CRATOS
Center for Research on the Applications of Telematics
to Organisations and Society

A Virtual Network Service for Integrated-Services Internetworks

Luca Delgrossi and Domenico Ferrari

{ldgrossi, dferrari}@pc.unicatt.it

Quaderni del CRATOS

Serie di Telematica

CTR-T97-002

Università Cattolica del Sacro Cuore – Piacenza, Italy

Abstract

Integrated-services internetworks are usually defined as those internetworks that are designed to carry a variety of traffic types and to satisfy the main requirements of each type of traffic. In this paper, we argue that the services needed to satisfy such requirements are not the only ones those internetworks will have to offer: in order to become a true basis for human society in the next century, they must be able to reproduce the most important features of networkless society, one of which is the provision of support to a variety of groups. We show how a service that will allow any interested user to create a virtual network for a group on top of any integrated-services internetwork can be designed. We call such a virtual network a “supranet”.

1. Introduction

The term “integrated-services” that is normally coupled with “network” or “internetwork” refers to the integration of services to be offered at the network and transport layers of a communication infrastructure’s architecture. These are the services needed by various types of traffic to be carried with adequate quality of delivery by the infrastructure; for example, best-effort service for data, statistically guaranteed-performance service for conferencing-quality video and voice, and so on. These are, however, higher-level services and same-level but different-purpose services that will also have to be offered by future infrastructures if they are to become a true basis for human society in the information age, i.e., to be able to reproduce the most important features of yesterday’s and today’s networkless societies. We believe that one important such feature is the ability to form groups. Groups are designated by a variety of terms, depending on their context and purpose: for example, association, alliance, congregation, conglomerate, corporation, consortium, club, and so on. When all the members of a group have access to an internetwork, the group could often benefit greatly from the use of that internetwork by the group’s members to communicate and/or collaborate with each other for the purposes of the group itself.

When a group wants a telematic infrastructure for its purposes, there is an alternative to creating a virtual network on top of the physical internetwork all of its members have access to: the establishment of a network private to a group. This solution generally has some advantages; for instance, it can achieve more easily a given level of security with respect to the world outside the group as well as a given level of control (note that the types of control that may be required are discussed in Section 2). However, it is bound to be more expensive and less flexible than a virtual network defined on top of a physical internetwork; low flexibility, for instance, is demonstrated by the cumbersome procedure to be followed for adding a new member to the group’s private network, or for changing any other characteristic of the group or of its network. Furthermore, the integration of this group-support service with the other services referred to alone is very convenient for the group’s members, who will not have to cope with two different networks for their group-related and non-group-related communications, respectively; this argument is especially convincing when one considers the predicament of a user who is a member of several groups, as most users indeed are.

There are virtual networks that cannot satisfy the requirements of a given group. In the next section, we discuss and define the six types of requirements that groups to different extents have. We call a virtual network capable of satisfying such requirements a *supranet* (from the Latin *supra*, which means “on top of”). Section 3 discusses the design of a service enabling users to create supranets to their specifications. Section 4 concludes the paper.

2. Requirement Types

A thorough analysis of a wide variety of groups and of their needs has resulted in the following list of requirement types:

1) Membership: only “members” should be allowed to access the network, i.e., to make use of the available services. Restricted membership implies the existence of an admission control policy, possibly enforced by a central authority or committee ruling the network. Members may belong to one of several classes, differing from the others in privileges and obligations.
2) Topology: the paths (consisting of nodes and links) that connect members in the network may be explicitly designed and tailored to specific user needs. This provides the ability to decide what paths should messages follow for security, control, or other reasons. It is desirable that the network’s topology be dynamically adjustable based on new needs or convenience.
3) Capacity: as for paths, it should be also possible for the network designer to define the capacity of the resources involved in the communications. Such resources include the bandwidth of the links, as well as the CPU power and memory size of each node. Resource management functions that enable a dynamic adjustment of the resources’ capacity are also desirable.
4) Security: communications among members may need to be kept private. In such cases, both the intrusion of outsiders and the eavesdropping by other members must be avoided. Network users may need to be guaranteed on the identity of their interlocutors. Different applications may have different requirements: some may request services that guarantee the authentication or the anonymity of the sender, others may need mechanisms that help avoid the denial or the repudiation of service.
5) Connectivity: this requirement consists of the ability to control which services a member is allowed to benefit from and which other members he can reach and communicate with. This allows the

network authority to "hide" services from a member, e.g., when he does not possess appropriate privileges, or to hide the presence of other members, e.g., for confidentiality reasons. Ideally, it should be possible to dynamically establish or tear down connectivity among members.

6) Multicast: this is the ability to send messages to a subgroup of the members in a secure manner, so that the messages cannot reach a member which is not included in the subgroup. As an alternative, a message reaching an undesired destination should be useless, e.g. indecipherable, for that destination.

