

MEXT Working Group
Internet-Draft
Intended status: Informational
Expires: January 11, 2009

B. Lim
C. Ng
K. Aso
Panasonic
S. Krishnan
Ericsson
July 10, 2008

Verification of Care-of Addresses in Multiple Bindings Registration
draft-lim-mext-multiple-coa-verify-02

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on January 11, 2009.

Abstract

Using multiple care-of address registration, there is a possibility that a malicious mobile node could create multiple care-of address bindings that does not belong to the mobile node at its home agent. The home agent would accept these bindings without verifying them due to the trust relationship it has with the mobile node. With these bindings, the mobile node can launch attacks by asking the home agent to flood the victims of these care-of addresses with useless packets. To mitigate such a problem, this memo introduces a few possible verification mechanisms that the home agent would use in order to

verify the care-of addresses for the mobile node before using them for packet routing.

Table of Contents

1.	Introduction	3
2.	Problem Scope	3
2.1.	Initiating the Attack	4
2.2.	Launching the Attack	6
3.	Analysis: Possible Solution Approaches	7
3.1.	Restrict Multiple Care-of Addresses Registration in a Single Binding Update Message	7
3.2.	Using Cryptographically Generated Addresses	9
3.3.	Using a Third Party Trusted Entity	10
3.4.	Using Message Exchanges with a Mobile Node	13
4.	Conclusion	17
5.	Acknowledgements	17
6.	Security Considerations	17
7.	IANA Considerations	17
8.	References	17
8.1.	Normative References	17
8.2.	Informative References	17
Appendix A.	Change Log	18
	Authors' Addresses	19
	Intellectual Property and Copyright Statements	21

[1.](#) Introduction

IP mobility allows a mobile node to maintain an ongoing communication session regardless of where the mobile node roams in the Internet. To achieve this constant connectivity, [\[1\]](#) permits the mobile node to register a transient IP address (e.g. care-of address) at an anchoring point (e.g. home agent). The mobile node sends a binding update (BU) message to the home agent for registering its care-of address. With this registration, the home agent could then forward any packets addressed to the mobile node at the mobile node's care-of address.

Security has always been a key consideration in user communications. This is especially so in IP mobility where the mobile node frequently changes its point of attachment. One such attack in IP mobility is the registration of an invalid care-of address at the home agent. Such registration allows the attacker to direct its traffic to other nodes. Such an attack is previously ignored since the attacker can only use one fake care-of address, thereby losing communications with its home agent once the fake address is registered. However, [\[2\]](#) permits the registration of multiple care-of addresses by means of using an identification number called Binding Unique Identification (BID) number to distinguish between multiple bindings to a home address. Unless proper security methodologies are implemented, the use of multiple care-of address registration heightens the chances of allowing an invalid care-of address to be registered at the home agent.

In [Section 2](#), this document first expands on the problem briefly described in [\[2\]](#) and [\[3\]](#) about the security threat with the binding of fake care-of addresses. Next, possible approaches to mitigate such threat is analyzed in [Section 3](#).

[2.](#) Problem Scope

For Mobile IPv6, it is assumed that the mobile node and its home agent establish some form of mutual authentication such that the home agent trusts the mobile node and will hence accept the registration of any care-of address that the mobile node presents to the home agent. This trust relationship is further strengthened when one assume that ingress filtering is being used such that when the home agent receives a binding update message from the mobile node stating its care-of address as the source address, the home agent trusts that the incoming packets do indeed originate from the specified source address. In addition, the home agent also trusts the routing infrastructure that packets forwarded by the home agent would be sent to the intended destination.

However, such a trust is no longer valid when the mobile node utilizes a single binding update message to register its multiple care-of address at the home agent. This technique is explained in [2] with the aim of introducing some optimization when registering multiple care-of address for a mobile node. Such optimization technique is useful in scenarios when resources (e.g. bandwidth) are scarce on some of the mobile node's interfaces, since it allows the mobile node to send a binding update message containing multiple care-of address to the home agent from an interface that does not have such resource constraint. Moreover, in Mobile IPv6, the use of the alternate care-of address option permits the mobile node to achieve the same effect of registering a care-of address for another interface via a specific interface. This introduces the risk of having fake care-of addresses registered at the home agent and compromise the security of the network. One applicable sceanrio that might encounter such a problem is the Long Term Evolution (LTE) system that is being worked on by the Third Generation Partnership Program (3GPP).

