ecma INTERNATIONAL

Technical Report

ECMA TR/95

1st Edition / June 2008

Next Generation Corporate Networks (NGCN) - General

Rue du Rhône 114 CH-1204 Geneva T: +41 22 849 6000 F: +41 22 849 6001

is the registered trademark of Ecma International

COPYRIGHT PROTECTED DOCUMENT

Technical Report

ECMA TR/95

1st Edition / June 2008

Next Generation Corporate Networks (NGCN) - General

Introduction

This Ecma Technical Report is the first of a series of Ecma publications that explore IP-based enterprise communication involving Corporate telecommunication Networks (CNs) (also known as enterprise networks) and in particular Next Generation Corporate Networks (NGCN). The series particularly focuses on inter-domain communication, including communication between parts of the same enterprise, between enterprises and between enterprises and carriers. This particular Ecma Technical Report provides general information on the subject, defines some architectural concepts, identifies various communication scenarios, and provides a framework in support of other publications that provide greater detail on particular topics.

This Technical Report is based upon the practical experience of Ecma member companies and the results of their active and continuous participation in the work of ISO/IEC JTC1, ITU-T, ETSI, IETF and other international and national standardization bodies. It represents a pragmatic and widely based consensus. In particular, Ecma acknowledges valuable input from experts in ETSI TISPAN.

This Ecma Technical Report has been adopted by the General Assembly of June 2008.

Table of contents

1 Scope

This Ecma Technical Report is part of a series of publications that provides an overview of IP-based enterprise communication involving Corporate telecommunication Networks (CNs) (also known as enterprise networks) and in particular Next Generation Corporate Networks (NGCN). The series particularly focuses on session level communication based on the Session Initiation Protocol (SIP) [6], with an emphasis on inter-domain communication. This includes communication between parts of the same enterprise (on dedicated infrastructures and/or hosted), between enterprises and between enterprises and public networks. Key technical issues are investigated, current standardisation work and gaps in this area are identified and a number of requirements are stated.

This particular Technical Report provides general information on the subject, defines some architectural concepts, identifies various communication scenarios, and provides a framework in support of other publications that provide greater detail on particular topics. At the time of publication of this Technical Report, one further document in the series has been published, on the subject of identification and routing [3].

The scope of this Technical Report is limited to communications with a real-time element, including voice, video, real-time text and instant messaging.

Further details on mobility in an NGCN environment are to be found in ECMA TR/92 [2].

2 References

[1] ECMA-269 Services for Computer Supported Telecommunications Applications (CSTA) Phase III

[2] ECMA TR/92 Corporate Telecommunication Networks - Mobility for Enterprise Communication

[3] ECMA TR/96 Next Generation Corporate Networks (NGCN) - Identification and Routing

[4] ITU-T Recommendation H.248 Gateway control protocol

[5] ITU-T Recommendation H.323 Packet-based multimedia communications systems

[6] IETF RFC 3261 SIP: Session Initiation Protocol

[7] IETF RFC 3550 RTP: A Transport Protocol for Real-Time Applications

[8] IETF RFC 4566 SDP: Session Description Protocol

[9] SIP Forum sf-adopted-twg-IP_PBX_SP_Interop-sibley-sipconnect "IP-PBX / Service Provider Interoperability - SIPConnect 1.0 Technical Recommendation"

[10] ETSI EG 201 017 Corporate Telecommunication Networks (CN); Standardization plan

[11] ETSI TR 180 000 Telecommunications and Internet converged Services and Protocols for Advanced Networking (TISPAN); NGN Terminology

[12] IEEE 802.1x Port Based Network Access Control

3 Definitions

For the purposes of this Technical Report the following definitions apply:

3.1 Corporate telecommunication Network (CN) (ETSI EG 201 017 [10])

Telecommunication network serving a corporation, i.e. a single organization, an extended enterprise, or an industry application group as defined by the International Chamber of Commerce (ICC).

NOTE
Sets of equipment [Customer Premises Equipment (CPE) and/or Customer Premises Networks (CPN)] are

typically located at geographically dispersed locations and are interconnected to provide networking services to a defined group of users. A CN can employ connection-oriented and connectionless technology.

3.2 Domain

Session level capabilities within a single administrative area.

NOTE
A domain may or may not correspond to a DNS domain.

3.3 Enterprise network

A CN comprising session level capabilities and optionally application layer capabilities hosted on one or more infrastructures.

NOTE
Infrastructures can include the enterprise's own infrastructure (dedicated NGCN), the infrastructure of one or more hosting NGNs, the infrastructure of one or more hosting NGCNs or any combination of these.

3.4 Home server

For a given user, as identified by a SIP address of record, the SIP intermediary that contains registrar and proxy functionality in support of that user.

NOTE
It is therefore the SIP intermediary with which the user's UAs register.

3.5 Transport service provider (TSP)

A business or organisation separate from an enterprise that provides services for transporting data based on the use of IP at the network layer, thereby allowing the enterprise to communicate with entities outside the enterprise or with geographically dispersed parts of the enterprise.

NOTE 1
Communication can but need not be via the public Internet.

NOTE 2
A TSP should not intervene above the transport layer. This does not preclude a business or organisation that acts as a TSP also acting as the provider of higher level services, e.g., as an SSP.

3.6 Medium

A given type of payload transported between session users (e.g., audio, video, text), separate from any signalling used for session establishment.

3.7 Next Generation CN (NGCN)

That part of an enterprise network that is not based on public network infrastructure, that is designed to take advantage of emerging IP-based communications solutions and that can have its own applications and service provisioning.

NOTE.
An NGCN can be an entire enterprise network if none of that network is based on public network infrastructure.

3.8 Next Generation Network (NGN)

The definition in [11] applies.

NOTE.
This defines an NGN as follows: "A Next Generation Network is a packet-based network able to provide services including Telecommunication Services and able to make use of multiple broadband, QoS-enabled transport technologies and in which service-related functions are independent from underlying transport-related technologies. It offers unrestricted access by users to different service providers. It supports generalized mobility which will allow consistent and ubiquitous provision of services to users." It also goes on to list some fundamental aspects that characterise NGN.

3.9 Private network traffic

Signalling for session level communications that is handled according to rules specific to an enterprise network.

3.10 Public network traffic

Signalling for session level communications that is handled according to rules for public networks.

3.11 Roaming

The use of session capabilities of a visited domain to allow a user to access session level services at his home domain.

NOTE 1
This usually requires a roaming agreement between the operators of the domains concerned.

NOTE 2
This definition of roaming reflects the concept of roaming as found in mobile telephone networks, for example. It does not encompass certain other common uses of the term, e.g., concerning transport service provision.

3.12 Roaming hub

A network or other entity with which an enterprise domain has a roaming agreement, allowing enterprise users to visit other domains that have a roaming agreement with the roaming hub but not directly with the enterprise domain.

3.13 Session service provider (SSP)

A business or organisation separate from an enterprise that provides communication capabilities at the session layer using SIP and thereby allows the enterprise to communicate using SIP with entities outside the enterprise or with geographically dispersed parts of the enterprise.

3.14 SIP intermediary

Any intermediate entity involved either actively or passively in SIP signalling between two UAs, including but not limited to proxies, Back-to-Back User Agents (B2BUAs), Application Layer Gateways (ALGs) and Session Border Controllers (SBCs).

4 Abbreviations

ALG	Application Layer Gateway
B2BUA	Back-to-Back User Agent
CN	Corporate telecommunication Network
DNS	Domain Name System
IP	Internet Protocol
IPPBX	IP Private Branch eXchange
IPSEC	Internet Protocol Security
ISDN	Integrated Services Digital Network
LAN	Local Area Network
NAT	Network Address Translator
NGCN	Next Generation Corporate Network
NGN	Next Generation Network
PSTN	Public Switched Telephone Network
QoS	Quality of Service

RTP	Real Time Protocol
SBC	Session Border Controller
SDP	Session Description Protocol
SIP	Session Initiation Protocol
SRTP	Secure Real Time Protocol
SSP	Session Service Provider
TLS	Transport Layer Security
TSP	Transport Service Provider
UA	User Agent
UAC	User Agent Client
UAS	User Agent Server
URI	Universal Resource Identifier
WAN	Wide Area Network

5 Background

Many enterprises and other organisations require their own telecommunications capabilities to support their own internal communications as well as supporting communications with the outside world. This avoids incurring unnecessary charges and provides added value in terms of services and features available, integration with other enterprise applications, etc. These capabilities are provided through enterprise telecommunication networks (or corporate telecommunication networks, CN, or simply enterprise networks). Many administrations do not apply the same licensing conditions or regulation to enterprise networks and their internal traffic as they do to public networks and their public network traffic. Many public networks also offer optional services to corporate customers, such as hosted (Centrex) services and the leasing and maintenance of customer premises equipment.

There has been a major evolution in enterprise telecommunications during the last few years. Prior to that, enterprise network were based on 64 kbit/s circuit-switched technology, which had synergy with corresponding technology deployed in public Integrated Services Digital Networks (ISDN) and traditional analogue services. Those enterprise networks primarily delivered a voice or telephony service to their users, although in principle they were capable of other services too, including video and various types of data service. For communication outside the enterprise, enterprise networks were able to interwork with public ISDNs across standardized interfaces (at the T-Reference Point). The T reference point formed the demarcation point between standards and regulations applicable to public networks and standards and regulations applicable to enterprise networks and their PBXs.

With the advent of technologies for transmitting voice and other real-time media over the Internet Protocol (IP) (e.g., based on Real Time Protocol (RTP) [7]) and corresponding new signalling protocols (e.g., H.323 [5], the Session Initiation Protocol (SIP) [6]), there was potential for providing telephony and other real-time person-to-person services in the public Internet. Moreover, such services also became possible in the IP-based "intranets" already deployed in enterprises for data services such as corporate email, file transfer, corporate web services and access to the world wide web. Enterprises saw advantages such as savings on infrastructure costs (e.g., one wire to the desk) and the introduction of innovative services that exploited the convergence of real-time and data communication. The traditional PBX (Private Branch Exchange) was replaced by or evolved to an "IP-PBX" or "soft switch" that supported IP connectivity to the desktop and IP connectivity between nodes. Direct IP-based transmission of multimedia between endpoints meant that switching capabilities were no longer required, except gateways for interworking with "legacy" circuit-switched networks and media servers for conference bridging, announcements, etc. The "IP-PBX" or "soft switch" was just required to handle signalling.

IP-based enterprise networks are continuing to evolve, to support additional services, improved security, improved Quality of Service (QoS), etc. Moreover, SIP has become the dominant signalling protocol. An enterprise network that fully embraces IP technology and uses SIP as the signalling protocol is referred to here as a Next Generation CN (NGCN). An NGCN could still contain some components that are not based on IP (e.g., traditional PISN components) or that use signalling other than SIP (e.g., H.323 [5], H.248 [4]), but it would also include SIP components and be able to interface externally using SIP.

Until recently, NGCNs generally fell back to legacy circuit-switched techniques for standardised communication outside the enterprise, e.g., using public ISDN or circuit-switching over leased lines. Gateways provided the necessary interworking of signalling and media. This was sometimes the case also for communication between different parts of the same enterprise.

We are now witnessing a period when NGCNs are extending IP-based communication externally by interfacing to public IP-based networks. This permits IP-based communication between:

- enterprise users supported by different NGCNs (i.e., different enterprises),
- enterprise users supported by different parts of the same NGCN at geographically dispersed locations (different sites);
- enterprise users supported by NGCN and users supported by hosted enterprise services provided on public network infrastructure;
- enterprise users supported by NGCNs and individual users of public networks (fixed or mobile); and
- enterprise users supported by an NGCN and users of legacy networks via a gateway outside the NGCN.

Enterprise users in this context includes users of terminals based on IP and SIP and also users of legacy terminals connected via gateways and other legacy equipment such as PBXs.

With this the NGCN no longer needs gateways to external legacy networks (except where required by existing investment or economic considerations) and can enjoy the benefits of end-to-end IP-based communication with appropriately equipped communication partners.

The public Internet is one example of a public IP-based network that an NGCN can use for external communications. This can be used for direct connection between two enterprises, between two parts of the same enterprise (e.g., between the main office and a branch or home office), or between an enterprise and a public telecommunications network. In addition, some public network providers are starting to offer public IP-based networks that offer improvements compared with the public Internet (e.g., in terms of QoS, security, mobility, applications, etc.) and are even basing their entire telephony capabilities (including emulation of the PSTN/ISDN) on IP-based technology. These value added public IP-based networks are collectively known as Next Generation Networks (NGN).

At present there are no defined standards for direct connection between enterprises. For connection between an enterprise and a public network, various public network providers are offering their own specifications, based on SIP, and there are some standardisation activities, but these are at risk of being mutually incompatible. This Technical Report analyses enterprise requirements for IP- and SIP-based interworking with other networks, with a view to influencing standardisation and regulation. It also discusses communication between different parts of an enterprise, including parts supported by enterprise infrastructure and parts supported through hosting on other network infrastructures, including that of an NGN.

6 General concepts

6.1 Basic communication architecture

Communication networks in general and NGCNs in particular can be viewed as providing capabilities at three levels, as shown in Figure 1:

- transport;

- session;
- application.

