SeaMoby Working Group
Internet-Draft
Expires: August 20, 2001

O. H. Levkowetz
ABNW
P. R. Calhoun
Sun Microsystems, Inc
G. Kenward
H. M. Syed
Nortel Networks
J. Manner
University of Helsinki
M. Nakhjiri
Motorola
G. Krishnamurthi
R. Koodli
Nokia
K. S. Atwal
Zucotto Wireless
M. Thomas
Cisco
M. Horan
COM DEV
P. Neumiller
3Com

February 19, 2001

Problem Description: Reasons For Doing Context Transfers Between Nodes in an IP Access Network

<draft-ietf-seamoby-context-transfer-problem-stat-00.txt>

Status of this Memo

This Internet-Draft will expire on August 20, 2001.

Copyright Notice

Abstract

There are a large number of IP access networks that support mobility of hosts. For example, wireless Personal Area Networks (PANs) and LANs, satellite and cellular WANs. The nature of this mobility is such that the communication path to the host changes frequently, and rapidly.

This Internet-Draft aims at expressing the problems occurring during such mobility which require the exchange of IP flow context between different nodes in the access network.

A reference architecture is described and the central terms "Context" and "Context Transfer" defined. Some explicit problems which benefits from context transfer are listed.

Table of Contents

1. Introduction 5
2. Reference Architecture 5
2.1 Definitions 6
2.2 Architectural Diagram 8
3. Context Defined 8
3.1 Basic definitions 8
3.1.1 Microflows 9
3.1.2 Feature 9
3.1.3 Feature State 10
3.1.4 Feature Context 10
3.1.5 Microflow Context 10
3.1.6 Context 10
3.2 Some Types of Context 10
3.2.1 Security and AAA information 11
3.2.2 Header Compression State 11
3.2.2.1 Terminology 11
3.2.2.2 Discussion 12
3.2.3 Diffserv / Intserv 13
3.2.4 QoS Policy Management 13
3.2.5 Buffers 14
3.2.6 Sub-network layer state 15
4. Context Transfer Defined 15
4.1 Context Transfer -- Alternative Approaches 15
4.1.1 No context transfer 15
4.1.2 Mobile Updates new Access Router 15
4.1.3 Access Routers Exchange State 16
4.1.4 Central repository 16
4.1.5 Each Application is aware of handovers and access routers 17
5. Security Considerations 17
References 17
Authors' Addresses 18
A. Context Transfer Issues to Consider 20
A.1 Selection of a generic architecture for context transfer . 20
A.2 Definition of Sender and Receiver 21
A.3 Transferring the context information 21
A.4 Knowledge of neighbour capabilities 22
A.5 Per packet or real-time context information transfer . . . 22
A.6 Partially Failed Context Transfer 22
A.7 Different security environments 23
A.8 Dynamic trust relationships 23
A.9 Context Transfer Security Issues 24
A.10 Mobile Routers 24
A.11 Characteristics of a Context Transfer Transport Protocol . 24
Full Copyright Statement 26

1. Introduction

There are a large number of IP access networks that support mobility of hosts. For example, wireless Personal Area Networks (PANs) and LANs, satellite and cellular WANs. The nature of this mobility is such that the communication path to the host changes frequently, and rapidly. In many situations, the change of communications path includes a change in communications media between the host and access networking, including changes from a wireless to a wired connection.

This Internet-Draft aims at expressing the problems occurring during such mobility which require the exchange of IP flow context between different nodes in the access network.

In networks where hosts are mobile, the success of real-time sensitive services like VoIP telephony, video, and others rests heavily on the matter of how seamless (Section 2.1) a handover can be made. Perfect seamlessness would mean that mobility will not give the user of IP based services any reduction in the quality of service received. There exist a number of impediments to perfectly seamless handovers if only existing protocols and technology are used. Some of these are listed in Section 3.2, and include set-up of AAA, header compression, Diffserv/Intserv, policies, and possibly lower layers, e.g. PPP, to be done after each handover. Context transfers reduce the effect of handovers on real-time applications, by minimizing the time needed to attain the level of service provided to the mobile node at the previous access router.

In the rest of the draft, a reference architecture and definitions of the terms "Context" and "Context Transfer" will be given, during this we will list in more detail a number of cases where explicit context transfer would be advantageous, and why. In Appendix A we mention some issues that need to be considered in designing context transfer protocols.

2. Reference Architecture

The reference architecture described with definitions and diagrams below is a functional map, and may map in many ways to physical implementations. In particular, one physical container may well include several different functional elements.

In general, the definitions of draft-manner-seamoby-terms-00.txt [1] are applicable to this document. In particular, the following terms are useful:

2.1 Definitions

Seamless
The absolute reference definition for a seamless handover is one in which there is no change in service capability, security, or quality.

