Entropy Source and DRNG Manager

Stephan Müller <smueller@chronox.de>

October 9, 2022

Abstract

Almost all parts of cryptography rest on a random number generator which produces data that are indistinguishable from a perfect generator. This also implies that the random number generator is seeded with sufficient entropy and the entropy is maintained during processing. Also a reseed at proper intervals is important to maintain the security strength. Finally, different users and use cases have different requirements regarding seeding and the deterministic processing. All these tasks are non-trivial in nature. The Entropy Source and DRNG Manager (ESDM) provides a flexible and extensible framework. The ESDM also provides a set of well-known interfaces: an API and ABI compliant drop-in replacement for Linux /dev/random, /dev/urandom and getrandom(2) system call.

Contents

List of Figures

List of Tables

1 Introduction

The ESDM originated from the Linux Random Number Generator (LRNG) patch series. Considering the limitations of maintaining a proper entropy source management inside the confinements of a kernel and the fact that such actions do not need the privileges of executing inside the kernel, the LRNG patch series was converted into user space to form the ESDM. Even in user space, the ESDM can execute completely unprivileged and thus adds to the overall system security to isolate the operation.

Before discussing the design of the ESDM, the goals of the ESDM design are enumerated:

1. The ESDM manages the proper seeding and reseeding of DRNGs. In addition, it provides internal entropy sources which the ESDM fully controls as well as interfaces to obtain data from external entropy sources.

2. The ESDM provides a full API and ABI drop-in replacement for the Linux device files of /dev/random and /dev/urandom as well as the Linux getrandom(2) system call and the getentropy(3) library calls. All code re-

lated to these interfaces is solely implemented in user space and executed with limited or no privileges.

3. The ESDM is modular consisting of a central server which is implemented with minimal dependencies. This allows the ESDM to execute on different operating systems. The aforementioned Linux interfaces are implemented with separate processes and libraries and therefore their absence do not affect the ESDM operation on operating systems other than Linux.

4. All user-visible behavior implemented by the kernel /dev/random – such as the per-NUMA-node DRNG / per-CPU instances are provided by the ESDM as well.

5. The ESDM must not use locking in hot code paths to limit the impact on massively parallel systems.

6. The ESDM must handle modern computing environments without a degradation of entropy. The ESDM therefore must work in virtualized environments, with SSDs, on systems without HIDs or block devices and so forth.

7. The ESDM must provide a design that allows quantitative testing of the entropy behavior.

8. The ESDM must use testable and widely accepted cryptography for conditioning.

9. The ESDM must allow the use of cipher implementations backed by architecture specific optimized assembler code or even hardware accelerators. This provides the potential for lowering the CPU costs when generating random numbers – less power is required for the operation and battery time is conserved.

10. The ESDM must separate the cryptographic processing from the entropy source maintenance to allow a replacement of these components.

11. The ESDM shall offer flexible configurations allowing vendors to apply the settings applicable to their environment.

1.1 Properties Offered by the ESDM

Apart from the fact that a user does not need to manage the DRNG and its seeding status, the ESDM provides the following properties making the ESDM a contemporary and future-proof entropy source and DRNG management framework:

- Pure user space implementation:
  - The entire ESDM as well as the Linux interfaces are fully implemented in user space.
  - Two additional but optional entropy sources are provided for resource-constrained systems: the scheduler-based entropy source and the interrupt-based entropy source. Both, however, uses a kernel extension as it hooks into the scheduler / the interrupt handling code path.

  - The ESDM server executes without any privilege. The Linux interfaces providing /dev/random, /dev/urandom and the associated proc files execute with limited privileges.

- Sole use of crypto for data processing:
  - Exclusive use of a hash operation for conditioning entropy data with a clear mathematical description as given section 2.2 – non-cryptographic operations like LFSR are not used.
  - The ESDM uses only properly defined and implemented cryptographic algorithms unlike the use of the SHA-1 transformation in the Linux kernel /dev/random implementation that is not compliant with SHA-1 as defined in FIPS 180-4.
  - Hash operations use on-stack hash instances to benefit large parallel systems.
  - ESDM uses limited number of data post-processing steps as documented in section 2.2 compared to the large variation of different post-processing steps in the legacy /dev/random implementation that have no apparent mathematical description.

- Performance
  - High-performance implementation of scheduler entropy source requiring only a few cycles to collect entropy.
  - High-performance implementation of interrupt entropy source requiring only a few cycles to collect entropy which is significantly faster than the existing kernel interrupt handler.

- Testing
  - Availability of run-time health tests of the raw unconditioned entropy source data of the scheduler-based entropy source to identify degradation of the available entropy. Such health tests are important today due to virtual machine monitors reducing the resolution of or disabling the high-resolution timer.
  - Heuristic entropy estimation for the scheduler-based entropy source is based on quantitative measurements and analysis following SP800-90B.
  - Power-on self tests for critical deterministic components (DRNG, hash implementation, and entropy collection logic).
  - Availability of test interfaces for all operational stages of the ESDM including boot-time raw entropy event data sampling as outlined in section 2.15.
  - The ESDM offers a test interface to validate the used software hash implementation and in particular that the ESDM invokes the hash correctly, allowing a NIST ACVP-compliant test cycle – see section 2.15.
  - Availability of stress testing covering the different code paths for data and mechanism (de)allocations and code paths covered with locks.

  - Availability of fully automated regression testing covering different ESDM functions in different configurations allow an unattended functional verification testing.

- Entropy collection of the internal interrupt and scheduler-based entropy source
  - The ESDM scheduler-based and interrupt-based entropy sources are fully compliant to SP800-90B requirements and are shipped with a full SP800-90B assessment and all required test tools.
  - Full entropy assessment and description is provided with chapter 3, specifically section 3.3.6.
  - The ESDM provides a configuration to be compliant to SP800-90C following RBG2(NP) as well as RBG2(P) construction methods.
  - The ESDM provides an interface allowing the request of random numbers with prediction resistance, i.e. requiring an immediate reseed and only returning as many bits of data as entropy was available. This shall support chaining of DRBGs compliant to SP800-90C.
  - The ESDM provides a configuration to be compliant to the German AIS 20/31 NTG.1 considering the aforementioned prediction resistance interface.

- Configurable
  - ESDM build-time configuration allows a flexible adoption of the ESDM to different use cases. Non-compiled additional code is folded into no-ops.
  - Configurable seeding strategies are provided following different concepts.

1.2 Document Structure

This paper covers the following topics in the subsequent chapters:

- The design of the ESDM is documented in chapter 2. The design discussion references to the actual implementation whose source code is publicly available.

- The statistical testing of the internal interrupt entropy source including the SP800-90B compliance assessment is provided in chapter 3.

- The statistical testing of the internal scheduler entropy source including the SP800-90B compliance assessment is provided in chapter 4.

- The discussion of various configurations offered by the ESDM is given in chapter 5.

- The various appendices cover miscellaneous topics supporting the general description.

2 ESDM Design

The ESDM can be characterized with figure 2.1 which provides a big picture of the ESDM processing and components.

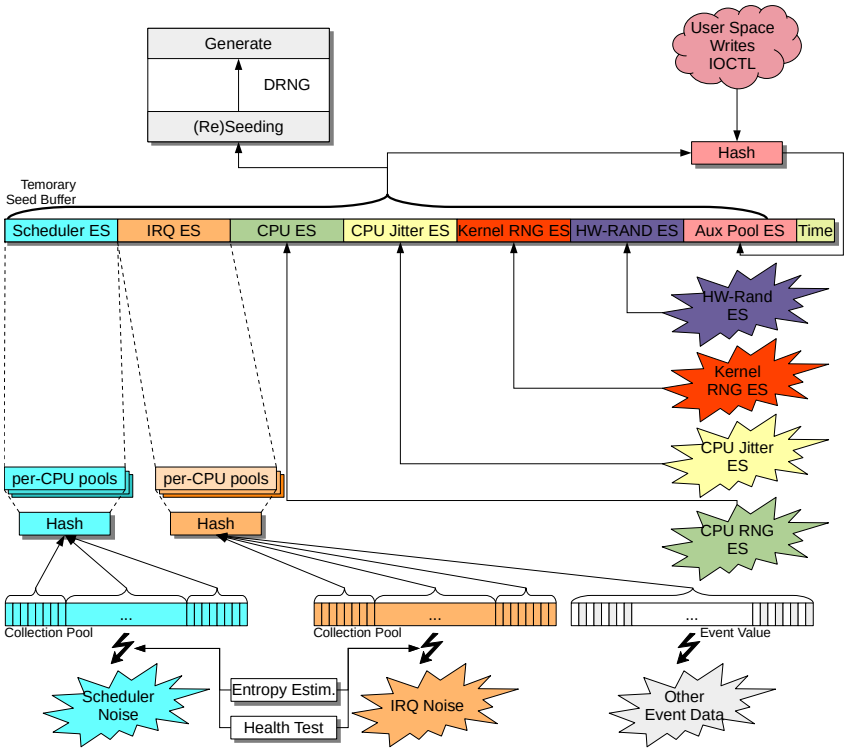

Figure 2.1: ESDM Big Picture

The colors indicate the different entropy sources managed by the ESDM.

The ESDM introduces the concept of slow and fast entropy sources. Fast entropy sources provide entropy at the time of request. A slow entropy source collects data over time into an entropy pool.

The entropy source that is under full control of the ESDM, also called the internal entropy source, comprises of the:

- The indigo marked parts refer to the scheduler entropy source that feeds per-CPU collection pools. These collection pools are hashed into per-CPU entropy pools when seeding of the DRNG is to be performed.

- The orange marked parts refer to the interrupt entropy source which has a very similar architecture compared to the scheduler-based entropy source: data is fed in per-CPU collection pools to hash them into per-CPU entropy pools.

The following external entropy sources are present. These entropy sources are expected to be fully self-contained. The ESDM only requests data from them and expects that the entropy estimate provided with the data is correct:

- The auxiliary pool collects data from external entropy sources which deliver data at times not controllable by the ESDM.
- The CPU entropy source obtains data from a potentially existing source in the CPU like RDSEED on Intel CPUs.
- The Jitter RNG entropy source is another external entropy source.
- The random number generator provided with the `getrandom` system call or the /dev/random device file can be enabled at compile time to serve as an entropy source.
- Using the Linux hardware random number generator (/dev/hwrng) as seed source that may obtain data from various hardware sources including TPM v2.0.

The ESDM treats all external and internal entropy sources equally. It can handle the situations where one or more entropy sources returns little or no entropy.

The the scheduler entropy source is assessed in chapter 4 which provides its complete entropy assessment. All other entropy sources are expected to provide their own entropy assessment supporting the claim of the supplied entropy that is credited by the ESDM for these sources. The ESDM treats these additional entropy sources as black boxes and take their claimed entropy rate at face value. The ESDM, however, guarantees that all entropy sources are processed in compliance with defined standards.

The different colors used in figure 2.1 depict the different entropy sources mentioned before.

The ESDM offers various interfaces to obtain random numbers as depicted in figure 2.2.

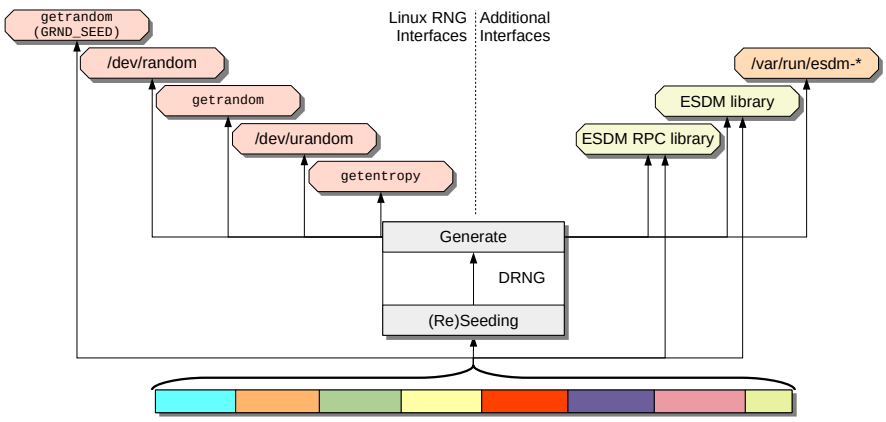

Figure 2.2: ESDM Interfaces To Obtain Random Numbers

Based on this figure, the following types of interfaces are available which are marked with the different colors in figure 2.2. Each type of interface can be selectively enabled and disabled at compile time:

- Linux RNG interfaces: The common interfaces found from today's Linux kernels include /dev/random, /dev/urandom, the `getrandom` system call

or the `getentropy` C-library wrapper. These interfaces behave identically to the Linux RNG and thus allow the ESDM to act as a drop-in replacement. When not compiling those, the ESDM can be operated in parallel with the Linux RNG and all user space software divert to use the ESDM.

- The ESDM also exports a set of Unix Domain Sockets to /var/run which can be used directly, albeit it is not advised to be used as the protocol offered by those sockets is defined to be not stable, i.e. it can change from one version to another.

- To wrap the mentioned Unix Domain Sockets, the ESDM RPC library is available which offers a stable API.

- The ESDM library provides the heart of the ESDM functionality. It could be directly used by a calling application. Yet, it is advised to use the RPC interface as it is provided by a daemon that uses the ESDM library already.

The ESDM consists of the following components:

- `libesdm.so`: The ESDM provides a library with the core of the ESDM. The DRNG and entropy source managers together with all entropy sources and cryptograpphic algorithm implementations are implemented with this library. This library is wrapped by the `esdm-server` listed below. The API for using the library is exported by and documented with `esdm.h`.

  The library is provided in case a user wants to employ the ESDM in his projects instead of the `esdm-server`. Yet, the `esdm-server` provides a wrapping daemon to the ESDM library that is intended to be commonly used.

  In addition, the ESDM can be configured with the options specified in `esdm_config.h`. This would need to be performed by the consuming application. When using the tools below, this library can be ignored.

- `esdm-server`: The ESDM server provides the RPC server that encapsulates the ESDM with its random number generator and the entropy source management. When starting the server, the mentioned Unix Domain Sockets are created that allows clients to request services including random numbers.

  The ESDM server can either be started manually or with the provided (and installed) systemd unit file. When using systemd, start the server with 'systemctl start esdm-server'.

  A wrapper library to access the ESDM server RPC interface is provided with `libesdm_rpc_client.so`. Its API is specified and documented with `esdm_rpc_client.h`.

  The `esdm-server` is the backend to all of the following ESDM components.

  Note: The Unix domain sockets of the `esdm-server` are only visible in the respective mount namespace. If you have multiple mount namespaces, you need to start the daemon in each mount namespace or make the files otherwise available if its services shall be available there.

- `esdm-cuse-random`: The ESDM CUSE daemon creates a device file that behaves identically to /dev/random. It must be started as root. Reading, writing and IOCTLs are implemented in an ABI-compatible way.

  The ESDM CUSE daemon can either be started manually or with the provided (and installed) systemd unit file. When using systemd, start the server with 'systemctl start esdm-cuse-random'. Although the daemon creates a /dev/random device, the actual visible operation is atomic (a bind mount) for both creation and destruction of the new device file which implies that the daemon can be started and stopped at any time during runtime of the Linux OS.

  Note: The bind mount is only visible in the respective mount namespace. If you have multiple mount namespaces, you need to start these daemons in each mount namespace or make the files otherwise available if their services shall be available there.

- `esdm-cuse-urandom`: Same as esdm-cuse-random but behaving like /dev/urandom.

- `esdm-proc`: This FUSE file system implements all files found on a Linux system under /proc/sys/kernel/random but pointing to the ESDM server. This process is required to ensure that all interfaces are provided by ESDM. For details about the provided files, see the `random(4)` man page. Note, the kernel exports the /proc/sys/kernel/random information also as `sysctl(8)`. This interface is not covered by 'esdm-proc'.

- `libesdm-getrandom.so`: The library provides a wrapper to the `getrandom` and `getentropy` Linux C-library calls. To use the library for other consumers, use one of the following considerations:

  - Use `LD_PRELOAD="/path/to/libesdm-getentropy.so"` with the intended application.
  - Compile the application or library with the following options:
    - `LDFLAGS += -Wl,--wrap=getrandom,--wrap=getentropy`
    - `LDFLAGS += -lesdm-getrandom`

2.1 ESDM Components

The ESDM consists of the following components shown in figure 2.1:

1. The ESDM implements a DRNG. Unless using the interface providing prediction resistance, the DRNG always generates the requested amount of output. When using the SP800-90A terminology it operates without prediction resistance. The DRNG maintains a counter of how many bytes were generated since last re-seed and a timer of the elapsed time since last re-seed. If either the counter or the timer reaches a threshold, the DRNG is seeded from the entropy sources with the available entropy.

   In case the Linux kernel detects a NUMA system, one DRNG instance per NUMA node is maintained.

   Depending on the used interface to request data from the DRNG, the caller may be put to sleep until the ESDM is fully seeded:

(a) All interfaces using the `esdm_get_entropy_bytes` function always generates data including when the ESDM is not properly seeded.

(b) All interfaces using the `esdm_get_entropy_bytes_full` function generate data only when the ESDM is fully seeded and fully initialized.

(c) When using the `esdm_get_entropy_bytes_pr` function, the special DRNG instance which operates with prediction resistance is used such that it is reseeded from the entropy sources and generates at most as much data as entropy was added. This implies the DRNG operates with prediction-resistance (SP800-90A terminology) or is NTG.1 compliant (German BSI AIS 20/31 from 2011). In addition, when the ESDM executes in FIPS mode, the behavior complies with all FIPS requirements allowing the output of the API to be claimed to be from a vetted conditioning component. I.e. the caller may safely use this data to seed another DRBG. To achieve this, the following checks are applied in the function `esdm_drng_get`:

- Require a reseed - in FIPS mode, the ESDM requires the availability of at least 256 bits of entropy in its entropy sources for the reseed to commence.
- Produce at most only 256 bits of random bits from the DRNG (i.e. the security strength of the DRNG).

2. The DRNG is seeded by concatenating the data from the following sources in case they are enabled at compile time:

(a) the output of the auxiliary pool,

(b) the output of the per-CPU interrupt entropy pools,

(c) the output of the per-CPU scheduler entropy pools,

(d) the Jitter RNG if available,

(e) the CPU-based entropy source such as Intel RDSEED if available,

(f) the kernel RNG output, and

(g) the Linux kernel HWRAND framework (/dev/hwrng).

The entropy estimate of the data of all entropy sources are added to form the entropy estimate of the data used to seed the DRNG with. The ESDM ensures, however, that the DRNG after seeding is at maximum the security strength of the used DRNG implementation of 256 bits.

The ESDM is designed such that none of these entropy sources can dominate the other entropy sources to provide seed data to the DRNG due to the following:

(a) During boot time, the amount of available entropy is the trigger point to (re)seed the DRNG following the explanation in the next section.

(b) At runtime, the the DRNG reseed is triggered by either the DRNG due to hitting the aforementioned thresholds or by a user space caller. The reseed is never triggered by the entropy sources.

3. To support backward secrecy, the following steps are applied:

(a) The temporary seed buffer holding the concatenation of data from all entropy sources to seed the DRNG is injected into the auxiliary pool like other data by hashing it together with the existing auxiliary pool data to form the new auxiliary pool content. The injection of the temporary seed buffer will not alter the entropy estimation of the auxiliary pool.

(b) The message digest created for each per-CPU entropy pool of the scheduler and interrupt entropy sources is inserted into the corresponding per-CPU entropy pool.

The ESDM allows the DRNG mechanism and the used hash to be changed at compile time. Per default, an SP800-90A Hash DRBG implementation along with a SHA-512 hash implementation is available. Both are standard C implementations which were tested against NIST's ACVP service.

The following subsections cover the different components of the ESDM from the bottom to the top.

2.2 ESDM Data Processing

The processing of entropic data from the different entropy source before injecting them into the DRNG is performed with the following mathematical operations. The operation $SHA()$ refers to the hash operation using the message digest implementation that is currently present, i.e. SHA-512.

2.2.1 Scheduler and Interrupt Entropy Sources

1. Truncation: The time stamps received by the IRQ as well as the scheduler entropy sources are truncated to 8 least significant bits (or 32 least significant bits during boot time) – note the GCD is a value calculated during initialization and is fixed thereafter which implies that the time stamp divided by the GCD is the raw entropy value: t_8 (or t_{32})

2. Concatenation: The time stamps received and truncated by the IRQ and scheduler entropy sources as well as auxiliary 32 bit words a_{32} are concatenated to fill the per-CPU collection pool that is capable of holding 1,024 8-bit words[1] - the order of the data a_{32} or t_8 present in the concatenation depends on the occurrence of events - the following formula depicts one possible order for illustration - the implementation is provided with functions `_esdm_pcpu_array_add_u32` and `esdm_pcpu_array_add_slot`:

$$CP = t_{8_{n-1019}}||a_{32_n}||t_{8_{n-1018}}|| \ ... \ ||t_{8_n} \tag{2.1}$$

   Note: In case the continuous compression operation is disabled for the IRQ entropy source, the auxiliary 32 bit words a_{32} are discarded and are not injected into the collection pool. This approach is taken to prevent non-entropy data to potentially overwrite entropy data in the collection pool when the array wraps. The scheduler entropy source only records time stamps.

