INTERNATIONAL TELECOMMUNICATION UNION

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

M.2140

(02/2000)

SERIES M: TMN AND NETWORK MAINTENANCE: INTERNATIONAL TRANSMISSION SYSTEMS, TELEPHONE CIRCUITS, TELEGRAPHY, FACSIMILE AND LEASED CIRCUITS

International transport network

Transport network event correlation

ITU-T Recommendation M.2140

(Formerly CCITT Recommendation)

ITU-T M-SERIES RECOMMENDATIONS

TMN AND NETWORK MAINTENANCE: INTERNATIONAL TRANSMISSION SYSTEMS, TELEPHONE CIRCUITS, TELEGRAPHY, FACSIMILE AND LEASED CIRCUITS

Introduction and general principles of maintenance and maintenance organization	M.10–M.299
International transmission systems	M.300–M.559
International telephone circuits	M.560–M.759
Common channel signalling systems	M.760–M.799
International telegraph systems and phototelegraph transmission	M.800–M.899
International leased group and supergroup links	M.900–M.999
International leased circuits	M.1000–M.1099
Mobile telecommunication systems and services	M.1100–M.1199
International public telephone network	M.1200–M.1299
International data transmission systems	M.1300–M.1399
Designations and information exchange	M.1400–M.1999
International transport network	**M.2000–M.2999**
Telecommunications management network	M.3000–M.3599
Integrated services digital networks	M.3600–M.3999
Common channel signalling systems	M.4000–M.4999

For further details, please refer to the list of ITU-T Recommendations.

ITU-T Recommendation M.2140

Transport network event correlation

Summary

This ITU-T Recommendation is a functional description of the operations functions required to achieve Network Event Correlation in the surveillance of international transport networks. The capability described is to correlate notifications of faults and impairments from Network Elements and to reveal the nature, location and severity of the faults and impairments and their true effect on the network.

This Recommendation is based on the TMN Functional Architecture as specified in ITU-T Recommendation M.3010. This Recommendation identifies the TMN Management Function Sets from ITU-T Recommendation M.3400 that are involved in Transport Network Event Correlation, provides functional descriptions of each Management Function Set, and illustrates how they interact to achieve Transport Network Event Correlation.

Source

ITU-T Recommendation M.2140 was prepared by ITU-T Study Group 4 (1997-2000) and approved under the WTSC Resolution 1 procedure on 4 February 2000.

Keywords

Alarm Correlation, Event Correlation, Fault Management, Functional Architecture, International Transmission Entities, Network Management, Performance Management, Surveillance, Telecommunications Management Network (TMN).

FOREWORD

The International Telecommunication Union (ITU) is the United Nations specialized agency in the field of telecommunications. The ITU Telecommunication Standardization Sector (ITU-T) is a permanent organ of ITU. ITU-T is responsible for studying technical, operating and tariff questions and issuing Recommendations on them with a view to standardizing telecommunications on a worldwide basis.

The World Telecommunication Standardization Conference (WTSC), which meets every four years, establishes the topics for study by the ITU-T study groups which, in turn, produce Recommendations on these topics.

The approval of ITU-T Recommendations is covered by the procedure laid down in WTSC Resolution 1.

In some areas of information technology which fall within ITU-T's purview, the necessary standards are prepared on a collaborative basis with ISO and IEC.

NOTE

In this Recommendation, the expression "Administration" is used for conciseness to indicate both a telecommunication administration and a recognized operating agency.

INTELLECTUAL PROPERTY RIGHTS

ITU draws attention to the possibility that the practice or implementation of this Recommendation may involve the use of a claimed Intellectual Property Right. ITU takes no position concerning the evidence, validity or applicability of claimed Intellectual Property Rights, whether asserted by ITU members or others outside of the Recommendation development process.

As of the date of approval of this Recommendation, ITU had not received notice of intellectual property, protected by patents, which may be required to implement this Recommendation. However, implementors are cautioned that this may not represent the latest information and are therefore strongly urged to consult the TSB patent database.

CONTENTS

Introduction

Planning, developing and supporting operations functions in Network Elements (NEs) and management systems, and linking them together via well-defined open interfaces, are critical steps to introducing new technologies and services into the network and providing network operations support across boundaries between network/service providers.

One of the key areas for network/service providers to automate with management systems supporting open interfaces is the area of service assurance. Service assurance is increasingly driven by the ability of the network to detect and report service affecting conditions before the customer reports a trouble. The network provides a wealth of information about its performance and the occurrence of service affecting conditions. Existing Operations Systems can process this information to some degree, but not completely, and rarely across different layers of a layered network or across traditional operations centre boundaries. As a result, service assurance today requires significant manual effort to properly interpret the information reported by NEs. The key to driving cost out of service assurance and improving its response time is to further automate the interpretation and correlation of events in complex layered networks.

The intended audience of this Recommendation includes:

- network operators interested in deploying management systems that perform Network Event Correlation within a TMN (operations and maintenance planners, transport network maintenance staff);
- OS vendors interested in providing management systems that perform Network Event Correlation, including element management systems (systems engineers and software developers);
- NE suppliers needing to understand how notifications that NEs produce will be used by the TMN (systems engineers and software developers).

Background

TMN emphasizes the use of open, standard interfaces among management systems as well as between management systems and NEs. TMN also provides a logical framework for describing the Management Application Functions (MAFs) associated with the management of the network and the management of services supported by the network. ITU-T Recommendation M.3400, *TMN Management Function*s, is a guide for planning and deploying the management systems necessary for timely and efficient management of networks and services. The Management Function Sets of M.3400 serve as atomic units of network management functionality. The Management Functions in each set cannot easily stand alone. The Management Function Sets can stand alone or they can be packaged together into management systems at the discretion of network/service providers. Each Management Function Set can be considered to be its own Management Application Function (MAF) as defined in ITU-T Recommendation M.3010. This Recommendation treats each Management Function Set as an MAF.

The functional decomposition of operations functions into the MAFs that correspond to the Management Function Sets of ITU-T Recommendation M.3400 helps achieve consistency of functionality among Operations Systems from various suppliers used by different network/service providers managing various network technologies. This functional consistency along with TMN open, standard interfaces (defined in ITU-T Recommendation M.3010) allows interoperability between these Operations Systems and the rapid integration of new emerging technologies into a comprehensive Network Event Correlation capability.

ITU-T Recommendation M.2140

Transport network event correlation

1 Scope

The purpose of this Recommendation is the use of the Telecommunications Management Network (TMN) to correlate notifications of conditions (faults and impairments) from Network Elements and to reveal the conditions themselves and their true effect on the network. This Recommendation is based on the TMN Functional Architecture as specified in ITU-T Recommendation M.3010 [3]. This Recommendation identifies the TMN Management Application Functions (MAFs) corresponding to the Management Function Sets of ITU-T Recommendation M.3400 [4] that are involved in Transport Network Event Correlation. It provides functional descriptions of each MAF, and illustrates how these MAFs interact to achieve Transport Network Event Correlation. This Recommendation also includes timing criteria for event correlation.

Transport Network Event Correlation relies on current configuration information from Configuration Management MAFs. It also is controlled in the specifics of its operation by policies in Business Management Layer MAFs. These interactions are mentioned in appropriate places, but are outside of the scope of this Recommendation.

This Recommendation does not require the exchange of any particular operator's information about its national network.

Algorithms for event correlation are important to the successful implementation of the MAFs described in this Recommendation. Detailed specification of these algorithms is outside of its scope.

2 References

The following ITU-T Recommendations and other references contain provisions which, through reference in this text, constitute provisions of this Recommendation. All Recommendations and other references are subject to revision; all users of this Recommendation are therefore encouraged to investigate the possibility of applying the most recent edition of the Recommendations and other references listed below. A list of the currently valid ITU-T Recommendations is regularly published.

[1] CCITT Recommendation M.20 (1992), *Maintenance philosophy for telecommunication networks.*

[2] ITU-T Recommendation M.60 (1993), *Maintenance terminology and definitions.*

[3] ITU-T Recommendation M.3010 (2000), *Principles for a telecommunications management network.*

[4] ITU-T Recommendation M.3400 (2000), *TMN management functions.*

3 Terms and definitions

This Recommendation defines the following terms:

3.1 **active**: Not cleared.

3.2 **clear**: The end of a fault; the termination of a standing condition.

3.3 **data collection interval**: The period over which performance parameters are accumulated to compute each stored measurement and detect maintenance threshold crossings.

3.4 **degradation**: The presence of anomalies or defects in the absence of a fault.

3.5 **event**: An instantaneous occurrence that changes the global status of an object. This status change may be persistent or temporary, allowing for surveillance, monitoring and performance measurement functionality, etc. Events may or may not generate reports; they may be spontaneous or planned; they may trigger other events or may be triggered by one or more other events. (Recommendation M.60.)

3.6 **event correlation process**: A process that accepts events as input, performs one or more event correlation subprocesses, and reports events as output.

3.7 **event correlation subprocess**: A single step in an event correlation process (see 5.1 for the five types).

3.8 **event report message**: A message sent from one physical system to another that contains information about an event.

3.9 **event set**: The set of all events that are grouped by a selection process for direct comparison or patterning.

3.10 **impairment**: A condition that causes anomalies and defects without causing a failure (degrades the performance of a resource without causing a failure).

3.11 **indication**: An intermediate output of the event correlation process. A notification, indicating a persistent network detected trouble condition. The three types of indication are fault indication, impairment indication and trend indication.

3.12 **independent event**: An event that is not currently superseded by another event.

3.13 **notification**: Information emitted by a managed object relating to an event that has occurred within the managed object (Recommendations M.60, X.710; Information passed from one MAF to another regarding an event).

