A Web Service-based Network Composition Architecture for the Next Generation Internet

Luciano B. de Paula[1], Rodolfo Villaça[1], Fábio L. Verdi[1], Maurício F. Magalhães[1], and Annikki Welin[2]

[1] Department of Computer Engineering and Industrial Automation (DCA), School of Electrical and Computer Engineering (FEEC), State University of Campinas (UNICAMP), Campinas, Brazil
`{luciano, villaca, verdi, mauricio}@dca.fee.unicamp.br`
[2] Ericsson Research, Torshamnsgatan 23, 16480, Stockholm, Sweden
`annikki.welin@ericsson.com`

Abstract. Network Composition is, in a rough way, the process which one network can offer/request services to/from another network with little or even no human interference. To explore the full potential of Network Composition, a new Internet architecture is needed to provide support for node and domain mobility, security and heterogeneity. This paper presents an architecture for Network Composition and its implementation on top of a prototype of a next generation Internet infrastructure in order to show how this can be achieved and the advantages of this combination. Experimental scenarios based on Web Services were implemented to validate the model.

1 Introduction

Imagine a scenario where a user with a laptop enters in an urban train and wants to get Internet access. The train offers this service through a hotspot located in the train station, so the user can access the service by paying some fee. When the train starts its trip and goes far from the initial station, its connectivity with the hotspot becomes weak. Automatically, the Internet access infrastructure of the train changes its connectivity to the Internet from the station hotspot to the city wireless network. The fee charged from the user is updated. Some minutes later, the train enters in the coverage area of the hotspot located in the next train station. This hotspot offers Internet access by a smaller fee. The train network realizes it and changes, automatically, its connectivity with the city wireless network to the station hotspot. The fee charged from the user is updated again, decreasing the value to be paid. All this process is not noticed by the user.

The scenario described above illustrates four network compositions: one between the user and the train, two others between the train and the stations and one more between the train and the city wireless network. As defined by the Ambient Network (AN) project [1], a network composition is, in a rough way, the process by which one network can offer/request services to/from another network. The overall process formed by the negotiation, configuration, and establishment of the 'composition agreement' is designed to be done automatically.

The composition process makes possible the user to use the train's network, the train's network to accesses the stations services and so on.

A key concept in the AN project is the composition of different ANs or domains providing access to services through the establishment of agreements. The composition of ANs allows a domain to have the possibility of choosing different alternatives before using a given service. The composition between ANs is performed by the Ambient Control Space (ACS) of the ANs involved in the process. The main challenge is to define how such process will occur and how it happens in a real scenario.

The objective of this paper is to propose and evaluate an implementation architecture in order to support the composition of domains on top of a next generation network infrastructure. This next generation network has to support mobility of nodes and domains, layer 3 heterogeneity and security. Currently these requirements have been solved by using narrow solutions that sometimes do not work together, such as the Network Address Translation (NAT), Mobile IP mechanisms [2] and IPSec. We have designed, implemented and tested a prototype that naturally supports flat and global identifiers, domain heterogeneity, and node and domain mobility [3][8]. For this paper, the prototype was improved with domain identifiers. The challenge is to make the prototype of the next generation network and the composition architecture presented to work together.

We implemented the modules of our composition architecture using Web services which is a technology indicated by the AN project as a possible mechanism to instantiate the functional entities belonging to the ACS [1]. The contribution of this paper is twofold: First, the composition architecture is running on top of a next generation infraestructure that contemplates all the requirements of the future Internet architecture and second, the composition architecture was totally developed using Web services and evaluated in order to verify the feasibility of such tecnology as an option for implementing this architecture.

The paper is organized as follows: the next section shortly presents the background on network composition. Section 3 details the proposed composition architecture, the modules and the functionalities that are necessary to perform the composition of ANs. Section 4 presents the scenarios used to test and evaluate the proposed architecture. Finally, Section 5 concludes the paper.

2 Basic Concepts on Network Composition

One of the key concepts for the future networks is the composition, the process which one network can offer/request services to/from another network by negotiating a Composition Agreement (CA) between these networks, with little or even no user effort. This agreement between two domains defines the services that one domain will be allowed to use from the other one, the characteristics of the services and how they will be accessed and used.

As defined by AN project, every domain has several Functional Entities (FEs). These FEs are part of the ACS and deal with some functionality that exists in the AN as, for example, the Management-FE (that takes care of the

general management of the ACS), the InQA-FE (that takes care of the QoS aspects inside the ACS) and so on [1]. In this paper, the general term for a FE is X-FE or Y-FE. These terms represent "Any"-FE.

