ANA Project

Autonomic Network Architecture

Sixth Framework Programme
Priority FP6-2004-IST-4
Situated and Autonomic Communications (SAC)

Project Number: FP6-IST-27489

Deliverable D.4.1

First draft of the testbed architecture

ANA Project

Autonomic Network Architecture

Project Number	FP6-IST-27489
Project Name	ANA - Autonomic Network Architecture
Document Number	FP6-IST-27489/WP4/D.4.1
Document Title	First draft of the testbed architecture
Workpackage	WP4
Editor	Christophe Jelger (UBasel)
Authors	Christophe Jelger (UBasel) Igor Brusic (TA)
Reviewers	David Hutchison (ULanc)
Dissemination level	Public
Contractual delivery date	31^{st} December 2006
Delivery date	12^{th} February 2007
Version	Version 1.0

Abstract:

This document presents a first draft of the testbed architecture that will be deployed to study, analyse and validate the autonomic network architecture currently being developed by the ANA project. The goal of the document is to highlight the main components and organization of the testbed, and to propose milestones for its development and deployment.

Keywords:

ANA, testbed architecture, design/deployment milestones.

Executive Summary

Following the Internet tradition that networking software matures through implementation, we have planned in the ANA project to implement a prototype and deploy a testbed at an early stage. The motivation is that the feedback obtained from preliminary experimental results will be used to steer and refine the architectural design while the project is still active. Inside the four year program we have planned for two such prototyping cycles.

The first prototype will concentrate on the autonomic forming of networks and will take mobility into account. It will be based on the predominant infrastructure of Ethernet switches and wireless access points. The goal is to demonstrate complete self-organization of individual nodes into a network. We target large network meshes in the range of 10^5 active (routing) elements. Obviously, the project consortium alone will not have resources to build a network of 10^5 nodes. In order to demonstrate potential for scalability, four approaches are envisaged: a) overlay for interconnecting the participating sites, b) deployment of multiple virtual nodes per hardware node[1], c) a distributed open collaborative approach similar to successful initiatives such as “SETI@Home” and “PlanetLab”, to include external experimentators and to disseminate ANA results, and d) simulations.

Based on the experiences from the first testbed we will then be able to revise the architecture and extend networking functionality to cover the federation of previously independent networks, both physical as well as virtual networks. In the second testbed, the focus is on the self-organization of networks into a global network, taking into account multiple administrative domains and the cross-domain support for mobile and intermittently connected nodes.

[1]Note that this approach was not listed in the initial description of work of the project. This has been considered during the first year of the project.

Table of contents

1 Introduction

As described in the description of work [1] of the project, the technological objective of ANA is to build an experimental autonomic network architecture in order to demonstrate the feasibility of autonomic networking. For this purpose, and following the Internet tradition that both networking software and paradigms mature through implementation, ANA is organised around two prototyping cycles over the four years of the project. That is, and in contrast to other initiatives for which prototyping is often the closing outcome, in ANA we want to develop a first prototype after the first half of the project.

The feedback obtained from this preliminary phase will be used to steer and refine the architectural design and will eventually lead to the development of an updated and mature prototype. This pragmatic approach is advantageous for two main reasons: first, it permits to validate and upgrade the core architectural concepts and mechanisms while the project is still active, and second, it increases the credibility of the architecture thanks to the early release of a prototype. Another advantage of having an early prototype is that it permits to attract external experimenters during the rise of the project when fruitful collaborations are most likely to take off. We indeed plan to rapidly offer to the research community an open access to the ANA testbed in a similar manner than the PlanetLab [2] initiative. The objective is to turn ANA into an appealing and flexible research platform for wide-scale experimentations in autonomic networking.

The text below, taken from the project's description of work, gives a brief overview of what we want to achieve during the two prototyping cycles of ANA.

