure that the initiated activities are taken according to the defined procedures, initiates workarounds, and other activities leading to the resolution of the emergency problem. After the problem gets resolved, she decides when the normal operation should be resumed.

Emergency Closure

After the problem has been resolved, the *Task Force Leader* should produce a report on the emergency problem. The report should be based on the information collected in the emergency log. It should contain (1) time when the problem first occurred, (2) description of what happened and why, (3) description of the impact, (4) measures taken to limit the impact, (5) time stamp when the problem got resolved, (6) description of the measures taken in order to resolve the problem, (7) status of the emergency procedures used, (8) action list for changes to the emergency procedures and (9) suggestions for how to prevent similar situations.

During this phase, the software organisation may decide to restart the problem resolution process, in cases when it is deemed necessary. This time, however, the problem resolution will follow the planned and scheduled maintenance procedures.

Post-Emergency – Emergency Follow Up

During the *Post-Emergency – Follow-Up* phase, the *Task Force Leader* together with the *Emergency Manager* makes additional investigations of the problem and its causes. If the problem undergoes a planned and scheduled problem resolution, then they should follow its resolution.

In this phase, the *Task Force Leader* has regular meetings with the relevant roles and organisations or departments during which they follow up all problems of high severity. The goal is to find ways to avoid future emergency situations. Hence, a vital goal of this phase is to specify measures to prevent the problems from occurring. These measures should be recorded and delivered to the process improvement process.

3. STATUS AT SAS

In this section, we describe the emergency process at SAS. Due to the fact, that *CM[3]*: *Emergency Problem Management* has been induced from the SAS process, we only describe the differences.

3.1 Identification of Collaborating Organisations

SAS outsources development, evolution and maintenance activities to a separate organisation called Computer Sciences Corporation (CSC). However, they keep the ownership of all systems. Regarding problem management, SAS conducts it on *Support Line levels 1* and *2* and CSV on *Support Line 3* (Kajko-Mattsson, 2003). For this reason, as depicted in Figure 2, there are at least two organisations involved in all emergency cases at SAS. They are SAS and CSV. There may also be other organisations involved. One of the most important ones is Amadeus (Amadeus, 2005).

In cases when a problem affects Amadeus or comes from Amadeus and affects SAS, then Amadeus is involved in the emergency process as well. Otherwise, if the emergency problem at

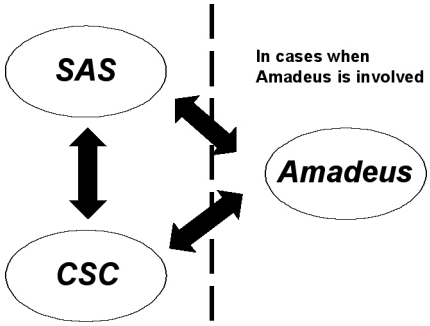

Figure 2: SAS and its collaborating organisations

SAS does not affect the Amadeus systems, then Amadeus is not involved in the emergency process at all. Amadeus is not even informed about the problem.

3.2 Identification of Products at SAS

At SAS, some systems are more critical than others. The most critical ones are those having commercial and business value. Examples of them are *Reservation Systems*, *Pricing*, *Ticketing*, *Internet Booking Platform* and *Telephony Systems*. These systems are used worldwide, 365 days per year, seven days a week, and 24 hours a day.

Emergency at SAS is invoked when problems of *Severity 1* or *2* are encountered. Such problems have the highest impact on the user community and system stability. SAS uses the severity definitions as defined by Amadeus. According to these definitions, a problem of *Severity 1* implies that a complete system, application or functionality is severely corrupted or degraded in service for a significant group of users. An example of such a problem might be the fact that one cannot book, or access, or obtain data from the *SAS Reservation System*, for instance, one cannot book a single SAS flight in the world. Such problems are very rare. Their resolution shall be managed at any time of the day, irrespective of whether it is during or outside the office hours.

A problem of *Severity 2* implies that a system, application, or functionality is severely degraded in service for a significant number of users. An example of such a problem might be the case when Japanese customers, defined as *Area Japan*, cannot book their flights. The resolution of such a problem shall be worked on within office hours.

3.3 Designation of Roles

SAS has designated their roles as defined by *CM³: Emergency Problem Management*. Their naming is different however. For instance, the role of the *Emergency Administrator* is called *HB System Support Function Copenhagen*. It consists of ten people. The role of the *Emergency Manager* is called *Production Management*. It consists of two people. The remaining roles are the same. Only one *CM³* role is missing. It is the *Programmer*. This is due to the fact that this role is active on the CSV side.

