

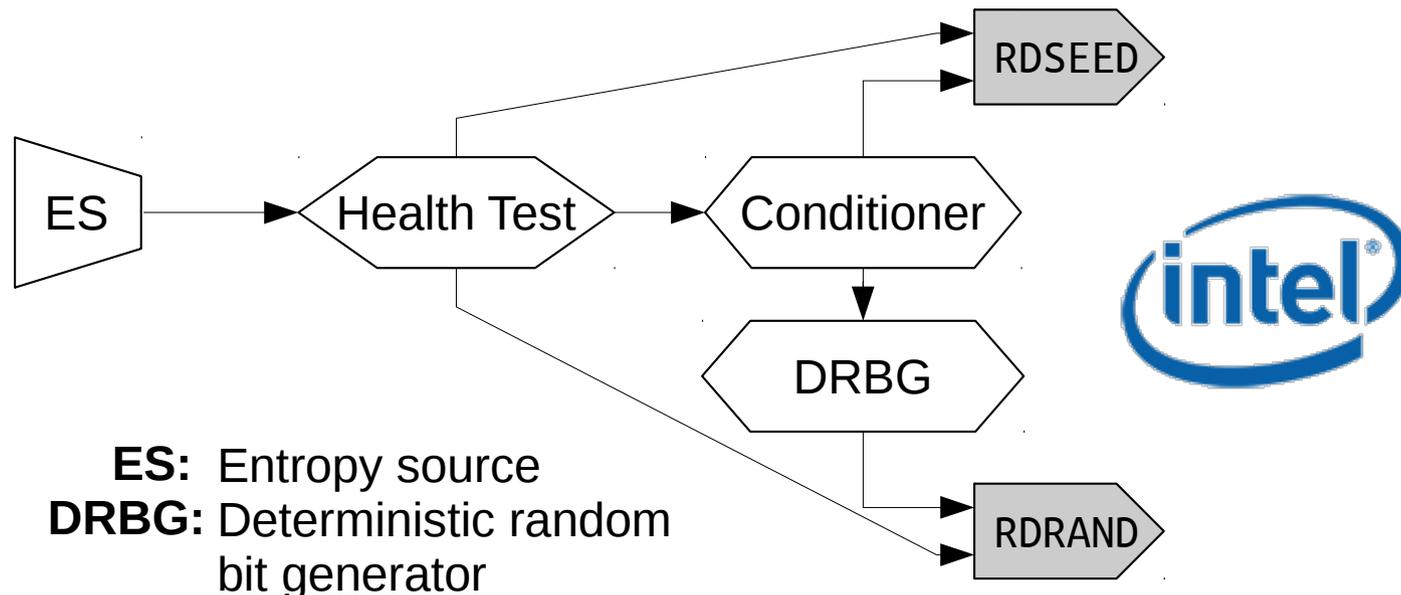
# A Provable Security Analysis of Intel's Secure Key RNG



Thomas Shrimpton & **Seth Terashima**  
Portland State University

# The Intel RNG

- New hardware random-number generator on all recent Intel chips (Ivy Bridge +)
- Two new instructions:
  - **RDRAND**: Fetch pseudo-random bits
  - **RDSEED**: Fetch “truly random” bits (Broadwell +)