By considering the first three types of requirements, (1), (2), and (3), it is easy to realise that they correspond to the specifications that are needed to design any network: the members of the group correspond to the hosts[1], whose set is restricted to those explicitly listed; and the locations and sizes of the virtual links and routers are determined by the topology and resource capacities assigned by the network designer[2]. Of course, when the description provided for the virtual network to be built is to be mapped onto a physical network, the set of members and their interactions must be protected from outside encroachments. Thus, security from external intruders must be guaranteed. These protections are implicit and compulsory requirements; they are not dependent on the preferences of the supranet's creator. Once the creator has expressed his wishes with respect to (1), (2), and (3), and the implicit protections have been implemented, a virtual network has been fully specified.

The three requirement types (4), (5), and (6) are those that make the virtual network a true supranet. Since the null requirement is acceptable for any of these three types (while it is not for (1), (2) or (3)), there may be supranets that are just simple virtual networks without any special properties. We consider these as degenerate supranets, and those for which at least one of the types (4) through (6) has a non-null requirement as true supranets. Thus, supranets are virtual networks, but not all virtual networks are supranets (for example, the Mbone [1] is a virtual network but is not a supranet).

3. Service Design

The service discussed in this paper can be offered if the integrated-services internetwork is equipped with

- an appropriate architectural support for supranets,
- tools users can exploit to
    - create a supranet and
    - manage it.

Here, we can only sketch a description of the architectural support. The full paper will include treatment of the supranet creation and management tools that will be required.

3.1 Address Space

A supranet has its own address space. Supranet addresses are assigned to supranet hosts and routers by the creator. These virtual network components are to be mapped onto physical components; a one-to-one mapping will usually be preferred for the sake of simplicity. Thus, a physical host that is also a supranet host has two addresses, a physical one and a virtual one. Furthermore, a supranet is endowed with its own name space and needs a centralised or distributed name server that translates supranet names (e.g., supranet URLs) into supranet addresses.

---

[1] A "supranet host" is a supranet node with a "member", e.g., a supranet user, on it. If there are no users associated with it, a node is called "supranet router". In this paper, for the sake of simplicity, we assume that each supranet host corresponds exactly to one "member".

[2] The network designer of a supranet is called the "supranet creator". The creator determines the characteristics of the specific supranet to be created and is responsible for its actual construction, regulation and activation. The creator may be an individual as well as a team or committee of more individuals.

3.2 Architecture

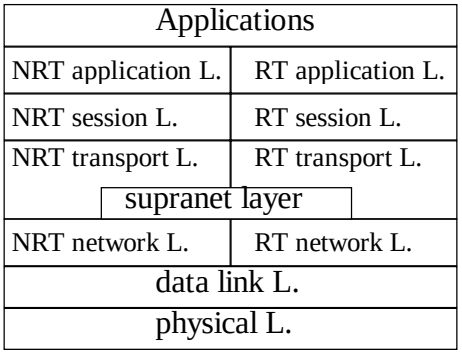

Figure 1: Double-stack architecture of an integrated-services internetwork
(SN = supranet layer; L = layer; NRT = non-real-time; RT = real-time)

In a supranet there will generally be non-real-time applications and real-time applications. For economic and good-quality communications, each application will have to use service types suitable for its needs. To show how a supranet service can be added to those already provided by an integrated-services internetwork, we assume that the internetwork's architecture is based on two stacks: one for non-real-time traffic and one for real-time traffic. Figure 1 shows one possible architectural solution for the supranet (SN) layer: that between the network and the transport layer.

3.3 Routing

The degree of topological and connectivity control normally desired in a supranet makes it essential for all routes to be fixed. Having fixed routes in a supranet means that, for instance, to reach supranet host 'E' from supranet host 'A', supranet packets always visit supranet routers 'B', 'C' and 'D' in this order (see Figure 2).

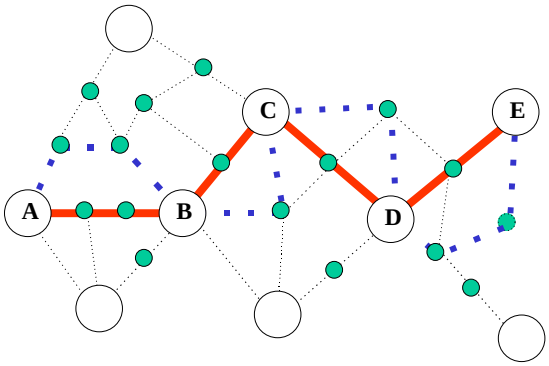

Figure 2: A supranet route connecting hosts 'A' and 'E' through routers 'B', 'C', and 'D'. Large circles represent supranet (and Internet) routers; small circles are Internet routers only. Note that paths AB, BC, CD and DE are supranet (i.e., virtual) links; to each link there correspond multiple physical paths.

When the supranet application is non-real-time and the network protocol in the non-real-time stack is connectionless, the degree of control over topology is weaker: for example, router 'B' may be reached from host 'A' through different physical routes (hence, different physical routers) for different packets. If, however, all physical routers are equipped with SN layer software, they can be "promoted" to supranet routers by the creator; the static routing tables stored in each supranet router are sufficient to fix routes even for non-real-time traffic when this is necessary.