Every composition follows specific rules and policies, generally established by the domain administrator. The phases defined by the AN project are [1]: *Phase 1 - Media Sense:* in this phase the ANs find each other. Examples of media sense include an access point that gets turned on in a network, a wired or wireless link is set up between two devices of two different domains, or even a new domain that becomes alive and appears to other ANs; *Phase 2 - Discovery/Advertisement:* in this phase, the ANs are able to offer services to other ANs, as well as to discover services provided by different ANs. These two actions can be done actively or passively; *Phase 3 - Security and Internetworking Establishment:* when the previous phase results in an AN candidate for composition, the two ANs establish basic security and Internetworking connectivities. One way to do that is by trading certificates authenticated by a Trusted Third Party (TTP). This secure interconnection will be used for every negotiation and information exchanging between the two ANs; *Phase 4 - Composition Agreement Negotiation:* having received the information from other ANs, a given AN can then decide if the composition is desired or not. Also in this phase, the ANs can negotiate as a way of filling the requisites of both ANs. The end of this phase can be the creation of a CA formalizing the composition; *Phase 5 - Composition Agreement Realization:* this phase only occurs if the CA was defined in the previous phase. This phase "prepares" the functional elements that will be involved in the shared services. The necessary FEs are properly configured in this phase in order to reflect the negotiated CA. Then, the composition is completed.

There are two main FEs related to the composition phases: the Network Advertisement and Discovery FE (NAD-FE), and the Composition FE (C-FE). The NAD-FE is responsible for receiving the information (in the form of profiles containing the AN identification) related to the services offered by other ANs, as well as for advertising its services offered to other ANs. The NAD-FE can receive triggers from other FEs in the following situations: 1 - A given X-FE can ask the NAD-FE to be announced as an available service. The X-FE sends to the NAD-FE the service characteristics, as well as some advertisements information as, for example, its periodicity. So, the NAD-FE assumes the responsibility for advertising the service to the other ANs respecting the characteristics informed by the X-FE; 2 - A given X-FE can ask the NAD-FE to be informed when the NAD-FE receives the advertisement of an specific service with certain characteristics. When the NAD-FE receives an advertisement from other ANs of a service matching the characteristics informed by the X-FE, the NAD-FE notifies the X-FE indicating which AN is offering the requested service.

When an X-FE receives the information from the NAD-FE about an AN offering the desired service, the X-FE triggers the C-FE to start a composition process. Based on the profile sent by the X-FE, the C-FE starts the composition process with the peer AN. The C-FEs negotiate a Composition Agreement (CA) regarding the services the X-FE wants to use. When the CA is defined, the C-FE

of each AN configures the FEs involved in the composition to reflect the CA. The C-FE is responsible for the decomposition process when the CA is cancelled or it expires and is not re-negotiated.

Another important issue is how the composition affects the naming and the addressing mechanisms of the composing networks [4]. The AN project defines the Ambient Naming System (ANS) which is responsible for the name registration and resolution framework. This framework enables the FEs and other entities in the ANs to find the details of the peer entities, either in the local AN or in a remote AN. The following are examples of control plane objects to be named in the AN architecture: ANs and Security Domains, ACS services, users and other legal entities. Also, the ANS must support legacy user-space naming systems such as DNS, LDAP (Lightweight Directory Access Protocol) and next-generation naming systems like CoDoNS (Cooperative Domain Name System) [5], INS (Intentional Naming System) [6] and any DHT-based naming systems.

There are two main elements in the ANS: the ACS Registry-FE and the DEEP (Destination Endpoint Exploration Protocol) protocol [7]. The first one is used for intra-domain registration and resolution of name-object bindings, and the second one defines the ANS name resolution over all ANs.

The DEEP name resolution is typically a gradual process which resolves a name into a locator (or identifier) in a distributed approach. It involves multiple DEEP nodes and name resolution services that can be independent of each other. The name is resolved in several steps by sending the same query message from one DEEP resolver to the next one until the name is resolved. DEEP nodes provide two interfaces: a local name resolution interface for accessing local resolvers and an external interface for exchanging messages with other DEEP nodes.

