INTERNET DRAFT

P. Mutaf
C. Castelluccia
INRIA

February, 2002

IP Paging Threat Analysis
<draft-mutaf-paging-threats-00.txt>

Status of This Memo

Abstract

This document is an analysis of threats that arise from using link layer paging technologies or IP paging in the Internet where denial-of-service attacks are common and easy. These problems fall in the scope of IP paging, since link layer paging technologies do not have provisions for repelling such threats and the source of an attack may be anywhere in the Internet. In addition, vulnerabilities that may be added by IP paging are also discussed.

Table of Contents

1. INTRODUCTION

Currently, the unique role of the IP routing system is routing packets to their destinations. IP paging adds a new responsibility to the routing system: detecting communication attempts of behalf of hosts which have limited energy. As a result, hosts will rely on this new service for their accessibility. Attackers may want to exploit this for rendering hosts inaccessible.

Secondly, the network will offer this service at the considerable bandwidth cost of "paging", assuming that this will not be too frequent. In the Internet, this assumption can be

abused by attackers. Paging which is intended for optimizing bandwidth use, may become a bandwidth threat. This problem arises from adding the paging functionality (existing link layer paging technologies or IP paging) to the operation of the Internet. Hence it is independent of IP paging. However, this problem falls

in the scope of IP paging since existing link layer technologies do not attempt to solve this problem. In addition, the attackers may be anywhere in the Internet, hence IP layer solutions may be necessary.

Finally, in an IP network, attacks that drain a target host's battery by forcing it to continuously receive packets are easy. This problem also falls in the scope of IP paging since currently no other mechanism is intended for protecting battery.

RFC 3154 [2] outlines some of these threats. This document is a detailed analysis of the techniques that can be used by attackers for implementing attacks that exploit and sabotage the IP paging service. Additionally the challenges of IP paging security and possible solutions are discussed.

2. TERMINOLOGY

Host (H) — An Internet host which implements IP paging.

Correspondent Node (CN) — A node residing anywhere in the Internet.

Attacker, attacking host or attacking CN — A rogue CN or H. The attacker generally uses forged (random) source addresses, except when impersonating another node.

Paging — IP paging or link layer paging technologies unified by IP paging as described in [1].

Paging Area — A set of cells where a host is likely to be found (subnet->cell mappings are arbitrary) as described in [1][2].

IP Paging functions — Dormant Monitoring Agent (DMA), Tracking Agent (TA), Paging Agent (PA) as defined in [2].

Dormant Mode Binding (DMB) — The association between a dormant host's state and its identity.

Wait-For-Sleep (WFS) timer — Period of time a host waits before renewing its DMB (when the host is not in active IP communication).

T The value of the WFS timer.

C	Attacker's DoS (Denial-of-Service) capacity (in packets/second). Limited by the available bandwidth to the target region and processing power.
N	The number of DMBs exploited by the attacker.
S	Paging area size (in number of link layer cells). For simplicity it is assumed that all paging areas have the same size. S is the amplification factor of paging, since one IP packet triggers S link layer frames each received by a different access point (AP). This factor is independent of IP layer concepts such as "subnet".
G_dos	Malicious signaling gain of paging.

Additionally, the definitions of Mobile IP and Mobile IPv6 terms such as home agent, home address, home subnet, care-of-address (CoA), can be found in [3][4].

3. ASSUMPTIONS

In the following threat analysis, the following assumptions are made:

- A global security infrastructure is not necessarily available. Hence, authentication and authorization of CNs residing anywhere in the Internet is difficult.
- The details of the IP paging protocol are not defined yet. According to recent IP paging proposals, the IP paging functions may be implemented on a single network element or separately in different network elements [5][6][7][8][9]. Therefore, the cases where the IP paging functions are in different network elements, are also analyzed.

These assumptions imply that strong authentication and authorization of CNs is difficult. Secondly, if two IP functions are implemented in different network elements residing in different administrative domains, the authentication and authorization of these elements may be difficult.

4. MALICIOUS PAGING SIGNALING IN IP NETWORKS

An attacker can exploit link or IP layer paging for amplifying the impacts of DoS attacks that reduce the available bandwidth on a

target cellular region. The amplification factor is due to the paging process which consists of searching a destination dormant mode host in multiple cells.