> As a first step, a network based on the predominant infrastructure of Ethernet switches and wireless access points will be built. The goal is to demonstrate self-organization of individual nodes into a network. The design of such network should potentially scale to large network meshes in the range of 10^5 active (routing) elements. Obviously, the consortium alone will not have resources to literally build a network of 10^5 nodes. In order to show scalability, three approaches are envisaged: a) overlay for interconnecting the participating sites, b) simulations, and c) a distributed open collaborative approach similar to successful initiatives such as "SETI@Home", "Folding@Home", to include external experimentators and to disseminate ANA results.
>
> The second step, using insights from the first effort, will loosen the constraints and permit wired and multihop wireless heterogeneous devices to be integrated in an autonomic way. Here the focus is on the self-organization of networks into a global network. The rationale for a two phase approach is that an architecture can only be developed and its quality be validated if more than one case is explored.

2 ANA prototyping

2.1 Overview of the testbed

From a practical point of view, the ANA testbed will mainly consist of physical devices (PCs, switches, wireless access points) interconnected by dedicated management links. In a project partner's site, the testbed devices will typically be interconnected by legacy link-layer technologies such as Ethernet or 802.11 (wifi). However having direct access to layer2 links to interconnect different partners' sites is not trivial so we will initially also rely on tunnels over the existing IP infrastructure. Note that although we initially plan to use tunnels to interconnect the different testbed sites, we would be interested to use lower-layer links such as optical lightpaths or long-range Gigabit Ethernet links. The ANA consortium will study this possibility with european NREN (National and Research Educational Networks) although this is not a priority during the first testbed deployment phase where we would already integrate wireless links and bridging technologies locally in partners' sites.

During the initial deployment phase, each partner of the project will host at least one device and some partners are already willing to quickly deploy up to 10 devices. To manage this physical infrastructure, we plan to develop a distributed web-based platform that will permit to check in real-time the status of the devices and the management tunnels interconnecting them. This management platform will also eventually permit to dynamically schedule the creation and removal of experimentation-specific tunnels in order to emulate link failures and topology changes on a per-experimentation basis. The core advantage is that this feature will allow each experiment to run in its private *slice* of the testbed.

In practice, the devices or "ANA hosts" will mainly initially consist of Linux-based or *BSD-based computers running open source software. This alternative provides the highest flexibility in terms of development (e.g. possibility to extend or modify kernel code) and permits to have access to the very large number of free software developed by the open source community. In addition and of very high interest, many devices such as Ethernet switches and wireless access points (e.g. Linksys WRT54GL) now support variants of the Linux operating system which are being actively maintained. This is clearly an interesting opportunity which permits to expand and diversify the ANA testbed with off-the-shelf network devices other than traditional computers. It is an ambitious development challenge since the ANA software will have to be able to run on devices with low CPU power and limited storage capabilities.

Note that including switches and wireless devices in the ANA testbed is essential to fulfill the prototyping objectives as described in the project's description of work. The main goal of the first development phase, beside the validation of the core architectural concepts of ANA, is indeed to demonstrate the self-organization of individual nodes into a network. Clearly with ANA-enabled switches and wireless devices, we will be able to demonstrate the self-organizing principles of ANA components in bridged networks with wired and wireless hosts but also in e.g., wireless multi-hop mesh networks. Note that thanks to the ability to create virtual links between distant sites of the testbed, it will also be possible to introduce more heterogeneity with e.g. various link delays and bandwidths. It is worth mentioning here that while the number of physical devices in the first core testbed may well remain below 50, we will introduce the possibility to run multiple instances of the ANA node software on a given physical device. This extra feature will permit to greatly increase the number of (virtual) nodes in the testbed and overcome the inherent technical limitations of deploying a large number of physical devices. This extension is described in more details in section 2.2.