The supranet header, which, in the case in Figure 2, for transmissions from 'A' to 'E' is assembled by 'A' and examined by 'B', 'C', 'D' and 'E', must include the supranet addresses of the source (A) and the destination (E), as well as an "upper protocol" field. The sizes, in bits, of the address fields are dictated by the sum of the maximum number of supranet hosts, the maximum number of supranet routers, and the maximum number of supranet multicast groups. All of these maximum numbers are to be specified at the outset by the supranet's creator.

As shown in the architecture in Figure 1, supranet packets will travel between supranet hosts and supranet routers encapsulated into network protocol packets. Tunnelling is therefore an important technique in the construction of supranets.

3.4 Security

The security requirements of supranets will cover a wide spectrum, given the wide variety of applications of such networks and the dependency of security needs on network's specific applications. We believe that supranet creators should have access to several different security mechanisms satisfying the many possible requirements and should be allowed to select those they need and combine them for their purposes. We also believe that the most widespread security requirements in supranets will be:

a) authentication
b) confidentiality
c) integrity

Authentication mechanisms will prevent supranet members or outsiders from successfully impersonating a supranet member. Confidentiality and integrity must always be enforced with respect to outsiders (except for messages that are totally non-sensitive) and sometimes even with respect to insiders. To simplify our discussion, here we will only mention other types of security requirements, that are expected to be important in some supranets: non-repudiation, resistance to traffic analysis, and anonymity.

The creator will have to specify what types of security are needed in the supranet being defined. Appropriate supranet construction tools will then implement the creator's specifications. Furthermore, access to supranet-owned information must be controlled. The creator is in charge of setting up the access control list for each repository of such information, taking into account the duties and privileges of the various classes (if any) of supranet members. Coupled with good authentication mechanisms, access control lists should make unauthorised access impossible.

3.5 Multicasting

As supranets provide a controlled communications environment, it has to be possible to achieve control also over multicast forms of communications. To this purpose, the supranet creator is allowed to specify a number of multicast groups, and to assign a supranet multicast address to each of them. Since supranet membership is restricted, it is possible in general to build multicast trees that define the paths from each member of the multicast group to the others. When building such paths, appropriate mechanisms exist that can be used to minimise the overall group traffic. Additional security mechanisms, such as private keys for each group, can be used to provide security services at the multicast group level.

3.6 Rules and Sanctions

In the human society, groups and associations are governed by a number of rules. In a similar way, the operation of a supranet is governed by rules dictated by its creator. Rules are concerned with admission, access rights, use of supranet resources, relationships among members, etiquette, and so on. For instance, leaking protected information to the outside will generally be prohibited. Compliance of members with these rules will be monitored, and sanctions for disobeying them will be decided by the creator or by a committee set up by the group for this purpose.

3.7 Supranet Modifications

The creator's initial specifications for all user requirements cannot be assumed to be perfect or immutable. Changes will have to be made to them during the supranet lifetime due to the addition of new members, the departure of old members, the discovery of new needs or of mistakes in the previous specifications, and so on. The creator will have to be provided with tools facilitating all reasonable modifications of the requirements on which the supranet has been based.

4. Conclusions

This paper introduced the notion of "supranet" as a virtual network for groups built on top of a physical integrated-services internetwork, and argued for designing into such internetworks the ability to create supranets as well as tools users will need both for the creation and the management of supranets.

In an integrated-services internetwork whose resources can be partitioned among users or classes of users [3], it would be possible to associate physical resources with a supranet, which would then allocate those resources to its members on a static or dynamic basis. Otherwise, supranet members

who have real-time traffic to transmit will request from the internetwork resources at the time they need them, or reserve them in advance if this is allowed [4] [5].

5. References

[1] M. R. Macedonia, D. P. Brutzman: "*Mbone provides Audio and Video across the Internet*", IEEE Computer, Vol. 27, No. 4, April 1994.

[2] D. Ferrari: "*Should an Integrated Services Internetwork be Connectionless or Connection-Oriented ?*", 6th International NOSSDAV Workshop, pp: 3-4, Zushi, Japan, April 1996.

[3] A. Gupta and D. Ferrari: "*Resource Partitioning for Real-Time Communications*", IEEE/ACM Transactions on Networking Vol. 3, Num. 5, pp: 501-508, October 1995.

[4] D. Ferrari, A. Gupta, G. Ventre: "*Distributed Advance Reservation of Real-Time Connections*", Proceedings of NOSSDAV'95, 5th International Workshop on Network and Operating System Support for Digital Audio and Video, Durham (New Hampshire), April 1995.

[5] L. Delgrossi, S. Schaller, H. Wittig, L. Wolf: "*Issues of Reserving Resources in Advance (ReRA)"*, Proceedings of NOSSDAV'95, 5th International Workshop on Network and Operating System Support for Digital Audio and Video, Durham (New Hampshire), April 1995.

[6] D. Ferrari: "The Tenet Experience and the Design of Protocols for Integrated-Services Internetworks", to appear on the ACM Multimedia Systems Journal.