Each role is described and assigned clearly defined tasks. Hence, each person at SAS knows his role and responsibilities. In addition, SAS has created emergency call lists, listing all key persons who should be involved in responding to an emergency. A summarised list of tasks and responsibilities assigned to *HB System Support Function Copenhagen* (*Emergency Administrator* at SAS) at all alert levels is presented in Table 1.

HB System Support Function Copenhagen

Responsibilities

During the whole emergency process, they have the following responsibilities:

- Report on any problem by use of agreed means (SPAR, PTR, Telex, Phone).
- Inform all the parties involved.
- Liaise with SAS First Level Support, SAS International Sales Offices, CSC, Amadeus, and other organisations.
- Escalate the problem to the *Increased Attention – Alert Level 2* phase within SAS (after 30 minutes of problem occurrence).
- Notify Amadeus of any major interruption of the planned maintenance.
- Provide the first point of contact towards the users during the whole emergency process.
- Conduct and control the emergency process.

Involvement of Operational Level 1 within the process

Alert Level 1 – Normal Operation

At this level, the following should be done:

- Investigate the problem and identify and verify its severity. If the problem is of *Severity 1 or 2*, the support personnel or other role should immediately contact the HB System Support Function in Copenhagen and hand over the emergency handling.
- Record and report the problem to the vendor, either CSC or other companies.
- After five minutes of problem occurrence, update the relevant web pages, the so called HOT pages, on which important information is distributed to all parties concerned. The recorded information is the following:
  - Time when the problem occurred for the first time. Usually, it is the time stamp when the problem got reported in Copenhagen.
  - Cause of the problem.
  - Expected impact due to the problem.
  - Actions undertaken. At this early stage, one reports that the investigation is in progress, that a report has been created, and that the problem has been identified.
  - Possible work arounds, if any.
  - Expected time of closure.
  - Time for sending information about the problem next time to all the groups concerned. The goal of this telex is to make all the relevant parties aware of the problem.
- Establish internal emergency log for tracking purposes.
- Update the log with the copy of the above mentioned data and the logging time.

Alert Level 2 – Increased Attention

At this level, the following should be conducted:

- Provide up-to-date information to all parties concerned.
- Contact end-user submitters, via phone, in order to obtain a picture of the situation.
- Inform SAS Emergency Task Force Leader.
- Update the log with various information such as time, names of persons called (the users), the information recorded on the HOT pages, and the like.
- Update the HOT pages with information on problem cause, expected impact, actions undertaken, possible workarounds, expected time of closure, time when next information will be sent.
- Liaise with Production Management and System Management and provide up-to-date information to them.
- In case of ticketing problems when the problem is of Severity 1 and it is 40 minutes old, inform Amadeus about the problem, actions taken, and estimated time of problem resolution.
- Escalate the problem to the Operational Level 2 within SAS, or more exactly to Production Management.

Alert Level 3 – Emergency Situation

At this level, HB System Support Function Copenhagen should conduct the following:

- Contact the end-users in order to get a report on the problem status and to get a better understanding of the problem. This is only done when need arises. Sometimes, the users may be contacted every 30 minutes.
- Update the log with the relevant information such as (1) time of occurrence, (2) the time the action has taken, (3) action description, (4) copy of all mails, information pages and texts that have been sent, (5) the identification of the logger, and (6) names of people who can help collect data and who can provide important information.
- Provide up-to-date information to SAS users worldwide, via the HOT pages and group telex. This information is to be updated every 30 minutes and it should include the following: (1) latest status of the situation, (2) information that the HOT pages have been updated, (3) time when information about the problem will be sent next time.
- Communicate with CSC Support Centre and Amadeus Scandinavia about the status of the problem.
- Stay as a focal point of contact during the entire emergency situation.

Table 1: Responsibilities of HB System Support Function Copenhagen at three alert levels

3.4 Point of Contact at SAS

At SAS, the focal point of contact is *HB System Support Function Copenhagen*. If the problem is encountered during office hours, then *HB System Support Function Copenhagen* is contacted first (see Figure 3). Outside office hours, it is *CSV Support Centre* who is responsible for supervising all the emergency cases (see Figure 3).

As already mentioned, CSV is the first point of contact for all serious problems reported outside office hours. If a problem gets reported to them, then they should immediately report it to the manager of the *HB System Support Function Copenhagen* group who guards all the emergency calls coming in outside office hours and who is the first point of contact for *CSV Support Centre*.

