

comp.os.vms: Re: Process's PreciseMail AntiSpam Gateway – any experience so far ?

Re: Process's PreciseMail AntiSpam Gateway – any experience so far ?

Source: <http://unix.derkeiler.com/Newsgroups/comp.os.vms/2003-09/2184.html>

From: Don Sykes (*anonymous_at_pacbell.net*)

Date: 09/27/03

Date: Sat, 27 Sep 2003 19:27:21 GMT

John Santos wrote:

>
> On Fri, 26 Sep 2003 david20@alpha2.mdx.ac.uk wrote:
>
>> In article <vn5s22g2nrds0e@news.supernews.com>, "John Vottero" <John@mvpsi.com> writes:
>>> <david20@alpha2.mdx.ac.uk> wrote in message
>>> >news:bkuhsq\$dk3\$1@news.mdx.ac.uk...
>>>> In article <3F71D664.D92AAC37@pacbell.net>, Don Sykes
>>>> <anonymous@pacbell.net> writes:
>>>>>
>>>>> >david20@alpha2.mdx.ac.uk wrote:
>>>>>>
>>>>>> In article <3F70934A.3C36DD45@pacbell.net>, Don Sykes
>>>>>> <anonymous@pacbell.net> writes:
>>>>>>>
>>>>>>>
>>>>>>> >At this point I'm fairly convinced that the implementation of fees via
>>>>>>> >central gateways &/or routers is not workable. So I have come up with a
>>>>>>> >protocol that implements e-mail in 2 phases: a meta phase and a data
>>>>>>> >phase. In phase 1, all the info about the email is sent to the open,
>>>>>>> >listening port of the receiver. Then the link is dropped, by both.
>>>>>>> >Phase 2 must be initiated by the receiver, so they are in complete
>>>>>>> >control of the transmission and final delivery and at that point they
>>>>>>> >can also charge a fee.
>>>>>>>>
>>>>>>>> >A first draft is available at <http://alphase.com/vms/FBEProtocol.html>
>>>>>>>>>
>>>>>>>>> >Serious suggestions are more than welcome, but please no nit-picking.
>>>>>>>>>> >This is a early, early draft. A suggestion, if you will
>>>>>>>>>>>
>>>>>>>>>>>
>>>>>>>>>>> >I haven't had a chance to look at your link yet but one thing strikes me
>>>>>>>>>>>> >about
>>>>>>>>>>>>> >your suggestion straight away. How are you going to deal with Natt'd
>>>>>>>>>>>>>> >clients ?

Re: Process's PreciseMail AntiSpam Gateway – any experience so far ?

comp.os.vms: Re: Process's PreciseMail AntiSpam Gateway – any experience so far ?

> > > > *If you drop the connection then there is no guarantee that the public
> > > address
> > > that the sender first used will still be valid when the receiver tries to
> > > reopen the connection.
> > >>
> > >
> > > There is no need to be concerned about NAT. This proposal is a replacement
> > > for SMTP servers. They already need special consideration when used with
> > > NAT, as do all listening servers.
> > >
> > >
> > >
> > It does matter because your example which uses a client on the 10 address space
> > (10.11.12.1) contacting a server on the 21.22 network will not work in general.
> >
> > (As an aside the address of the receiver 21.22.0.0 is invalid since it is a
> > network address – you should never use .0 or .255 in the final octet).
> >
> > The 10 address is a private address hence must use NAT to contact systems on
> > the public internet.
> >
> > So in the real world you have a client on a small home network connecting to an
> > ISP using dynamic NAT with port overloading.
> >
> > 10.11.12.1 is the clients real address and it opens a connection from its port
> > 32100 this is mapped to 21.22.5.20 port 7521 on the public side of his home
> > NAT/firewall. (21.22.5.20 is the single public address given out to this user by
> > his ISP).
> >
> > This connection connects to the IPS's receiver on 21.22.0.10 (10 rather than 0
> > to make it a valid address) for your phase 1.
> >
> > Negotiation proceeds as you describe on your link and the receiver sends back to
> > say it will contact the sender on port 1398. Then the link is dropped.
> >
> > 10.11.12.1 listens on port 1398.
> >
> > Receiver (21.22.0.10) attempts to open connection to 21.22.5.20 on port 1398.
> > Attempt fails. There is either no entry in the NAT mapping table for
> > 21.22.5.20 port 1398 or if there is it would be accidental and might point at
> > another machine or port on the user's home network.
> > The connection is dropped.
> >
> > With dynamic NAT with port overloading (which is the most common form of NAT
> > used on home networks where the home user has multiple machines hiding behind
> > one external address) there is no preservation of port numbers – unless a port
> > number has been placed in the NAT mapping table by an internally initiated
> > connection to an external machine having been made or by the user explicitly
> > setting up a manual mapping then an externally initiated connection cannot
> > be made to it.
> >
> >*

Re: Process's PreciseMail AntiSpam Gateway – any experience so far ?

comp.os.vms: Re: Process's PreciseMail AntiSpam Gateway – any experience so far ?

> > *Your system falls apart.*
>
> *I don't think I understand the point of the second, reverse connection[1],*
> *but in any case, there is an alternate method that might work.*
>
> *Instead of negotiating a port number for the 2nd connection, you could*
> *use a well-known port that the NAT firewall forwards to the client*
> *system, and negotiate a magic key (perhaps encrypted). The client*
> *listens on the well-known port, the server establishes the call-back,*
> *and sends the key. The client can accept the key and continue or*
> *reject it and close the connection.*
>
> *Both ends could be behind NAT'ing firewalls, and this would still*
> *work, since the client could accept connections from anywhere, then*
> *immediately close those that didn't provide a valid outstanding*
> *key.*
>
> *Both sides would know who they were talking to because of the key*
> *exchange.*
>
> *If you have multiple systems behind a NAT firewall, each could be*
> *permanently assigned a port, which would be port-forwarded to its*
> *well-known port, and that port could be included in the negotiations,*
> *so the server would know what port to connect to. That port could*
> *be statically configured in the NAT box and the client system, or*
> *could be configured dynamically using an extension to DHCP.*
>
> *[1] Is this second reverse connection just so the client and server*
> *can go away for awhile to meditate on whether to allow the message*
> *to propagate, or is it so the actual data message can be postponed*
> *to a less busy time, if necessary (middle of the night), or is it*
> *just to allow the server some measure of "control" over the connection?*

It's for ALL 3 reasons and more. But it doesn't allow "some measure of control", it gives complete control to the receiver, which is a major facet of the FBEM protocol. If the receiver decides for ANY reason it doesn't want to deal with the incoming request, it doesn't have to do anything and the message dies on the vine.

* please post future comments on this issue in the thread titled: Fee Based Email (From Re: Process's PreciseMail AntiSpam...) TIA

--
Have VMS, Will Travel
Wire paladin, San Francisco
(paladinATalphaseDOTcom)