As mentioned above, the bandwidth efficiency of the paging service is based on the assumption that hosts do not receive incoming sessions most of the time (in the order of 90% of time). This assumption has proven true for years and paging has been successfully keeping an equilibrium between paging and registration signaling. However, in a IP network where DoS attacks are very common, this assumption will be probably abused. Attackers may ruin the natural equilibrium between paging and registration signaling. The paging service intended for optimizing the bandwidth use, may become not only sub-optimal but also a bandwidth threat.

Paging support is already available in current link layer technologies, hence the amplification factor mentioned above will exist regardless of whether IP paging is deployed or not. The additional threat rather arises from the deployment of an all-IP cellular system with millions of IP hosts. First, in an IP network, sending a single packet is interpreted as a communication attempt (e.g. a TCP SYN segment). If the destination host is dormant, then that packet should initiate paging. An attacker can generate millions of malicious packets, each initiating paging.

In conclusion, in an all-IP cellular system, the amplification gain of paging can be exploited much more easily. This problem falls in the scope of IP paging, since link layer paging technologies do not have provisions for defending against such attacks. Furthermore, the source of a DoS attack may be anywhere in the Internet, hence IP layer solutions may be necessary.

As a result, it is crucial to understand the methods that attacker can adopt for launching DoS attacks that exploit paging and define solutions if they are indeed feasible. Below is the analysis of the possible techniques for implementing bandwidth attacks that abuse the paging service.

4.1. Exploiting Legitimate DMBs

A malicious CN residing anywhere in the Internet can send many packets to many different dormant hosts on a target region. Each malicious packet will be intercepted by one or more DMA(s) which informs one or more PAs. The PA(s) will in turn create signaling for paging the destination host in S cells. In other words, one malicious packet will create S link layer frames each received by a different AP in the paging area. It is noteworthy that the malicious packets have random and different source addresses. Thus, limiting the number of simultaneous pages per CN, does not help. The intervening DMA(s) cannot differentiate between malicious and legitimate packets. Hence, every packet (malicious or legitimate) initiates paging.

S is the amplification factor, however this is not necessarily equal to the malicious signaling gain. When paging is exploited, the attacker can not transmit at a rate higher than N/(T+k) packets/second (where, k is DMB update latency). This is because

the attacker has to wait for a previously paged host to renew its DMB. This limits the intensity of the attack. Therefore, the overall malicious gain of this technique should be defined as:

$$G_{dos} = \frac{N \times S}{(T + k) \times C} \quad ; \quad N/(T+k) <= C$$

The actual value of T is currently an open issue. Its exact definition is challenged by packet arrival irregularities observed in datagram networks. If T is too small a host may unnecessarily enter IP dormancy upon unimportant but frequent bandwidth degradations during a same session. If T is too large, this will reduce the benefit of paging since the host will enter IP dormancy too late. This issue is out of the scope of this document.

The only parameter which is under the attacker's control is N. Therefore whether this technique arises malicious signaling threat depends on an attacker's ability to locate relatively large numbers of dormant mode hosts. If the attacker has high DoS capacity but is not able to locate a large number of dormant hosts, then this attack may result in loss (i.e., G_dos < 1). This is because the attacker will not be able to reap the benefit of high DoS capacity.

The attacker's capacity is limited to C packets/second. Then, a full capacity attack will require C x (T + k) dormant mode bindings. For example, if the attacker is capable of pumping 20,000 packets/second, and if T is 10 seconds, then the attacker has to locate approximately 200,000 dormant hosts in order to obtain the maximum gain.

In a cellular system, hosts are in dormant mode most of the time (in the order of 90%). Therefore, rather than searching or monitoring dormant mode binding updates, the attacker can simply send many packets to many different hosts in a cellular domain. Probably, most of the destination hosts will be in dormant mode, hence N will be large. Therefore, the attacker's problem is reduced to finding host identifiers.

The home addresses and DNS names are attractive host identifiers since they are they are permanent, so the cellular host address collection work needs to be done only once. Home addresses may be even more attractive since they are transmitted over the network in home address options or in routing headers as clear-text. As a result they can be also detected by sniffers.