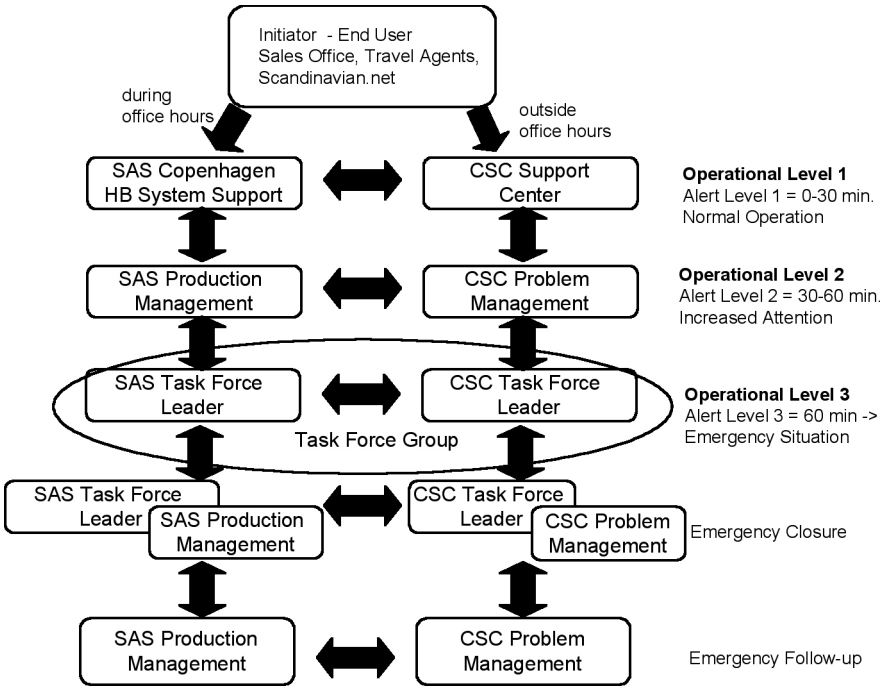

Figure 3: Operational levels at SAS and CSV

3.5 Processes at SAS

SAS has, in collaboration with CSV and Amadeus Distribution System, defined a common process model, so that each organisation can uniformly act in all emergency situations. SAS is however, the originator and main author of the process model.

With minor adjustments, the CM³ process phases reflect the process phases as conducted by SAS and its collaborating organisations. This is because the preliminary version of *CM³: Emergency Problem Management* presented in this paper is strongly based on the SAS process.

3.6 Operational Levels at SAS

The operational levels at SAS follow similar procedures as described in *CM³: Emergency Problem Management*. However, the process is run by at least two organisations SAS and CSV. In Figure 3, we illustrate how the operational levels at the two companies co-operate with one another. There, one may see that the CSV organisation has defined the corresponding roles at each operational level. In addition, the *Task Force Team* consists of roles coming from the two organisations. Due to space restrictions, we cannot describe these levels. However, we welcome interested readers to study our report on the emergency process at SAS (Kajko-Mattsson, 2005).

4. DISCUSSION

It has taken many years for SAS to achieve the present status of their process. When building and improving it, SAS worked in very small steps, made sure that these steps were possible to accomplish, that they were reliable enough to reflect the actual activities, and reliable enough to

allow SAS to continue with improvements. We feel that the SAS emergency process is now well established and that it appropriately mirrors its activities. It provides a substantial aid for monitoring the progress of their emergency corrective maintenance. Thanks to the well defined and established process, SAS has benefited in the following:

- *Gained control of the emergency process*. This has been achieved by defining the process model and by establishing different operational levels where each level has clearly defined roles and responsibilities.
- *Gained control of vendor emergency activities*. By commonly defining the process model with CSV and Amadeus, and by meticulously recording the process data, SAS has achieved good visibility into the process. Hence, SAS may easily compile various types of statistics. For example, in year 2004, SAS encountered seven problems of *Severity 1* and 61 of *Severity 2*. Out of the total of 68 problems, only five of them required *Task Force* management.
- SAS was able to focus on the impact on its business.
- *Lowered problem resolution costs*. Due to the fact that the system downtime is very costly (3 MSEK/hour for only one system ResAid), SAS has decreased investigation and correction time of the emergency problems, and hence, saved a lot of resources.
- *Increased availability of the emergency process*. The process is defined in such a way so that it is available 24 hours within 365 days.
- *Implemented preventive actions*. During the post-emergency phases, SAS conducts thorough analyses of the encountered problems and identifies measures to prevent them from happening in the future.
- *Increased availability of their systems*. By improving the quality of the emergency process, regular problem management process, and other development and enhancement processes, SAS has substantially increased the uptime of most of its business critical systems, such as the availability of Reservation system (>99,73), of Amadeus (> 99,75), and of Internet (> 99,50).

