

Design and evaluation of dynamic policy-based flow redirection for multihomed mobile networks

Tobias Hof

System Engineering & Architectures
Thales Communications S.A.
92704 Colombes, France
tobias.hof@fr.thalesgroup.com

Qi Wang, Robert Atkinson, John Dunlop

Mobile Communications Group
University of Strathclyde
Glasgow G1 1XW, UK
{qwang, ratkinso, j.dunlop}@eee.strath.ac.uk

Abstract—This paper presents the design, implementation and evaluation of a solution for dynamic redirection of traffic flows for multihomed mobile networks. The solution was developed for a mobile user that disposes of a Personal Area Network (PAN) with a Personal Mobile Router (PMR), in order to achieve Always Best Connected (ABC) service by distributing flows belonging to different applications among the most appropriate access networks. Designed in a modular way for a NEMO-based mobility and multihoming support, the proposed flow redirection solution can be easily coupled with and controlled by dynamic traffic policies that come from advanced network intelligence, according to the currently available network resources and user and application requirements. A prototype implementation was validated and assessed on a testbed as proof-of-concept.

Keywords: *NEMO; multihoming; MCoA; dynamic flow handoff; traffic policies, Always Best Connected*

I. INTRODUCTION

Next-generation wireless communication systems will incorporate multiple heterogeneous access networks such as cellular networks, wireless local area networks and WiMAX into a uniform IP-based platform, evolving towards IPv6. In order to provide Always Best Connected (ABC) service, research activities now focus on multimode devices with several interfaces, enabling mobile users to be multihomed in multiple heterogeneous access networks. Consequently, service provisioning can be enhanced by routing traffic flows of different applications to the most appropriate access networks, where the term *appropriate* can be interpreted in different ways according to the various ABC metrics such as shortest delay or high bandwidth, but also minimal cost.

In the European project (IST-FP6) MULTINET [1][2], we address this challenge of advanced service delivery for nomadic users by building a complete system that integrates monitoring of access networks, measurement-based generation of ABC policies to map traffic flows to appropriate radio access networks and the actual redirection of the flow data packets to the desired access routers.

MULTINET targets mobile workers in charge of the

This work has been partly funded by the European research project *MULTINET: Enabler for Next Generation Service Delivery* (IST FP6 No. IST-2005-027437). The authors would like to thank all the MULTINET project partners for their contributions during the development of various ideas presented in this paper

maintenance of high-tech and complex machines all over the world. In this context, a worker is equipped with a range of devices such as laptop, mobile phone, webcam or sensors of any kind, all of which form a Personal Area Network (PAN). Typically, these workers need to connect to their back office regularly for remote support, transmitting data or downloading instructions. It is obvious that the devices in the PAN generate multiple and manifold data flows with different importance and QoS needs. With the increasing deployment of wired and wireless overlay networks, the mobile workers can very likely find themselves in the coverage of multiple access networks.

Based on the above scenario, MULTINET develops a communication system that provides simultaneous, intelligent and traffic-dependant usage of multiple access networks for mobile users carrying a PAN. Thus, the MULTINET mobility support is based on the Network Mobility (NEMO) Basic Support Protocol [3], with the PAN being connected via a Personal Mobile Router (PMR), and on related IETF drafts for multihoming support. Additionally, one of the major innovations of MULTINET is the dynamic redirection of traffic flows to appropriate access networks based on traffic policies in order to take advantage of the multihoming support.

The present paper focuses on this flow redirection functionality and presents the design, implementation and evaluation of the proposed solution, especially the modules for the reception and enforcement of traffic policies for a Linux-based NEMO implementation.

The remainder of the paper is organized as follows: Section II introduces the overall MULTINET system as the context of the presented work and the problem statement. Then, in section III, after a brief review of the state-of-the-art of handling multihoming and flow policies for mobile devices, we present our solution for policy-based, intelligent multihoming and traffic flow management. Section IV describes the proof-of-concept implementation, together with experimental results. The paper is concluded in section V.