There is a fundenmental difference between a mobile node using Mobile IPv6 to flood a victim with useless packets and a mobile node using Monami6 to flood a victim with useless packets. The difference lies in that using Monami6, the mobile node can bind a real care-of address at the home agent and use this care-of address to control the flow of attacks to the victims (e.g. to increase the packet transmission rate to the vicitm). In Mobile IPv6, the mobile node loses this means of control as the mobile node can only refresh the binding at the home agent. The example in [Section 2.2](#) illustrates the esclated threat of Monami6 in more details.

[2.1.](#) Initiating the Attack

This problem is best understood through a specific example. Let us assume that a malicious mobile node, MN, sends a binding update message to its home agent, HA, to register multiple care-of addresses. This is illustrated in the Figure 1 below.

Lim, et al.

Expires January 11, 2009

[Page 4]

Internet-Draft

Verifying Multiple CoA Bindings

July 2008

[Start of packet header]

Source Address : CoA
Destination Address : HA's address

[Mobility Options]

Binding Unique Identifier: BID1

Binding Unique Identifier: BID2
Care-of Address : V1's address

Binding Unique Identifier: BID3
Care-of Address : V2's address

[End of packet header]

Figure 1: Binding Update Message for Multiple Care-of Addresses
Registration

CoA is a valid care-of address owned by MN. MN is attempting to bind

addresses of two victims, V1 and V2, at HA in order to launch an attack towards the victims. The binding update message is secured using an IPsec security association to protect the integrity and authenticity.


When HA receives this binding update message, it will accept this binding update message based on the following. First, the binding update message is deemed authorized as the correct IPsec association key is used for the message. Second, the trust relationship that HA has with the routing infrastructure allows it to understand that this binding update message is sent from MN. Finally, after the first two checks have succeeded, the trust relationship that HA has with the MN permits it to trust the care-of addresses that are specified in this binding update message. Hence, the binding cache at HA will record three bindings for MN tied to MN's home address (HoA) as shown in Figure 2 below.

Binding 1 [HoA, CoA , BID1]
 Binding 2 [HoA, V1's address, BID2]
 Binding 3 [HoA, V2's address, BID3]

Figure 2: Binding Cache at Home Agent

2.2. Launching the Attack

With the bindings created at HA, the malicious mobile node can now attempt to flood the victims with useless packets. This type of flooding can cause disruption of services at the victims' end as their devices will be busy processing these packets for meaningless data. This is illustrated in Figure 3.






Figure 3: Packet flooding attack scenario

The binding cache at HA contains MN's care-of address along with V1's and V2's addresses. MN uses real-time transport protocol (RTP) and real-time transport control protocol (RTCP) to initiate a video stream from a streaming server (CN). Furthermore, MN could employ flow filtering techniques to assist in the forwarding of packets from HA. Such techniques are described in [4], [5] and [6].

Prior to starting the video flow, MN uses CoA to set filter rules at HA. One of the filter rule would tell HA that the control packets (shown as "a") for the video stream would be forwarded to BID1. Another filter rule informs HA that the data packets (shown in "b") for the video stream would be multicast to BID2 and BID3. When HA receives packets from CN, it would adhere to the filter rules created and forward the packets accordingly. This would cause V1 and V2 to be flooded with unintended data packets. Furthermore, MN can utilize care-of address1 to maintain the attack to V1 and V2 by sending binding update message to constantly renew the lifetime of the bindings. Likewise, MN can also use CoA1 to launch new attacks by binding more victim's addresses at HA and modifying the filter rule to include them in the reception of the data packets.

Although both the home agent and the mobile node are mutually authenticated, this does not stop an authenticated mobile node from acting maliciously. Given time, the home agent might be able to detect that the mobile node is acting maliciously and may then deny

mobility services to such a mobile node. One such way is when the home agent starts receiving requests from those victims to stop forwarding packets to it. With a substantial amount of such requests, the home agent might launch an investigation to determine if the mobile node was acting maliciously. Hence, we see that a home agent might not be oblivious to such an attack.

