

ENTROPY, RANDOM NUMBERS AND KEYS: WHAT'S GOOD ENOUGH?

John Leiseboer
QuintessenceLabs

Security in
knowledge



In the news

Security Analysis of Pseudo-Random Number Generators with Input:
/dev/random is not Robust
Yevgeniy Dodis¹, David Pointcheval², Sylvain Ruhault³, Damien Vergnaud

Stealthy Dopant-Level Hardware Trojans *
T. Becker¹, Francesco Regazzoni
and Wayne P. Burleson

On the Possibility of a Back Door
in the NIST SP800-90 Dual Ec
Prng
Dan Shumow
Niels Ferguson
Microsoft

MIT Research: Encryption Less Secure Than We Thought
Posted by **Soulskill** on Wednesday August 14, 2013 @02:50PM
from the but-still-pretty-darn-secure dept.
A group of researchers from MIT and the University of Ireland has presented [a paper](#) (PDF) showing that cryptographic security is wrong. As a result, certain encryption-breaking methods explains, is that information-theoretic analyses of secure systems have gotten on so-called [Shannon entropy](#), named after the founder of information theory.

Fatal crypto flaw in some government-certified smartcards makes forgery a snap
With government certifications this broken, the NSA may not need backdoors.
by Dan Goodin - Sept 17 2013, 1:25am +1000



Deliberately flawed? RSA Security tells customers to drop NSA-related encryption algorithm
Published time: September 20, 2013 10:31
Edited time: September 20, 2013 13:15
[Get short URL](#)

— Introduction

- ▶ Relationships
 - ▶ Entropy
 - ▶ Random numbers
 - ▶ Cryptographic keys
- ▶ Random bit generator construction
- ▶ Entropy sources
- ▶ Measuring and testing entropy and randomness
- ▶ Putting it all together

— Cryptographic strength

- ▶ Number of bits of security
 - ▶ Algorithm
 - ▶ Key size
- ▶ Security strength relies on
 - ▶ Robust algorithms
 - ▶ Correct implementation
 - ▶ No feasible mathematical attack or advances
 - ▶ Brute force attack infeasible


— Random use cases

- ▶ RNG seed
- ▶ Key generator seed
- ▶ IV
- ▶ Nonce
- ▶ Random challenge
- ▶ Authentication
- ▶ DSA signing
- ▶ One-time pad cipher
- ▶ Zero knowledge proof
- ▶ E-voting
- ▶ Random beacon
 - ▶ Transaction protection
 - ▶ Contract signing
 - ▶ PII protection
 - ▶ Cloud entropy