In practice, some degradation in service is to be expected. The definition of a seamless handover in the practical case should be that the end user does not detect any change in service capability, security or quality.

Since the user's ability to detect change is subjective and conditioned by many environmental conditions, this definition is extremely difficult to quantify. Characterization of end user perception of seamlessness is beyond the scope of an IETF working group. Thus, the reference definition, although stringent, is the best working definition for Seamless Mobility.

Mobile Node (MN)
An IP node capable of changing its point of attachment to the network. The MN can be either a mobile end-node or a mobile router serving an arbitrarily complex mobile network.

Mobile Router
A mobile node can be a router, which is responsible for the mobility of one or more entire networks moving together, perhaps on an airplane, a ship, a train, an automobile, a bicycle, or a kayak. The nodes connected to a network served by the mobile router may themselves be fixed nodes or mobile nodes or routers. In this document, such networks are called "mobile networks". [2]

Mobile Host (MH)
An IP node capable of changing its point of attachment to the network. The MH only refers to an end-node without further routing support.

Access Point (AP)
A layer 2 device that is connected to one or more Access Routers and offers the wireless link connection to the MH. Access Points are sometimes called 'base stations'. Note that this usage differs from that used by some Access Router vendors, who call their boxes 'Access Points'.

Access Router (AR)
An IP router residing in an Access Network and connected to one or more access points. An AR offers connectivity to MNs.

The router may include intelligence beyond simple forwarding service offered by ordinary IP routers. An AR communicates with one or more Access Points.

Access Network Gateway (ANG)
An IP gateway that separates the Access Network from a third party network.

Access Network (AN)
An IP network that includes one or more ARs and ANGs.

Administrative Domain(AD)
Administrative Domain: A collection of networks under the same administrative control and grouped together for administrative purposes. [3]

Radio Cell
An area associated with each AP, where there is radio coverage, i.e. where radio communication between a MN and the specific AP is possible.