However, even if the home agent bars a single identity, the attacker

may have multiple other identities available at his/her disposal. Take for example, an attacker is able to purchase multiple pre-paid cards for his/her mobile device. This implies that the attacker would have multiple accounts at its disposal. Using these multiple accounts, the attacker can start a spamming attack by flooding victims with useless data. The attacker knows that eventually his/her account would get banned, but the attacker does not care. The purpose of the attacker has been achieved and that is to disrupt the services of those victims. Furthermore, the attacker can continue to purchase more pre-paid cards to replace those accounts that got banned. Thus, in this sense, having the identity known does not hinder the above problem from occurring.

[3.](#) Analysis: Possible Solution Approaches

This section will analyse possible solutions to counter the threat described in the previous section. One obvious way is for the home agent to stop a mobile node from sending a binding update with multiple care-of addresses embedded in it. Another obvious way is for the home agent to use some mechanisms to verify those care-of addresses embedded in the binding update message. The focus of such verification solutions would be more towards ensuring that the care-of addresses carried within the binding update message itself are correct (and not the source address of the binding update message). This draft assumes the ingress filtering would be deployed in most networks, which allow the home agent to be sure that the source address of the binding update message is correct. Such assumption can be considered valid when we take in account that most operators deploying the 4G cellular networks would likely run ingress filtering within their networks as a security precaution to reduce the chances of packet spoofing.

[3.1.](#) Restrict Multiple Care-of Addresses Registration in a Single Binding Update Message

To prevent such flooding attack threat caused by allowing a mobile node to bind fake care-of addresses at a home agent, it is possible for the home agent to reject the use of a binding update message that specify multiple care-of addresses. This type of logic is described

in [\[7\]](#) where a home agent would reject a binding update message if

the home agent detects that the source address of the binding update message is different from the care-of address specified in the binding update message. The home agent rejects the binding update message as it is unsure if the care-of address actually belongs to the mobile node. It could be that the mobile node is malicious and spoof the care-of address to launch a flooding attack on a victim. Thus, with such logic at the home agent, the threat of a mobile node launching a flooding attack could be adverted. This is illustrated in Figure 4.




Figure 4: Scenario on policy at home agent

A mobile node (MN) attempts to bind a care-of address (CoA2) using bulk registration via another care-of address (CoA1) at its home agent (HA). When HA receives the binding update message from the mobile node, it notices that the source address (CoA1) of the binding update message is not the same as the care-of address (CoA2) specified in the binding update message. HA is unsure if MN is reachable at CoA2. Thus, HA rejects the binding of CoA2 as it fears it might initiate a flooding attack to a victim.

However, by implementing such logic at a home agent, it really limits the amount of optimization a person could achieve on the signaling load between a mobile node and its home agent. For example, in [8], it shows a savings in terms of signaling overheads when the mobile node is able to use a single binding update message to register multiple care-of addresses as opposed to sending them individually. If the home agent restricts the mobile node from utilizing a single binding update message to register multiple care-of addresses, this implies that the mobile node has to send an individual binding update message for each care-of address to the home agent. Hence, such

action removes the optimization benefit of performing the bulk registration operation as described in [2].

It is possible to employ such logic at a home agent to prevent the flooding attack threat without losing all of the optimization benefit of bulk registration. This would require the relaxation of the logic at the home agent to reject the use of specifying multiple care-of addresses in a single binding update message. Of course, such relaxation cannot sacrifice the security requirement in preventing a mobile node to launching a flooding attack. An example of such relaxation could be a policy at the home agent that allows it to identify if a particular care-of address is from another trusted network (e.g. a trusted foreign 3GPP network). This type of trust relationship could be found in 3GPP deployments where they establish trust relationship for roaming purposes. Hence, when the home agent identifies such care-of addresses, it would accept the care-of address. Likewise, if the home agent understands that the care-of address is not from a trusted foreign network, it might employ some care-of address verification methods (e.g. ping the care-of address) or even reject the registration of that care-of address.