3 The Composition Architecture

The composition architecture presented in this work focuses on three of the five phases described in Section 2: phase 2, phase 4 and phase 5. Phase 1 is not considered in this paper since it depends on the media type used by each AN and can varies from one AN to another. Phase 3 deserves specific attention since security is a quite important item to be considered during the composition. Such item will be addressed in future papers.

The functionalities performed by the NAD-FE correspond to the phase 2 (Discovery and Advertisement) as described in Section 2 and those performed by the C-FE correspond to the phase 4 (Composition Agreement Negotiation) and phase 5 (Composition Agreement Realization). So, this two FEs, together with the naming system modules (DEEP and ACS Registry), form the key components for performing network composition.

Figure 1 shows the overall view of the architecture. The arrows represent the interactions between the modules. The X-FEs and the C-FE must be registered in the ACS Registry. At the same time, the ACS Registry is used by the DEEP module as a database. The X-FEs and the C-FE query the DEEP module for name resolution. The X-FEs interact with the NAD-FE to request services and to receive service offerings. The relations between the X-FEs and the C-FE occur in

order to start the composition process. Internally to each module there are sub-modules performing small tasks to completely execute the composition process. The arrows between the ANs represent the composition negotiation (between the C-FEs), name resolution (between the DEEP modules) and the AN profile exchange (between NAD-FEs, requesting or advertising services).
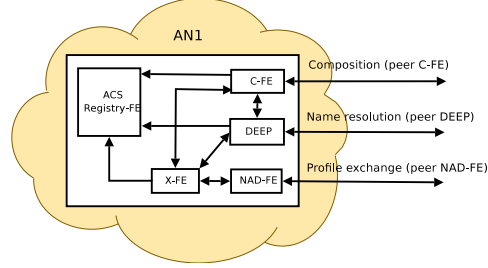

Fig. 1. General view of the architecture.

All the interactions between the FEs located inside the same AN, as well as the communication between modules located in different ANs are performed using XML traded by Web services. The profile that is exchanged between NAD-FEs of different ANs should contain enough information to make the composition possible. Typical information includes the name of the composition entity and the description of the services offered by the AN. Fig. 2 shows an excerpt of an AN profile defined in this work.

```
<?xml version="1.0" encoding="UTF-8"?>
<AN_Profile>
    <ANname>AN2</ANname>
    <compositionEntity>Composition@AN2</compositionEntity>
    <service>
        <serviceDescription>Internet Access</serviceDescription>
        <type>service</type>
        <provider>InternetAccess@AN2</provider>
        <characteristics>
            <bandwidth>270kbps</bandwidth>
            <cost>$0.5/min</cost>
            ...
        </characteristics>
        ...
    </service>
</AN_Profile>
```

Fig. 2. Example of an AN profile.

In the AN profile *ANname* identifies the name of the AN; the *compositionEntity* holds the name of the entity that is responsible for the composition in the AN; the *service* describes all the services and their features (such as its type and provider) offered by the AN; and *characteristics* presents other features like price, QoS, time duration and any other can be defined inside this element. In this example, the AN2 is offering Internet access having a bandwidth of 270kbps and the cost to use the service is $0.5 per minute.

The composition process was modelled by a sequence of messages representing the interactions between the modules. Fig. 3 shows these messages without those that represent the name resolution process which will be explained later.
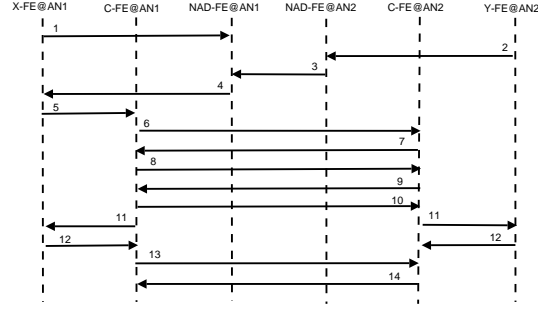
Fig. 3. The composition process.