The attacker may probe many home addresses by sending ICMP Echo Request or other packets. Locating a home subnet will be easy.

However, guessing valid suffixes among 2^64 possibilities on a subnet may be a difficult task if home addresses are unpredictable e.g. configured using privacy extensions[10], or more similar techniques. Another possibility is to install a network sniffer placed strategically on a link across which many mobile nodes'

packets are routed. This can allow the attacker to obtain many home addresses (found in routing headers and home address options). Privacy extensions could solve this problem since they are intended for not revealing the true home addresses of hosts. However, this ensures privacy during sessions initiated by mobile nodes. A sniffer placed on the path between correspondent nodes and the home subnet can help detect many sessions destined for many mobile nodes, hence their home addresses. Alternatively, the attacker can exploit the DNS by launching a brute force analysis on the name space of a cellular service provider. There may be a continuity in the naming pattern for efficient use name space and the attacker can easily obtain many corresponding home addresses.

An important factor which will challenge the attacker is the density of the destination hosts of which the dormant mode bindings are exploited. This can be defined as the number of destination hosts compared to the size of the target region (in number of cells). If the destination hosts are far away from each other, their paging areas may not superpose and the impacts of the attack may not be felt. In order to ensure high destination host density, the attacker can benefit from host location predictability. For example, a majority of the hosts served by a same home agent may be owned by users living in the same region. More importantly, the attacker can sample the CoA of many hosts using their home addresses, get an idea of their whereabouts and select the ones which move more or less in the same region. Probably, this information will remain valid for longtime (possibly years) unless users physically move to other regions, which is unlikely.

4.2. Creating Fake DMBs

Alternatively, the attacker can organize a two-party attack where a malicious host creates dormant mode bindings each pointing to a different and "fake" host, and another malicious host (possibly under the control of the same attacker) transmits page trigger packets. Assuming that both hosts transmit C packets/second (not necessarily in a synchronous fashion), it is possible to obtain a much more important malicious gain:

$$G\_dos = S \times P$$

where P is the number of times a host should be paged before it can be assumed unreachable. This factor is due to the fact that the attacker probably does not reply to pages. The T factor is omitted since the attacker does not need to use a dormant mode

binding twice.

Fake dormant mode bindings can be directly created by a malicious cellular host. However, it may be also possible to spoof the TA->DMA traffic (if TA and DMA functions are implemented in different network elements). The attacker can impersonate a real

TA, or imitate the TA function, then send one or more DMAs many packets that report fake hosts entering dormant mode.

The attacker may be also able to redirect the malicious paging traffic to a target region. In the RFC3154 architecture, the paging area information is provided by the hosts. In this case, the attacking host can issue many dormant mode binding updates as described above, but pointing to one or more paging areas away from its actual location.

4.3. Creating Fake Paging Requests

The attacker may impersonate a real DMA or imitate the DMA function and request one or more PAs to page fake hosts. The advantage of this technique is that, the attacker does not need to locate nor create dormant mode bindings. The malicious gain is:

$$G\_dos = S \times P$$

for the same reasons described in the previous section.

5. ACCESSIBILITY THREATS

This section is an analysis of possible malicious techniques that exploit IP paging for rendering hosts inaccessible.

The attacks discussed below consist of impersonating one of the functional entities, or imitating the IP paging functions such as TA and DMA. Impersonation of an authorized functional entity can be detected by authenticating its IP address. Whereas, detecting an attacker acting as a TA for example, requires strong binding between the TA's identity and its authorization to act as a TA.

In the following discussion, it is assumed that the attacker is able to monitor the packet exchanges between the functional entities, inject packets for impersonating or imitating them, but not capable of dropping packets.

5.1. Imitating the IP Paging Functions

If informing the DMA when a host enters dormant mode is under the responsibility of TA, an attacker can pretend TA ability to visiting hosts (assuming that hosts are dynamically informed of the existence and addresses of TAs). A victim host believing to have registered a valid dormant mode binding may become completely

unreachable until it obtains a valid binding.

The impacts may be much more important than a case where a single access router (AR) is compromised and shut down. A victim host may

be unreachable in other subnets as well (covered by its current paging area). The unreachability duration may be extended up to several hours depending on whether the host frequently crosses paging area boundaries. The threat is considerable if the host does not change its paging area for longtime. Depending on the paging policy, the host may be spoofed repeatedly by the same attacker still pretending TA ability.