3.14 **standing condition**: A condition that has duration, beginning with a failure and ending with a clear.

3.15 **threshold crossing alert**: A transient condition declared when a performance monitoring parameter reaches or exceeds a preset threshold.

3.16 **transient condition**: A condition that has no duration, a one time report.

3.17 **trouble**: The perception of a fault or degradation that is judged to require maintenance attention.

4 Abbreviations

This Recommendation uses the following abbreviations:

ISM	In-Service Monitoring
MAF	Management Application Function
MFA	Management Functional Area
NE	Network Element
OS	Operations System
PDH	Plesiochronous Digital Hierarchy
SDH	Synchronous Digital Hierarchy
TCA	Threshold Crossing Alert, also Threshold Report
TMN	Telecommunications Management Network

5 Principles

Telecommunications networks are highly complex systems that must be very reliable. Yet they are made up of millions of parts that can fail and they are subject to a wide variety of unpredictable stresses. The solution that the telecommunications industry has arrived at is to build into the network the ability to detect and report stresses, failures and performance statistics. These real-time reports should allow corrective action to begin quickly and proceed rapidly with the information supplied in the notifications from the network.

Each Network Element (NE) has many ways in which it can fail, or be stressed, and it has the capability to report as many of these conditions as can be foreseen. The only difficulty is that one condition in the network may result in many notifications from multiple NEs for each trouble that needs to be fixed. The number of messages is increased further because the NEs have been designed to send notifications to each other and to both report the condition. This is done to allow limited automatic corrective action and just in case one of the NEs is cut off from the communications with the TMN.

For example, a few troubles in a network can result in hundreds of notifications being generated. Transport Network Event Correlation serves the vital business need to process the hundreds of notifications in real-time and to distinguish each independent trouble. It must also provide the best information from all of the notifications about the nature, location and severity of each trouble.

5.1 Event correlation processes

An event correlation process accepts notifications from Network Elements (or other Operations Systems) and eliminates redundant and symptomatic notifications. It also eliminates notifications of conditions that are not persistent or otherwise do not require maintenance action. It does this by sorting and comparing notifications, consulting databases about the layout and hierarchy of the network, and constantly updating its view of the state of the network. An event correlation process produces a single notification for each distinct condition with the best information from all of the notifications about the nature, location and severity of each condition. This single notification is often referred to as the root cause within the scope of the event correlation process.

An event correlation process can be described as a combination of selecting notifications, comparing or patterning them, evaluating dependencies and checking to see that the condition that triggered the notifications is persistent. These five subprocesses can be linked up in a process flow to create sophisticated event correlation processes. Each of these subprocesses is described below.

5.1.1 Selection of events

The event correlation process starts with a stream of incoming notifications. One of the fundamental event correlation tools is the selection of notifications. Related notifications are selected for direct comparison or patterning. This selection may be done based on any attributes of an event including the type of notification, the source of the notification, or the dependencies between the sources of notifications.

The set of all events that are grouped by a selection process for direct comparison or patterning is called an event set. Event sets must be defined so that the event correlation rules can always produce a single event to characterize each underlying trouble condition, through filtering and combination of any possible members of the event set, that is meaningful and without losing valuable surveillance information.

5.1.2 Comparison of events

An event comparison subprocess ensures that there is always at most one active, independent event for each underlying condition within the domain of the MAF. It does this by using event correlation rules that are specific to each stage of event correlation and often specific to the types of resources being monitored.

An event is active if it has not been cleared and has not expired by falling outside of a correlation window. An event is independent if it has not been superseded by another event.

The comparison of events occurs when a new event notification is received and selected into the same event set as another active, independent event. The attributes of the new event are compared with the attributes of the independent events already in the event set. Correlation rules are used to determine whether the new event supersedes or is superseded by the independent events already in the event set. If there is an independent event, related to the same underlying condition, that cannot supersede or be superseded by the new event, it will be combined with the new event. This combination is done to preserve relevant information from both events into a single synthesized event that supersedes both to become a new independent event within the event set.

In some cases, events can be ordered by severity or by significance. This is often because the condition reported by one notification implies the condition reported by another. For example, a critical alarm of a given type and source supersedes a major alarm of the same type and source. An environmental alarm indicating high humidity is superseded by one indicating a flood. When one event is superseded by another it is filtered out when the persistence time is reached. If no persistence analysis is performed within the event correlation process, then superseded events are filtered out immediately.

In other situations notifications are of equal severity and significance, but are clearly related. Examples of this are, two equipment alarms from the same equipment component or two environmental alarms of the same severity from equipment at the same location. In these situations two events must be combined into a single synthesized event with multiple cause codes.

If there are unrelated notifications then they may be the result of independent underlying conditions in the network and so they are not correlated. All independent notifications are sent forward to the next subprocess of the event correlation process.

The rules for comparing and combining notifications will differ slightly from one technology to another, but should be stable for each technology (set of managed resources), and should not require much administration.

5.1.3 Persistence analysis

Persistence analysis looks at one or more types of notification, which can be an output of another event correlation subprocess, and correlates those types of notification over time.

The purpose of persistence analysis is to distinguish between persistent conditions, which require further analysis and attention, and non-persistent conditions that should be ignored and filtered out. Non-persistent conditions are infrequent and of short duration. When a condition is not persistent there is no reason to believe that a subsequent condition will occur or that any intervention will improve the expected performance of the affected network component in the future. Non-persistent conditions usually have external causes such as human activity or adverse weather conditions. Persistent conditions are long duration or intermittent but frequent occurrences of the same condition. When a persistent condition is observed there is reason to believe that the condition will continue or recur, that some network component is at fault, and that the trouble can be found and corrected.

There are two fundamentally different types of conditions for the purpose of persistence analysis. A standing condition, such as a fault or state value exists for the entire period from when it is raised (reported by an alarm or state change notification) to when it is cleared (reported by an alarm clear or state change notification). All types of indications defined in 5.1.3 are standing conditions. A transient condition, such as a threshold crossing [reported by a Threshold Crossing Alert (TCA)], carries information that applies only to a particular moment in time and does not have an associated clear event.

For standing conditions, persistence analysis consists of measuring the length of time that passes after the event (a condition is raised or cleared or a state value changes) and after enough time has passed, declaring the event persistent. When the event is declared persistent, a notification is sent forward. If the event is reversed before enough time has passed, then it is filtered out and no notification is sent forward.

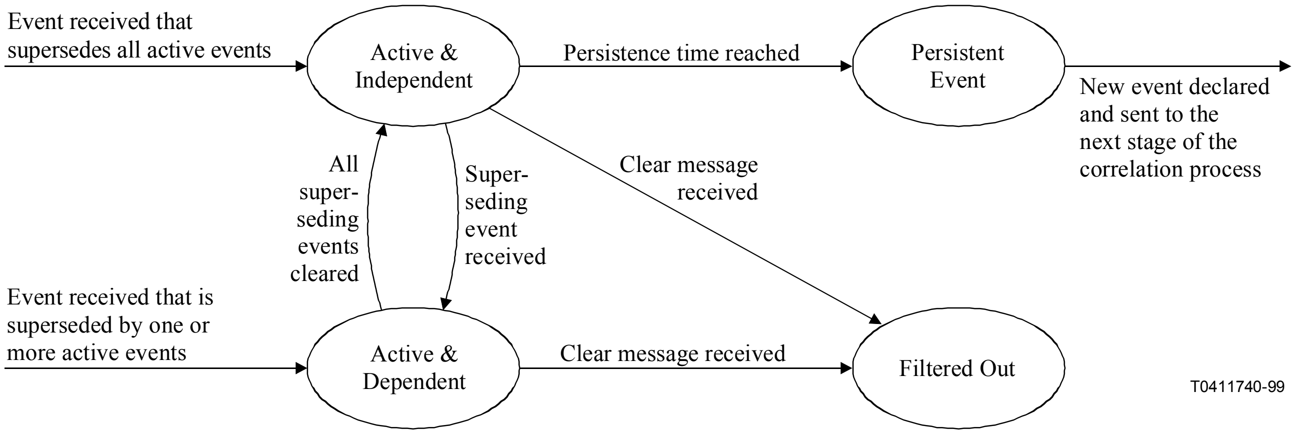

Figure 1/M.2140 – Persistence analysis for standing conditions

In the case of an intermittent phenomenon, a standing condition may be raised and cleared repeatedly at the earliest stage of correlation. However, the generic event correlation model in 5.2 defines four stages of event correlation with each successive stage operating over longer time periods. This staging of the correlation process allows intermittent standing conditions can be declared persistent.

For transient conditions, persistence analysis consists of measuring the frequency of events and when the frequency is high enough, declaring the condition persistent and sending an event notification forward.

Persistence analysis is not appropriate for all types of conditions. If a certain condition type or state value is unlikely to be cleared or changed soon after it is reported, then persistence analysis can be bypassed, by setting the required persistence time for that condition type to zero.

5.1.4 Dependencies

Dependencies are relationships between equipment components and/or transmission media determined by the topology of the network and the way that trails are provisioned. These relationships can be used to determine where the failure or degradation of one resource will result in the failure or degradation of other dependent entities. Dependency relationships can also be used to determine where events detected at one resource indicate a condition present at another resource. For example a trail is dependent on the equipment that terminates the trail at either end and on any server trails that carry it. It is these dependency relationships that allow events detected at one resource to be correlated with events detected at dependent resources.