Referring back to Figure 4, assuming that HA has establish a trust relationship with non-3GPP network. When HA receives a binding update from CoA1 specifying the registration of CoA2, HA would accept CoA2. On the other hand, if HA does not have a trust relationship with non-3GPP network, HA would inform MN of the rejection of CoA2. This rejection would then force MN to send an individual binding update message for CoA2.

Also, such knowledge of trusted networks could be made known to the mobile node from the home agent. For example, MN could be pre-configured with a list of trusted networks that it is connected to at the moment. As and when MN roams, HA would update this list to allow MN to know which access networks that it is currently connected to are trusted. This action permits the mobile node to know which care-of addresses could be used for bulk registration.

[3.2.](#) Using Cryptographically Generated Addresses

Another approach to prevent such flooding attack threat is to have all care-of addresses to be formed using the cryptographically generated addresses (CGA) technique [9]. With CGA in place, it implies that a mobile node's care-of address would have to be configured using its public key along with the subnet prefix. Furthermore, this care-of address would need to be signed by the mobile node's private key. With such system in place, the mobile

node would not be able to successful in binding a spoof care-of address at a home agent. Even if the mobile node has successfully

obtain the public key of another mobile node, the mobile node would not be able to get the private key of this same mobile node to sign the care-of address. Without the proper signature, the home agent would not accept the care-of address for the mobile node. Thus, this prevents the mobile node from binding fake care-of addresses at the home agent.

However, introducing cryptographically generated care-of addresses increases the complexity of the mechanism to achieve multiple bindings. It is unclear how a message can contain multiple CGA signatures for each of the care-of addresses. Furthermore, complexity is increased by the fact that additional addresses are not found in the source address field but somewhere in the extension header of the packet (i.e. mobility header). This requires substantial integration between the CGA module and Mobility Support module in a network stack implementation. Additionally, each CGA Parameters structure which is at least 73 octets in length must be added to the binding update message, further increasing its size. Finally, the method of using CGA does not prevent the malicious mobile node to launch a flooding attack against a subnet by generating multiple non-existent care-of addresses in that subnet using the mobile node's own public keys.

To reduce the overheads from carrying the CGA parameters data structure for each care-of address that uses a CGA, it is possible to use of the same hash modifier value and public key for each of the care-of addresses in the binding update message. This allows the receiving home agent to use the same modifier and public key to verify that the mobile node (which sent the BU) is in fact the owner of all the care-of addresses present in the binding update message. Hence, it is possible for the home agent to just verify the reachability of any one of the care-of addresses to be assured that the mobile node is not faking the other care-of addresses and redirecting traffic to other victims. However, there are some privacy issues with using this method of generating multiple care-of addresses with the same modifier since they are easily linkable. Also, if there are a substantial number of nodes on one of the links hosting one of the care-of addresses, the collision count may be incremented and this mechanism will cease to work.

3.3. Using a Third Party Trusted Entity

Yet another approach to solve the flooding attack threat is to use a trusted anchor to verify the care-of addresses of the mobile node. For example, in [10], a mobile node is able to establish a symbiotic relationship to an access router (which is a third party entity) in order to permit the home agent to verify the care-of address of the mobile node. The mobile node will tell its home agent of such

Lim, et al.

Expires January 11, 2009

[Page 10]

Internet-Draft

Verifying Multiple CoA Bindings

July 2008

symbiotic relationship establishment by providing the IP address of access router for the home agent to contact. This prompts the home agent to contact the access router to verify the care-of address that the mobile node claims to be using. This is illustrated in Figure 5.