In the first message (1) the X-FE at AN1 asks the NAD-FE of its AN to be informed when the NAD-FE receives a profile from an AN offering a service with an specific name and possessing some specific characteristics; the Y-FE at AN2 sends a message (2) to the NAD-FE of its AN informing it to advertise a service with some characteristics. This event can even happen before message 1; the NAD-FE at AN1 receives the profile of AN2 from the NAD-FE (3). The profile contains the services offered by AN2, including the one offered by Y-FE. Then the NAD-FE at AN1 realizes that the Y-FE at AN2 offers the service with the characteristics desired by the X-FE; the NAD-FE at AN1 tells X-FE that AN2 is offering the service that it had asked for. The NAD-FE passes to the X-FE the name of the C-FE at AN2 (4); the X-FE tells the C-FE at AN1 that it wants to use the service offered by AN2 (5). This means that a composition between AN1 and AN2 is needed. This message contains the name of the C-FE at AN2 (C-FE@AN2), so the C-FE at AN1 can locate the C-FE at AN2; the C-FE at AN1 contacts the C-FE at AN2 to initiate the composition process (6); the C-FE at AN2 sends to C-FE at AN1 a template of a CA (7); the C-FE at AN1 fills in the template informing all the necessary information, like the service(s) that it wants to use from AN2 and some desired characteristics. The template is sent back to the C-FE at AN2 (8); If AN2 agrees with the template that it has just received, an ACK message is sent (9). This message tells to AN1 that the terms of the composition were accepted; AN1 sends to AN2 one message informing it to "prepare" the composition (10), which means that AN1 and AN2 have to configure all the elements that will be involved in the composition. This configuration is done to reflect what was agreed in the CA; the C-FEs at AN1 and at AN2 send the necessary configuration to the FEs that will take part of the composition (X-FE and Y-FE, respectively) (11); the FEs notify the C-FE that the configuration was completed (12); the C-FE at AN1 notifies the conclusion of the composition to the C-FE at AN2 (13); the C-FE at AN2 sends a "committing" message to the C-FE at AN1 and the composition is properly established (14). Only the name of the composition entity (C-FE), followed by the list of services available, appears in the profile exchanged by the ANs. In the example of Fig. 2, the name of the entity is *Composition@AN2*. The name resolution service is used by the AN-enabled applications, as well as by all the FEs inside the ACS, as explained below.

3.1 The DEEP resolution

The first step for the C-FE to start the composition is to resolve the name of the composition entity and obtain its contact information. This is done by querying the local DEEP gateway, which in turn has to analyze the name to be resolved. Based on the local AN technology, it decides whether it can solve the query locally or if it is necessary to forward the EXPLORE message to another DEEP Gateway in the DEEP Overlay Network towards the network whose name belongs to (in the example of Fig. 2, it is the network named AN2). In our architecture, this decision is based on the AN name[3]: if the destination AN is equal to the AN in which the DEEP Gateway is authoritative, the query is solved locally, otherwise the EXPLORE message is forwarded to the next one.

Our solution for the DEEP architecture implementation is based on two different entities: the DEEP Resolver and the DEEP Gateway. The DEEP Resolver is located in every node participating in an AN and does not participate in the DEEP Overlay Network. The DEEP Resolver always acts as a source DEEP, generating EXPLORE messages, or acts as a destination DEEP, generating RESPONSE messages. The DEEP Gateway implements all the features of the DEEP Resolver plus the capacity to forward EXPLORE messages to another DEEP Gateway in the DEEP Overlay Network, so it can either sends a RESPONSE XML-based message to the DEEP Source or forwards the EXPLORE message to another DEEP Gateway.

Figure 4 depicts a complete scenario formed by four ANs, each one with one DEEP Gateway, one Deep Resolver and one ACS Registry. The ellipses are representations of each AN and the cloud represents the DEEP Overlay Network in which all DEEP Gateways participate. An slashed arrow represents a Web service invocation for the simplified EXPLORE message from the client to its local DEEP Resolver, and another slashed arrow represents the ACS Registry query from the destination DEEP Gateway to the ACS Registry Database. The dotted arrows denote the paths of the EXPLORE and RESPONSE messages. The EXPLORE message is forwarded by DEEP intermediaries (DEEP Resolver and DEEP Gateway for AN1). The RESPONSE is sent directly by the AN2 DEEP Gateway (DEEP Destination) to the DEEP Resolver (DEEP Source).

After the composition is done, the C-FE notifies the respective X-FE about the finalization of the composition. Before using the service, the X-FE needs to resolve the service name using the DEEP protocol. In our example (Fig. 2), the service name *InternetAccess@AN2* is resolved and the contact information is returned to the application which in turn can access the service.

4 Implementation and Evaluation

The architecture presented in the previous section can be seen as a general framework of the main modules necessary to perform the composition of ANs. However, to make all the modules and FEs work together, it is necessary to

---

[3] A name in this work is composed by two parts: the object name and the AN name, which are separated by the symbol '@', for example: Composition@AN3.