Dependency information is created during the engineering process and the provisioning process by Configuration Management. It must be available to the appropriate Performance Management and Fault Management MAFs to enable Network Event Correlation. Interfaces between Configuration Management MAFs and Performance and Fault Management MAFs provide this dependency information. Fault Management MAFs need dependency information in real-time. Performance Management MAFs need dependency information in near real-time. Both require near real-time updates, when dependency information changes. The interfaces that provide dependency information must be automated.

Dependencies are evaluated by following these steps:

1) adding dependency relationship information to event information received in notifications;
2) then selecting all notifications for entities that support the same dependent resource;
3) performing an event comparison to filter out redundant notifications;
4) and combining related notifications of the same importance in light of the relationship information that was added.

5.1.5 Patterning

Patterning is the correlation of events across the many supported resources of a supporting resource to infer the occurrence of unreported events on the supporting resource. Patterning is done to compensate for potential loss of communication with NEs or other OSs in the TMN and the resulting loss of notifications for the supporting resource. This method can also be used to infer events affecting unmonitored, common resources in the network, such as cables, conduits, and power sources.

When a pre-set fraction of all provisioned supported resources have indications of a given type logged against them, then an indication is generated on the supporting resource which may supersede the individual indications on the supported resources.

Patterning requires a database of the dependency relationships between managed resources including unmonitored resources. Patterning may be done recursively to infer trouble with a resource such as a conduit or power supply from indications in many indirectly supported resources (as depicted in Figure 2).

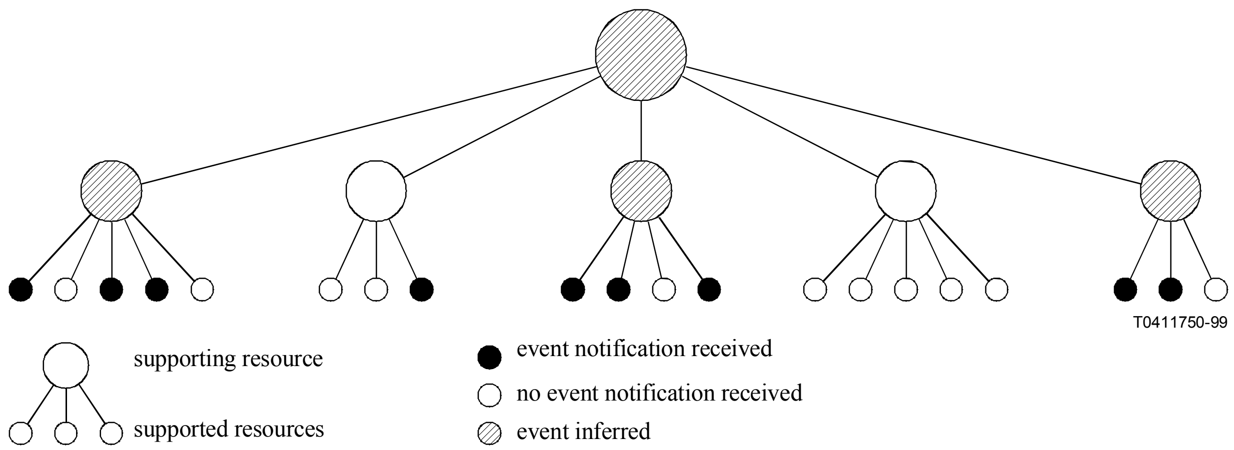

Figure 2/M.2140 – Patterning of events

5.2 Generic event correlation model

To correlate events of diverse types from multiple technologies across a large managed network the event correlation processes described in 5.1 must be performed many times with the results of one process becoming the input to the next process. As correlation proceeds, dependency information is added to events and some event detail is removed resulting in increasing levels of abstraction.

The following generic event correlation model defines four stages of event correlation processing. Figure 3 shows the correlation processes present at each stage. It also shows the data stored at each stage and the data passed upward from one stage to the next.

5.2.1 Stage 0

At the lowest level, anomalies, defects, and State Changes are detected in monitored resources on the order of milliseconds and usually reported to the appropriate Stage 1 correlation process within one second or sooner.

5.2.2 Stage 1

At the first stage, anomaly, defect, and state change information is received. Anomalies are accumulated. Defect detections and terminations may be subjected to a brief persistence analysis. Failures are declared and faults are cleared based on the persistence or persistent absence of defects. State Changes may also be subjected to a brief persistence analysis. Performance parameters are calculated based on the occurrence of anomalies and defects. Intermediate values of performance parameters, during the data collection interval, are compared to thresholds periodically as they are accumulated and TCAs are generated if the threshold is reached or exceeded. The Stage 1 correlation process generates the following sorts of events:

- Failure declarations.
- Fault clears.
- State Changes.
- Performance parameter values.
- TCAs on performance parameters.

5.2.3 Stage 2

At the second stage of the correlation process, events in the same category within a local domain are received and dependency information is added to them based on information stored in a database containing the dependency relationships between equipment components and in some cases local network connectivity. These events are then selected for comparison based on common dependencies (such as receiving power from the same immediate source or a payload signal from the same source) and correlated. A local persistence analysis is performed and transient or redundant events are filtered out. The result of the local correlation process is called an indication and should be stored, available to other processes, and reported automatically to the third stage of the correlation process. These indications provide a reliable abstraction of the nature of a persistent fault or impairment (degradation) within the local domain. Indications may also be exchanged between network operators to facilitate fault localization.

5.2.4 Stage 3

At the third stage of the correlation process, indications in the same category are received from each local domain and dependency information is added to them based on information stored in a database containing the network-wide dependency relationships between NEs or subnetwork connections and end-to-end facilities or circuits. These indications are then selected for comparison based on common dependencies (such as receiving a payload signal from the same source) and correlated. Redundant indications are filtered out. No persistence analysis is required at Stage 3. Instead updated indications are sent when new information changes the root cause information for a fault or impairment. Indications that are not filtered out should be stored, available to other processes, and reported automatically to the fourth stage of the correlation process. These indications provide a reliable abstraction of the nature of a persistent fault or impairment (degradation) based on current information from a broad category of events.

5.2.5 Stage 4

At the fourth and final stage of the correlation process, independent, non-redundant indications from all categories are received from the network-wide Stage 3 correlation processes. Dependency information is added to them based on information stored in a database containing the network-wide dependency relationships between end-to-end facilities or circuits and the physical resources, such as cables, conduits and rights-of-way. These indications are then selected for comparison based on common dependencies (such as common routing) and correlated. Trouble patterning is done to infer troubles on common resources that are not directly monitored, or on resources that may be unmonitored due to failures in the TMN. Real-time information input from external trouble reports and other relevant sources may be included in this final stage of event correlation. A final persistence analysis is performed and transient or redundant events are filtered out. Indications that are not filtered out should be stored, available to other processes, and reported automatically to the trouble ticket administration process. Additionally, if the trouble is localized to an external network operator then it should be reported automatically to the trouble report handling process of the external network operator. The indications that survive the Stage 4 correlation process provide a reliable and definitive statement of the nature of a persistent fault or impairment (degradation) based on current information from all sources.

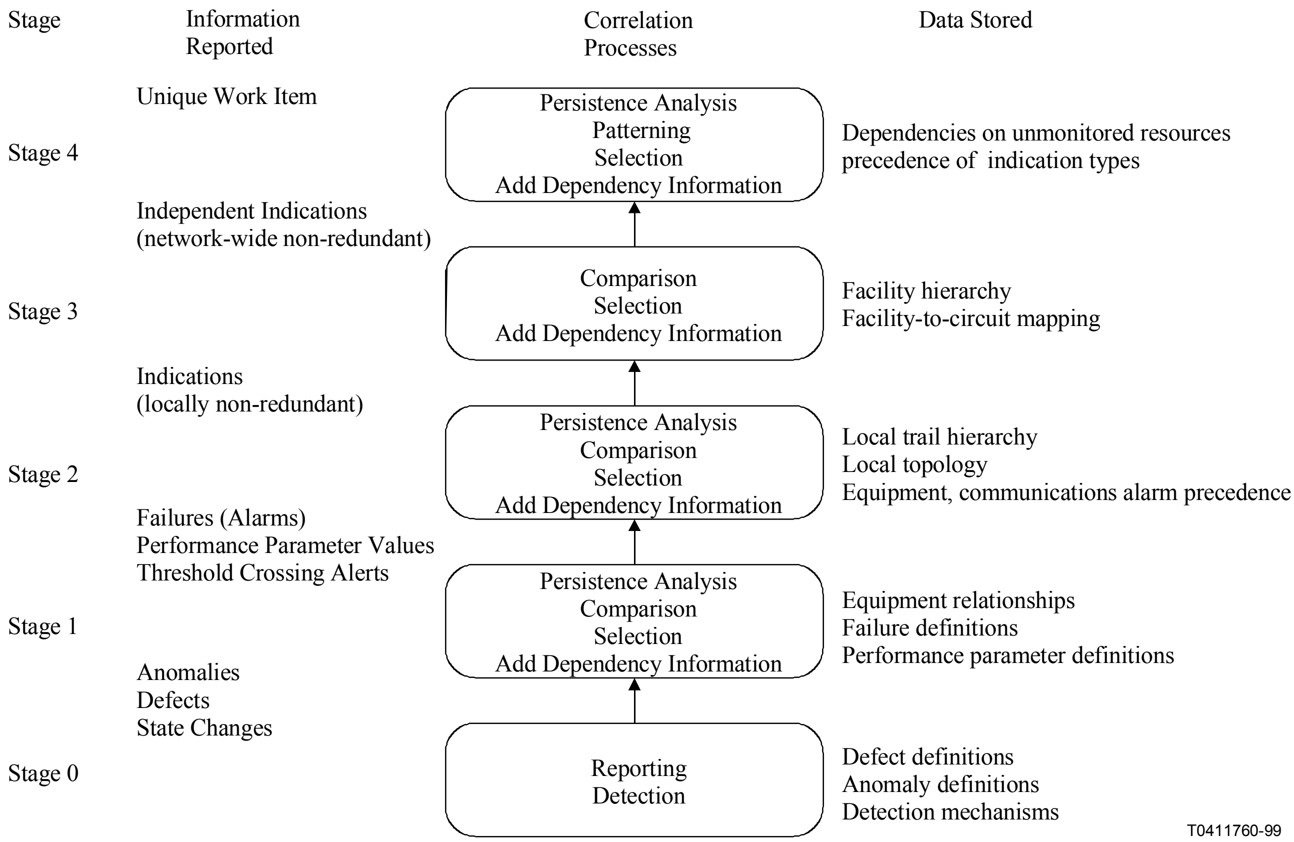

Figure 3/M.2140 – Generic event correlation model

5.3 Surveillance operations flows

Information from NEs of interest to Transport Network Event Correlation can be classified into three categories.