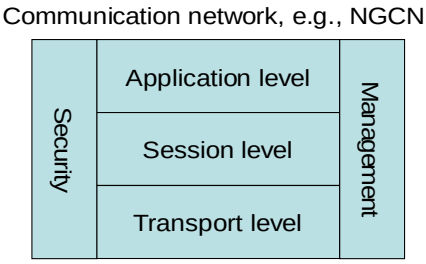

Figure 1 — Basic communication architecture

Capabilities at the transport level provide basic IP connectivity between physical items of equipment, such as servers, PCs, phones, PDAs, etc., including connectivity to the outside world such as the public Internet and NGNs. This can be used to transport signalling, media and other data. IP connectivity can be based on IP version 4 or version 6 and can operate over different infrastructures (e.g., fixed or wireless, LAN or WAN) with different security mechanisms (e.g., IEEE 802.1x [12], WiFi security). Capabilities present at the transport level include but are not limited to routing, switching, firewall, network address translation (NAT), quality of service (QoS) provision, security, etc. Connectivity may be limited to enterprise employees or may be extended to guests.

The session level provides capabilities in support of user-to-user communication, generally with a real-time element. Communications can use a variety of media, including audio, video, real-time text, messaging, fax, etc., or a combination of these. Within an NGCN (and likewise within an NGN) SIP signalling is assumed for establishment of communications. Users are not necessarily human users - in some cases automata (applications) can take the role of users.

For the purposes of this document, the application level comprises applications that have some relevance to communication, as provided by the session level. This includes:

- applications that enhance communication capabilities in some way, e.g., advanced conferencing, presence;
- applications that use communication capabilities, e.g., a business process application that uses communication capabilities for communicating with customers, suppliers, partners, etc.

This document focuses on the session level, but also includes some discussion on the application level. The transport level is mentioned only to the extent of how it impacts the session level, e.g., providing connectivity between session level entities, the impact of NATs and firewalls.

Each level has its own security and management considerations and capabilities. In addition, security and management have to be co-ordinated across all three levels, NGCN-wide.

The broad architectural framework described above can be realised in a number of ways. One example utilises the IP Multimedia Subsystem (IMS) specified by 3GPP, in which case the IMS layer corresponds to the session level described above. The IMS layer runs on top of a transport layer and provides support to a service/applications layer. The IMS layer itself comprises a number of functional components and inter-relationships. This Technical Report makes no assumptions regarding the use of IMS or any other architecture (beyond the framework described above and the session level architecture described in 6.2) in NGCNs.

The three level architecture is applicable also to communication between networks. Two networks can each provide functionality at each of the three levels, and interworking can take place potentially at all three levels, as shown in Figure 2. Interworking does not need to occur right up to the application layer. For example, if applications are local to the two networks, interworking could be just at the transport and session levels. If there is no real-time communication session involved

(e.g., data communication such as email or web, which is outside the scope of this Technical Report), interworking would be only at the transport level.

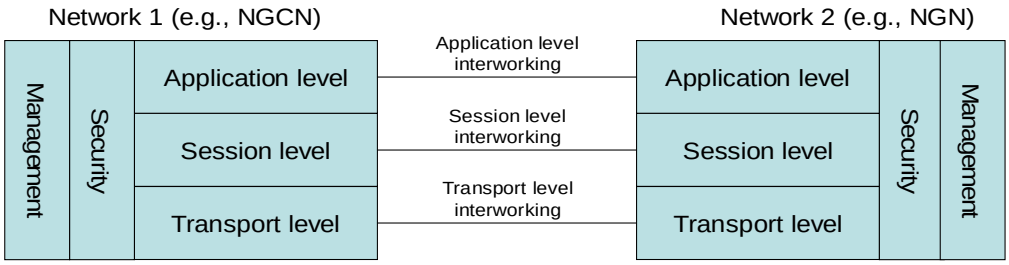

Figure 2 — Overall architecture for inter-network operation

Similar principles can be extended to multiple networks in series, where intermediate networks might not be involved at the higher levels. Figure 3 shows an example where networks 1 and 4 communicate up to the application level, with network 2 involved only at the transport level (e.g., a transport service provider, TSP) and network 3 involved only up to the session level (management and security omitted for clarity).

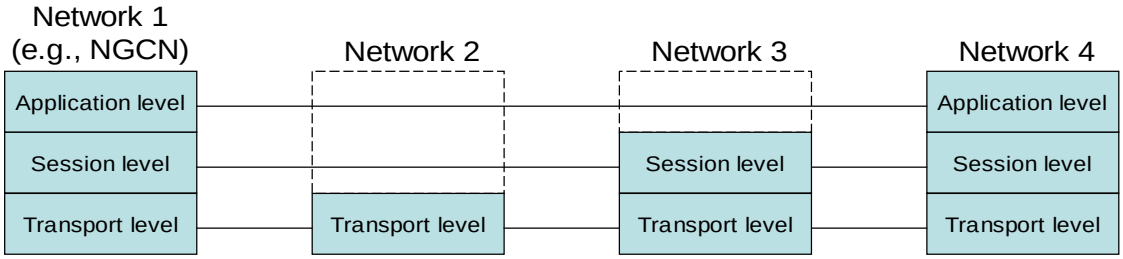

Figure 3 — Example of operation through multiple networks (management and security omitted for clarity)

6.2 Session level architecture

During a session, two users (more in conferencing arrangements) exchange media information. To do this each user has an endpoint device, which in the case of a human user can be a dedicated phone (fixed or wireless) or a PC or PDA running a soft client or some higher level application. Other types of endpoint include application servers (e.g., media applications, conferencing applications, presence applications) and gateways to non-IP-based or non-SIP-based networks or equipment, etc. Establishment, clear down and maintenance of a session are achieved using signalling, the protocol being SIP.

Signalling and media are transported separately between endpoints. Signalling typically passes through one or more intermediaries, as shown in Figure 4, whereas media is often transported directly between endpoints (intermediaries such as NATs and firewalls can be involved up to the transport layer). However, see also 6.2.2 for further discussion of the media path.

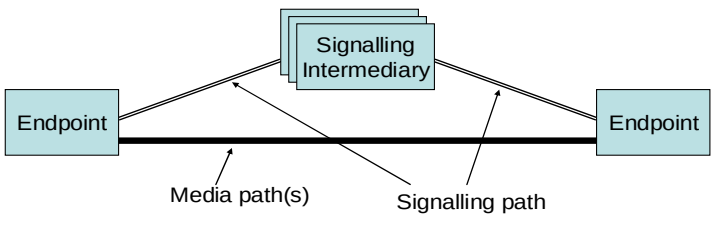

Figure 4 — Session architecture

6.2.1 Signalling using SIP

SIP messages are used to exchange session descriptions using the embedded Session Description Protocol (SDP) [8]. Session descriptions indicate factors such as the media to be used, codecs, packet rates, security contexts and IP addresses and port numbers for media reception.

Using SIP terminology, each endpoint contains a SIP user agent (UA). In principle, SIP can operate directly between two UAs. However, the SIP standards do define other SIP entities that can assist. In particular the main SIP standard [6] defines the concept of a location service at which UAs can register to assist in locating UAs representing a particular user. Related to this are the concepts of a registrar (a SIP entity that registers UAs at the location service) and a proxy (a SIP entity that queries the location service in order to assist in the forwarding of requests to appropriate UAs).

NOTE
Also there is a redirect, which is similar to a proxy but redirects requests (by instructing an upstream entity to remake the request to a different destination) rather than forwarding requests. For the purposes of this Technical Report a redirect is treated as a proxy unless otherwise stated.

The handling of a SIP request from one UA (the UA client, UAC) to another UA (the UA server, UAS) typically passes through a proxy in the destination domain (as determined by the domain part of a SIP URI), which queries the location service in order to locate UAs that might be able to handle the request. Also the UAC often sends requests to a local proxy in the originating domain (known as an outbound proxy) in order to reduce the routing burden on the UAC. This gives rise to the typical SIP trapezoid involving two UAs and two proxies (see Figure 5).

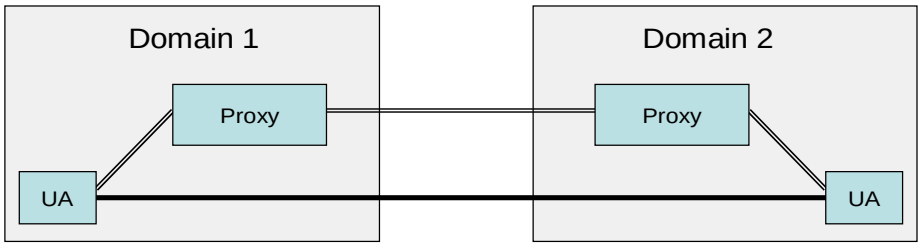

Figure 5 — Typical SIP trapezoid

Sometimes a SIP request can result in the formation of an association between two UAs, in the context of which further SIP requests in either direction can be sent. This is known as a dialog. The most important dialog in SIP is that initiated by an INVITE request, because such a dialog has an associated session, comprising a collection of media between the UAs. An INVITE-initiated dialog and its associated session can be regarded as a call. Proxies involved in routing the dialog-initiating request may or may not remain involved in the dialog for routing subsequent requests.

Typically a proxy is collocated with a location service and a registrar, so the whole entity tends to be known as a proxy. The behaviour of a proxy is standardised in [6] and other SIP RFCs and its behaviour is quite tightly constrained by the standards in order not to jeopardise the end-to-end character that is fundamental to much of SIP.

Although not standardised in [6], the market has found the need for an entity known as a Back-to-Back User Agent (B2BUA). In some respects a B2BUA acts like a proxy, being an intermediate entity on a signalling path between two UAs, but in other respects a B2BUA acts as two UAs back-to-back. The behaviour of a B2BUA is not standardised, and because it is not constrained by standards a B2BUA can behave more flexibly than a proxy. This allows a B2BUA to fulfil certain roles that a proxy cannot (e.g., by offering transcoding capabilities for media). On the other hand by not behaving as a proxy a badly designed B2BUA can behave in ways that are inconsistent with the expectations of UAs and can cause problems.

Other forms of entity that can intervene either actively or passively in a SIP signalling path include Application Layer Gateways (ALGs) and Session Border Controllers (SBCs). These have knowledge of the SIP protocol and can monitor or modify it for certain purposes, e.g., control of NATs and firewalls, security, privacy, etc. They often behave as B2BUAs in some respects.

In the rest of this document, the term SIP intermediary is used to refer to any intermediate entity involved either actively or passively in SIP signalling between two UAs, including but not limited to home servers, other proxies or B2BUAs, ALGs and SBCs. In this document, a SIP intermediary that contains registrar and proxy functionality (whether or not it is a pure proxy/registrar or a B2BUA) is called a home server, since it serves as the "home" for those SIP addresses of record for which it provides a registrar function. For mobile users, even when roaming to distant access points, their UAs still need to register at the home server.

As stated above, SIP can in principle operate directly between UAs, and indeed there is a lot of interest at present in peer-to-peer (P2P) SIP where SIP intermediaries are eliminated. Whether and how soon this ideal will be realised is open to speculation (perhaps initially within very small enterprises), but the trend is expected to be away from the circuit switching concept of signalling passing through an arbitrary number of intermediaries (e.g., associated with switches) towards a lightweight solution with a minimal number of SIP intermediaries.

In many enterprises a single SIP intermediary (home server) can suffice for communication within the enterprise network (or within a domain of the enterprise network, see 6.3), with additional intermediaries involved when communicating outside the enterprise network or domain. Even multi-national enterprises will not necessarily deploy a SIP intermediary in every country, and UAs in one country can be allowed to register at a home server in another country. In many enterprises these intermediaries are in the form of B2BUAs, rather than proxies, in that they do not conform to all the rules for proxies.

Additional SIP intermediaries are often deployed at the border of a network or domain, normally described as session border controllers (SBCs). These provide functions such as topology hiding, QoS, firewall traversal and NAT traversal, and generally the media flows through these too.

6.2.2 Media path

In principle, media is transported directly between the two UAs without intervention above the transport layer. Intervention can occur, for example, when there is a need to perform some function on the media such as transcoding, translating, recording or bridging. Equipment through which media is routed for such purposes may or may not be equipment involved in the signalling path.

Even when there is no intervention above the transport layer, although the media path in principle is independent of the signalling path (and can differ between media) there can be constraints governing the routing of media. When SBCs are involved on the signalling path and perform functions on media such as NAT or firewall functions, charging or QoS measures, media will also be required to pass through those SBCs.

An enterprise operator might have a policy on the routing of media, e.g., to minimise charging, to gain appropriate QoS, or to ensure proper security. One technique for achieving this is to force signalling to follow a particular path and media to follow that path. Another way is to set aside certain infrastructure resources for media.

6.2.3 Example

Figure 6 shows an example of two domains communicating at the session level, each with an SBC in addition to a normal SIP intermediary (proxy or B2BUA). It is assumed that media as well as signalling pass through the SBCs.

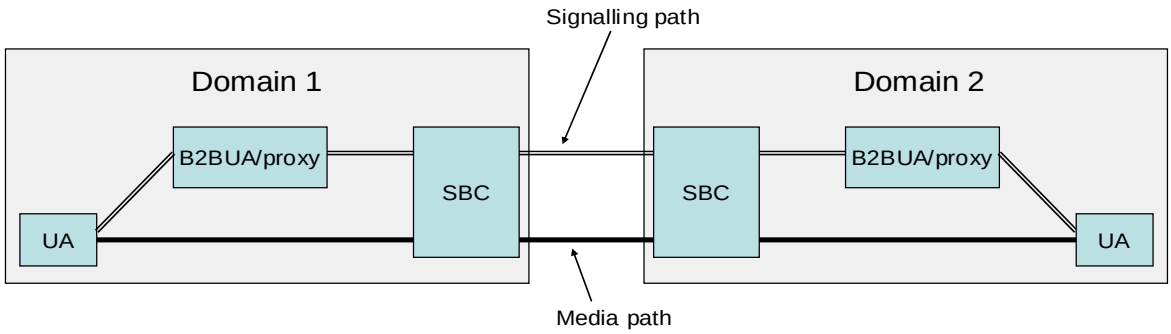

Figure 6 — Example of session involving two domains with SBCs

6.3 Domains

When discussing session capabilities in this document, the term "domain" is used to refer to session capabilities within a single administrative area.