2.2 Architectural Diagram

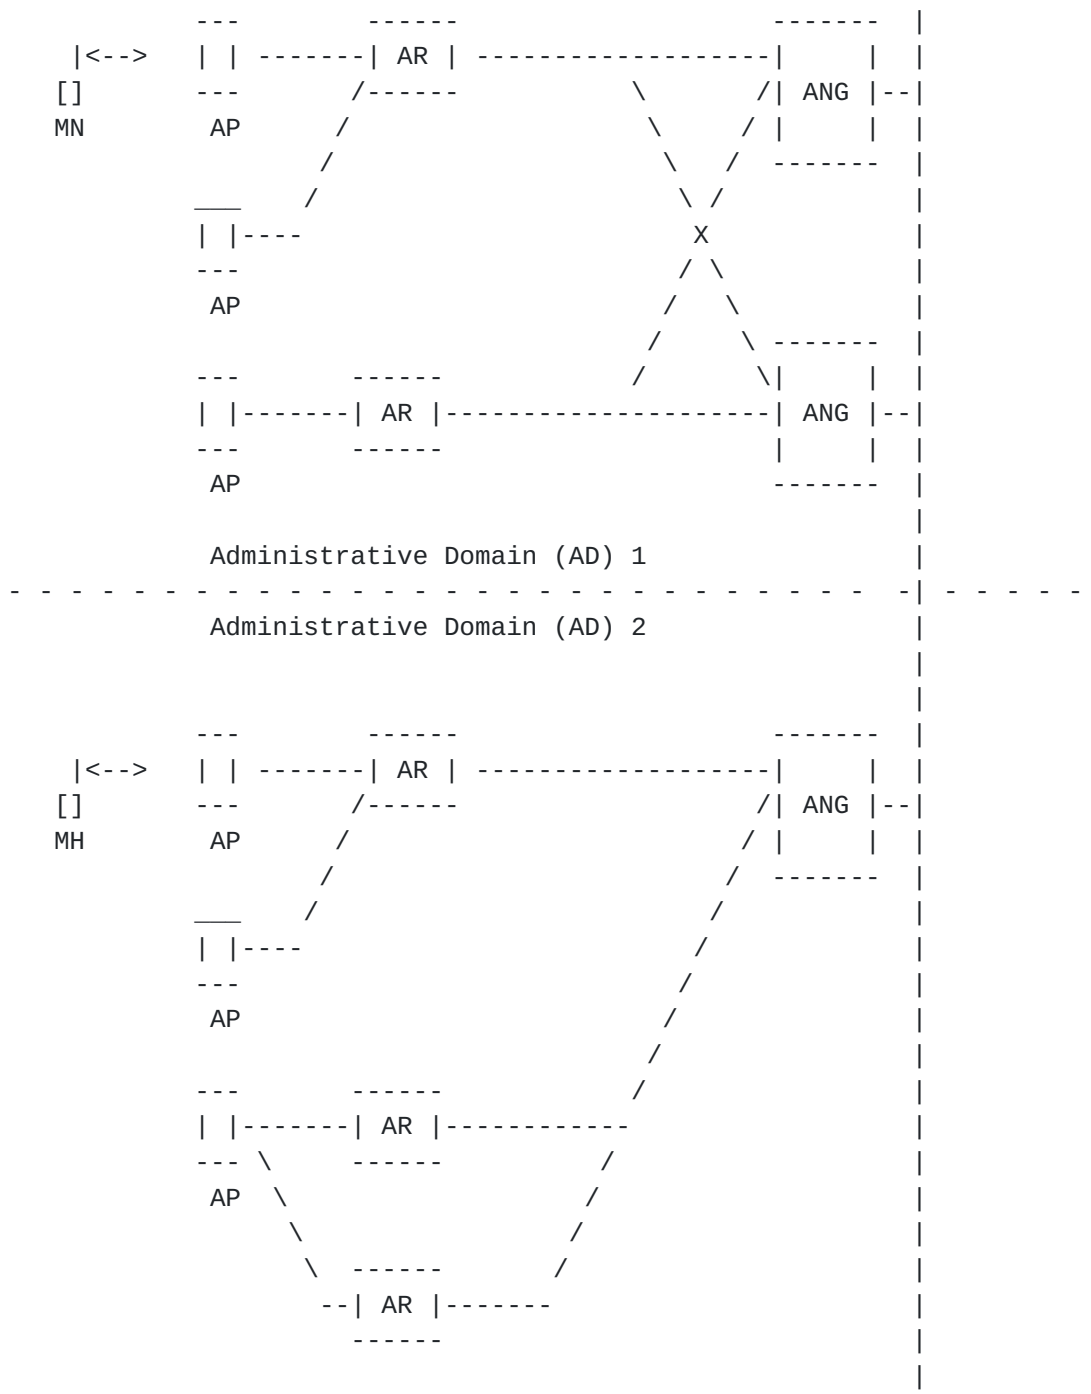

3. Context Defined

3.1 Basic definitions

In this section we define the components of context, and context itself, as used in this document.

3.1.1 Microflows

The fundamental unit of IP service is the microflow. IP microflows may be bundled, or aggregated, for a variety of reasons. As examples, Differentiated Services are provided to aggregates of IP microflows, and authentication is typically associated with the all the IP microflows with the same source address. However, the smallest component of traffic sent to and from a given MN that may have a distinct context is an IP microflow.

RFC 2475[4] defines microflow as 'A single instance of an application-to application flow of packets destined to or originated from a mobile device (MN or MH), which is identified by source address, source port, destination address, destination port and protocol id.'

RFC 2207[5] also provides a definition of a microflow based on the IPsec SPI found in the AH or ESP header. Other handles for micro flows may be forthcoming in the future, and in particular the use of the IPv6 flow label may be standardized.

Two or more IP microflows may be collectively supporting a higher layer application. There is an association or co-dependency between these microflows that should also be preserved during handover. However, this co-dependency relationship between microflows is unknown to the network, and thus cannot be explicitly preserved by the network.

It would seem reasonable that higher layer context would be preserved implicitly if seamless handover is achieved. There is cause for concern when a completely seamless handover cannot be achieved: degradation of the context for even one of a set of co-dependent microflows may have a disproportionate impact on the function or performance of the application.

3.1.2 Feature

A feature is an IP network functionality offered to a mobile node that is of interest to context transfers.

Each feature, in general, is an amalgamation of one or more protocols and the associated mechanisms that support them. It is worthwhile to observe the multiplicity of these "protocol mechanisms" underlying each feature.

As an example, header compression is a feature that consists of various "mechanisms", such as IP/TCP compression and IPv6/UDP/RTP compression, operating according to well-defined "protocols". Thus, an instance of a feature reflects the state associated with

protocols and the corresponding mechanisms.

3.1.3 Feature State

The condition or state of a particular feature instance, representing the current evolution of the behaviour of that feature. At a given instant in time, the state of a feature instance is represented by the values of the data elements associated with the feature instance. Some of these data elements are time invariant and represent the configuration of the feature instance, while others elements, often called "state variables" change value over time in support of various protocol operations related to that feature.

3.1.4 Feature Context

This is the state information associated with a particular feature for a microflow. When a feature is supported by multiple protocol mechanisms, the state information has to be specific for each such protocol mechanism. As an example, header compression is a feature that may support IP/TCP compression and IPv6/UDP/RTP compression. The feature context for header compression must clearly specify the state information for each supported mechanism.

A feature context is the smallest unit of context transfer.

3.1.5 Microflow Context

A Microflow Context is a collection of various feature contexts associated with a microflow.

3.1.6 Context

A set of all microflow contexts representing all the microflows associated with a mobile node.