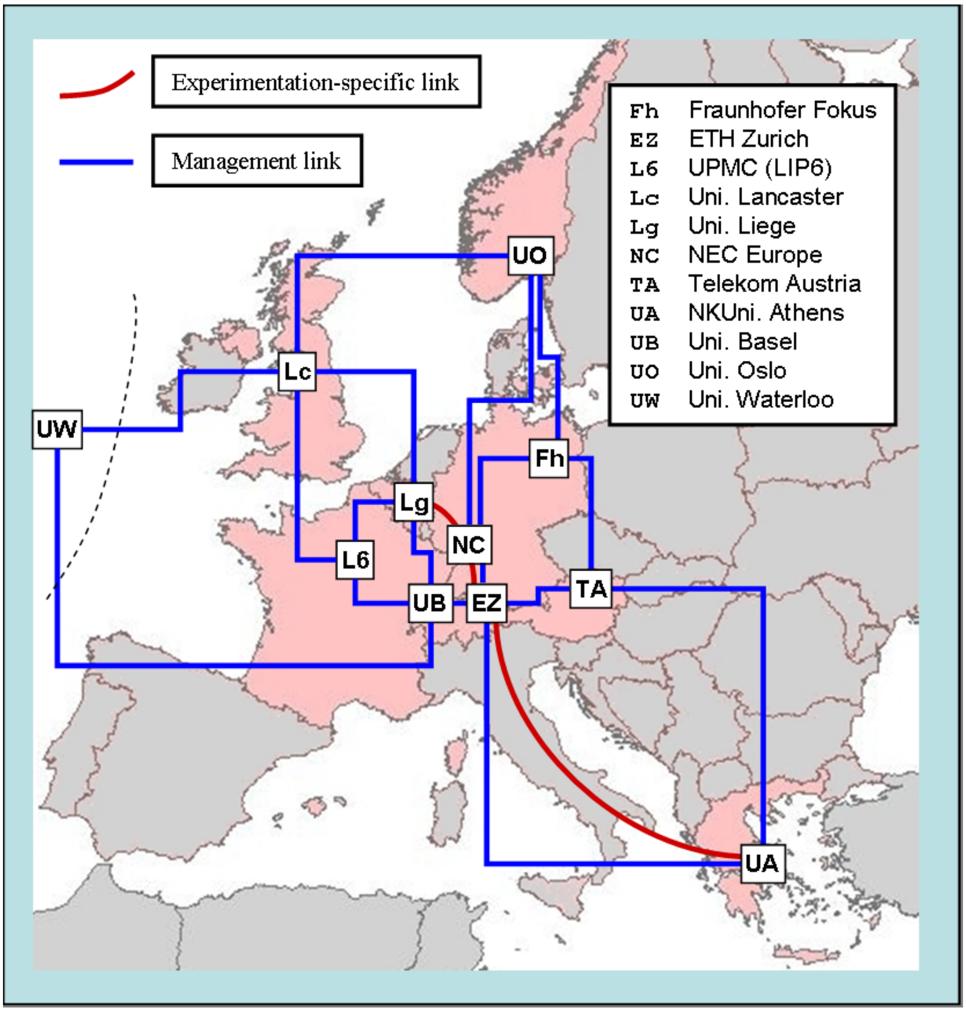

Figure 1: ANA testbed: example topology.

Figure 1 illustrates a possible topology for the ANA testbed. It shows a number of management links (e.g. IP tunnels) that connect the hosts located in the project partners premises. These links will not be used for experimentations but solely for the daily management and operation of the testbed infrastructure (i.e. to get a remote access to all the hosts of the testbed). Also routing on the management links of the testbed could be performed via an existing protocol such as RIP. For experimentations, we envisage the possibility to create dedicated links between particular ANA nodes that could e.g., be turned up and down on-demand in order to simulate topology changes. The goal is that each experimenter can create its own *slice* in the testbed on top of the shared testbed infrastructure.

2.2 Network emulation

In order to increase the number of nodes in the ANA testbed and to provide more flexibility in its deployment, we plan to develop the prototype software in a way such that it will be possible to run multiple ANA nodes on the same physical device. This feature will allow an experimenter to run multiple experimentations in parallel in a similar manner as PlanetLab allows multiple *slices* to coexist on the same set of physical devices.