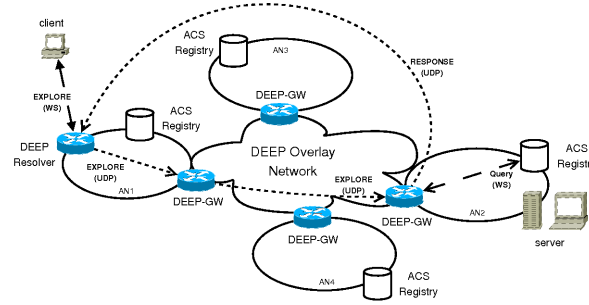**Fig. 4.** DEEP implementation scenario.

specify how such modules communicate to each other. The implementation of the architecture using Web services is a way to validate and experiment how such technology can help in the composition scenario in a next generation Internet.

The composition architecture, as well as the legacy applications used in this work are entirely running on our implementation of a next generation Internet. All modules presented in Fig. 1 were implemented using Web services. The interface of each Web service was defined by following the required parameters specified by the AN project for each FE and mapped to an XML message. First of all, a given X-FE, acting in name of an AN-enabled application, registers its interest for service(s) in the NAD-FE. So, the NAD-FE interface consists in offering methods to send and receive solicitation and announcing messages to/from other ANs, and receiving the register for interest from local X-FEs.

The C-FE implements the methods responsible for the composition negotiation, composition realization and composition deletion. These three methods are implemented through three Web services methods: the *"caneg"* method corresponds to the composition agreement negotiation; the *"carlz"* corresponds to the composition agreement realization and the *"cadel"* method corresponds to the composition agreement deletion.

4.1 Evaluated scenarios

In order to test and evaluate our composition architecture over the next generation Internet prototype, we created two scenarios described below. These scenarios were chosen since they can show two different aspects of the composition process. The first one presents a situation where the composition is done by the desire of a user located in a given domain. The second scenario presents a situation where there is a mobile event (a domain mobility), and the composition is triggered by this event, composing the mobile domain with the static one. These two scenarios are representative and enough to evaluate all the functionalities of the architecture. Any other scenario can be defined as a variation of these two situations. In both scenarios, the DEEP resolution (Section 3.1) is part of the process. The tests were done using Linux Debian boxes (Intel Core 2 Duo, 4 GB of RAM), Ethernet 100 Mb as the wired network and wireless 802.11g.

Scenario 1 - Composition by demand: In this scenario a home network has Internet access through a CA established with a domain offering the access

service (the domain of a provider in the city, for example). Trying to lower the costs, the FE responsible by the management of the home network (called in this scenario Management-FE) asks the network ACS to find another domain offering the accesss service by a lower cost. Later, another domain supporting the Internet access by a cheaper cost shows up (a new provider in the city, for example) and starts to advertise the service. The ACS of the user's home network realizes it and, due to the interest of the home network Management-FE, it starts a composition process with this new provider. After the composition process is finished, all the user devices in the home network are notified and they can start to access the Internet via the new provider.

In our implementation, the scenario described above can be translated into a topology as the one presented in Fig. 5. Each domain is related to a different Ambient Network, named AN1, AN2, AN3 and AN4. The AN1 has two gateways named GW1_AN1 and GW2_AN1 that can be used for Internet access. The scenario includes a client located at AN1 accessing the Internet for a Server located at AN2. This client can illustrate all the other ones in the home network.

To use the access service, the home network is charged by a certain value per minute. Due to the management decision, the home network tries to minimize the costs with the usage of the access service. The Management-FE of the home network informs this interest to the ACS of AN1 by sending a message to the NAD-FE. Later on, another domain (AN4, the new provider) begins to offer Internet access to it. As the cost of usage of the service offered by AN4 is cheaper than the one offered by AN3, the ACS of AN1 starts to compose with AN4. After the composition process is finished, the AN4 acts as a transit domain between AN1 and the Internet. The client's devices are warned by the NAD-FE about this new option, so they change their Internet accesses from AN3 to AN4, changing their default gateway from GW1_AN1 to GW2_AN2. In this scenario, we are interested in measuring the amount of time needed for the composition process to be performed. Two sets of tests were done (600 composition processes in each set) between two adjacent domains. The first set was done using two domains connected by a wired link. The second set was done using two domains connected by a wireless link. The time of the name resolution done by the DEEP modules and the entire composition process were recorded for each test. The average results obtained were:
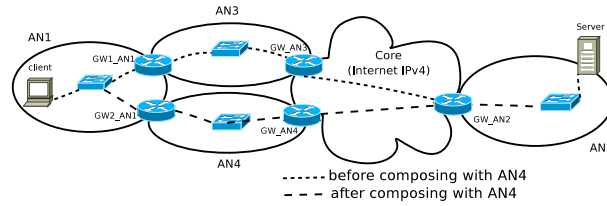