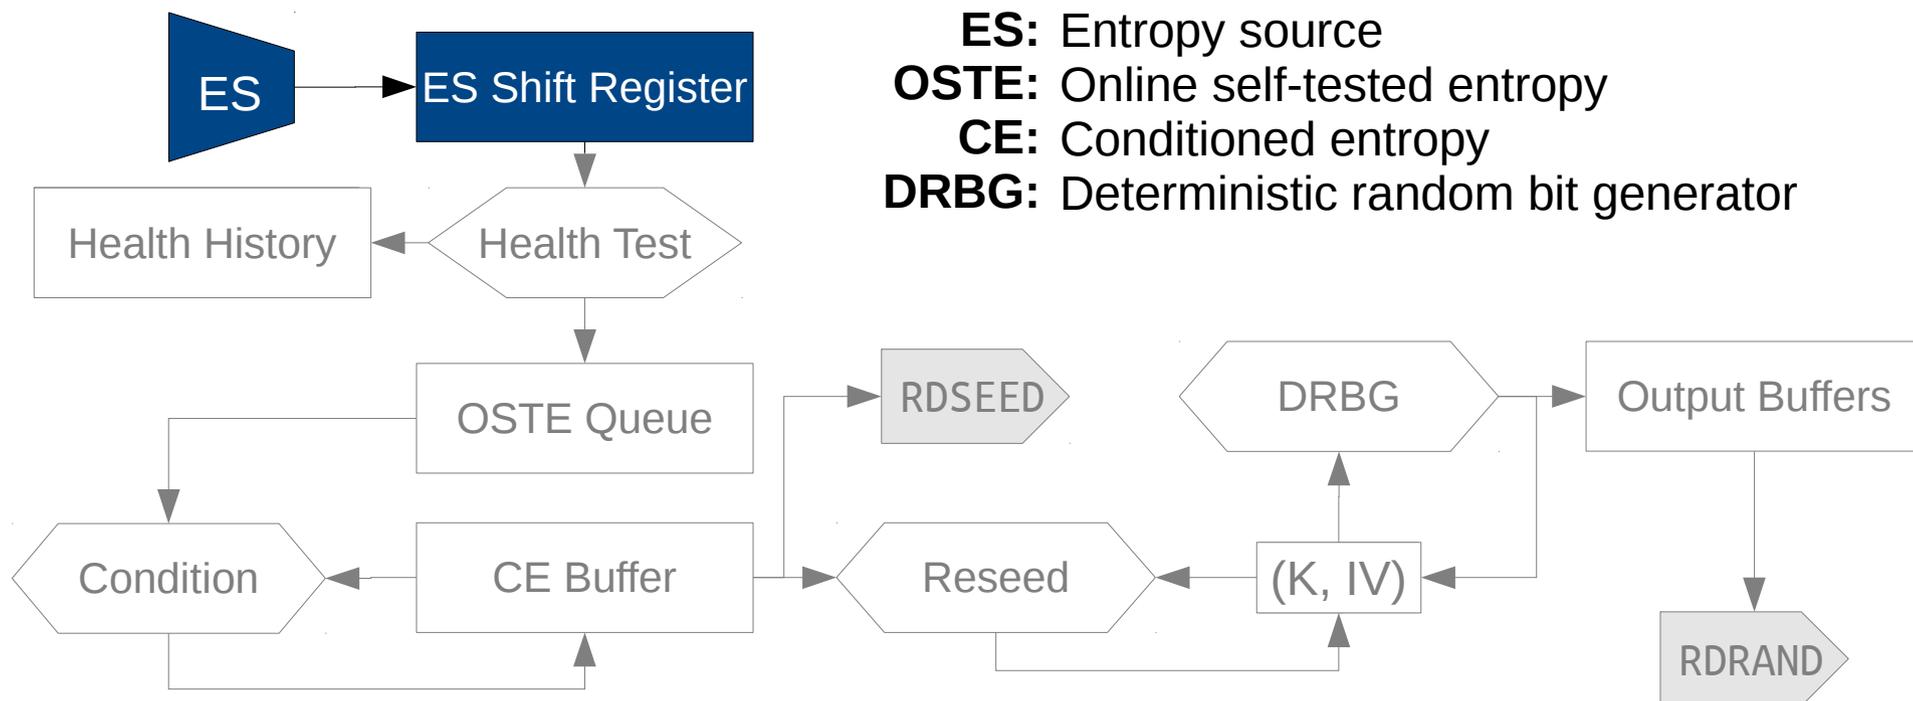
# Agenda

- Intel's Secure-Key RNG design
- The Model : “PRNGs With Input” (PWIs)
- Analysis

Not on the Agenda:

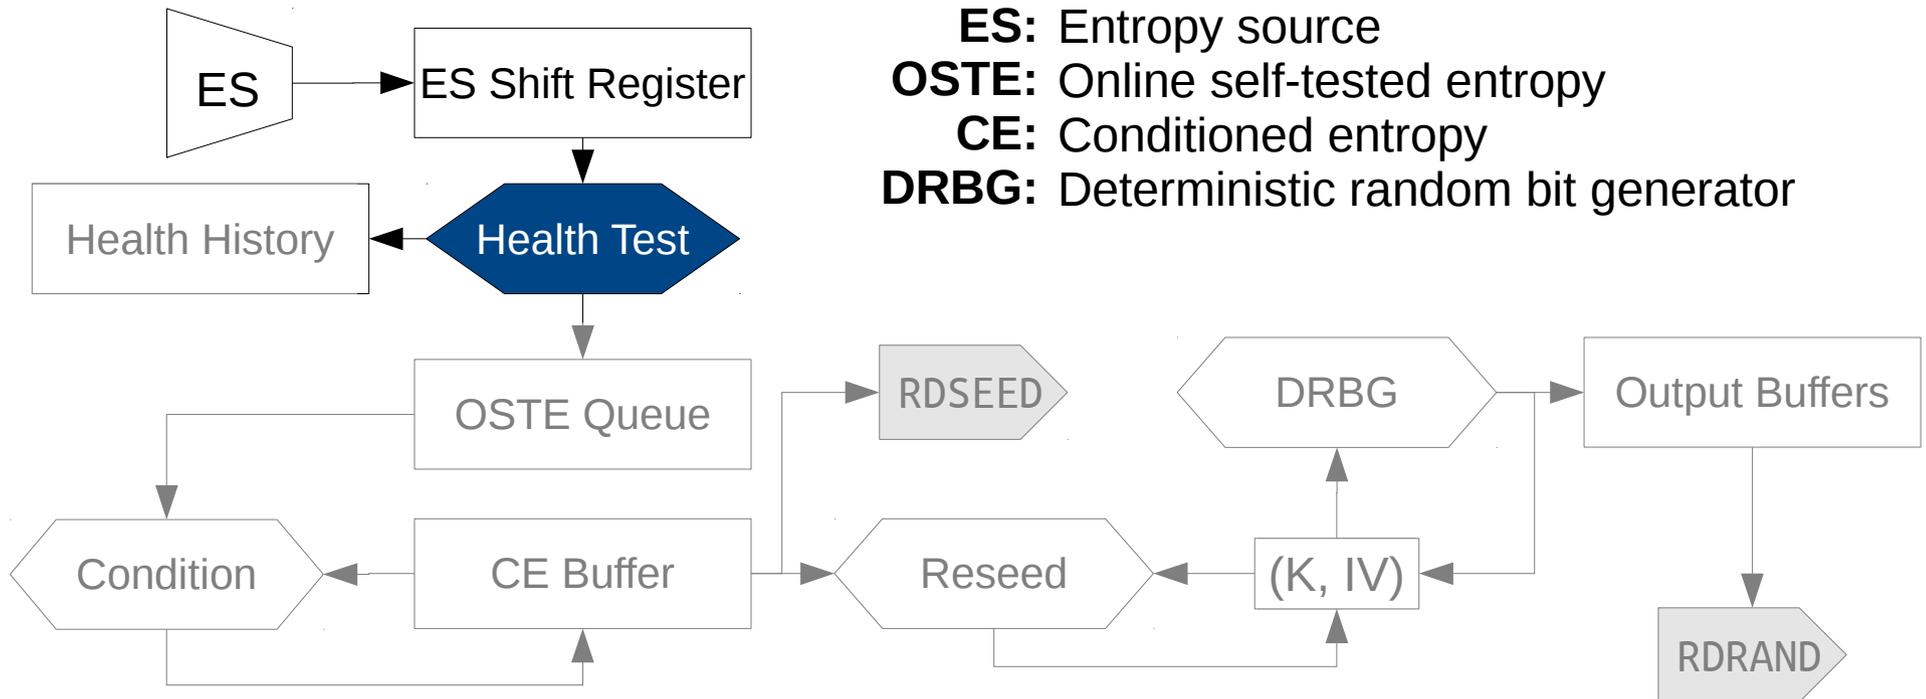


# Entropy Source



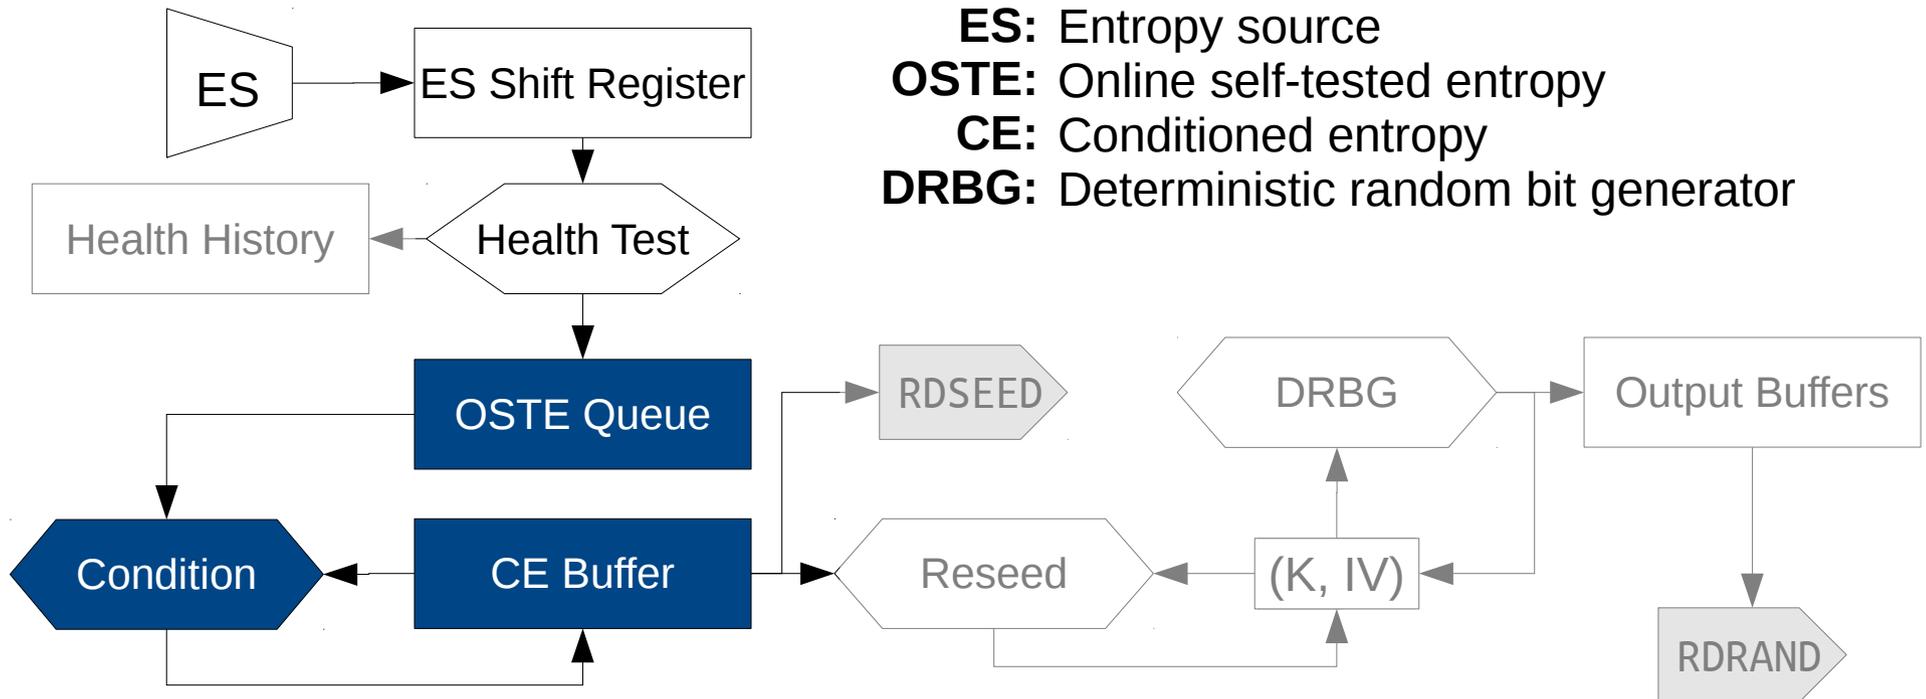
- Uses thermal noise to generate random bits
- Analysis of empirical data by Cryptographic Research, Inc. (Hamburg, Kocher, Marson '12)