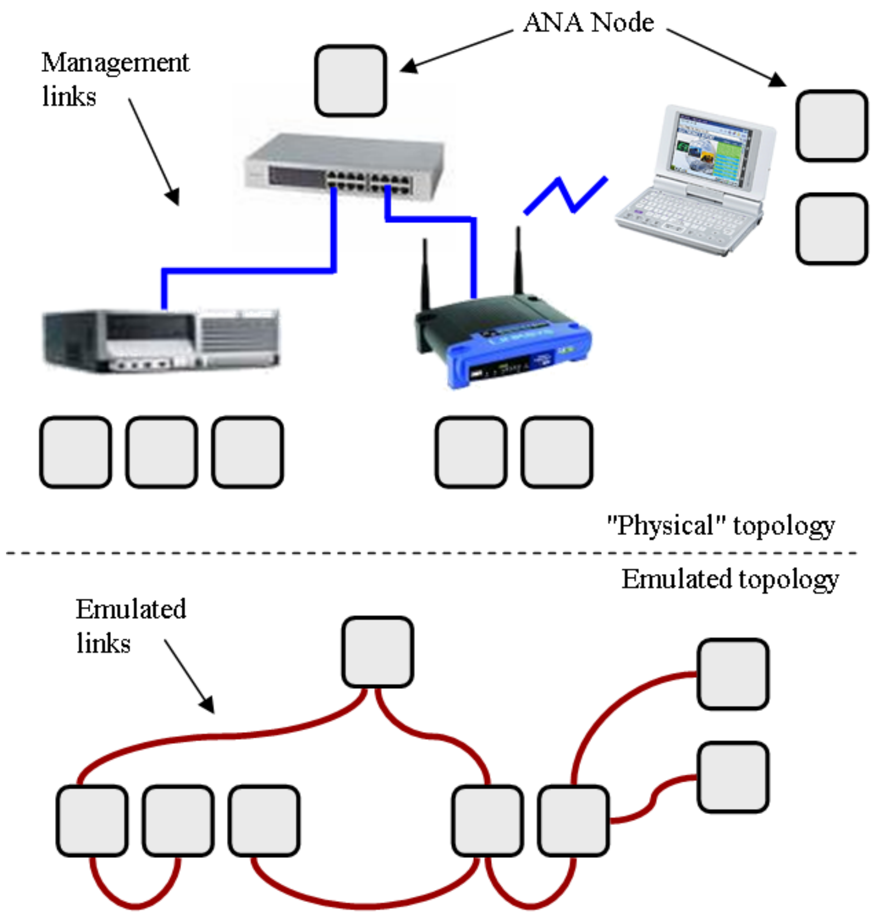

Figure 2: Network emulation in ANA.

In addition, the ANA prototype will permit the arbitrary creation of virtual links among the multiple ANA nodes running on a given device. With this powerful feature, it will be possible to emulate any network topology without any restriction imposed by the underlying connectivity between the physical devices of the testbed. Moreover, we want to extend the possibility of creating virtual links between ANA nodes running on distant devices: this will for example be useful to emulate redundant links and dynamically introduce link failures in order to e.g. study the self-healing performance of routing protocols or distributed monitoring. The emulation concept is illustrated by Figure 2 showing four devices running eight ANA nodes.

2.3 Development and integration strategy

A central and challenging task in the development of a prototype is the integration of the various modules being developed by the different partners of the project. In ANA, we envisage to handle the integration process by introducing a main integration site/partner that will be responsible to centralise stable developments. That is, while the development and validation of a given module and its integration with the existing stable code will be handled separately by the responsible partner(s), the stable modules will be centralised at the integration site which will then be in charge of periodically distributing the stable software to all the other ANA sites.