II. OVERALL PRESENTATION OF THE MULTINET SYSTEM

MULTINET intends to bring advanced connectivity functions to next-generation services and to optimize communications in multiple dimensions such as cost, quality and functionality. Therefore, the MULTINET communication system will be capable of providing the required network and application functionalities to handle multiple wireless access

networks seamlessly and simultaneously. Consequently, a comprehensive chain has been designed for a systematic support of advanced service delivery based on the decision about the best usage of the available wireless access networks.

As can be seen in Figure 1, the first step is to monitor and measure the state of the access networks. This can be accomplished jointly by the PMR and the access points or base stations, as each of them can have access to different kinds of information concerning the condition of the radio link.

In the next step, the collected dynamic information is processed with the user preferences, the applications' characteristics and the operator constraints by an intelligent Network Selection Algorithm (NSA) to make a decision whether a flow redirection should be triggered.

It can be deduced from Figure 1 that selecting the "best" access network for a given traffic may not necessarily mean an optimization in terms of purely technical QoS parameters, but can also stand for cost reduction or security level selection, while covering the respective operator rules.

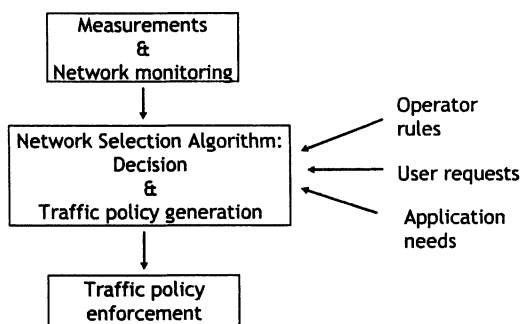


Figure 1. Information and decision chain of the system

If a flow redirection is needed, the NSA generates new flow binding policies (i.e., traffic policies) for the application flows over the active access networks. The NSA then sends these policies to the Home Agent (HA) and maybe also to the PMR (depending on implementation choices; see Section III B). The HA and the PMR then enforce the policies, i.e. redirect the traffic flows according to the policies.

This crucial step of enforcing the policies is executed by the Policy Reception and Enforcement Modules (PREM) that has been developed during the project and constitutes the focus of the remainder of this paper.

III. STATE OF THE ART AND DESCRIPTION OF THE PROPOSED SOLUTION

A. Current standardization activities for mobility, multihoming and traffic flow redirection

Mobile IPv6 (MIPv6) [4] has enabled connection continuity for Mobile Nodes (MN) visiting foreign networks. A topologically correct Care-of Address (CoA) is used for traffic rerouting, and a long-term home address (HoA) is used as the host identifier. NEMO [3] enhances the mobility functionality by moving it from a mobile node to a mobile router that serves as the gateway for all the mobile network nodes (MNNs).

Furthermore, the Multiple CoAs draft [5] proposes a MIPv6 and NEMO extension that allows a mobile node (either a mobile host or a mobile router) to register multiple CoAs for a single HoA, and thus enables the potential simultaneous use of the multiple interfaces for a multihomed mobile node. A Binding Unique Identifier (BID) is introduced to distinguish the multiple pairs of (HoA, CoAn) registrations.

The Flow Binding draft [6] defines options to identify flows, e.g. through the five tuple {source address, source port, destination address, destination port, protocol}, and to specify flow binding policies (a binding between the identified flows and appropriate interfaces represented by either CoAs or BIDs) in the Binding Update message sent from the mobile node to its HA. Consequently, this approach supposes that flow redirections are initiated exclusively by the mobile node.

Compared with [6], the Flow Distribution draft [7] decouples flow binding policies from MIPv6 and NEMO mobility messages and thus flow redirections can be initiated in a more flexible way since either a mobile host, the HA or yet another entity can generate flow binding policies. Moreover, the XML-based policy code is easier to analyze and extend in contrast to binary code used in mobility messages. Nevertheless, considerable investigations and developments are still needed to substantiate the proposed framework.