Similarly, if hosts are dynamically informed of the existence and addresses of DMAs. An attacker imitating DMA function can spoof visiting hosts believing to register their DMBs with an authorized DMA. The impacts are the same as above.

5.2. Impersonating the Functional Entities

An attacker on the path H->DMA or DMA->H can monitor many DMBs and host identities. It is also possible to exploit the fact that a given host is probably in dormant mode and guess the address of its DMA. Then, the attacker can impersonate a dormant host and send a dormant mode binding deregistration request. Upon reception of this request the DMA will remove the dormant mode binding of the host or change it to an active mode binding. The host which is still in dormant mode (hence, accessible only via paging) will become unreachable. Unreachability duration may be extended to several hours as described above. It is important to note that the acknowledgement of DMA sent to the host will be probably lost, since it is not a paging request nor sent to a PA. Alternatively, if the attacker is able to monitor the H->TA or TA->H traffic, it is possible to impersonate a host and send its TA a message indicating that the host changed its paging area. Upon receipt of a packet from a CN, the host will be paged in a wrong paging area, hence unreachable. The duration of unreachability may be extended to several hours.

Secondly, an attacker on the path DMA->PA, can monitor the paging request messages sent by the DMA. Then, the attacker can impersonate the PA and send the DMA an immediate negative response indication for a host which is probably being paged (in the mean time) by the real PA, but the host has not yet replied due to paging delay. Upon reception of this message, the DMA may return an ICMP "Destination Unreachable" message to the CN. At that point, the host may reply to its real PA, which will in turn send a positive response indication to the DMA. However, the CN's application or transport layer will have already interrupted the session on receipt of the ICMP error message.

Finally, an attacker on the path DMA->TA, can monitor the messages

sent by the DMA informing the TA that a packet has arrived for the dormant host. Then, by impersonating the TA, the attacker can reply to the DMA with a message indicating a wrong PA's address. The host may be paged in a wrong paging area, hence become unreachable.

6. BATTERY DRAINING ATTACKS

The attacks described in this section consists draining the battery of a target host by injecting packets to its link layer. End-to-end authentication does not solve this problem since the target is forced to receive packets in order to check their authenticity. Attacks that flood a target host with outstanding requests already exist. However, when the target is battery powered, this arises a more important threat: the host will become unreachable and incapable of running applications until the user has the possibility to re-charge it's battery. This problem falls in the scope of IP paging since no other mechanism is intended for protecting battery.

Below is the analysis of the malicious techniques that an attacker can adopt for draining a target's battery. It is assumed that the attacker and target reside in different subnets.

6.1. Awakening a Dormant Host

An attacker can periodically awaken a host by sending 1/T packets/second. Each packet will be intercepted by a DMA, which will in turn start the process of paging the target host. This attack is not different than the bandwidth attack described in Section 4.1. The malicious goal is different but the impacts are the same.

6.2. Flooding an Active Host

An attacker can also send an active host C packets/second in order to prevent that host from beginning IP dormancy. Malicious packets may reset the WFS timer, hence the target host may not be capable of beginning IP dormancy.

6.3. Flooding a Dormant Host

6.3.1. Target is Stationary

An attacker may try to flood a dormant host by injecting packets to its last known subnet. In this section it is assumed that the target host has not moved to another subnet and the attacker has knowledge of the last CoA used by the target before entering IP dormancy (i.e. the attacker is capable of monitoring its traffic).

The attacker's packets may trigger frames destined to the target's link layer if the neighbor cache hold by the AR still contains a mapping between the host's last configured CoA and link layer address.

If a paging channel support is available, the target host's link layer will be no longer listening to the traffic channel, then the attack is defeated.

In time-slotted dormant mode, the frames triggered by the attacker's packets will be received by the host's link layer forced to continuously listen to the traffic channel, hence consume battery. The host's link layer cannot switch to or remain in low-power mode.

6.3.2. Target is Mobile

The target host may move to another subnet while the attack continues. As time passes, the attacker's uncertainty of the exact subnet of the target will increase. However, the attacker can try to flood the dormant host in several subnets where it is likely to be found.