[1]The ESDM collection size is compile-time configurable where 1,024 is a default value. When configuring a different value, the number of the concatenated data must be adjusted as needed. However, this modification has no impact to the illustration of the data processing.

2.2.2 Interrupt Entropy Source

1. Hashing: For the IRQ entropy source all concatenated time stamp data received from the interrupts since the last output generation of the per-CPU entropy pool $EP_{CPU_{n-1}}$ are hashed together with that last output $EP_{CPU_{n-1}}$ to generate new per-CPU entropy pool output of EP_{CPU_n}. The following steps are performed:

   (a) One filled per-CPU collection pool for the interrupt entropy source CP_{IRQ_m} is inserted into the per-CPU entropy pool using a hash update operation.

   (b) To generate data from the entropy pool EP_{CPU_n} as used by function 2.3, a hash final operation is performed.

   (c) Once a hash final operation is performed it is followed by an immediate re-initialization of the hash state with a hash init operation and adding the just calculated message digest with the first hash update.

   The implementation is provided with function `esdm_pcpu_array_compress` together with the function `esdm_pcpu_pool_hash_one` generating data from the per-CPU entropy pool:

   $$EP_{CPU_n} = SHA(EP_{CPU_{n-1}}||CP_{IRQ_{m-(n-1)}}||...||CP_{IRQ_{m-1}}||CP_{IRQ_m}) \tag{2.2}$$

   Note: The hash update operation is performed at the following occasions:

   (a) Continuous compression enabled: The hash update is performed every time the collection pool is full. This operation therefore is performed in interrupt context. In addition, the operation is performed at the time operation of equation 2.3 is invoked which is in process context.

   (b) Continuous compression enabled and disabled: The hash update is performed at the time the operation of equation 2.3 is invoked, i.e. at the time the DRNG is reseeded. This operation therefore is performed in process context. This guarantees that all unprocessed entropy data in the collection pool is added to the entropy pool at the time the entropy pool is requested for random data.

   This implies that in case of disabled continuous compression, the oldest entries in the collection pool are overwritten with newer entropy event data when more entropy events are collected than can be held in the collection pool between DRNG reseeds.

2. Hashing: For the IRQ entropy source, a message digest of all per-CPU entropy pools is calculated. This message digest is used to fill the interrupt entropy source output buffer S discussed in the following - the implementation is provided with function `esdm_pcpu_pool_hash`:

   $$EP_{all_n} = SHA(EP_{CPU0_n}||EP_{CPU1_n}|| \ ... \ ||EP_{CPUX_n}) \tag{2.3}$$

3. Truncation: For the interrupt entropy pool, the most-significant bits (MSB) defined by the requested number of bits (commonly equal to the security strength of the DRBG) or the entropy available transported with the

buffer (which is the minimum of the message digest size and the available entropy in all entropy pools), whatever is smaller, are obtained from the interrupt entropy source output buffer S - the implementation is provided with function `ESDM_pcpu_pool_hash`:

$$E_n = MSB_{min(entropy,security\ strength)}(EP_{all_n}) \tag{2.4}$$

2.2.3 Scheduler Entropy Source

1. Hashing: For the scheduler-based entropy source, all concatenated time stamp data received from the interrupts since the last output generation of the per-CPU entropy pool $SP_{CPU_{n-1}}$ are hashed together with that last output $SP_{CPU_{n-1}}$ to generate new per-CPU entropy pool output of SP_{CPU_n}. The following steps are performed:

   (a) One filled per-CPU collection pool for the interrupt entropy source CP_{SCHED_m} is inserted into the per-CPU entropy pool using a hash update operation.

   (b) To generate data from the entropy pool SP_{CPU_n} as used by function 2.6, a hash final operation is performed.

   (c) Once a hash final operation is performed it is followed by an immediate re-initialization of the hash state with a hash init operation and adding the just calculated message digest with the first hash update.

   The implementation is provided with function `esdm_sched_pool_hash_one` generating data from the per-CPU entropy pool:

$$SP_{CPU_n} = SHA(SP_{CPU_{n-1}}||CP_{SCHED_{m-(n-1)}}||...||CP_{SCHED_{m-1}}||CP_{SCHED_m}) \tag{2.5}$$

   Note: The hash update is performed at the time the operation of equation 2.6 is invoked, i.e. at the time the DRNG is reseeded. This operation therefore is performed in process context. This guarantees that all unprocessed entropy data in the collection pool is added to the entropy pool at the time the entropy pool is requested for random data. This implies that he oldest entries in the collection pool are overwritten with newer entropy event data when more entropy events are collected than can be held in the collection pool between DRNG reseeds.

2. Hashing: For the scheduler-based entropy source, a message digest of all per-CPU entropy pools is calculated. This message digest is used to fill the interrupt entropy source output buffer S discussed in the following - the implementation is provided with function `esdm_sched_pool_hash`:

$$SP_{all_n} = SHA(SP_{CPU0_n}||SP_{CPU1_n}||\ ...\ ||SP_{CPUX_n}) \tag{2.6}$$

3. Truncation: Just like the interrupt entropy source, the scheduler entropy source applies a truncation to the generated data as implemented by the function `esdm_sched_pool_hash`:

$$S_n = MSB_{min(entropy,security\ strength)}(SP_n) \tag{2.7}$$

2.2.4 Auxiliary Entropy Pool

1. Hashing: When new data D_m is added to the auxiliary pool AP, the data is inserted into the auxiliary pool with a hash update operation - the implementation is provided with function `ESDM_pool_insert_aux`. The message digest generation is performed at the time entropy from the auxiliary pool is requested. To ensure backward secrecy, the temporary seed buffer T_{n-1} that holds among others the auxiliary pool digest from the previous generation round as depicted with equation 2.12 is concatenated with the received data:

$$AP_n = SHA(T_{n-1}||D_{m-(n-1)}||...||D_{m-1}||D_m) \tag{2.8}$$

2. Truncation: The MSB of the auxiliary pool of the size of the DRNG security strength are used for the seed buffer:

$$A_n = MSB_{min(digest\,size,security\ \ strength)}(AP_n) \tag{2.9}$$

2.2.5 CPU Entropy Source

1. Hashing: If the CPU entropy source provides less than full entropy, a message digest of the amount of data to be requested from it is calculated:

$$C_{cond} = SHA(C_1||\,...\,||C_m) \tag{2.10}$$

2. Truncation: If the CPU entropy source provides less than full entropy, the MSB defined by the requested number of bits (commonly equal to the security strength of the DRBG) or the applied message digest size, what ever is smaller, are obtained - the implementation is provided with function `ESDM_get_arch_data_compress` – otherwise C_n is the data obtained directly from the CPU entropy source:

$$C_n = MSB_{min(digest\,size,security\,strength)}(C_{cond}) \tag{2.11}$$

2.2.6 Temporary Seed Buffer Construction

1. Concatenation: The temporary seed buffer T used to seed the DRNG at the time n is a concatenation of one or more of the following entropy source data sets, depending on the compile-time configuration:

   (a) the auxiliary pool entropy source A,
   (b) the interrupt entropy source buffer E,
   (c) the scheduler entropy source buffer S,
   (d) the Jitter RNG output J,
   (e) the CPU entropy source output C,
   (f) the kernel RNG entropy source output K,
   (g) the Linux HWRAND framework H,and
   (h) the current time t

   with the implementation is provided with function `esdm_fill_seed_buffer`:

$$T_n = A_n||E_n||S_n||J_n||C_n||K_n||H_n||t \tag{2.12}$$

2.3 ESDM Architecture

Before going into the details of the ESDM processing, the concept underlying the ESDM shown in figure 2.1 is provided here.

The entropy derived from the slow entropy sources is collected and accumulated in the entropy pools which contain already compressed entropy data, supported by the collection pools which contain uncompressed, but only concatenated entropy data.

At the time the DRNG shall be seeded, the all entropy pools, any non-compressed data in the collection pools and the auxiliary pool are processed with a cryptographic hash function which can be chosen at runtime.

For the entropy pool, if the digest of the hash and the available entropy are larger than requested by the caller, the digest is truncated to the appropriate size. For the auxiliary pool, always 256 bits of data are returned irrespective of the entropy rate of this pool. This ensures that also data that is not credited with entropy but injected into the ESDM is used to stir the seed for the DRNG.

The DRNG always tries to seed itself with 256 bits of entropy, except during boot. In any case, if the entropy sources cannot deliver that amount, the available entropy is used and the DRNG keeps track on how much entropy it was seeded with. The entropy implied by the ESDM available in the entropy pool may be too conservative. To ensure that during boot time all available entropy from the entropy pool is transferred to the DRNG, the hash function always generates 256 data bits during boot to seed the DRNG. During boot, the DRNG is seeded as follows:

1. The DRNG is reseeded from the entropy sources if all entropy sources collectively have at least 256 bits of entropy available.

At the time of the reseeding steps, the DRNG requests as much entropy as is available in order to skip certain steps and reach the seeding level of 256 bits. This may imply that one or more of the aforementioned steps are skipped.

In all listed steps, the DRNG is (re)seeded with a number of random bytes from the entropy pool that is at most the amount of entropy present in the entropy pool. This means that when the entropy pool contains 128 or more bits of entropy, the DRNG is seeded with that amount of entropy as well.

Before the DRNG is seeded with 256 bits of entropy in the last step, requests of random data from the blocking interfaces are not processed.

At runtime the DRNG operates as deterministic random number generator with the following properties:

- The maximum number of random bytes that can be generated with one DRNG generate operation is limited to 4096 bytes. When longer random numbers are requested, multiple DRNG generate operations are performed. The used DRNGs implement an update of their state during the generation operation for backward secrecy.

- The DRNG is reseeded with whatever entropy is available, but at least 128 bits (256 bits if SP800-90C compliance is enabled) – in the worst case where no additional entropy can be provided by the entropy sources, the DRNG is not re-seeded and continues its operation to try to reseed again after again the expiry of one of these thresholds:

– If the last reseeding of the DRNG is more than 600 seconds ago[2], or

– 2^{20} DRNG generate operations are performed, whatever comes first, or

– the DRNG is forced to reseed before the next generation of random numbers if data has been injected into the ESDM by writing data into /dev/random or /dev/urandom or the respective RPC interfaces.

The chosen values prevent high-volume requests from user space to cause frequent reseeding operations which drag down the performance of the DRNG[3].

- If the DRNG was not reseeded for the last 2^{30} DRNG generate operations – i.e. the reseeding requests discussed in the previous bullets were unsuccessful – the DRNG reverts back to an unseeded state. This applies that the DRNG will not produce random numbers when accessed via the blocking interfaces. In this case, the DRNG behaves like during boot time.

With the automatic reseeding after 600 seconds, the ESDM is triggered to reseed itself before the first request after a suspend that put the hardware to sleep for longer than 600 seconds.

2.3.1 Minimally Versus Fully Seeded Level

The ESDM's DRNG is reseeded when the first 128 bits / 256 bits of entropy are received during boot as indicated above. The 128 bits level defines that that the DRNG is considered "minimally" seeded whereas reaching the 256 bits level is defined as the DRNG is "fully" seeded.

Both seed levels have the following implications:

- Upon reaching the minimally seeded level, the kernel-space callers waiting for a seeded DRNG via the API calls of either `esdm_get_random_bytes_min` is woken up.

- When reaching the fully seeded level, the user-space callers waiting for a fully seeded DRNG via the `getrandom` system call or /dev/random are woken up. Using the ESDM API of `esdm_get_random_bytes_full`, the caller is waiting synchronously until the fully seeded level is reached.

2.3.2 NUMA Systems

To prevent bottlenecks in large systems, the DRNG will be instantiated once for each NUMA node. The instantiations of the DRNGs happen all at the same time when the ESDM is initialized.

[2]Note, this value will not empty the entropy pool even on a completely quiet system. Testing of the ESDM was performed on a KVM without fast entropy sources and with a minimal user space, where only the SSH daemon was running, During the testing, no operation was performed by a user. Yet, the system collected more than 256 bits of entropy from the interrupt entropy source within that time frame, satisfying the DRNG reseed requirement.

[3]Considering that the maximum request size is 4096 bytes defined by `ESDM_DRNG_MAX_REQSIZE` (i.e. each request is segmented into 4096 byte chunks) and at most 2^{20} requests defined by `ESDM_DRNG_RESEED_THRESH` can be made before a forced reseed takes place, at most $4096 \cdot 2^{20} = 4,294,967,296$ bytes can be obtained from the DRNG without a reseed operation.

The question now arises how are the different DRNGs seeded without re-using entropy or relying on random numbers from a yet insufficiently seeded ESDM. The ESDM seeds the DRNGs sequentially starting with the one for NUMA node zero – the DRNG for NUMA node zero is seeded with the approach of 256 bits of entropy stepping discussed above. Once the DRNG for NUMA node 0 is seeded with 256 bits of entropy, the ESDM will seed the DRNG of node one when having again 256 bits of entropy available. This is followed by seeding the DRNG of node two after having again collected 256 bits of entropy, and so on. Figure 2.3 illustrates the seeding strategy showing that each DRNG instance is freshly seeded with a separate seed buffer.

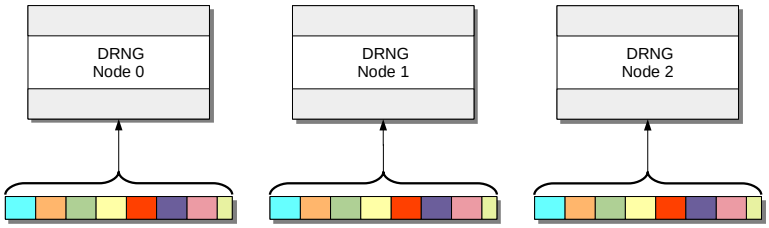

Figure 2.3: DRNG Instances on NUMA systems with seeding strategy

When producing random numbers, the ESDM tries to obtain the random numbers from the NUMA node-local DRNG. If that DRNG is not yet seeded, it falls back to using the DRNG for node zero.

Note, to prevent draining the entropy pool on quiet systems, the time-based reseed trigger, which is 600 seconds per default, will be increased by 100 seconds for each activated NUMA node beyond node zero. Still, the administrator is able to change the default value at runtime.

2.3.3 Flexible Design

Albeit the preceding sections look like the DRNG and the management logic are highly interrelated, the ESDM code allows for an easy replacement of the DRNG with another deterministic random number generator. This flexible design allowed the implementation of the ChaCha20 DRNG if the SP800-90A DRBG is not desired.

To implement another DRNG, all functions in `struct esdm_drng_cb` in "esdm_crypto.h" must be implemented. These functions cover the allocation/deallocation of the DRNG as well as its usage. Similarly, all functions in `struct esdm_hash_cb` from "esdm_crypto.h" must be implemented to provide the the conditioning hash.

The implementations can be changed at compile time. The default implementation is the SP800-90A Hash-DRBG using a software-implementation of the used SHA-512 message digest for accessing the entropy pools.

In addition, the ESDM allows the addition of new entropy sources.

2.4 ESDM Data Structures

The ESDM uses the following main data structures:

- The data from the interrupt entropy source is processed with a per-CPU entropy pool. In addition, a per-CPU collection pool that can hold the concatenated time stamps is maintained. Both are accessed lockless since the currently executing CPU's entropy pool and collection pool is used. During access to the entropy pool, the ESDM though takes a lock since the entropy pool is also read when the hash is calculated for filling the seed buffer. As the filling of the seed buffer is very infrequently (see above for the reseed periods of the DRNG), the lock is hardly contented which allows the conclusion that the entropy collection operates quasi-lockless.

- The scheduler entropy source defines per-CPU collection pools. The entropy pools are not maintained as compression of scheduler-based entropy sources is only performed when the DRNG shall be (re)seeded. The entire scheduler-based entropy source operates lock-less.

- The cryptographic algorithm data structures hold the reference to the DRNG instance and the hash instance and associated meta data needed for its operation. When using the DRNG, a full read/write lock is used to guard (a) against replacement of the DRNG reference while operating on the DRNG state, and (b) to read/write the DRNG state. Contrarily when using the hash, only a read-lock is used to guard against the replacement of the hash reference. This implies that the hash state is kept on the stack of the calling application.

2.5 Interrupt Processing - ESDM-internal Entropy Source

The ESDM hooks a callback into the bottom half interrupt handler at the same location where the legacy /dev/random places its callback hook.

The ESDM interrupt processing callback is a void function that also does not receive any input from the interrupt handler. That interrupt processing callback is the hot code path in the ESDM and special care is taken that it is as short as possible and that it operates without locking.

Figure 2.4 illustrates the interrupt processing performed by the ESDM. The figure specifies which parts of the interrupt processing execute in IRQ context and which executes in process context. The operations executed in interrupt context are all completely listed in this section. All steps executed in process context are illustrated in section 2.11.

The figure depicts the example when one interrupt arrives on CPU 0. If an interrupt arrives on another CPU, the same operation is applied, but the respective CPU-local collection pool and entropy pool is used. The entropy pools from other CPUs in the figure therefore are filled with the same processing steps, which, however, are not shown.

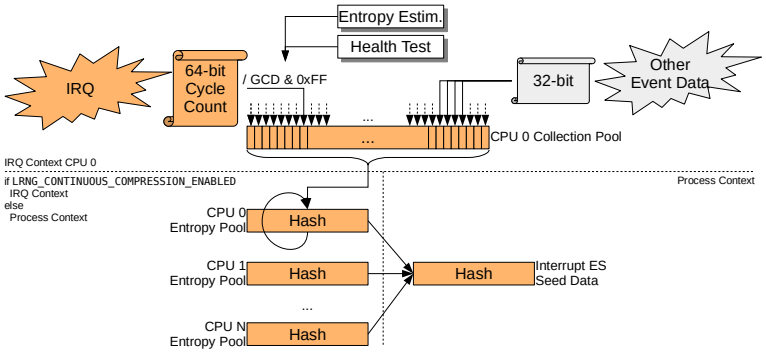

Figure 2.4: Interrupt Processing

The following processing happens when an interrupt is received and the ESDM is triggered:

1. A high-resolution time stamp is obtained using the service `random_get_entropy` kernel function. This integer value is divided by a GCD to eliminate bits that do not change. Although that function returns a 64-bit integer, only the bottom 8 bits, i.e. the fast moving bits, are used for further processing. To ensure fast processing, these 8 bits are concatenated and stored in the operating CPU's data collection pool. After the receipt of 1,024 time stamps, the data collection pool with all concatenated time stamps is inserted into the currently executing CPU's entropy pool. During boot time until the ESDM completed the calculation of the GCD, the 32 least significant bits of the data are directly inserted into the CPU's entropy pool. Entropy is contained in the variations of the time of events and its time delta variations. Figure 2.1 depicts the time stamp array holding the 8-bit time stamp values.

2. The health tests discussed in section 2.6.2 are performed on each received time stamp where the truncated time stamp value is forwarded to the health test. Unless 1,024 time stamps have been received, the processing of an interrupt stops now.

3. The per-CPU collection pool is added to the same CPU's entropy pool by performing a hash update operation. This approach works as the per-CPU entropy pool is managed as the message digest state. When data of the per-CPU entropy is to be extracted, a hash final operation is performed followed by an immediate re-initialization of the state buffer using the message digest of the previous extraction. This operation is depicted in figure 2.5 for the entropy pool maintained by CPU 0. The other CPUs perform the same processing with their independent copy of the collection pool and the entropy pool. In case the continuous compression support is disabled, the hash operation is not performed. Instead, the oldest entropy values in the collection pool are overwritten with the latest entropy value. In case the continuous compression operation is disabled, the hash update operation is conducted in process context at the time of obtaining random numbers from the entropy pool requested to seed the DRNG documented in section 2.11.

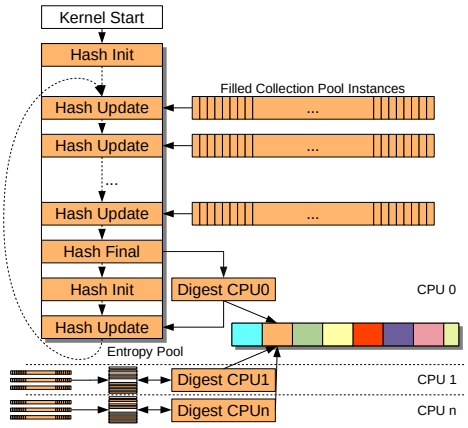

Figure 2.5: Collection Pool Processing

4. The ESDM increases the per-CPU counter of the received interrupt events by the number of healthy interrupts stored in the per-CPU collection pool. This counter is translated into an entropy statement when the ESDM wants to know how much entropy is present in the entropy pool. This counter is also adjusted when reading data from the entropy.

5. The ESDM server implements a monitor thread that polls the interrupt ES for its entropy level until all DRNGs become fully seeded. While the poll monitor is active, every time the interrupt ES entropy rate reaches the security strength of the ESDM, a reseed is triggered as discussed in section 2.11.

6. If all DRNG instances are fully seeded, the poll monitor stops. This implies that only during boot time the next step is triggered. At runtime, the entropy sources will not trigger a reseeding of the DRNG.

The entropy collection mechanism is available immediately after the kernel RNG reached full entropy. This shall guarantee that the kernel RNG is fully seeded before hijacking its main entropy source. To prevent starving the kernel RNG, the ESDM server starts a thread that reseeds the kernel RNG once every 2 minutes with 256 bits of data.