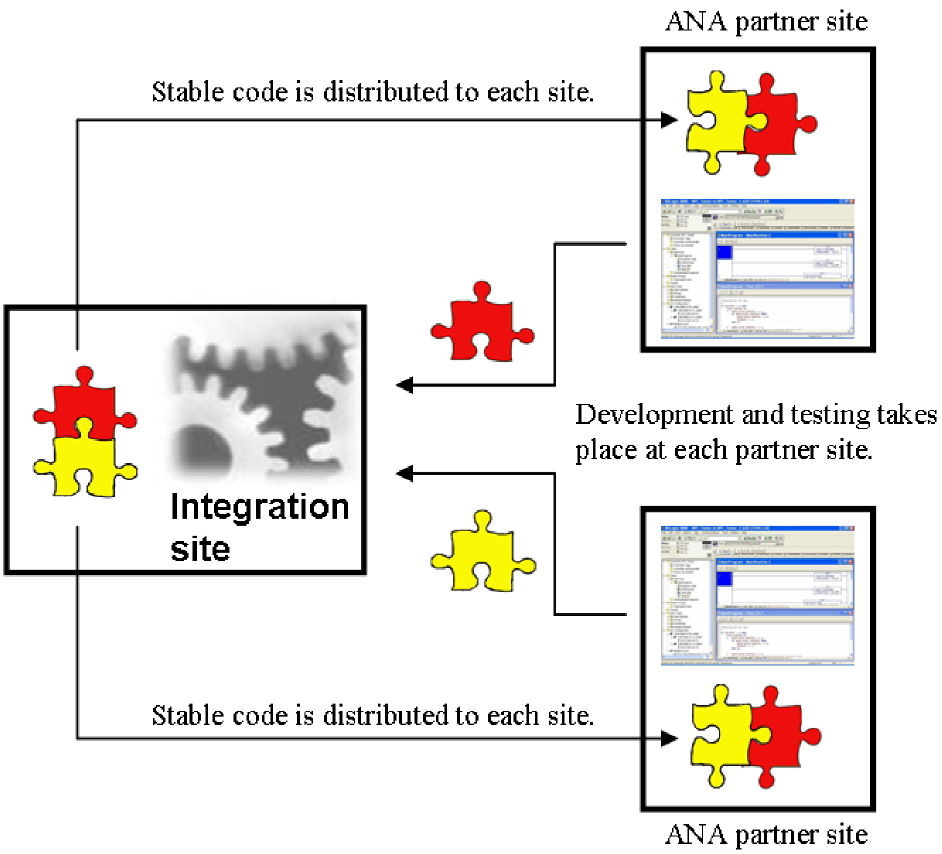

Figure 3: Integration procedure of ANA prototype.

The partner responsible for the integration process will also maintain a set of devices running the current stable code in order to detect integration problems. In order to minimise the occurrence of such problems, each partner will be responsible for developing a test procedure that will stress its code. Because typically such tests are successful when the experimenter is also the developer of the tests, we will require that a module is cross-tested by at least one other partner before being integrated in the stable release. Note that to minimise the overhead for the partner performing this extra validation step, the developer will be responsible to provide the entire test procedure and code.

Finally, the core software of ANA such as the MINMEX and the abtraction and adaptation layers (see Blueprint deliverable) will require more careful testing and validation before a new version is integrated in the stable release. For this purpose, we foresee a strict validation procedure to ensure that the new core software does not introduce any backwards compatibility problems. The details of this procedure will be specified at the beginning of the active prototyping phase. Figure 3 illustrates the integration procedure of the ANA prototype.

2.4 Assessment and methodology

A critical aspect of the prototyping phase is to assess the operation of the enabling components of ANA, namely the core ANA software and the underlying testbed system. All the developments demonstrating advanced autonomic properties indeed rely on both the ANA Node and the testbed and hence cannot be accurately evaluated if the behavior of these two core components is not deterministic. We believe that such a "quality control" procedure is critical in order to improve the credibility and reliability of the software developed by the ANA project.

In order to perform this assessment, the consortium plans to reuse the expertise of Telekom Austria (TA) as a major network operator in Europe. TA is indeed used to performing a careful and systematic testing of network equipments and software before they are deployed in TA's production network. TA will use similar methodologies and procedures to validate the operation, on one side, of the core ANA node software, and on the other side, of the testbed infrastructure and software. Shall these tests permit to detect faults or inconsistent operation, precise feedback will be provided to the appropriate implementation team.

TA will also perform "stress tests" of both the ANA node and the testbed. That is, prototypes will be excessively loaded (e.g., with traffic, control operations, etc) in order to "push them" to operational limits where unexpected failures may occur. This procedure will ensure that the software developed by ANA will be reliable when performing intensive experiments involving e.g., a large number of ANA nodes or generating a high traffic load. Note that this second testing phase will only start when the basic operation of the core components will have been validated.