1) Failure declarations, fault clears and State Changes.

2) Threshold Crossing Alerts (TCAs).

3) Transport performance parameter values.

These three categories of information are the source material for three surveillance operations flows. At Stage 2 each of these categories of data are received by different MAFs and correlated separately within each local domain. A local domain is the span of control of each Management System that performs that MAF. For each of the three flows, the Stage 2 MAFs correlate the notifications from Stage 1 in the local domain to produce an intermediate output, called an indication, for further correlation at Stage 3.

The concept of an indication has been used for each Interaction between Stage 2 and Stage 3 MAFs and is consistent across the three surveillance operations flows. Table 1 introduces the indication terminology.

Table 1/M.2140 – Indication types

Surveillance Flow	Indication Type
Alarm Surveillance	Fault Indication
TCA Surveillance	Impairment Indication
Trend Analysis	Trend Indication

In each flow, the Stage 2 MAF will produce one indication for each underlying condition in the network. The indication will contain the best information about the nature, location and severity of the condition available from the notifications within the local domain.

The definition of indications is crucial for the interoperation of local and network-wide systems. The indication is not just an NE notification that has made it through a filter at Stage 2. It is a statement by the applicable Stage 2 MAF that it has observed a significant condition in its domain. The information in the indication may be drawn directly from one or more NE notifications or it may be very indirectly related to the individual NE notifications. In either case, the Stage 3 recipient of the indication (whether a person or a machine) must be prepared to receive the indication and act on it appropriately. Without a common definition of indications, the job of Network Event Correlation cannot be divided between Stage 2 and Stage 3 and an efficient exchange of reliable information between network operators will not be possible. This Recommendation proposes a common definition for each of these indications, allowing a partitioning of Network Event Correlation MAFs into stages.

At Stage 3 Indications of each category are received from multiple local domains. The indications are correlated based on all of the information available from all local domains.

The three surveillance operations flows are brought together at Stage 4 that culminates in a single work item from Network Fault Localization. Table 2 presents suggested timing criteria for each of the three surveillance operations flows. These criteria specify bounds on the persistence times to be used at each stage of event correlation. The essential features of the suggested criteria are that each of the three flows occurs at different time-scales and that later stages have longer persistence times than earlier stages within each flow.

Table 2/M.2140 – Suggested event timing criteria

Flow	Alarm Surveillance	TCA Surveillance	Trend Analysis
Stage 0	anomaly counts and defect detection < 1 second	anomaly counts and defect detection < 1 second	anomaly counts and defect detection < 1 second
Stage 1	equipment < 3 seconds communications 2-3 seconds environmental < 3 seconds fault clears < 20 seconds	15-minute TCA < 1 minute	15-min value 15 minutes 24-hour value 24 hours
Stage 2	fault indication 10-30 seconds	impairment indication 1-3 hours	trend indication 2 hours – 30 days
Stage 3	no persistence analysis	no persistence analysis	no persistence analysis
Stage 4	work item 30-150 seconds	work item 1-5 hours	work item 2 hours – 30 days

5.4 Implementation considerations

For large public networks, any implementation of event correlation must have the following characteristics:

Scalable: Some network troubles can cause the generation of thousands of event reports and these event reports can come from multiple Network Elements. This need implies also that the event correlation processes need to be both fast and use computing resources efficiently.

Robust: Some event notifications that should be generated by a given trouble may not be generated or may not be received by a management system. False notifications may also be generated under some circumstances. The event correlation implementation must be robust enough to deal with missing or false event reports.

Layered: TMN-compliant management system architectures depend on multiple systems such as element managers, domain managers and inter-domain managers. To reduce the volume of event notifications that must be processed by higher level and multi-domain management systems, domain-specific event correlation should take place at each stage of event correlation and only summarized results should be sent to higher level management systems.

Flexible: Large public networks require constant reconfiguration. This activity may be caused by customer orders, network growth or technology changes. The event correlation process, which must have knowledge of the network configuration, must be able to accommodate these changes and keep up-to-date with the current network configuration. If maintenance activity, such as removal of transport links from service, would affect event correlation, then this information must be available to the event correlation process and must be utilized by that process.

6 The transport network event correlation MAF interaction diagram

The following diagram (Figure 4) is presented in the style of Appendix I/M.3400 [4]. Each object represents a single MAF. The first line of text in each object is the function group to which the MAF belongs. The number in each object is the number of the subclause pertaining to the MAF, the text following the number is the name of the MAF. Each arrow corresponds to the flow of information between MAFs. The numbers on the arrows correspond to numbered descriptions of each interaction that follow the diagram. The stages of event correlation are represented by the vertical alignment of the MAF objects.

Following each interaction description is a list of Associated Data, which contains a list of the most important data items that are conveyed in each interaction.

The Transport Network Event Correlation MAF Interaction Diagram shows the relationships among the MAFs for the purpose of event correlation. It includes all MAFs involved in event correlation for network surveillance and maintenance. It includes MAF from Performance Management, Fault Management, and Configuration Management to create a complete picture of Network Event Correlation.

At Stage 0 there are three sources of events for correlation. Three distinct operations flows can be seen:

- the Alarm Surveillance Flow: interactions 1, 2, 3, 4, 5, 6;
- the TCA Surveillance Flow: interactions 7, 8, 9, 14;
- the Trend Analysis Flow: interactions 10, 11, 12, 13, 14.

These three flows come together in Stages 3 and 4. The Trend Analysis Flow is filtered by Data Aggregation and Trending and passed on to Network Performance Monitoring Event Correlation and Filtering where it is correlated with the TCA Surveillance Flow. All network degradation and failure information is correlated by Network Fault Localization.

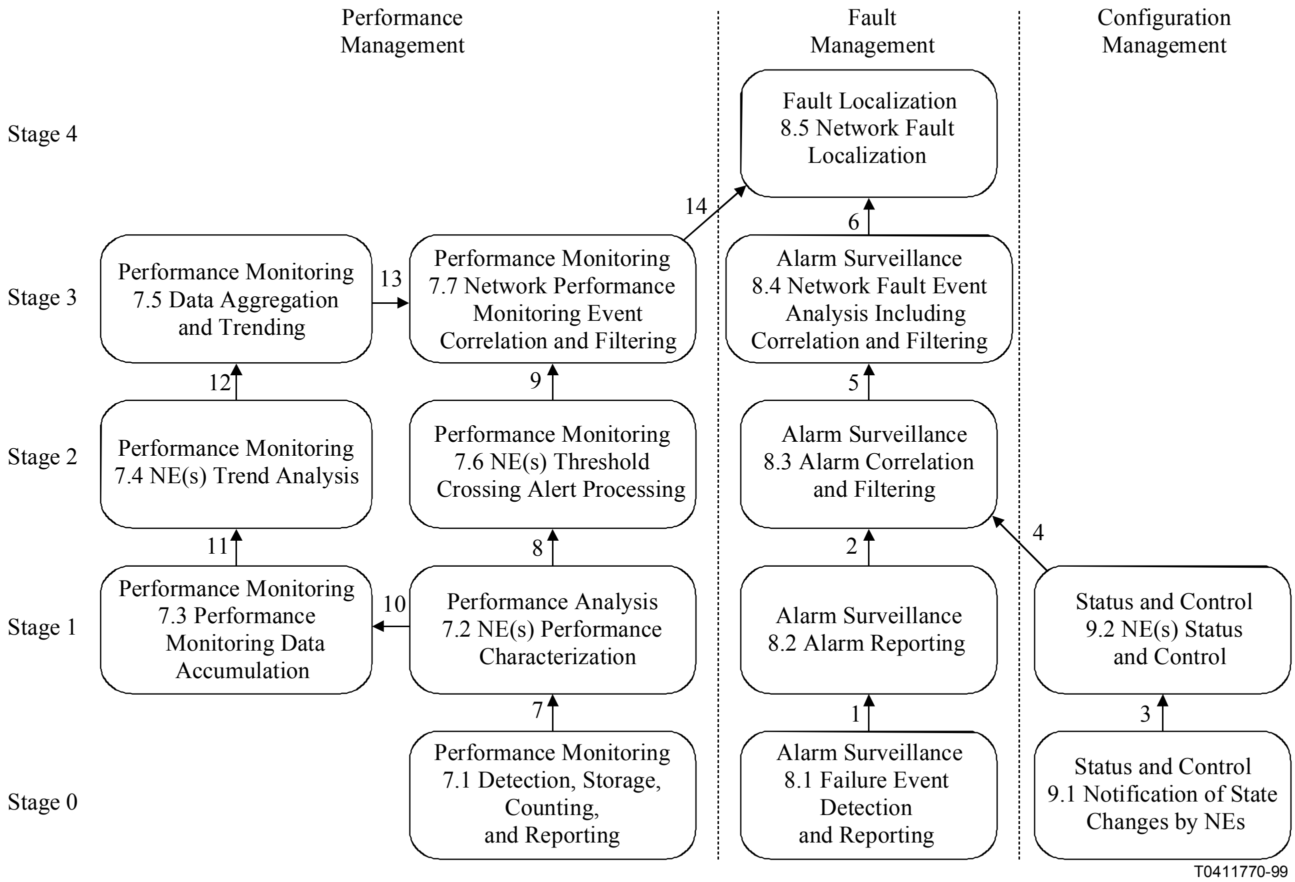

Figure 4/M.2140 – Transport network event correlation interactions

1) Defect detections and terminations are sent by Failure Event Detection and Reporting to Alarm Reporting.