- 256-bit samples buffered in shift register

# Health Tests



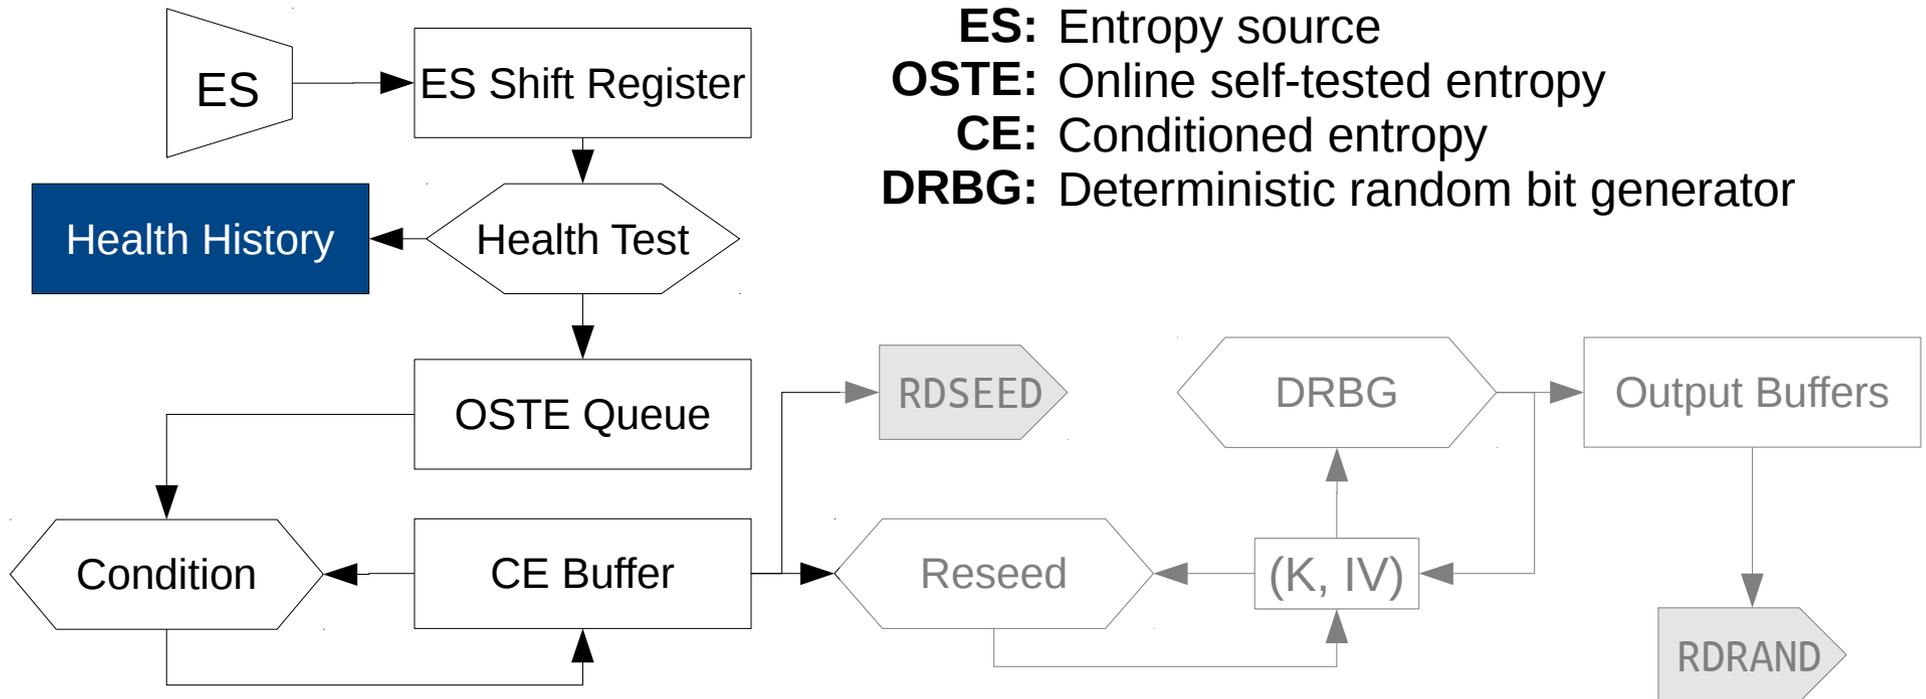
- Heuristic tests for catastrophic ES failure
- 1% false-positive rate on ideal random source