Eventual feature state modifications, needed for proper operation or maximizing the utility of a feature at the new network, are outside the scope of the context definition. The same applies to cases where protocols at the original network are not supported at the new network; matching equivalent protocols at different sides of a handover is outside the scope.

3.2 Some Types of Context

Some examples of context that might need to be transferred is given below. This does not constitute a complete list of possible context to be transferred.

3.2.1 Security and AAA information

Examples of security services are those provided to a MH, such as user data encryption, user data integrity protection, and MH location privacy. Examples of network security requirements are network topology and policy protection. Example of AAA mechanisms are MH-to-network authentication, authorization of MH for access to a specific service type, and accounting, including network usage records for billing and traffic engineering purposes.

Security and AAA context include the information regarding trust relationships to provide security services and the data to maintain AAA functionality. However, the context only includes the information that already exists at the source AR prior to handover, i.e. the information needed to re-establish the trust of AAA services at the target AR is not part of the context that must be transferred. This is explained in further detail as a security issue in Appendix A.

Among the most important security and AAA context data are security associations (SA), which include encryption or authentication keys and algorithms, identification data necessary for authentication , and authorization of usage privileges. For seamless handover, the security and AAA portion of the context, needs to be transferred as part of the context transfer process in order to expedite the resumption of the security and AAA services at the target AR.

3.2.2 Header Compression State

Compression of IP and transport layer headers is crucial over low-bandwidth links in order to achieve efficient utilization of the link capacity to deliver useful payload to applications. Header compression requires the maintenance of state information at the network periphery, the AR, and at the MN.

When terminal mobility is involved, relocation of the compression context(s) is needed in order to avoid surges in bandwidth consumption associated with compression context re-establishment, which causes interruptions to the smooth delivery of real-time packets. This overhead can be avoided, and thus the seamless operation of header compression facilitated, by a simple transfer of context variables from an access router's compression engine to that of the terminal's new access router.

3.2.2.1 Terminology

The following terminology is used in describing the need for header compression context relocation.

HC:
Header Compression

Full Header (FH) packet:
Contains the full IP and transport protocol headers, plus the HC-specific fields

First Order (FO) [packet]:
Contains only those fields that change from packet to packet (e.g. RTP timestamp), and does not contain fields that do not change at all.

Second Order (SO):
Contains HC-specific header and sequence number. The rest of the fields are constructed from the sequence number info and the information maintained by the decompressor.

3.2.2.2 Discussion

The motivation for HC context relocation is best understood by examining the typical header compression operation.

A compressor starts by sending Full Header (FH) packets with HC-specific headers in them and waits for few of the FHs to be reliably propagated (e.g., ACK-based or 'n' number of headers transmitted). Note that in bandwidth-constrained links, the link latency as well as higher error probabilities could force the transmission of many FHs before confirming reliable propagation of header information. Similarly, many FO packets are sent before confirming that transition to the SO state is possible. This process of "reference state" establishment is expensive. E.g., on a 60 ms cellular link and with 20 ms packetisation for voice, it would take 6 FHs (assuming no errors and dropped packets) to establish the FH state that allows transition to FO, and 6 more FOs to establish the FH+FO reference state that allows transition to SO state. The total overhead is thus 240 ms plus the processing overhead associated with the header compression operation. Perhaps a value of 300 ms may be deemed typical. Since this context establishment needs to be done for each unidirectional packet stream, the overhead gets worse with multiple packets streams belonging to the same mobile node (approximately 600 ms for bi-directional voice). When Mobile IPv6 is used with Home Address option, FH = 84 bytes (excluding HC-specific fields) for a payload of about 20-30 bytes, and FO could be 8 bytes. In comparison, the overhead is 1 byte when operating in the SO state.

The expensive overheads associated with context re-establishment can be avoided by relocating the appropriate context between access

routers. For each type of compression mechanism used (e.g.,IPv6/UDP/RTP, IPv6 only, IPv6/TCP wtc), a Compression Profile Type (CPT) identifies the particular type, and defines the state variables. Thus, the combination of CPT, and the associated state variables (along with a suitable identifier) constitutes the context for transfer purpose. This transfer (presumably) occurs synchronously with handover signalling associated with terminal mobility.

3.2.3 Diffserv / Intserv

Integrated services and Differentiated services are the two proposed QoS-enabling frameworks in the IP networks. A resource reservation protocol (RSVP) enables the Integrated services framework. Both of the mechanisms (RSVP and Diffserv) are stateful and requires certain information to be maintained at the access router. For example, in a pure RSVP-enabled session, the access router requires the flow classification information and the bandwidth requirements of a particular flow to make a packet forwarding decision. The flow classification state is composed of the 5-tuples that uniquely identifies a flow (source/destination IP address, source/destination port and the protocol ID).