In case the underlying system does not support a high-resolution time stamp, step 2 in the aforementioned list is changed to fold the following 32 bit values each into one bit and XOR all of those bits to obtain one final bit:

- IRQ number,
- High 32 bits of the instruction pointer,
- Low 32 bits of the instruction pointer,
- A 32 bit value obtained from a register value – the ESDM iterates through all registers present on the system.

2.5.1 Entropy Amount of Scheduling Events

The discussion of the entropy content of interrupt events in section 2.6.1 applies to the scheduler-based entropy. The only difference is the use of a different compile-time entropy value of `ESDM_IRQ_ENTROPY_RATE`.

2.5.2 Health Tests

The health tests documented in section 2.6.2 are applied to the scheduler entropy source as well. The ESDM ensures that the health tests for the interrupt and the scheduler entropy sources are strictly separated. This separation applies to the SP800-90B as well.

2.6 Scheduler Events - ESDM-internal Entropy Source

The ESDM hooks into the scheduling operation of the Linux kernel which is triggered every time a context switch is performed as initiated by the scheduler. The ESDM handling of a scheduling event is depicted with figure 2.6.

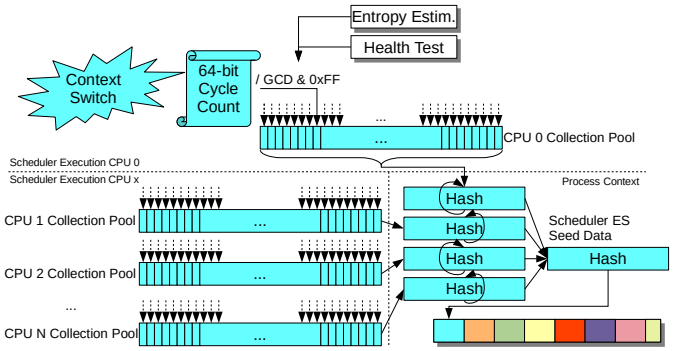

Figure 2.6: Scheduler-event Processing

When a context switch occurs, the ESDM callback is invoked which obtains a high-resolution time stamp that is concatenated into the CPU-local collection pool. When the CPU-local collection pool is full, the oldest entries are overwritten by the latest time stamp.

At the time the DRNG shall be (re)seeded, a message digest of each scheduler per-CPU collection pool is calculated to fill the per-CPU scheduler entropy pool as outlined in figure 2.7. This figure shows that first the collection pool is inserted into a per-CPU scheduler entropy pool which then is inserted into the temporary seed buffer. This guarantees that "unused entropy" is appropriately protected by the per-CPU scheduler entropy pool. The data processing concepts between the scheduler ES and the interrupt ES are identical when comparing the figures in this section to section 2.5.

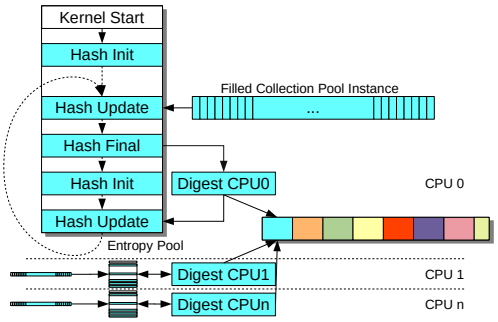

Figure 2.7: Scheduler per-CPU Entropy Pool Management

This entire message digest calculation is performed in the process context. Therefore, the hashing operation is not performed as part of the scheduling operation. The message digest is truncated to the entropy available in all scheduler per-CPU collection pools or the requested amount of entropy, whatever is smaller. Note, the requested amount of entropy is always smaller or equal to the size of the message digest of the used hashing algorithm.

As a side note, the scheduler entropy source potentially has some relationship with the IRQ entropy source because an IRQ may trigger a scheduler event (e.g. with the flag `TIF_NEED_RESCHED`). Therefore, both entropy sources cannot be used at the same time and credit entropy to both. It is permissible to use both, but only one is credited with entropy.

2.6.1 Entropy Amount of Interrupts

The question now arises, how much entropy is generated with the scheduler entropy source. The current implementation implicitly assumes one bit of entropy per time stamp obtained for one scheduling event.

With the kernel compile time parameter of `ESDM_SCHED_ENTROPY_RATE` the number of interrupts that must be collected to obtain 256 bits of entropy can be specified. This value is forced by the ESDM to be at least the aforementioned limit, i.e. 256 interrupts.

The entropy of high-resolution time stamps is provided with the fast-moving least significant bits of a time stamp which is supported by the quantitative measurements shown in section 4. Although only one bit of entropy is assumed to be delivered with a given time stamp the ESDM uses the 8 least significant bits (LSB) of the time stamp to provide a cushion for ensuring that at any given time stamp there is always at least one bit of entropy collected on all types of environments.

However, the question may be raised of why not use more data of the time stamp, i.e. why not using 32 bits or the full 64 bits of the time stamp to increase that cushion? There main answer is performance and memory consumption. The collection of a time stamp is performed as part of an scheduling function. Therefore, the performance of the ESDM in this code section is highly performance-critical. To limit the impact on the scheduler, the ESDM concatenates the 8 LSB of 1,024 time stamps received by the current CPU. The second aspect is that the higher bits of the time stamp must always be considered to have zero bits of entropy when considering the worst case of a skilled attacker.

As the ESDM cannot identify whether it is under attack by a skilled attacker, it always assumes it is under attack.

The Linux kernel allows unprivileged user space processes to monitor the scheduling events of interrupts by reading the process listing with a certain degree of accuracy. The ESDM uses a high-resolution time stamp that executes with nanosecond precision on 1 GHz systems. Local attackers are expected to be measure the occurrence of a scheduling event with a microsecond precision. The distance between a microsecond and a nanosecond is 2^{10}. Thus, when the attacker is assumed to predict the interrupt occurrence with a microsecond precision and the time stamp operates with nanosecond precision, 10 bits of uncertainty remains that cannot be predicted by that attacker. Hence, only these 10 bits can deliver entropy.

To ensure the ESDM interrupt handling code has the maximum performance, it processes time stamp values with a number of bits equal to a power of two. Thus, the ESDM implementation uses 8 LSB of the time stamp (after the time stamp was divided by its GCD).

During boot time, the presence of attackers is considered to be very limited as no remote access is yet possible and no local attack applications are assumed to execute. On the other hand, the performance of the interrupt handler is not considered to be very critical during the boot process. Thus, the ESDM uses the 32 LSB of the time stamp that is injected into the per-CPU collection pool when the time stamp is collected – the ESDM still awards this time stamp one bit of entropy. Once the ESDM completed the calculation of the GCD the aforementioned runtime behavior of concatenating the 8 LSB of 1,024 time stamps before mixing them into the per-CPU entropy pool is enabled.

2.6.2 Health Tests

The ESDM implements the following health tests:

- Stuck Test
- Repetition Count Test (RCT)
- Adaptive Proportion Test (APT)

Those tests are detailed in the following sections.

Please note that these health tests are only performed for the internal entropy sources. Other entropy sources like the entropy sources feeding the auxiliary pool, the Jitter RNG, or the CPU-based entropy sources are not covered by these tests as they are fully self-contained entropy sources where the ESDM does not have access to the raw noise data and does not include a model of the entropy source to implement appropriate health tests. The ESDM considers both as external entropy source. Thus, the user must ensure that either those other entropy sources implement all health tests as needed or the kernel must be started such that these entropy sources are credited with zero bits of entropy. Not crediting any entropy to these other entropy sources can be achieved with the following kernel configuration options:

- Sources feeding the auxiliary entropy pool: The interface functions to provide entropy data have to be invoked with the value 0 for the entropy rate.

- CPU-based entropy source: `ESDM_CPU_ENTROPY_RATE=0`
- Jitter RNG: `ESDM_JENT_ENTROPY_RATE=0`

These options ensure that random data from the entropy sources are pulled, but are not credited with any entropy.

The RCT, and the APT health test are only performed when the kernel is booted with `fips=1` and the kernel detects a high-resolution time stamp generator during boot.

In addition, the health tests are only enabled if a high-resolution time stamp is found. Systems with a low-resolution time stamp will not deliver sufficient entropy for the interrupt entropy source which implies that also the health tests are not applicable.

Stuck Test The stuck test calculates the first, second and third discrete derivative of the time stamp to be processed by the per-CPU collection pool. Only if all three values are non-zero, the received time delta is considered to be non-stuck. The first derivative calculated by the stuck test verifies that two successive time stamps are not equal, i.e. are "stuck". The second derivative calculates that there is no linear repetitive signal.

The third derivative of the time stamp is considered relevant based on the following: The entropy is delivered with the variations of the occurrence of interrupt events, i.e. it is mathematically present in the time differences of successive events. The time difference, however, is already the first discrete derivative of time. Now, if the time difference delivers the actual entropy, the stuck test shall catch that the time differences are not stuck, i.e. the first derivative of the time difference (or the second derivative of the absolute time stamp) shall not be zero. In addition, the stuck test shall ensure that the time differences do not show a linear repetitive signal – i.e. the second discrete derivative of the time difference (or the third discrete derivative of the absolute time stamp) shall not be zero.

Repetition Count Test The ESDM uses an enhanced version of the Repetition Count Test (RCT) specified in SP800-90B [2] section 4.4.1. Instead of counting identical back-to-back values, the input to the RCT is the counting of the stuck values during the processing of received interrupt events. The data that is mixed into the entropy pool is the time stamp. As the stuck result includes the comparison of two back-to-back time stamps by computing the first discrete derivative of the time stamp, the RCT simply checks whether the first discrete derivative of the time stamp is zero. If it is zero, the RCT counter is increased. Otherwise, the RCT counter is reset to zero.

The RCT is applied with $\alpha = 2^{-30}$ compliant to the recommendation of FIPS 140-2 IG 9.8.

During the counting operation, the ESDM always calculates the RCT cut-off value of C. If that value exceeds the allowed cut-off value, the ESDM will trigger the health test failure discussed below. An error is logged to the kernel log that such RCT failure occurred.

This test is only applied and enforced in FIPS mode.

Adaptive Proportion Test Compliant to SP800-90B [2] section 4.4.2 the ESDM implements the Adaptive Proportion Test (APT). Considering that the entropy is present in the least significant bits of the time stamp, the APT is applied only to those least significant bits. The APT is applied to the four least significant bits.

The APT is calculated over a window size of 512 time deltas that are to be mixed into the entropy pool. By assuming that each time stamp has (at least) one bit of entropy and the APT-input data is non-binary, the cut-off value $C = 325$ as defined in SP800-90B section 4.4.2.

This test is only applied and enforced in FIPS mode.

Runtime Health Test Failures If either the RCT, or the APT health test fails irrespective whether during initialization or runtime, the following actions occur:

1. The entropy of the entire entropy pool is invalidated.
2. All DRNGs are reset which imply that they are treated as being not seeded and require a reseed during next invocation.
3. The SP800-90B startup health test are initiated with all implications discussed in section 2.6.2. That implies that from that point on, new events must be observed and its entropy must be inserted into the entropy pool before random numbers are calculated from the entropy pool.

SP800-90B Startup Tests The aforementioned health tests are applied to the first 1,024 time stamps obtained from scheduler events. In case one error is identified for either the RCT, or the APT, the collected entropy is invalidated and the SP800-90B startup health test is restarted.

As long as the SP800-90B startup health test is not completed, all ESDM random number output interfaces that may block will block and not generate any data. This implies that only those potentially blocking interfaces are defined to provide random numbers that are seeded with the interrupt entropy source being SP800-90B compliant. All other output interfaces will not be affected by the SP800-90B startup test and thus are not considered SP800-90B compliant.

To summarize, the following rules apply:

- SP800-90B compliant output interfaces
  - `/dev/random`
  - `getrandom(2)` system call when called with a flag that does not include `GRND_INSECURE`
  - `esdm_get_random_bytes_full` API call
  - `esdm_get_random_bytes_pr` API call
- SP800-90B non-compliant output interfaces
  - `/dev/urandom`
  - `getrandom(2)` system call when called with `GRND_INSECURE`
  - `esdm_get_random_bytes` API call
  - `esdm_get_random_bytes_min` API call

2.7 Auxiliary Entropy Pool - ESDM-external Entropy Sources

The ESDM also supports obtaining entropy from the following data sources and entropy sources that are external to the ESDM. The data is injected into the auxiliary pool.

During the reseeding operation of the DRNG, any user-space entropy provider waiting via select(2) are triggered to provide one buffer full of data. This data is mixed into the auxiliary pool. This approach shall ensure that the ESDM-external entropy sources may provide entropy at least once each DRNG reseed operation.

2.7.1 Injecting Data From User Space

User space can take the following actions to inject data into the DRNG:

- When writing data into /dev/random or /dev/urandom, the data is added to the auxiliary pool and triggers a reseed of the DRNGs at the time the next random number is about to be generated. The ESDM assumes it has zero bits of entropy.

- When using the privileged IOCTL of `RNDADDENTROPY` with /dev/random, the caller can inject entropic data into the auxiliary pool and define the amount of entropy associated with that data.

2.7.2 Auxiliary Pool

The auxiliary pool is maintained as a separate entropy source that eventually is concatenated with all other entropy sources in compliance with SP800-90C.

The auxiliary pool is processed with the available hash as follows:

1. Data is inserted the same way as data is added into the per-CPU entropy pools. The auxiliary pool technically is the message digest state where new data is inserted into the pool by performing a hash update operation.

2. When entropy is to be extracted from the auxiliary pool, a hash final operation is performed which is immediately followed by a hash init operation to initialize the hash context for new data.

3. The generated message digest is truncated to the amount of data requested by the DRNG (e.g. either 256 or 384 bits) and returned to the caller. Note, the auxiliary pool output is not truncated to the amount of entropy the data contains because the entropy provider may add data to the auxiliary pool without entropy, e.g. by simply writing to /dev/random. Thus, it may be possible that the auxiliary pool contains zero bits of entropy but yet contains data that should be used to "stir" the DRNG state.

Figure 2.8 illustrates the auxiliary pool operation for the case when user space inserts three separate buffers.

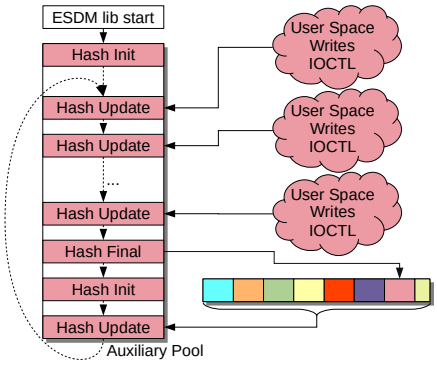

Figure 2.8: Auxiliary Pool Processing

In addition, the ESDM maintains an entropy estimator for the auxiliary pool counting the received entropy. The entropy estimator is capped to a maximum of the digest size of the used hash as this hash cannot maintain more entropy.

The auxiliary pool message digest is copied into the seed buffer when generating random numbers to seed the DRNG. The entire seed buffer is mixed back into the auxiliary pool for backward secrecy as shown in figure 2.8. The copy operation as well as the backtracking operation is atomic with respect to the auxiliary pool. This implies that both operations will always be fully completed before the next operation can commence. This ensures that the same auxiliary pool state can only be used once for a given seeding operation. Thus, both, the entropy pool and the auxiliary pool, are simultaneously used as noise data provider to seed the DRNG.

The entropy estimator is decreased by the amount of data read via the message digest.

2.8 Jitter RNG - ESDM-external Entropy Source

The Jitter RNG is treated as an external entropy source which is requested for random bits at the time the DRNG shall be seeded.

2.8.1 Entropy of CPU Jitter RNG Entropy Source

The CPU Jitter RNG entropy source is assumed provide 16th bit of entropy per generated data bit. Albeit studies have shown that significant more entropy is provided by this entropy source, a conservative estimate is applied.

The entropy value can be altered by changing the meson configuration option of `es_jent_entropy_rate`.

2.9 CPU-base Entropy Source - ESDM-external Entropy Source

The CPU-based entropy source is treated as an external entropy source which is requested for random bits at the time the DRNG shall be seeded. Depending on the underlying CPU, only one such source is available like RDSEED on Intel x86 (or RDRAND if RDSEED is not available), the POWER CPU DARN instruction, etc.

Depending whether the the CPU entropy source is documented to full entropy, the following data collection methods are applied. This approach is orthogonal to the amount of entropy the ESDM awards to the CPU entropy source.

- CPU entropy source provides full entropy: The CPU entropy source is queried for the amount of data which is stored in the temporary seed buffer.
- CPU entropy source provides less than full entropy: For this entropy source, the ESDM contains the information about how much data must be fetched from the CPU entropy source to get full entropy. The required amount of data is pulled from the CPU entropy source and conditioned with the hash currently in use by the ESDM. The calculated message digest is truncated to the requested amount of data which is stored in the temporary seed buffer.

2.9.1 Entropy of CPU Entropy Source

The entropy source of the CPU is assumed to have one 32th of the generated data size – 8 bits of entropy. The reason for that conservative estimate is that the design and implementation of those entropy sources is not commonly known and reviewable. The entropy value can be altered by changing the meson configuration option of `es_cpu_entropy_rate`.

2.10 Kernel RNG Entropy Source - ESDM-external Entropy Source

The ESDM offers the use of the kernel RNG as an entropy source. The kernel RNG derives its entropy from sampling of interrupts.

The legacy RNG is queried for random numbers using the `getrandom(2)` system call.

2.10.1 Entropy of Kernel RNG Entropy Source

When the lernel RNG is enabled, the ESDM applies the entropy rate defined at compile time with `es_kernel_entropy_rate` which is a value between 0 and 256 bits of entropy when 256 data bits are pulled from the legacy RNG.

If the ESDM is operated in FIPS mode, i.e. the kernel command line contains "fips=1", the kernel RNG entropy source's entropy rate is set to zero. The reason is that the legacy RNG is known to not comply with FIPS 140 rules like SP800-90B and thus must be assumed to provide no entropy.

When the interrupt entropy source is present, the kernel entropy source entropy rate is capped to `ESDM_ES_IRQ_MAX_KERNEL_RNG_ENTROPY` bits of entropy. This is due to the fact that the presence of IRQ entropy source takes away the kernel RNG's main entropy source. Yet, some other entropy sources are present which may or may not deliver entropy. Thus a safe entropy value is applied by the ESDM.

2.11 DRNG Seeding Operation

The seeding operation obtains random data from all available entropy sources.

The (re)seeding logic tries to obtain 256 bits of entropy from the entropy sources. However, if less entropy can only be delivered, the DRNG reseeding is only performed if at least 128 bits of entropy collectively from all entropy sources can be obtained.

For efficiency reasons, the seeding operation uses a seed buffer depicted in figure 2.1 that is the following set of blocks of 256 bits each. If SP800-90C compliance is enabled, the initial seeding of the DRNG is seeded with a seed buffer that pulls 384 bits in the following blocks. Each block is dedicated to an entropy source. As each entropy source can be disabled at compile time, the different bullets in the following list only applies if the corresponding entropy source is enabled.

1. One block is filled with the message digest from the auxiliary pool.
2. One block contains the message digest calculated from all per-CPU interrupt entropy pools as depicted in figure 2.4 That buffer receives as much data from the hash operation as entropy can be pulled from the entropy pools. In the worst case when no new interrupts are received a zero buffer will be injected into the DRNG. This is performed by iterating over all per-CPU entropy pools and:
   (a) Perform a hash update operation to inject the current content of the per-CPU collection pool into the per-CPU entropy pool.
   (b) Perform a hash final operation on the per-CPU entropy pool to obtain the message digest.
   (c) That message digest is used to re-initialize the per-CPU entropy pool with a hash init and hash update operation to ensure backward secrecy.
   (d) Also, the message digest is fed into the hash operation to collect the output from all entropy pools.

   Once the message digest from all is obtained, it is truncated to the amount of entropy present in all entropy pools.
3. Another block is filled with the message digest of all scheduler per-CPU entropy pools with the same approach as outlined for the interrupt ES.
4. One block is filled with the data from the kernel RNG entropy source.
5. The next block is filled by the Jitter RNG entropy source.
6. Finally, a block is filled by the fast entropy source of the CPU entropy source.

Finally, also a 32 bit time stamp indicating the time of the request is mixed into the DRNG. That time stamp, however, is not assumed to have entropy and is only there to further stir the state of the DRNG.

The filled seed buffer is handed to the DRNG as a seed string. In addition, the seed buffer is inserted back into the auxiliary pool for backward secrecy. The seed buffer will not alter the entropy estimation of the auxiliary pool.

2.11.1 DRNG May Become Not Fully Seeded

The ESDM maintains a counter for each DRNG instance how many generate operation are performed without performing a reseed that has full entropy. If this counter exceeds the threshold of 2^{30} generate operations, i.e. the DRNG did not receive a seed with full entropy for that many generate operations, the DRNG is set to not fully seeded. This setting implies that the DRNG instance will not be used any more for generating random numbers until the ESDM received sufficient entropy to reseed the DRNG with full entropy.

If the DRNG that becomes not fully seeded is the initial DRNG instance that was seeded during boot time as outlined in section 2.3.1, the entire ESDM is marked as not operational. This setting blocks all blocking interfaces just like during boot time when the ESDM is not yet fully seeded.

