

3GPP TSG-SA5 (Telecom Management)
Meeting #12, Rome, 5-9 June 2000

SA5#12(00)0331

CHANGE REQUEST		<i>Please see embedded help file at the bottom of this page for instructions on how to fill in this form correctly.</i>	
32.111	CR	003	Current Version: 3.0.1
<i>GSM (AA.BB) or 3G (AA.BBB) specification number ↑</i>		<i>↑ CR number as allocated by MCC support team</i>	
For submission to: SA#8 <small><i>list expected approval meeting # here ↑</i></small>	for approval <input checked="" type="checkbox"/> for information <input type="checkbox"/>	strategic <input type="checkbox"/> non-strategic <input type="checkbox"/>	<i>(for SMG use only)</i>

Form: CR cover sheet, version 2 for 3GPP and SMG The latest version of this form is available from: <ftp://ftp.3gpp.org/Information/CR-Form-v2.doc>

Proposed change affects: (U)SIM ME UTRAN / Radio Core Network
(at least one should be marked with an X)

Source: SA5#12 **Date:** 20 June 2000

Subject: Split of TS - Part 1: Main part of spec - Alignment of FM requirements with IRP, etc.

Work item: 32.111 3G Fault Management

Category:	F Correction <input checked="" type="checkbox"/> A Corresponds to a correction in an earlier release <input type="checkbox"/> B Addition of feature <input type="checkbox"/> C Functional modification of feature <input type="checkbox"/> D Editorial modification <input type="checkbox"/>	Release:	Phase 2 <input type="checkbox"/> Release 96 <input type="checkbox"/> Release 97 <input type="checkbox"/> Release 98 <input type="checkbox"/> Release 99 <input checked="" type="checkbox"/> Release 00 <input type="checkbox"/>
------------------	--	-----------------	--

(only one category shall be marked with an X)

Reason for change: The following changes are proposed to be introduced in TS 32.111 Ver 3.0.1

1. Alignment of the FM requirements with the IRP actually defined for Release '99
2. Introduction of Test Management requirements
3. Editorial and error corrections.

Clauses affected: 3.1, 4, 4.1.1, 4.1.2, 4.1.3, 4.3, 5.1, 5.2.*, 5.3.*, 5.4 and 5.5

Other specs affected:	Other 3G core specifications <input type="checkbox"/> Other GSM core specifications <input type="checkbox"/> MS test specifications <input type="checkbox"/> BSS test specifications <input type="checkbox"/> O&M specifications <input type="checkbox"/>	→ List of CRs: → List of CRs: → List of CRs: → List of CRs: → List of CRs:
------------------------------	---	--

Other comments:

3 Definitions and abbreviations

3.1 Definitions

For the purposes of the present document, the following terms and definitions apply:

Active alarm: an alarm that has not been cleared. An alarm is active until the fault that caused the alarm is corrected and a clear alarm is generated

ADAC Faults = faults that are “Automatically Detected and Automatically Cleared” by the system when they occur and when they are repaired

ADMC Faults = faults that are Automatically Detected by the system when they occur and Manually Cleared by the operator when they are repaired

Alarm: an alarm is an abnormal network entity condition which categorises an event as a fault

Alarm notification: a notification used to inform the recipient about the occurrence of an alarm

Clear alarm: an alarm where the severity value is set to "cleared"

Event: this is a generic term for any type of occurrence within a network entity. A notification or event report may be used to inform one or more OS(s) about the occurrence of the event

Fault: a deviation of a system from normal operation. This deviation may result in the loss of operational capabilities of the element or the loss of redundancy in case of a redundant configuration

Notification: information message originated within a network entity to inform one or more OS(s) about the occurrence of an event

~~**Steady fault:** a steady fault is characterised by well defined conditions for the declaration of its presence or absence, i.e. fault occurrence and fault clearing conditions. This implies that the fault can be both detected and cleared automatically by the fault management functions of the network entity~~

~~**Unsteady fault:** an unsteady fault is characterised by a defined condition for the declaration of the fault, but no clearing condition exists. This implies that the fault can be detected but not cleared automatically by the fault management functions of the network entity~~