Similarly, the Diffserv-enabled routers require the classification, packet forwarding and traffic conditioning information to perform an appropriate scheduling of the flows. In addition to the classification information, a Diffserv code point (DSCP) is also required as a state information at the access router. Meter values for an MN used to police and shape microflows also form part of the MN context.

In a handover situation, all candidate access routers will need the QoS context used by the old access router to support an MN's microflows. Once the information is available (via context transfer) at the potential new access router, it can be processed to make any decisions on whether or not the whole (or partial) QoS context can be supported with the available capabilities of the access router. The target capabilities may include support of the QoS mechanism (for example the target router may not have Diffserv support), size of the queues, policy rules at the router, available resources on particular interfaces etc.

3.2.4 QoS Policy Management

The authorization of the QoS requests against the user profile can be performed based on the service level agreements set up between the user's applications and the network. These agreements are translated into the network policies and controlled by policy

servers responsible for the domain policy control. A policy-based

admission control framework has already been defined and standardized at the IETF [3]. The common open policy services (COPS) [6] is the protocol that carries the network policies in the form of device configuration parameters from a policy server to the actual device for enforcement. COPS is a stateful mechanism that keeps a synchronized copy of all decisions at both client and the server.

There are two popular flavours of COPS protocol. The outsourcing model of COPS [7] allows the network devices to outsource the policy decisions to the policy server by encapsulating the RSVP message objects into COPS-RSVP protocol. The dynamic admission control decisions are based on a per RSVP request and the decision states are maintained at the client as well as the server. Any change in the decision or the device configuration is propagated to either side in a way that the client and server states are always mirrored.

In the provisioning model [8], The Policy server provision the device policies when the device is introduced in its domain. These decisions contain both user and device specific policies. Again, the decision states are maintained at both client and server and any change is propagated to each other.

In case of change in access router (due to user mobility), the new serving access router need the policy context created and maintained at the source access router. The new device may or may not reside in the same policy domain. In either case, the policy context would help the new access router or the policy server responsible for it to make any decision on whether the mobile's context can be supported at the device, or a subset of the whole context could be allowed. If the access router belongs to a different policy domain, the mobile's active sessions may need to be re-authorized against its profile in the new policy domain.

3.2.5 Buffers

A requirement for seamless handovers is to minimize the packet loss when a MN moves from the old AR to the new AR. Thus, buffered packets must be transferred to provide seamless handovers in mobile networks. A research effort that supports this claim is [9]. The incoming packets to a MN are buffered at the old AR and are transferred to the new AR when the new AR is made known to the old AR. The new AR buffers the packets until they can be forwarded to the MN. A buffering context for the MN would comprise the packets that the MN receives at the old AR. Issues like buffer capacity at the new AR are to be considered to ensure successful buffering context transfer during a handover. The new AR, therefore, would need information about the buffer requirements for the MN.

3.2.6 Sub-network layer state

Sub-network Layer State information may include PPP state information. To prevent reestablishment of a connection during a handover from one AR to another this information may be transferred. Examples of PPP context are standard LCP parameters including Max Receive Unit, Authentication Protocol, Magic Number, and header compression. Examples of IPCP (NCP) parameters include IP address header compression and DNS information. However, it should be noted that some information, such as DNS, may change when moving to a new access router and therefore transferring this may be less than useful; instead renegotiation may be needed.

4. Context Transfer Defined

Context transfer is a mechanism for establishing sufficient conditions at one or more ARs to fully support the microflow(s) of a mobile node. After completion of a context transfer, an AR will be capable of forwarding the IP packets to and from the mobile node without disruption of the established service.

4.1 Context Transfer -- Alternative Approaches

There are many alternatives when considering context transfer between two Access Routers. The following sections discusses some of these alternatives and the considerations that need to be addressed.

4.1.1 No context transfer

No context transfer is essentially how today's Mobile IP networks operate. When a Mobile Node moves to a New Access Router, it re-establishes any state with the Access Router. Each application will use its own protocol (signalling) to update the New Access Router.

- Pro: No changes, no additional protocol.
- Con: Long latency (service interruption) during handover
- Con: Use of radio channel bandwidth to re-establish context
- Con: Every application needs to adapt to changing service levels from the network.

4.1.2 Mobile Updates new Access Router

Another alternative would allow a Mobile to update the new Access Router with its state. However, for the Mobile to have an accurate snapshot of the current context, it would be necessary to periodically transfer the AR context to the Mobile (e.g. AAA or QOS state may change periodically on the current Access Router).

- Pro: No connectivity required between ARs.
- Con: Added state synchronization needed between AR and MN
- Con: Long latency (service interruption) during handover if the context transfer cannot be done proactively.
- Con: Use of radio channel bandwidth to re-establish context
- Con: Security problems, MN may not be a trusted entity
- Con: New protocol needed for context transfer

4.1.3 Access Routers Exchange State

In this alternative, when Access Routers notice that a handover is occurring or imminent, context information would be sent to the candidate Access Routers. It is assumed that a generalized protocol would carry all of the context for all of the mobile node's microflows.