The ESDM automatically tries to recover from it when it received sufficient entropy.

2.12 Cryptographic Primitives Used By ESDM

The following subsections explain the cryptographic primitives that may be used by the ESDM.

2.12.1 DRBG

If the SP800-90A DRBG implementation is used, the default DRBG used by the ESDM is the CTR DRBG with AES-256. The reason for the choice of a CTR DRBG is its speed. The source code allows the use of other types of DRBG by simply defining a DRBG reference using the kernel crypto API DRBG string – see the top part of the source code for examples covering all types of DRBG.

All DRNGs are always instantiated with the same DRNG type.

The implementation of the DRBG is taken from the Linux kernel crypto API. The use of the kernel crypto API to provide the cipher primitives allows using assembler or even hardware-accelerator backed cipher primitives. Such support should relieve the CPU from processing the cryptographic operation as much as possible.

The input with the seed and re-seed of the DRBG has been explained above and does not need to be re-iterated here. Mathematically speaking, the seed and re-seed data obtained from the entropy sources and the ESDM external sources are mixed into the DRBG using the DRBG "update" function as defined by SP800-90A.

The DRBG generates output with the DRBG "generate" function that is specified in SP800-90A. The DRBG used to generate two types of output that are discussed in the following subsections.

/dev/urandom and `esdm_get_random_bytes` Users that want to obtain data via the /dev/urandom interface or the `esdm_get_random_bytes` API are delivered data that is obtained from the DRNG "generate" function. I.e. the DRNG generates the requested random numbers on demand.

Data requests on either interface is segmented into blocks of maximum 4096 bytes. For each block, the DRNG "generate" function is invoked individually. According to SP800-90A, the maximum numbers of bytes per DRBG "generate"

request is 2^{19} bits or 2^{16} bytes which is significantly more than enforced by the ESDM.

In addition to the slicing of the requests into blocks, the ESDM maintains a counter for the number of DRNG "generate" requests since the last reseed. According to SP800-90A, the number of allowed requests before a forceful reseed is 2^{48} – a number that is very high. The ESDM uses a much more conservative threshold of 2^{20}requests as a maximum. When that threshold is reached, the DRBG will be reseeded by using the operation documented in section 2.11 before the next DRNG "generate" operation commences.

The handling of the reseed threshold as well as the capping of the amount of random numbers generated with one DRNG "generate" operation ensures that the DRNG is operated compliant to all constraints in SP800-90A.

/dev/random and `esdm_get_random_bytes_full` The random numbers to be generated for /dev/random as well as `esdm_get_random_bytes_full` are defined to have a special property: it only provides data once at least 256 bits of entropy have been collected by the ESDM. In addition, the ESDM must be fully initialized before random numbers are generated, including the completion of the SP800-90B heath test if entropy from internal entropy sources is gathered.

2.12.2 ChaCha20 DRNG

If the SP800-90A DRBG is not desired, the ESDM can use a standalone C implementations for ChaCha20 to provide a DRNG.

The ChaCha20 DRNG is implemented with the components discussed in the following section. All of those components rest on a state defined by [1], section 2.3.

The operation of the ChaCha20 DRNG can be characterized with figure 2.9. This figure outlines the initialization of the DRNG, its seeding using the state update operation and the invocation of one generate operation that is requested to obtain more than 512 bits of data.

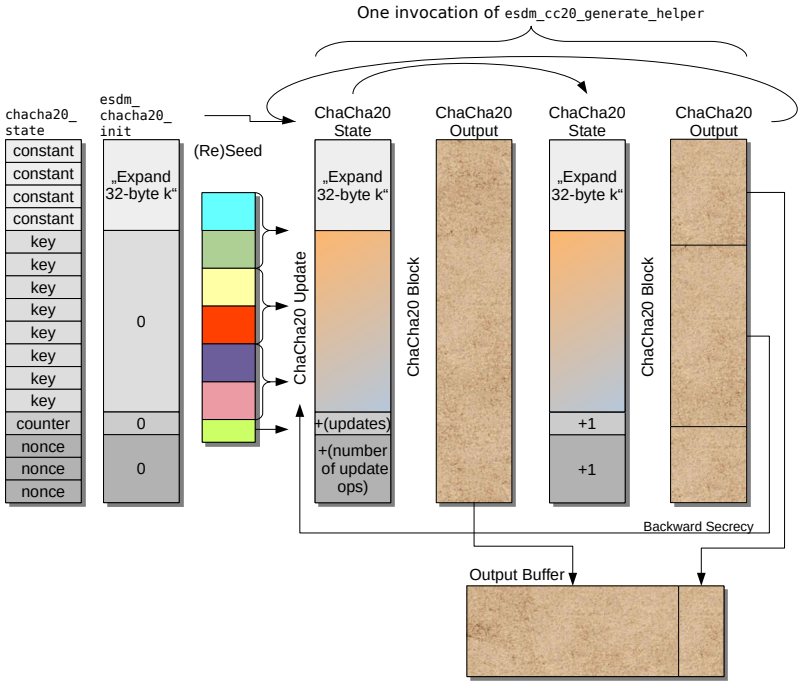

Figure 2.9: ChaCha20 DRNG Operation

State Update Function The state update function's purpose is to update the state of the ChaCha20 DRNG. That is achieved by

1. generating one output block of ChaCha20,
2. partition the generated ChaCha20 block into two key-sized chunks,
3. and XOR both chunks with the key part of the ChaCha20 state.

In addition, the nonce part of the state is incremented by one to ensure the uniqueness requirement of [1] chapter 4.

Seeding Operation The seeding operation processes a seed of arbitrary lengths. The seed is segmented into ChaCha20 key size chunks which are sequentially processed by the following steps:

1. The key-size seed chunk is XORed into the ChaCha20 key location of the state.
2. This operation is followed by invoking the state update function.
3. Repeat the previous steps for all unprocessed key-sized seed chunks.

If the last seed chunk is smaller than the ChaCha20 key size, only the available bytes of the seed are XORed into the key location. This is logically equivalent to padding the right side of the seed with zeroes until that block is equal in size to the ChaCha20 key.

The invocation of the state update function is intended to eliminate any potentially existing dependencies between the seed chunks.

Generate Operation The random numbers from the ChaCha20 DRNG are the data stream produced by ChaCha20, i.e. without the final XOR of the data stream with plaintext. Thus, the DRNG generate function simply invokes the ChaCha20 to produce the data stream as often as needed to produce the requested number of random bytes.

After the conclusion of the generate operation, the state update function is invoked to ensure enhanced backward secrecy of the ChaCha20 state that was used to generate the random numbers.

2.13 ESDM External Interfaces

The following ESDM interfaces are provided:

- ESDM library API: The API is documented in 'esdm.h'.
- ESDM server RPC API: The RPC API is documented in 'esdm_rpc_client.h'.
- /dev/random and /dev/urandom device files: See the Linux man page `random(4)`. In addition, when opening the /dev/random device with with the O_SYNC flag, the DRNG is managed with prediction resistance enabled. When using /dev/random with `O_SYNC`, the ESDM API call of `esdm_get_entropy_bytes_pr` is used to service the request. Therefore all considerations outlined for this API apply to /dev/random and `O_SYNC` as well.
- /proc/sys/kernel/random/* interfaces: See the Linux man page random(4).
- getrandom and getentropy APIs: See Linux man pages `getrandom(2)` and `getentropy(3)`. When `getrandom(2)` is invoked with the flag of `GRND_RANDOM`, the ESDM API call of `esdm_get_entropy_bytes_pr` is used to service the request. In addition when `getrandom(2)` is invoked with `GRND_SEED`, the ESDM API call of `esdm_get_seed` is used to fulfill the request.

2.14 ESDM Self-Tests

All cryptographic primitives are self-tested during the startup phase of the ESDM.

2.15 ESDM Test Interfaces

During kernel compilation, the following interfaces may be enabled allowing direct access to non-deterministic aspects. It is not advisable to enable these interfaces for production systems. Yet, these interfaces are considered to be protected against misuse by allowing only the root user to access them. In addition, any data obtained through these interfaces is not used by the ESDM to feed the entropy pool. Thus, even when leaving these interfaces enabled on production systems, the impact on security is considered to be limited.

- Interrupt Entropy Source:

- The interface `/sys/kernel/debug/esdm_testing/esdm_raw_hires` allows reading of the raw unconditioned noise data collected while the read operation is in progress by providing the time stamps of the events collected by the ESDM that otherwise are injected into the entropy pool. When booting the kernel with the kernel command line option `esdm_testing.boot_raw_hires_test=1`, the time stamps of the first 1,024 events recorded by the ESDM are stored. The first read of the `esdm_raw_hires` file after boot provides this data in this case.
- The interface `/sys/kernel/debug/esdm_testing/esdm_raw_jiffies` allows reading of the raw unconditioned Jiffies collected while the read operation is in progress by providing the Jiffies values collected by the ESDM that otherwise are injected into the entropy pool (if no high-resolution time stamp is detected). When booting the kernel with the kernel command line option `esdm_testing.boot_raw_jiffies_test=1`, the time stamps of the first 1,024 events recorded by the ESDM are stored. The first read of the `esdm_raw_jiffies` file after boot provides this data in this case.
- The interface `/sys/kernel/debug/esdm_testing/esdm_raw_irq` allows reading of the raw unconditioned interrupt numbers collected while the read operation is in progress by providing the interrupt number values collected by the ESDM that otherwise are injected into the entropy pool (if no high-resolution time stamp is detected) or into the random32 PRNG external to the ESDM. When booting the kernel with the kernel command line option `esdm_testing.boot_raw_irq_test=1`, the time stamps of the first 1,024 events recorded by the ESDM are stored. The first read of the `esdm_raw_irq` file after boot provides this data in this case.
- The interface `/sys/kernel/debug/esdm_testing/esdm_raw_irqflags` allows reading of the raw unconditioned interrupt flags collected while the read operation is in progress by providing the interrupt flag values collected by the ESDM that otherwise are injected into the entropy pool (if no high-resolution time stamp is detected) or into the random32 PRNG external to the ESDM. When booting the kernel with the kernel command line option `esdm_testing.boot_raw_irqflag_test=1`, the time stamps of the first 1,024 events recorded by the ESDM are stored. The first read of the `esdm_raw_irqflags` file after boot provides this data in this case.
- The interface `/sys/kernel/debug/esdm_testing/esdm_raw_retip` allows reading of the raw unconditioned return instruction pointer collected while the read operation is in progress by providing the instruction pointer 32 LSB values collected by the ESDM that otherwise are injected into the entropy pool (if no high-resolution time stamp is detected) or into the random32 PRNG external to the ESDM. When booting the kernel with the kernel command line option `esdm_testing.boot_raw_retip_test=1`, the time stamps of the first 1,024 events recorded by the ESDM are stored. The first read of the `esdm_raw_retip` file after boot provides this data in this case.

  - The interface `/sys/kernel/debug/esdm_testing/esdm_raw_regs` allows reading of the raw unconditioned interrupt register state collected while the read operation is in progress by providing the selected register 32 LSB values collected by the ESDM that otherwise are injected into the entropy pool (if no high-resolution time stamp is detected). When booting the kernel with the kernel command line option `esdm_testing.boot_raw_regs_test=1`, the time stamps of the first 1,024 events recorded by the ESDM are stored. The first read of the `esdm_raw_regs` file after boot provides this data in this case.
  - The interface `/sys/kernel/debug/esdm_testing/esdm_raw_array` allows reading of the raw noise data that has been stored in the per-CPU collection pool collected while the read operation is in progress. When booting the kernel with the kernel command line option `esdm_testing.boot_raw_array=1`, the array content of the first 1,024 events recorded by the ESDM are stored. The first read of the `esdm_raw_array` file after boot provides this data in this case.
  - The interface `/sys/kernel/debug/esdm_testing/esdm_irq_perf` allows reading of the number of cycles used to process one interrupt event. This allows measuring the performance impact of the ESDM on the interrupt handler. When booting the kernel with the kernel command line option `esdm_testing.boot_irq_perf=1`, the performance data of the first 1,024 events recorded by the ESDM are stored. The first read of the `esdm_irq_perf` file after boot provides this data in this case.

- Scheduler Entropy Source:
  - The interface `/sys/kernel/debug/esdm_testing/esdm_raw_sched_hires` allows reading of the raw unconditioned noise data collected while the read operation is in progress by providing the time stamps of the events collected by the ESDM that otherwise are injected into the entropy pool. When booting the kernel with the kernel command line option `esdm_testing.boot_raw_sched_hires_test=1`, the time stamps of the first 1,024 events recorded by the ESDM are stored. The first read of the `esdm_raw_sched_hires` file after boot provides this data in this case.
  - The interface `/sys/kernel/debug/esdm_testing/esdm_raw_sched_pid` allows reading of the PID value of the task that are about to be scheduled to collected while the read operation is in progress. The PID values are extracted which are collected by the ESDM that otherwise are injected into the entropy pool (if no high-resolution time stamp is detected). When booting the kernel with the kernel command line option `esdm_testing.boot_raw_sched_pid_test=1`, the time stamps of the first 1,024 events recorded by the ESDM are stored. The first read of the `esdm_raw_sched_pid` file after boot provides this data in this case.
  - The interface `/sys/kernel/debug/esdm_testing/esdm_raw_starttime_pid` allows reading of the start time value of the task that are about to be

scheduled to collected while the read operation is in progress. The start time values are extracted which are collected by the ESDM that otherwise are injected into the entropy pool (if no high-resolution time stamp is detected). When booting the kernel with the kernel command line option `esdm_testing.boot_raw_sched_starttime_test=1`, the time stamps of the first 1,024 events recorded by the ESDM are stored. The first read of the `esdm_raw_starttime_pid` file after boot provides this data in this case.

- The interface `/sys/kernel/debug/esdm_testing/esdm_raw_nvcsw_pid` allows reading of the numbers of context switches of the task that are about to be scheduled to collected while the read operation is in progress. The context switch values are extracted which are collected by the ESDM that otherwise are injected into the entropy pool (if no high-resolution time stamp is detected). When booting the kernel with the kernel command line option `esdm_testing.boot_raw_sched_nvcsw_test=1`, the time stamps of the first 1,024 events recorded by the ESDM are stored. The first read of the `esdm_raw_nvcsw_pid` file after boot provides this data in this case.
- The interface `/sys/kernel/debug/esdm_testing/esdm_sched_perf` allows reading of the number of cycles used to process one scheduler event. This allows measuring the performance impact of the ESDM on the scheulder. When booting the kernel with the kernel command line option `esdm_testing.boot_sched_perf=1`, the performance data of the first 1,024 events recorded by the ESDM are stored. The first read of the `esdm_sched_perf` file after boot provides this data in this case.

The helper tool getrawentropy.c is provided to read the files and format the data for post-processing.

3 Interrupt Entropy Source Assessment

This section documents the entropy assessment and the compliance with various standards of the interrupt entropy source. All other entropy sources are treated as black boxes by the ESDM and must therefore deliver their own entropy assessment.

The entropy sources of the auxiliary pool, the Jitter RNG and the CPU-based entropy source are all not related to the interrupt entropy source. Thus, the combination of all those entropy sources is subject to the SP800-90C assessment provided in section 5.1

3.1 Noise Source Behavior

The noise source component of an entropy source contains the non-deterministic, entropy-producing activity. For the scheduler entropy source this is the timing of the arrival of interrupts. The general behavior of the noise source can be characterized by analyzing the time stamps of the arrival time. Before performing this analysis, a recap of the noise source is important:

- When a scheduling event occurs, the ESDM obtains a high-resolution time stamp. Technically it may be a cycle counter offered by the CPU that produces the time stamp which is unrelated to time, but provides a fast-moving counter.
- The time stamp is divided by the greatest common divisor (GCD) that is calculated during boot time by using the first 100 time stamps.
- After the division with the GCD, the 8 least-significant bits of the time stamp are used. The higher-order bits are discarded. These 8 bits are now subject to post-processing via the collection pool, the entropy pool and the DRNG.

Considering these static and never-changing steps, the noise source rests on the gathering of the 8 bits from the last step. An analysis of the noise source therefore focuses on these 8 bits. These 8 bits are subsequently referenced as "raw noise source data".

3.1.1 Distribution of Raw Data

If the noise source would operate perfectly, an equi-distribution of the raw noise source data is to be expected. Various tests have been conducted on most major CPUs as outlined in appendix C. Using the runtime data, a common pattern of the raw noise source data emerges that can be seen with figure 3.1. This figure uses the 1,000,000 traces of the raw noise source data from the IRQ entropy source on the RISC-V, i.e. a system with less-then-average entropy rate.

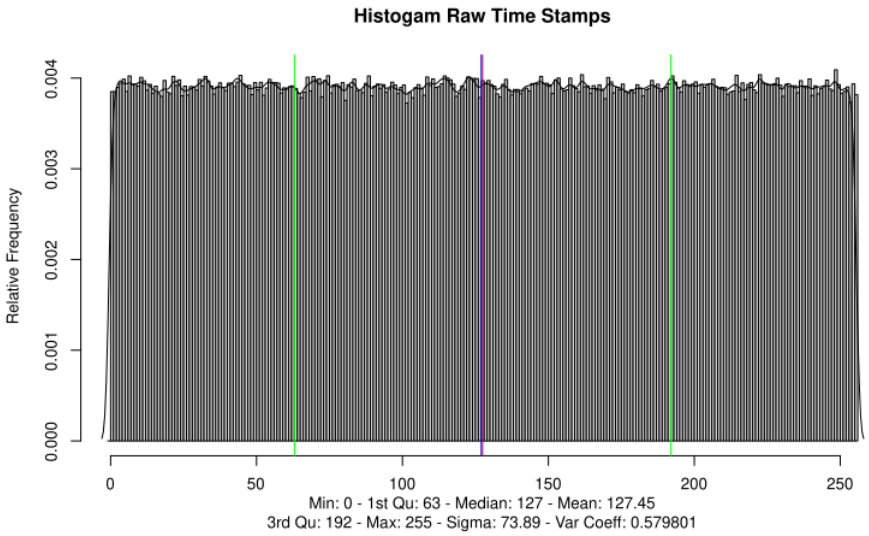

Figure 3.1: RISC-V Raw Noise Source Data Distribution

Figure 3.1 shows the following characteristics:

- A histogram of the number of time stamps received for each of the 256 possible time stamp values is depicted.

- The 25% and 75% quartiles of the distribution are marked with the two green lines. They are at the time stamp value of 63 and 192, respectively
- The mean of the distribution is marked with the red line which is 127.45.
- The median of the distribution is marked with the blue line at 128.
- The black line marks the distribution of the time stamps.
- The standard derivation is not depicted, but it is at 73.89.
- The variation coefficient is not shown in the figure, but it is 0.579801.

The figure together with the mentioned statistical values clearly shows that it is close to an equi-distribution. This means that the raw noise source data for an entropy source instance on a less-than-ideal hardware still exhibit a distribution that is close to the expected entropy source for a perfect operation. The distributions of most tested systems are always close to an equi-distribution as shown with the table below.

To verify that the distribution is an equi-distribution, a Chi-Squared Goodness-of-Fit test is applied. For the mentioned RISC-V system, the following values are observed:

- Asymptotic significance P: 0.2343
- Degrees of freedom: 255

The degrees of freedom shows that indeed all 256 possible timing values are covered. Applying an alpha of 5%, the Chi-Squared test indicates that the observed data set is an equi-distribution.

With the table below, the statistical properties for the different tested systems are listed.

Test System	25% Quartile	Median	Mean	75% Quartile	Std. Deriv.	Var. Coeff.	χ^2 P	χ^2 DF
AMD Ryzen 5950X - 64-bit KVM environment	64	128	127.6	192	73.88	0.58	0.0523	255
AMD EPYC Milan 7713 2 sockets 128 cores 8-way NUMA	63	127	127.48	192	73.91	0.58	0.4458	255
ARMv7 rev 5	68	128	121.05	191	72.95	0.60	0	255
ARMv7 rev 5 (Freescale i.MX53)[4]	63	128	127.52	192	74	0.58	0.4693	255

[4]USBArmory MK I

Test System	25% Quartile	Median	Mean	75% Quartile	Std. Deriv.	Var. Coeff.	χ^2 P	χ^2 DF
ARMv7 rev 5 (Freescale i.MX6 Ultralite)[5]	63	127	127.15	191	73.96	0.58	0	255
ARM 64 bit AppliedMicro X-Gene Mustang Board	63	128	127.54	192	73.94	0.58	0.7229	255
Intel Atom Z530 – using GUI	64	128	127.58	192	73.93	0.58	0.3181	255
Intel Sandy Bridge Clang Compile	62	125	125.86	187	73.24	0.58	0.9347	252
Intel i7 8565U Whiskey Lake – 32-bit KVM environment	64	128	127.53	192	73.91	0.58	0.2893	255
Intel i7 8565U Whiskey Lake	63	127	127.37	191	73.86	0.58	0	255
Intel Xeon E7 4870 8 sockets 160 CPUs 8-way NUMA	64	127	127.44	191	73.89	0.58	0	255
Intel Xeon Gold 6234	64	126	126.75	190	73.88	0.58	0	127[6]
IBM POWER 8 LE 8286-42A	63	127	127.66	191	73.86	0.58	0	255
IBM POWER 7 BE 8202-E4C	64	128	130.07	192	73.96	0.57	0	255
IBM System Z z13 (machine 2964)	56	120	118.27	184	73.58	0.62	0	235
IBM System Z z15 (machine 8561)	58	122	124.3	192	73.54	0.59	0	239
MIPS Atheros AR7241 rev 1[7]	64	128	127.55	191	73.88	0.58	0.5646	255
MIPS Lantiq 34Kc V5.6[8]	64	127	127.49	191	73.85	0.58	0	255

[5]USBArmory MK II
[6]Tested without GCD.
[7]Ubiquiti Nanostation M5 (xm)
[8]AVM Fritz Box 7490

Test System	25% Quartile	Median	Mean	75% Quartile	Std. Deriv.	Var. Coeff.	χ^2 P	χ^2 DF
Qualcomm IPQ4019 ARMv7[9]	63	127	127.4	191	73.9	0.58	0	255
SiFive HiFive Unmatched RISC-V U74	63	127	127.45	192	73.89	0.58	0.2343	255

The Chi-Squared asymptotic significance shows for some systems that an equi-distribution is not applicable, i.e. when the χ^2 P value is less than 0.05. Yet, when considering the other statistical values, the actual distribution is neither skewed nor otherwise loop-sided. When looking at the distribution, it becomes evident that some time stamps have a higher likelihood than others which hint to special properties of the system. This observation applies to all measurements which do not follow an equi-distribution based on the Chi-Squared Goodness-of-Fit test. This allows the conclusion that the min-entropy estimates given in appendix C can be considered to illustrate the real entropy rate.