Associated Data:
Managed Resource
Defect Type
Detection/Termination Indicator

2) Alarms are sent by instances Alarm Reporting in one or multiple NEs to Alarm Correlation and Filtering.

Associated Data:
NE ID
Managed Resource
Event Type
Probable Cause
Time & Date Stamp

3) State Change Notifications are sent to NE(s) Status and Control.

Associated Data:
NE ID
Managed Resource
Administrative State
Operational State

4) State Changes of interest to Network Event Correlation are forwarded to Alarm Correlation and Filtering in Fault Management.

Associated Data:
NE ID
Managed Resource
Administrative State
Operational State

5) Fault indications are sent from multiple instances of Alarm Correlation and Filtering to Network Fault Event Analysis Including Correlation and Filtering.

Associated Data:
NE ID
Managed Resource
Event Type
Service Affect
Probable Cause
Severity
Fault Location
Time & Date Stamp

6) Independent, non-redundant fault indications are sent from Network Fault Event Analysis Including Correlation and Filtering to Network Fault Localization.

Associated Data:
NE ID
Managed Resource
Event Type
Service Affect
Probable Cause
Severity
Fault Location
Time & Date Stamp

7) Performance anomaly and defect information is provided to NE(s) Performance Characterization within each NE.

Associated Data:
NE ID
Termination Point ID
Performance Anomaly
Performance Defect
Direction Of Transmission

8) Threshold Crossing Alerts are sent from multiple NEs to NE(s) Threshold Crossing Alert Processing.

Associated Data:
NE ID
Termination Point ID
PM Parameter Name
PM Parameter Value
Direction Of Transmission

9) Impairment indications are sent from multiple instances of NE(s) Threshold Crossing Alert Processing to Network Performance Monitoring Event Correlation and Filtering.

Associated Data:
NE ID
Termination Point ID
PM Parameter Name
Direction Of Transmission
Observed TCA Frequency

10) Performance parameter values are sent from multiple instances of NE(s) Performance Characterization in various NEs to Performance Monitoring Data Accumulation.

Associated Data:
NE ID
Termination Point ID
PM Parameter Name
Direction Of Transmission
Data Collection Interval End Time
PM Parameter Value

11) Performance Monitoring Data Accumulation sorts and organizes the performance parameter values and makes them available to NE(s) Trend Analysis.

Associated Data:
NE ID
Termination Point ID
PM Parameter Name
Direction Of Transmission
Data Collection Interval End Time
PM Parameter Value

12) Trend indications are sent from multiple instances of NE(s) Trend Analysis to Data Aggregation and Trending.

Associated Data:
NE ID
Termination Point ID
Pattern Code
Pattern Descriptors

13) Independent, non-redundant trend indications are sent on to Network Performance Monitoring Event Correlation and Filtering

Associated Data:
NE ID
Termination Point ID
Pattern Code
Pattern Descriptors

14) Independent, non-redundant impairment indications and trend indications that are not tied to redundant impairment indications are forwarded to Network Fault Localization.

Associated Data:

NE ID	or	NE ID
Termination Point ID		Termination Point ID
PM Parameter Name		Pattern Code
Direction Of Transmission		Pattern Descriptors
Observed TCA Frequency		

7 Performance management

Transport Performance Management uses transport performance monitoring data to detect degraded performance of transport facility resources that is either service affecting already, or that is potentially service affecting in the near future. When degraded performance is detected and persists, proactive maintenance is begun to fix the facility component that is causing the degraded performance.

Proactive maintenance is triggered by one of the following occurrences in the performance monitoring measurements:

- a level of degradation that directly impacts service performance;
- measurements that exceed performance quality guarantees or objectives;
- a change for the worse observed;
- a pattern indicates impending onset of more severe degradation or complete failure.

Transport Performance Management deals with degradation that has not yet interrupted service completely and measurements may vary from one moment to the next. Frequently, when degraded performance is detected it quickly fades, caused by some temporary phenomena external to the network. Because of these characteristics of transport performance data, events are examined over a longer time than they are when service is interrupted. Evidence of degradation is withheld at each stage of event correlation until there is confidence that the true severity of the problem is known, that the degradation is persistent, or that the source of the problem is clear, before the next stage is notified. The philosophy is that delayed reliable information is preferable to immediate unreliable information. This is in contrast to the Alarm Surveillance philosophy stated in clause 8.

There are two approaches to Transport Performance Management. The Trend Analysis approach looks at all of the performance parameter values. The TCA Surveillance approach only considers periods of exceptionally poor performance. The Trend Analysis approach is more complete, but is more costly. The TCA Surveillance approach is less costly, but less flexible, less sensitive and less sophisticated. Both have their place in different network contexts. Often times, the Trend Analysis approach is applied only for high signal rate facilities or premium services. In other cases, Trend Analysis may be triggered by the detection of exceptionally poor performance by the TCA Surveillance approach. Both approaches are covered in this clause. The TCA Surveillance Flow and the Trend Analysis Flow and their use of the MAFs defined in this clause are presented in this clause.

7.1 Detection, counting, storage and reporting

Detection, Counting, Storage, and Reporting is a Stage 0 function that monitors the status and behaviour of resources within an NE for conditions that impact performance. For transport networks, this MAF monitors the incoming (and in some cases outgoing) signals at each port in a Network Element to detect and track anomalies and defects in the signal. Not all monitored trails must be terminated at the NE.

An anomaly is a discrepancy between the actual and desired characteristics of an item. For this MAF the items are monitored signals. An anomaly may or may not affect the ability of an item to perform a required function.

Examples of anomalies are:

- bit error detection anomalies: line coding errors, parity or other error detection code errors;
- synchronization related anomalies: controlled slips for PDH paths and pointer justifications for SDH paths.

A defect is a limited interruption in the ability of an item to perform a required function. It may or may not lead to maintenance action depending on the results of additional analysis. Successive anomalies causing a decrease in the ability of an item to perform a required function may result in the detection of a defect. Examples of defects are:

- loss of information defects: Loss Of Signal (LOS), Loss Of Pointer (LOP), Alarm Indication Signal (AIS);
- framing defects: Out Of Frame (OOF), Severely Errored Frame (SEF) and Loss Of Frame (LOF);
- incorrect information defects: incorrect or invalid variable values in frame or cell overhead in successive frames or cells.

7.2 NE(s) performance characterization

NE(s) Performance Characterization is a Stage 1 function that is usually performed by the NE. This MAF uses the anomalies and defects from Detection, Counting, Storage, and Reporting to compute various Performance Monitoring parameters such as Errored Seconds, Severely Errored Seconds and Unavailable Seconds. This is done for each monitored trail. NE(s) Performance Characterization accumulates each of the parameters over 15-minute intervals and separately over 24-hour intervals. These counts are updated at the end of each second based on the anomaly and defect information from Detection, Counting, Storage and Reporting. Once an interval ends, the total count of each parameter is stored temporarily as historical Performance Monitoring data. The NE can be instructed to report current or historical data on demand or on a scheduled basis. These performance parameter values are used by NE(s) Trend Analysis.

For each parameter, thresholds are set. There is one for 15-minute intervals and one for 24-hour intervals. If a count reaches or exceeds its threshold, then a Threshold Crossing Alert is generated and sent to the NE(s) TCA Processing MAF.

7.3 Performance monitoring data accumulation

Performance Monitoring Data Accumulation is a Stage 1 function that receives performance parameter values for each data collection interval, sorts them by monitoring point and stores it for use by other MAFs. This data is then used by NE(s) Trend Analysis as well a other MAFs (not included in Network Event Correlation) involved in characterizing long-term network performance and producing quality of service reports.

7.4 NE(s) trend analysis

NE(s) Trend Analysis is a Stage 2 function that analyses the raw performance monitoring data gathered from NEs to detect trends, or more generally, patterns of performance monitoring data that indicate the presence or impending onset of a service affecting impairment that requires maintenance action to correct. For most applications, 15-minute data is used to allow rapid detection of trends.

NE(s) Trend Analysis concentrates on uncovering patterns in the performance parameter values that indicate the cause of the impairment in more detail than NE(s) TCA Processing can provide. For example, a certain pattern may indicate that a specific component of a repeater is going bad. It also concentrates on cases when the data values are not off-normal enough to trigger TCAs, but the pattern indicates that the problem will worsen.

NE(s) Trend Analysis examines the performance parameter values from NEs with no reference to the particular trail that the monitoring point supports. The only information that is needed is the type of trail that the monitoring point supports. The data is analysed by one or more pattern recognition algorithms to search for matches to rules, profiles, or other criteria that indicate the presence or impending onset of a condition that requires maintenance action. When a match is made, a trend indication is issued.

