

Method for managing service policy in a telecommunication network

The present invention concerns the field of service quality control in a data telecommunication network.

5 Modern communication networks consist of a plurality of communication nodes connected together by communication links. These communication nodes include terminal equipment such as computer stations, mobile communication terminals or application servers. They also include intermediate equipment such as relays and routers that serve for interconnection and routing of the data traffic between the items of terminal equipment.

10 Communication links are based on a wide range of physical transmission technologies. Cabled technologies such as optical fibre, Ethernet links or the like and radio technologies according to a large panel of protocols such as WiFi, WiMax, satellite links or the like, can be cited. Each of these technologies leads to different properties concerning the transmission channel both in terms of available bandwidth and error rate or variability of this available bandwidth. At the present time, the
15 dominant technology is a packet communication technology, typically the IP (Internet protocol) network.

These communication networks are used to implement a great diversity of services, such as voice communication exchange, typically so-called VoIP (Voice over IP) services, broadcasting of multimedia programs, typically television services, or access to videos on demand or access to data, typically access to web services.

5 These applications have bandwidth and latency requirements that are particular thereto. Packet transfer networks historically do not guarantee a given bandwidth or latency to a particular service. On the other hand, circuit communication networks are based on the establishment of a communication circuit between the end items of communicating equipment that guarantees bandwidth and latency between these items
10 of equipment, as soon as the circuit is established. One example of such circuit networks is the switched telephone network.

To enable the development of services having constrained requirements in terms of bandwidth and/or latency on packet communication networks, a set of technologies have been developed and are known, among other things, by the term QOS (Quality
15 Of Service). These technologies are based on the definition of a set of service categories subject to the communication constraints. The network is partitioned into security domains applying a service quality policy. The data packets entering the domain are marked according to the category of the service that sends them (colouring), the data streams are admitted or refused at the entry, and their rate may be
20 regulated (policing). Priorities are attributed to the various categories and the intermediate items of equipment of the network that deal with congestion use these priorities in order to eliminate the ones with the lowest priority and favour the transmission of packets belonging to the data streams with the highest priority. These technologies are supplemented by other mechanisms such as access control lists
25 (ACL), routing rules, services admission and the like, in order to form a service policy. This service policy defines the admission, broadcasting and bandwidth required for a set of services on a communication domain. This service policy is implemented at various points in the network, points the behaviour of which is at least partly driven by this policy. Among these points are routers, applications or others. All
30 these points constitute points for implementing the service policy, or PEPs (Policy Enforcement Points).

The physical links that are the most tricky to manage are radio links. These links are subject to high variations in their bandwidth according to their environment and are subject to attenuation, interference or the like, which dynamically degrades their

transmission capacities. The variations in bandwidth of these radio links have an impact on the service policy. They cause a variable congestion level in the link. The service policy must therefore advantageously take into account this congestion level. In the absence of such consideration, it is necessary to configure the service policy for
 5 a minimum bandwidth and therefore to under-use the radio link. This is particularly detrimental for satellite radio links where the bandwidth is precious and expensive.

It is also usual to interconnect various domains each having their service policy and able to be subject to different security constraints. Fig. 1 illustrates an example of network architecture illustrating this mechanism. A first communication network 1.6
 10 includes a radio link 1.1. This radio link, typically a satellite link but any other radio technology could be used, operates by means of two modulators/demodulators or modems 1.2 and 1.3, one located before the link and the other after. Before the modem 1.2 there is a router 1.4. The same applies downstream, the router 1.5. The services therefore pass through the domain in the form of data streams using the radio
 15 link. Other routers, not shown, may be connected to these routers. It is considered here that the router 1.4 is the boundary of the domain 1.6. In this regard, it is one of the points implementing the service policy. In this regard, it imposes the rules of the current service policy on the various streams and requests entering the domain at its level. The radio link 1.1, therefore subject to bandwidth variations, has an influence
 20 on the actual ability of the incoming services at the router 1.4 to properly pass through the domain 1.6. It is therefore advantageous to pass on the information on bandwidth actually available on the link 1.1 at the level of the service policy applied at the point 1.4. This can be done by feeding back information on the congestion level from the modem 1.2 upstream of the radio link to a module 1.12 managing the service policy.
 25 This module, then being informed of the changes in the bandwidth actually available on the link 1.1, can accordingly modify the service policy and reflect its modifications in particular to the router 1.4. These modifications will also have an impact on all the actors in this service policy, the application level among other things, this not being shown in the figure.