3.2 Abbreviations

-
-
-
-

4 Fault Management concept

Any evaluation of the network elements' and the overall network health status will require the detection of faults in the network and, consequently, the notification of alarms to the OS (EM and/or NM). Depending on the nature of the

fault, it may be combined with a change of the operational state of the logical and/or physical resource(s) affected by the fault. Detection and notification of these state changes is as essential as it is for the alarms. A list of active alarms in the network and operational state information as well as alarm/state history data are required by the system operator for further analysis. Additionally, test procedures can be used in order to obtain more detailed information if necessary, or to verify an alarm or state or the proper operation of NEs and their logical and physical resources.

The following subclauses explain the detection of faults, the handling of alarms and states changes and the execution of tests.

Only those requirements covered by section 5 and related IRPs shall be considered as valid requirements for compliance to the standard defined by this TS.

4.1 Faults and alarms

-
-
-
-

4.1.1 Fault detection

When any type of fault described above occurs within a 3G network, the affected network entities must be able to detect them immediately.

The network entities accomplish this task using autonomous self-check circuits/procedures, including, in the case of ~~NEs~~NEs, the observation of measurements, counters and thresholds. The threshold measurements may be predefined by the manufacturer and executed autonomously in the NE, or they may be based on performance measurements administered by the EM, cf. [4]. The fault detection mechanism as defined above shall cover both active and standby components of the network entities.

The majority of the faults will have well-defined conditions for the declaration of their presence or absence, i.e. fault occurrence and fault clearing conditions. Any such incident shall be referred to in the present document as a steady ADAC fault. The network entities ~~must~~ should be able to recognise when a previously detected steady ADAC fault is no longer present, i.e. the clearing of the fault, using similar techniques as they use to detect the occurrence of the fault. For some faults, no clearing condition exists. For the purpose of the present document, these faults shall be referred to as ~~unsteady~~ ADMC faults. An example of this is when the network entity has to restart a software process due to some inconsistencies, and normal operation can be resumed afterwards. In this case, although the inconsistencies are cleared, the cause of the problem is not yet corrected. Manual intervention by the system operator will always be necessary to clear ~~unsteady~~ ADMC faults since these, by definition, cannot be cleared by the network entity itself.

For some faults there is no need for any short-term action, neither from the system operator, nor from the network entity itself, since the fault condition lasted for a short period of time only and then disappeared. An example of this is when an NE detects the crossing of some observed threshold, and in the next sampling interval, the observed value stays within its limits.

For each fault, the fault detection process shall supply the following information:

- the device/resource/file/functionality/smallest replaceable unit as follows:
 - for hardware faults, the smallest replaceable unit that is faulty;
 - for software faults, the affected software component, e.g. corrupted file(s) or databases or software code;
 - for functional faults, the affected functionality;
 - for faults caused by overload, information on the reason for the overload;

- for all the above faults, wherever applicable, an indication of the physical and logical resources that are affected by the fault if applicable, a description of the loss of capability of the affected resource.
- the type of the fault (communication, environmental, equipment, processing error, quality of service) according to [9];
- the severity of the fault (indeterminate, warning, minor, major, critical), as defined in [9];
- the probable cause of the fault;
- the time at which the fault was detected in the faulty network entity;
- the nature of the fault, i.e. steady-ADAC or unsteadyADMC;
- any other information that will help understanding the cause and the location of the abnormal situation (system/implementation specific).

For some faults, additional means, such as test and diagnosis features, may be necessary in order to obtain the required level of detail. See subclause 4.3 for details.

4.1.2 Generation of alarms

For each detected fault, appropriate alarms shall be generated by the faulty network entity, regardless of whether it is a steady-ADAC or unsteadyADMC fault. Such alarms shall contain all the information provided by the fault detection process as described in subclause 4.1.1.