The trend indication will include the pattern code for the trend or pattern detected, parameters (not PM parameters) characterizing the pattern detected and measures of the goodness of fit to the pattern detected. The trend indication will apply only to a single monitoring point. No correlation across monitoring points is included in this MAF.

7.5 Data aggregation and trending

Data Aggregation and Trending is a Stage 3 correlation function that receives trend indications from multiple instances of NE(s) Trend Analysis in different local domains. Trending is done at Stage 2 by NE(s) Trend Analysis on data from individual monitoring points to find patterns in raw performance monitoring data that may have unique significance for a specific type of equipment or transmission technology. When an impairment occurs, performance monitoring data will be collected at multiple points along the impaired trail. One or more instances of NE(s) Trend Analysis will try to find a pattern in the data at each of these monitoring points.

Data Aggregation and Trending determines what trail is affected by each trend indication that it receives. If it receives more than one for a given trail, it also performs correlation along a trail to arrive at the most descriptive and diagnostic trend information for each impairment on a trail. These trend indications are sent to Network Performance Monitoring Event Correlation and Filtering.

No correlation among different trails is included in this MAF.

7.6 NE(s) TCA processing

NE(s) TCA Processing is a Stage 2 function that receives Threshold Crossing Alerts (TCAs), based on 15-minute interval performance monitoring data, from NE Performance Characterization at Stage 1. It correlates the TCAs across multiple monitoring points on a trail and over time. It filters out redundant or symptomatic TCAs and sectionalizes impairments. For each impairment, it forwards a single reliable impairment indication to Network Performance Monitoring Event Correlation and Filtering at Stage 3 for correlation on a network-wide basis.

NE(s) TCA Processing correlates TCAs generated by NEs within a single period to filter out redundant or symptomatic TCAs related to a given trail. It also correlates TCAs from low order (lower rate) trails to the high order (higher rate) supporting trails within its local domain. It will suppress TCAs from a low order trail when a TCA from a higher order trail indicates an impairment that would cause the low order trail TCA as a symptom. For each 15-minute interval, NE(s) TCA Processing will allow only one root cause TCA to pass through its filter for each impairment.

The root cause TCA for a given impairment should be the one that best indicates the cause, location, and severity of the impairment. Fortunately, for the vast majority of impairments, one TCA, out of the many that may be received, is clearly the best on all three criteria.

TCAs apply to one of the two directions of a bidirectional service. The two directions can be impaired independently and NE(s) TCA Processing must distinguish the affected direction of transmission. The different types of TCA can be ordered by severity. Finally, TCAs can be generated at multiple points along a trail. For this reason, NE(s) TCA Processing must record where TCAs come from and construct a notion of upstream and downstream along a trail in one or the other direction of transmission.

7.6.1 TCA correlation rules

The following rules specify how TCAs are correlated within each 15-minute interval.

1) Any TCA on a trail shall suppress downstream near-end TCAs and far-end TCAs relating to the same direction of transmission on lower order trails carried by that trail.
2) Among TCAs for the same trail, TCAs of greater severity shall suppress TCAs of lower severity relating to the same direction of transmission.
3) Among TCAs of the same type (parameter name) for the same trail, near-end TCAs shall suppress downstream near-end TCAs and all far-end TCAs relating to the same direction of transmission.
4) Among TCAs of the same type (parameter name) for the same trail and in the absence of near-end TCAs, the first received far-end TCA shall suppress subsequent far-end TCAs relating to the same direction of transmission.

Far-end data for a given trail is read from the overhead, not directly observed. As a result, far-end data values for the same direction of transmission will be the same at the trail termination point and at all intermediate performance monitoring points along the trail. All far-end TCAs are equally valid but only one should be allowed to become a root cause TCA, and then only if there are no near-end TCAs for that direction of transmission on the trail and no TCAs on supporting higher order trails.

7.6.2 Persistence analysis

The purpose of persistence analysis is to distinguish between persistent impairments, which should result in maintenance action, and non-persistent impairments, which should be ignored.

Non-persistent impairments are infrequent and of short duration (e.g. affecting only one or two 15-minute intervals). When an impairment is not persistent there is no reason to believe that a subsequent impairment will occur or that any repair action will improve the expected quality of the transmission system in the future. Non-persistent impairments are usually caused by external events such as human activity or adverse weather conditions.

Persistent impairments are long duration or intermittent but frequent occurrences of the same impairment. When a persistent impairment is observed there is reason to believe that the impairment will continue or recur and that some component of the transmission system is at fault.

NE(s) TCA Processing checks to see that an impairment is persistent by using a sliding window algorithm. It works by recording the presence of an impairment as represented by the root cause TCA in successive intervals and requiring the impairment to be present in three or more intervals within a 4 to 12 interval window before declaring an impairment indication. The window length and number of impaired intervals required can be set based on performance monitoring policy, but they should not vary from trail to trail. If greater or lesser sensitivity to the degree of impairment is desired, the threshold settings on the individual PM parameters in NE(s) Performance Characterization should be modified. The impairment indication is cleared again when the number of impaired intervals in the window falls below the number required.

This persistence analysis turns many root cause TCAs into a single, reliable standing condition. The impairment indications and associated clears are sent to Network Performance Monitoring Event Correlation and Filtering.

7.7 Network performance monitoring event correlation and filtering

Network Performance Monitoring Event Correlation and Filtering is a Stage 3 function that has two event inputs. It receives impairment indications from multiple instances of NE(s) TCA Processing. It receives trend indications from Data Aggregation and Trending. Network Performance Monitoring Event Correlation and Filtering has a larger span of control than NE(s) TCA Processing does. For this reason, it performs some similar correlation functions.

This MAF does three types of event correlation. The first is correlation of impairment indications from different local domains relating to the same trail. The second is correlation of impairment indications from one trail to indications on another trail that carries it. The third is correlation of any trend indication that exists to the impairment indication for the trail that is determined to be the root cause of an impairment.

The output of this MAF is root cause indications (either impairment indications or trend indications) that contain the best information about the severity, location and cause of degradation on a trail. These root cause indications are sent to Network Fault Event Analysis, Including Correlation and Filtering.

7.7.1 Correlation of impairment indications along a trail

This is similar to the correlation of individual TCAs from different monitoring points (within a local domain) on the same trail that takes place in NE(s) TCA Processing. The difference is that Network Performance Monitoring Event Correlation and Filtering correlates impairment indications (which are a product of persistence analysis of many TCAs) and it correlates over a larger domain. The process, however, is the same.

The two directions of transmission for a trail are treated separately. Impairment indications that are of a lower or equal severity than another, declared upstream, are superseded. If there are multiple impairment indications based on far-end data, then all but the first received are superseded.

7.7.2 Correlation of impairment indications among related trails

Correlation of impairment indications among trails that have a carrying/carried relationship also has an analogous function in NE(s) TCA Processing. Just as in NE(s) TCA Processing, correlation along a trail and among related trails can be done simultaneously.

Impairment indications on a lower order trail are superseded, if there is an impairment indication on a higher order trail that carries the lower order trail.

The procedure for one impairment indication superseding another depends on the order in which the impairment indications are received and processed. If the superseding impairment indication is processed earlier, the subsequent superseded impairment indication will be filtered out and no message will be sent to Network Fault Event Analysis, Including Correlation and Filtering for the superseded impairment indication. If the superseding impairment indication is processed later, an updated impairment indication is set to Network Fault Event Analysis, Including Correlation and Filtering reflecting the information in the superseding impairment indication.

7.7.3 Correlation of trend indications with impairment indications

In order to understand the recommended approach to the correlation of trend indications to impairment indications, it is important to consider that for a given impairment, it is possible to receive a trend indication by itself, an impairment indication by itself, or both. If both are triggered, then they may be triggered at widely different times. For this reason, no persistence analysis is recommended to hold on to one indication while waiting for another. Instead, the two kinds of indication will be reported separately. This may seem contrary to the spirit of correlation, but delaying either category of performance indication long enough to correlate them effectively would withhold the information too long to support proactive maintenance. Furthermore, trend indication information and impairment indication information are so different in structure and semantics that combination of the two would be an artificial packaging exercise.

Even without persistence analysis, redundant trend indications can be filtered out by correlating them with superseded impairment indications. Network Performance Monitoring Event Correlation and Filtering gets trend indications from Data Aggregation and Trending. A trend indication is said to correspond to an impairment indication if it is from the same direction of transmission of the same trail as the impairment indication. There are three cases to consider when correlating trend indications with impairment indications.

1) A trend indication arrives before any corresponding impairment indication. In this case the trend indication is forwarded to Network Fault Event Analysis, Including Correlation and Filtering.
2) A trend indication arrives and corresponds to a root cause impairment indication. In this case the trend indication is forwarded to Network Fault Event Analysis, Including Correlation and Filtering.
3) A trend indication arrives and corresponds only to one or more superseded impairment indications. In this case the trend indication is superseded.

8 Fault management

Fault Management deals with failures that have interrupted service completely or adverse conditions that are immediate threats to service. These fault events often require prompt maintenance action. The Alarm Surveillance Flow is designed to operate quickly to provide near-real time information for the maintenance process. Persistence analysis, required to filter out redundant or non-persistent events, is kept to a minimum (in the order of one minute). It is likely that the alarms that contain the best root cause information and are most significant will be received first. The Alarm Surveillance Flow takes advantage of this. However, there are cases when one stage of event correlation will provide incomplete information to the next stage and then update that information when more complete information is available. The philosophy is that, when service continuity is lost or under threat, prompt incomplete information is preferable to delayed complete information. This in contrast to the Transport Performance Management philosophy stated in clause 7.