Please note that, although MULTINET refers explicitly to a PMR, the concepts presented in this paper are valid for all NEMO usages, therefore the general term Mobile Router (MR) will be used in the remainder of this paper.

B. The MULTINET approach for flow handling

Based on the state-of-the-art as summarized above, in MULTINET we have chosen NEMO [3], the Multiple CoA draft [5] and the Flow Distribution draft [7] as theoretical building blocks to achieve the desired flow redirection functionality, along with further design implementation considerations. In this subsection, we concentrate on the introduction of the core signaling and operation procedures for network-controlled policy-based flow redirection. We have analyzed the following two schemes:

In the first scheme, the NSA generates the flow binding policies and sends the downlink and uplink policies in a message to the HA and the MR, respectively. The message sent to the HA contains the downlink policies for traffic addressed to the mobile network, and the one for the MR contains uplink policies for outgoing traffic from the mobile network. The policies are then enforced by the PREM at the HA and the MR, respectively. (The detailed design of the PREM is presented in Section IV.) Meanwhile, an acknowledgement is sent to the NSA. Since the uplink and downlink policies can be same or different, this scheme suits both symmetric and asymmetric policies between the HA and the MR.

Alternatively, the second scheme assumes symmetric policies and employs a different signaling procedure. Once the downlink policies are generated by the NSA, a message is sent from the NSA to the HA only; subsequently the HA maps symmetric policies for the uplink traffic flows and sends a resultant message to the MR. This scheme requires an

additional (but fairly simple) functionality for policy mapping at the HA and is limited to symmetric policies. On the other hand, it can reduce the policy generation cost (in terms of processing time and computational consumption) at the NSA.

Figure 2 illustrates the first scheme which was finally chosen for MULTINET. The XML-coded policies are enclosed in SOAP (Single Object Access Protocol) messages over either HTTP or directly on UDP, depending on the implementation (the HTTP case is shown here).

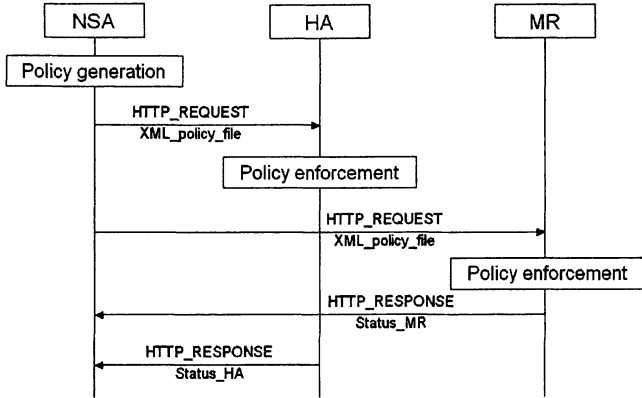


Figure 2. Flow redirection signalling and operations

IV. PROOF-OF-CONCEPT IMPLEMENTATION AND EXPERIMENTAL RESULTS

In this section, we present the actual implementation of the PREM as complementary module of an existing NEMO implementation as well as its main functioning principles. Furthermore, we present some first evaluation results of this prototype implementation. The PREM was designed to run on the HA as well as on the MR.

A. Design of the implementation

The flow redirection execution functionality was designed to collaborate with a NEMO implementation with integrated MCoA support from Nautilus6 [8]. This implementation consists of both user-space modifications of the basic NEMO implementation for Linux (NEPL, [9]) and a kernel patch.

With respect to the MCoA-enabled NEPL implementation, the PREM represents a complementary module as it enables advanced use of the multiple CoAs. From the application traffic point of view, the multiple tunnels between the HA and the MR can be seen as multiple bit pipes provided by MCoA whereas the PREM allows the dynamic mapping of traffic flows to these bit pipes.