For example, the distribution for the USB Armory Mark II system shown with figure 3.2 indicates that all time stamp values divisible by 8 are only chosen two-thirds as often as the rest.

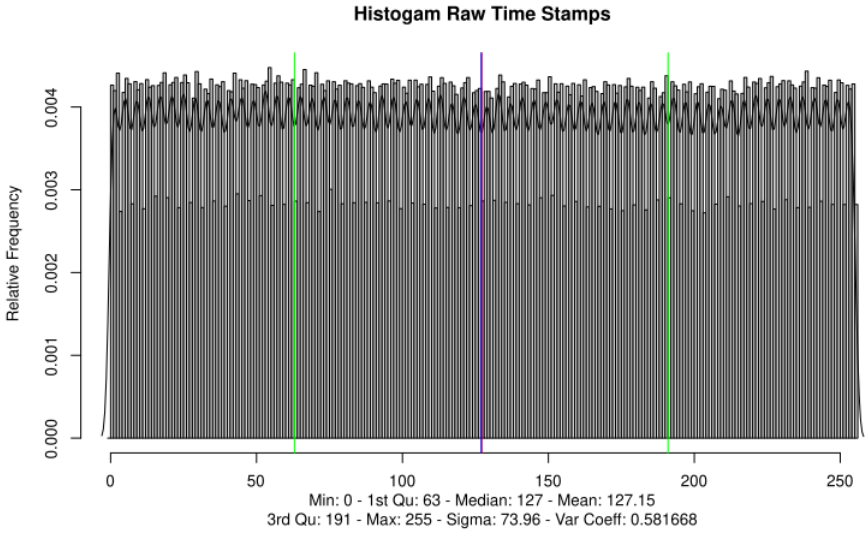

Figure 3.2: USB Armory Mark II: Raw Noise Source Data

Another example is a specific ARMv7 system which contains a very periodic timer interrupt. Figure 3.3 shows that the time stamp has a pattern but still exhibits a distribution that does not contradict the expectation to deliver entropy

[9] AVM Fritz Box 7520

at the rate calculated in appendix C. In the worst case that the timer interrupt only causes the raw noise source data which would exhibit a clear pattern, the stuck health test would identify this pattern with the second discrete derivative and disregards time stamp for entropy collection.

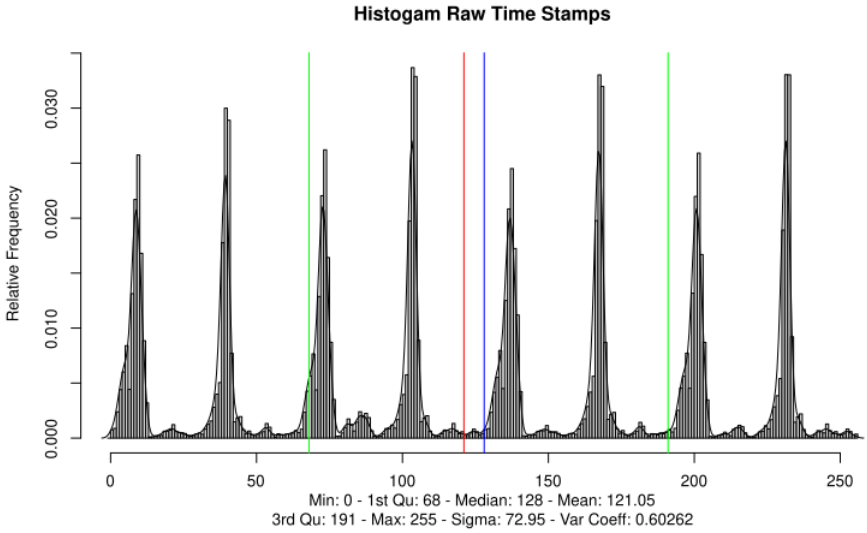

Figure 3.3: Periodic timer interrupt: Specific ARMv7 System Raw Noise Source Data

A similar effect is visible on an IBM System Z z13 system shown with figure 3.4. Again, the shown pattern does not contradict the entropy rate calculated for this system in appendix C.

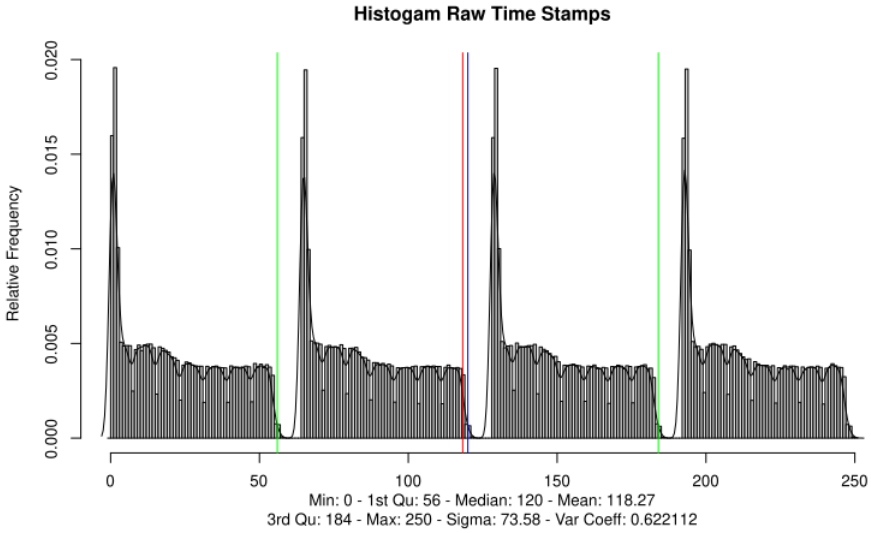

Figure 3.4: IBM System Z Raw Noise Source Data

3.1.2 Greatest Common Divisor Assessment

The behavior of the GCD application can be clearly seen with the following figures and numbers obtained for an Intel Atom Z530 system whose GCD is 4.

Without the application of the GCD, the distribution of the time stamp is given with figure 3.5

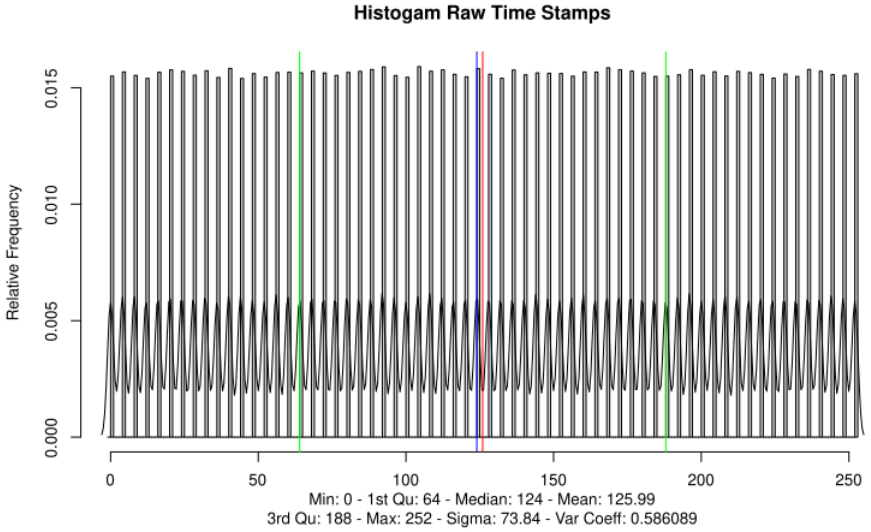

Figure 3.5: Without GCD - Raw Noise Source Data Distribution

Figure 3.5 shows the following characteristics:

- A histogram of the number of time stamps received for each of the 256 possible time stamp values is depicted.
- The 25% and 75% quartiles of the distribution are marked with the two green lines. They are at the time stamp value of 63 and 188, respectively
- The mean of the distribution is marked with the red line which is 125.99.
- The median of the distribution is marked with the blue line at 128.
- The standard derivation is not depicted, but it is at 73.84.
- The variation coefficient is not shown in the figure, but it is 0.586089.
- χ^2 P value for equi-distribution Goodness-of-Fit test: 0.4263
- χ^2 degrees of freedom for equi-distribution Goodness-of-Fit test: 63

When applying the GCD, and obtaining a new measurement, the distribution shown with figure 3.6 emerges:

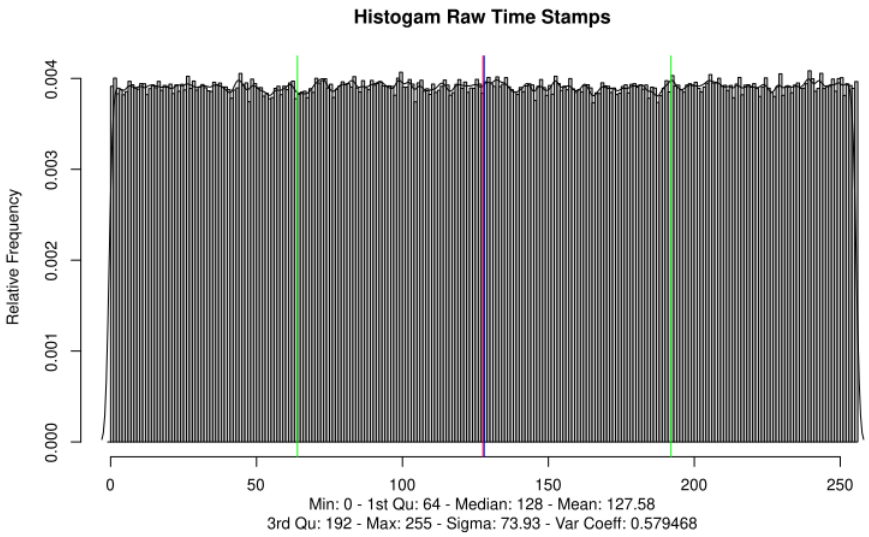

Figure 3.6: With GCD - Raw Noise Source Data Distribution

Figure 3.6 shows the following characteristics:

- A histogram of the number of time stamps received for each of the 256 possible time stamp values is depicted.
- The 25% and 75% quartiles of the distribution are marked with the two green lines. They are at the time stamp value of 64 and 192, respectively
- The mean of the distribution is marked with the red line which is 127.58.
- The median of the distribution is marked with the blue line at 124.

- The standard derivation is not depicted, but it is at 73.93.
- The variation coefficient is not shown in the figure, but it is 0.579468.
- χ^2 P value for equi-distribution Goodness-of-Fit test: 0.3181
- χ^2 degrees of freedom for equi-distribution Goodness-of-Fit test: 255

This shows that the application of the GCD removes the "unused" time stamp values without changing the overall distribution.

3.1.3 Worst and Regular Case Distribution

The measurements shown in appendix C commonly are obtained by applying a worst-case which triggers as much interrupts as possible in the shortest amount of time.

To compare the distributions of the time stamp between the worst case and a regular case of normal system use, the USB Armory mark I system was tested twice with the following result values:

- Worst case
  - 25% quartile: 64
  - Median: 128
  - Mean: 127.6
  - 75% quartile: 192
  - Standard derivation: 73.91
  - Variation coefficient: 0.579221
  - χ^2 P value for equi-distribution Goodness-of-Fit test: 0.9445
  - χ^2 degrees of freedom for equi-distribution Goodness-of-Fit test: 255
- Regular case
  - 25% quartile: 63
  - Median: 128
  - Mean: 127.52
  - 75% quartile: 192
  - Standard derivation: 74
  - Variation coefficient: 0.580276
  - χ^2 P value for equi-distribution Goodness-of-Fit test: 0.4693
  - χ^2 degrees of freedom for equi-distribution Goodness-of-Fit test: 255

The values indicate no statistical significant difference allowing the conclusion that both distributions are very similar.

3.2 FIPS 140-2 Compliance

FIPS 140-2 specifies entropy source compliance in FIPS 140-2 IG 7.18. This section analyzes each requirement for compliance. The general requirement to comply with SP800-90B [2] is analyzed in section 3.3.

3.2.1 FIPS 140-2 IG 7.18 Requirement For Statistical Testing

The ESDM is provided with the following testing tools:

- Raw Entropy Tests: The tests obtain the raw unconditioned and unprocessed noise information and records it for analysis with the SP800-90B non-IID statistical test tool. The test tool includes the gathering of raw entropy for one execution run as well as for the restart tests required in SP800-90B section 3.1.4. The tool adjusts the data to be processed by the SP800-90B statistical test tool. The test tool provides the SP800-90B minimum entropy values.

In particular the first test covers the test requirement of FIPS 140-2 IG 7.18.

3.2.2 FIPS 140-2 IG 7.18 Heuristic Analysis

FIPS 140-2 IG 7.18 requires a heuristic analysis compliant to SP800-90B section 3.2.2. The discussion of this SP800-90B requirement list is given in section 3.3.

3.2.3 FIPS 140-2 IG 7.18 Additional Comment 1

The first test referenced in section 3.2.1 covers this requirement.

The test collects the time stamps of interrupts as they are received by the ESDM. Instead of having these interrupts processed by the ESDM to add them to the entropy pool, they are sent to a user space application for storing them to disk.

The collection of the interrupt data for the raw entropy testing is invoked from the same code path that would otherwise add it to the ESDM entropy pool. Thus, the test collects the exact same data that would otherwise have been used by the ESDM as noise data. Thus, the testing does not alter the ESDM processing.

However, the tester performing the test should observe the following caveat: the raw entropy data obtained with the user space tool should be stored on "disk space" that will not generate interrupts as otherwise the testing would itself generate new interrupts and thus alter the measurement. For example, a ramdisk can be used to store the raw entropy data while the test is ongoing. On common Linux environments, the path `/dev/shm` is usually a ramdisk that can readily be used as a target for storing the raw entropy data. If that partition is non-existent, the tester should mount a ramdisk or use different backing store that is guaranteed to not generate any interrupts when writing data to it.

3.2.4 FIPS 140-2 IG 7.18 Additional Comment 2

The lowest entropy yield is analyzed by gathering raw entropy data received from interrupts that come in high frequency. In this case, the time stamps would be close together where the variations and thus the entropy provided with these time stamps would be limited.

The extreme case would be to send a flood of ICMP echo request messages with a size of only one byte to the system under test from a neighboring system that has a close proximity with very little network latency. Each ICMP request would trigger an interrupt as it is processed by the network card. The most extreme case can be achieved when executing the ESDM in a virtual machine

where the VMM host sends a ping flood to the virtual machine. In this case, network latency would be reduced to a minimum. In the subsequent sections, test results are shown which are generated with an ESDM executing in a virtual machine where the host sends a flood of ICMP echo request messages to trigger a worst case measurement.

The entropy is not considered to degrade when using the hardware within the environmental constraints documented for the used CPU. The online health tests are intended to detect entropy source degradation. In case of online health test failures, section 2.6.2 explains the applied actions.

3.2.5 FIPS 140-2 IG 7.18 Additional Comment 3

The ESDM uses the following conditioning components:

- For collecting of entropy data from the entropy source, an approved message digest operation is used.
- For reading the entropy pool and compressing the entropy data, the hash operation is used. The security strength of the ESDM is the minimum of the DRBG security strength and the security strength of the hash following [2] section 3.1.5.1.1 table 1. All ciphers can be tested via ACVT.

3.2.6 FIPS 140-2 IG 7.18 Additional Comment 4

The restart test is covered by the first test documented in section 3.2.1.

3.2.7 FIPS 140-2 IG 7.18 Additional Comment 6

The entropy assessment usually shows this conclusion – tests performed on Intel x86-based systems show the following conclusions:

The entropy rate for all devices validated with the raw entropy tests outlined in section 3.2.1 show that the minimum entropy values are always above one bit of entropy per four data bits. The data bits are the least significant bits of the time stamp generated by the raw noise.

Assuming the worst case that all other bits in the time delta have no entropy, that entropy value above one bit of entropy applies to one time stamp.

The ESDM continuously gathers time stamps to be combined with a hash which is entropy preserving. The hash operation function providing data to the DRNG gathers only as much bits as time stamps were received. For example, if the ESDM only received 16 time stamps, the ESDM will only deliver 2 bytes of data to the DRNG. This effectively implies that the ESDM assumes that one bit of entropy is received per time stamp.

As the ESDM maintains an entropy pool, its entropy content cannot be larger than the pool itself. Thus, the entropy content in the pool after collecting as many time stamps as the entropy pool's size in bits is the maximum amount of entropy that can be held. Yet, as new time stamps are received, they are mixed into the entropy pool. In case the entropy pool is considered to have fully entropy, existing entropy is overwritten with new entropy.

This implies that the ESDM data generated from the entropy pool has (close to) 1 bit of entropy per data bit.

3.2.8 FIPS 140-2 IG 7.18 Additional Comment 9

N/A as the raw entropy is a non-IID source and processed with the non-IID SP800-90B statistical tests as documented in section 3.2.1.

3.3 SP800-90B Compliance

This chapter analyzes the compliance of the ESDM to the SP800-90B [2] standard considering the FIPS 140-2 implementation guidance 7.18 which alters some of the requirements mandated by SP800-90B.

3.3.1 SP800-90B Section 3.1.1

The collection of raw data for the SP800-90B entropy testing documented in section 3.2.1 uses 1,000,000 consecutive time stamps obtained in one execution round.

The restart tests documented in section 3.2.1 perform 1,000 restarts collecting 1,000 consecutive time stamps.

3.3.2 SP800-90B Section 3.1.2

The entropy assessment of the raw entropy data including the restart tests follows the non-IID track.

3.3.3 SP800-90B Section 3.1.3

Please see section 3.2.7: The entropy of the raw noise source data is believed to have more than one bit of entropy per time stamp to allow to conclude that one output block of the ESDM has (close to) one bit of entropy per data bit. Yet, this rate can be configured at compile time to be lower than one bit of entropy per interrupt event.

The first test referenced in section 3.2.1 performs the following operations to provide the SP800-90B minimum entropy estimate:

1. Gathering of the raw entropy data of the time stamps.

2. Obtaining the four least significant bits of each time stamp and concatenate them to form a bit stream.

3. The bit stream is processed with the SP800-90B entropy testing tool to gather the minimum entropy.

For example, on an Intel Core i7 Skylake system executing the ESDM in a KVM guest, the SP800-90B tool shows the following minimum entropy values when multiplying the SP800-90B tool bit-wise minimum entropy by four since four bits are processed: 3.452064.

3.3.4 SP800-90B Section 3.1.4

For the restart tests, the raw entropy data is collected for the first 1,000 interrupt events received by the ESDM after a reboot of the operating system. That means, for one collection of raw entropy the test system is rebooted. This

implies that for gathering the 1,000 restart samples, the test system is rebooted 1,000 times.

Each restart test round stores its time stamps in an individual file.

After all raw entropy data is gathered, a matrix is generated where each line in the matrix lists the time stamp of one restart test round. The first column of the matrix, for example, therefore contains the first time stamp for each boot cycle of the Linux kernel with the ESDM.

The SP800-90B minimum entropy values column and row-wise is calculated the same way as outlined above:

1. Gathering of the raw restart entropy data of the time stamps.

2. Obtaining the four least significant bits of each time stamp either row-wise or column-wise and concatenate them to form a bit stream. There are 1,000 bit streams row-wise, and 1,000 bit streams column-wise boundary generated.

3. The bit streams are processed with the SP800-90B entropy testing tool to gather the minimum entropy.

In a following step, the sanity check outlined in SP800-90B section 3.1.4.3 is applied to the restart test results. The steps given in 3.1.4.3 are applied.