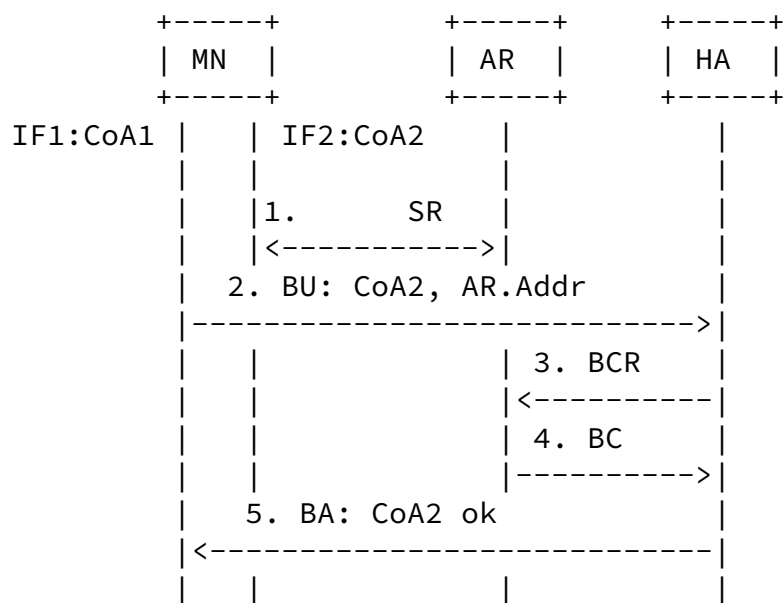


Figure 5: Message sequence diagram on using symbiotic relationship

A mobile node (MN) has two interfaces (IF1, IF2). On IF1, the MN is using CoA1 for communication. Likewise, for IF2, the MN is using CoA2 for communication. MN intends to use bulk registration via IF1 to register CoA2 at its home agent (HA). At first, MN establishes a symbiotic relationship with an access router (AR) that IF2 is associated with (step 1). Once the symbiotic relationship has been

setup, MN sends a binding update (BU) message from IF1 to register CoA2 at HA (step 2). Also, in this BU message, MN tells HA regarding the symbiotic relationship it has setup with AR (e.g. AR's IP address). When HA receives this information, HA will send a Binding Confirm Request (BCR) message to query the AR if MN is using CoA2 (step 3). If the AR identifies that MN is indeed using CoA2, AR will reply back to HA with a Binding Confirm (BC) message (step 4). With positive confirmation that MN is using CoA2, HA will response back to MN indicating that the registration of CoA2 has been successful (step 5). Such message exchanges between the HA and the AR could be secured via public keys methodology (PKI).

Typically, the verification of a mobile node's care-of addresses usually involves only the mobile node and its home agent. However, the use of a third party entity to verify care-of addresses for the mobile node implies the need for an infrastructure support. The

purpose of using the third party entity is to offload the work of verifying care-of addresses from the mobile node to the third party entity. This act of testing the mobile node's reach-ability would result in the reduction of signaling the mobile node has to perform with its home agent. However, such shift of responsibility from the mobile node to the access router means that workload for the access router would increase. This is especially so if the access router has a number of mobile nodes associated to it that requires such validation services. This implies that the access router would have to perform the message exchanges to the respective home agents on behalf of all these mobile nodes. Such action might significantly increase the signaling load of the access router and could be seen as a tradeoff for the simplicity at the mobile node.

It is possible to minimize this tradeoff effect by reducing the amount of signaling that an access router has to do for the mobile nodes associated to it. For example, a certificate could be issued to a mobile node by the foreign network that the mobile node is currently connected to. This certificate would contain information pertaining to the validity of the mobile node's care-of addresses specified in the foreign network. When the mobile node notifies the home agent regarding the symbiotic relationship established with the foreign network, it would further send this certificate to the home agent. If there is sufficient trust relationship between the home agent and the foreign network, the home agent could use the

certificate as prove of the mobile node's care-of addresses. This would help remove the need for the home agent to do message exchanges with the foreign network, thus reducing the signaling load. On the other hand, if the home agent does not feel that it has sufficient trust relationship with the foreign network, it could perform the message exchange sequence with the foreign network to test the reachability of the mobile node's care-of addresses. This is illustrated in Figure 6.