30 If a second domain were to be considered, here the domain 1.7, connected to the domain 1.6 as shown in the figure, the important point is here the security of this domain. The domain 1.7 is a domain having a higher security level than the domain 1.6, for which it uses the services. The high-security network is referred to as red network and the low-security network is referred to as black network. The red network

1.7 communicates with the black network 1.6 through security gateways 1.8 and 1.9 in charge of filtering and encrypting the traffic. The red traffic coming from the router 1.11 is therefore encrypted by the security gateway 1.8. It next passes through the black network 1.6 and in particular through the radio link 1.1. It is decrypted by the security gateway 1.9 at the output from the black network towards a red router 1.10. The domain 1.7 must therefore be considered to be a special domain in the sense of services having its own service policy managed by a service policy management module 1.13. This policy is applied by means of instructions 1.15 *vis-à-vis* policy implementation points including the router 1.11 upstream of the black network and the link 1.1. The domain 1.7 is also special because of its high security level making it a red network.

When the bandwidth variation on the radio link 1.1 by the service policy both of the red network and of the black network is considered, it should be noted that this requires to be coordinated. As a matter of fact, seen from the black network, the radio link will pass services that are also qualified as black services but some of which are in fact streams transporting, in encrypted form, streams coming from red services. The red services are therefore subject on one hand to the service policy of the red network in the red domain and on the other hand to the service policy of the black domain when they pass through this domain. The consideration, by the manager 1.13 of the service policy of the red domain, of the bandwidth variations on the radio link in the black domain requires a feedback of information 1.14 between the modem 1.2 and this manager 1.13. Alternatively, we could consider that the service policy manager 1.13 of the red domain communicates with that one 1.12 of the black domain in order to harmonise their policy. Another alternative consists of designing the service policy manager 1.13 of the red domain and that one 1.12 of the black domain as a whole, in order to harmonise their service policy. These embodiments require a communication of information from the black domain to the red domain, which constitutes a security weakness.

The invention aims at solving the above problems by means of a method of managing the service policy applied to a black network comprising a radio link and a red network connected to this black network managed by a predetermined set of profiles. Congestion on the radio link is detected by any means by the black domain. Congestion on the radio link is detected by a congestion detection method operating by means of exchanges in the red domain. A service policy profile then established

corresponds for each domain to a detected congestion level. This method is applicable to a network with more than two security levels.

In this way, by means of the predefined profiles and the independent detection of congestion on the radio link in each domain, coherence is enabled for the service policies implemented in the two domains according to the congestion level on the radio link with no need to exchange information fed back from the black domain to the red domain, guaranteeing in fact perfect isolation of the security domains in question.

The invention concerns a telecommunication network, said network consisting of a first security domain, called black domain, comprising at least one radio link, at least one second domain with higher security, called red domain, connected to the black domain, characterised in that it includes: a list of profiles, each containing a set of parameters applicable to the management of the service policy for the black domain, each profile being adapted to a given congestion level on the radio link; a list of profiles, each containing a set of parameters applicable to the management of the service policy for each red domain, each profile being adapted to a given level of congestion on the radio link; means for detecting the congestion level on the radio link in the black domain; means for detecting the congestion level on the radio link in each red domain, said detection being based solely on exchanges of messages in said red domain; means for selecting a new profile adapted to the congestion level of the radio link in the case of detection of a change in said congestion level in the black domain, and means for selecting a new profile adapted to a congestion level on the radio link in the case of detection of a change in said congestion level in each red domain.