NIST's RBG standards

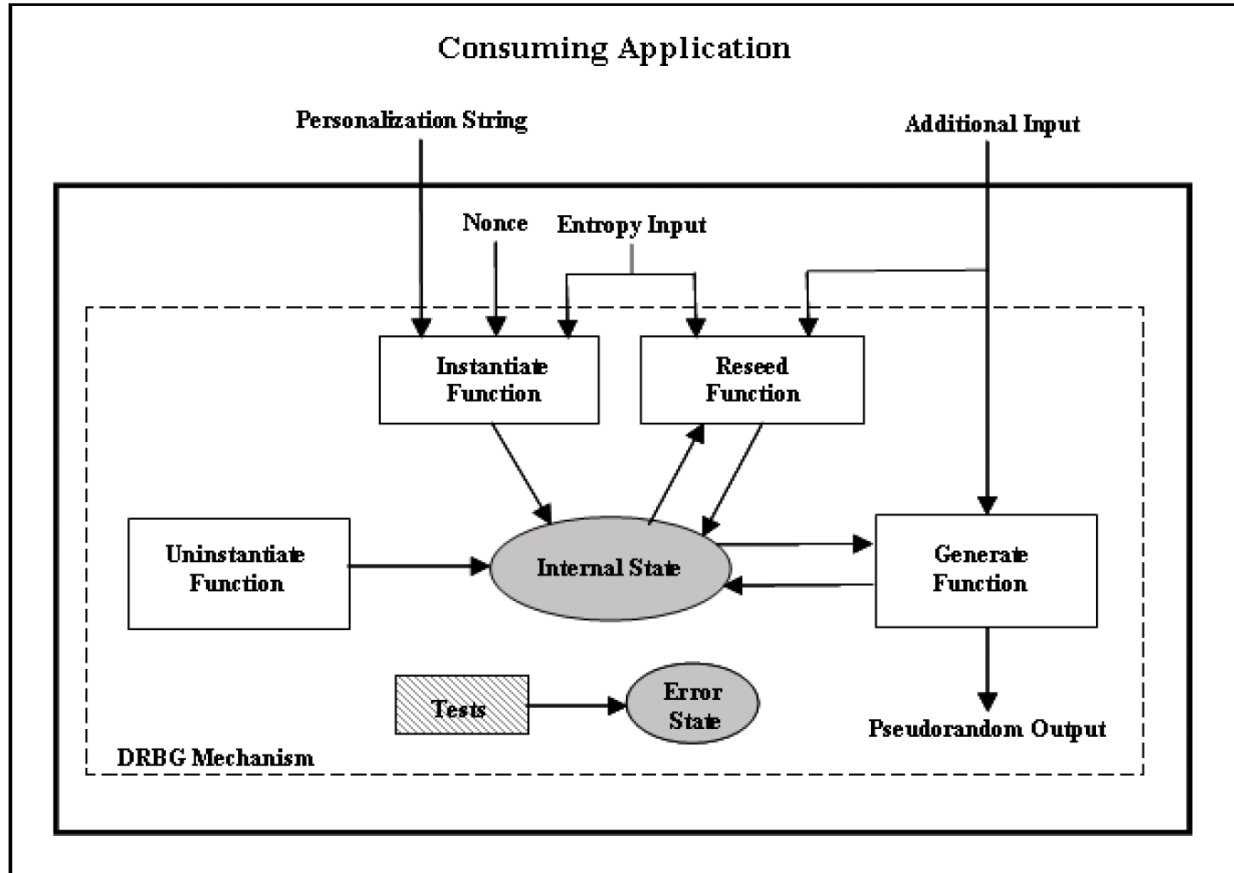
- ▶ NIST SP 800-90A: DRBG
- ▶ NIST SP 800-90B: Entropy source
- ▶ NIST SP 800-90C: RBG = Entropy source + DRBG

- ▶ Example:



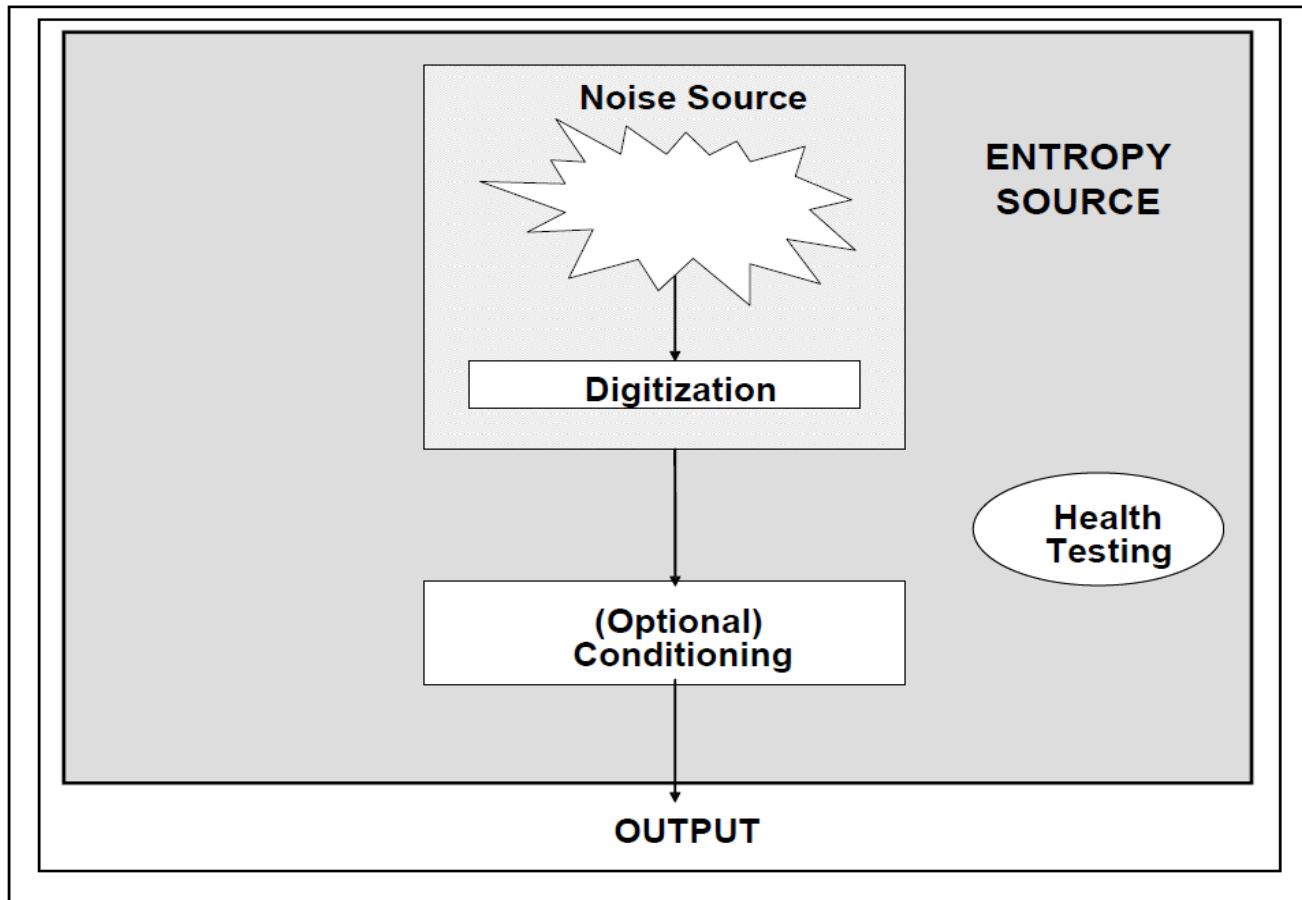
NIST SP 800-90A

DRBG Functional Model



NIST SP 800-90B


Entropy Source Functional Model




NIST SP 800-90C

RBG Construction



Enhanced NRBG – XOR Construction



Enhanced NRBG – Oversampling Construction



Entropy sources



Estimating entropy

- ▶ 9, 9, 9, 9, 9, 9, 9, 9, 9, 9, 9, 9, ...
- ▶ 3, 1, 4, 1, 5, 9, 2, 6, 5, 3, 5, 9,
- ▶ Independent and identical distribution (IID)

Non-i.i.d. Min-Entropy Estimation Test Results:

Collision test	4.310757 bits per 8-bit symbol
Partial collection test	2.467009 bits per 8-bit symbol
Markov test	5.692803 bits per 6-bit symbol (remapped due to test limit)
Compression test	3.401208 bits per 8-bit symbol
Frequency test	6.699898 bits per 8-bit symbol

Sanity Check Results:

Compression test	passed - (840168 bits, 840800 bits, 840336 bits, 839440 bits, 840688 bits, 839776 bits, 840312 bits, 839688 bits, 839800 bits, 840016 bits)
Collision test	passed - (0 13-symbol values with a count of 3 or more, 0 colliding 13-symbol values in total)

Min-entropy estimate is 2.467009 bits per 8-bit symbol, based on a 95% confidence interval.

— Random number tests

▶ NIST STS

- ▶ <http://csrc.nist.gov/groups/ST/toolkit/rng/index.html>
- ▶ “...no set of statistical tests can absolutely certify a generator as appropriate...”

▶ Dieharder

- ▶ <http://www.phy.duke.edu/~rgb/General/dieharder.php>
- ▶ “dieharder is a tool designed to permit one to push a weak generator to unambiguous failure ...”