Fig. 5. Scenario 1

– **C-FE name resolution**. The name resolution of the C-FE of the peer AN done by the AN that initiates the composition process using the DEEP modules: 46.11 ms for the wired scenario and 52.18 ms for wireless scenario.

– **Total composition process**. The entire composition process, including the C-FE name resolution by DEEP modules, until the establishment of the CA (after all messages in Fig. 3 be performed): 125.48 ms for the wired scenario and 162.42 ms for the wireless scenario.

As showed above, the time of the resolution and the total time of the composition are very short. Such time is mostly related to the Web services engine and XML processing. Since the composition in this scenario was done before the breaking of the previous connectivity (a *make-before-break* approach), the composition process does not cause any loss of data. The flow of the data just changes its path (from GW1_AN1 to GW2_AN2) after the composition is completely done.

Scenario 2 - Domain mobility and composition: A user with a laptop steps inside a train. The train's AN offers Internet access service to all the users inside the train. This service is offered through the connectivity and based on the CA established between the ACS of the train and the ACS of the train station. As the train starts to move, its connectivity with this station becomes weaker, but the signal of connectivity with the AN of the next train station starts to get stronger. Due to these events, the train ACS starts a composition process with the next station ACS as a tentative of not losing the connectivity with the Internet. As the train gets closer to the next station, the composition process is already established in order that the users inside the train can continually access the Internet through the next station domain, without losing the connections and connectivity. All this process is transparent to the users inside the train/domain.

In our implementation, this scenario can be translated into a topology like the one represented in Fig. 6. The AN1 (the train) is composed with domain AN2 (the train station domain) and Internet access is provided to AN1 by AN2 as formalized in a CA. The traffic goes through GW_AN1 to GW_AN2, crosses the core (the Internet) and reaches its destination (AN4). In this scenario, the availability of the Internet access service is controlled by a FE, called here Connectivity-FE. When AN1 moves, it connects to AN3 (the domain of the next station). This is done by GW_AN1 connecting with the network of AN3. To keep providing Internet access to its users, the Connectivity-FE realizes that AN1 needs to compose with another AN providing the same service. So, the Connectivity-FE sends to the NAD-FE at AN1 a message informing the desire to find a new Internet access provider. The NAD-FE at AN1 starts to send request messages, asking about Internet access service. The NAD-FE of AN3 receives these messages from NAD-FE of AN1 and, as it offers the asked service, it responds with an advertisement message. As the NAD-FE of AN1 realizes that AN3 provides this service, it sends a message to the Connectivity-FE telling it that AN3 offers the service. The Connectivity-FE sends a message to the C-FE of AN1, informing that a composition with AN3 is necessary. The DEEP modules resolves the name of C-FE of AN3. Both C-FEs (AN1's C-FE and AN3's C-FE) negotiate a CA performing the composition process. In the tests, before the composition process, the GW_AN3 is configured to block the

traffic originated or destinated to/from any domain that is not composed with AN3 or is not represented by a domain that is composed with AN3. During this composition process, this configuration is changed, and GW_AN3 starts to accept traffic to/from AN1 (this is the "commit" step represented by the message number 14 in Fig. 3). After the composition process is finished, all the users inside AN1 can continue using the Internet access, but now through AN3. The traffic goes from GW_AN1 to GW_AN3. All this process is transparent to the users inside AN1.
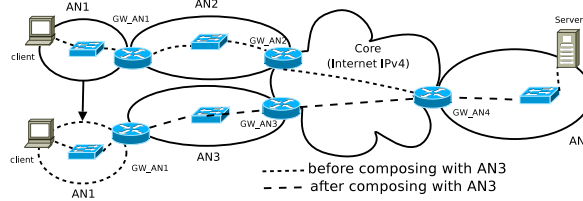

Fig. 6. Scenario 2.