The design of our implementation and the operation mode was kept align with NEPL to facilitate the deployment of these two modules in parallel. Consequently, the PREM is implemented as a Linux daemon in C. The policies are expressed in XML according to the syntax proposed in [7] and handled using the libxml2 library. The policy transmission from the NSA to the HA and the MR is based on a web service. The web service modules for client (NSA) and server (HA and MR) are based on the gSOAP library for C++ [10].

The operation cycle of the PREM is shown in Figure 3. Upon receiving a SOAP message containing policies, the PREM parses the message and extracts the policies and the associated actions. Finally, the updated policies are enforced and ongoing or future application flows that are identified in the policies are forwarded over the desired interfaces. Currently, the two basic actions, policy insertion and deletion, have been implemented, nevertheless additional actions such as policy replacement are considered to be added later.




Figure 3. Policy reception and enforcement module.

The policy parsing engine is implemented in a modular way to respect the structure of XML-coded messages. The PREM is thus easily extensible or modifiable in case of the addition or modification of options in the flow definition syntax.

B. Policy processing and interaction with mobility support

Figure 4 shows in detail the complementarity of the PREM regarding its interaction with NEPL. Moreover, the figure indicates the compatible design style of the two daemons regarding their interaction with the native Linux routing engine and the native packet filter, respectively, in the control plane.

The user-space parts of NEPL and PREM are only involved in the handling of signaling messages, i.e. parsing the essential information and passing it to the relevant kernel-space modules, using native system calls or command line calls.

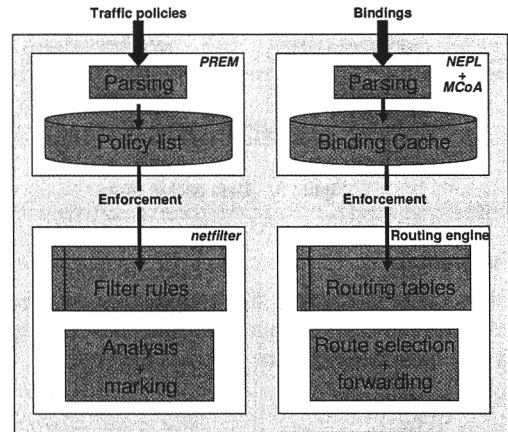


Figure 4. Reception and processing of flow handoff messages

Binding Updates with the BID sub-option are parsed and validated by the MCoA-enabled NEPL daemon. The resulting tunnels are mounted and the corresponding routes to the Mobile Router and the Mobile Network Prefixes are set up into the Linux routing table whose number corresponds to the BID. This means that in contrast to the basic NEPL implementation,

where all NEMO-related routes are installed in the main routing table, the MCoA implementation activates one routing table per BID, which is directly addressed by the BID number.

With regard to the PREM, it enforces the policies by passing them over to the native Linux IPv6 packet filter *netfilter* through its command line interface *iptables*. To do so, the PREM parses incoming policy files and constructs the command for *iptables* with the correct syntax.

C. Forwarding of data packets –redirection execution

Regarding the user data, incoming packets are analyzed, marked and forwarded according to the currently active filter rules and bindings as can be seen in Figure 5. *netfilter* analyses incoming packets and if they match one of the traffic rules in the PREROUTING chain of *netfilter*, they are marked with the corresponding BID number, using the MARK target. Thanks to the kernel patch provided by the MCoA implementation, this marker can be taken into account for the packet forwarding and the packet can be redirected to the routing table whose index is the BID. Consequently, the route lookup is done on that dedicated routing table and the packets are forwarded to the desired CoA. Note that on Home Agents, there can be multiple routes for the same BID in a given routing table, as different MRs might use the same BID. Nevertheless, the packets are forwarded correctly by using the destination address as search key amongst the different routes for the same BID.