Dieharder RNGs and tests

```
#####  
#   Id Test Name      | Id Test Name      | Id Test Name      #  
#####  
| 000 borosh13        |001 cmrg           |002 coveyou        |  
| 003 fishman18       |004 fishman20      |005 fishman2x      |  
| 006 gfsr4           |007 knuthran       |008 knuthran2      |  
| 009 knuthran2002    |010 lecuyer21      |011 minstd         |  
| 012 mrg             |013 mt19937        |014 mt19937_1999   |  
| 015 mt19937_1998    |016 r250           |017 ran0           |  
| 018 ran1            |019 ran2           |020 ran3           |  
| 021 rand            |022 rand48         |023 random128-bsd  |  
| 024 random128-glibc2|025 random128-libc5|026 random256-bsd  |  
| 027 random256-glibc2|028 random256-libc5|029 random32-bsd   |  
| 030 random32-glibc2 |031 random32-libc5 |032 random64-bsd   |  
| 033 random64-glibc2 |034 random64-libc5 |035 random8-bsd    |  
| 036 random8-glibc2  |037 random8-libc5  |038 random-bsd     |  
| 039 random-glibc2   |040 random-libc5   |041 randu          |  
| 042 ranf            |043 ranlux         |044 ranlux389      |  
| 045 ranlxd1         |046 ranlxd2        |047 ranlxs0        |  
| 048 ranlxs1         |049 ranlxs2        |050 ranmar         |  
| 051 slatec          |052 taus           |053 taus2          |  
| 054 taus113         |055 transputer     |056 tt800          |  
| 057 uni             |058 uni32          |059 vax            |  
| 060 waterman14      |061 zuf            |                   |  
#####  
| 200 stdin_input_raw |201 file_input_raw |202 file_input     |  
| 203 ca              |204 uvag           |205 AES_OFB        |  
| 206 Threefish_OFB   |207 XOR (supergenerator)|208 kiss          |  
| 209 superkiss       |                   |                   |  
#####  
| 400 R_wichmann_hill |401 R_marsaglia_multic. |402 R_super_duper  |  
| 403 R_mersenne_twister |404 R_knuth_taocp    |405 R_knuth_taocp2 |  
#####  
| 500 /dev/random     |501 /dev/urandom     |                   |  
#####
```

Dieharder RNGs and tests

```
#####  
#   Id Test Name           | Id Test Name           | Id Test Name           |  
#####  
| 000 boro  
| 003 fish  
| 006 gfsr  
| 009 knut  
| 012 mrg  
| 015 mt19  
| 018 ran1  
| 021 rand  
| 024 rand  
| 027 rand  
| 030 rand  
| 033 rand  
| 036 rand  
| 039 rand  
| 042 ranf  
| 045 ranl  
| 048 ranl  
| 051 slat  
| 054 taus  
| 057 uni  
| 060 wate  
#####  
| 200 stdi  
| 203 ca  
| 206 Thre  
| 209 supe  
#####  
| 400 R_wi  
| 403 R_me  
#####  
| 500 /dev/random      | 1501 /dev/urandom      |  
#####
```

Diehard "Birthdays" test (modified).
Each test determines the number of matching intervals from 512
"birthdays" (by default) drawn on a 24-bit "year" (by
default). This is repeated 100 times (by default) and the
results cumulated in a histogram. Repeated intervals should be
distributed in a Poisson distribution if the underlying generator
is random enough, and a a chisq and p-value for the test are
evaluated relative to this null hypothesis.

It is recommended that you run this at or near the original
100 test samples per p-value with -t 100.

Two additional parameters have been added. In diehard, nms=512
but this CAN be varied and all Marsaglia's formulae still work. It
can be reset to different values with -x nmsvalue.
Similarly, nbits "should" 24, but we can really make it anything
we want that's less than or equal to rmax_bits = 32. It can be
reset to a new value with -y nbits. Both default to diehard's
values if no -x or -y options are used.