# Conditioning



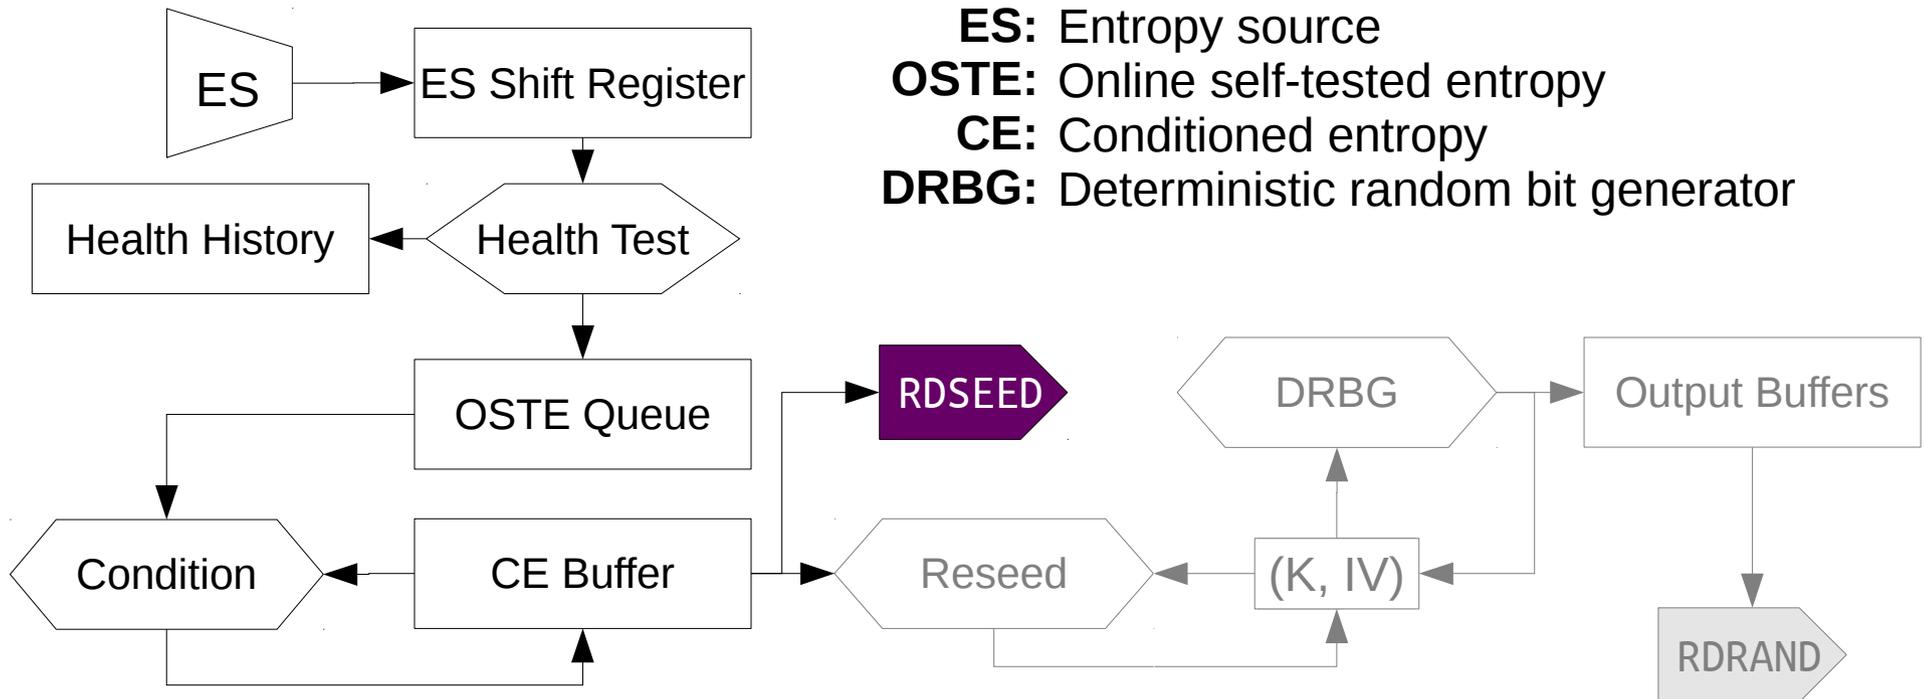
- ES bits are assumed to be biased, correlated
- Fed into streaming CBC-MAC computation to “condition” them into (hopefully) uniform random bits