According to a particular embodiment of the invention, the black domain includes a black router upstream of the radio link that constitutes one of the points implementing the service policy of the black domain and that aggregates the traffic received from the red domain and the traffic coming from the black domain in order to produce a data stream towards the radio link.

According to a particular embodiment of the invention, the black domain includes a black router upstream of the radio link that constitutes one of the points implementing the service policy of the black domain, and a switch that aggregates the traffic received from said router of the black domain and the traffic coming from the red domain in order to produce a data stream towards the radio link.

According to a particular embodiment of the invention, the traffic coming from the red domain passes through a red router that constitutes one of the points implementing the service policy of the red domain and then through a security gateway in charge of encrypting said traffic.

5 The invention also concerns a method for managing the service policy in a telecommunication network, said network consisting of a first security domain, called black domain, comprising at least one radio link, at least one second domain with higher security, called red domain, connected to the black domain, characterised in that it includes a step of defining a list of profiles, each containing a set of parameters
10 applicable to the management of the service policy for the black domain, each profile being adapted to a given congestion level of the radio link; a step of defining a list of profiles each containing a set of parameters applicable to the management of the service policy for each red domain, each profile being adapted to a given congestion level on the radio link; a step of detecting the congestion level on the radio link in the
15 black domain; a step of detecting the congestion level on the radio link in each red domain, said detection being based solely on exchanges of messages in said red domain; a step of selecting a new profile adapted to the congestion level on the radio link in the case of detection of a change in said congestion level in the black domain and a step of selecting a new profile adapted to a congestion level on the radio link in
20 the case of detection of a change in said congestion level in each red domain.

The features of the invention mentioned above, as well as others, will emerge more clearly from a reading of the following description of an example embodiment, said description being given in relation to the accompanying drawings, among which:

25 Fig. 1 illustrates an example of architecture of a network interconnecting two security domains.

 Fig. 2 illustrates an example embodiment of the admission of services upstream of the radio link.

 Fig. 3 illustrates another example embodiment of the admission of services upstream of the radio link.

30 Fig. 4 illustrates an example of a network architecture based on a multipoint-to-point radio link.

 Fig. 5 illustrates an example of a network architecture comprising two red networks.

Fig. 6 illustrates the flow diagram of the method for managing the service policy according to an example embodiment of the invention.

Fig. 2 illustrates an example embodiment of the admission of services upstream of the radio link 2.1. The elements described in Fig. 1 are presented again with the same minor numbers. It can be seen that the black router 2.4 upstream of the modulator 2.2 of the link receives two types of traffic. A first type of traffic referenced 2.15 is black traffic coming from the black domain. Although shown by a single arrow, this traffic may be composed of multiple links at the input of the router 2.4. This traffic is present again as a component of the traffic 2.17 at the output from the router. It also receives traffic 2.16 coming from the security gateway and prior to a red router 2.11. This traffic is therefore red before the protected gateway. After 2.16, it forms part of the black domain and is therefore also considered to be black. However, black traffic transports red traffic in encrypted form. Seen from the black router 2.4, it receives as an input a set of services 2.15 and 2.16, which it aggregates as an output in order to produce a data stream 2.17 towards the link. For it, all these services are black and it applies a defined service policy thereto. A potential differentiation may then be made in the configuration of the service policy. In this diagram, the red services undergo a first so-called red service policy applied in the red domain by the red router 2.11 or any other means upstream of the passage in the black domain. These same services undergo a second black service policy applied in the black domain by the black router 2.4, or any other means, to the black services that transport them.

Advantageously, the service policy includes an admission control service where the admission of the services depends on the granularity of the topology of the network, the link or links and/or partitioning of the topology of the link.