- Pro: No use of radio bandwidth to re-establish context
- Pro: A large reduction of latency (service interruption) during handover compared to the case in Section 4.1.1 (assuming radio bandwidth is much less than fixed bandwidth)
- Pro: Possibly elimination of latency due to context transfer, as it may be done before and / or during handover
- Con: Protocol needed for context transfer
- Con: A mechanism to choose candidate ARs and where to finally handover is needed.

4.1.4 Central repository

The context transfer between access points/routers is controlled by a central entity in the network. This entity could be a policy server, one of the access network gateways, or one of the access routers. This entity keeps the context information for the mobiles that are registered under the domain of that entity and "re-installs" the context at the new access point/router as the mobile moves. This entity may also "delete" the context from the old access point, if required.

- Pro: No use of radio bandwidth to re-establish context
- Pro: A large reduction of latency (service interruption) during handover compared to the case in Section 4.1.1 (assuming radio bandwidth is much less than fixed bandwidth)
- Pro: Possibly elimination of latency due to context transfer, as it may be done before and / or during handover
- Pro: one clear decision point (similar to PDP) and information storage. Knows the state of the whole (sub) network.
- Con: New protocol needed for context transfer
- Con: As in Section 4.1.2, added synchronization is needed, this time between AR and repository

- Con: May have scalability problems in terms of the number of

mobiles registered within the domain of the entity and the number of active sessions per mobile. Can be even worse if multiple access networks are involved.
- Con: Contract relationship must pre-exist or be established

4.1.5 Each Application is aware of handovers and access routers

A last alternative would be to define the requirements for context transfer, and modify all applications to arrange for state to be moved between Access Routers.

- Pro: Context transfer more fully under application control
- Con: Context transfer additions needed for every single higher protocol -- multiplied complexity
- Con: Only applications specifically adapted for context transfer would be able to take advantage of seamless mobility
- Con: wasted bandwidth (if all application transfer their context information individually)

5. Security Considerations

This type of non-protocol document does not directly affect the security of the Internet. (However, for some comments on possible security issues with the implementation of context transfer, see Appendix A.7, Appendix A.8 and Appendix A.9).

References

[1] Manner, et al., "Mobility Related Terminology", draft-manner-seamoby-terms-00 (work in progress), January 2001.

[2] Perkins, C., "IP Mobility Support", RFC 2002, October 1996.

[3] Yavatkar, et al., "A Framework for Policy-based Admission Control", RFC 2753, January 2000.

[4] Blake, et al., "An Architecture for Differentiated Services", RFC 2475, December 1998.

[5] Berger & O'Malley, , "RSVP Extensions for IPSEC Data Flows", RFC 2207, September 1997.

[6] Durham, et al., "The COPS Common Open Policy Services protocol", RFC 2748, January 2000.

[7] Herzog, et al., "COPS Usage for RSVP", RFC 2749, January 2000.

[8] Chan, et al., "COPS Usage for Policy provisioning", draft-ietf-rap-pr-05 (work in progress), October 2000.

[9] Caceres, R. and V.N. Padmanabhan, "Fast and scalable wireless handovers in support of mobile Internet audio", Mobile Networks and Applications 3, 1998, pp. 351-363, 1998.

Authors' Addresses

O. Henrik Levkowetz
A Brand New World
Ùster÷gatan 1
S-164 28 Kista
SWEDEN

Phone: +46 8 477 9942
EMail: henrik@levkowetz.com

Pat R. Calhoun
Network and Security Research Center, Sun Labs
15 Network Circle
Menlo Park CA 94025
USA

Phone: +1 650-786-7733
EMail: pcalhoun@eng.sun.com

Gary Kenward
Nortel Networks
3500 Carling Avenue
Nepean, Ontario K2G 6J8
CANADA

Phone: +1 613-765-1437
EMail: gkenward@nortelnetworks.com

Hamid Syed
Nortel Networks
100 Constellation Crescent
Nepean Ontario K2G 6J8
CANADA

Phone: +1 613 763-6553
EMail: hmsyed@nortelnetworks.com

Jukka Manner
Department of Computer Science, University of Helsinki
P.O. Box 26 (Teollisuuskatu 23)
FIN-00014 Helsinki
FINLAND

Phone: +358-9-191-44210
EMail: jmanner@cs.helsinki.fi

Madjid Nakhjiri
Motorola
1501 West Shure Drive
Arlington Heights IL 60004
USA

Phone: +1 847-632-5030
EMail: madjid.nakhjiri@motorola.com

Govind Krishnamurthi
Communications Systems Laboratory, Nokia Research Center
5 Wayside Road
Burlington MA 01803
USA