The goal of this test was to analize and evaluate the domain mobility followed by the composition process. By doing this, we could verify the real impact of the composition process as a part of the whole handover process. To evaluate such scenario, we used the JTG traffic generator to send traffic from the server to the client (Fig. 6) and then collect numbers related to the total handover time, time for the composition process, number of packets transmitted, number of packets received and number of packets lost. The traffic consisted in having a 60 seconds UDP transmission. The bandwidth was 20 Mb/s and the size of the packets was set up to 1300 bytes plus headers. The handover was done using a wireless interface in GW_AN1, connecting it from AN2 to AN3. The results are shown in Tab. 1 after collected the results of 5 tests.

Table 1. Results

Test	Pckts. Transm.	Pckts. Recv.	Pckt. Loss	Handover (s)	Composition (ms)
1	113850	108403	5447 (4.78%)	2.70	191
2	113930	107837	6093 (5.35%)	2.57	168
3	113830	109421	4409 (3.87%)	2.31	172
4	113707	108409	5298 (4.66%)	2.34	163
5	113908	110007	3901 (3.43%)	2.05	153
Average	113845	108815.4	5029.6 (4.41%)	2.39	169.4

The total handover time includes the time that is necessary to perform the reconfiguration of the underlying Internet infrastructure [3]. On average, 4.41% of the packets are lost due to a handover of 2.39 seconds. The time of the total composition process is of 169.4 miliseconds on average, included in the total handover time in the wireless link (GW_AN1 goes from its connectivity with AN2 to AN3). This means that total composition process is responsible for only 7.08% of the handover time, causing a loss of ≈352 packets.

The purpose of the two scenarios evaluated above was not to analyze aspects related to scalability neither to stress the architecture with a set of exhaustive tests. The main goal was to evaluate and verify the behaviour of the proposed composition architecture running over a next generation Internet infrastructure.

The numbers collected show that the Web services technology is well indicated as a solution for implementing the composition process. Future works include a deeper evaluation of the Web services and the SOAP message including other tests such as the composition with not adjacent domains.

5 Conclusion

The proposed composition management architecture presented in this work represents a general framework that considers some of the specifications defined by the AN project and contributes to a platform that can be extended to perform other specific functionalities. This paper presented a model for the composition process, implemented using Web services and XML. This model was tested on top of a prototype of a next generation Internet architecture that provides full mobility of nodes and domains.

The results have shown that the composition process is possible and it will be an important role in the next generation of networks. Although it is necessary the establishment of standards to provide real integration between different domains, the use of Web services interfaces is an alternative to speed up the process of standardization becoming a key mechanism to support composition in future communication networks, establishing new ways of relationship between different administrative domains.

References

1. Ambient Networks, *http://www.ambient-networks.org* (accessed in Feb. 13, 2008)
2. C. Perkins. IP Mobility Support. *RFC 2002.* October 1996.
3. W. Wong, R. Villaca, L. B. de Paula, R. Pasquini, F. L. Verdi and M. F. Magalhães. An Architecture for Mobility Support in a Next Generation Internet. *The 22nd IEEE International Conference on Advanced Information, Networking and Applications (AINA 2008).* Okinawa, Japan. March 2008.
4. C. Prehofer, N. Papadoglou, M. Johnsson. On the Evolution of a Naming and Addressing Architecture in view of Dynamic Network Interconnection. *Proceedings of the Third IASTED International Conference on Communications and Computer Networks* Marina del Rey, CA, USA, 2005.
5. V. Ramasubramanian, E. G. Sirer. The Design and Implementation of a Next Generation Name Service for the Internet. *ACM Annual Conference of the Special Interest Group on Data Communication 2004 (SIGCOMM04)* Portland, OR, USA, August 30 - September 03, 2004.
6. W. Adjie-Winoto, E. Schwartz, H. Balakrishnan, J. Lilley. The Design and Implementation of an Intentional Naming System. *SOSP '99: Proceedings of the seventeenth ACM symposium on Operating systems principles* New York, NY, USA, December, 1999.
7. P. Pöyhönen, R. Campos, P. Pääkkönen, N. Akhtar, C. Kappler, D. Zhou. DEEP - A Generic Name Resolution Protocol for Heterogeneous Networks. *Information and Communication Technologies, 2006. ICTTA '06. 2nd* Damascus, Syria, April 24-28, 2006.
8. W. Wong, F. L. Verdi and M. F. Magalhães. A Next Generation Internet Architecture for Mobility and Multi-homing Support. ACM CoNext Student Workshop. New York, NY, December 2007.