The AR in the new visited subnet will receive the attacker's packets and trigger Neighbor Discovery. If a paging channel support is available, the target host's link layer will not be listening to the traffic channel, hence the attack is defeated.

In time-slotted dormant mode, the Neighbor Discovery packets will be buffered by the APs. Next time the target host's link layer wakes up, it will receive this packet. At this point, if the host's IP module is ON and if it has configured a CoA on that subnet (and given that the host's suffix remains the same), the host will reply to the request, revealing its link layer address to the AR. Then, the subsequent link layer frames triggered by the attacker will be continuously received by the host's link layer, which can no longer save power.

6.4. Impersonating or Imitating a PA or DMA

The attacker may try to impersonate a PA or imitate the PA function in order to page the target host.

Alternatively. the attacker may try to impersonate a DMA or imitate the DMA function and send paging request packets to the target's current PA.

7. DISCUSSION OF SOLUTIONS

7.1. Defeating Attackers Residing Anywhere in the Internet

In the threats 4.1 and 6.1 the attacker is a malicious CN anywhere in the Internet. Source addresses of malicious packets are random, hence intervening DMAs cannot differentiate between malicious and legitimate packets.

These attacks can be defeated by the authentication and authorization of CNs before starting the paging process. However, it is assumed that a global security infrastructure is not available, and pre-shared trust relationships with millions of CN anywhere in the Internet will not scale. Hence, strong authentication and authorization using certificates that cryptographically prove identity and access rights are not feasible. Therefore other solutions are necessary.

7.1.1. Weak Authentication Embedded in the IP Paging Protocol

Assuming that an attacking CN will not want to reveal its real IP address, a level of security near to strong authentication can be obtained by weak authentication (it is noteworthy that this does not have any bearing on the access rights of a CN). By current DoS practice, the assumption that attacker use random source addresses, generally holds.

In this scheme, the intervening DMA sends a special message to the CN. This might be an ICMP "Destination Dormant" message containing a cookie or puzzle request. Then, the session is allowed by the DMA if only the correspondent node satisfies the authentication rule (correct cookie or puzzle reply). The definition of this protocol extension is out of the scope of this document. Other methods for implementing this DMA<->CN handshake prior to starting the paging process, may be also possible.

The main problem of this approach is that it requires the modification of CNs. Additionally, care should be taken in order not to add any vulnerability to the CN side. There may be also privacy problems. Whether a given user is in active communication at a given time is possibly important information. This approach may be problematic if upper layers do not already reveal such information.

Points in favor of this approach are simplicity, independence of any paging policy, and the existence of administrative or per-host security policies (i.e. a level of security is embedded in the IP paging protocol). There are also motivations other than security. For example, the diagnostic tools such as "ping" can be changed such that it does not unnecessarily cause paging cost and give information of the "actual state" of a destination host without changing its state (i.e. dormant).

7.1.2. Host Defined Access Control & Adaptive Paging

Another approach can be the enforcement of per-host access control with the aid of adaptive paging area sizes. This approach separates the threats 4.1, 4.2, 4,3 and 6.1 into:

- o Threats 4.1, 4.2, 4,3 Bandwidth attacks (network's problem)
- o Threat 6.1 Battery draining attack (host's problem)

In adaptive paging, each time a host is paged, its paging area size is reduced (on the contrary, it is augmented each time the host crosses the boundaries of a paging area). In [12], paging area sizes are controlled hosts. However if paging area sizes are controlled by the network, the bandwidth attacks described in Section 4 can be defeated, since paging area sizes of the destination hosts will be eventually very small (due to very frequent paging, compared to host movement rate).

In threat 4.1, the hosts of which the dormant mode bindings are exploited will have excessively small paging areas, hence can not reap the benefit of paging. This approach says that this is the hosts' problem (just like the threat 6.1). If a host desires protection against these attacks, then this host should register access control rules with its DMA in order to reliably benefit from the paging service.

This approach has three problems. First, access control interferes with a host's ability to receive packets from new legitimate CNs. Hence, it is not always desirable. Secondly, per-host access control may not scale well if each host has many trust relationships with many CNs. Third, there exists a possible counter argument saying that paging is a service offered by the network (just like routing), hence paging area availability and helping preserve energy should be under the responsibility of the network. The network should not penalize a given host for not enforcing access control.