Dieharder RNGs and tests

```
#####  
#   Id Test Name           | Id Test Name           | Id Test Name           |  
#####  
| 000 boro  
| 003 fish  
| 006 gfsr  
| 009 knut  
| 012 mrg #=====  
| 015 mt19 #           Diehard Overlapping 5-Permutations Test.  
| 018 ran1 # This is the OPERM5 test. It looks at a sequence of one mill-  
| 021 rand # ion 32-bit random integers. Each set of five consecutive  
| 024 rand # integers can be in one of 120 states, for the 5! possible or-  
| 027 rand # derings of five numbers. Thus the 5th, 6th, 7th,...numbers  
| 030 rand # each provide a state. As many thousands of state transitions  
| 033 rand # are observed, cumulative counts are made of the number of  
| 036 rand # occurrences of each state. Then the quadratic form in the  
| 039 rand # weak inverse of the 120x120 covariance matrix yields a test  
| 042 ranf # equivalent to the likelihood ratio test that the 120 cell  
| 045 ranl # counts came from the specified (asymptotically) normal dis-  
| 048 ranl # tribution with the specified 120x120 covariance matrix (with  
| 051 slat # rank 99). This version uses 1,000,000 integers, twice.  
| 054 taus #  
| 057 uni #  
| 060 wate #  
#####  
| 200 stdi #  
| 203 ca # Note that Dieharder runs the test 100 times, not twice, by  
| 206 Thre # default.  
| 209 supe #=====  
#####  
| 400 R_wi  
| 403 R_me  
#####  
| 500 /dev/random |501 /dev/urandom |  
#####
```

Dieharder RNGs and tests

```
#####  
#   Id Test Name           | Id Test Name           | Id Test Name           |  
#####  
| 000 boro  
| 003 fish  
| 006 gfsr  
| 009 knut  
| 012 mrg  
| 015 mt19  
| 018 ranl  
| 021 rand  
| 024 rand  
| 027 rand  
| 030 rand  
| 033 rand  
| 036 rand  
| 039 rand  
| 042 ranf  
| 045 ranl  
| 048 ranl  
| 051 slat  
| 054 taus  
| 057 uni  
| 060 wate  
#####  
| 200 stdi  
| 203 ca  
| 206 Thre  
| 209 supe  
#####  
| 400 R_wi  
| 403 R_me  
#####  
| 500 /dev/random      | 1501 /dev/urandom      |  
#####
```

Diehard 32x32 Binary Rank Test

This is the BINARY RANK TEST for 32x32 matrices. A random 32x32 binary matrix is formed, each row a 32-bit random integer. The rank is determined. That rank can be from 0 to 32, ranks less than 29 are rare, and their counts are pooled with those for rank 29. Ranks are found for 40,000 such random matrices and a chisquare test is performed on counts for ranks 32,31,30 and <=29.

As always, the test is repeated and a KS test applied to the resulting p-values to verify that they are approximately uniform.

Dieharder RNGs and tests

```
#####  
#   Id Test Name           | Id Test Name           | Id Test Name           #  
#####  
| 000 boro #  
| 003 fish # Diehard Bitstream Test.  
| 006 gfsr # The file under test is viewed as a stream of bits. Call them  
| 009 knut # b1,b2,... . Consider an alphabet with two "letters", 0 and 1  
| 012 mrg # and think of the stream of bits as a succession of 20-letter  
| 015 mt19 # "words", overlapping. Thus the first word is b1b2...b20, the  
| 018 ran1 # second is b2b3...b21, and so on. The bitstream test counts  
| 021 rand # the number of missing 20-letter (20-bit) words in a string of  
| 024 rand # 2^21 overlapping 20-letter words. There are 2^20 possible 20  
| 027 rand # letter words. For a truly random string of 2^21+19 bits, the  
| 030 rand # number of missing words j should be (very close to) normally  
| 033 rand # distributed with mean 141,909 and sigma 428. Thus  
| 036 rand # (j-141909)/428 should be a standard normal variate (z score)  
| 039 rand # that leads to a uniform [0,1) p value. The test is repeated  
| 042 ranf # twenty times.  
| 045 ranl #  
| 048 ranl #  
| 051 slat #  
| 054 taus #  
| 057 uni # NOTE WELL!  
| 060 wate #  
#####  
| 200 stdi # The test is repeated 100 times by default in dieharder, but the  
| 203 ca # size of the sample is fixed (tsamples cannot/should not be  
| 206 Thre # varied from the default). The sigma of this test REQUIRES the  
| 209 supe # use of overlapping samples, and overlapping samples are not  
#####  
| 400 R_wi # independent. If one uses the non-overlapping version of this  
| 403 R_me # test, sigma = 290 is used instead, smaller because now there  
#####  
| 500 /dev # are 2^21 INDEPENDENT samples.  
#####
```