Figure 6: Message sequence diagram on using Certificate

A mobile node (MN) has two interfaces (IF1, IF2). On IF1, the MN is using CoA1 for communication. MN associates with an access router (AR) to obtain a care-of address (CoA2) for communication for IF2 (step 1). Also, AR forwards a certificate to MN to allow MN to use it as proof to its home agent (HA) regarding the use CoA2 (step 2). In this example, as per symbiotic relationship, we assume that AR and HA has exchange the necessary PKI parameters (e.g. public keys). Hence, the certificate would contain CoA2 along with a timestamp signed using AR's private key. The purpose of the timestamp is to prevent MN from performing a replay attack to HA by claiming to own the same CoA2 at a later time. When MN decides to use bulk registration to bind CoA2 at HA, MN sends a binding update (BU) message from IF1 along with the token to register CoA2 at HA (step 3). At HA, it authenticates that the certificate is indeed from AR and thus, reply to MN that the binding for CoA2 has been accepted (step 4). Hence, compared to Figure 5, the certificate allows the HA to skip the message exchange sequence with the AR.

[3.4.](#) Using Message Exchanges with a Mobile Node

Alternatively, the flooding attack threat may be tackled using a series of message exchanges between a mobile node and its home agent to verify the reach-ability of the mobile node's care-of addresses. For example, the home agent send some encrypted information to the mobile node to verify the reach-ability for that care-of address. If the home agent is able to received the decrypted version of the information from the mobile node, it would prove to the home agent that the mobile node is using that particular care-of address. This technique is described in [\[11\]](#) where the use of a cookie proves to the home agent that the mobile node is addressable at the specified

care-of address. However, the optimization benefit of sending multiple care-of address within a single binding update is greatly reduced as it would require the mobile node to respond via all these care-of addresses. In this case, the mobile node might as well send the binding update message individually for each care-of address.

To optimize the inefficiency of using [\[11\]](#) for multiple care-of addresses case, it is possible for the home agent to send a notification to a mobile node via one unverified care-of address and ask the mobile node to respond to the reception of the notification via another care-of address. This utilizes the concept of multiple

care-of addresses whereby the mobile node need not respond back using the same path as the reception of the request. Such a concept is particularly useful in the event that the mobile node has several unverified care-of addresses that needs to be tested. By asking the mobile node to respond via another unverified care-of address, in one round trip time, the home agent would be able to verify two care-of addresses. Figure 7 illustrates this.




Figure 7: Operation of Home Agent

In Figure 7, the mobile node MN has three interfaces. A first interface associates with the cellular system and uses a care-of address CoA1 for communication. A second interface connects to the General Packet Radio System (GPRS) and uses a care-of address CoA2 for communication. A third interface associates with the 802.11 Wireless Local Area Network (WLAN) and uses a care-of address CoA3 for communication. MN sends a binding update message from CoA1 that is attempting to register all its care-of addresses at its home agent, HA. Thus, when HA receives the binding update message, it binds CoA1 in its binding cache as a valid routing path to MN as it has proven to HA that MN is using that care-of address. However, for CoA2 and CoA3, HA would need to verify their addressability before using them for packet routing. Hence, HA uses the notification to verify both of these care-of addresses. Comparing this to the fact

the other hand, if HA would to test the care-of addresses in pairs, then a total of only two messages would be required.

It is possible for the home agent to choose when to perform such verification process for the mobile node's bindings. One way is to have the home agent trigger the verification process immediately upon the reception of binding update message registering multiple care-of address from a mobile node. This is mostly helpful when it comes to supporting the concept of registering multiple care-of addresses. The purpose of having multiple care-of addresses is to allow the home agent to route packets to the mobile node via any of the multiple paths. Having the home agent perform the verification process immediately permits the home agent to quickly utilize all of the mobile node's care-of addresses for packet routing. For example, when the home agent receives the binding update message on CoA1 that specifies the binding of both CoA2 and CoA3, the home agent can choose to immediately verify these care-of addresses in order to start using them for packet routing to the mobile node. Another way could be that the home agent triggers the verification process just before it needs to use an unverified routing path to the mobile node. This is especially so in the event that the mobile node sets a filter rule at the home agent specifying the routing of certain packets via the unverified routing path. When the home agent receives packets that match the filter rule condition, the home agent would trigger the verification process for the unverified path prior to using it. An example would be that the mobile node specify in a filter rule at the home agent that data packets would be routed via CoA3. When the home agent receives such a packet, it notes that CoA3 has yet been verified. Thus, the home agent performs the verification process to verify CoA3.