The main argument in favor of this approach is that, bandwidth can be protected even if the attacking CN uses its real IP address. Secondly, the motivation of paging area size optimization is not limited to security, it is by all means advantageous (however, this is more a paging policy issue).

7.2. Authentication and Authorization of Hosts and Functional Entities

The threat discussed in Section 4.2 requires the authentication of hosts so that a host can not have more than one DMB registered with the network. This may be possible if an AAA infrastructure exists. However, whether IP paging should rely on the existence of AAA is

not clear.

Defending against the accessibility threats discussed in the Section 5 and the bandwidth threat discussed in Section 4.3, requires the authentication and authorization IP paging functional

entities when communicating with each others. This may be difficult however it is hard to exactly define the problem before the IP paging protocol is designed.

Defending against the threats discussed in Section 5.1 requires that hosts verify the authorization of DMA and TAs before registering their dormant mode bindings. However, the exact problem is hard to define exactly, before the IP paging protocol designed.

7.3. Defeating Battery Draining Attacks

The threat 6.1 is already addressed in Section 7.1. In the following discussion it is assumed that this problem is solved.

The most serious battery draining threat is 6.2, i.e. flooding an active host in order to prevent it from beginning IP dormancy. It is desirable to have an IP paging protocol that allows hosts to begin IP dormancy under attack. If in IP dormancy a host's battery is protected, in this case the battery draining threat can be solved.

However, beginning IP dormancy under attack may be difficult. With IPsec [11], this might be possible depending on how IP paging is implemented. If packets dropped by IPsec do not reset the WFS timer, then the IP paging module can timeout and begin IP dormancy. Otherwise IPsec can not help begin IP dormancy under attack. This implies that, if the WFS timer is implemented below IP (i.e. at the link layer), IPsec is not helpful. Other solutions may be possible.

The attacks 6.3 and 6.4 can be defeated to some extent if the dormant mode binding update message is encrypted and contains a random 64 bits IPv6 suffix (as part of the encrypted payload) that will identify the host in dormant mode. The advantage of this scheme is that, the defense is independent of how IP paging is implemented i.e., whether IP is ON in IP dormancy, or hosts configure CoAs in dormant mode (in this case the host would use its secret suffix for configuring CoAs in dormant mode).

This scheme ensures that, only an authorized party can awaken the host (given that the host's current DMA is authorized). An attacker on the path H->DMA can not obtain the secret suffix that identify the target host. It is also noteworthy that no trust relationship is necessary between a given PA and the ARs in its paging area. An attacker can not impersonate a PA nor imitate the PA function in order to awaken a target host, since

the dormant host's identity can be only provided by an authorized DMA. This also prevents an attacker from imitating the DMA function in order to awaken a dormant host.

However, DMA's authorization should be verified by the host before the dormant mode binding is registered. This may be difficult as mentioned in the previous section.

In summary, battery security can be ensured by the following three mechanisms:

1) Host's capability of starting IP dormancy under attack,

2) Host's capability to check that a given DMA is authorized,

3) Host's capability to secretly share with a DMA its dormant mode identity (e.g. suffix).

8. CONCLUSION

This document analyzed the threats that arise from using paging or IP paging in the Internet. Three classes of attacks were described. Attacks that exploit paging for degrading bandwidth on a target cellular region, attacks that drain a target host's battery and attacks that impersonate or imitate the IP paging functions for rendering hosts inaccessible.

The bandwidth attacks that exploit paging are especially difficult to deal with, since the attacker can be anywhere in the Internet. DMAs cannot differentiate between malicious and legitimate packets, hence every packet will initiate paging. A global security infrastructure is not available, and pre-shared trust relationships with millions of CN anywhere in the Internet will not scale. This attack, many seriously reduce the motivation of using paging in the Internet. Paging is intended for optimizing battery. However, if it is exploited well, it may become a bandwidth threat.

The above threat has also important implications on battery consumption on target hosts'. The remaining battery draining threats can be possibly defeated more easily. This document discussed several solutions, others may be possible.

IP paging may also add new accessibility threats. Depending on how the IP paging protocol is designed, these problems may be difficult to deal with. However, these issues are currently not clear since the interactions between the IP paging functions are not yet defined.