8.1 Failure event detection and reporting

Failure Event Detection and Reporting is a Stage 0 function that monitors resources within an NE for defects. Monitored resources include equipment components, software, incoming signals, power and timing signals. Defects are reported as soon as possible after their detection, often within milliseconds and always within a second. The specific defects and their definitions are dependent on the technology being monitored.

8.2 Alarm reporting

Alarm Reporting is a Stage 1 function that receives defect information for resources within an NE from Failure Event Detection and Reporting. Defects may be subjected to a brief persistence analysis. Failures are declared based on the persistence of defects. Faults are cleared based on the persistent absence of defects.

When a failure is detected, it is reported, usually as an alarm to Alarm Correlation and Filtering at Stage 2. These failures may directly affect service if there is no backup resource to assume the affected resource's function. If there is a backup resource, service will continue without immediate manual intervention[1]. Alarm Reporting is responsible for indicating in the alarm message whether or not the failure is directly service affecting or not. It also provides other information such as the type of fault condition, its severity, its location, and other attributes.

When a fault is cleared, it is reported with a clear message to Alarm Correlation and Filtering at Stage 2.

Alarm Reporting may perform some filtering of alarms and other events. For example, when there is a Loss of Signal (LOS) failure on a high rate signal the NE may suppress the reporting of Loss of Signal (LOS) and Alarm Indication Signal (AIS) alarms on carried lower rate signals.

There are three different types of monitoring and therefore three types of fault.

8.2.1 Equipment faults

Equipment monitoring monitors the various components of the NE, indicating when a component has failed or been removed. Equipment faults are always detected by the affected NE and are almost always reported by the affected NE. The NE usually sends an equipment alarm immediately without performing any persistence analysis.

8.2.2 Communication faults

Communication monitoring monitors the incoming (and in some cases outgoing) signals at each port in the Network Element to detect failure conditions. Some examples of communication faults are Loss of Signal (LOS), Loss of Frame (LOF) and Alarm Indication Signal (AIS). All communications faults are declared as the result of a persistence analysis performed by Alarm Reporting. For instance, a Loss of Signal is declared when there has been no valid signal for 2-3 seconds (the actual definition is quite precise and out of the scope of this Recommendation). Some failures are based on measurements detected directly, others are based on measurements made by the NE at the other end of the trail and reported in the overhead part of the signal in the reverse direction of transmission.

Communication faults always indicate that the problem is somewhere else, either in the transmission media, in an equipment component in an upstream NE, or occasionally another component of the same NE performing an upstream function.

8.2.3 Environmental conditions

Environmental monitoring monitors environmental conditions, such as temperature, humidity, and the presence of smoke and compares the levels of these measurements to limits for each measurement. If the limit is exceeded a notification (usually an alarm) is reported. These conditions are external threats to the proper functioning of the entire Network Element. Some of the conditions monitored by an NE are related to electrical power (either standard or battery backup). Other conditions relate to the physical security or data security of the NE. For the purposes of this Recommendation, these are considered part of environmental monitoring.

Environmental alarms are triggered independently of the detection of any equipment or communication fault. The limits should be set to trigger an environmental alarm before an equipment or communication fault occurs. This allows environmental alarms to provide early warning of trouble, so that proactive maintenance can occur before service is affected.

8.3 Alarm correlation and filtering

Alarm Correlation and Filtering is a Stage 2 function that receives alarms (and other fault notifications, including State Changes) from multiple NEs within a local domain. It may correlate alarms based on network topology and circuit layout information. It filters out redundant or

---

[1] If service restoration requires manual intervention, then the resources used do not fit the definition of a backup resource and the fault is considered service affecting.

symptomatic alarms and any non-persistent alarms. Its output is one fault indication for each trouble in its local domain, which contains the best information from the alarms that indicate the location, severity and type of the failure.

8.3.1 Processing of equipment alarms

Since equipment alarms report the exact component that is at fault, they are the most precise of all alarms. When many alarms are received, it is usually an equipment alarm that provides the best root cause information.

Equipment alarms from different NEs are always independent even if the failed components support the same trail. This is true even if an environmental fault has caused both equipment faults, because the two equipment components still must both be fixed. If more than one equipment alarm is received from the same NE, then they may not be independent. NEs usually prevent multiple equipment alarms from being triggered by the same underlying condition, but in some cases Alarm Correlation and Filtering must correlate equipment alarms from the same NE. This is done based on the equipment dependency relationships among all of the NE's components (bays, shelves, power supplies, timing generators, line cards, etc.).

Alarm Correlation and Filtering must have this equipment dependency information for all of the NEs that it monitors in its local domain. Included in this information should be trail termination points directly supported by each equipment resource and the physical location of each NE. The dependency relationships will be slightly different for each NE due to the differences in the components installed to meet different customer needs.

8.3.2 Processing of communication alarms

Communication alarms are detected when there is a problem with an incoming signal. The problem can be caused by an equipment fault, which will trigger an equipment alarm for the failed equipment from the NE. The problem can also be caused by a fault in a fibre, cable, or other transmission medium, which will not cause anything but communication alarms. Given this understanding, communication alarms are only important when they occur without related equipment alarms. If an equipment alarm is received, all downstream communication alarms on that trail must be superseded. This includes far-end communication alarms that are a result of the equipment fault.

If no equipment alarms related to the trail are received, but communication alarms are received, then the cause of the fault is either:

- outside of the local domain;
- some unmonitored piece of equipment;
- the transmission medium that carries the trail between equipment.

In any of these three cases, the most upstream communication alarm will stand as an independent fault and will be the best root cause information available within the local domain. The other communications alarms will be filtered out.

8.3.3 Processing of environmental alarms

Environmental alarms relate to an entire NE or all NEs at a given location, so they relate to many equipment components and many trails. For this reason environmental alarms should be given high priority. However, most environmental alarms only detect a potential threat and not the immediate effect on individual components or trails. Environmental alarms do not imply that any equipment has failed or any service has been affected. Service may be affected, but the environmental alarm itself will not state that. Also, if an environmental condition is affecting network components, it may not affect all of them, and it is important to retain information about network components that are in trouble. Thus it is important to separate environmental alarms from other alarms and not to filter out equipment and communication alarms that may or may not be caused by the environmental condition.

Alarm Correlation and Filtering must maintain environmental conditions as independent faults, however, equipment faults and communication faults will be related or linked to potentially related environmental conditions by Network Fault Localization at Stage 4.

8.3.4 Processing of state changes

In addition to alarms from NEs, Alarm Correlation and Filtering also receives administrative and operational State Changes from NE(s) Status and Control, a Stage 1 Configuration Management MAF. State Changes may provide additional information about the root cause of a condition in the network. State Changes also can serve as a back-up method of detecting faults. Administrative and operational State Changes should be treated like equipment alarms, since they always indicate a change in the ability of a resource to fulfil its function.

The fact that one or more resources have entered the "locked" or "shutting down" administrative state, which can only occur though maintenance intervention, should be recorded in any related fault indication. This will allow a technician to distinguish the activity of another technician, when the administrative State Change is the root cause of a condition, from a resource failure.

Operational State Changes from "enabled" to "disabled" result from a failure condition and will usually be accompanied by an alarm. In some cases an operational state change message will contain additional status information that may be useful for fault isolation and that is not contained in the accompanying alarm notification. If so, this information should be combined with the alarm information in the fault indication that is generated for the condition. If no such useful information is present in an operational state change notification, the notification can be superseded in the presence of alarms that explain the change in operational state.

8.3.5 Correlation of alarms among related trails

In addition to correlating alarms along a trail, Alarm Correlation and Filtering includes correlation of alarms among related trails wherever possible. The opportunity to correlate among related trails exists whenever an instance of the Alarm Correlation and Filtering MAF will monitor more than one type of trail. This is usually true. The other requirement is that the instance of this MAF has access to trail configuration data that reveals how trails are related. Alarms on a supported trail will be superseded when they are correlated to alarms on an upstream supporting trail. There are two types of correlation between facilities.

Alarm Correlation and Filtering correlates alarms of any type with any downstream alarms of any type on supported (client) facilities. This is done to eliminate symptomatic alarms. If the alarm at the lower level is less severe, then it is superseded. If the alarm at the lower level is more severe, then it remains an independent alarm because it is not fully explained by the higher level alarm. Another underlying condition is the likely cause of the lower level alarm.

8.3.6 Persistence analysis

Persistence analysis has two purposes. One is to filter out alarms that are quickly followed by a clear notification from the NE. The other is to filter out alarms that are quickly explained by another related alarm that contains better information on the root cause and severity of the underlying condition. Once an alarm state has persisted for the required time without being cleared or superseded by the arrival of a related alarm that explains it, a fault indication is declared and sent to Network Fault Event Analysis, Including Correlation and Filtering at Stage 3.

The required persistence should be long enough to allow for the correlation, at Stage 2, of all of the alarms triggered by a fault condition within the local domain in most circumstances and yet short enough to provide timely reporting to Stage 3. The range of 10 seconds to 10 minutes should allow enough flexibility, but should not be taken as a requirement. The trade-off between maximizing correlation and minimizing delay will be different depending on the number of affected customers or affected transmission capacity. The required persistence can be made shorter for faults that affect more capacity and longer for faults that affect less capacity.

8.3.7 Declaring fault indications

Alarms that survive the persistence analysis are ready to be reported out of the Alarm Correlation and Filtering MAF to Network Fault Event Analysis, Including Correlation and Filtering at Stage 3. Alarm Correlation and Filtering generates a fault indication for each independent persistent alarm. If multiple related alarms were received that have the same severity or otherwise could not be filtered out then they are combined into a single fault indication. Specifically, when the first of the multiple alarms achieves the required persistence, all related, active alarms that have not been superseded will be combined into a single fault indication.