Also, the transmission path of the notification packet from the home agent can vary based on how strict the security policy is enforced. Strict security policy can forbid the home agent in using an unverified care-of address for any packet routing. In this sense, the home agent can only send the notification to the mobile node via a verified care-of address. Such policies are beneficial in deterring any such flooding attacks being launch from the mobile node. For example, since CoA1 is the only verified care-of address to the mobile node, the home agent would send a message via CoA1 to notify the mobile node to respond back via CoA2. This action will cause the mobile node to understand that the home agent is trying to verify CoA2. Hence, the mobile node would send a packet to the home agent via CoA2, thereby proving to the home agent that it is

addressable at CoA2.

On the other hand, if the home agent employs techniques such as credit-based authorization (CBA) [12], it permits the home agent to route packets to the mobile node via an unverified care-of address. These packets would be limited to the amount of credits associated to that particular binding. The advantage of using an unverified path to transmit the notification is that it allows the home agent to verify at best two unverified care-of addresses. This is especially useful in the event that the mobile node has several unverified care-of addresses that needs to be tested. An example would be that the home agent sends a message via CoA2 to notify the mobile node to respond back via CoA3. This action allows the mobile node to understand that the home agent is trying to verify both CoA2 and CoA3. Hence, the mobile node would send a packet to the home agent via CoA3, thereby proving to the home agent that it is addressable at both CoA2 and CoA3.

Furthermore, it is possible that the home agent does not specify a return unverified care-of address when it notifies the mobile node. This would allow the mobile node to respond back to the home agent via any of the mobile node's care-of address. By permitting the mobile node to choose the respond path to the home agent, it supports the situation where the uplink path maybe resource constraint (e.g. limited uplink bandwidth). In the first place, this could be the reason why the mobile node decides to send a single binding update message to bind its multiple care-of address from a path that does not limit its resources. For example, the home agent wants to verify that CoA2 is addressable to the mobile node. Thus, the home agent notifies the mobile node via CoA2 to verify this care-of address. Since GPRS has limited uplink bandwidth, the mobile node decides to send the respond via the cellular path (CoA1). Hence, upon receiving the respond from the mobile node, the home agent is able to verify that the mobile node is addressable at CoA2.

Finally, how the notification information is carried to the mobile node may differ at the home agent. One possible way is that such notification is carried using a dedicated message format. This makes the notification packet small in size and easily recognized by the mobile node. On the other hand, since the notification is specifying just a care-of address that the mobile node should respond back by, an optimization to this would be to have the notification to be tagged to packets that are being sent to the mobile node. One example could be the binding acknowledgment message that the home agent replies to the mobile node. Another example could be data packets that are addressed to the mobile node. By tagging it to existing packets, the home agent saves on the overhead incurred from

the transmission the notification packet to the mobile node.

[4.](#) Conclusion

This draft explains that the use of multiple care-of address registration breaks the trust relationship between a home agent and a mobile node. With this relationship broken, it permits a malicious mobile node to bind multiple victims' care-of addresses at the home agent. With these bindings, the mobile node can launch attacks by flooding the victims with useless packets. To reduce the chances of having such a situation, we briefly analyse a few possible solutions approaches that would be able to solve the problem. The draft list each approach with their own advantages and dis-advantages to allow an implementer to decide upon which approach would be best suited for their deployment. Also, the draft attempts to provide some recommendations to reduce the dis-advantages for each approach.

[5.](#) Acknowledgements

The authors would like to thank George Tsirtsis for suggestions to improve upon this document.

[6.](#) Security Considerations

This draft mentions new security requirements when implementing Multiple Care-of Address Registration [[2](#)].

[7.](#) IANA Considerations

This document does not require any action from the IANA.

[8.](#) References

[8.1.](#) Normative References

- [1] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", [RFC 3775](#), June 2004.

- [2] Wakikawa, R., "Multiple Care-of Addresses Registration", [draft-ietf-monami6-multiplecoa-05](#) (work in progress), January 2008.