For example, on an Intel Core i7 Skylake system executing the ESDM in a KVM guest, the SP800-90B tool shows the following minimum entropy values when multiplying the SP800-90B tool bit-wise minimum entropy by four since eight bits are processed:

- Using the 8 least significant bits of the time stamps in column-wise assessment – lowest entropy value of all 1,000 column entries: 3.455504

- Using the 8 least significant bits of the time stamps in row-wise assessment – lowest entropy value of all 1,000 column entries: 3.393808

- Sanity check of the 1,000 x 1,000 matrix passes with value of one

With the shown values, the restart test validation passes according to SP800-90B section 3.1.4.

3.3.5 SP800-90B Section 3.1.5

The conditioning component applied to the interrupt entropy source are performed at different stages as outlined in section 2.1. Although the hashing operation is used for different stages, the following discussion is applicable to all use cases.

Truncation The truncation operation ensures that the entropy in that data is at maximum the truncated hash.

The truncation of operation (1) listed in section 2.2 is not affected by the capping of the entropy, because the quantitative measurement of the existing entropy using the SP800-90B tool set is performed using that truncated input data. The ESDM implies an entropy of 1 bit per truncated time stamp and zero bits of entropy per arbitrary 32-bit word size which means that the entropy present in the data is always smaller as the data size.

The truncation operation of step (6) listed in section 2.2 verifies that the truncated data contains at most the amount of entropy as the generated data size. The remaining part of the truncated data is not exported to any external entity but remains in the per-CPU entropy pools - when new random data is generated involving the entropy pools, the current entropy pool states are always hashed. This is a deviation from SP800-90B section 3.1.5.1.2 which requires a relative reduction of entropy. This statement is considered inconsistent with the statement implied in table 1 [2] and therefore wrong depicted with the following analogy: Assume to have a buffer of 512 bits of data having 256 bits of entropy. When hashing it with SHA-512, the resulting message digest of 512 bits has 256 bits of entropy. When truncating the digest to 256 bits, SP800-90B states the entropy is 128 bits. However, SP800-90B section 3.1.5.1.1 table 1 states that full entropy is given to approved hash functions. Assume to use a SHA-512/256 which has a digest size of 256 bits and thus could transport 256 bits of entropy following table 1. This SHA-512/256 hash operation calculates a SHA-512 hash truncated to 256 bits. Albeit the cryptographic operation of SHA-512/256 is identical to the ESDM-applied truncation[10], SP800-90B table 1 awards 256 bits of entropy to SHA-512/256 but at the same time SP800-90B would apply only 128 bits to the ESDM-applied truncation. Due to this inconsistency, the ESDM applies the entropy behavior implicitly specified in table 1, i.e. the entropy is the minimum of the available entropy and the message digest size. Furthermore, applying the Output_Entropy formula for a vetted conditioning component of a truncated hash, the following calculation applies. This calculation shows the entropy rate of a SHA-512 hash processing a buffer with 1024 bits that contains, say, 384 bits of entropy and truncating it to 256 bits. This means, the formula for $h_{out_{SHA-512\ trunc}} = Output\_Entropy_{SHA-512\ trunc}(1024, 256, 512, 384)$ following [2] section 3.1.5.1.2 is calculated:

$$P_{high} = 2^{-384}$$

$$P_{low} = \frac{(1 - 2^{-384})}{2^{1024} - 1} \approx 2^{-1024}$$

$$n = min(256, 512) = 256$$

$$\psi = 2^{1024-256} \cdot 2^{-1024} + 2^{-384} = 2^{-256} + 2^{-384} \approx 2^{-256}$$
$$U = 2^{1024-256} + \sqrt{2 \cdot 256 \cdot (2^{1024-256}) \cdot ln(2)} = 2^{768} + \sqrt{2^{777} \cdot ln(2)} \approx 2^{768}$$
$$\omega = 2^{768} \times 2^{-1024} = 2^{-256}$$

$$\underline{Output\_Entropy_{SHA-512\ trunc}(1024, 256, 512, 384) = -log_2(max(2^{-256}, 2^{-256})) = 256}$$

Even after calculating other entropy rates using the same formula, the following conclusion for the truncation can be applied:

$$\underline{Output\_Entropy_{trunc}(n_{in}, n_{out}, nw, h_{in}) = min(n_{in}, n_{out}, nw, h_{in})}$$

This function is applied by the ESDM for hash truncation.

[10] Depending on the runtime configuration the ESDM uses a hash of SHA-512 and fills a buffer of the DRNG security strength size, i.e. 256 bits.

Concatenation When applying a concatenation operation, the ESDM simply adds the entropy delivered with each data entry part.

Hash The input of the hash n_{in} is fixed as it processes the existing per-CPU entropy pool(s), auxiliary pool and the per-CPU collection pools.

The output of the hash n_{out} is usually fixed to the message digest size. The on exception is the output of the hash n_{out} to provide the seed to the DRNG: it is the minimum of either the digest size of the used hash or the amount of entropy available in the processed entropy pools based on the number of "unprocessed" time stamps held in the per-CPU entropy pools.

The following hashes are used for the hash function depending on the loaded DRNG:

- ChaCha20: SHA-256 in normal case, SHA-1 if kernel is not compiled with `CONFIG_CRYPTO`
- SP800-90A Hash DRBG: SHA-512
- SP800-90A HMAC DRBG: SHA-512
- SP800-90A CTR DRBG: SHA-512

In the following, the different hash operations specified in section 2.2 are applied as follows:

- Compression of entropy delivered from the interrupt entropy source when adding the entropy into the per-CPU entropy pools.
- Compression of entropy delivered from one or more ESDM-external entropy sources when adding the entropy into the auxiliary pool.
- Compression of entropy delivered by the different per-CPU entropy pools when "reading" the entropy pools.

The requirement of [2] section 3.1.6 states that when combining two or more noise sources using a vetted conditioning component, only one noise source is to be credited with entropy. This requirement is met as follows: according to [2] section 2.2 a noise source is the phenomenon delivering entropy. The noise source data is post-processed with conditioning components and health tests to form an entropy source. Based on this statement, the collection of the per-CPU entropy pools together form one entropy source that is compliant to SP800-90B.

Note, the reason for hashing the per-CPU entropy pools together with the auxiliary pool is to ensure backward secrecy when calculating the next round of random numbers used to fill the seed buffer used to seed the DRBG from the entropy sources sources.

Approach for Calculating Entropy Although the aforementioned sections explain that the input and output sizes may not be fixed, in regular operation they are quasi-fixed. In order to reseed a DRNG, 256 bits of entropy are to be generated from the noise source. Although the per-CPU collection pools receive interrupt time stamps continuously, only the entropy from 256 time stamps are required as illustrated below. Only when all per-CPU entropy pools have received too little interrupt time stamps to satisfy the 256 bit entropy

request, less output data is generated. This commonly happens during boot or at runtime when too much entropy is requested. Though, during boot time, the DRNG will receive a (re)seed with 256 bits of entropy before the ESDM is considered fully operational. Therefore, the prior boot-time (re)seed events with less entropy may even be disregarded for the entropy assessment.

With the given combination of the hash as outlined above, the following approach for the entropy calculation is taken for each of the data processing steps outlined in section 2.2:

- Function 2.2:
  - $n_{in_{per-CPU\ pool}}$ equals to 8,192 bits as the per-CPU entropy pool obtains its input data from the collection pool that has a size of 8,192 bits (1,024 * 8 bits)[11] .
  - $n_{out_{per-CPU\ pool}}$ is the message digest size in bits.
  - $nw_{per-CPU\ pool}$ is the message digest size in bits.
- Function 2.8:
  - $n_{in_{aux\ pool}}$ equals to 512 bits as the auxiliary pool pool size is 512 bits in size plus the provided input data.
  - $n_{out_{aux\ pool}}$ is the message digest size in bits.
  - $nw_{aux\ pool}$ is the message digest size in bits.
- Function 2.3:
  - $n_{in_{hash\ pools}}$ equals to $\sum_{a=0}^{max\ CPU} n_{out_{per-CPU\ pool_a}} + n_{out_{hash\ aux}}$
  - $n_{out_{hash\ pools}}$ is the message digest size in bits.
  - $nw_{hash\ pools}$ is the message digest size in bits.

3.3.6 SP800-90B Section 3.1.5.1

The hash operation is either SHA-512, SHA-256, or SHA-1 as outlined above is considered to be a vetted conditioning component. Thus the entropy rate of the hash output is calculated as follows using the aforementioned variables for the hash function. In addition, the following consideration applies:

- The entropy content of the input $h_{in_{per-CPU\ pool}}$: The input entropy of the hash used to process the per-CPU entropy pool is equal to the entropy provided by the per-CPU collection pool and the entropy already present in the per-CPU entropy pool considering that both data components are hashed at the same time to form a new per-CPU entropy pool state. Of course, the entropy held in the per-CPU entropy pool will never be larger than the digest size of the used hash which is compliant to [2] section 3.1.5.1.1 table 1.

[11] Section **??** outlines that the ESDM collection size can be modified at compile time where the default is 1,024. When a different collection size is chosen, the value needs to be adjusted accordingly. Yet, such modified value has no impact to the entropy analysis.

- The entropy content of the input $h_{in_{aux\ pool}}$: The input entropy of the hash used to process the auxiliary pool is equal to the entropy provided by the noise source and the already collected entropy in the auxiliary pool considering that both data components are hashed at the same time to form a new auxiliary pool state. Of course, the entropy held in the auxiliary pool will never be larger than the digest size of the used hash which is compliant to [2] section 3.1.5.1.1 table 1.

- The entropy content of the input $h_{in_{hash\ pools}}$: The input entropy of the hash used to process the entire entropy pool is equal to the entropy found in all per-CPU entropy pools managed by the hash operation and the auxiliary pool. Again, the entropy generated by the hash will never be larger than the digest size of the used hash which is compliant to [2] section 3.1.5.1.1 table 1.

Function 2.2 Output_Entropy To perform a calculation of the Output_Entropy of a conditioning component, the input entropy must be considered. The heuristic input entropy awarded for one time stamp processed by the ESDM is given in equation 3.2. Due to the concatenation operation of time stamps, the entropy of multiple time stamps can be added.

For the function 2.2, the entropy when 1,024 time stamps are received is $h_{out_{SHA-512}} = Output\_Entropy_{SHA-512}(8192, 512, 512, 1024)$ following [2] section 3.1.5.1.2. Therefore, the following calculation is applicable:

$$P_{high} = 2^{-1024}$$

$$P_{low} = \frac{(1 - 2^{-1024})}{2^{8192} - 1} \approx 2^{-8192}$$

$$n = min(512, 512) = 512$$

$$\psi = 2^{8192-512} \cdot 2^{-8192} + 2^{-1024} = 2^{-512} + 2^{-1024} \approx 2^{-512}$$

$$U = 2^{8192-512} + \sqrt{2 \cdot 512 \cdot (2^{8192-512}) \cdot ln(2)} = 2^{7680} + \sqrt{2^{7690} \cdot ln(2)} \approx 2^{7680}$$

$$\omega = 2^{7680} \times 2^{-8192} = 2^{-512}$$

$$\underline{Output\_Entropy_{SHA-512}(8192, 512, 512, 1024) = -log_2(max(2^{-512}, 2^{-512})) = 512}$$

As a complement, the same calculation is provided when only one time stamp is received for the formula $h_{out_{SHA-512}} = Output\_Entropy_{SHA-512}(8, 512, 512, 1)$

$$P_{high} = 2^{-1}$$

$$P_{low} = \frac{(1 - 2^{-1})}{2^{8} - 1} \approx 2^{-9}$$

$$n = min(512, 512) = 512$$

$$\psi = 2^{8-512} \cdot 2^{-9} + 2^{-1} = 2^{-513} + 2^{-1} \approx 2^{-1}$$

$$U = 2^{8-512} + \sqrt{2 \cdot 512 \cdot (2^{8-512}) \cdot ln(2)} = 2^{-504} + \sqrt{2^{-494} \cdot ln(2)} \approx 2^{-247} \cdot \sqrt{ln(2)} \approx 2^{-248}$$

$$\omega = 2^{-248} \times 2^{-9} = 2^{-257}$$

$$\underline{Output\_Entropy_{SHA-512}(8, 512, 512, 1) = -log_2(max(2^{-257}, 2^{-1})) = 1}$$

The calculation can be generalized with the following formula:

$$h_{out_{SHA-512}} = Output\_Entropy_{SHA-512} = min(h_{in}, n_{out_{SHA-512}})$$

When using SHA-256, the same type of calculation can be provided. The first set of formulas show the case when 1,024 time stamps are received and thus for $h_{out_{SHA-256}} = Output\_Entropy_{SHA-256}(8192, 256, 256, 1024)$:

$$P_{high} = 2^{-1024}$$

$$P_{low} = \frac{(1 - 2^{-1024})}{2^{8192} - 1} \approx 2^{-8192}$$

$$n = min(256, 256) = 256$$

$$\psi = 2^{8192-256} \cdot 2^{-8192} + 2^{-1024} = 2^{-256} + 2^{-1024} \approx 2^{-256}$$

$$U = 2^{8192-256} + \sqrt{2 \cdot 256 \cdot (2^{8192-256}) \cdot ln(2)} = 2^{7936} + \sqrt{2^{7945} \cdot ln(2)} \approx 2^{7936}$$

$$\omega = 2^{7936} \times 2^{-8192} = 2^{-256}$$

$$\underline{Output\_Entropy_{SHA-256}(8192, 256, 256, 1024) = -log_2(max(2^{-256}, 2^{-256})) = 256}$$

As a complement, the same calculation is provided when only one time stamp is received for the formula $h_{out_{SHA-512}} = Output\_Entropy_{SHA-256}(8, 256, 256, 1)$

$$P_{high} = 2^{-1}$$

$$P_{low} = \frac{(1 - 2^{-1})}{2^{8} - 1} \approx 2^{-9}$$

$$n = min(256, 256) = 256$$

$$\psi = 2^{8-256} \cdot 2^{-9} + 2^{-1} = 2^{-257} + 2^{-1} \approx 2^{-1}$$

$$U = 2^{8-256} + \sqrt{2 \cdot 256 \cdot (2^{8-256}) \cdot ln(2)} = 2^{-248} + \sqrt{2^{-239} \cdot ln(2)} \approx 2^{-120} \cdot \sqrt{ln(2)} \approx 2^{-121}$$

$$\omega = 2^{-121} \times 2^{-9} = 2^{-130}$$

$$\underline{Output\_Entropy_{SHA-256}(8, 256, 256, 1) = -log_2(max(2^{-130}, 2^{-1})) = 1}$$

Again, the calculation can be generalized with the following formula:

$$h_{out_{SHA-256}} = Output\_Entropy_{SHA-256} = min(h_{in}, n_{out_{SHA-256}})$$

Comparing the conclusions for SHA-512 and SHA-256, both allow to draw the general conclusion that underlies the entire entropy assessment and therefore data entropy management applied by the ESDM for all vetted conditioning operations:

$$h_{out_{vetted}} = Output\_Entropy_{vetted} = min(h_{in}, n_{out_{vetted}}) \tag{3.1}$$

Function 2.8 The function 2.8 uses the same hash operation as the discussed in the preceding section. Thus, the conclusion drawn with equation 3.1 applies here as well.

Function 2.3 The function 2.3 uses the same hash operation as the discussed in the preceding section. Thus, the conclusion drawn with equation 3.1 applies here as well.

Conclusions for Output_Entropy As stated in [2] section 3.1.5.1.2, vetted conditioning components are allowed to claim full entropy. In case of full entropy, the following is applied which matches exactly analysis and conclusion of equation 3.1:

- $h_{out_{SHA-512}} = min(h_{in}, n_{out_{SHA-512}})$,
- $h_{out_{SHA-256}} = min(h_{in}, n_{out_{SHA-256}})$, or
- $h_{out_{SHA-1}} = min(h_{in}, n_{out_{SHA-160}})$.

Based on that conclusion, the entropy rate for each processing step given in section 2.2 can be illustrated in the following. This entropy assessment uses n_{out} which depends on the chosen hash operation with the respective value listed above for the chosen hash.

Heuristic Entropy Assessment The heuristic entropy value for the individual time stamps is defined with the following equation applicable when a high-resolution timer is present – the absence of a high-resolution timer automatically implies the ESDM is treated as non-compliant to SP800-90B:

$$h_{t_8} = h_{t_{32}} = 1 \tag{3.2}$$

Note, in order to assess whether the ESDM heuristic entropy value is appropriate, it must be compared with the entropy analysis result received from practical measurements such as outlined in sections 3.3.3 and 3.3.4. This comparison must show that the heuristic entropy value is always lower and thus more conservative than what the measurements show on target devices.

The entropy present in the arbitrary 32 bit word that may be added to the per-CPU collection pool is defined with:

$$h_{a_{32}} = 0 \tag{3.3}$$

The entropy in the concatenated time stamps found in the interrupt as well as scheduler per-CPU collection pool is calculated as the sum of all time stamps (truncated or not) present in the interrupt as well as scheduler per-CPU collection pool of 1,024 bytes per default – if a different collection pool size is used, the right-hand value of the following equation must be adjusted accordingly:

$$h_{per-CPU\ CP} = min(\sum_{n=0}^{number\ time\ stamps} (h_{t_{\{8,32\}}})_n, 1024) \tag{3.4}$$

For the maintenance of the interrupt per-CPU entropy pool as specified by equation 2.2, the following entropy rate applies when continuous compression support is enabled. This formula implies that each output of the interrupt per-CPU entropy pool holds the sum of the entropy of the received per-CPU collection pool since last generation of the per-CPU output data and the entropy remained in the per-CPU entropy pool capped by the message digest size. This operation implies that the used hash compresses of the entropy available in the different input data.

$$h_{per-CPU\ pool_n} = min(\sum_{m=0}^{m=n-1} h_{per-CPU\ CP_m} + h_{per-CPU\ pool_{n-1}}, n_{out}) \quad (3.5)$$

When continuous compression support is disabled as well as for the scheduler-based entropy source, the per-CPU entropy pool maintenance specified by equation 2.2 shows the following entropy rate. The formula implies that the maximum amount of entropy that can be held depends on the size of the collection pool depicted with equation 2.1 since additional entropy received by the collection pool overwrites old entropy data. The collection pool can hold the maximum amount of entropy event data as defined with its size. After converting the number of received entropy event data into an entropy statement using equation 3.2, the maximum amount of entropy held in the collection pool is available.

$$h_{per-CPU\ pool_n} = min(h_{per-CPU\ CP_n} + h_{per-CPU\ pool_{n-1}}, n_{out}) \quad (3.6)$$

Similarly, the following equation applies to the entropy of the auxiliary pool maintenance as specified by equation 2.8. Note, although this entropy source is not considered to be modeled in this chapter, the formula is still provided illustrating the use of a vetted conditioning component. This formula implies that auxiliary pool holds the sum of the entropy of the received data capped by the message digest size. Again, this operation implies that the used hash compresses the entropy available in the different input data.

$$h_{aux\ pool} = min(h_{in_{aux\ pool}} + h_{aux\ pool}, n_{out}) \quad (3.7)$$

The following equation applies when calculating the interrupt and scheduler-based entropy source output buffer before its truncation as specified by equation 2.3. The formula implies that the interrupt entropy source buffer before truncation holds the sum of the entropy of all per-CPU entropy pools plus the auxiliary pool capped by the message digest size. Again, this operation implies that the used hash compresses the entropy available in the different input data.

$$h_{hash\ pools} = min(\sum_{c=0}^{max\ CPU} h_{per-CPU\ pool_c}, n_{out}) \quad (3.8)$$

The entropy present in the truncated interrupt and scheduler-based entropy source buffer is the minimum of the entropy found in the pools and the requested amount of bits which is equal to the security strength of the DRBG:

$$requested\ size_s = security\ strength = 256 \quad (3.9)$$

$$h_s = min(h_{hash\ pools}, requested\ size_s) \tag{3.10}$$

The entropy of the temporary seed buffer following equation 3.1 is simply an addition of the entropy values credited for each of the entropy source with h_A denominating the auxiliary entropy entropy rate, h_E references the IRQ ES entropy rate, , h_S references the scheduler ES entropy rate, , h_J covers the Jitter RNG ES entropy rate, h_C references the CPU ES entropy rate, h_K references the kernel ES entropy rate, and h_H denominates the HWRAND entropy source:

$$h_T = h_A + h_E + h_S + h_J + h_C + h_K + h_H \tag{3.11}$$

The result of the formulas show that the entropy is simply a sum of the entropy of all input events capped to the message digest size of the used hash operation.

When generating the random numbers filling the interrupt entropy source buffer, the entropy is debited in the following steps. First the entropy estimator of the auxiliary pool is reduced as much as possible: either by h_s or at most to zero. If not all entropy of h_s could have been debited from the auxiliary pool entropy estimator, then the yet not debited part of h_s is debited from the per-CPU entropy pool entropy estimators.

For example, assume that after the generation of random numbers and filling the slow noise source buffer its entropy is $h_s = 256$. Assume further, the per-CPU entropy pools of the assumed 2 CPUs contain $h_{per-CPU\ pool_{CPU0}} = 185$ and $h_{per-CPU\ pool_{CPU1}} = 123$. The debit operation performs:

1. $h_{per-CPU\ pool_{CPU0}} = 185 - 185 = 0$ leaving $h_{s_{not\ debited}} = 256 - 185 = 71$
2. $h_{per-CPU\ pool_{CPU1}} = 123 - 71 = 52$

3.3.7 SP800-90B Section 3.1.6

The ESDM uses the following noise sources for the interrupt entropy source:

- The noise source of the timing of the occurrence of interrupts. The entire SP800-90B analysis covers this one noise source. Thus, the requirements in this section for the interrupt noise source are trivially met.
- Auxiliary data delivered as part of interrupts: HID event data. This data is concatenated with the interrupt time stamps into the collection pool. Yet it is always credited with zero bits of entropy.