The term "enterprise network" is used to refer to session capabilities provision for the entire enterprise, and this may comprise a single domain or multiple domains. In the latter case, the individual domains can also be regarded as sub-domains of the domain representing the entire enterprise network.

The term "public network" is used to refer to session capabilities provision for the general public. Generally it is sufficient to regard a public network as a single domain from the perspective of an enterprise network.

Figure 7 shows an example of an enterprise domain (comprising two sub-domains) and a public network domain.

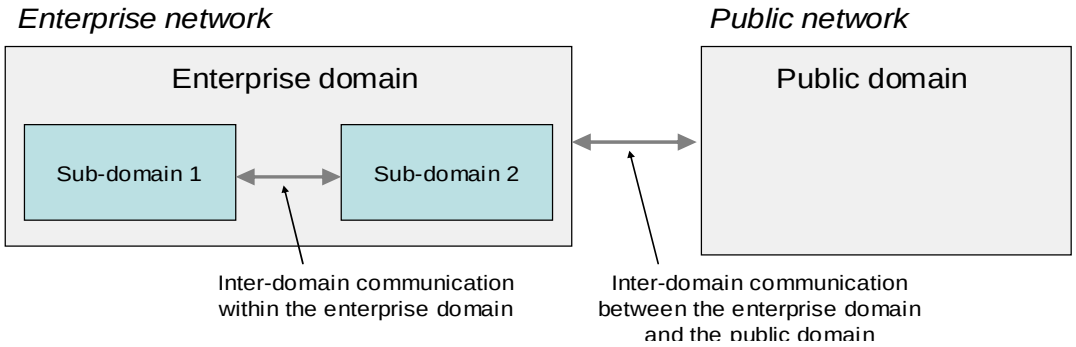

Figure 7 — Example of domains and inter-domain communications

Inter-domain communications therefore refers either to communications involving two or more domains within an enterprise network or to communications involving an enterprise network and one or more other networks (public networks and/or other enterprise networks).

Frequently a domain as described above will coincide with a DNS domain, e.g., example.com. However, this is not necessarily so. For example, all administrative domains within a single enterprise network could share the domain name example.com.

The domain name for an enterprise will normally be used within identifiers for enterprise users of session level services, in order to ensure that such identifiers are fully qualified (globally unique). This is particularly important for identifiers that are not based on E.164 numbers, i.e., identifiers based on numbers in accordance with a private numbering plan (PNP) or non-numeric identifiers (e.g., names). In the case of a SIP URI, the domain name forms the domain part of the URI. This is a complex area and will be the subject of a separate publication in this series.

Typically an enterprise network is fully interconnected, although in some cases there might be islands that are not directly connected and have to communicate with each other via other networks (e.g., via an NGN). Each of these islands is then treated as a separate domain. If the interconnecting network is involved at the session level, that too is considered as a domain.

For communications between domains, signalling passes through at least one SIP intermediary in the originating domain, in the terminating domain, and in any intermediate (transit) domains. In addition, SBCs are often involved on either side of domain boundaries.

6.4 Mobility

An NGCN intrinsically supports a basic mobility capability by allowing users to register with their home server from anywhere within the domain and obtain the same services. The same applies to enterprise users hosted on public network infrastructure. This caters for situations where the user logs onto different fixed terminals in different locations as well as to the case where a terminal moves with the user.

The ability to register and obtain services from different locations can be extended to the case where a user connects to a different transport infrastructure, with transport connectivity to the home domain's transport infrastructure by some means, e.g., VPN. This has particular implications for geographic location provision and emergency calls.

In some cases, performance considerations might limit the effectiveness of real-time communications when visiting a remote location, e.g., when using a cellular access network. In such cases it may be beneficial to use session level capabilities of a visited domain, more closely associated with the visited transport infrastructure. This might also apply if, for any reason, VPN is not available as a means of traversing NATs and firewalls. Registration would still be with the NGCN user's home server, but would be via one or more SIP intermediaries (proxies) in the visited domain and perhaps also in intermediate domains. Likewise, signalling for outgoing and incoming communications would pass between the endpoint and the home server via one or more SIP intermediaries in the visited domain and any intermediate domains. This is based on the same principle as that used to support mobility in IMS.

The ability to obtain services when visiting another domain with the support of session capabilities of the visited domain is known as roaming. Roaming generally requires a roaming agreement between the operators of the domains concerned.

When a wireless terminal is used for mobility, including roaming, some form of session continuity may also be available, thereby giving a more complete mobility experience. Session continuity is outside the scope of this Technical Report but will be the covered by a separate document in this series on the subject of mobility.

6.4.1 Roaming of enterprise users outside their home domain

Figure 8 shows an example of an enterprise user A whose home domain is domain 1 roaming into visited domain 2 and from there being involved in a call with user B in domain 3.

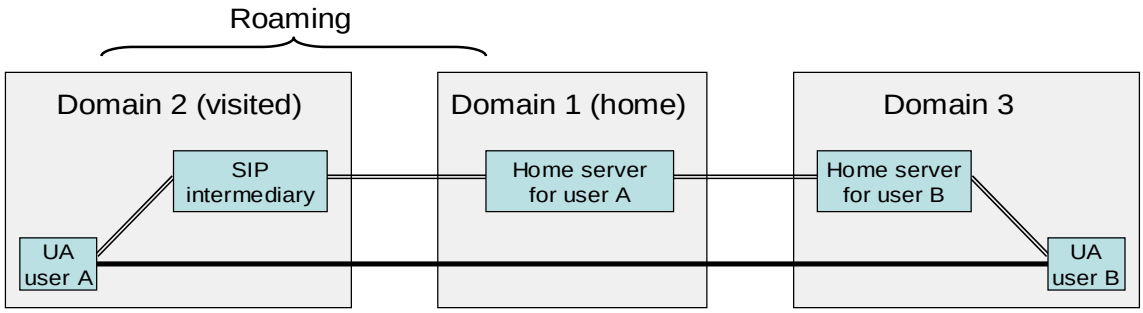

Figure 8 — Roaming example

To permit roaming, there will in general need to be a roaming agreement between the home and visited domain. However, having a separate roaming agreement with every potential visited domain might be impracticable, and therefore reliance on indirect roaming agreements is likely.

For example, an NGCN might have a roaming agreement with a particular NGN, which in turn has roaming agreements with many other NGNs and mobile networks in different countries and somehow advertises reachability. The practice of one network having a roaming agreement with an entity that in turn has agreements with other networks is common practice in the cellular world. The entity in the middle is often referred to as the roaming hub. Applying this to the case of an enterprise user roaming to other domains, the domain with which the enterprise has a roaming agreement acts as the roaming hub. Figure 9 illustrates the use of a roaming hub and assumes that signalling between the visited domain and the home domain passes through the hub (although other arrangements might be possible).

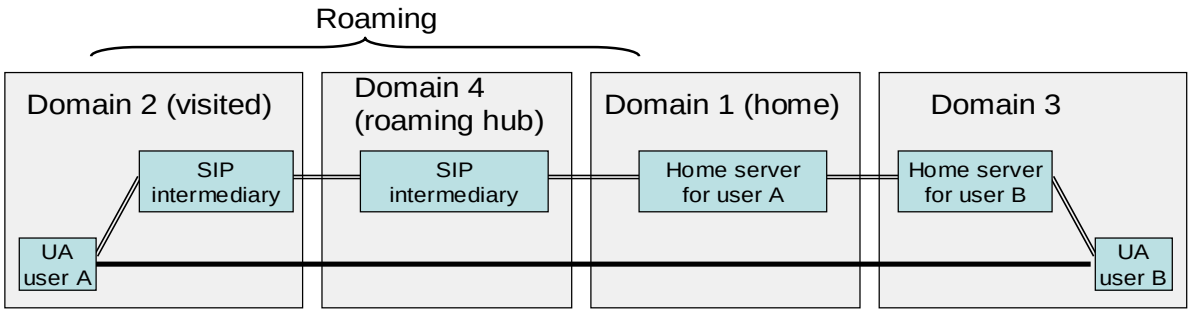

Figure 9 — Indirect roaming example

This concept is equally applicable to NGCN users and enterprise users hosted by a public network infrastructure (see 6.5). Various scenarios involving roaming are shown in 7.5.

6.4.2 Accommodating guest users on an NGCN

Conversely, an NGCN can support guest users, who do not have regular accounts with the NGCN either at the transport level or at the session level. At the transport level, such users are normally accommodated in a way that clearly separates them from regular enterprise users and servers, e.g., using a virtual LAN or WAN that allows them to access the public Internet. This also allows the guest users to connect to their home infrastructure using VPN. A guest user can conduct session services with his/her normal domain using this transport level capability, and therefore the NGCN would not be involved at the session level. There may be performance considerations, e.g., a user roaming visiting an NGCN and accessing conference facilities also within that NGCN may experience performance degradation if media are routed via the user's home domain rather than locally within the visited NGCN. However, this should be no worse than when visiting one domain and accessing conference facilities in a third domain. At present there does not seem to be any motivation for allowing a guest user to access his/her home server via NGCN SIP servers, i.e., for no need for roaming support.

NOTE
This does not preclude an NGCN home server providing a visitor with a temporary account, allowing him/her to use NGCN session level services.

6.5 The hosting concept

Within a single site or campus, an enterprise will have its own infrastructure, built on technologies such as Ethernet, Wireless LAN, etc. and providing at least a transport level capability. This can be extended between sites using techniques such as leased line or virtual private network (VPN), the latter going across public network infrastructure, e.g., the public Internet. Methods for achieving an infrastructure supporting transport level capabilities are outside the scope of this Technical Report, which focuses on communications at the session level and above.

There are several basic options for providing session capabilities for an enterprise. One possibility is to provide dedicated session capabilities on enterprise premises using the enterprise infrastructure, the capabilities being managed by the enterprise (or possibly by a third party). This is analogous to the traditional PBX.

However, there is an increasing market for hosted session services, whereby session capabilities for an enterprise are hosted on a third party infrastructure. The third party infrastructure could be

that of a public network, for example an NGN. It could also be an enterprise infrastructure that has spare capacity or exists primarily for the purpose of providing hosted capabilities. Regardless of the network that hosts session capabilities for an enterprise, a terminal accessing session capabilities will typically be on the enterprise's own infrastructure (directly or via VPN) or with roaming guest status on a third party infrastructure (e.g., a home network, a hotel network or a wireless hotspot).

A hosting infrastructure will typically host session capabilities for multiple different enterprise networks. Furthermore, different parts of the same enterprise can be supported by different hosting infrastructures. There may even be a mixture of hosting and dedicated infrastructures supporting session capabilities for a single enterprise. Each of these parts of the enterprise network are regarded as separate domains. Of course, public network infrastructure will also support public session capabilities, i.e., provision of services to the general public. The general concept of a hosting infrastructure supporting session capabilities is illustrated in Figure 10.

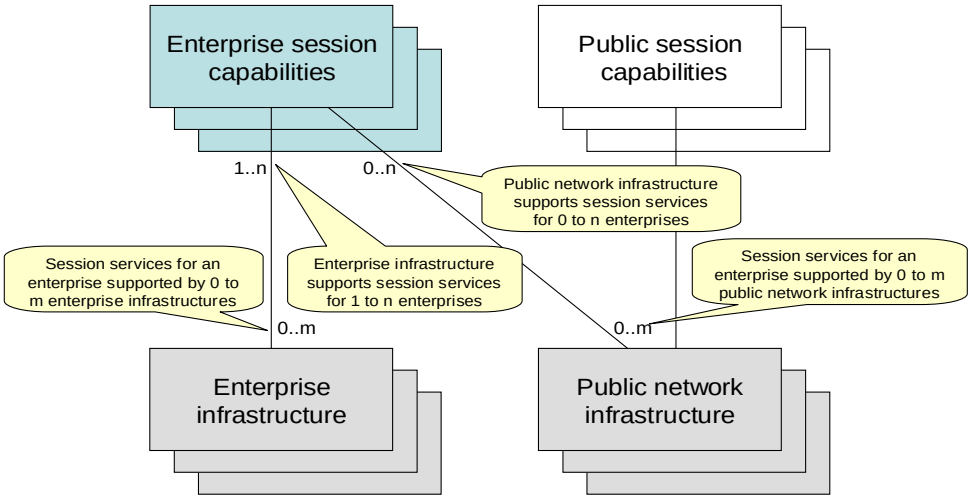

Figure 10 — Session capabilities supported by different infrastructures

NOTE
Although a fourth possibility can be identified (public session capabilities supported on enterprise infrastructure), no use case has been identified. One possibility might be to accommodate roaming users with guest status from outside the enterprise, but typically transport level capabilities are sufficient to support such users.

An NGCN comprises enterprise infrastructure with session capabilities support for one or more enterprises. An NGN comprises public network infrastructure with session capabilities support for public communications as well as for zero or more enterprises. Session capabilities for an enterprise could be supported by one or more NGCNs and/or one or more NGNs. Thus different users use session capabilities from different infrastructures, often but not necessarily on a geographic basis.

Similar hosting options exist for applications. For example an NGN could host enterprise session capabilities and a number of enterprise applications.

Examples of specific arrangements derived from the general hosting concept described above are described below.

6.5.1 Dedicated NGCN

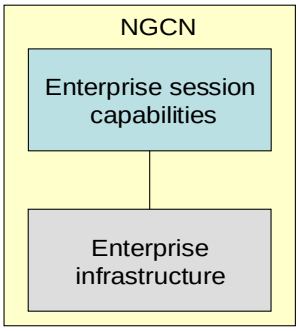

Figure 11 — Dedicated NGCN

In this example, session capabilities for an enterprise are provided on top of the enterprise's own infrastructure. This is a simple NGCN and is analogous to the traditional PBX. The NGCN can be viewed as an enterprise network comprising a single domain (possibly divided into sub-domains for administrative reasons).