In order to ease the fault localisation and repair, the faulty network entity should generate for each single fault, one single alarm, also in the case where a single fault causes a degradation of the operational capabilities of more than one physical or logical resource within the network entity. An example of this is a hardware fault which affects not only a physical resource but also degrades the logical resource(s) that this hardware supports. In this case the network entity ~~shall~~ should generate one single alarm for the faulty resource (i.e. the resource which needs to be repaired) and a number of events related to state management (cf. subclause 4.2) for all the physical/logical resources affected by the fault, including the faulty one itself.

In case a network entity is not able to recognise that a single fault manifests itself in different ways, the single fault is detected as multiple faults and originates multiple alarms. In this case however, when the fault is repaired the **network entity** ~~must~~ should be able to detect the repair of all the multiple faults and clear the related multiple alarms.

When a fault occurs on the connection media between two NEs or between a NE and an OS, and affects the communication capability between such NE/OS, each affected NE/OS will detect the fault as described in subclause 4.1.1 and generate its own associated communication alarm toward the managing OS. In this case it is the responsibility of the OS to correlate alarms received from different NEs/OSs and localise the fault in the best possible way.

Within each NE, all alarms generated by that NE shall be input into a list of active alarms. The NEs must be able to provide such a list of active alarms to the OS when requested.

4.1.3 Clearing of alarms

The alarms originated in consequence of faults need to be cleared. To clear an alarm it is necessary to repair the corresponding fault. The procedures to repair faults are implementation dependent and therefore they are out of the scope of the present document, however, in general:

- the equipment faults are repaired by replacing the faulty units with working ones;
- the software faults are repaired by means of partial or global system initialisations, by means of software patches or by means of updated software loads;

- the communication faults are repaired by replacing the faulty transmission equipment or, in case of excessive noise, by removing the cause of the noise;
- the QOS faults are repaired either by removing the causes that degraded the QOS or by improving the capability of the system to react against the causes that could result in a degradation of the QOS;
- Solving the environmental problem repairs the environment faults (high temperature, high humidity, etc.).

It is also possible that a ~~steady-ADAC~~ fault is spontaneously repaired, without the intervention of the operator (e.g. a threshold crossed fault). In this case the NE behaves as for the ~~steady-ADAC~~ faults repaired by the operator.

In principle, the NE uses the same mechanisms to detect that a fault has been repaired, as for the detection of the occurrence of the fault. However, for ~~unsteady-ADMC~~ faults, manual intervention by the operator is always necessary to clear the fault. Practically, various methods exist for the system to detect that a fault has been repaired and clear alarms and the faults that triggered them. For example:

- The system operator implicitly requests the NE to clear a fault, e.g. by initialising a new device that replaces a faulty one. Once the new device has been successfully put into service, the NE will clear the fault(s). Consequently, the NE will clear all related alarms.
- The system operator explicitly requests the clearing of one or more alarms. Once the alarm(s) has/have been cleared, the NE will detect that the fault condition has ceased.
- The NE detects the exchange of a faulty device by a new one and initialises it autonomously. Once the new device has been successfully put into service, the NE will clear the fault(s). Consequently, the NE will clear all related alarms.
- The NE detects that a previously reported threshold crossed alarm is no longer valid. It will then clear the corresponding active alarm and the associated fault, without requiring any operator intervention. The details for the administration of thresholds and the exact condition for the NE to clear a threshold crossed alarm are implementation specific and depend on the definition of the threshold measurement, see also subclause 4.1.1.
- ~~Unsteady-ADMC~~ faults/alarms can, by definition, not be cleared by the NE autonomously. Therefore, in any case, system operator functions shall be available to request the clearing of ~~unsteady-ADAC~~ alarms/faults in the NE. Once an ~~unsteady-ADMC~~ alarm/fault has been cleared, the NE will clear the associated ~~unsteady-ADAC~~ fault/alarm.

Details of these mechanisms are system/implementation specific.

Each time an alarm is cleared the NE shall generate an appropriate clear alarm event. A clear alarm is defined as an alarm, as specified in subclause 4.1.2, except that its severity is set to "cleared". The relationship between the clear alarm and the active alarm is established:

- by re-using a set of parameters that uniquely identify the active alarm (cf. subclause 4.1.2); or
- by including a reference to the active alarm in the clear alarm.