All data is processed by the vetted conditioning component of the hash before it is injected as seed data into the DRNG. Thus, this operation complies with the last paragraph of section 3.1.6.

3.3.8 SP800-90B Section 3.2.1 Requirement 1

This entire document is intended to provide the required analysis.

3.3.9 SP800-90B Section 3.2.1 Requirement 2

This entire document in general and chapter 3 in particular is intended to provide the required analysis.

3.3.10 SP800-90B Section 3.2.1 Requirement 3

There is no specific operating condition other than what is needed for the operating system to run since the noise source is a complete software-based noise source.

The only dependency the noise source has is a high-resolution timer which does not change depending on the environmental conditions.

3.3.11 SP800-90B Section 3.2.1 Requirement 4

This document explains the architectural security boundary.

The boundary of the implementation is the source code files provided as part of the software delivery. This source code contains API calls which are to be used by entities using the ESDM.

3.3.12 SP800-90B Section 3.2.1 Requirement 5

The per-CPU entropy pools as processed by the hash is the output of the interrupt noise source. I.e. the entropy pools maintained by the hashing operation holds the data that is given to the DRNG when requesting seeding.

The noise source output without the hashing operation is accessed with specific tools which add interfaces that are not present and thus not usable when employing the ESDM in production mode. These additional interfaces are used for gathering the data used for the analysis documented in section 3.3.3. These interfaces perform the following operation:

1. Switch the ESDM into raw entropy generation mode. This implies that each raw entropy event is fed to the raw entropy collection interface and not processed by the per-CPU collection pool or otherwise used.

2. When an interrupt event is received, forward the time stamp holding the entropy to a ring buffer. This operation is performed repeatedly until the ring buffer is full or the user space application read that ring buffer.

3. When an application requests the reading of the ring buffer, the data is extracted from the kernel and the ring buffer is cleared.

With this approach, the actual interrupt events which would be processed by the ESDM are obtained.

The kernel interface is only present if the kernel is compiled with the option `CONFIG_ESDM_RAW_HIRES_ENTROPY`. This option should not be set in production kernels.

3.3.13 SP800-90B Section 3.2.1 Requirement 6

Please see section 3.2.3 for details how and why the raw entropy extraction does not substantially alter the noise source behavior.

3.3.14 SP800-90B Section 3.2.1 Requirement 7

See section 3.3.4 for a description of the restart test.

3.3.15 SP800-90B Section 3.2.2 Requirement 1

This entire document provides the complete discussion of the noise source.

3.3.16 SP800-90B Section 3.2.2 Requirement 2

The noise source is based on the receipt of interrupts. The receipt of interrupts follows the usage of the system. The more I/O is performed with the system, the more interrupts are received by the ESDM. The entropy rate only is a function of the received I/O events and the timer and does not depend on any other system property such as physical characteristics (e.g. temperature variations or voltage/current variations). This finding is consistent with the fact that the noise source is a pure software-based noise source which relies on the presence of a high-resolution timer. Note, the used timer is a cycle counter that increments with a given rate.

3.3.17 SP800-90B Section 3.2.2 Requirement 3

See sections 3.3.6 for a discussion of the entropy provided by the interrupt noise source.

A stochastic model is not provided.

3.3.18 SP800-90B Section 3.2.2 Requirement 4

The noise source is expected to execute in the kernel address space. This implies that the operating system process isolation and memory separation guarantees that adversaries cannot gain knowledge about the ESDM operation.

3.3.19 SP800-90B Section 3.2.2 Requirement 5

The output of the noise source is non-IID as it rests on the execution time of a fixed set of CPU operations and instructions.

3.3.20 SP800-90B Section 3.2.2 Requirement 6

The noise source generates the data via the hash generation function as outlined in section 3.3.5.

Although the hash commonly generates a fixed-length string, this string length may be reduced by the amount of available entropy as outlined in section 3.3.6.

3.3.21 SP800-90B Section 3.2.2 Requirement 7

N/A as no additional noise source is implemented with the interrupt entropy source.

Though, the ESDM employs complete self-contained other entropy sources which may be compliant to SP800-90B by itself. To seed the DRNG maintained by the ESDM, the output of all entropy sources are concatenated compliant to SP800-90C as outlined in section 5.1.

3.3.22 SP800-90B Section 3.2.3 Requirement 1

The conditioning component is the hash operation. See section 3.3.5 for a discussion of the input and output sizes.

3.3.23 SP800-90B Section 3.2.3 Requirement 2

The used hash implementations for the conditioning components functions are all ACVP-testable. The ESDM offers an ACVP interface to ensure also the built-in SHA-256 and SHA-1 implementations are testable.

3.3.24 SP800-90B Section 3.2.3 Requirement 3

For the defined hashes, no key is required.

3.3.25 SP800-90B Section 3.2.3 Requirement 4

For the defined hashes, no key is required.

3.3.26 SP800-90B Section 3.2.3 Requirement 5

The conditioning component is the hash operation. See section 3.3.6 for a discussion of the narrowest internal width and the output block size.

3.3.27 SP800-90B Section 3.2.4 Requirement 1

Test tools for measuring raw entropy are provided at the ESDM web page. These tools can be used by everybody without further knowledge of the ESDM.

3.3.28 SP800-90B Section 3.2.4 Requirement 2

The operation of the test tools for gathering raw data are discussed in section 3.3.3. This explanation shows that the raw unconditioned data is obtained.

3.3.29 SP800-90B Section 3.2.4 Requirement 3

The provided tools for gathering raw entropy contains exact steps how to perform the tests. These steps do not require any knowledge of the noise source.

3.3.30 SP800-90B Section 3.2.4 Requirement 4

The raw entropy tools can be executed on the same environment that hosts the ESDM. Thus, the data is generated under normal operating conditions.

3.3.31 SP800-90B Section 3.2.4 Requirement 5

The raw entropy tools can be executed on the same environment that hosts the ESDM. Thus, the data is generated on the same hardware and operating system that executes the ESDM.

3.3.32 SP800-90B Section 3.2.4 Requirement 6

The test tools are publicly available at LRNG web page allowing the replication of any raw entropy measurements.

3.3.33 SP800-90B Section 3.2.4 Requirement 7

Please see section 3.2.3 for details how and why the raw entropy extraction does not substantially alter the noise source behavior.

3.3.34 SP800-90B Section 4.3 Requirement 1

The implemented health tests comply with SP800-90B sections 4.4 as described in section 3.3.43.

3.3.35 SP800-90B Section 4.3 Requirement 2

When either health test fails, the kernel:

- Emits a failure log,
- Resets the noise source, and
- Restarts the SP800-90B startup health tests.

This implies that no data is produced by the ESDM (including its DRNG) when using the SP800-90B compliant external interfaces.

Both health test failures are considered permanent failures and thus trigger a full reset.

3.3.36 SP800-90B Section 4.3 Requirement 3

The following false positive probability rates are applied:

- RCT: The false positive rate is $\alpha = 2^{-30}$ and therefore complies with the recommended false positive probability.
- APT: The cut-off value is set to 325 compliant to SP800-90B section 4.4.2 for non-binary data at a significance level of $\alpha = 2^{-30}$ with time stamp is assumed to at least provide one bit of entropy, i.e. $H = 1$[12].

3.3.37 SP800-90B Section 4.3 Requirement 4

The ESDM applies a startup health test of 1,024 noise source samples. Additional tests are applied. The collected noise source samples are re-used for the generation of random numbers if the startup test was successful.

3.3.38 SP800-90B Section 4.3 Requirement 5

The noise source supports on-demand testing in the sense that the caller may restart the kernel.

[12] Note, the referenced Excel function seems to be imprecise when calculating the value. The data has been obtained using R-Project with the formula of $1 + qbinom(1 - 2^{-30}, 512, 2^{-1})$.

3.3.39 SP800-90B Section 4.3 Requirement 6

The health tests are applied to the raw, unconditioned time stamp data directly obtained from the noise source before they are injected into the per-CPU collection pool and further processed with the hash conditioning component.

3.3.40 SP800-90B Section 4.3 Requirement 7

The health tests are documented with section 2.6.2.

The tests are executed as follows:

- During startup, the RCT and the APT are applied to 1,024 samples. The startup test can be triggered again when the caller reboots the kernel.
- At runtime, the RCT is applied to each received time stamp. The APT collects 512 time stamps. The APT is calculated over all 512 time stamps. If the test fails, the entire ESDM is reset to drop all existing entropy and the startup testing is performed again.

3.3.41 SP800-90B Section 4.3 Requirement 8

There are no currently known suspected noise source failure modes.

3.3.42 SP800-90B Section 4.3 Requirement 9

N/A as the noise source is pure software. The software is expected to execute on hardware operating in its defined nominal operating conditions.

3.3.43 SP800-90B Section 4.4

The health tests described in section 2.6.2 are applicable to cover the requirements of SP800-90B health tests.

The SP800-90B compliant health tests are implemented with the following rationale:

RCT The Repetition Count Test implemented by the ESDM compares two back-to-back time stamps to verify that they are not identical. If the number of identical back-to-back time stamps reaches the cut-off value of 30, the RCT test raises a failure that is reported and causes a reset the ESDM. The RCT uses the a cut-off value that is based on the following: $\alpha = 2^{-30}$ compliant to FIPS 140-2 IG 9.8 and compliant to SP800-90B which mandates this value to be in the range $2^{-20} \leq \alpha \leq 2^{-40}$. In addition, one time stamp is assumed to at least provide one bit of entropy, i.e. $H = 1$. When applying these values to the formula given in SP800-90B section 4.4.1, the cut-off value of 31 is calculated.

When the RCT passes, the counter is set to zero for the next time delta to arrive. In mathematical terms, the verification of back-to-back values being not identical is the calculation of the first discrete derivative of the time stamp to show that it is not zero. In addition, the ESDM enhances the RCT by calculating also the second and third discrete derivative of the time stamp to be concatenated with the per-CPU collection pool. With that, up to 8 consecutive time stamp values are assessed. All derivatives

must always be non-zero in order to pass the RCT. If one discrete derivative shows a zero, the RCT counter is increased. Thus, the addition of the second and third derivative makes the RCT even more conservative. Hence, the first discrete derivative is considered to be identical to the "approved" RCT specified in SP800-90B section 4.4. In addition, linear and exponential patterns are identified with the second and third discrete derivative, respectively. As the additional pattern recognition do not invalidate the mandatory pattern recognition, this RCT approach therefore is considered to be an enhanced version of the "approved" RCT and thus meets the requirement (a) of SP800-90B section 4.5.

APT The ESDM implements the Adaptive Proportion Test as defined in SP800-90B section 4.4.2. As explained in other parts of the document, one time stamp value is assumed to have (at least) one bit of entropy. Thus, the cut-off value for the APT is 325 compliant to SP800-90B section 4.4.2 for non-binary data with a significance level of $\alpha = 2^{-30}$. The APT is calculated using the four least significant bits of the time stamp. During initialization of the APT, a time stamp is set as a base. All subsequent time stamps are compared to the base time stamp. If both values are identical, the APT counter is increased by one. The window size for the APT is 512 time stamps. The implementation therefore provides an "approved" APT.

3.4 NIST Clarification Requests

In addition to complying with the requirements of FIPS 140-2 and SP800-90B, NIST requests the clarification of the following questions.

3.4.1 Sensitivity of Interrupt Timing Measurements

The question that needs to be answered is whether the logic that measures the interrupt timing is sensitive enough to pick up the variances of the interrupt timing.

The sensitivity implies that timing variations are picked up and measured. This is enforced by the stuck test enforced on each interrupt time stamp. That stuck test requires that the first, second and third discrete derivative of the time stamp must always be non-zero to accept that time stamp. Therefore, the time stamp must vary for the received and processed interrupts which implies that the ESDM health test ensures that the sensitivity of the time stamp mechanism is sufficient.

3.4.2 Dependency Between Interrupt Timing Measurements

Another question that is raised by NIST asks for a rationale why there are no dependencies between individual Jitter measurements.

The interrupts are always created by either explicit or implicit human actions. The ESDM measures the time stamp of the occurrence of these interrupts. Thus, the ESDM measures the effects of operations triggered by human interventions. With the presence of a high-resolution time stamp that operates in the nanosecond range and the assumption that only one bit of entropy is present in one nanosecond time stamp of one interrupt event, the dependency discussion

therefore focuses on the one (or maybe up to four) least significant bit of the nanosecond time stamp. With such high-resolution time stamps and considering that only the least significant bit(s) is/are relevant for the ESDM, dependencies are considered to be not present for these bits.

3.5 SP800-90B Compliant Configuration

In order to use the ESDM SP800-90B compliant, the following configurations and settings must be made. These settings are cover requirements for the compile-time options found in the ESDM meson configuration framework.

The following compile-time settings must be observed:

- `fips140` must be enabled.
- `es_irq` must be enabled to enable the interrupt entropy source to which the entropy discussion of this chapter applies to.
- Set the configuration option of `es_irq_entropy_rate` to the rate resulting from the entropy assessment outlined in section 3.6.
- The ESDM Linux kernel module must be enabled and compiled as documented in addon/linux_esdm_es/README.md. The Makefile in this directory must have the interrupt entropy source enabled.

The following requirements apply to the runtime configuration:

- The kernel must be booted with the kernel command line option of `fips=1` to enable the SP800-90B health test.

To verify that the SP800-90B compliance is achieved, the file /proc/sys/kernel/random/esdm_type or the `esdm_rpcc_status` RPC API call provides an appropriate status indicator.

To achieve a compliant configuration to SP800-90A and SP800-90B, the following requirements must be met:

- All requirements for SP800-90B documented in section 3.5 must be met.
- The ESDM compile time option `drng_hash_drbg` must be enabled which provides the SP800-90A DRBG.

Only data obtained from the potentially blocking output interfaces of the ESDM are SP800-90B compliant. If SP800-90C compliance is requested, these interfaces are also providing SP800-90C compliant output. Finally, the following interfaces are DRG.3 compliant:

- /dev/random,
- `getrandom` system call invoked with a zero flag value,
- invoking the in-kernel `esdm_get_random_bytes_full` API,
- invoking the in-kernel `esdm_get_random_bytes_pr` API that provides access to the DRBG operated with prediction resistance.

Any other interface is not considered to provide SP880-90B compliant data.

3.6 Reuse of SP800-90B Analysis

To reuse the SP800-90B analysis provided in this document the following steps must be performed on the target platform:

1. Obtain raw noise data through the raw noise source interface on the intended target platform as explained in section 3.3.3. The obtained raw noise data must be processed by the SP800-90B tool to obtain an entropy rate which must be above the entropy rate per time delta that is configured with `es_irq_entropy_rate`: the entropy rate must be above $^{256}/_{\texttt{es\_irq\_entropy\_rate}}$.

2. Obtain the restart noise data through the raw noise source interface on the intended target platform as explained in section 3.3.3. The obtained raw noise data must be processed by the SP800-90B tool to verify:

   (a) the sanity test to apply to the noise restart data must pass, and

   (b) the minimum of the row-wise and column-wise entropy rate must not be less than half of the entropy rate from measurement (1) and the entropy assessment of the noise source based on the restart data must be at least entropy rate per time stamp mentioned in (1).

If these steps are successfully mastered the user would now satisfy all SP800-90B criteria and thus does not need to prepare his own SP800-90B analysis since the document we discuss here covers all other aspects of the SP800-90B analysis.

The tool set provided as part of ESDM library code distribution provides the measurements and validation tools.

4 Scheduler Entropy Source Assessment

TBD

5 ESDM Specific Configurations

The ESDM offers a secure and appropriate set of features with the default configuration. Yet, use cases may arise where the ESDM should exhibit a different behavior. The flexibility of the ESDM allows a various configurations that are intended to meet different requirements.

5.1 SP800-90C Compliance

The specification of SP800-90C defines construction methods to design non-deterministic as well as deterministic RNGs. As the specification is currently in draft form, the latest available draft from January 21, 2021 is applied.

The specification defines different types of RNGs where the following mapping to the ESDM applies:

- The ESDM follows the construction of RBG2(NP) with the following entropy sources whose outputs are concatenated:

- Auxiliary external entropy sources, such as user-space rngd, if available. The administrator is responsible to guarantee that this entropy source is compliant to SP800-90B if it alters the entropy estimator maintained by the ESDM via the `RNDADDENTROPY` IOCTL. Depending on the selected entropy source, this may be a physical or non-physical entropy source. Note, all data maintained by the auxiliary pool are processed with a vetted conditioning component. Thus, to achieve full SP800-90C compliance for such entropy sources, only one should feed data credited with entropy into the auxiliary pool.
- Interrupt entropy source is the only entropy source that is fully maintained as part of the ESDM and is subject to a full entropy analysis following section 3.3. Furthermore, it is claimed to be fully SP800-90B compliant.
- Scheduler entropy source is the only entropy source that is fully maintained as part of the ESDM and is subject to a full entropy analysis following 4. Furthermore, it is claimed to be fully SP800-90B compliant.
- The Jitter RNG entropy source is used by the ESDM. This entropy source is a fully self-contained SP800-90B entropy source. Its SP800-90B compliance must be assessed separately. If SP800-90B compliance cannot be demonstrated, it must be awarded to credit zero bits of entropy with the configuration documented in section 2.8.1 or by completely disabling it by not selecting the meson configuration option of `es_jent`.
- The CPU entropy source is used by the ESDM, if available. On Intel, this uses RDSEED. This entropy source is a fully self-contained SP800-90B entropy source. Its SP800-90B compliance must be assessed separately. If SP800-90B compliance cannot be demonstrated, it must be awarded to credit zero bits of entropy with the configuration documented in section 2.9.1 or by completely disabling it by not selecting the meson configuration option of `es_cpu`.
- The Linux HWRAND entropy source is used by the ESDM, if available. Its SP800-90B compliance must be assessed separately. If SP800-90B compliance cannot be demonstrated, it must be awarded to credit zero bits of entropy with the configuration documented in section 2.9.1 or by completely disabling it by not selecting the meson configuration option of `es_hwrand`.

By applying Method 2 of section 3.3 in SP800-90C, the entropy provided by all entropy sources can be added which is applied when the ESDM constructs the temporary seed buffer as shown with equation 2.12.

- During instantiation of the DRBG, the ESDM tries to seed the DRBG with at least (DRBG security strength) + 128 bits = 384 bits of entropy. Only when this amount of entropy was obtained to seed the DRBG is considered to be fully seeded and is allowed to produce output.

  If the SP800-90C construction is to be used as the "randomness source" following bullet 5 of section "Note to Reviewer" in the SP800-90C document to seed another DRBG with a security strength of 256 bits, either

the hash DRBG should be used as they are capable of transporting up to 512 bits of entropy and are initially seeded with 384 bits of entropy. Thus, the SP800-90C seeding requirement of providing 384 bits of entropy can be satisfied.

- Reseeding of the DRBG can be triggered by either writing data into /dev/random or by the `RNDRESEEDCRNG` IOCTL. The ESDM guarantees that the reseed operation is only performed if at least 256 bits of entropy is available if operated compliant to SP800-90C. If this is not available, the reseed is attempted during the next generate operation. Yet, the generate operation is conducted in any case. Nonwithstanding, the ESDM always attempts to obtain 256 bits of entropy for reseeding. This is considered appropriate because the upper limit when a reseed is ultimately triggered is 2^{20} generate operations or after 10 minutes, whatever is reached earlier. If the DRNG cannot be reseeded it will continue to operate until the next time this threshold is reached. The DRNG will revert to an unseeded stage if it cannot be reseeded by the time it serviced 2^{30} generate requests since the last successful seeding operation.

The requirements from section 6.3 SP800-90C are met as follows:

1. The administrator must use the SP800-90A DRBG ESDM extension as mentioned above to satisfy the requirement.

2. The DRBG can be ACVTS-tested to show compliance to SP800-90A. The entropy sources are to be assessed pursuant to SP800-90B as outlined in the above listing.

3. See the above listing for the reseeding support.

4. If an entropy source is not validated, its entropy estimation must be set to zero as outlined in the above listing.

5. N/A as the ESDM is claimed to conform with RBG2(NP).

6. The entropy sources are listed above. By using concatenation of the output of all entropy sources, the Method 2 SP800-90C is implemented. The entropy of all entropy sources are added.

7. Technically it is possible that all DRBG security strengths can be chosen as the DRBG supports all security strengths. Yet, the ESDM interfaces currently only support the highest security strength of 256 bits to ensure that it can be used for all use cases.

8. The entropy source output is destroyed immediately after it was used to (re)seed the DRBG. Note, the use of the seed for backward secrecy by injecting it into the auxiliary pool via the vetted conditioning operation is considered to not violate the requirement as the seed data is unrecoverable. Furthermore, the seed data is not credited with any entropy during the backward secrecy operation. Therefore, the seed data is only used to further mix the internal state of the ESDM.