6.5.2 Enterprise hosted by a single public network infrastructure

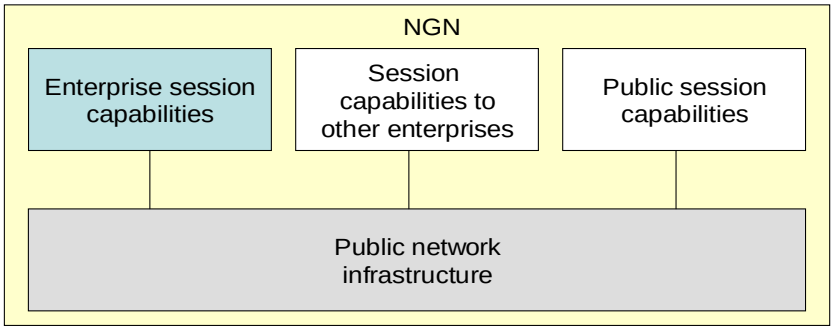

Figure 12 — Enterprise hosted by a single public network

In this example, session capabilities for an enterprise are provided by an NGN, which also provides session capabilities to other enterprises and public session capabilities. For a given enterprise, the enterprise session capabilities can be regarded as an enterprise network comprising a single domain (possibly divided into sub-domains for administrative reasons).

6.5.3 Enterprise hosted by multiple public network infrastructures

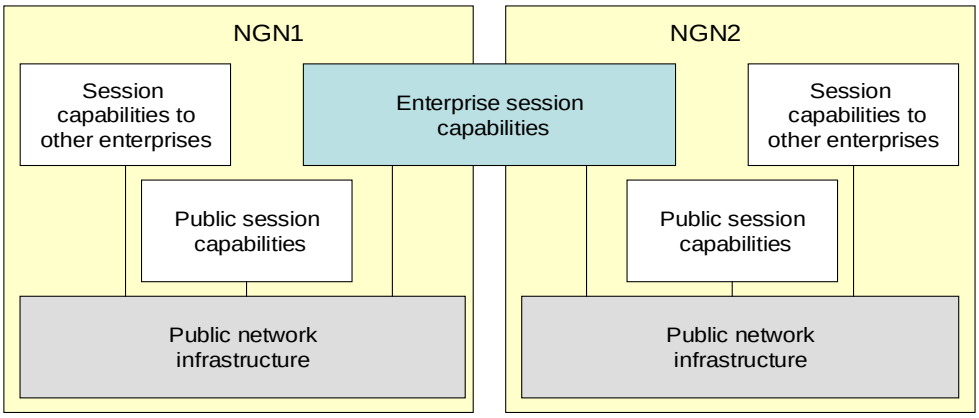

Figure 13 — Enterprise hosted by two public networks

In this example, session capabilities for an enterprise are provided by two or more NGNs, e.g., in different countries, serving users based in the respective countries. The enterprise session capabilities can be regarded as an enterprise network comprising two domains, one per NGN (each possibly divided into sub-domains for administrative reasons).The two NGNs need to communicate in a similar way to two nodes of an NGCN and support features specific to the enterprise network concerned, e.g., private numbering plan.

6.5.4 Enterprise hosted by a combination of enterprise infrastructure and a public network infrastructure

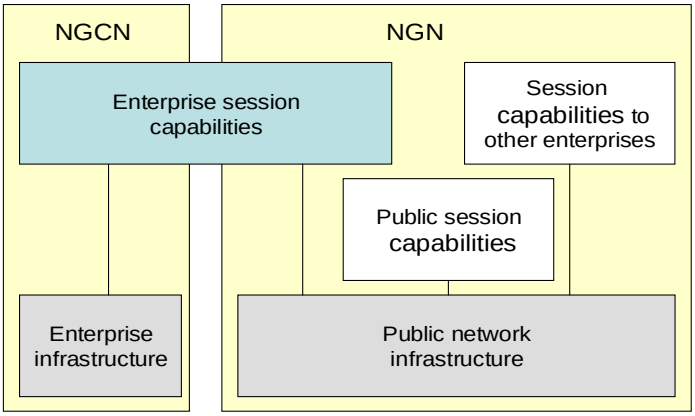

Figure 14 — Enterprise hosted by own and public network infrastructure

In this example, session capabilities for an enterprise are provided by an NGCN (on the enterprise's own infrastructure) and an NGN. The NGCN might serve users based at the main office and the NGN might serve users based at branch offices or at home. The enterprise session capabilities can be regarded as an enterprise network comprising two domains, one being the NGCN, the other in the NGN (each possibly divided into sub-domains for administrative reasons). The NGCN and the NGN need to communicate in a similar way to two nodes of an NGCN and support features specific to the enterprise network concerned, e.g., private numbering plan.

To a first approximation the services received by enterprise users hosted by the NGN are similar to services received by enterprise users hosted by the NGCN but roaming (on a long term basis) into the NGN (see 6.4.1). There could, of course, be detailed differences in offerings, and commercially the two arrangements might be quite different and might attract different charges. Also the mechanisms differ, in that signalling (and perhaps media too) for calls involving roaming NGCN-hosted users always needs to be routed via the home server in

the NGCN even when communicating with another party in the public network or a party supported by public network infrastructure. With NGN-hosted enterprise users, this longer path is avoided. On the other hand, the use of roaming may be easier to set up, since it reduces the amount of configuration needed at the NGN (e.g., it avoids the need to support the enterprise private numbering plan).

6.6 Private network traffic and public network traffic

Within an NGN infrastructure that hosts both enterprise networks and a public network, signalling for session level communications can be handled according to rules specific to an enterprise network or rules applicable to the public network. Rules for an enterprise network can differ in a number of respects, e.g.:

- resources available and their priorities;
- routing based on identifiers specific to the enterprise (e.g., private numbering plan);
- charging;
- the applicability of services such as call recording;
- regulations applicable to public networks may not apply.

Within such an NGN the term private network traffic applies when signalling is handled according to rules specific to an enterprise network and the term public network traffic applies when signalling is handled according to rules for the public network.

A call between two users of the enterprise network will generally be regarded as private network traffic end-to-end, including any part that is routed through NGN infrastructure, e.g., between two NGCN domains, between two hosted domains or between an NGCN domain and a hosted domain.

A call from a user of an enterprise network to a user of a public network is said to break out of the enterprise network. Between the calling user and the break-out point the call is considered private network traffic, and beyond the break-out point it is considered public network traffic.

For a call originating in an NGCN and going via NGN to a public network user, the break-out point can be located either in the NGCN or in the NGN. If located in the NGCN, the NGCN determines that the call has to break out, performs any necessary actions such as changing identifiers, and indicates to the NGN that the call is public network traffic. If the break-out point is located in the NGN, the NGCN indicates to the NGN that the call is private network traffic and the NGN determines that it needs to break-out.

A call from a user of a public network to a user of an enterprise network is said to break into the enterprise network. Between the calling user and the break-in point the call is considered public network traffic, and beyond the break-in point it is considered private network traffic.

For a call terminating in an NGCN and coming via NGN from a public network, the break-in point can be located either in the NGN or in the NGCN.

NOTE
Although when media happens to follow the same path as signalling (in terms of the domains it passes through) the media too might be considered to comprise segments that are private network traffic and segments that are public network traffic, this does not work well when media is routed independently of signalling. Therefore the concept is not applicable to media.

Further study is required to determine whether this concept will be sufficient to address all regulatory considerations related to enterprise networks, e.g., whether break-out within an NGN would have impact on the NGCN, such as requiring the NGCN to perform lawful interception.

6.7 Other technical considerations

In addition to the architectural and other concepts described above, there are many other technical considerations associated with the scenarios identified in Clause 7, including but not limited to:

- identification and routing (including identification presentation and restriction);

- NAT and firewall traversal;
- security;
- regulatory considerations (including emergency calls, lawful interception, prioritised communications such as for disaster relief);
- quality issues;
- location-based services, etc.

Discussion of these technical considerations is outside the scope of this Technical Report. Some of these topics will be covered by other Ecma publications in this series.

7 Scenarios for session-based communications

This Clause introduces various intra-domain and inter-domain scenarios for communication. In each case the architectural impact is considered, Other technical considerations are identified, but discussion is generally outside the scope of this Technical Report and will be addressed by other documents in this series.

7.1 Session-based intra-domain communications

This sub-clause discusses session-based communications within a single-domain enterprise network or within a single domain of a multi-domain enterprise network. The focus is on dedicated NGCNs (based on an enterprise's own infrastructure) and on domains supported by hosting NGCNs. Similar principles apply to domains hosted by NGNs, except that they will be based largely on the same NGN technology that is used to support public networks.

Intra-domain communications are considered private network traffic.

Geographically dispersed parts of a single enterprise domain will often share a common enterprise infrastructure based on technologies such as leased line and VPN. At the session level, the transport level capabilities can be used to link SIP intermediaries in different locations and to allow SIP endpoints to register with and otherwise communicate with their home servers. VPN connections can be established on a semi-permanent basis (e.g., between a branch office and a main office) or can be established on demand (e.g., from an employee's PC at home, in hotels, when travelling, etc.). Locations connected in this way form an essentially homogenous infrastructure or domain, when viewed from the session level or above.

However, as discussed in 6.4.1, remote access at the transport level by a roaming user will not always deliver the desired performance, and consequently session level services of a visited domain might need to be used. Thus communication is no longer intra-domain.

When a domain is hosted, UAs may be supported on a separate infrastructure from that of the hosting network, e.g., the enterprise's own infrastructure or a home network infrastructure, without necessarily having a VPN connection into the hosting infrastructure. For the purposes of this Technical Report, such UAs are still considered part of the hosted domain, provided they signal directly with a SIP intermediary in the hosting network.

Within a domain, there is no fundamental reason for more than one home server, even when an enterprise spans different countries and continents. However, considerations such as the following will lead to more than one home server in many cases:

- resilience;
- server scalability;
- regulation (e.g., if a country forbids international VoIP traffic without using a public network);
- bandwidth utilisation on long distance links;
- propagation time,

- the presence of NATs and firewalls inside the enterprise network (as opposed to at borders) (unusual);
- country- or region-independence (one country not impacted if the network in another country fails or links fail);
- cost considerations (e.g., if interconnection via a session service provider is cheaper than leased lines);
- distributed management responsibility;
- company mergers and acquisitions, etc.

Where a cluster of home servers is provided for resilience, UAs will perhaps register with multiple home servers, or register with one and be prepared to register with another if that one fails.

In other cases, users will be assigned to particular home servers (normally on the basis of geographic location) and UAs will always register with the home server concerned. In such cases the home servers can be considered to be in different (sub-)domains and scenarios in 7.2 apply for communication between users served by different home servers.

Typically registration can occur from anywhere within the domain (if the user is mobile), except where NATs or firewalls or non-technical reasons prevent this.

In addition to home servers, there can also be other SIP intermediaries involved in intra-domain communications, e.g., edge proxies (acting as front end concentrators into a home server), SBCs (allowing UAs on enterprise infrastructure to access the hosting infrastructure).

In the simplest case, session signalling will pass through just a single SIP intermediary, the home server serving both the calling and called users, as shown in Figure 15. In other cases it will pass through other types of SIP intermediary, e.g., SBCs between UAs and the home server (see Figure 16).

Note that the use of one or a small number of SIP intermediaries in a geographically dispersed domain should not impact the routing of media traffic, which should still flow directly between UAs. Different considerations apply to entities such as media servers and gateways, which might need to be provided in each country or region to prevent excessive bandwidth demands on long distance links.

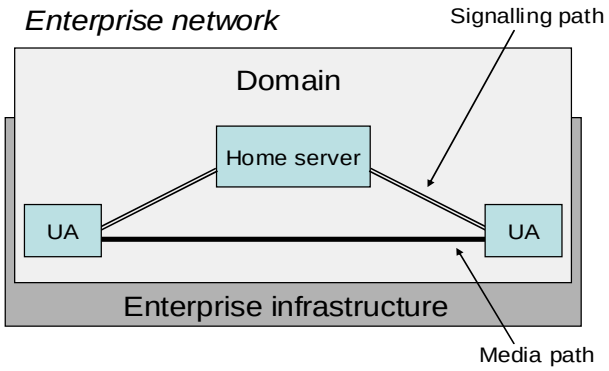

Figure 15 — Example of an intra-domain session through a single home server

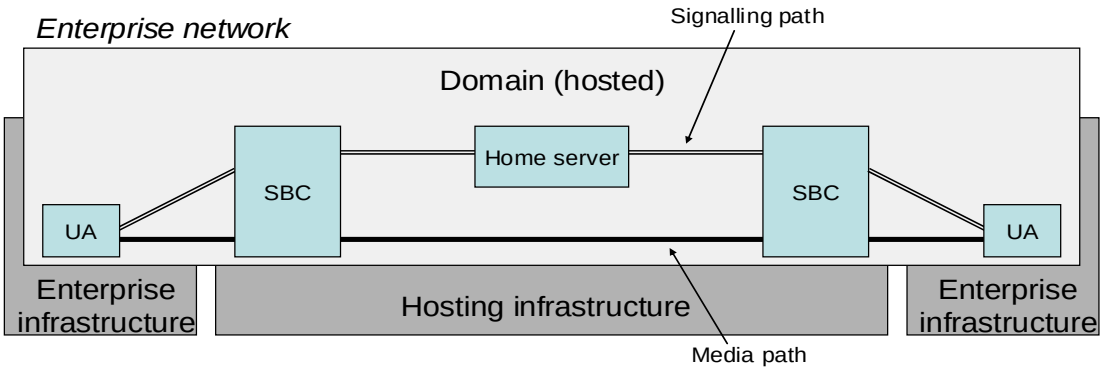

Figure 16 — Example of session within a hosted domain employing SBCs

7.2 Session-based inter-domain communications within a single enterprise network

This sub-clause discusses session-based communications between domains of a single enterprise's network.