When a clear alarm is generated the corresponding active alarm is removed from the active alarm list.

4.1.4 Alarm forwarding and filtering

-
-
-
-

4.3 ~~4.3~~ Test management

This management function provides capabilities that can be used in different phases of the fault management. For example:

- when a fault has been detected and if the information provided through the alarm report is not sufficient to localise the faulty resource, tests can be executed to better localise the fault;
- during normal operation of the NE, tests can be executed for the purpose of detecting faults;
- once a faulty resource has been repaired or replaced, before it is restored to service, tests can be executed on that resource to be sure that it is fault free.

However, regardless of the context where the testing is used, its target is always the same: verify if a system's physical or functional resource performs properly and, in case it happens to be faulty, provide all the information to help the operator to localise and correct the faults.

Testing is an activity that involves the operator, the managing system (the OS) and the managed system (the NE). Generally the operator requests the execution of tests from the OS and the managed NE autonomously executes the tests without any further support from the operator.

In some cases, the operator may request that only a test bed is set up (e.g. establish special internal connections, provide access test points, etc.). The operator can then perform the real tests which may require some manual support to handle external test equipment. Since the "local maintenance" and the "inter NE testing" are out of the scope of this TS, this aspect of the testing will not be treated any further.

The requirements for the testing service component are based on ITU-T X.745 [14], where the testing description and definitions are specified.

x Fault management requirements

-
-
-
-

5 N interface

5.1 Fault Management concept of Itf-N

An operations system on the network management layer (i.e. the NM) provides fault management services and functions required by the 3G operator on top of the element management layer.

~~As pointed out in clause 5,~~ The N interface (Itf-N) may connect the network management system either to EMs or directly to the NEs. This is done by means of IRPs. In the following, the term "subordinate entities" defines either EMs or NEs, which are in charge of supporting the N interface.

This clause describes the properties of an interface enabling a NM to supervise a 3G-telecommunication network including - if necessary - the managing EMs. To provide to the NM the fault management capability for the network implies that the subordinate entities have to provide information about:

- events and failures occurring in the subordinate entities;

- events and failures of the connections towards the subordinate entities and also of the connections within the 3G network;
- the network configuration (due to the fact that alarms and related state change information are always originated by network resources, see [1]). This is, however, not part of the FM functionality.

Therefore, for the purpose of fault management the subordinate entities send notifications to a NM indicating:

- alarm reports (indicating the occurrence or the clearing of failures within the subordinate entities), so that the related alarm information can be updated;
- state change event reports, so that the related (operational) state information can be updated. This is, however, not part of the FM functionality.

The forwarding of these notifications is controlled by the NM operator using adequate filtering mechanisms within the subordinate entities.

The Itf-N provides also means to allow the NM operator the storage ("logging") and the later evaluation of desired information within the subordinate entities.

The retrieval capability of alarm-related information concerns two aspects:

- retrieval of "dynamic" information (e.g. alarms, states), which describes the momentary alarm condition in the subordinate entities and allows the NM operator a synchronisation of its alarm overview data;
- retrieval of "history" information from the logs (e.g. active/clear alarms and state changes occurred in the past), which allows the evaluation of events that may have been lost, e.g. after an Itf-N interface failure or a system recovery.

As a consequence of the requirements described above, both the NM and the subordinate entity must be able to initiate the communication.

5.2 Management of alarm and state change event reports

5.2.1 Mapping of alarm and related state change event reports

The alarm and state change reports received by the NM relate to functional objects in accordance with the information model of Itf-N. This information model tailored for a multi-vendor capability is different from the information model of the EM-NE interface (if an EM is available) or from the internal resource modelling within the NE (in case of direct NM-NE interface), thus a mapping of alarm and related state change event reports is performed by a mediation function within the subordinate entity.

The mediation function translates the original alarm/state change event reports (which may contain proprietary parameters or parameter values) taking into account the information model of the Itf-N.