# Health History



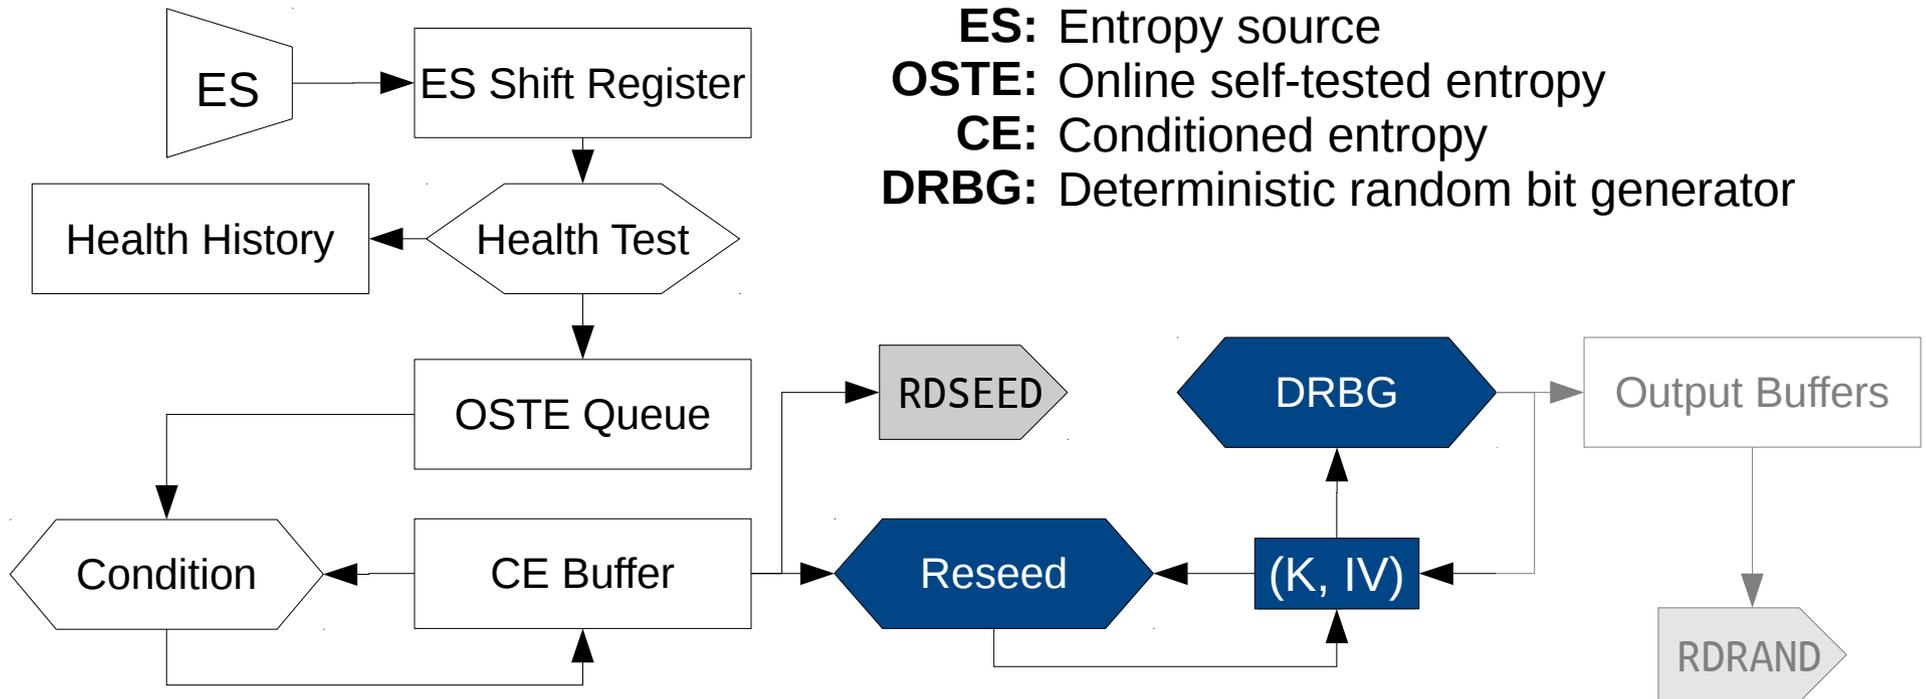
- At least two/three (Ivy Bridge/Broadwell) healthy samples needed before CE buffer is “available”
- But samples don't count unless at least half of the past 256 samples were healthy

# RDSEED



- RDSEED instruction grabs bits from CE buffer
- Buffer not cleared, but flagged as “unavailable”
- Will be made available again after sufficient number of healthy samples generated, conditioned

# (Re)Seeding the DRBG



- CE buffer also used to reseed traditional deterministic PRNG (CTR-AES based)
- Reseeding makes CE buffer unavailable



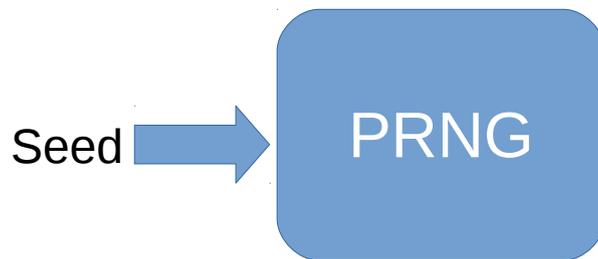
# Agenda

- Intel's Secure-Key RNG design
- **The Model : “PRNGs With Input” (PWIs)**
- Analysis

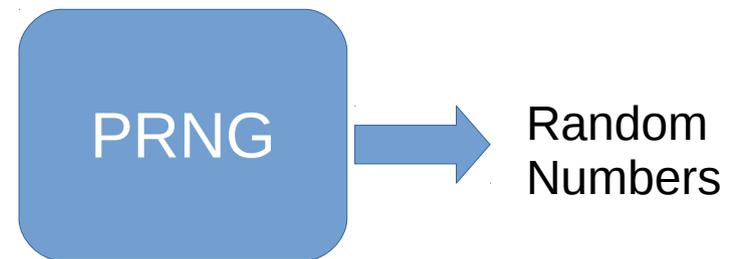


# PRNGs (the traditional view)

**Step 1.** Provide seed.

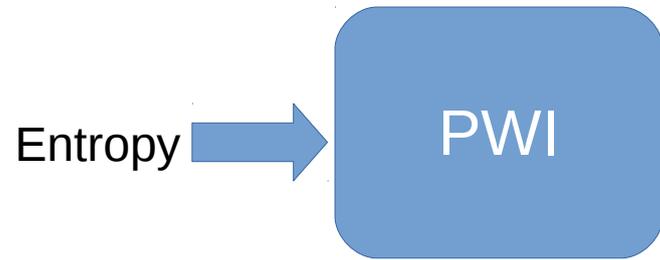


**Step 2.** Get random numbers.

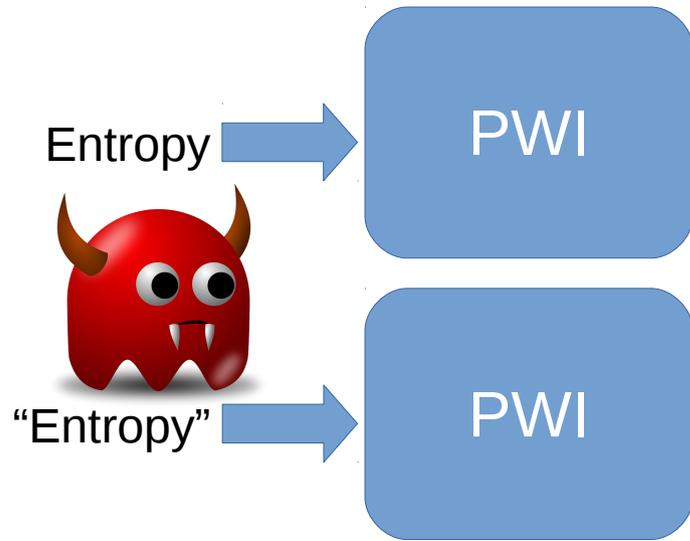


This isn't how  
{/dev/[u]random, OpenSSL RNG, RDRAND}  
work.