A third area where the expertise of Telekom Austria will be used concerns the management tools of the testbed. As a large-scale network operator, TA is familiar with evolved network management tools and will hence be able to provide valuable input and feedback with respect to the specification, design, and refinement of the management features of the ANA testbed. Note that we foresee this activity to become really effective during the first intensive testing phase that will start in 2008 where the testbed management tools will be intensely used on a daily basis.

3 Joining and linking the ANA testbed

3.1 Public access and ANA@Home

A recurrent problem in networking research is that prototype testbeds do not reflect the reality of production networks that have to deal with a large number of users and heterogeneous devices and usage patterns. Note that a typical business-case for the first prototyping cycle of ANA can be a mid-size company with a few hundred devices (including end-user computers) up to a university network with a few thousand devices.

Clearly with the emulation feature described in section 2.2, the ANA consortium could experiment with networks up to e.g. one thousand devices, with for example 10 devices per partner with each device running 10 ANA nodes. Note that while such a network size is feasible, the management overhead (to create virtual nodes and links) might already become too large and would prevent easy manual re-configuration of the emulated topology. A solution to this problem is to develop an automatic procedure to create virtual nodes and links according to some algorithms that would generate a topology with tunable parameters such as e.g. a certain power-law connectivity distribution or the model of a multi-tier hierarchy. Note that we plan to provide an experimental version of this feature in the first version of the prototype.

However, the ANA consortium alone will not have the resources to build and manage a network larger than 1,000 (virtual) nodes. In addition, while the emulation scheme is an appealing solution to increase the size of the network and validate the operation of protocols, it can only mimic end-user behaviors and may not accurately reflect the network traffic patterns caused by real users. It might hence not be possible to study the performance of experimental protocols with a realistic network load.

To overcome these two limitations, we plan to develop a distributed and open collaborative approach similar to successful initiatives such as “SETI@Home” and peer-to-peer networks such as Skype or eMule which attract a large number of participants. With this initiative named “ANA@Home”, we plan to offer to the research community an open access to the ANA testbed in a similar manner than the current PlanetLab [2] initiative and the upcoming GENI [3] infrastructure. The objective is to turn ANA into an appealing and flexible research platform for wide-scale experimentations in autonomic networking. While the ANA@Home software will mainly be a standard ANA Node, the initiative will investigate and develop the open platform required to first incorporate external participants as "satellites" of the core ANA testbed. The basic idea is to provide ANA-POPs (Points of Presence) via which users running ANA@Home could join the main testbed. During the second prototyping phase when more experience will have been gained, we could provide direct access to the main testbed if this appears a reasonable and technically sound choice.

In addition to purely technical issues, ANA@home introduces two extra challenges. First, we must develop an appealing usage model and propose innovative features that will efficiently attract external participants. Second, we must ensure that the resources of the testbed are fairly shared among the participants. These two aspects are detailed in the follwing two paragraphs.

A key observation is that while the research community is clearly interested to have a large testbed for networking research, most researchers have little time to configure and maintain the devices and software needed to join such an initiative. For example for many researchers, the overhead of registering, configuring, and maintaining PlanetLab nodes (a customized Linux distribution) may be an inhibitory entry barrier to PlanetLab. In ANA we believe that the key feature of ANA@Home must be simplicity of installation. First it must be straightforward for a novice ANA user to install the software on top of existing systems (e.g. Linux) and second, the user interface must be intuitive enough such that using the basic features does not require a discouraging and excessive learning phase. It is indeed well known that while using a software for the first time, the first minutes are critical for a user (albeit a motivated one) to decide whether this is worth spending more time or not.

The second important aspect for ANA@Home to be successful is that the resources of the testbed must be fairly shared among all the participants. We here envisage to implement a scheme where, during peak usage times, the share of resources allocated to a given user is proportional to the resources (e.g. number of ANA nodes) that this user has offered over a certain period of time up to the present time. This share model is hence also an incentive for users to continuously provide resources (ANA nodes) eventhough they may not actively perform any experimentation. The goal is to prevent selfish behaviors where users only activate ANA nodes when they want to run an experimentation.