If a mediation application function is needed, it works according to the following principles:

- Every alarm notification generated by a functional object in a subordinate entity is mapped to an alarm report of the corresponding ("equivalent") functional object at the Itf-N. If the functional object generating the original alarm notification has not a direct corresponding object at the Itf-N, the mediation functions maps the alarm to the next superior functional object in accordance with the containment tree of the Itf-N.
- Every state change notification generated by a functional object in a subordinate entity is mapped to a state change report of the corresponding ("equivalent") functional object at the Itf-N. If the functional object generating the original state change notification has not a direct corresponding object at the Itf-N, the mediation functions maps the alarm to the next superior functional object in accordance with the containment tree of the Itf-N.

Every alarm notification generated by a manufacturer-specific, equipment-related object in the subordinate entity is mapped to an alarm report of a generic logical object, which models the corresponding equipment-related resource.

~~NOTE:—In some cases a failure or the locking of an equipment related object implies also the change of the operational state of its corresponding functional object within the NE or EM (if EM is available). The mapping of this state change notification to an alarm or state change notification of the corresponding functional object at the Itf-N is subject of further study.~~

~~On the Itf-N the correlation between functional related and the generic logical objects (modeling equipment related network resources) is performed explicitly by means of a relationship attribute in the functional object class definition.~~

~~With regard to the multi-vendor capability of the Itf-N, this mapping concept combines the following requirements:~~

- ~~—Precise information about manufacturer specific, equipment related failures for the NM operator in charge of network maintenance (this information is provided in some parameters of alarm reports mapped to the generic logical objects).~~
- ~~—If functionality is affected, an additional alarm report concerning the related functional object is provided for the NM operator in charge of network's quality of service.~~

~~If possible, the two types of alarm reports generated by the mediation function shall be correlated.~~

5.2.2 Real-time forwarding of event reports

If the Itf-N is in normal operation (the NM connection to the subordinate entities is up), ~~alarm and related state change event~~ reports are forwarded in real-time to the NM via appropriate filtering located in the subordinate entity. These filters may be controlled either locally or remotely by the managing NM (via Itf-N) and ensure that only the event reports which fulfil pre-defined criteria can reach the superior NM. In a multi-NM environment each NM must have an own filter within every subordinate entity which may generate notifications.

5.2.3 Alarm clearing

On the Itf-N, alarm reports containing the value "cleared" of the parameter perceivedSeverity are used to clear the alarms. The correlation between the clear alarm and the related active alarms is performed by means of unambiguous identifiers.

This clearing mechanism ensures the correct clearing of alarms, independently of the (manufacturer-specific) implementation of the mapping of alarms/state change events in accordance with the information model of the Itf-N.

5.3 Retrieval of alarm ~~and state~~ information

The retrieval of alarm ~~and state~~ information comprises two aspects:

a) Retrieval of current information

This mechanism shall ensure data consistency about the current alarm/~~state change~~ information between the NM and its subordinate entities and is achieved by means of a so-called synchronisation ("alignment") procedure, triggered by the NM. The synchronisation is required after every start-up of the Itf-N, nevertheless the NM may trigger it at any time.

b) Logging and retrieval of history information

This mechanism offers to the NM the capability to get the alarm/~~state change~~ information stored within the subordinate entities for later evaluation.

5.3.1 Retrieval of current alarm information on NM request

The present document defines a flexible, generic synchronisation procedure, which fulfils the following requirements:

- The alarm information provided by means of the synchronisation procedure shall be the same (at least for the mandatory parameters) as the information already available in the alarm list. The procedure shall be able to assign the received synchronisation-alarm information to the correspondent requests, if several synchronisation procedures triggered by one NM run at the same time.
- The procedure shall allow the NM to trigger the start at any time and to recognise unambiguously the end and the successful completion of the synchronisation.
- The procedure shall allow the NM to discern easily between an "on-line" (spontaneous) alarm report and an alarm report received as consequence of a previously triggered synchronisation procedure.

~~NOTE: This requirement is for further investigation.~~

- The procedure shall allow the NM to specify filter criteria in the alignment request (e.g. for a full network or only a part of it).
- The procedure shall support connections to several NM and route the alignment-related information only to the requesting NM.
- During the synchronisation procedure new ("real-time") alarms may be sent at any time to the managing NM.