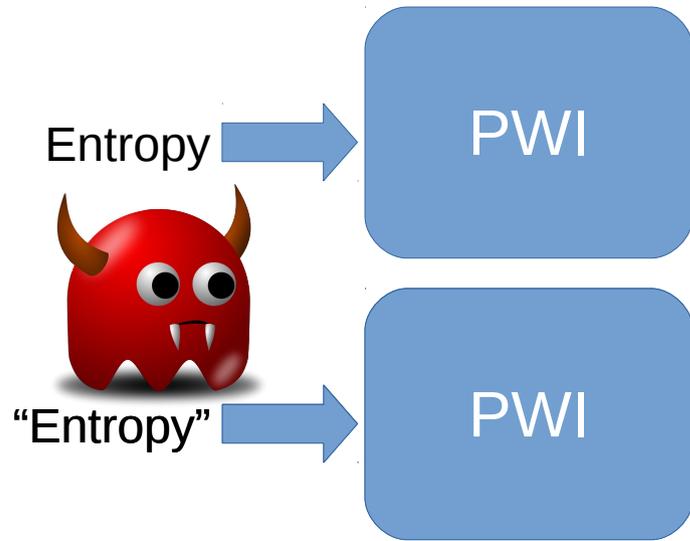
# PRNGs with Input



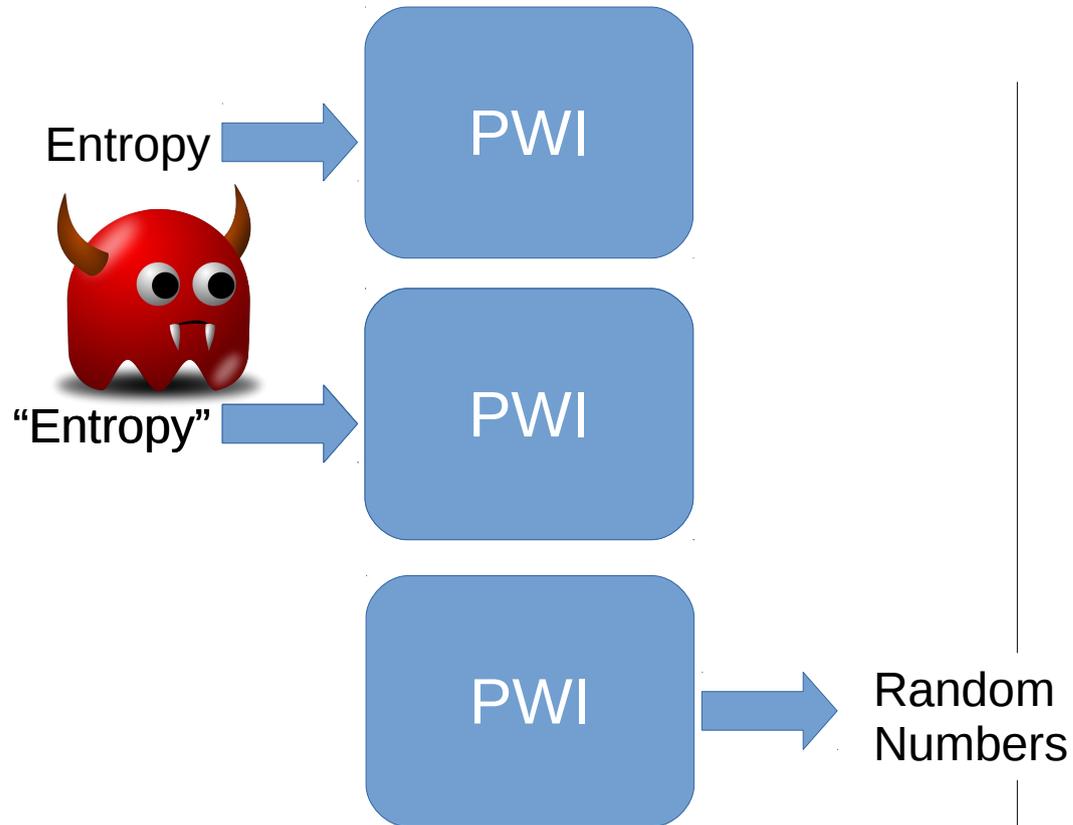
# PRNGs with Input



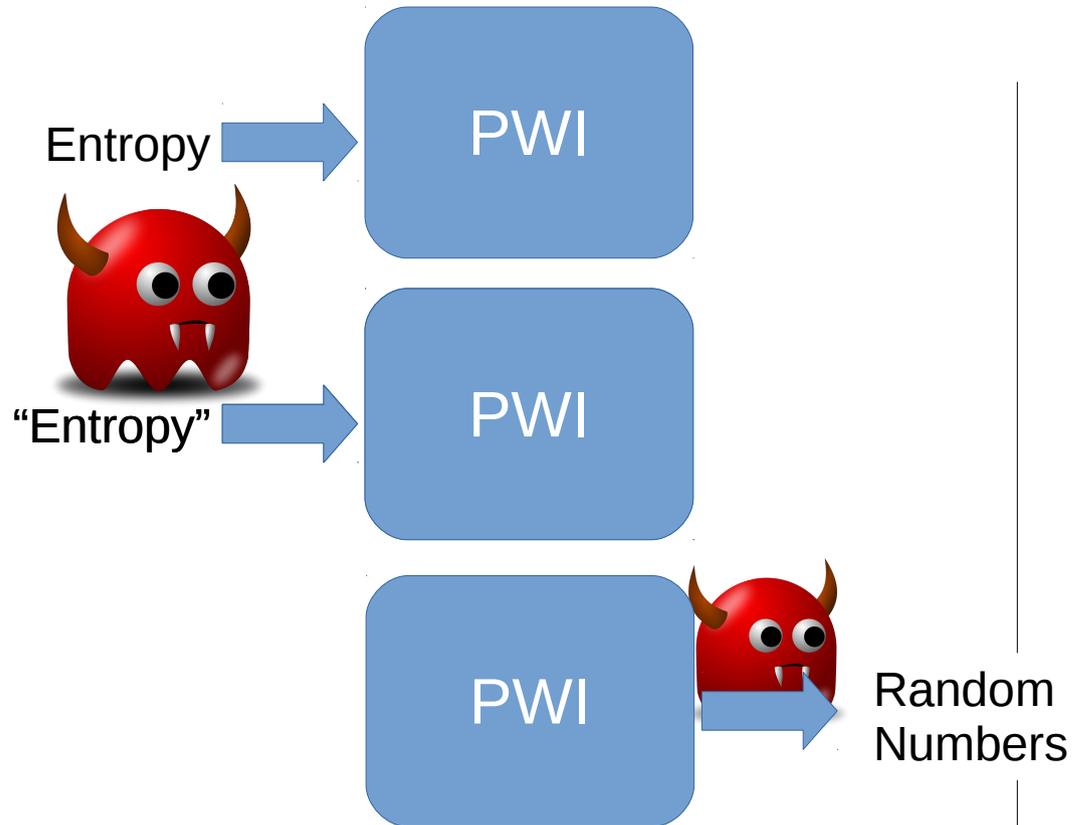