Phone: +1 781 993 3627
EMail: govind.krishnamurthi@nokia.com

Rajeev Koodli
Communications Systems Lab, Nokia Research Center
313 Fairchild Drive
Mountain View CA 94043
USA

Phone: +1 650 625 2359
EMail: rajeev.koodli@nokia.com

Kulwinder S. Atwal
Zucotto Wireless Inc.
Ottawa Ontario K1P 6E2
CANADA

Phone: +1 613 789 0090
EMail: kulwinder.atwal@zucotto.com

Michael Thomas
Cisco Systems
375 E Tasman Rd
San Jose CA 95134
USA

Phone: +1 408 525 5386
EMail: mat@cisco.com

Mat Horan
COM DEV Wireless Group
San Luis Obispo CA 93401
USA

Phone: +1 805 544 1089
EMail: mat.horan@comdev.cc

Phillip Neumiller
3Com Corporation
1800 W. Central Road
Mount Prospect IL 60056
USA

EMail: phil_neumiller@3com.com

Appendix A. Context Transfer Issues to Consider

Context transfer may come in many flavours and implementations. This section lists some issues that have come to light during the formulation of the problem statement for context transfer.

A.1 Selection of a generic architecture for context transfer

Two architectures can be discussed here: Centralized Vs Distributed

The first architecture assumes a single central entity in the network or in a defined domain, which maintains the updated context information on behalf of a set of access routers and is also responsible for distributing it amongst the potential receivers. This role can be applied, for example, to the access network gateway, an entity like a policy server or an elected access router itself. This approach may suffer with scalability issues as the number of microflow information per mobile per access router could be a very large data to process and transfer in real-time with minimum delays.

The second architecture distributes the role of context maintenance

and distribution to the access routers itself. This approach seems to be more appropriate as the access routers hold most of the context information (QoS, Policy etc). Moreover, the access router is responsible for the context information or the microflows of the mobiles that are connected through them, which is scalable.

A.2 Definition of Sender and Receiver

Defining the sender and the potential receiver(s) of the context information and "events" that enable access routers become the members of one "context group":

A mobile may have radio connectivity (at least it may receive RF signals) from more than a single AP. The case where the APs are connected to the same access router AR, no context transfer is required but when the potential APs are connected to different ARs, the access routers may expect the mobile's QoS sessions to arrive. These access routers, therefore, become the potential receivers of the context information. A "context group" can be defined as the group of potential access routers in the network that have, or need to have, the context information related to a mobile. There are few issues to resolve in the 'dynamic' formation of a context group;

- The addition and deletion of the members to a context group: This could be triggered by certain network events. These events may be network policy/configuration conditions that dictate the access routers to join a context group or could be an event generated by the mobile's movement prediction that dictates an access router to become the member of a context group.
- How does a (potential) new member identify the context group associated with the mobile?: This requires an 'identifier' for each context group per mobile and the information of the identifier should be propagated to any potential member of the context group

A.3 Transferring the context information

- How does a new member of a context group know about the sender of the context? The context information of the mobile are maintained by some access router in the network or in other words, there is one potential sender per context group. The new member needs to know who to contact (within the context group) in order to retrieve the current context.
- When should the context information be transferred? Two possible approaches are "reactive" and "proactive or make-before-break". In the later approach, the context information is transferred to any new member as soon as it joins the context group irrespective of the fact that the mobile may not choose the access router for data connectivity.

- How updates in the context propagate to the members of the

context group, and probably any associated events that trigger the context update need to be understood.

A.4 Knowledge of neighbour capabilities

In some circumstances, knowledge of neighbouring ARs can lead to better handover strategies, as well as help with load balancing, etc. While not necessary, and possibly undesirable in some cases, it would be useful to have the capability to use topology knowledge when helpful.

A.5 Per packet or real-time context information transfer

There are pieces of information that an access router calculates and maintains on a per packet basis. Metering in a Diffserv-enabled router is one example of such an information. The context transfer solution must address how and when the per packet information could be transferred between the access routers and what are the trade-offs. For example, for a proactive case, the transferred metering information may become irrelevant or obsolete at the new access router because of the instant of it was last updated and the time when the router really needs this information could be large enough. Even for a reactive transfer mode, by the time the information is propagated to the new access router , it may become obsolete or irrelevant.

A.6 Partially Failed Context Transfer

[Editor: Rajeev had some comments on this text on Jan 18, these have not been taken into consideration here. Hamid or Rajeev, would you like to propose a changed text, please?]

