From:	Bassham, Lawrence E (Fed)
To:	Perlner, Ray A. (Fed); Moody, Dustin (Fed)
Subject:	Re: Someone Is Testing Our DRBG requirements
Date:	Monday, April 10, 2017 2:19:01 PM

We are allowing a non-NIST-approved DRBG if they give a rationale for it. In that case, we need the max value we specify to get some sort of timing data for whatever they use. Let's talk briefly tomorrow.

Larry

From: "Perlner, Ray (Fed)" <ray.perlner@nist.gov>
Date: Monday, April 10, 2017 at 2:14 PM
To: "Bassham, Lawrence E (Fed)" <lawrence.bassham@nist.gov>, "Moody, Dustin (Fed)" <dustin.moody@nist.gov>
Subject: RE: Someone Is Testing Our DRBG requirements

Those numbers would cover the DRBGs in 90A, I think. But what about the idea of just saying that randombytes may only be used as a seed for a NIST approved DRBG? Then we don't need to specify a maximum size for the seed.

I think if we're going to update the requirements on the use of randombytes, this would be the most sensible way to do it.

From: Bassham, Lawrence E (Fed)
Sent: Monday, April 10, 2017 2:06 PM
To: Moody, Dustin (Fed) <dustin.moody@nist.gov>; Perlner, Ray (Fed) <ray.perlner@nist.gov>
Subject: Re: Someone Is Testing Our DRBG requirements

I did. I'm fine with the entropy/randomness numbers that Ray suggests. The numbers as set now reflect just the security with respect to the PQ side of things. If you want to bump them up a bit to cover the CTR-DRBG requirements that's fine. Ray, would {256, 320, and 384} cover all DRBG's in 90A? If so, let's just do that.

I don't think you're going to get Dan to change anything. I think he will complain because he'll say that the testing software would have to try all the DRBG's on each platform to see what gives the best performance.

Larry

From: "Moody, Dustin (Fed)" <dustin.moody@nist.gov>
Date: Monday, April 10, 2017 at 1:55 PM
To: "Perlner, Ray (Fed)" <ray.perlner@nist.gov>, "Bassham, Lawrence E (Fed)" <lawrence.bassham@nist.gov>

Subject: RE: Someone Is Testing Our DRBG requirements

Larry,
Did you get this?

Dustin

From: Perlner, Ray (Fed)
Sent: Monday, April 10, 2017 9:08 AM
To: Bassham, Lawrence E (Fed) <lawrence.bassham@nist.gov>; Moody, Dustin (Fed) <dustin.moody@nist.gov>
Subject: RE: Someone Is Testing Our DRBG requirements

Several of the functions indisputably need secure randomness, in particular, key pair generation for all algorithms, and the encapsulate and encrypt functions. (Secure randomness is really optional for signature generation.) There are two ways to get secure random bits, a hardware source, and a DRBG that was previously seeded with a hardware source. The latter is cheaper, so all algorithms will need to get at least a few random bits from a DRBG. Now, it is entirely plausible that a small number of bits from a NIST approved DRBG might be expanded into a large number of bits, using a non-NIST-approved DRBG, and we should permit that. Banning the use of random-bytes directly does not prevent that however.

Another possibility I discussed with Dustin would be to simply ask Dan to use a real DRBG (probably AES counter DRBG) in place of the LFSR in random-bytes, and then we could just let random-bytes be called directly.

From: Bassham, Lawrence E (Fed)
Sent: Friday, April 07, 2017 7:53 PM
To: Perlner, Ray (Fed) <ray.perlner@nist.gov>; Moody, Dustin (Fed) <dustin.moody@nist.gov>
Subject: Re: Someone Is Testing Our DRBG requirements

But what about the people that feel they don't need a NIST-approved DRBG? I do like the idea though.

On: 07 April 2017 16:36, "Perlner, Ray (Fed)" <ray.perlner@nist.gov> wrote:

Might the fairest thing be to simply say that randombytes may only be used to seed a NIST approved DRBG? Then we don't have to specify a maximum length for the seed. Otherwise we're left in a situation where an algorithm that needs 384 bits, say, of secure random data gets them basically for free, while an algorithm that needs 385 bits of secure random data has to use a DRBG for all of them.