Dieharder test results

```
#=====#
#           dieharder version 3.31.1 Copyright 2003 Robert G. Brown           #
#=====#
  rng_name      |          filename          |rands/second|
file_input_raw|          /dev/urandom| 3.14e+06  |
#=====#
  test_name     |ntup| tsamples |psamples|  p-value |Assessment
#=====#
  diehard_birthdays| 0|    100|    100|0.74056677| PASSED
    diehard_operm5| 0| 1000000|    100|0.06667553| PASSED
  diehard_rank_32x32| 0|   40000|    100|0.02200848| PASSED
    diehard_rank_6x8| 0|   10000|    100|0.09107953| PASSED
  diehard_bitstream| 0| 2097152|    100|0.31603737| PASSED
    diehard_opso| 0| 2097152|    100|0.34660453| PASSED
    diehard_oqso| 0| 2097152|    100|0.37564518| PASSED
    diehard_dna| 0| 2097152|    100|0.36650750| PASSED
  diehard_count_1s_str| 0|  256000|    100|0.48806562| PASSED
  diehard_count_1s_byt| 0|  256000|    100|0.55880694| PASSED
  diehard_parking_lot| 0|   12000|    100|0.83939978| PASSED
    diehard_2dsphere| 2|    8000|    100|0.04126659| PASSED
    diehard_3dsphere| 3|    4000|    100|0.06311990| PASSED
    diehard_squeeze| 0|   100000|    100|0.50805779| PASSED
    diehard_sums| 0|    100|    100|0.08429191| PASSED
```

— Practical issues

- ▶ Entropy sources
- ▶ Hardware and software implementations
- ▶ Standards compliance
- ▶ Performance requirements
- ▶ Security requirements
- ▶ User experience – application developer and end-user
- ▶ Trust

/dev/random and /dev/urandom

▶ Try this:

▶ \$ cat /dev/random

```
jleiseboer@f16:~$ cat /dev/random | hexdump -C
00000000  81 5b 83 77 e6 61 de 27 e2 bc 1f 14 13 b1 5b 01 |. [.w.a.'.....[. |
00000010  05 c7 0a 1e 4c 70 a0 a2 bd d2 03 81 ec 03 2b c9 |.....lp.....+. |
```