9. SECURITY CONSIDERATIONS

This document is a security analysis of IP paging.

10. RELATED WORK

RFC3154 also mentioned a paging buffer overflow threat. For implementation reasons there may be an upper limit on the number of concurrent sessions on behalf of which a destination host can be paged. This upper limit will equal the size of a buffer where the first packets of the session initiating peers are hold (by the dormant monitoring agent). RFC3154 assumes an aggregated buffer, resulting in important threats. However, as will be shown below, when this buffer is per-host the threat is not serious.

An attacker can send many packets (with different source addresses) to a destination host which is in dormant mode, in order to overflow the buffer where its concurrent session requests are hold. The main factor is the paging latency, i.e. the time period between initiating the paging process (by the dormant monitoring agent) and receiving its reply. This delay may be important. During this period, if the paging buffer is full, all packets from legitimate correspondent nodes will be lost. However, the first packet of the attacker will initiate paging and the host will send a MIPv6 binding update to its home agent. Then, the next packets of the correspondent nodes will reach the host.

When paging buffers are per-host, this attack can be serious only when combined with the congestive impacts of another DoS attack causing packet loss (in particular, loss of paging and paging reply packets). In this case, the host may become unreachable instead of suffering from degraded performance. The vulnerability added by IP paging is the paging delay which is much more important than the processing delay of an ordinary IP router. As a result, additional risk exists, but it is difficult to exploit.

11. ACKNOWLEDGEMENTS

The authors would like to thank James Kempf who has reviewed the threat analysis, and also Yoshihiro Ohba and Erik Anderlind for their valuable comments on the ICMP "Destination Dormant" message proposal during Seamoby mailing list discussions. In particular, Erik Anderlind pointed out the possible privacy problems with this approach.

REFERENCES

[1] Kempf, J., Editor, "Dormant Mode Host Alerting ("IP Paging") Problem Statement", RFC 3132, June, 2001.

[2] Kempf, J., et. al. "Requirements and Functional Architecture for an IP Host Alerting Protocol," RFC 3154, August, 2001.

[3] Johnson, D. B. and Perkins C., "Mobility Support in IPv6", Work in Progress, July 2001.

[4] Perkins, C., "IP Mobility Support for IPv4", RFC 3220, January 2002.

[5] Faccin, S., et. al., "Dormant Mode Handover Support in Mobile Networks," draft-koodli-paging-00.txt, Work in Progress.

[6] Liebsch, M., Renker, G., and Schmitz, R., "Paging Concept for IP based Networks," draft-renker-paging-ipv6-01.txt, Work in Progress.

[7] Ohba, Y., Nakajima, N., and Zhang, T., "LH-DMHA - Last Hop DMHA (Dormant Mode Host Alerting) Protocol," draft-ohba-seamoby-last-hop-dmha-02.txt, Work in Progress.

[8] Sarikaya, B., et. al., "Mobile IPv6 Hierarchical Paging," draft-sarikaya-seamoby-mipv6hp-00.txt, Work in Progress.

[9] Gurivireddy, S., et. al., "Layer-2 aided mobility independent dormant host alerting protocol," draft-guri-seamoby-lahap-00.txt, Work in Progress.

[10] Narten T. and Draves R., "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", Work in Progress, July, 2001.

[11] Kent S. and Atkinson R., "Security Architecture for the Internet Protocol", RFC 2401, November 1998.

[12] Castelluccia C. and Mutaf P., "An Adaptive Per-Host IP Paging Architecture", ACM SIGCOMM CCR Special Issue on Wireless Extensions to the Internet, October 2001.

AUTHORS' ADDRESSES

Pars Mutaf
INRIA Rhone-Alpes
655 avenue de l'Europe
38330 Montbonnot Saint-Martin
FRANCE

email: pars.mutaf@inria.fr
phone: +33 4 76 61 55 07
fax: +33 4 76 61 52 52

Claude Castelluccia
INRIA Rhone-Alpes
655 avenue de l'Europe
38330 Montbonnot Saint-Martin
FRANCE

email: claude.castelluccia@inria.fr
phone: +33 4 76 61 52 15
fax: +33 4 76 61 52 52