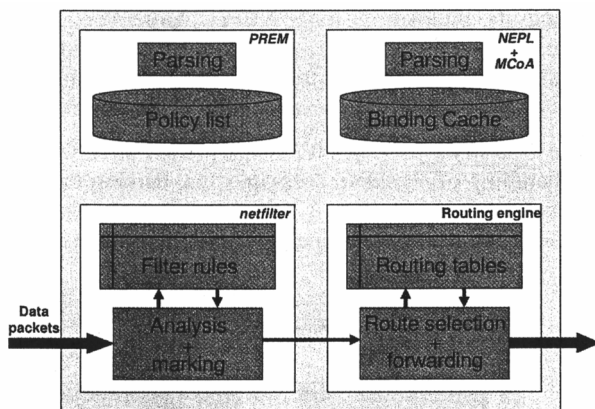


Figure 5. Data packet flow

D. Evaluation

This section introduces the architecture of the testbed that was used for the experiments and presents first experimental results that have been obtained on this testbed.

1) Testbed architecture

Figure 6 illustrates the testbed that allowed us to test the PREM implementation and to validate the interaction with the mobility and multihoming support. Moreover, the testbed permitted preliminary measurements on implementation performances. Currently, the multihomed MR is connected to the Access Routers (ARs) through two Ethernet cables. During the integration phase of the MULTINET project, this will be ported to an evolved wireless testbed (WLAN first, then WiMax and UMTS), although the current setup is sufficient for proof-of-concept of this IP-layer based protocol.




Figure 6. Testbed architecture

For the validation of the PREM, an NSA stub was realized, providing a GUI for policy generation and policy transmission to the PREM instances on the HA and the MR.

2) Comprehensive, qualitative handover scenario

To illustrate the flow redirection in a qualitative way from the point of view of the MR, we show a case study with several flow redirections of a HTTP-based video stream in Figure 7. Using the software Ethereal, we represent the IPv6 packets passing through the two egress interfaces of the MR. The time is expressed in seconds on the X axis, and the traffic is expressed in number of packets on the Y axis. For scalability reasons, the first traffic burst on eth0 has been cropped.

By analyzing Figure 7 step by step, we explain the different flow redirection processes. The video streaming traffic is sent from an HTTP server located on AR_3 to MNN1 in the PAN, and the TCP acknowledgements (ACKs) are returned by the MNN. In addition, we can see the signaling traffic whose volume is very low compared to the HTTP streaming and the ACKs, as expected. Signaling traffic is observed on both interfaces as it includes the periodical Router Advertisements sent by the ARs as well as the periodical Binding Update and Binding Acknowledgement messages.

At the beginning of the transmission (t_0), all traffic is sent over eth0 as it has been set as default egress interface of the MR. At t_1 , the policy generation engine sends a trigger message to the HA and to the MR with the new policies that binds the HTTP traffic to the interface eth1. The SOAP message that contains the traffic policies is also sent through eth0 since there is no explicit policy for this type of traffic.

At t_2 and then at t_3 , the policies are enforced at the HA and the MR, respectively, and the HTTP packets as well as the ACKs are redirected to eth1 accordingly. As the trigger message is sent to the MR via the HA, it arrives with an additional delay in the MR. Consequently, the ACKs routed by the MR are redirected later than the incoming HTTP packets routed by the HA. However, this temporarily unsymmetrical traffic flow does not impact the end-to-end TCP connection, as neither the original IP header nor the TCP payload is modified and all packets always go through the HA and the MR.

At t_4 , eth1 performs a handover from one access network to another. During the handover, the traffic is automatically redirected to eth0 in order to reduce the handover latency and to reduce its impact on the perceived QoS for the user. At t_5 , this interface handoff is completed and the traffic flows are switched back to eth1 according to the currently valid policies.

At t_6 and t_7 , all flow binding policies are deleted upon the reception of a correspondent traffic policy. Consequently, the packets are switched to the default interface eth0 again.