The idea behind this approach is that as long as one of the alarms has persisted, then the underlying condition is persistent and requires prompt maintenance attention. The other related alarms that are active (not cleared) and independent (not superseded) provide additional information about the persistent underlying condition and should be included in the fault indication along with the first alarm that persists. This reduces the message traffic between Stage 2 and Stage 3 and reduces the processing load for Network Fault Event Analysis, Including Correlation and Filtering.

8.4 Network fault event analysis, including correlation and filtering

Network Fault Event Analysis, Including Correlation and Filtering is a Stage 3 function. It receives fault indications from multiple instances of Alarm Correlation and Filtering for different local domains. Fault indications are correlated across the trail from end-to-end and among related trails.

There is no persistence analysis subprocess in Network Fault Event Analysis, Including Correlation and Filtering. This is consistent with the Alarm Surveillance philosophy, stated in the first paragraph of clause 8, that when service continuity is lost or under threat, prompt incomplete information is preferable to delayed complete information.

8.4.1 Correlation of fault indications along a trail

Any fault indication to reach Network Fault Event Analysis, Including Correlation and Filtering that is not related to an active indication already received, triggers a report to Network Fault Localization. Subsequent related indications will be superseded and filtered out if they are downstream on the same trail and of equal or lower severity. Subsequent related fault indications that are upstream on the same trail and of greater severity will supersede the previous fault indication and trigger an updated fault indication to be reported to Network Fault Localization. Subsequent related fault indications that are upstream on the same trail but of lower severity or downstream on the same trail, but of greater severity will remain independent and will trigger separate reports to Network Fault Localization.

Note that fault indications based on far-end data read from the overhead in the signal in one direction of transmission must be correlated with the other (near-end) fault indications related to the opposite direction of transmission. Fault indications based on far-end data are considered to be downstream of all fault indications based on near-end data, because far-end data reflects conditions detected at the terminating equipment for the opposite direction of transmission.

Clears to active independent fault indications are also passed through to Network Fault Localization. Clears to fault indications that have been filtered out in Network Fault Event Analysis, Including Correlation and Filtering will be filtered out and not forwarded to Network Fault Localization.

8.4.2 Correlation of fault indications among related trails

In addition to correlation along a trail, Network Fault Event Analysis, Including Correlation and Filtering correlates fault indications among trails (circuits, facilities).

A fault indication on a supporting trail supersedes downstream fault indications on supported (client) facilities. This is done to eliminate symptomatic fault indications from other local domains. If the fault indication at the lower level is less severe, then it is superseded. If the fault indication at the lower level is more severe, then it remains an independent fault indication, but the two indications are linked. Linking means that each indication is augmented with a pointer to the other as a related indication.

The procedure for one fault indication superseding another depends on the order in which the fault indications are received and processed. If the superseding fault indication is processed earlier, the subsequent superseded fault indication will be filtered out and no message will be sent to Network Fault Localization for the superseded fault indication. If the superseding fault indication is processed later, an updated fault indication is set to Network Fault Localization reflecting the information in the superseding fault indication.

8.5 Network fault localization

Network Fault Localization is a Stage 4 function. This is where all surveillance information comes together. Network Fault Localization receives fault indications from Network Fault Event Analysis, Including Correlation and Filtering. It receives impairment indications and trend indications from Network Performance Monitoring Event Correlation and Filtering. Real-time information input from external trouble reports and other relevant sources may be included in this final stage of event correlation.

All three types of indication are correlated together. Any indication supersedes others that are downstream on the same or supported trail and of equal or less significance. Fault indications are the most significant indications followed by impairment indications and trend indications.

If there is an active fault indication for a given trail, then related impairment indications and trend indications are superseded and prevented from generating independent trouble work items and are associated with the fault indication as detail information. If there is no fault indication for a trail, then the most severe and upstream impairment indication will stand as an independent indication with any corresponding trend indication information associated as detail information. If there is no impairment indication or fault indication, then a trend indication will stand as an independent indication.

8.5.1 Patterning of indications

Patterning is the correlation of indications across the many supported (client) resources of a supporting resource to infer the occurrence of unreported events on the supporting resource. Patterning is done to compensate for potential loss of communication with NEs or other OSs in the TMN and the resulting loss of notifications for the supporting resource. This method can also be used to infer events affecting unmonitored, common resources in the network, such as cables, conduits, and power sources. When a pre-set fraction of all provisioned supported resources have indications of a given type logged against them, then an indication is generated on the supporting resource and the individual indications on the supported resources are superseded (except in the case of environmental conditions, see 8.5.2). Indications generated as a result of patterning should be distinctively labelled.

An indication generated as a result of patterning is cleared, when enough of the contributing indications are cleared that the number of active contributing indications no longer meets or exceeds the pre-set fraction of all provisioned supported resources required to generate a pattern. An indication generated as a result of patterning is also cleared when a clear is received explicitly for the supported resource, following the recovery of a failed TMN communications link.

Patterning requires a database of the dependency relationships between managed resources including unmonitored resources. Patterning may be done recursively to infer trouble with a resource such as a conduit or power supply from indications in many indirectly supported resources (as depicted in Figure 2 located in 5.1.5).

8.5.2 Correlation of indications among related resources

In addition to correlation along and among trails, Network Fault Localization uses patterning to correlate indications on trails and other resources supported by the same unmonitored resource. This allows the generation of a single work item (trouble ticket) for a cable, or other unmonitored equipment or communications resource and the suppression of work items for each affected supported resource.

Indications reporting environmental conditions are treated differently from indications reporting an equipment or communications condition, because the environmental condition may or may not actually affect the dependent resources. When an indication reporting an environmental condition is correlated to an indication on a dependent resource, the two indications are linked. Each indication remains independent but with a pointer to the other added prominently as additional information.

Locations where NEs are housed can be treated as supporting the NEs and indirectly supporting facilities monitored by NEs. Equipment faults can be marked as related to environmental conditions existing at the location or NE that houses the reporting equipment component. Communication faults can be marked as related to any environmental condition existing at any location or NE that houses equipment that supports the facility impacted by the communications alarm.

9 Configuration management

Status and Control is the group of MAFs in Configuration Management that relate to the tracking of the status of the network and its components. The two MAFs from this group that are important to Network Event Correlation are described below. These two MAFs are considered part of the Alarm Surveillance Flow because the state change information that they provide may indicate that a resource is completely inoperable and these State Changes are correlated with alarms at Stage 2.

9.1 Notification of state changes by NEs

Notification of State Changes by NEs is a Stage 0 function that tracks state attributes of resources within a Network Element. When one of these states changes, this MAF sends a state change notification to NE(s) Status and Control at Stage 1.

9.2 NE(s) status and control

NE(s) Status and Control is a Stage 1 function that tracks the states of resources within a local domain. Administrative and operational State Changes are of interest to Network Event Correlation; other State Changes are not. When NE(s) Status and Control is notified of an operational or administrative State Change, this MAF sends a notification to Alarm Correlation and Filtering in Fault Management for correlation with alarms.

9.2.1 Operational state changes

A resource's operational State Changes from "enabled" to "disabled" when a condition occurs that renders the resource completely inoperable. Conditions that cause operational State Changes will almost always also trigger alarms for the same resource in Failure Event Detection and Reporting that will also be sent to Alarm Correlation and Filtering. Thus operational state change notifications will almost always be redundant with alarms. Yet State Changes enhance the robustness of Network Event Correlation by providing another backup method of detecting failures. Also there are cases when State Changes can provide additional useful information about a condition in the network.

9.2.2 Administrative state changes

Administrative State Changes can only occur due to maintenance action. When the administrative state for a resource changes from "unlocked" or "shutting down" to "locked", the resource will cease to perform its function. Thus an administrative State Change can be service affecting. The locking of one resource may impact other resources that are dependent or downstream of the locked resource. When used properly, the administrative state for dependent and downstream entities are locked first to prevent alarms from being generated. When used improperly, the impacted resources become disabled while remaining in the unlocked administrative state and alarms are triggered in Failure Event Detection and Reporting. It is important to send administrative State Changes to Alarm Correlation and Filtering so that maintenance activity can be identified as the root cause of resulting conditions.

The correlation of administrative State Changes with alarms will allow a technician to distinguish the activity of another technician from a resource failure in some situations.

SERIES OF ITU-T RECOMMENDATIONS

Series A	Organization of the work of ITU-T
Series B	Means of expression: definitions, symbols, classification
Series C	General telecommunication statistics
Series D	General tariff principles
Series E	Overall network operation, telephone service, service operation and human factors
Series F	Non-telephone telecommunication services
Series G	Transmission systems and media, digital systems and networks
Series H	Audiovisual and multimedia systems
Series I	Integrated services digital network
Series J	Transmission of television, sound programme and other multimedia signals
Series K	Protection against interference
Series L	Construction, installation and protection of cables and other elements of outside plant
Series M	**TMN and network maintenance: international transmission systems, telephone circuits, telegraphy, facsimile and leased circuits**
Series N	Maintenance: international sound programme and television transmission circuits
Series O	Specifications of measuring equipment
Series P	Telephone transmission quality, telephone installations, local line networks
Series Q	Switching and signalling
Series R	Telegraph transmission
Series S	Telegraph services terminal equipment
Series T	Terminals for telematic services
Series U	Telegraph switching
Series V	Data communication over the telephone network
Series X	Data networks and open system communications
Series Y	Global information infrastructure and Internet protocol aspects
Series Z	Languages and general software aspects for telecommunication systems

Geneva, 2001