8.2. Informative References

- [3] Montavont, N., "Analysis of Multihoming in Mobile IPv6", [draft-ietf-monami6-mipv6-analysis-04](#) (work in progress),

Lim, et al. Expires January 11, 2009 [Page 17]

Internet-Draft Verifying Multiple CoA Bindings July 2008

July 2007.

- [4] Soliman, H., "Flow Bindings in Mobile IPv6 and Nemo Basic Support", [draft-soliman-monami6-flow-binding-04](#) (work in progress), March 2007.
- [5] Larsson, C., "A Filter Rule Mechanism for Multi-access Mobile IPv6", [draft-larsson-monami6-filter-rules-02](#) (work in progress), March 2007.
- [6] Mitsuya, K., "A Policy Data Set for Flow Distribution", [draft-mitsuya-monami6-flow-distribution-policy-04](#) (work in progress), August 2007.
- [7] Arkko, J., "Verifying Correctness of Alternate Care-of Address Option", [draft-arkko-mext-rfc3775-altcoa-check-01](#) (work in progress), February 2008.
- [8] Kuparinen, M., Mahkonen, H., and T. Kauppinen, "Multiple CoA Performance Analysis", [draft-kuparinen-monami6-mcoa-performance-00](#) (work in progress), April 2006.
- [9] Aura, T., "Cryptographically Generated Addresses (CGA)", [RFC 3972](#), March 2005.
- [10] Haddad, W., "Care-of Address Test for MIPv6 using a State Cookie", [draft-haddad-mext-enhanced-reachability-test-00](#) (work in progress), February 2008.
- [11] Dupont, F. and J. Combes, "Care-of Address Test for MIPv6 using a State Cookie", [draft-dupont-mipv6-rrcookie-04](#) (work in progress), January 2007.

- [12] Arkko, J., Vogt, C., and W. Haddad, "Enhanced Route Optimization for Mobile IPv6", [RFC 4866](#), May 2007.

[Appendix A](#). Change Log

- o [draft-lim-mext-multiple-coa-verify-02](#):
 - * Revamp the analysis section [Section 3](#) to highlight the advantages and dis-advantages for each solution approach. Also, recommendations are given to try and reduce dis-advantages of each approach.

Lim, et al.	Expires January 11, 2009	[Page 18]
-------------	--------------------------	-----------

Internet-Draft	Verifying Multiple CoA Bindings	July 2008
----------------	---------------------------------	-----------

- o [draft-lim-mext-multiple-coa-verify-01](#):
 - * Reference latest MCoA draft that specify the problem.
 - * Added two more solutions (limiting care-of address binding at HA and CoA test using 'symbiotic' relationship) for analysis in [Section 3](#).
- o [draft-lim-mext-multiple-coa-verify-00](#):
 - * Initial version.

Authors' Addresses

Benjamin Lim
Panasonic Singapore Laboratories Pte Ltd
Block 1022 Tai Seng Ave #06-3530
Tai Seng Industrial Estate
Singapore 534415
SG

Phone: +65 65505478
Email: benjamin.limck@sg.panasonic.com

Chan-Wah Ng
Panasonic Singapore Laboratories Pte Ltd
Blk 1022 Tai Seng Ave #06-3530
Tai Seng Industrial Estate
Singapore 534415
SG

Phone: +65 65505420
Email: chanwah.ng@sg.panasonic.com

Keigo Aso
Matsushita Electric Industrial Co. Ltd. (Panasonic)
5-3 Hikarino-oka
Yokosuka, Kanagawa 239-0847
JP

Phone: +81 46 840 5123
Email: asou.keigo@jp.panasonic.com

Lim, et al.	Expires January 11, 2009	[Page 19]
-------------	--------------------------	-----------

Internet-Draft	Verifying Multiple CoA Bindings	July 2008
----------------	---------------------------------	-----------

Suresh Krishnan
Ericsson
8400 Decarie Blvd.
Town of Mount Royal, QC
Canada

Phone: +1 514 345 7900 x42871
Email: suresh.krishnan@ericsson.com

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS

OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.