A second diagram can be foreseen. This is the diagram illustrated by Fig. 3; the minor references correspond to Fig. 2. In this diagram, the router 3.4, which is one of the points implementing the service policy of the black domain, is not in series with the traffic coming from the red domain 3.16. The traffic coming from the red domain and the purely black traffic are connected by a switch 3.18, which is not a service policy implementation point. In this new diagram, the black service policy is not applied to the services carrying red traffic. The red services are then subject only to the red service policy implemented by the red router 3.11, which is one of the points implementing the service policy of the red domain.

Congestion on the radio link can be detected by any means in the black domain. It can be done in a traditional way by an explicit method such as the feedback of information from the modem 1.2 to the service policy management module 1.12. Any other method can also be foreseen, for example implicit methods such as a comparison
5 of the rates sent by the router 1.4 with the rate received by the router 1.5 or an analysis of the traffic jitter.

The problem is more tricky with regard to the detection of congestion in the red domain. To guarantee perfect partitioning of the domains, it is not desired to create a data link going back from the black domain to the red domain or any feedback of
10 black domain congestion information to the red domain via the security gateway. Sending back any congestion information available at the modem 1.2 to the service policy management module 1.13 of the red domain cannot therefore be foreseen.

It is therefore necessary to use a method for detecting congestion on the radio link based solely on exchanges of messages in the red domain. Such methods exist.
15 French patent application number 10/04220 of the same applicant can for example be cited. This method is based on a periodic measurement of the rate sent by the red router 1.15 to the security gateway 1.8 and the rate received by the red router 1.10 coming from the security gateway 1.9. These rate measurements are exchanged between the two red routers in the form of red messages. In practice, they are typically
20 messages of layer 3 of the OSI model. When a received rate less than the sent rate is detected, congestion of the radio link is deduced therefrom. Using such a congestion detection method therefore enables the module managing the service policy of the red domain to become aware of the state of congestion of the radio link in the black domain without requiring any transmission of information between the black domain
25 and the red.

Once the methods for detecting congestion on the radio link particular to each domain are established, it is also necessary to ensure consistency between the service policies applied in response to a congestion, in order to ensure harmonious management of the radio link. For this purpose, a set of service profiles is established
30 in each domain. Each service profile contains a set of parameters defining a service policy. At a given instant, one of these profiles is active. This means that the service policy applied at this instant by the module managing the service policy of the domain uses the parameters of said profile for managing the service policy. When a congestion on the radio link is detected, a change in the active profile and therefore a

change in the service policy of the domain is caused. In this way, a match is established between the congestion level detected and a given service policy profile. Consistency is thus afforded by the configuration of the different profiles.

Consistency is also due to the fact that a degradation of the available bandwidth at the radio link would cause a loss of packets in the traffic passing over this link. This loss of packets typically affects the various services in a uniform manner. Because of this, the congestion level is detected homogeneously by the various domains.

This configuration is done according to the topology of the black domain. In the case illustrated by Figure 2, where the router that implements the service policy of the black level is in series with the data stream of the red services, the parameterising of the profiles must take account of the fact that these red services are subject to the service policies of the two domains. Conversely, in a topology as illustrated by Fig. 3, this is not the case. The red services are subject only to the red service policy and then access the link through the black domain without any obstacle.

The method for detecting the congestion on the radio link implemented in the red domain means that a congestion is detected only in the case of red traffic passing over the link. This is because, if congestion occurs because of the black traffic and at that moment no red traffic is passing, this congestion will be detected by the black domain but not by the red domain. The active profile of the red domain will therefore remain a profile corresponding to a non-congested link. This is not problematic, because of the absence of red traffic. If red traffic appears, it will be subject to the service policy adapted to a link without congestion. Very rapidly, the congestion will be detected since it will produce a difference between the sent rate and the received rate. This detection will cause a switching of profile in the red domain. In the presence of red traffic, the service policy therefore adapts automatically to the current actual state of congestion of the link.