9. N/A as the ESDM does not use a CTR DRBG without derivation function.

10. The ESDM attempts to instantiate a DRBG with 3/2s bits of entropy. Only if this succeeds, the DRBG becomes available. The ESDM attempts to reseed a DRBG with s bits of entropy. See the rationale above for the discussion about the minimum entropy size of the reseeding operation of 128 bits.

11. The DRBG only provides output once it is fully seeded as mandated by SP800-90C.

12. An error occurring in the interrupt entropy source triggers a full reset of the ESDM as outlined in section 2.6.2. If the other entropy sources are subject to a health test failures, SP800-90B mandates that they do not produce entropy. Before the first initialization of the DRBG, it is subject to a power-on self test. The ESDM performs power-up self tests as outlined in section 2.14.

13. This requirement is implicitly met by the fact that the ESDM only provides DRBGs with the maximum security strength of 256 bits.

5.1.1 RBG2(P) Construction Method

It is possible to convert the ESDM into the SP800-90C type of RBG2(P). This approach requires that only physical entropy sources are credited with entropy. The following specific settings must be applied in addition to the general configurations listed in the next section:

- Configure the interrupt entropy source to not credited entropy: compile the ESDM with the kernel configuration option of `es_irq_entropy_rate=0`. This setting ensures that the interrupt entropy source still collects data, but it is not credited with entropy. If it shall be completely disabled, the compile time option of `es_irq` must be unset.

- Configure the scheduler entropy source to not credited entropy: compile the ESDM with the meson configuration option of `es_sched_entropy_rate=0`. This setting ensures that the scheduler entropy source still collects data, but it is not credited with entropy. If it shall be completely disabled, the compile time option of `es_sched` must be unset.

- Configure the Jitter RNG entropy source to not credit entropy: compile the ESDM with the meson configuration option of `es_jent_entropy_rate=0`. If the value is set to 0, the entropy source still collects data, but it is not credited with entropy. If it shall be completely disabled, the compile time option of `es_jent` must be unset.

- Configure the kernel RNG entropy source to not credit entropy: compile the ESDM with the meson configuration option of `es_kernel_entropy_rate=0`. If the value is set to 0, the entropy source still collects data, but it is not credited with entropy. If it shall be completely disabled, the compile time option of `es_kernel` must be unset.

- Adjust the entropy rate of the CPU entropy source as needed: compile the ESDM with the meson configuration option of `es_cpu_entropy_rate=256` when full entropy is assumed to be provided by the CPU entropy source.

In general, set any value between 0 and 256 if the default value is not appropriate. If the value is set to 0, still collects data, but it is not credited with entropy. If it shall be completely disabled, the compile time option of `es_cpu` must be unset.

- Adjust the entropy rate of the HWRAND entropy source as needed: compile the ESDM with the meson configuration option of `es_hwrand_entropy_rate=256` when full entropy is assumed to be provided by the configured HWRAND entropy source. Ensure that the intended HWRAND entropy source is set in /sys/devices/virtual/misc/hw_random/rng_current. In general, set any value between 0 and 256 if the default value is not appropriate. If the value is set to 0, still collects data, but it is not credited with entropy. If it shall be completely disabled, the compile time option of `es_hwrand` must be unset.

- If other hardware entropy sources are considered deliver entropy, they may inject data into the ESDM via the IOCTL `RNDADDENTROPY`. Note, only physical entropy sources must provide data that is credited with entropy. Any other data fed into the ESDM via those interfaces must not be credited with entropy.

Important note: The entropy rate provided by the CPU entropy source plus all other physical entropy sources together must ensure they provide sufficient entropy. "Sufficient entropy" is provided when the entropy rate equals the hash type used by the ESDM. Considering the default message digest of SHA2-512, 512 bits of entropy should be provided.

Naturally, the hardware entropy sources that are credited with entropy must be compliant to SP800-90B.

In case the RBG2(P) construction method is achieved, the following additional requirement from section 6.3 SP800-90C is met:

- Requirement 5: With the required configuration mentioned before the ESDM will only count the physical entropy sources towards fulfilling the requested amount of entropy.

5.1.2 SP800-90C Compliant Configuration

SP800-90C compliance is only achieved when all of the following settings are achieved.

The following compile-time settings must be observed:

- The meson configuration option `oversample_es` is set. This option guarantees the following:
  - The final conditioning operation applied to the interrupt entropy source as well as to the auxiliary pool require 64 additional bits of entropy when obtaining data for the temporary seed buffer. This complies with the requirement specified in section 4.3.2 SP800-90C about vetted conditioning components. This ensures the conditioning components provide full entropy.
  - When the DRNG is initially seeded, it is attempted to be seeded with 384 bits of entropy at least. This complies with the requirement

specified in section 6.2.1 bullet 2 of SP800-90C requiring 3/2s bits of entropy for the initial seeding. Note, the ESDM applies a stepwise seeding of 32, 128 and 256 bits of entropy during initialization. When the final step of 256 bits is to be performed, the ESDM will guarantee that at least 384 bits of entropy are collectively pulled from all entropy sources. Only if this is achieved, the SP800-90C compliant-marked interfaces of the ESDM specified in section 2.6.2 will produce random numbers.

- Enable `es_irq` to use the interrupt entropy source compliant to SP800-90B as outlined in section 3.5 if the interrupt entropy source is considered to be compliant to SP800-90B by accepting the assessment in chapter 3 and by applying the testing of the entropy source as outlined in section 3.6 on the target system. This also requires the compilation of the ESDM Linux kernel module documented in addon/linux_esdm_es/README.md.

- Enable `es_sched` configure the scheduler entropy source compliant to SP800-90B as outlined in section 4 if the scheduler entropy source is considered to be compliant to SP800-90B by accepting the assessment in chapter 4 and by applying the testing of the entropy source as outlined in this chapter on the target system. This also requires the compilation of the ESDM Linux kernel module documented in addon/linux_esdm_es/README.md.

- The ESDM must be compiled with the meson configuration option of `es_cpu_entropy_rate=0` or unset the option `es_cpu` unless the CPU-based entropy source (e.g. RDSEED on Intel) has an SP800-90B compliant entropy assessment and comply with all requirements from SP800-90B.

- The ESDM must be compiled with the meson configuration option of `es_hwrand_entropy_rate=0` or unset the option `es_hwrand` unless the HWRAND entropy source has an SP800-90B compliant entropy assessment and comply with all requirements from SP800-90B.

- The ESDM must be compiled with the kernel configuration option of `es_jent_entropy_rate=0` or unset the option `es_jent` unless the Jitter RNG entropy source has an SP800-90B compliant entropy assessment and comply with all requirements from SP800-90B[13].

The following requirements apply to the runtime configuration:

- The ESDM along with the scheduler and/or interrupt ES must be operated in FIPS mode to ensure the oversampling for the conditioning and the initial seeding of the DRBG is applied. This is achieved with the kernel command line option of `fips=1`.

- Filling up the ESDM with entropy using either a user-space RNGD via the IOCTL `RNDADDENTROPY` is allowed. However, the caller is only allowed to claim entropy associated with the data and thus increase the ESDM entropy estimation if the entropy source is SP800-90B compliant with its own entropy assessment.

[13] At the time of writing, the user space Jitter RNG is SP800-90B compliant. Patches ensuring the in-kernel variant is SP800-90B compliant as well when into the kernel for version 5.8.

To verify that the SP800-90C compliance is achieved, the file /proc/ESDM_type provides an appropriate status indicator. The SP800-90C compliant is ensured only for the respectively marked ESDM interfaces specified in section 2.6.2. All other interfaces are not providing SP800-90C compliant random numbers.

It is permissible to use the ESDM DRBG output to seed another DRBG when using the prediction resistance `esdm_rpcc_get_random_bytes_pr` RPC API call, using `getrandom(2)` with the `GRND_RANDOM` flag, or opening /dev/random with the `O_SYNC` flag. See the details given for the API call `esdm_get_entropy_bytes_pr` in section 3.

5.2 AIS 20 / 31

The German BSI defines construction methods of RNGs with AIS 20/31 with the revision from 2011. In particular, this document defines different classes of RNGs where the ESDM is capable of meeting class NTG.1. The NTG.1 compliance is achieved when using /dev/random with the `O_SYNC` flag during opening, by using `getrandom(2)` with the `GRND_RANDOM` flag, or by using the `esdm_rpcc_get_random_bytes_pr` RPC API call.

The NTG.1 from AIS 20/31 of 2011 requirements are met as follows:

- NTG.1.1: The ESDM offers interfaces to collect raw unconditioned entropy as outlined in section 2.15.

- NTG.1.2: The ESDM maintains an entropy estimator for the auxiliary pool as well as each entropy pool. The credit operation of the ESDM is documented in section 2.1. To outline the debit operation section 3.3.6 provides several examples showing the approach. For each of the fast entropy sources, the ESDM obtains the entropy estimate when fetching data from them and applies this estimate when seeding the DRNG.

- NTG.1.3: The ESDM provides a compile-time switchable DRNG support using an SP800-90A Hash DRBG with a SHA-512 core. Furthermore, the ESDM allows using a ChaCha20-based DRNG which is documented in section 2.12.2. All those DRNGs are characterized as DRG.3 since they all support backward secrecy and use contemporary cryptographic primitives for the state transition as well as output generation.

- NTG.1.4: Due to the blocking behavior and only returning as many bits of data as entropy is available via /dev/random opened with `O_SYNC`, or `getrandom(2)` with the flag `GRND_RANDOM` the ESDM meets the requirement.

- NTG.1.5: The output functions of the offered DRNGs are all based on contemporary cryptographic primitives. Therefore, the output is assumed that it cannot be distinguished from random number output sequences of an ideal RNG. This is supported by the fact that the SP800-90A DRBGs are successfully passing the NIST CAVP test suite. The ChaCha20 DRNG can be tested using the ChaCha20 DRNG stand-alone implementation which is derived from the ESDM implementation. The tester can verify that both operate identical because the power-on self test in the stand-alone ChaCha20 DRNG implementation is identical to the ESDM power-up self test found in ESDM_selftest.c.

- NTG.1.6: Using the entropy analysis tools such as the NIST SP800-90B test tools, results shown in appendix C can be obtained for the interrupt entropy source. For the second entropy source, such assessment must be performed as well. For example, the Jitter RNG is provided with a complete independent entropy assessment. When comparing the results with the heuristic entropy estimate, it is clear that the ESDM collects more entropy than it credits. Further, considering the mathematical operations processing the raw entropy data outlined in section 3.3.5, the processing is considered to not destroy entropy.

The NTG.1 requirements from AIS 20/31 from 2022 are identical with the following exception:

- When compiling the ESDM with the option `ais2031` enabled the initial seeding strategy of the ESDM with the 32/128/256 bit reseeds is altered to 32/128 bits followed by two entropy sources required to deliver 220 bits of entropy. Only when the last step is completed, the blocking output interfaces listed in section 3.5 are released to produce random numbers. This implies that those interfaces operate NTG.1-compliant.

5.2.1 NTG.1 (AIS 20/31 2011) Compliant Configuration

To achieve NTG.1 compliant operation, the following configuration must be applied. It is permissible to use the configuration in parallel with the SP800-90B and SP800-90C configurations.

The following compile-time settings must be observed:

- None.

The following requirements apply to the runtime configuration:

- The caller must use /dev/random and open it with `O_SYNC`, the `getrandom(2)` API call with the flag `GRND_RANDOM`, or the `esdm_rpcc_get_random_bytes_pr` RPC API call is used.

To verify that the NTG.1 compliance is achieved, the file /proc/sys/kernel/random/esdm_type provides an appropriate status indicator. The NTG.1 compliant is ensured only for the respectively marked ESDM interfaces specified in section 3.5. All other interfaces are not providing NTG.1 compliant random numbers.

5.2.2 NTG.1 (AIS 20/31 2022) Compliant Configuration

To achieve NTG.1 compliant operation, the following configuration must be applied. It is permissible to use the configuration in parallel with the SP800-90B and SP800-90C configurations.

The following compile-time settings must be observed for AIS 20/31 (2022):

- The ESDM compile time option of `ais2031` must be enabled.

- The ESDM must have access to two entropy sources which are capable of delivering 220 bits of entropy each. For example, the interrupt entropy source can deliver 220 bits of entropy unless being operated with SHA-1 on old Linux kernels (see section 2.2 for details on SHA-1). In

addition, the Jitter RNG entropy source is capable of delivering 220 bits of entropy when booting the kernel with the kernel compile-time option of `es_jent_entropy_rate=220` or a larger value. See section 2.8.1 for details on the Jitter RNG entropy rate.

The following requirements apply to the runtime configuration:

- None.

To verify that the NTG.1 compliance is achieved, the file /proc/sys/kernel/random/esdm_type provides an appropriate status indicator. The NTG.1 compliant is ensured only for the respectively marked ESDM interfaces specified in section 3.5. All other interfaces are not providing NTG.1 compliant random numbers.

5.2.3 DRG.4 / PTG.3 Compliant Configuration

The ESDM allows a configuration that would make the ESDM compliant to the requirements of a DRG.4 as well as a PTG.3. This is achieved by deconfiguring all entropy sources except one that provides the data from a PTG.2 entropy source. For example, an rngd can be created having access to a smart card that is provides a PTG.2 entropy source. In this case, all other entropy sources must be deactivated using the following runtime configuration.

The PTG.3 naturally can only be claimed if the PTG.2 entropy source is part of the validation. If the PTG.2 entropy source is not part of the validation, a DRG.4 must be claimed instead.

For achieving such compliance claim, the following compile-time configuration must be applied:

- The kernel command line must contain the following setting to ensure the interrupt entropy source is not credited with entropy: `es_irq_entropy_rate=0`.
- Configure the Jitter RNG entropy source to not credit entropy: boot the kernel with the ESDM compile-time option of `es_jent_entropy_rate=0`. If the value is set to 0, the entropy source is disabled.
- Adjust the entropy rate of the CPU entropy source to zero: boot the kernel with the ESDM compile-time option of `es_cpu_entropy_rate=0`. If the value is set to 0, the entropy source is disabled.
- Adjust the entropy rate of the HWRAN entropy source to zero: boot the kernel with the ESDM compile-time option of `es_hwrand_entropy_rate=0`. If the value is set to 0, the entropy source is disabled.

The following requirements apply to the runtime configuration:

- The PTG.2 entropy source delivers its entropy to the ESDM via either the `esdm_rpcc_rnd_add_entropy` RPC API call or via the /dev/random IOCTL `RNDADDENTROPY`.

A Thanks

Special thanks for providing input as well as mathematical support goes to:

- DJ Johnston
- Yi Mao
- Sandy Harris
- Dr. Matthias Peter
- Quentin Gouchet

B Source Code Availability

The source code, this document as well as the test code for all aforementioned tests is available at `http://www.chronox.de/esdm.html`.

C SP800-90B Entropy Measurements

The following table presents the SP800-90B entropy measurements indicating whether the found entropy is sufficiently high to support the entropy analysis given in section 3.3.5. Entropy values are given in bits and apply to the entropy found in one time stamp generated when receiving an interrupt event. The testing shown in this section provides the quantiative foundation of the entropy analysis compliant to sections 3.3.6 as well as all other assessments required for SP800-90B.

The testing collected raw unconditioned time stamps as delivered by the file `/sys/kernel/debug/esdm_testing/esdm_raw_hires`. The entropy calculation is based on 1,000,000 raw time stamps collected by the ESDM. To speed up the raw time stamp collection as well as to obtain a worst-case assessment, all test systems were either ping-flooded or within an SSH-session a `find /` was executed to generate a large number of interrupts in a short amount of time. The ping-flood generator was in close network proximity (e.g. KVM host, or a system at most one switch away from the test system).

The entropy result listing in the table below is generated as follows. The time stamps generated by the ESDM for each interrupt event is extracted and concatenated to form a bit-stream. This bit stream is processed by the NIST SP800-90B entropy analysis tool to obtain an entropy rate. This entropy rate is listed below. As the 8 least significant bits (LSB) of the time stamp are used and the other bits are ignored by the ESDM, the entropy rate applies to those 8 data bits. As discussed in sections 3.3.6, the ESDM assumes that each time stamp provides at least slightly more than one bit of entropy. As all values in the table below show significantly more entropy even with the worst-case measurement of 8 LSB, the ESDM underestimates the entropy existing in the respective system. Thus, the ESDM is considered to operate securely on these systems. The test complies with SP800-90B outlined in section 3.3.3.

Test System	Entropy of 1,000,000 Traces	Sufficient Entropy
AMD Ryzen 5950X - 64-bit KVM environment	4.531023	Y
AMD EPYC Milan 7713 2 sockets 128 cores 8-way NUMA	7.007947	Y
ARMv7 rev 5	1.9344	Y
ARMv7 rev 5 (Freescale i.MX53)[14]	7.07088	Y
ARMv7 rev 5 (Freescale i.MX6 Ultralite)[15]	6.638399	Y
ARM 64 bit AppliedMicro X-Gene Mustang Board	5.599128	Y
Intel Atom Z530 – using GUI	3.38584	Y
Intel i7 7500U Skylake - 64-bit KVM environment	3.452064	Y
Intel i7 8565U Whiskey Lake – 64-bit KVM environment	7.400136	Y
Intel i7 8565U Whiskey Lake – 32-bit KVM environment	7.405704	Y
Intel i7 8565U Whiskey Lake	6.871	Y
Intel Xeon E7 4870 8 sockets 160 CPUs 8-way NUMA	7.287790	Y
Intel Xeon Gold 6234	4.434168	Y
IBM POWER 8 LE 8286-42A	6.830712	Y
IBM POWER 7 BE 8202-E4C	4.233912	Y
IBM System Z z13 (machine 2964)	4.366368	Y
IBM System Z z15 (machine 8561)	5.691832	Y
MIPS Atheros AR7241 rev 1[16]	7.157064	Y
MIPS Lantiq 34Kc V5.6[17]	7.032740	Y
Qualcomm IPQ4019 ARMv7[18]	6.638405	Y
SiFive HiFive Unmatched RISC-V U74	2.387470	Y

Table 3: ESDM Entropy Testing Results on Different Hardware

Some of the tested systems are quite old or are small embedded devices demonstrating that even on older and smaller systems the ESDM does not overestimate the available entropy when applying worst case conditions.

I am looking for test data from all kinds of systems. The less common a system is the more I am interested in the data to verify that the basic entropy estimate underlying the ESDM is correct. If you want to provide support, please generate data using the ESDM test tool set specifically the test as documented in addon/.

[14] USBArmory MK I
[15] USBArmory MK II
[16] Ubiquiti Nanostation M5 (xm)
[17] AVM Fritz Box 7490
[18] AVM Fritz Box 7520

The effect of the application of the GCD can be clearly demonstrated with the Intel Atom Z530 listed in the above table. The table shows the measurement of without dividing the time stamp by the GCD. The GCD measurement during boot detects that all time stamps have a GCD of 4 which means that the low 2 bits are always unset. Re-running the entropy measurements again on the time stamp that is already divided by 4, the resulting entropy rate is 7.299 bits of entropy per the 8 LSB of the time stamp. This clearly shows that the now considered 2 additional bits of the 8 LSB time stamp after the division with the GCD provides additional entropy which again demonstrates that the ESDM heuristic entropy estimation is safe. It may be noted that by considering 2 additional bits that are now considered for the entropy rate seemingly provide more than 2 bits of entropy (before the GCD, the entropy rate was measured at 3.38 which would imply that by adding 2 bits that may provide full entropy, the rate cannot be higher than 5.38). This seeming inconsistency is due to the fact that a new test run was conducted to get new entropy data. The ping flood used to trigger the IRQ events may have been affected by network congestion adding some delays to the interrupts caused by the ping flood.

D Auxiliary Testing

In addition to the entropy testing additional functional tests are applied using the meson test infrastructure. For details, see the ‘tests‘ directory.

E Bibliographic Reference

References

[1] Y. Nir and A. Langley. ChaCha20 and Poly1305 for IETF Protocols. RFC 7539 (Informational), May 2015. URL http://www.ietf.org/rfc/rfc7539.txt.

[2] Meltem Sönmez Turan, Elaine Barker, John Kelsey, Kerry A. McKay, Mary L. Baish, and Mike Boyle. *NIST Special Publication 800-90B Recommendation for the Entropy Sources Uses for Random Bit Generation.* 2018.

F License

2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

3. The name of the author may not be used to endorse or promote products derived from this software without specific prior written permission.

ALTERNATIVELY, this product may be distributed under the terms of the GNU General Public License, in which case the provisions of the GPL are required INSTEAD OF the above restrictions. (This clause is necessary due to a potential bad interaction between the GPL and the restrictions contained in a BSD-style copyright.)

THIS SOFTWARE IS PROVIDED "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, ALL OF WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF NOT ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

G Change Log

Date	ESDM Version	Change
2022-05.10	v0.2.0	Initial public release
2022-07-02	v.0.4.0	Addition of new /proc interfaces Initial seeding is limited to availability of 256 bits Addition of interupt ES - addition of chapter 3 and appendix C NTG.1 compliance using /dev/random with O_SYNC per-CPU scheduler entropy pool Correct various left-over statements from LRNG
2022-08-02	v0.5.0	Addition of HWRAND entropy source Addition of GRND_RANDOM to getrandom(2) Addition of GRND_SEED to getrandom(2)
UNRELEASED	v0.6.0	Addition of AIS 20/31 (2022) support