During recent years, SAS has gained a lot of experience and learned many lessons. Some of them are:

- It is of utmost importance that the process and the supporting tools must be agreed upon by all the parties involved. This requirement sounds trivial, however it is difficult to achieve, especially with external vendors in an outsourcing situation. The emergency process, the severity levels and its supporting tools are commonly defined and coordinated with great care by SAS and CSC. Also, aggressive time windows with specific goals for each alert phase have been determined providing enough time for completion of work, but not long enough to allow the process to linger. They are 30 minutes for the alert phases 1 and 2, and indefinite time for the alert phase 3. They are formally agreed upon in a System Maintenance Agreement. To be able to improve the process, the organisations meet periodically to discuss improvements and mutually agree on all the process changes. This ensures that the vendor does not make changes to the process inconsistent with the SAS needs.
- At SAS, a new problem emerged when SAS was spilt into several separate and independent companies. These companies have their own IT systems and use some of the SAS IT systems. To manage common emergency problems, these companies must share a common understanding of their internal processes. However, they have to accept the authority of the Task Force. SAS has decided that only one Task Force should lead the emergency process. It is the Airline IT Task Force with a proper representation from the organisations involved. SAS does

not accept that other Task Forces work in parallel with the Airline IT Task Force. SAS however accepts that the organisations involved may have their internal Task Forces reporting to the Airline IT Task Force.

- All the necessary roles at SAS and CSC have been defined. Each person knows what his role and responsibilities are, and has accepted her placement in the decision hierarchy. The personnel works as a unit with a commonly shared goal of solving the emergency problems. Despite its formality, the SAS process allows free flow of ideas. SAS has however, learned one lesson with respect to the roles involved in the process. Despite the fact that the process is well defined and documented, the roles involved seem to forget the process. For this reason, SAS needs to establish a regular education and training program of the key roles involved.
- Before, for problems with multiple causes, it was hard to get an overview of the causes and their possible solutions. For this reason, SAS has developed a template for describing the problems. This template is quite simple; however it provides a necessary overview of the emergency problem resolution. The template consists of the fields such as (1) possible cause, (2) action/ fixes, (3) result.
- By conducting historical analysis and comparing its results to the agreements made in SLAs, SAS has found that the figures delivered from the vendor for *Severity 1* and *2* incidents strongly vary. This mainly concerns the amount of recorded minutes of system downtime. The amount does not always reflect the actual system downtime as perceived by the customers. Table 2 provides evidence on the amount of unscheduled downtime for the Internet system. The improper recording of time has resulted in the delay of the Task Force activation which should be exactly 60 minutes after the incident occurrence. Here, the lesson learned is that SAS must better follow up the vendor's process to make sure that the inter-organisational emergency process reflects the actual actions taken, and that the process timing is properly recorded and followed.
- SAS has discovered that they lack proper disaster and recovery procedures. To remedy this, SAS will soon introduce such procedures where a disaster is defined as "A disaster is a *Severity 1* incident where normal procedures no longer are sufficient to re-establish normal production via the established organisation. Therefore dedicated procedures and organisation must be established under management supervision to restore normal production".

		Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
CSC	Internet Weighted	99.99	99.99	99.96	100	99.96	99.98	100	99.87	99.99	100	100	100
	Unsched downtime minutes	4.38	4.38	17.52	0.00	17.52	8.76	0.00	56.98	4.38	0.00	0.00	0.00
SAS	SAS Liability	99.84%	99.08%	99.97%	100.00%	99.92%	99.85%	100.00%	98.84%	99.81%	100.00%	100.00%	100.00%
	SAS SPAR minutes	71	401	12	0	36	64	0	508	82	0	0	0
	Spar reg Internet Sev.1	61	0	12	0	0	64	0	25	82	0	0	
	Spar reg Internet SBP Sev.1	0	0	0	0	0	0	0	0	0	0	0	
	Spar reg Internet Sev.2	0	0	0	0	0	0	0	0	0	0	0	
	Spar reg Internet SBP Sev.2	10	401	0	0	36	0	0	483	0	0	0	

Table 2: Listing the scheduled and unscheduled downtime

5. FINAL REMARKS

In this paper, we have presented *CM³: Emergency Problem Management* that has been elicited from the SAS emergency process. We have also described the emergency process at SAS and its collaboration with other companies. To the authors' knowledge, there are no publications describing the domain of emergency software problem management. By creating the model, we hope to achieve the following goals:

- to transfer industrial experience in order to provide a basic reference point from which future research into software maintenance can proceed,
- to provide the students and researchers with the sense of reality, size and perspective to help them understand the spectrum of emergency corrective software maintenance, and
- to provide guidance to other industrial organisations world-wide in the process of building or improving their problem reporting and resolutions processes.