A part of the "definition of context transfer: problem statement" is the fact that the "context information" may not be completely supported by one or more of the receivers of the context. The reason could be a different set of capabilities available at the receivers of the context data, due to which one or more of the "sub-contexts" may not be supported. This may likely to happen in a heterogeneous "context group" where the members of the context group have different set of capabilities. One simple example of such a scenario is failure of admission control due to unavailability of resources required by a "sub-context". The impact could be a service disruption/degradation for one or more sessions of the mobile.

The capabilities of a receiver cannot be known without actual transfer of context. The receiver may then decide whether a complete support of the requirements indicated by the "context" is available or not. In case the outcome is negative, there is a chance of service degradation for one or more of the mobile's sessions.

The question is whether any solution investigation for this problem falls under "context transfer" activity or is beyond the scope of this forum. To explain it better, two possible scenarios can be considered;

1. There could be a mechanism that prevents a handover of bearer traffic to any receiver of context that provides a partial support to the "transferred context". In this scenario, the definition of such a mechanism really goes out of the context transfer activity. Only a feedback on the outcome of the decision on transferred context could be used to trigger such a mechanism. This scenario really based on the assumptions that all the session of mobile's should be handed over to a single receiver.

2. The second scenario may assume that different sub-contexts of the mobile may be supported by different receivers and therefore, the actual handover of the mobile's session would be done to multiple targets. This scenario may require some decisions to be made at the source access router to which sessions are to be moved to which target based on any feedback from the target access routers. This is a complicated situation and may require a substantial amount of work to be done both on context transfer and "partial handover" of sessions

Context transfer may also fail due to imperfect transport, over wired or wireless medium. This also need to be considered in a possible solution.

A.7 Different security environments

Depending on the design of the security provisioning systems and existing trust relationships (e.g. existence of public key infrastructures or AAA administrative domains), in some handover cases, some of data in the security portion of the context might be available at the target network. However, context transfer might not need to consider these (possibly special) cases and will include all the context data in the context transfer procedure in order to cover more general cases.

A.8 Dynamic trust relationships

According to the mobile IP model, at many instances, when a mobile moves into a new administrative domain, a re-authentication of the mobile to the new foreign network (and at times to the home network) is necessary. In a AAA based authentication, besides the static trust relationships, that already exist prior to the arrival of the mobile at the edge of each network (such as that between the target network AAA and home network AAA authorities and that between a

network mobility agent and its AAA authority), there are

relationships that have to be created dynamically. One such relationship is the one needed for re-authentication of MH to the target access agent (may or may not be AP/AR?). The delay involved in re-authentication triggered as a result of a handover might cause considerable and unacceptable latency and loss for many mobile applications.

Context transfer seeks to provide a faster mechanism for transfer of the STATIC security and AAA data than those transfer mechanism already available. However, due to the need for creation of dynamic trust relationships, a trade off in favour of seamlessness might lead to security and AAA compromises before completion the handover.

A.9 Context Transfer Security Issues

Context transfer should include a mutual authentication process, by which the party receiving the context and the party transmitting the context, each provides proof of legitimacy to the other side. Data integrity protection should provided, so that the context information is protected from tampering by a third party. Encryption of context data might be necessary, in case MH location privacy and network topology and policy protection are required. The mechanisms for establishing trust between the two parties involved in context transfer in order to transfer context data securely (as described above) should be provided.

The mobile node must ultimately retain control over whether it moves or not, and under what conditions it consents to have a network based move done on its behalf. In particular, the mobile node must have the ability to veto a move that could happen transparently, but may result in higher access charges, unexpected service degradation, loss of privacy, or other policy based exclusions.

A.10 Mobile Routers

In the case where the MN is actually a mobile network, which may contain it's own wireless Access Points, Access Routers and Mobile-IP entities, there may exist unsolved and even undiscovered issues related to Mobile-IP and mobile routers.

A.11 Characteristics of a Context Transfer Transport Protocol

Assuming that context transfer is needed, the actual protocol used to implement the transport needs to address some problems found in the micro/macro-mobility environments. It is fairly clear that the context transfer will likely need to be extensible since the context examples given in this draft are diverse and subject to change.

It is likely that the context transfer will need to be invoked just

prior to handoff decisions, if the very latest version of it is desired. This means that it will need to be efficient and offer a standard method for indicating success and/or failure of the transfer to the caller, which is likely to be a MIP mobility agent or a SeaMoby micro-mobility entity.

There are methods that COULD be used to decouple context transfer from mobility. One method COULD involve handoff neighbour ARs periodically updating each other irrespective of the mobile traffic that they are carrying. Another method COULD be to perform neighbour AR update whenever a micro-flow is added or dropped from an AR. This subject has not been debated by the SeaMoby WG to any large extent.

Another item of interest is the actual transport protocol used for the context transfer. The merits of reliable versus unreliable, and TCP or SCTP have not been debated extensively by the SeaMoby WG yet.

Full Copyright Statement

Acknowledgement

Funding for the RFC editor function is currently provided by the Internet Society.