A domain can be a dedicated NGCN or can be hosted on the infrastructure of another NGCN or an NGN. A domain could also be a part (sub-domain) of such a domain. An enterprise network can comprise a mixture of such domains.

There may be cases where an NGCN is spread over separate islands of enterprise infrastructure that, for NGCN purposes, are not interconnected. This includes cases where:

- there is no interconnection at all;
- interconnection exists but is reserved for other types of traffic;
- interconnection exists but has limited capacity;
- interconnection exists but has temporarily failed.

Each island can be treated as a separate domain. Communication between islands will need to be via public networks. Possibilities include:

- use of a TSP to carry signalling and media between NGCN domains;
- use of a hosting network (e.g., an NGN) to carry signalling and media as private network traffic between NGCN domains.
- use of a public network (e.g., an NGN) to carry signalling and media as public network traffic between NGCN domains.

This last case is considered to be communication outside the enterprise network and is covered in 7.4.

Inter-domain communications within a single enterprise network are considered private network traffic.

Generally signalling for inter-domain traffic will pass through a home server in the caller's domain and a home server in the called user's domain. In addition, it can pass through other SIP intermediaries, including those in any transit domains that might be involved. Some examples are given below.

7.2.1 Communication between domains supported by the same infrastructure

This is illustrated in Figure 17 for the case of two domains supported by a single enterprise infrastructure, i.e., two domains of an NGCN.

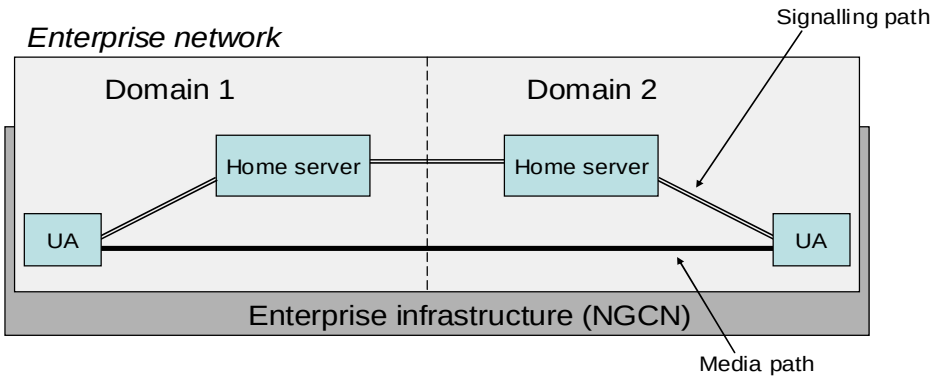

Figure 17 — Example of communication between domains within a single NGCN

7.2.2 Communications between domains of an NGCN via a TSP

This is illustrated in Figure 18.

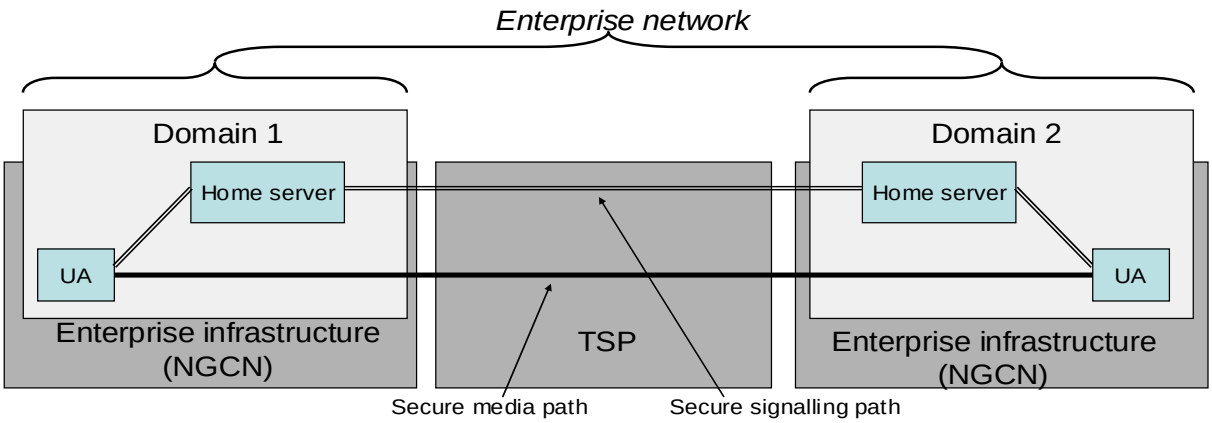

Figure 18 — Example of communication between domains of an NGCN via a TSP

This is a case of direct peering between parts of an NGCN, and is comparable with direct peering between enterprises as described in 7.3.2. It is essential that both signalling and media be secured (e.g., TLS or IPSEC for SIP, SRTP for real-time media). SIP intermediaries in the various parts (or at least those on the borders, e.g., SBCs) need to have public IP addresses to be reachable in this way. NAT and firewall traversal techniques need to be employed. SIP intermediaries can be configured only to accept traffic from known parts of the same enterprise, thereby avoiding handling unwanted traffic.

7.2.3 Communications between domains of an NGCN using private network traffic through a hosting NGN

This is illustrated in Figure 19.

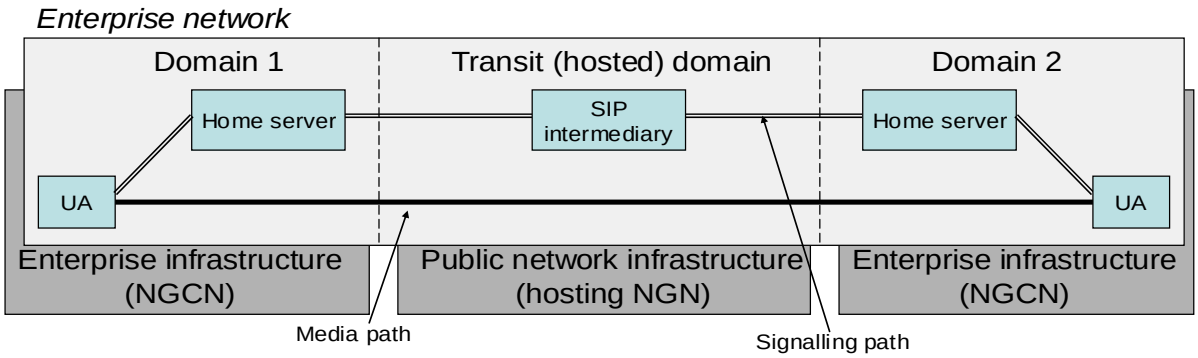

Figure 19 — Example of private network traffic between domains of an NGCN via an NGN

The NGN in this case provides enterprise transit functionality and treats the traffic as private network traffic. The SIP intermediary in the NGN routes requests to the appropriate part of the NGCN based on routing information in the request. In treating the traffic as private, enterprise features such as private numbering plans can be used, traffic is not subject to regulations applicable to public network traffic, and charging may be different.

Signalling between each NGCN domain and the transit domain should be as similar as possible to that between two NGCN domains (e.g., as shown in 7.2.1 or 7.2.2).

7.2.4 Communications between domains of an NGCN using public network traffic through a public SSP such as an NGN

In this case the two NGCN domains are treated as separate NGCNs (i.e., separate enterprises) by the NGN. It is therefore more like indirect peering between two enterprises, as described in 7.3.3.

Compared with the scenario described in 7.2.3 (using private network traffic), the advantage of this arrangement is that it can be entered into on an ad hoc basis if the NGCN domains already have access to the NGN for external traffic. Also different parts of the NGCN can use different NGNs without those NGNs having arrangements for routing private network traffic between them.

Routing considerations are different, because in this case the NGCN cannot take advantage of specific enterprise routing support in the NGN.

7.2.5 Communications between a dedicated NGCN domain and a domain hosted by an NGN

This is illustrated in Figure 20.

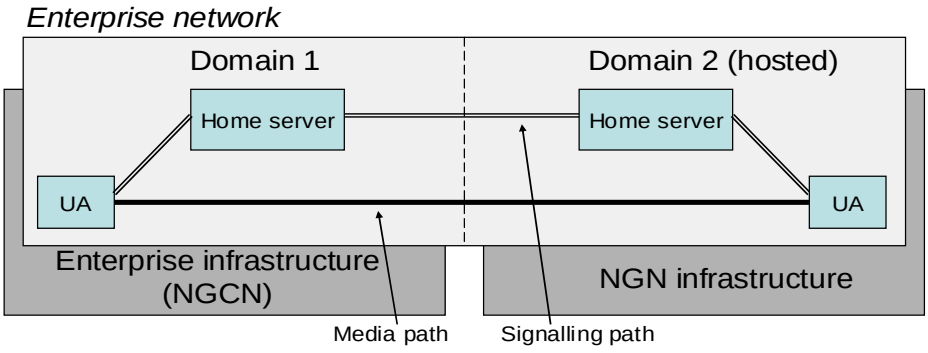

Figure 20 — Example of communication between a dedicated NGCN domain and an NGN-hosted domain

Signalling between the two domains should be as similar as possible to that between two NGCN domains (e.g., as shown in 7.2.1 or 7.2.2).

7.2.6 Communications between a dedicated NGCN domain and a domain hosted by an NGN as private network traffic via an intermediate NGN domain

This is illustrated in Figure 21.

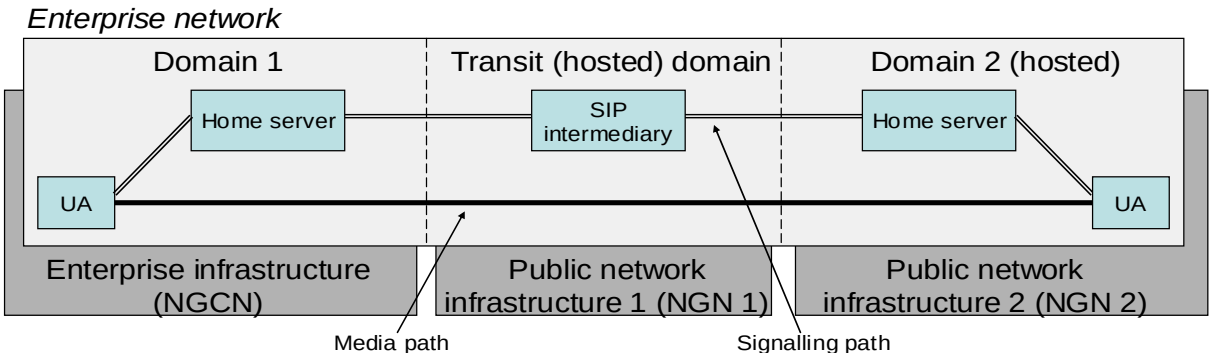

Figure 21 — Example of communication between a dedicated NGCN domain and an NGN-hosted domain as private network traffic via an intermediate NGN domain

Signalling between the NGCN domain and the transit domain and between two hosted domains should be as similar as possible to that between two NGCN domains (e.g., as shown in 7.2.1 or 7.2.2).

7.3 Session-based communication between two enterprises

This sub-clause discusses session-based communications between the networks of two enterprises. The focus is on enterprises supported by dedicated NGCNs (based on the enterprises' own infrastructures) or by hosting NGCNs. Similar principles apply to enterprises hosted by NGNs, except that they will be based largely on the same NGN technology that is used to support public networks. From the perspective of one enterprise, it should make no difference whether the peer enterprise is supported by dedicated infrastructure, by a hosting NGCN or by a hosting NGN.

Inter-enterprise communications are considered public network traffic.

7.3.1 Extension of enterprise network to include partner networks

A partner network infrastructure might be connected to an enterprise's network infrastructure using leased line or VPN connections through a public IP network., thereby forming an "extranet". This has similarities to communication between two parts of an enterprise, but user identifiers are likely to have different domain parts, making routing easier.

This arrangement may attract regulations applicable to public network traffic, since traffic is between two enterprises.

7.3.2 Direct peering

Two enterprises, each operating an NGCN, can agree to allow direct peering between their networks. via a public TSP (or a concatentation of TSPs). No transit SSP is involved. This is illustrated in Figure 22.

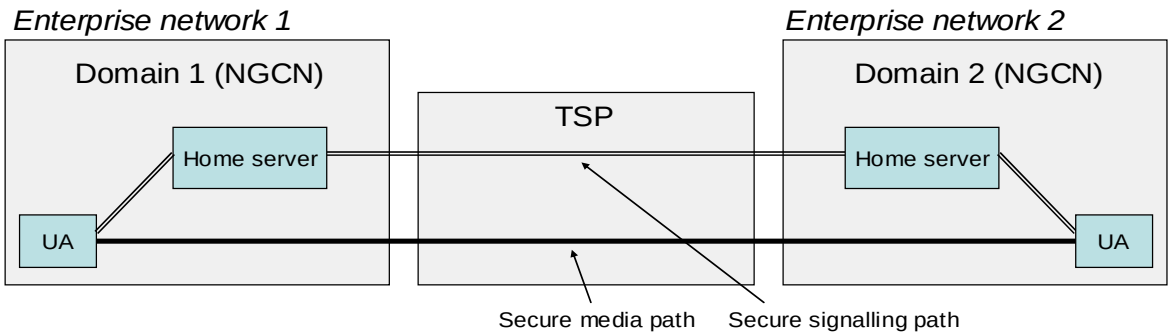

Figure 22 — Example of direct peering between enterprises

Signalling and media are both carried (securely) over the TSP network.