If applicable, an alarm synchronisation procedure may be aborted by the requesting NM. ~~(This requirement is for further investigation.)~~

~~5.3.2 Retrieval of current state change information on NM request~~

~~The requirements defined above for the alarm synchronisation procedure are valid analogously for the retrieval of current state change information as well.~~

~~Nevertheless the state change synchronisation procedure takes into account only the object instances whose state information is different from a combined default state. As combined default state the following values (according to [7]) shall be used:~~

- ~~—Operational state: enabled.~~
- ~~—Administrative state: unlocked.~~
- ~~—Usage state: idle (if supported).~~

5.3.3 Logging and retrieval of alarm and state change history information on NM request

The alarm/state change history information may be stored in the subordinate entities in dependence on the NM requirements. The NM is able to create logs for alarms/state change event reports and to define the criteria for storage of alarm/state change information according to [11].

~~The subsequent retrieval of stored information is possible on NM request in two different ways:~~

- ~~—via a read command with appropriate filtering;~~
- ~~—via bulk data transfer, using standardised file transfer procedures, as mentioned in subclause 5.1.2.~~

Nevertheless these particular requirements are not specific for alarm or state change information.

5.4 Co-operative alarm acknowledgement on the Itf-N

The acknowledgement of an alarm is a maintenance function that aids the operators in his day to day management activity of his network. An alarm is acknowledged by the operator to indicate he has started the activity to resolve this specific problem. In general a human operator performs the acknowledgement, however a management system (NM or EM) may automatically acknowledge an alarm as well.

The alarm acknowledgement function requires that:

- a) All involved OSs have the same information about the alarms to be managed (including the current responsibility for alarm handling).
- b) All involved OSs have the capability to send and to receive acknowledgement messages associated to previous alarm reports.

A co-operative alarm acknowledgement means that the acknowledgement performed at EM layer is notified at NM layer and vice versa, thus the acknowledgement-related status of this alarm is the same across the whole management hierarchy.

The co-operative alarm acknowledgement on Itf-N shall fulfil the following requirements:

- Acknowledgement messages may be sent in both directions between EMs and NM, containing the following information:
 - Correlation information to the alarm just acknowledged. - Acknowledgement history data, including the current alarm state (active | cleared), the time of alarm acknowledgement and, as configurable information, the management system (EM | NM) and the operator in charge of acknowledgement (the parameter operator name or, in case of auto-acknowledgement, a generic system name).
 - Acknowledgements notifications sent to NM shall be filtered with the same criteria applied to the alarms.
- ~~• The alarm acknowledgement procedure on the Itf-N shall cope with different customer requirements concerning the acknowledgement competence between operators working at EMs and NM. This matter may be managed by means of a "competence type" information, which may be controlled by every connected EM.~~
- ~~• Every time the communication between the two management systems is established, the NM is able to determine which OS is most suitable to handle the acknowledgement of alarms.~~
- Taking into account the acknowledgement functionality, the above described synchronisation procedure for retrieval of current alarm information on NM request may be extended. Additionally to the requirements defined in subclause 8.3.1, this extended synchronisation procedure relates not only to the active, but also to the "cleared and not acknowledged" alarms, which have still to be managed by the EM.

5.5 Overview of IRPs related to fault management

The N interface is built up by a number of IRPs. The basic structure of the IRPs is defined in [2] and [3].

For the purpose of Fault Management the following IRPs are needed:

- Alarm IRP, see part 2
- Notification IRP, see [1]
- Log IRP
(NOTE: This IRP ~~may is~~ not be part of Release 1999, therefore the requirements related to the log functionality are not valid for Release 1999)

Annex A (informative): Change history

Change history					
TSG SA#	Version	CR	Tdoc SA	New Version	Subject/Comment
S_07	2.0.0	-	SP-000013	3.0.0	Approved at TSG SA #7 and placed under Change Control
Mar 2000	3.0.0			3.0.1	cosmetic