This phenomenon may also occur in the black domain according to the used method for detecting traffic. In particular, if a detection method based, as in the red domain, on a measurement of actual rate is used, congestion will be detected only in the case of actual black traffic. And again, this occurs only in the so-called parallel topology illustrated by Fig. 3. In the so-called series topology illustrated by Fig. 2, in the absence of black traffic, congestion can be caused only by red traffic. Since this red traffic appears as black traffic to the black router 2.4, it automatically participates in the measurement of congestion in the black domain. One therefore observes that

congestion on the link is then automatically detected in the black domain whether or not there is any purely black traffic at this moment.

These considerations show that the moment of switching from one service policy profile to another in order to adapt it to a congestion level may be offset in time with respect to the moment of the congestion. This however occurs only in the absence of traffic in the domain concerned. As soon as traffic appears, congestion is detected and the profile is adapted. We therefore have a set of profiles selected at a given moment that is always adapted to a congestion level on the radio link and the reality of the traffic in the domain in question.

The described method operates whatever the topology of the radio link in question. It includes the multipoint-to-point topologies as illustrated by Fig. 4. This figure repeats the references in Fig. 1, except for the second chain comprising the second red router 4.11, the second security gateway 4.8, the black router 4.4 and the modem 4.2. In this topology, the red services follow two connection paths to the same destination. Red profiles are defined homogenously and apply to each of these communication paths. The method adapts using a method for detecting the congestion on the link that will add the rates leaving the red routers 1.11 and 4.11 and compare them with the rate entering the red router 1.10 on reception. The same type of adaptation is made in the congestion detection method used in the black domain. The same applies to a point-to-multipoint topology.

The method can also apply to a plurality of secure networks and therefore qualified as red domains, as illustrated by Fig. 5. In this figure, the references in Figure 1 are repeated. A second secure domain 5.7 is connected to the black domain 1.6. This domain includes a red upstream router 5.11 connected to a security gateway 5.8. This gateway is connected to the same black router 1.4 as the security gateway of the first red domain 1.7. Downstream, the traffic coming from the black domain is redirected according to its origin to one of the security gateways 1.9 for the first red domain 1.7 and to the security gateway 5.9 for the second. It is then redirected to the red router 5.10 for this second red domain. The security levels implemented in the two red domains may be different or equal depending on the embodiments. This example extends to any number of secure domains. Each of the domains will apply a service policy that is particular thereto. Each of these policies results in a set of service profiles implemented in said domain. The parameterising of these sets of profiles ensures consistency of the service policy between the various domains. Each domain

implements a method for detecting congestion on the radio link that is particular thereto and does not involve any exchanges of messages from the black domain to one of the red domains. The active profile at a given instant in each domain is always adapted to the level of congestion of the radio link and the reality of the traffic in the domain in question.

An example of a method for managing the service policy according to the invention is illustrated in Fig. 6. During a first step 6.1, a list of profiles is defined for the black domain. Each profile contains a set of parameters applicable to the management of the service policy. The result is therefore a list, the first profile of which corresponds to the service policy required for the black domain in the absence of congestion on the radio link. This will be the default active profile. The following profiles are adapted to a greater and greater congestion level on the link. During a second step 6.2, a list of profiles is likewise defined for the red domain. Once these profiles are consistently defined, they are established in each domain. According to the example embodiment of the invention illustrated by Fig. 6, there are two domains: a red domain 1.7 and a black domain 1.6. The module for managing the service policy of each domain selects a profile in the list of profiles defined for its domain during step 6.31 for the red domain and 6.32 for the black domain. Advantageously, the first profile intended for the policy required in the absence of congestion is selected by default as the active profile in each domain. Each domain then uses its method for detecting congestion. This is done during step 6.41 in the red domain and 6.42 in the black domain. The method for detecting congestion in the black domain can use any explicit or implicit detection technique. Direct measurement of the bandwidth at the modem, use of an error rate, the comparison of the sent and received rates or others can be cited. These methods are applied to the link whatever its topology among the following topologies: point-to-point, point-to-multipoint and multipoint-to-point.