This is the way SIP is intended to work, without any SIP intermediary other than in the originating domain and the terminating domain (the home servers for the respective users), and in theory all inter-domain communications could operate in this way. In practice it does not, because of concern about accepting unwanted traffic from the public Internet. Where direct peering does occur, an NGCN will generally accept communications only from networks with which it has a direct peering agreement in place. This might be done as part of a federation of enterprises that agree to allow direct peering between each other.

Direct peering between enterprises may attract regulations applicable to public networks. The NGCNs concerned will need to take this into account. It is unclear whether in some territories communication between closely related enterprises (e.g., partners) might still be considered private. Conceptually there would be a break-out function at the border of one NGCN and a break-in function at the border of the other NGCN (not shown in the Figure).

Where users of one or both enterprises are hosted (rather than being supported by an NGCN), direct peering between the enterprises will require appropriate support from the hosting network, e.g., NGN. This is not considered further in this Technical Report.

7.3.3 Indirect peering

Two enterprises, each operating an NGCN, can peer via a transit SSP. Each enterprise has a direct peering arrangement with the transit SSP, typically accessed via separate TSPs. Signalling is carried (securely) between each NGCN and the transit SSP. Media is carried (securely) end-to-end (not necessarily via the network infrastructure of the transit SSP). The need for security is particularly true when transport is via intervening TSPs (as opposed to direct connectivity between the network infrastructures of each NGCN and the transit SSP).

This is illustrated in Figure 23 for the case where media also flows through the transit SSP.

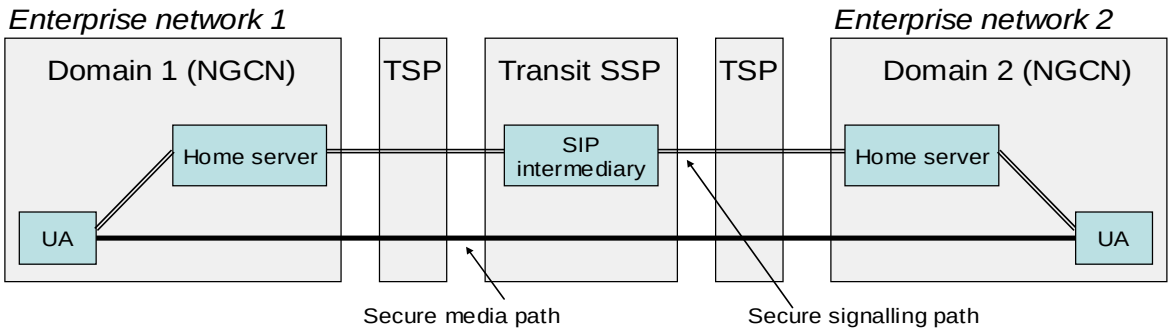

Figure 23 — Example of indirect peering between enterprises

Figure 24 shows a further example where media does not flow through the SSP's network but through a separate TSP.

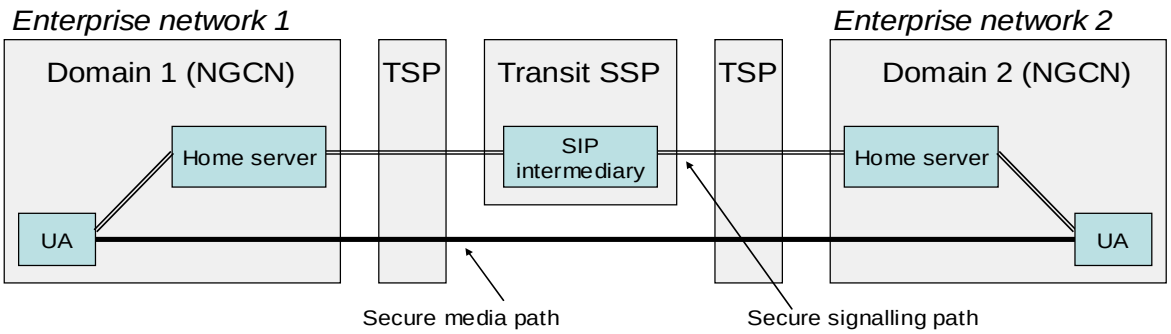

Figure 24 — Example of indirect peering between enterprises

The transit SSP can be another enterprise network (NGCN) or a public network (e.g., NGN). In some cases more than one transit SSP may be involved. For example, consider the case where NGCN1 has a direct peering relationship with NGN1, and NGCN2 has a direct peering relationship with NGN2. NGN1 and NGN2 also have peering relationships. Communications between NGCN1 and NGCN2 will go via NGN1 and NGN2.

Indirect peering may be carried out in the context of a federation of enterprises and other SSPs that agree to engage in peering. For example, a federation could involve a hub that acts as a transit SSP between other members of the federation. In another arrangement, a single member of a federation might have a peering relationship with SSPs outside the federation or with other federations, so sessions to/from those other entities would need to pass through that single member.

Indirect peering between enterprises may attract regulations applicable to public networks, since traffic is between two enterprises. It is unclear whether in some territories communication between closely related enterprises (e.g., partners) might still be considered private. It should be possible to rely on the transit SSP to take care of issues such as lawful interception, in particular if it is a public network (e.g., NGN). Conceptually there would be a break-out function at the border of one NGCN (either in the NGCN or in the transit SSP) and a break-in function at the border of the other NGCN (not shown in the Figure).

7.3.4 Third party assistance

Two enterprises, each operating an NGCN, can use the assistance of a third party to enable them to peer. The assisting entity is queried in order to establish communication to a peer and optionally in order to accept a communication request from a peer. However, the assisting entity does not form part of the resulting communication, which occurs via a TSP. This is illustrated in Figure 25.

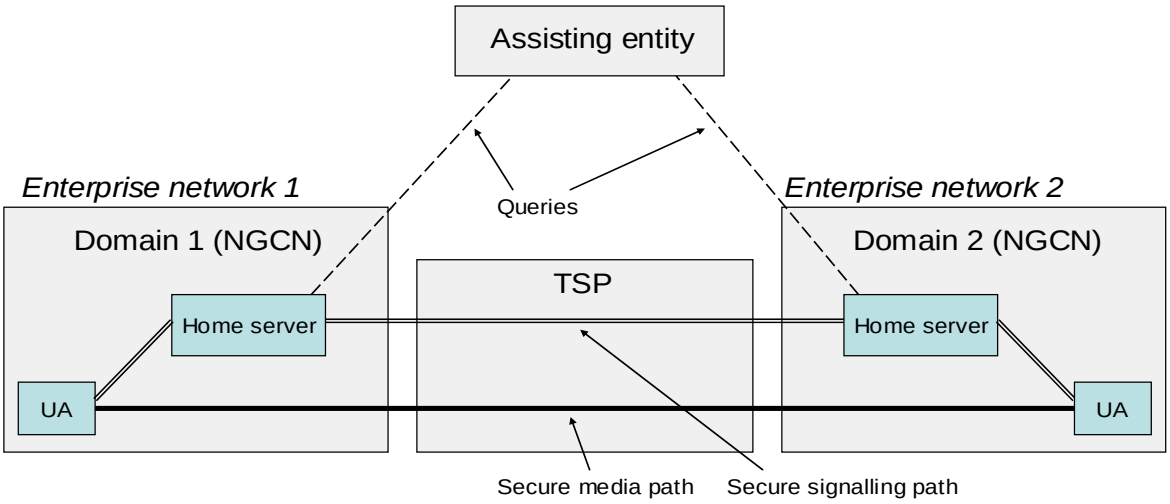

Figure 25 — Example of third party assistance for direct peering between enterprises

Signalling and media are both carried (securely) over the TSP network.

The result of a query by the originating NGCN will provide details on how to route a request in order to reach the required destination. It may also provide credentials that allow the terminating NGCN to accept the incoming request. The terminating NGCN might rely on information in the received request to authorise peering, or it might query the assisting entity

An originating NGCN can use SIP to make a query, in which case the assisting entity acts as a SIP redirect server and provides the result in a 3xx response. Alternatively other protocols can be used, e.g., HTTPS. SIP is not suitable for queries by the terminating NGCN (except possibly the SUBSCRIBE/NOTIFY mechanism).

Assistance is generally associated with federations, where the assisting entity handles queries from all members.

Conceptually there would be a break-out function at the border of one NGCN and a break-in function at the border of the other NGCN (not shown in the Figure).

Although, as shown here, the assisting entity is basically supporting direct peering between the two NGCNs, is could in principle be used in support of indirect peering between the two NGCNs. For example, the result of a query to the assisting entity could yield information for routing to a transit SSP.

7.3.5 Direct peering between two enterprises hosted by the same hosting enterprise infrastructure

Figure 26 illustrates a special case of direct peering where the two enterprises are hosted on the same NGCN infrastructure, often referred to as a "multi-tenant" arrangement. To comply with regulations in some territories, traffic between the two enterprises has to be treated as public network traffic, and hence a break-in or break-out function is needed on either side. Although logically there is a home server for each enterprise, in practice these could be the same physical server.

A similar arrangement would exist if the two enterprises were hosted by the same public network infrastructure.

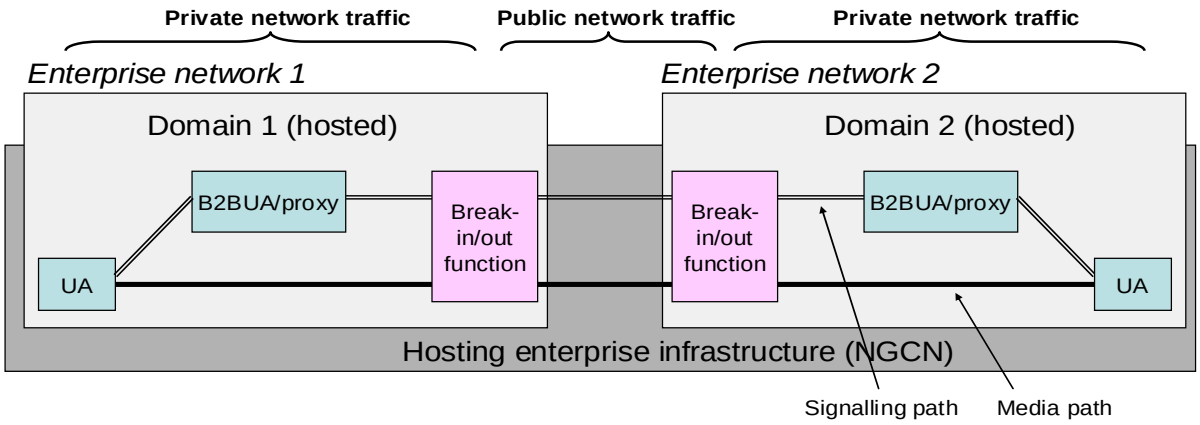

Figure 26 — Example of direct peering between two enterprises hosted by the same hosting enterprise infrastructure

7.4 Session-based communication with users of public networks

This sub-clause discusses session-based communication between users of enterprise networks and users of public networks. It focuses on the case where users of enterprise networks are NGCN users, since communication between NGN-hosted enterprise users and users of public networks is largely an internal matter for the NGN concerned.

A public network here is an NGN or any other public network based on SIP, e.g., a third generation mobile network, a network offering a SIP interface in accordance with the SIP Forum IPPBX service provider interface specification (see 10.2). It can also be a public network reached via a SIP-based public network, e.g., a PSTN/ISDN.

7.4.1 Communication between an NGCN user and an NGN public network user

This is illustrated in Figure 27. The NGN public network user is a residential or SOHO user.

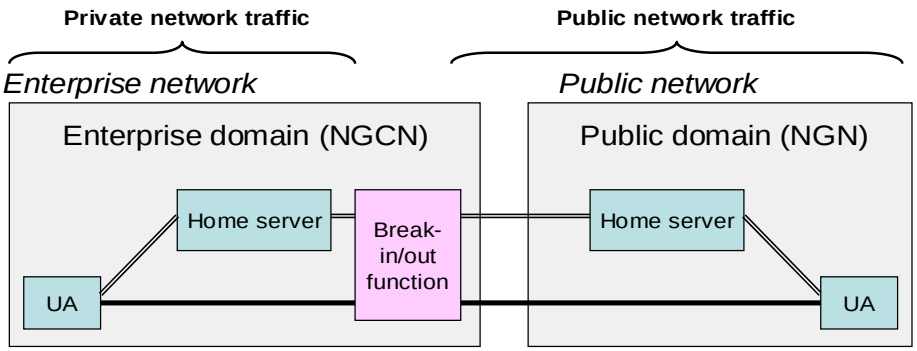

Figure 27 — Example of communication between an NGCN user and an NGN user

In this example the break-in or break-out function is located in the NGCN. For signalling, this function can be located in the Home server or in some other SIP intermediary such as an SBC. For a call from NGCN to NGN, the NGCN determines that the call is to leave the enterprise, inserts a break-out function and signals the call to the NGN as public network traffic. For a call from NGN to NGCN, the NGN signals the call to the NGCN as public network traffic and the NGCN inserts a break-in function. Alternatively the break-in or break-out function could be located in the NGN.

7.4.2 Communication between an NGCN user and an NGN public network user with break-in or break-out function in the NGN

This is illustrated in Figure 28. Effectively the NGN hosts a domain of the enterprise network, which in this particular scenario just provides the break-in or break-out function.

