Ruby master - Feature #4805

Add X509::Name#hash_old for 0.9.X compat

05/31/2011 09:42 PM - nahi (Hiroshi Nakamura)

Status:	Closed	
Priority:	Normal	
Assignee:	nahi (Hiroshi Nakamura)	
Target version:	1.9.3	

Description

X509::Name#hash with OpenSSL 1.0.0 returns different value than with OpenSSL 0.9.X.
Attached patch adds X509::Name#hash_old when you need MD5 based same X509_NAME_hash value as OpenSSL 0.9.X.

Martin, how do you think about adding it?

Associated revisions

Revision bf2e60cd - 06/23/2011 01:51 PM - nahi (Hiroshi Nakamura)

- ext/openssl/ossl_x509name.c: Add X509::Name#hash_old as a wrapper for X509_NAME_hash_old in OpenSSL 1.0.0. See #4805
- test/openssl/test_x509name.rb (test_hash): Make test pass with OpenSSL 1.0.0.
- NEWS: Add it.

git-svn-id: svn+ssh://ci.ruby-lang.org/ruby/trunk@32213 b2dd03c8-39d4-4d8f-98ff-823fe69b080e

Revision 32213 - 06/23/2011 01:51 PM - nahi (Hiroshi Nakamura)

- ext/openssl/ossl_x509name.c: Add X509::Name#hash_old as a wrapper for X509_NAME_hash_old in OpenSSL 1.0.0. See #4805
- test/openssl/test_x509name.rb (test_hash): Make test pass with OpenSSL 1.0.0.
- NEWS: Add it.

Revision 32213 - 06/23/2011 01:51 PM - nahi (Hiroshi Nakamura)

- ext/openssl/ossl_x509name.c: Add X509::Name#hash_old as a wrapper for X509_NAME_hash_old in OpenSSL 1.0.0. See #4805
- test/openssl/test_x509name.rb (test_hash): Make test pass with OpenSSL 1.0.0.
- NEWS: Add it.

Revision 32213 - 06/23/2011 01:51 PM - nahi (Hiroshi Nakamura)

- ext/openssl/ossl_x509name.c: Add X509::Name#hash_old as a wrapper for X509_NAME_hash_old in OpenSSL 1.0.0. See #4805

- test/openssl/test_x509name.rb (test_hash): Make test pass with OpenSSL 1.0.0.
- NEWS: Add it.

Revision 32213 - 06/23/2011 01:51 PM - nahi (Hiroshi Nakamura)

- ext/openssl/ossl_x509name.c: Add X509::Name#hash_old as a wrapper for X509_NAME_hash_old in OpenSSL 1.0.0. See #4805
- test/openssl/test_x509name.rb (test_hash): Make test pass with OpenSSL 1.0.0.
- NEWS: Add it.

Revision 32213 - 06/23/2011 01:51 PM - nahi (Hiroshi Nakamura)

- ext/openssl/ossl_x509name.c: Add X509::Name#hash_old as a wrapper for X509_NAME_hash_old in OpenSSL 1.0.0. See #4805
- test/openssl/test_x509name.rb (test_hash): Make test pass with OpenSSL 1.0.0.
- NEWS: Add it.

Revision 32213 - 06/23/2011 01:51 PM - nahi (Hiroshi Nakamura)

- ext/openssl/ossl_x509name.c: Add X509::Name#hash_old as a wrapper for X509_NAME_hash_old in OpenSSL 1.0.0. See #4805
- test/openssl/test_x509name.rb (test_hash): Make test pass with OpenSSL 1.0.0.
- NEWS: Add it.

History

#1 - 06/09/2011 06:14 AM - MartinBosslet (Martin Bosslet)

- Assignee changed from MartinBosslet (Martin Bosslet) to nahi (Hiroshi Nakamura)

Hi Hiroshi,

sorry for taking some time to answer, I was on vacation last week...
I think your patch is good! But there is one thing I don't like
about OpenSSL itself here - why do they hardcode the digest algorithm
in the first place?

There are situations where neither MD5 nor SHA-1 fits, OCSP requests
are a good example: The requested CertID is defined as

CertID ::= SEQUENCE {
hashAlgorithm AlgorithmIdentifier,
issuerNameHash OCTET STRING, -- Hash of Issuer's DN
issuerKeyHash OCTET STRING, -- Hash of Issuers public key
serialNumber CertificateSerialNumber }

This implies trouble for any SHA-2 family "hashAlgorithm".

In addition to applying your patch I'd favor a Name#hash implementation
that takes an optional OpenSSL::Digest that specifies the hash algorithm
to be used.
This would of course mean that we would have to implement the functionality
of X509_name_hash on our own. What do you think - would the benefit of a
cleaner solution outweigh the (partial) code duplication?

Regards,
Martin

#2 - 06/20/2011 12:18 PM - nahi (Hiroshi Nakamura)

- Due date deleted (05/31/2011)

#3 - 06/20/2011 12:23 PM - nahi (Hiroshi Nakamura)

Hi,

On Thu, Jun 9, 2011 at 06:14, Martin Bosslet
Martin.Bosslet@googlemail.com wrote:

> I think your patch is good! But there is one thing I don't like
> about OpenSSL itself here - why do they hardcode the digest algorithm
> in the first place?