\$ cat /dev/urandom

```
jleiseboer@f16:~$ cat /dev/urandom
00000000  6f 9d dd 83 0e d1 24 cb 2e 90 d2 83 bc 17 a3 ef |o.....$...... |
00000010  b4 76 9d fc b8 71 fc 0f 35 74 6e 9e bb 4b 2b d3 |.v...q..5tn..K+. |
00000020  65 f7 a8 60 d8 82 46 f0 47 b1 e2 a6 c4 c3 fa a7 |e..`...F.G..... |
00000030  1c 89 ab d2 e0 84 92 2e 64 86 c2 04 74 7a e9 bc |.....d...tz.. |
00000040  82 d9 96 1d ea 94 90 e8 90 43 3e 7f b4 3d ad 63 |.....C>..=..c |
00000050  12 2e 0b 17 5f f9 01 ca 51 88 fa 3d 2b 9e 60 f8 |....._Q..=+.`.. |
00000060  a6 8a 72 ee 66 31 8f f6 ef 41 b8 18 79 0b a0 e2 |...r.fi...A..y... |
00000070  a4 78 8f 03 15 d9 0e 6f aa bb 02 d1 41 48 96 a6 |.x.....o...AH... |
00000080  87 ac 8e 3c e8 c0 0d 76 df 21 02 7d 52 9b ea 28 |...<...v.!}R..(|
00000090  f8 0d 80 2a f9 69 0f b7 50 38 eb 7c 92 e9 a9 48 |...*.i...P8.|...H |
000000a0  f5 65 43 70 ad 0d d9 13 19 ea 2a 7f 30 ff 77 ec |.eCp.....*.0.w. |
000000b0  d7 b1 61 ec 4e 76 57 3b ee 3f 3d 6b 98 3b 65 b1 |...a.NvW;.?=k;.e. |
000000c0  99 ff 3c ff 3c 04 b3 67 7a 5d cf c9 c2 80 0a 3a |...<.<..gz]..... |
000000d0  50 6b 5e 77 2c 36 8f e6 5d c3 ad 2a ed e1 8d 81 |Pk^w,6...]*..... |
000000e0  4f 26 4a d5 6a e4 35 87 05 35 42 69 5f 51 c2 5f |o&J.j..5..5Bi_Q_|
000000f0  cd 47 03 9a 1d fd 4a b2 fe 30 a2 ac 2a 95 f9 68 |.G.....J..0.*..h |
00000100  d8 16 cf 77 06 3e 34 7a 6c 0e 19 80 ea 67 95 d4 |...w.>4zl....g.. |
00000110  14 82 bd 8f e5 b6 72 d6 bb 79 3b 01 1d 86 c0 81 |.....r..y;..... |
00000120  6b 84 38 2a 5c d3 8c e4 33 6d a6 c3 42 b0 5a 66 |k.8*\...3m..B.zf |
00000130  e5 80 5b 2f fc 31 b6 3f 49 00 34 2a 44 af 74 fb |...[/..1.?I.4*D.t. |
00000140  9a 81 c9 88 a0 a6 17 65 28 8f fa 06 6e 45 26 87 |.....e(...nE&. |
00000150  af ea 39 d0 ff 7e 78 fa 48 68 41 fb 89 c1 f8 75 |...9...-x.HhA.... |
00000160  2d 8d be a3 bc 4e d8 13 65 61 6f bd 53 f0 f6 07 |-....N...eao.S... |
```

▶ Implementations differ

▶ Blocking vs. non-blocking

▶ Entropy quality

Hardware RNG's

- ▶ Microcontrollers
- ▶ Intel RDRAND and RDSEED
- ▶ HSM's
- ▶ RNG appliances, cards and components
- ▶ DIY
- ▶ Entropy vs. random numbers
 - ▶ Deterministic output (DRNG/PRNG)
 - ▶ Non-deterministic (NRBG/TRNG)
 - ▶ “Full entropy” vs. cryptographic strength

Application interfaces

- ▶ PKCS#11: C_GenerateRandom()
 - ▶ Java: SecureRandom()
 - ▶ Microsoft CAPI and CNG: CryptGenKey()
 - ▶ OpenSSL: RAND_bytes()
 - ▶ Others
-
- ▶ OASIS KMIP – client/server network protocol

Conclusion

- ▶ Relationships
 - ▶ Entropy
 - ▶ Random numbers
 - ▶ Cryptographic keys
- ▶ Random bit generator construction
- ▶ Entropy sources
- ▶ Measuring and testing entropy and randomness
- ▶ Putting it all together

— Learning points

- ▶ Entropy
 - ▶ Identify sources of entropy and assess min entropy
- ▶ RBG construction
 - ▶ RBG contains required and approved components
- ▶ Seeding and re-seeding PRNGs
 - ▶ PRNG is seeded appropriately before use
 - ▶ PRNG is re-seeded as required to meet security requirements
- ▶ Correctly using the API
 - ▶ Use the correct functions in the correct order

Thank you!

John Leiseboer
QuintessenceLabs

jl@quintessencelabs.com
www.quintessencelabs.com



RSAC CONFERENCE
EUROPE 2013