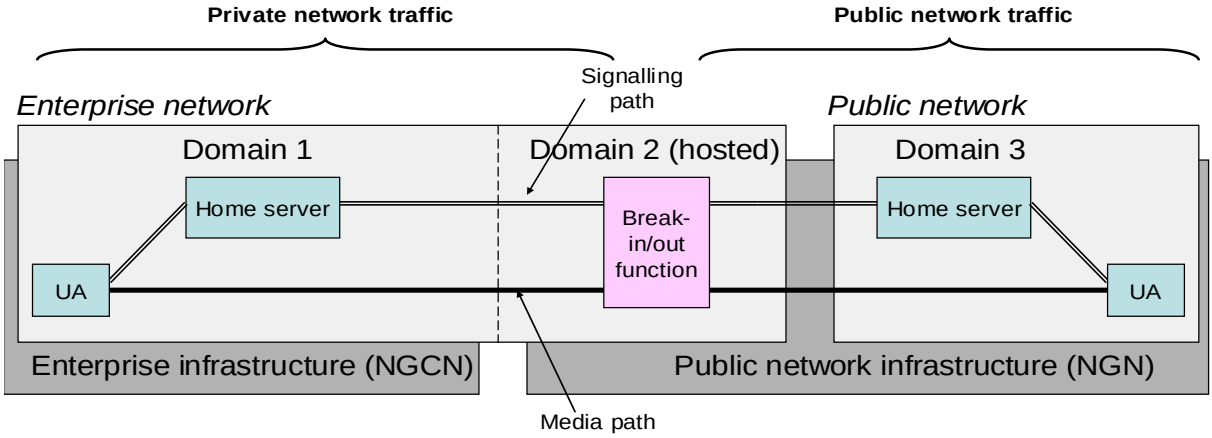

Figure 28 — Example of communication between an NGCN user and an NGN public network user with break-in or break-out function in the NGN.

For a call from NGCN to NGN, the NGCN signals the call to the NGN as private network traffic. The NGN (hosted domain of the enterprise network) determines that the call is to leave the enterprise, inserts a break-out function and signals the call onwards as public network traffic. For a call from NGN to NGCN, the NGN signals the call to the hosted domain of the enterprise network, which inserts a break-in function and signals the call onwards into the NGCN domain.

7.4.3 Communication between an NGCN user and a public network user of a remote NGN

This is illustrated in Figure 29. The NGN user is a residential or SOHO user of NGN 2, which is not the NGN with which the NGCN has an interconnection arrangement (NGN 1). Therefore the communication has to be transited through one or more SIP intermediaries in NGN 1. The break-in or break-out function is located in the NGCN.

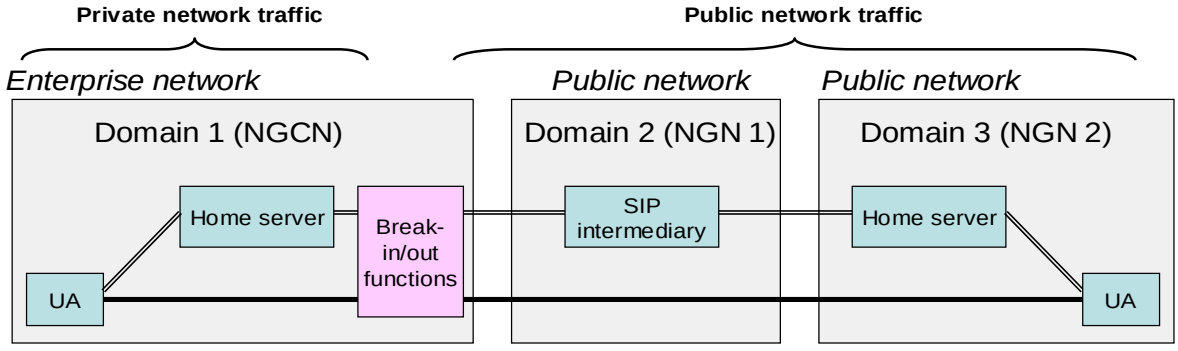

Figure 29 — Example of communication between an NGCN user and a public network user of a remote NGN

7.4.4 Communication between an NGCN user and a public network user of a remote NGN with break-in or break-out function in the local NGN

This is illustrated in Figure 30. Effectively the local NGN hosts a domain of the enterprise network, which in this particular scenario just provides the break-in or break-out function.

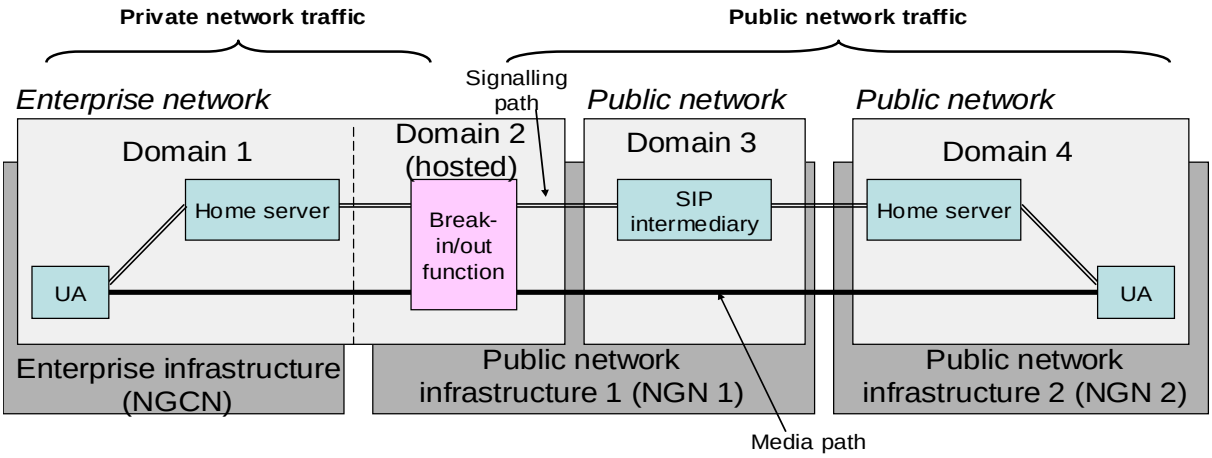

Figure 30 — Example of communication between an NGCN user and a public network user of a remote NGN with break-in or break-out function in the local NGN.

7.4.5 Communication between an NGCN user and a public network user of a remote NGN with break-in or break-out function in the remote NGN

This is illustrated in Figure 31. Effectively the local and remote NGNs each host a domain of the enterprise network, which in this particular scenario just provides a transit or break-in / break-out function respectively.

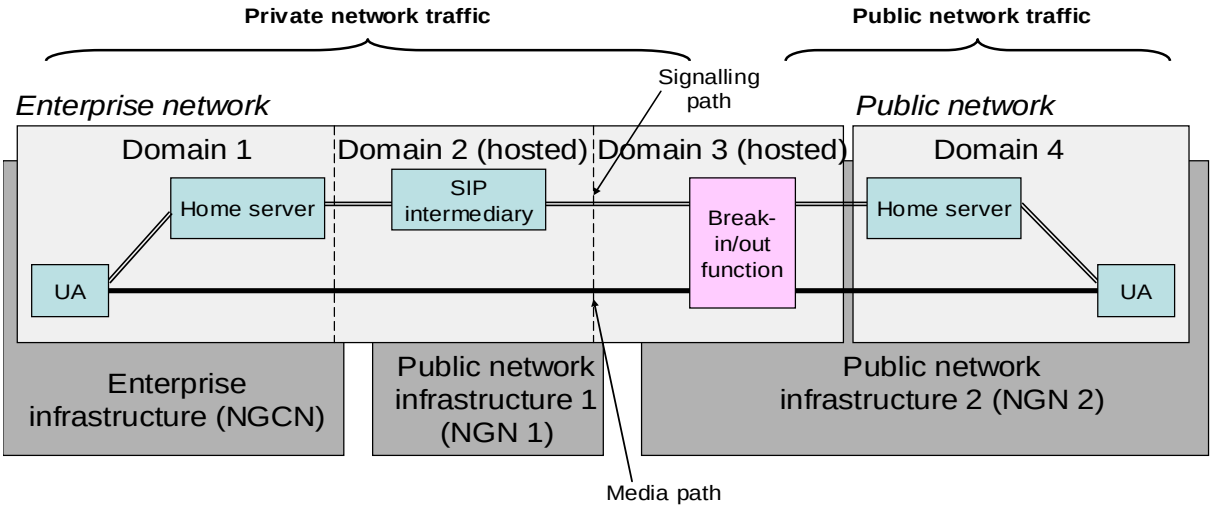

Figure 31 — Example of communication between an NGCN user and a public network user of a remote NGN with break-in or break-out function in the remote NGN.

7.4.6 Communication between an NGCN user and a PSTN/ISDN user via an NGN

This is illustrated in Figure 32. The break-in or break-out function (not shown) is located either in the NGCN or in the NGN.

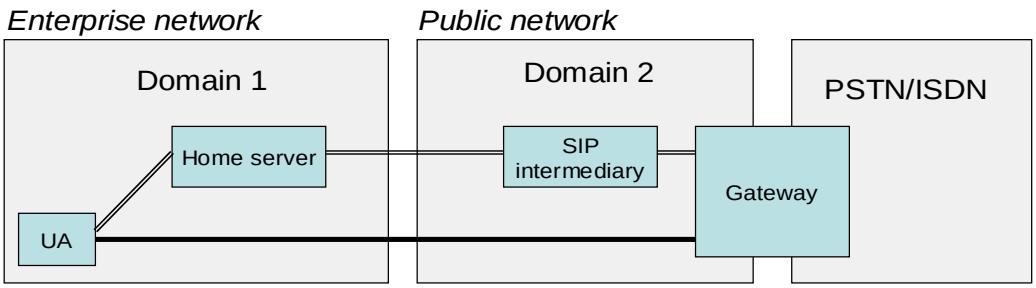

Figure 32 — Example of communication between an NGCN user and PSTN/ISDN user

Other variants of this scenario are possible but not shown, including:

- communication between an NGCN user and a PSTN/ISDN user with break-in or break-out function located in the NGN;
- communication between an NGCN user and a PSTN/ISDN user via a remote NGN;
- communication between an NGCN user and a PSTN/ISDN user via a remote NGN with break-in or break-out function located in the local NGN.

7.4.7 Communication between an enterprise user hosted by a public network infrastructure and a public network user of that same infrastructure

This is illustrated in Figure 33. The NGN-hosted domain of the enterprise network provides the break-in or break-out function.

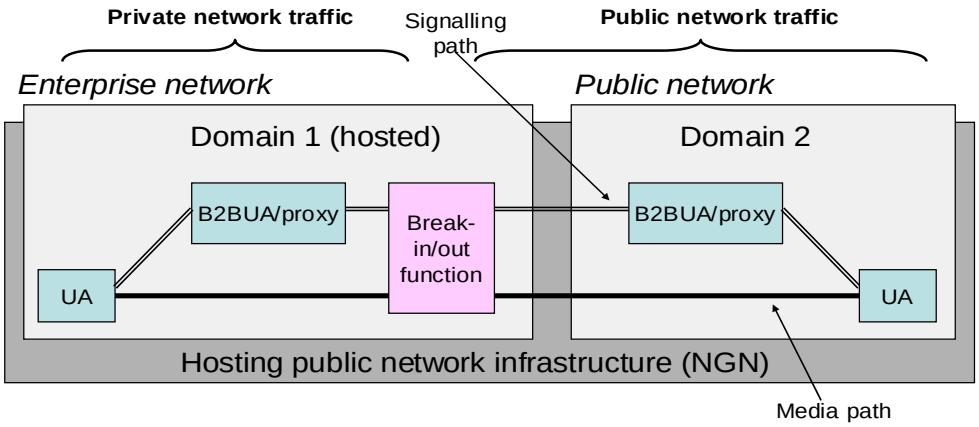

Figure 33 — Example of communication between an enterprise user hosted by a public network infrastructure and a public network user of that same infrastructure

Other variants of this scenario are possible but not shown, including communication between an enterprise user hosted by an NGN and:

- a public network user of a remote NGN;
- a PSTN/ISDN user accessed via a gateway on that same public network infrastructure;
- a PSTN/ISDN user accessed via a gateway on a remote public network infrastructure.

7.5 Session-based roaming

In the scenarios below, only the signalling path between the roaming user and the home domain for registration and for inbound and outbound communications is shown. Media will not

necessarily follow the same path, particularly if the communication partner is not in the home domain.

7.5.1 An NGCN user at a visited sub-domain of the NGCN

In the example in Figure 34, the visited domain has direct connectivity to the home domain over the enterprise infrastructure.

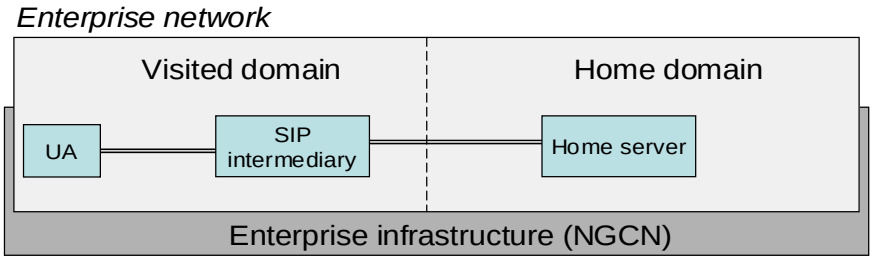

Figure 34 — Example of visiting a different sub-domain of an NGCN

7.5.2 An NGCN user at a visited sub-domain of the NGCN using private network traffic through an NGN

In the example in Figure 35, the NGN provides session level assistance for roaming into the visited NGCN domain. Signalling for this purpose is treated as private network traffic.