They're using the hash of name for c_rehash. You see files something
like hex encoded in certs dir of OpenSSL;

d2adc77d.0@
d537fba6.0@
d78a75c7.0@
d8274e24.0@
ddc328ff.0@

(e.g. /etc/ssl/certs/ in Ubuntu)

For that purpose, algorithm should be fixed so they don't get Digester
as a parameter for X509_NAME_hash and X509_NAME_hash_old I guess.

Besides this, I don't know the reason why they change base digester
from MD5 to SHA1 at the version bump from 0.9.8 to 1.0.0.

> There are situations where neither MD5 nor SHA-1 fits, OCSP requests
> are a good example: The requested CertID is defined as
>
> CertID ::= SEQUENCE {
> hashAlgorithm AlgorithmIdentifier,
> issuerNameHash OCTET STRING, -- Hash of Issuer's DN
> issuerKeyHash OCTET STRING, -- Hash of Issuers public key
> serialNumber CertificateSerialNumber }
>
> This implies trouble for any SHA-2 family "hashAlgorithm".

I can understand it but it should be different problem for them I
think. issuerNameHash has variable length, not fixed to 32bit integer.

> In addition to applying your patch I'd favor a Name#hash implementation
> that takes an optional OpenSSL::Digest that specifies the hash algorithm
> to be used.
> This would of course mean that we would have to implement the functionality
> of X509_name_hash on our own. What do you think - would the benefit of a
> cleaner solution outweigh the (partial) code duplication?

- Hide quoted text - I like 'X509::Name#hash' to be a wrapper of 'X509_NAME_hash' and 'X509::Name#hash_old' is for 'X509::Name::hash_old'. I prefer to have another name for hashing X509::Name if it's needed.

Regards,
// NaHi

#4 - 06/20/2011 06:43 PM - nahi (Hiroshi Nakamura)

- Assignee changed from nahi (Hiroshi Nakamura) to MartinBosslet (Martin Bosslet)

#5 - 06/23/2011 08:53 PM - MartinBosslet (Martin Bosslet)

- Assignee changed from MartinBosslet (Martin Bosslet) to nahi (Hiroshi Nakamura)

Hiroshi NAKAMURA wrote:

> They're using the hash of name for c_rehash. You see files something
> like hex encoded in certs dir of OpenSSL;
>
> d2adc77d.0@
> d537fba6.0@
> d78a75c7.0@
> d8274e24.0@
> ddc328ff.0@
>
> (e.g. /etc/ssl/certs/ in Ubuntu)
>
> For that purpose, algorithm should be fixed so they don't get Digester
> as a parameter for X509_NAME_hash and X509_NAME_hash_old I guess.

OK, I see, thanks for the info! X509_NAME_hash(_old) has a special meaning
internally. I checked, there also exists X509_NAME_digest for the general
purpose usage I was thinking of. It's also what they use internally to
create the CertIDs in their OCSP implementation.

> Besides this, I don't know the reason why they change base digester
> from MD5 to SHA1 at the version bump from 0.9.8 to 1.0.0.

I could imagine that they changed for some sort of security reasons - a lot
of official recommendations/guidelines advise to refrain from using MD5, so
it could be a political reason...

> I like 'X509::Name#hash' to be a wrapper of 'X509_NAME_hash' and
> 'X509::Name#hash_old' is for 'X509::Name::hash_old'. I prefer to have
> another name for hashing X509::Name if it's needed.

Now that I have a better understanding of the context I completely agree.
If we feel the need for a general-purpose method, we could probably also use
X509::Name#digest in analogy to OpenSSL.

If you'd like me to apply the patch, please feel free to reassign to me!

Regards,
Martin

#6 - 06/23/2011 11:04 PM - nahi (Hiroshi Nakamura)

- Status changed from Open to Closed

On Thu, Jun 23, 2011 at 20:54, Martin Bosslet Martin.Bosslet@googlemail.com wrote:

> > For that purpose, algorithm should be fixed so they don't get Digester
> > as a parameter for X509_NAME_hash and X509_NAME_hash_old I guess.
>
> OK, I see, thanks for the info! X509_NAME_hash(_old) has a special meaning
> internally. I checked, there also exists X509_NAME_digest for the general
> purpose usage I was thinking of. It's also what they use internally to
> create the CertIDs in their OCSP implementation.

Ah, X509::Name#digest sounds good. I didn't know OCSP impl already used. Since ext/openssl is a wrapper, I don't like to break class hierarchy, name, etc. of OpenSSL. Original API should be defined under Security::Crypto or something and it should be OpenSSL free... someday.

> > I like 'X509::Name#hash' to be a wrapper of 'X509_NAME_hash' and
> > 'X509::Name#hash_old' is for 'X509::Name::hash_old'. I prefer to have
> > another name for hashing X509::Name if it's needed.
>
> Now that I have a better understanding of the context I completely agree.
> If we feel the need for a general-purpose method, we could probably also use

X509::Name#digest in analogy to OpenSSL.

Thanks. Applied it at r32213.

Files

X509-Name-hash_old.diff	3.16 KB	05/31/2011	nahi (Hiroshi Nakamura)