3.2 Interconnection with other testbeds and networks

While this is not a priority of the first prototyping phase, the mature testbed developed during the second prototyping phase will eventually interconnect with other research infrastructures such as e.g. PlanetLab, future projects such as FIRE [4] and GENI (but not in the near future), but also with the legacy Internet. As described in one of the Blueprint deliverables of the project (i.e. D.1.4), the core architectural concept of ANA is the network compartment, i.e. a powerful abstraction which offers a kind of *wrapper mechanism* that allows accommodation of legacy network technologies within ANA. This enables the formulation of network architecture transition strategies, for both backwards compatibility as well as the hosting of future network instances in ANA. The compartment abstraction also serves as the unit for the federation of networks into global-scale communication systems. For example, interworking across heterogeneous network domains can be achieved by deploying a common overlay compartment on top of the disparate compartments that must be federated.

In a similar way as GENI envisages to offer a generic hardware *substrate* upon which multiple networks can be built, ANA offers a generic and flexible software substrate that can host and federate multiple network compartments. That is, ANA provides the least common denominator technologies upon which heterogeneous network compartments run and interconnect. Since the ANA software will eventually be able to run on top of any hardware platform, we envisage that it could be used on top of a future research infrastructure such as GENI, and could interconnect with the legacy Internet via the deployment of an "IP compartment" inside ANA. Note that other research initiatives and prototypes such as e.g. Haggle [5] and Cascadas [6] could potentially also be executed on top of ANA if a suitable compartment is developed.

4 Development and deployment milestones

This section presents a list of estimated milestones for both the prototype development and the testbed deployment for the second year of the project. Note that there are four main activities:

- The development of the testbed, the management tools, and the deployment of the testbed.
- The development of the ANA node prototype.
- The development of elaborated compartment modules.
- The development of ANA@Home.

The list below shows the estimated milestones that the project would like to reach for the four activities listed above. Note that these dates are indicative and may change during the course of the project. We however expect that the consortium will be able to mobilize the necessary resources in order to achieve these challenging goals in time.

- February 2007
  - Development of ANA node software starts.

  - Development of testbed software starts.

- April 2007
  - Hosts deployment in partners premises starts. Creation of IP tunnels.
  - Development of the management tools for the testbed infrastructure starts. E.g. web interface (connectivity status, hosts status, etc).

- July 2007
  - First experimental version of the ANA node software available.
  - Active development of elaborated compartment modules (functional blocks) starts. E.g. routing, service discovery, monitoring.

- September 2007
  - A first experimental compartment with basic functionalities is ready. It will serve as an early example for further developments and to already refine development templates and tools.
  - Specific development for ANA@Home initiative starts.
  - First testbed management suite ready. Experimental testbed is ready for first experiments.

- January 2008
  - Integration of elaborated compartment modules starts.
  - First experimental version of ANA@Home available.

- March 2008
  - First integrated version of testbed, management tools, ANA node, and ANA@Home is ready.
  - Evaluation phase begins until July 2008 where core architectural concepts start being refined.

5 Conclusion

Prototyping and testbed deployment are two crucial activities for the second year of the ANA project. The key objective is on one side, to provide a least common denominator architecture upon which different network compartments can be executed, and on the other side, to develop the substrate upon which autonomic networking protocols and algorithms will be executed and evaluated. The prototyping phase is clearly challenging in ANA, especially since we will also have to offer a public (but initially restricted) access to the testbed. However shall this openess strategy be succesful, we may benefit in the future from efforts and developments provided by external participants. We have to make sure we do not miss such a stimulating opportunity.

References

[1] Autonomic Network Architecture. Integrated Project FP6-IST-27489. Sixth Framework Programme - Situated and Autonomic Communications (SAC), October 2005.

[2] PlanetLab - http://www.planet-lab.org/.

[3] GENI - http://www.geni.net/.

[4] Future Internet Research and Experimentation (FIRE) initiative - http://cordis.europa.eu/ist/fet/comms-fire.htm.

[5] Haggle Project - http://www.haggleproject.org/.

[6] Cascadas Project - http://www.cascadas-project.org/.