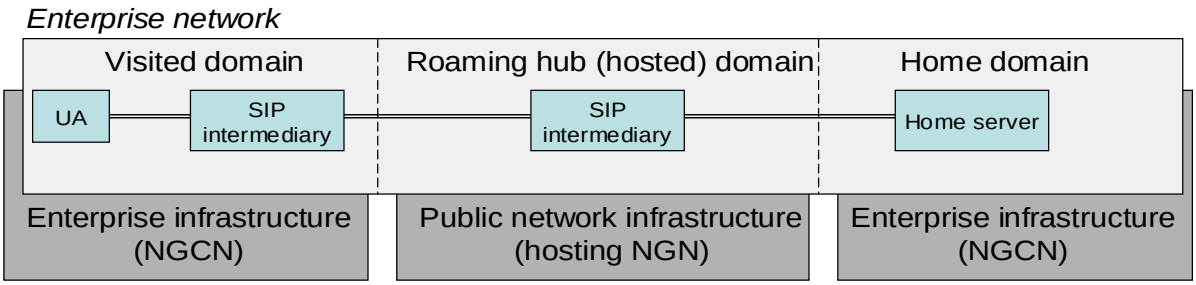

Figure 35 — Example of visiting a different sub-domain of the NGCN via an NGN

7.5.3 An NGCN user at a visited NGN

In the example in Figure 36, the NGCN user visits an NGN with which the enterprise has a roaming agreement.

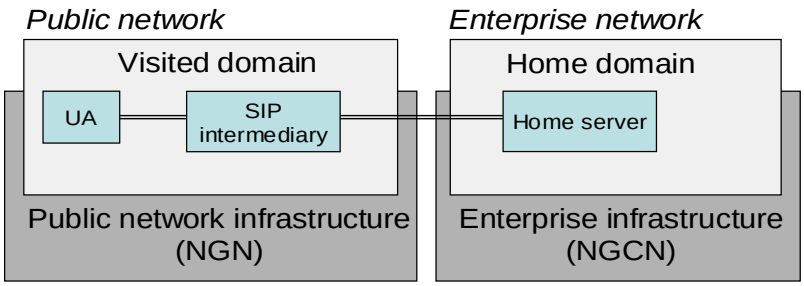

Figure 36 — Example of NGCN user visiting an NGN

7.5.4 An NGCN user at a visited NGN using indirect roaming

In the example in Figure 37, the NGCN user visits an NGN with which the enterprise does not have a roaming agreement, and hence signalling is via an NGN with which the enterprise does have a roaming agreement (the roaming hub).

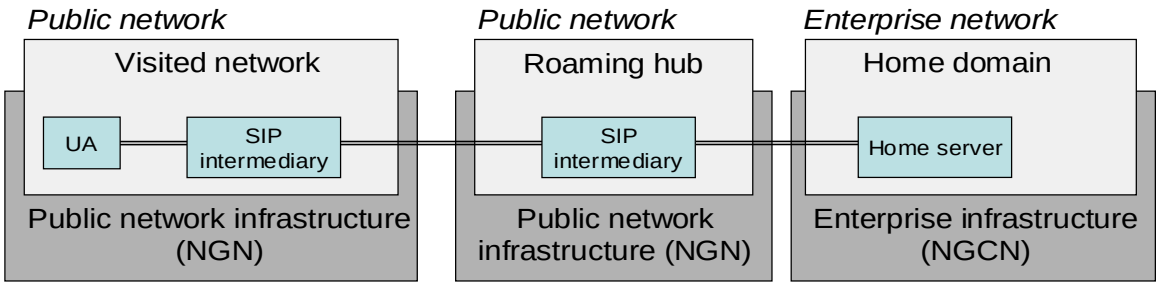

Figure 37 — Example of NGCN user visiting an NGN using indirect roaming

7.5.5 An NGN-hosted enterprise user at a visited NGN

In the example in Figure 38, the hosted enterprise user visits a different NGN. There is no NGCN involvement.

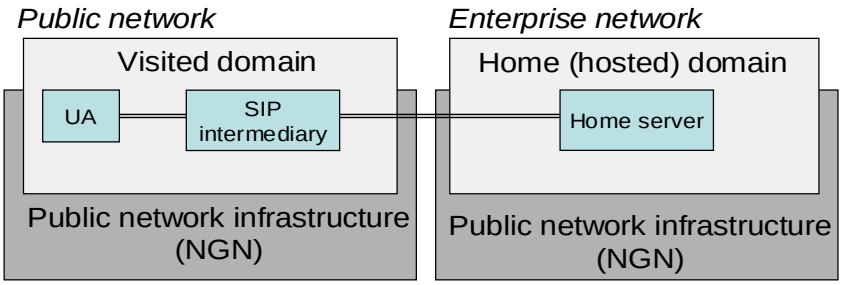

Figure 38 — Example of NGN-hosted user visiting another NGN

7.5.6 An NGN-hosted enterprise user at a visited NGCN

In the example in Figure 39, the hosted enterprise user visits an NGCN within the same enterprise.

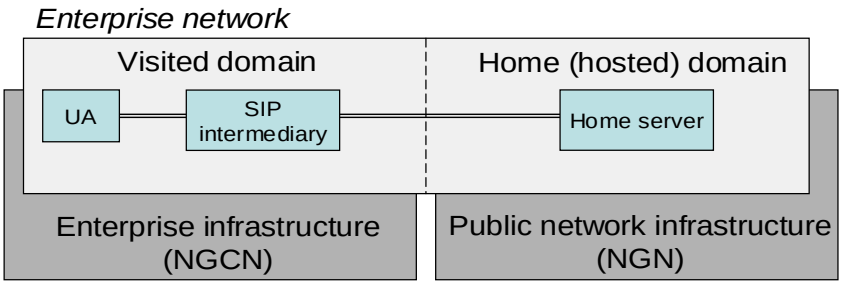

Figure 39 — Example of NGN-hosted user visiting another NGN

8 NGN considerations

Many of the scenarios described above involve NGCN-NGN interaction in some way. In this Clause these scenarios are summarised and some initial thoughts on interfacing NGCN to NGN are expressed.

8.1 Summary of scenarios involving NGN

The following is a list of the main types of scenarios identified. Many variations on these are possible.

8.1.1 Scenarios involving NGN as a TSP

The following scenarios involve NGN only as a TSP. Similar functionality may be available from other types of TSP.

- communication between users of two domains of an NGCN using VPN or other security means;
- communication between an NGCN and a remote NGCN user using VPN or other security means;
- direct peering between two enterprises.

In all cases there could be more than one NGN involved.

8.1.2 Scenarios involving NGN as a SSP

The following scenarios involve NGN as an SSP. Some of these scenarios may also be achievable using other types of SSP.

- communication between users of two domains of an NGCN using private network traffic through an NGN;
- indirect peering between two NGCNs or two domains of an NGCN using public network traffic through an NGN;
- communication between an NGCN user and a public network user;
- communication between an NGCN user and a PSTN/ISDN user;
- communication between an NGCN user and an NGN-hosted user of the same enterprise using private network traffic;
- communication between an NGCN user and a hosted user of another enterprise;
- roaming NGCN user.

In all cases there could be more than one NGN involved.

8.1.3 Roaming scenarios involving NGN as a SSP

The following roaming scenarios involve NGN as an SSP. Some of these scenarios may also be achievable using other types of SSP.

- NGCN user visiting another domain of the same NGCN using private network traffic through an NGN;
- NGCN user visiting an NGN;
- NGN-hosted enterprise user visiting an NGCN of the same enterprise;
- NGN-hosted enterprise user visiting another NGN (does not involve NGCN).

In each case there could be a further NGN involved as a roaming hub, in addition to the NGN acting as home or visited network.

8.2 Interfacing NGCN to NGN

At the time of publication, work was in progress in ETSI TISPAN to define ways in which NGN can support enterprise communications, and in particular on "business trunking" for interfacing an NGCN to an NGN. Two main approaches are considered.

8.2.1 Subscription-based business trunking

With this approach the NGCN appears rather like an ordinary subscriber, even though it supports a large number of end users. The NGCN equipment is treated like a User Equipment (UE). Special arrangements are needed for purposes such as:

- supporting a large number of end users that are managed in the NGCN, meaning that the NGN will not be aware of individual users and their identifiers, but only a range of identifiers assigned to the NGCN;
- discrimination between public network traffic and private network traffic;
- special handling of private network traffic, e.g., support private numbering plans, different regulatory impact;
- accommodation of a SIP proxy, as opposed to a SIP UA, on the NGCN side of the interface.

The NGCN has to use the SIP registration procedure to obtain service from the NGN.

8.2.2 Peering-based business trunking

With this approach the NGCN appears more like another NGN. Again some special arrangements may be needed, particularly in the following areas:

- security and trust;
- discrimination between public network traffic and private network traffic;
- special handling of private network traffic, e.g., support private numbering plans, different regulatory impact.

8.2.3 Roaming

Roaming support at the interface between an NGCN and an NGN is being progressed in ETSI TISPAN in a manner that is independent of the method of achieving business trunking. The interface needs to provide:

- a means for an NGCN user to register with its home server in the NGCN when roaming outside the NGCN; and
- a means for a user of another domain of the enterprise to register from the NGCN when visiting the NGCN.

9 Application considerations

Session-based communications, whilst useful on their own for basic communications between users, are also increasingly used by applications to give added value to those applications. In other words, session-based communications provide a service available to applications, which act as the users of communication sessions. See also the overall architecture of NGCN in 6.1.

Examples of added value include:

- the ability to establish, manipulate and tear down sessions;
- to obtain information about sessions, e.g., identity of communication partners, the geographic location of a communication partner, security status;
- to transfer application-specific information between communication partners in a reliable and secure manner.

Such a service could be realised within a Service-Oriented Architecture (SOA). A SOA defines how two or more entities interact in such a way as to enable one entity to perform a unit of work (a

service) on behalf of another entity. SOA helps integration with other services that are used in business processes.

Some applications may in turn provide higher level communications services to other applications. Examples include presence services, messaging services, conferencing services and services providing mediation between users and devices with different capabilities.

NOTE
This concept of applications is not to be confused with the IMS concept of Application Servers, which can be deployed even just as a means of accomplishing communication services and not necessarily for providing applications on top of those communication services.

A communications service can incorporate a SIP UA as a means of participating in sessions. Alternatively it can incorporate a SIP B2BUA and use third party call control techniques in SIP to establish, manipulate and tear down sessions. Another possibility is to incorporate a SIP proxy or redirect server if it just needs to provide limited functionality for manipulating incoming requests, e.g., in support of a scripting application for handling incoming calls.

CSTA [1] provides an example of functionality that can be exposed by a communications service. CSTA can be realised in a number of ways, including an API, a web service and an XML-based protocol.

Examples of applications that rely on communication services include:

- Unified Communications
  - One identifier for all forms of communication
  - Directory services
- Messaging
  - Combined IP messaging (CPM)
  - Push-based Email
  - Instant messaging
  - Voice mail
  - File transfer
- Presence
  - Availability
  - Location
- Location-based services
  - Tracking of vehicles and goods
  - Location-dependent services
- Collaboration
  - Audio/video conferencing
  - Application sharing
  - Meeting planning
- Corporate Information
  - Announcements
  - Television
- Business processes support
  - Electronic Data Interchange (EDI)
  - Customer Relations Management (CRM)

- Logistics and Transport
- Enterprise Resource Planning (ERP)

- Device Management

10 Current standards and standardisation efforts

10.1 IETF Real-time Applications and Infrastructure (RAI) area

The RAI area of the IETF includes several working groups of relevance.

The SIP (Session Initiation Protocol) and SIPPING (Session Initiation Proposal Investigation) working groups are responsible for the evolution of SIP. Any extensions to the SIP protocol and associated security are specified by the SIP working group, whereas SIPPING carries out investigations, defines some SIP event packages, and produces some Best Current Practice RFCs.

The MMUSIC (Multiparty Multimedia Session Control) working group is responsible for SDP.

The SPEERMINT (Session Peering for Multimedia Interconnect) working group addresses peering between SSPs using SIP. Although initially addressing only public SSPs, it now appears that its output will also be applicable to enterprise peering.

10.2 SIP Forum

The SIP Forum has published an initial version of a specification [9] for connection of an IPPBX (in other words, an NGCN) to a SIP service provider, essentially for the purpose of communicating with the PSTN, although in principle it can easily be extended to cover connection to a SIP service provider for other purposes (communication with other SIP users or networks). A number of SIP service providers are basing their interfaces to NGCN on that specification, rather than on TISPAN standards, for example.

This specification uses basic peering principles, although it is not sufficiently symmetrical to allow application to other peering situations (e.g., direct peering between NGCNs).

10.3 ETSI TISPAN

ETSI TISPAN, as part of its work on NGN, is responsible for several aspects of support for business communications, including:

- hosted enterprise services (HES), for support of hosted enterprise users on NGN;
- business trunking (connection of an NGCN to an NGN - subscription-based and peering based approaches);
- roaming of enterprise users; and
- virtual leased line (support of private network traffic between NGCN domains).

Deliverables on these subjects are part of TISPAN Release 2 and are based on requirements submitted by Ecma.

10.4 3GPP

3GPP is responsible for IMS, which is the basis for NGN. Any enhancements to IMS to support the business communications work of ETSI TISPAN will need to be handled by 3GPP.

10.5 ITU-T Study Group 13

ITU-T Study Group 13 is responsible for recommendations relating to the architecture, evolution and convergence of next generation networks including frameworks and functional architectures, signalling requirements for NGN, NGN project management coordination across study groups and release planning, implementation scenarios and deployment models, network and service capabilities, interoperability, impact of IPv6, NGN mobility and network convergence and public data network aspects.