From: Bassham, Lawrence E (Fed)

Sent: Friday, April 07, 2017 4:15 PM
To: Moody, Dustin (Fed) <dustin.moody@nist.gov>
Cc: Perlner, Ray (Fed) <ray.perlner@nist.gov>
Subject: Re: Someone Is Testing Our DRBG requirements

Might be interesting to know in practice if a call for extra entropy bits is more expensive than a few AES calls on a system with an AES instruction. If we bump them up we can say it is a maximum entropy size so that the choice is theirs. We shouldn't require them to use the larger entropy if they aren't using a NIST-approved DRBG. Since there's a bit of complaint about the whole RNG aspect of things it might be best to bump the numbers up and say that now that encompass the entropy for the security category and the (maximum) requirements of the NIST DRBG. (Is the word "maximum" necessary?)

Larry

From: "Moody, Dustin (Fed)" <dustin.moody@nist.gov>
Date: Friday, April 7, 2017 at 2:17 PM
To: "Bassham, Lawrence E (Fed)" <lawrence.bassham@nist.gov>
Cc: "Perlner, Ray (Fed)" <ray.perlner@nist.gov>
Subject: Fw: Someone Is Testing Our DRBG requirements

Larry,
What do you think?

From: Perlner, Ray (Fed)
Sent: Friday, April 7, 2017 2:02 PM
To: Moody, Dustin (Fed)
Subject: RE: Someone Is Testing Our DRBG requirements

After reading 90A slightly more the situation appears to be the following:

You can use less than seedlength bits for CTR DRBG, but then you have to use a KDF to generate the seed when you instantiate the DRBG, which induces some additional cost (2 or 3 AES operations). Some of the other DRBGs have a larger seed length, but it looks like they always require a KDF when instantiated (in which case the entropy input length can be anything you want as long as it exceeds the security level.)

As for specifying a range of values, I think we intend the random input lengths as maximum values. Some of the functions which can take random inputs (in particular the sign operation) don't necessarily need any random input. So we're already allowing a range of values. In summary, I see the situation as follows: Either we can leave the maximum entropy input lengths as is, or we can bump them up by another 64 bits to accommodate implementations that want to instantiate CTR-DRBG without using a derivation function. It probably doesn't matter all that much either way.

From: Moody, Dustin (Fed)
Sent: Friday, April 07, 2017 1:14 PM
To: Perlner, Ray (Fed) <ray.perlner@nist.gov>
Subject: Re: Someone Is Testing Our DRBG requirements

What would be the pros and cons of this? If we set it to 256, 320, and 384 are we basically telling people to use AES CTR DRBG? Can they still use AES CTR DRBG if we don't change the values? Should we have a small range of values for each security target?

From: Perlner, Ray (Fed)
Sent: Friday, April 7, 2017 11:16:13 AM
To: Moody, Dustin (Fed); Alperin-Sheriff, Jacob (Fed); Bassham, Lawrence E (Fed)
Cc: Chen, Lily (Fed)
Subject: RE: Someone Is Testing Our DRBG requirements

Huh, AES 128 AES 192 and AES 256 CTR DRBG have seed lengths of 256, 320, and 384 bits respectively according to SP 800-90a, maybe we should use those values for the lengths of randombytes in our API.

From: Moody, Dustin (Fed)
Sent: Friday, April 07, 2017 10:36 AM
To: Alperin-Sheriff, Jacob (Fed) <jacob.alperin-sheriff@nist.gov>; Perlner, Ray (Fed) <ray.perlner@nist.gov>; Bassham, Lawrence E (Fed) <lawrence.bassham@nist.gov>
Cc: Chen, Lily (Fed) <lily.chen@nist.gov>
Subject: Re: Someone Is Testing Our DRBG requirements

From their conclusion:

For generic implementations of discrete Gaussian sampling CTR-DRBG with AES is a safe option, particularly if AES hardware acceleration is available or you want to enter a NIST competition. It possesses a balance of performance and security, is well understood and is accepted by the security community.

I enjoyed the part "or if you want to enter a NIST competition".

From: Moody, Dustin (Fed)
Sent: Friday, April 7, 2017 10:28:35 AM
To: Alperin-Sheriff, Jacob (Fed); Perlner, Ray (Fed); Bassham, Lawrence E (Fed)
Cc: Chen, Lily (Fed)
Subject: Re: Someone Is Testing Our DRBG requirements

Sure. Just make it more of an email from you, and not an official NIST message from our PQC team.

They must have been already working on this.

From: Alperin-Sheriff, Jacob (Fed)

Sent: Friday, April 7, 2017 9:14:46 AM
To: Moody, Dustin (Fed); Perlner, Ray (Fed); Bassham, Lawrence E (Fed)
Cc: Chen, Lily (Fed)
Subject: Someone Is Testing Our DRBG requirements

That was relatively fast given that we only just finished clarifying it.

https://eprint.iacr.org/2017/298

Can I reach out to them and let them know we're very happy to see this kind of work being done?

—Jacob Alperin-Sheriff