

Re: OpenSSH, Telnet, Windows Authentication and double-hops

Source: <http://www.derkeiler.com/Newsgroups/comp.security.ssh/2007-10/msg00061.html>

- *From:* "Richard E. Silverman" <res@xxxxxxxx>
 - *Date:* 12 Oct 2007 13:56:47 -0400
-

"JM" ==
jmartzoo-google@yahoo
com
<jmartzoo-google@xxxxxxxx>
writes:

JM> On Oct 11, 10:03 pm, "Richard E. Silverman" <r...@xxxxxxxx> wrote:
>> > We're looking for a solution to create a secure single-signon >
>> deployment on a Windows network. The chain of connections looks
>> like > this:
>>
>> > Client: Telnets through SSH Tunnel to --->
>>
>> Does this mean that you are setting SSH port forwarding (e.g. with
>> ssh -L)

JM> Exactly, yes. On the client I execute the two following commands
JM> in separate DOS console windows in this order:

JM> 1) ssh -v -N -L 2023:server1:23 server1 2) telnet localhost
JM> 2023

You're doing this -- instead of just logging in directly with SSH --
because of the problems you're having described here:

http://groups.google.ca/group/comp.security.ssh/browse_thread/thread/c3fb90fd3b747553/#

.... correct?

>> and using telnet through that (that's what "tunnel" usually means),
>> or that you are actually logging into Server1 via SSH? Sometimes
>> people use "telnet" as a generic word referring to any kind of
>> remote login (which it isn't). I'm assuming the latter, because
>> otherwise your description doesn't make sense to me.

JM> No, I'm truly using a telnet client.

Re: OpenSSH, Telnet, Windows Authentication and double-hops

>>> Server1: that runs our application locally and connects to ---->
>>> Server2: that serves up the database using SQL Server
>>
>>> We're hitting an issue when we reach the SQL Server machine.
>> Logging > in to SQL Server, the network has deprecated our logon
>> down to 'NT > AUTHORITY\ANONYMOUS LOGON' . The database kicks us
>> out.
>>
>>> We've learned that this is commonly referred to as the
>> "double-hop" > issue and is well known with web development. There
>> are mechanisms in > IIS to set up delegation and impersonation and
>> caching etc, to get > past this.
>>
>>> We want to continue using public/private key authentication for
>> the > SSH.
>>
>> You will not be able to, at least not alone. SSH publickey
>> authentication does not provide the other side with either a
>> Kerberos ticket, or your password in order to acquire one. You
>> need GSSAPI/Kerberos authentication with ticket forwarding
>> ("delegation").

JM> I'm going to look into this right away. I'm aware of Kerberos,
JM> but have never seen the term GSSAPI before.

GSS is a protocol for abstractly negotiating security services. These days, Kerberos is usually used as a GSSAPI method, so that it can be made available to any GSS-aware application via a system shared library, rather than the application needing explicit Kerberos support.

JM> Is the 'delegation' a
JM> configuration that needs to be made on the Kerberos service? I
JM> expect that this will need to be done at the domain layer, is this
JM> correct? Will we have to rely on someone with the authority to
JM> make changes at the domain level to be able to do this? I've read
JM> the term "domain controller" in reference to the Kerberos service,
JM> does that fit into what you are describing to me?

In Windows, the domain controllers act as the Kerberos authentication servers (KDCs), among many other services. I believe that in Windows Kerberos, you do indeed need to set a flag for the service principal indicating that it is trusted for delegation (corresponds to a Microsoft-defined "DELEGATE-OK" bit in tickets), and if this is correct you'll need to get an admin to make the change.

>>> We've tried the -A switch when starting up the SSH tunnel with >
>> no avail.
>>
>> This has to do with forwarding the OpenSSH key agent, which has
>> nothing to do with Kerberos tickets.

Re: OpenSSH, Telnet, Windows Authentication and double-hops

JM> Ok. Sounds like there's no need for us to use this switch then.

>> > Is there anything that SSH can offer us so that we can maintain
>> the > authentication ticket over the second hop?

>>

>> You mean, maintain it over the first hop so it's available for the
>> second.

JM> Yes, that's what I'm hoping for.

>> You haven't said what SSH implementation you are using. Some
>> provide

JM> We're focusing on the OpenSSH for Windows distribution. Client

JM> side, we're going to need 2 flavors... one for Windows desktop

JM> machines, and another for Windows CE devices.

>> integration with Windows Kerberos (SSPI) and can do this for you,
>> e.g. Tectia from ssh.com and VShell from VanDyke. You *might* get
>> it to work with OpenSSH/Cygwin. I have compiled OpenSSH on Windows
>> with Kerberos support and had it work; however, the ccache is an
>> MIT one and has nothing to do with the Windows ccache. The MIT KfW
>> system does integrate with SSPI; however, I don't know if there's a
>> way to hook OpenSSH into it.

JM> Did you make code changes to the OpenSSH sources to achieve this?

No. I don't remember, however, if I go the GSSAPI/Kerberos libraries from
Cygwin, or had to compile them myself as well.

JM> Is the MIT ccache (certificate cache?) available to all of us?

Not sure what the question means. "ccache" is credentials cache; it's
where a process' Kerberos tickets are stored as it acquires them for
authentication. In Unix Kerberos, this is typically just a file in /tmp.
Windows employs a process-based ccache, so that without special support,
they have nothing to do with each other, and you'd have to authenticate
(kinit) separately using MIT Kerberos, rather than using the credentials
you already have as a result of logging into Windows (single-signon).

JM> This brings up an interesting segue actually... we'd love to
JM> remove the dependency for the Telnet server/client and are
JM> thinking ahead to integrating an SSH server in our application.
JM> Essentially, we'd like to run our application as a service that
JM> opens a secure port with the SSH authentication and that limits
JM> user interaction to streaming data from our application to the
JM> client. I'm getting the impression that it would be possible to
JM> do this with the OpenSSH source code, am I thinking correctly?
JM> Ultimately, I'd like to have our service use core functionality
JM> from the OpenSSH libraries without changing any of the core
JM> code... maybe just make calls into the cygcrypto-0.9.7.dll (or

Re: OpenSSH, Telnet, Windows Authentication and double-hops

JM> whichever one is truly pertinent)... does this seem possible?

The OpenSSH code isn't really structured as a library for other applications to use. There is libssh:

<http://freshmeat.net/projects/libssh/>

JM> We'd even be thrilled to just get rid of Telnet now without our

JM> own links into the SSH core, but we've run into a couple of

JM> behavioral differences that are blocking us from doing it. I've

JM> thrown one up to the group for help in this posting:

JM> http://groups.google.ca/group/comp.security.ssh/browse_thread/thread/c3fb90fd3b747553/#

Try ssh -t.

JM> and the other has to do with killing orphaned sub-processes and

JM> I've just found this thread to examine:

JM>

http://groups.google.ca/group/comp.security.ssh/browse_thread/thread/51c4ad8e9494a35f/0b8854eb4a1b32a7?lnk=gs

JM> Richard, many, many thanks for taking the time to help me out.

JM> Your feedback is helping me understand what's possible and driving

JM> out the shadows! :~)

JM> Regards, John

>> -- Richard Silverman r...@xxxxxxxx

--

Richard Silverman

res@xxxxxxxx

.