During step 6.51 in the red domain and 6.52 in the black domain, the state of congestion on the link is tested. This test may be done in a regular manner or be asynchronously triggered by a modification of the state of the link, typically upon interrupt. The test consists of detecting a modification of the congestion level. When such a modification is detected, step 6.61 is performed if in the red domain or step 6.62 if in the black domain, in order to select a new profile taking account of the new congestion level.

Thus, the profile used by the service policy of each domain is always consistent with the congestion level detected on the link. According to the congestion detection method used, it is possible that the two domains do not detect congestion and therefore do not perform the profile switching in a synchronised manner. This may
5 occur in the case of absence of traffic in the domain. In this case, any delay in the detection has no effect since precisely there is no traffic on the link in the concerned domain.

The invention therefore enables managing a consistent service policy over all the security domains using a radio link, without requiring a transfer of information
10 between the domains and therefore without weakening the security mechanisms established based on the security gateways.

CLAIMS

1. Telecommunication network, said network consisting of a first security domain, called black domain (1.6), comprising at least one radio link (1.1), at least
5 one second domain (1.7) with higher security, called red domain, connected to the black domain, characterised in that it includes:
 - a list of profiles, each containing a set of parameters applicable to the management of the service policy for the black domain, each profile being adapted to a given congestion level on the radio link;
 - 10 - a list of profiles, each containing a set of parameters applicable to the management of the service policy for each red domain, each profile being adapted to a given congestion level on the radio link;
 so that a service policy profile corresponds, for each domain, to a detected congestion level, and in that it also includes:
 - 15 - means for detecting the congestion level on the radio link in the black domain;
 - means for detecting the congestion level on the radio link in each red domain, said detection being based solely on exchanges of messages in said red domain;
 - means for selecting a new profile adapted to the congestion level of the radio link in the case of detection of a change in said congestion level in the black domain;
 - 20 - means for selecting a new profile adapted to a congestion level on the radio link in the case of detection of a change in said congestion level in each red domain.

2. Telecommunication network according to claim 1, characterised in that the black domain (1.6) includes a black router (2.4) upstream of the radio link, which
25 constitutes one of the points implementing the service policy of the black domain and which aggregates the traffic (2.16) received from the red domain (1.7) and the traffic (2.15) coming from the black domain (1.6) in order to produce a data stream (2.17) towards the radio link (2.1).

- 30 3. Telecommunication network according to claim 1, characterised in that the black domain (1.6) includes a black router (3.4) upstream of the radio link (3.1), which constitutes one of the points implementing the service policy of the black domain, and a switch (3.18) that aggregates the traffic received from said router of the

black domain and the traffic (3.16) coming from the red domain in order to produce a data stream (3.17) towards the radio link.

4. Telecommunication network according to one of claims 1 to 3, characterised
5 in that the traffic coming from the red domain passes through a red router (1.11) that constitutes one of the points implementing the service policy of the red domain and then through a security gateway (1.8) in charge of encrypting said traffic.

5. Method for managing the service policy in a telecommunication network, said
10 network consisting of a first security domain, called black domain, comprising at least one radio link, at least one second domain with higher security, called red domain, connected to the black domain, characterised in that it includes:

- a step (6.1) of defining a list of profiles, each containing a set of parameters applicable to the management of the service policy for the black domain, each profile
15 being adapted to a given congestion level of the radio link;

- a step (6.2) of defining a list of profiles, each containing a set of parameters applicable to the management of the service policy for each red domain, each profile being adapted to a given congestion level on the radio link;

so that a service policy profile corresponds, for each domain, to a congestion
20 level detected, and in that it also includes:

- a step (6.42) of detecting the congestion level on the radio link in the black domain;

- a step (6.41) of detecting the congestion level on the radio link in each red domain, said detection being based solely on exchanges of messages in said red
25 domain;

- a step (6.62) of selecting a new profile adapted to the congestion level on the radio link in the case of detection (6.52) of a change in said congestion level in the black domain;

- a step (6.61) of selecting a new profile adapted to a congestion level on the
30 radio link in the case of detection (6.51) of a change in said congestion level in each red domain.