Figure 7. Qualitative illustration of a flow handover scenario

This comprehensive scenario validates the PREM’s capability to dynamically redirect network flows over multiple interfaces of the MR according to traffic policies and also the possibility to use one egress interface as backup for another. From a user’s point of view, no perceptible disruptions of the video were experienced during the whole experiment.

3) Policy processing performance

To further assess the performance of the PREM, the policy processing time was measured and illustrated in Figure 8.




Figure 8. Effect of number of policies on processing time

The measurements were gathered from a Pentium M (1.73 GHz, 2GB RAM) laptop acting as MR. Two to ten homogeneous policies in the form of a protocol number

mapped to a BID were applied to exclude the effects of difference in policy complexity on the delay. As shown in Figure 8, the resulting processing time of the PREM increases in a roughly linear way with the number of policies that are sent in the SOAP message. The reason for this linearity is that, for each policy, the PREM builds a separate command, which is passed to *iptables* for the policy enforcement.

A complete theoretical analysis of the different delays in the whole process as shown in Figure 1 can be found in [11].

V. CONCLUSION AND OUTLOOK

In this paper, we presented a policy-based solution that can dynamically redirect traffic flows generated from or destined to a multihomed mobile network over multiple access networks. Then, we showed the implementation of an original Policy Reception and Enforcement Module (PREM), designed for a NEMO-based approach. Finally, experimental results are given that were obtained during the validation of the implementation on a proof-of-concept testbed.

The next validation step of the PREM is its integration in the complete MULTINET platform, featuring the signaling chain as shown in Figure 1 with an NSA, different wireless access technologies, applications and all other elements.

Future work is the integration of the PREM on the ANEMONE testbed [12] where it can be validated at a large-scale and where its performances will be compared with the ongoing implementation of an alternative approach [6].

It is expected that the proposed dynamic flow redirection mechanisms, integrated with QoS measurement means and intelligent network selection algorithms, constitute the base for providing efficient ABC services as a commercial offer.

REFERENCES

- [1] Lazaro, O., Aginako, L., Hof, T. et al. (November 2006). *MULTINET: enabler for next generation services*, proceedings of Wireless World Research Forum (WWRF) 17, Heidelberg (Germany).
- [2] MULTINET project web site: <http://www.ist-multinet.org/>.
- [3] Devarapalli, V., Wakikawa, R., Petrescu, A. & Thubert, P. (2005). Network Mobility (NEMO) basic support protocol. *IETF RFC 3963*.
- [4] Johnson, D. B., Perkins, C. & Arkko, J. (2004). Mobility support in IPv6. *IETF RFC 3775*.
- [5] Wakikawa, R., Ernst, T. & Nagami, K. (2006). *Multiple care-of addresses registration*. IETF Internet Draft, <draft-wakikawa-mobileip-multiplecoa-05.txt>, work in progress.
- [6] Soliman, H., Montavont, N., Fikouras, N. & Kuladinithi, K. (2006). *Flow bindings in mobile IPv6*. IETF Internet Draft, <draft-soliman-monami6-flow-binding-02.txt>, work in progress.
- [7] Mitsuya, K., Tasaka, K. & Wakikawa, R. (2006). *A schema fragment for flow distribution*. IETF Internet Draft, <draft-mitsuya-monami6-flow-distribution-00.txt>, work in progress.
- [8] Nautilus6 project: <http://www.nautilus6.org/nemo/>.
- [9] Nemo Platform for Linux (NEPL): <http://mobile-ipv6.org>.
- [10] gSOAP: C/C++ library for Web Service servers and clients: <http://gsoap2.sourceforge.net/>.
- [11] Wang, Q., Hof, T. et al. *QoS-Aware Network-Supported Architecture to Distribute Application Flows over Multiple Network Interfaces for B3G*. Accepted for the *Wireless Personal Communications* journal (Springer).
- [12] IST project for a large-scale testbed for mobility in IPv6: <http://www.ist-anemone.eu>.