The *CM³: Emergency Problem Management* process model presented in this paper is in its very initial phase. A lot of work remains to be done. First, the model will have to be mapped on other industrial processes. Second, the flow of data among the operational levels will have to be defined, the decision making procedures will have to be more detailed, measurements will have to be specified, and feedback should be defined between the emergency and post-emergency phases. Finally, the model will have to be integrated with the planned scheduled corrective maintenance process model. All this implies that a lot of work is awaiting us in the future.

ACKNOWLEDGEMENTS

We thank The Swedish Research Council (in Swedish Vetenskapsrådet) and Vinnova for their financial support of the projects: Software Maintenance Laboratory: Upfront Maintenance and SERVIAM. This study has been made possible thanks to their recognition of the importance of this area.

We also thank the SAS managers Björn Fagerstedt and Magnus Clarving for allowing us study their organisational processes. Finally, a great big thanks to the SERVIAM project leader, Peter Söderström, for effectively managing the SERVIAM project and for providing an excellent working milieu and conditions for conducting this study.

REFERENCES

AMADEUS (2005): Information about the AMADEUS organisation is available on http://www.amadeus.com/en/5060.jsp, 2005.

ANDREWS, R.A. (1994): An ounce of prevention: Guidelines for preparing a disaster recovery plan. *Proceedings of the IEEE National Aerospace and Electronics Conference*, 802–806.

BAMMIDI, P. and MOORE, K.L. (1994): Emergency management systems: A systems approach, *Proceedings of the IEEE Conference on Systems, Man, and Cybernetics*, 1565–1570.

IEEE, (1998): IEEE Standard for Software Maintenance, IEEE Std 1219–1998. *The Institute of Electrical and Electronics Engineers*, Inc., 1998.

ITIL (2003): ITIL service support, Version 2.0, *TSO for OGC (Office of Government Commerce)*, Crown.

KAJKO-MATTSSON, M. (2001): Towards a business maintenance model, *Proceedings of the International Conference on Software Maintenance*, IEEE Computer Society Press: Los Alamitos, CA, 500–509.

KAJKO-MATTSSON, M. (2003): Infrastructures of virtual IT enterprises, *Proceedings of the International Conference on Software Maintenance*, IEEE Computer Society Press: Los Alamitos, CA, 199–208.

KAJKO-MATTSSON, M., WINTHER, P., VANG, B. and PETERSEN, A. (2005): Emergency procedures at SAS, *Technical Report*, 2005-024, DSV, Stockholm University/KTH, Forum 100, SE-164 40 Kista, Sweden.

SAS (2005): Scandinavian Airline Systems (SAS), www.sas.se.

WEBSTER'S (1986): *Webster's Third New International Dictionary*, Merriam-Webster Inc., Publishers, Springfield, Massachusetts, USA.

BIOGRAPHICAL NOTES

Mira Kajko-Mattsson is a senior lecturer at Stockholm University and Royal Institute of Technology in Stockholm, Sweden. She received her PhD in software engineering in 2001. Dr Kajko-Mattsson is the creator of Corrective Maintenance Maturity Model (CM[3]). Her research interests include software evolution and maintenance processes. Her email is mira@dsv.su.se.

Mira Kajko-Mattsson

Claus W. L. Nielsen is the manager of Reservation and Distribution Systems at Commercial Systems at SAS Airline IT. Mr Nielsen has been working at SAS since 1975. He started at SAS Data and then he went over to Commercial Systems in 1993. He has been working as Test Manager, Project Leader, System Manager and Production Manager. His email is Claus.Nielsen@sas.dk.

Claus W.L. Nielsen

Per Winther is the application service manager at Commercial Systems at SAS Airline IT. Mr Winther has been working at SAS since 1978. He started at Commercial Systems in 1992 as System Supporter and Test Co-ordinator and has been working at the Production Management group since 2000. His email is Per.Winther@sas.dk.

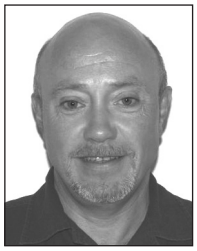

Per Winther

Anne Lylloff Petersen is the application manager at Commercial Systems at SAS Airline IT. Mrs Petersen has been working at SAS since 1983. She started as a System Manager in Commercial Systems in 1997 and has been working as an application manager since April 2005. Her email is Anne-L.Petersen@sas.dk.

Anne Lylloff Petersen