# PRNGs with Input



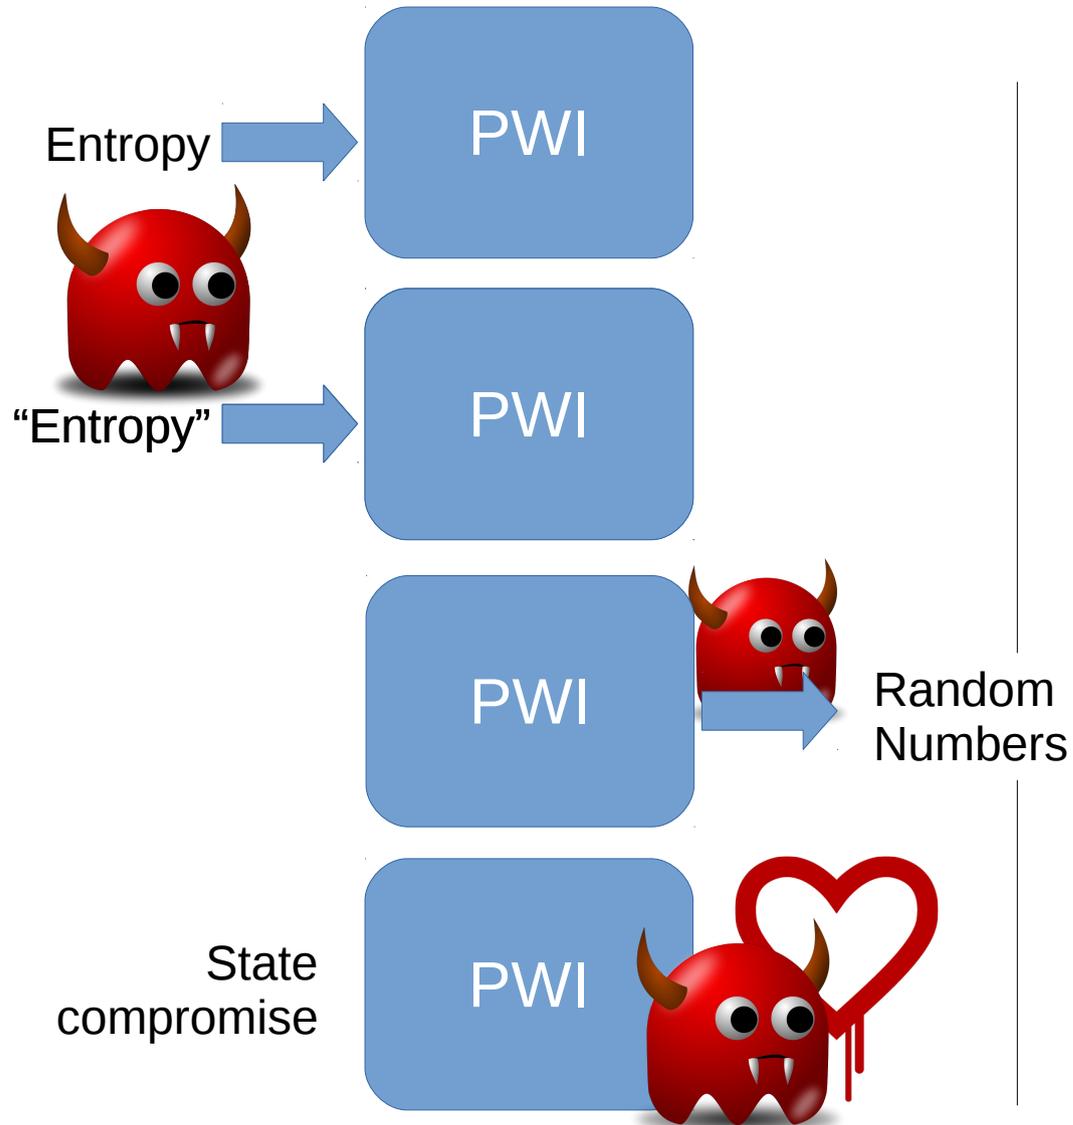
# PRNGs with Input



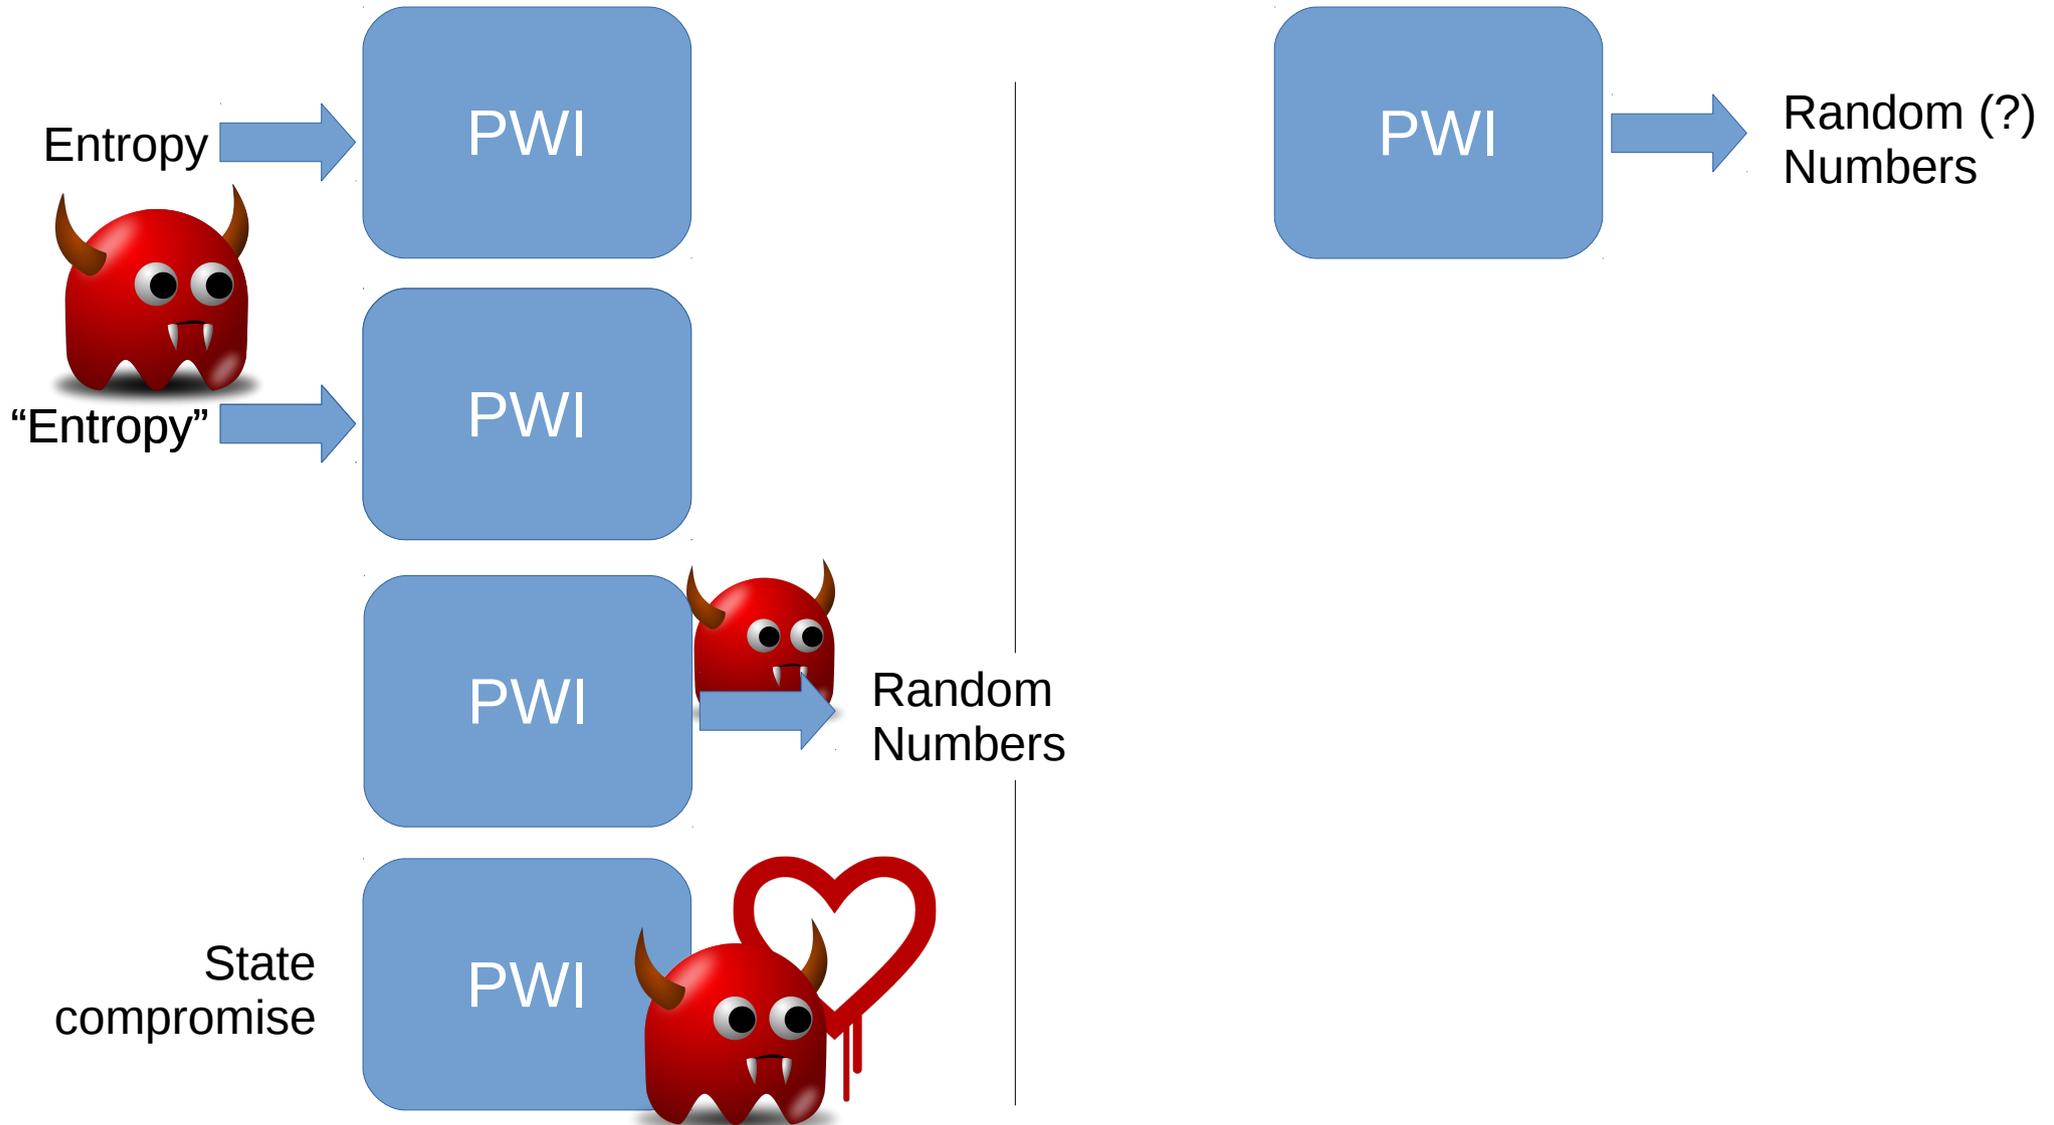
# PRNGs with Input



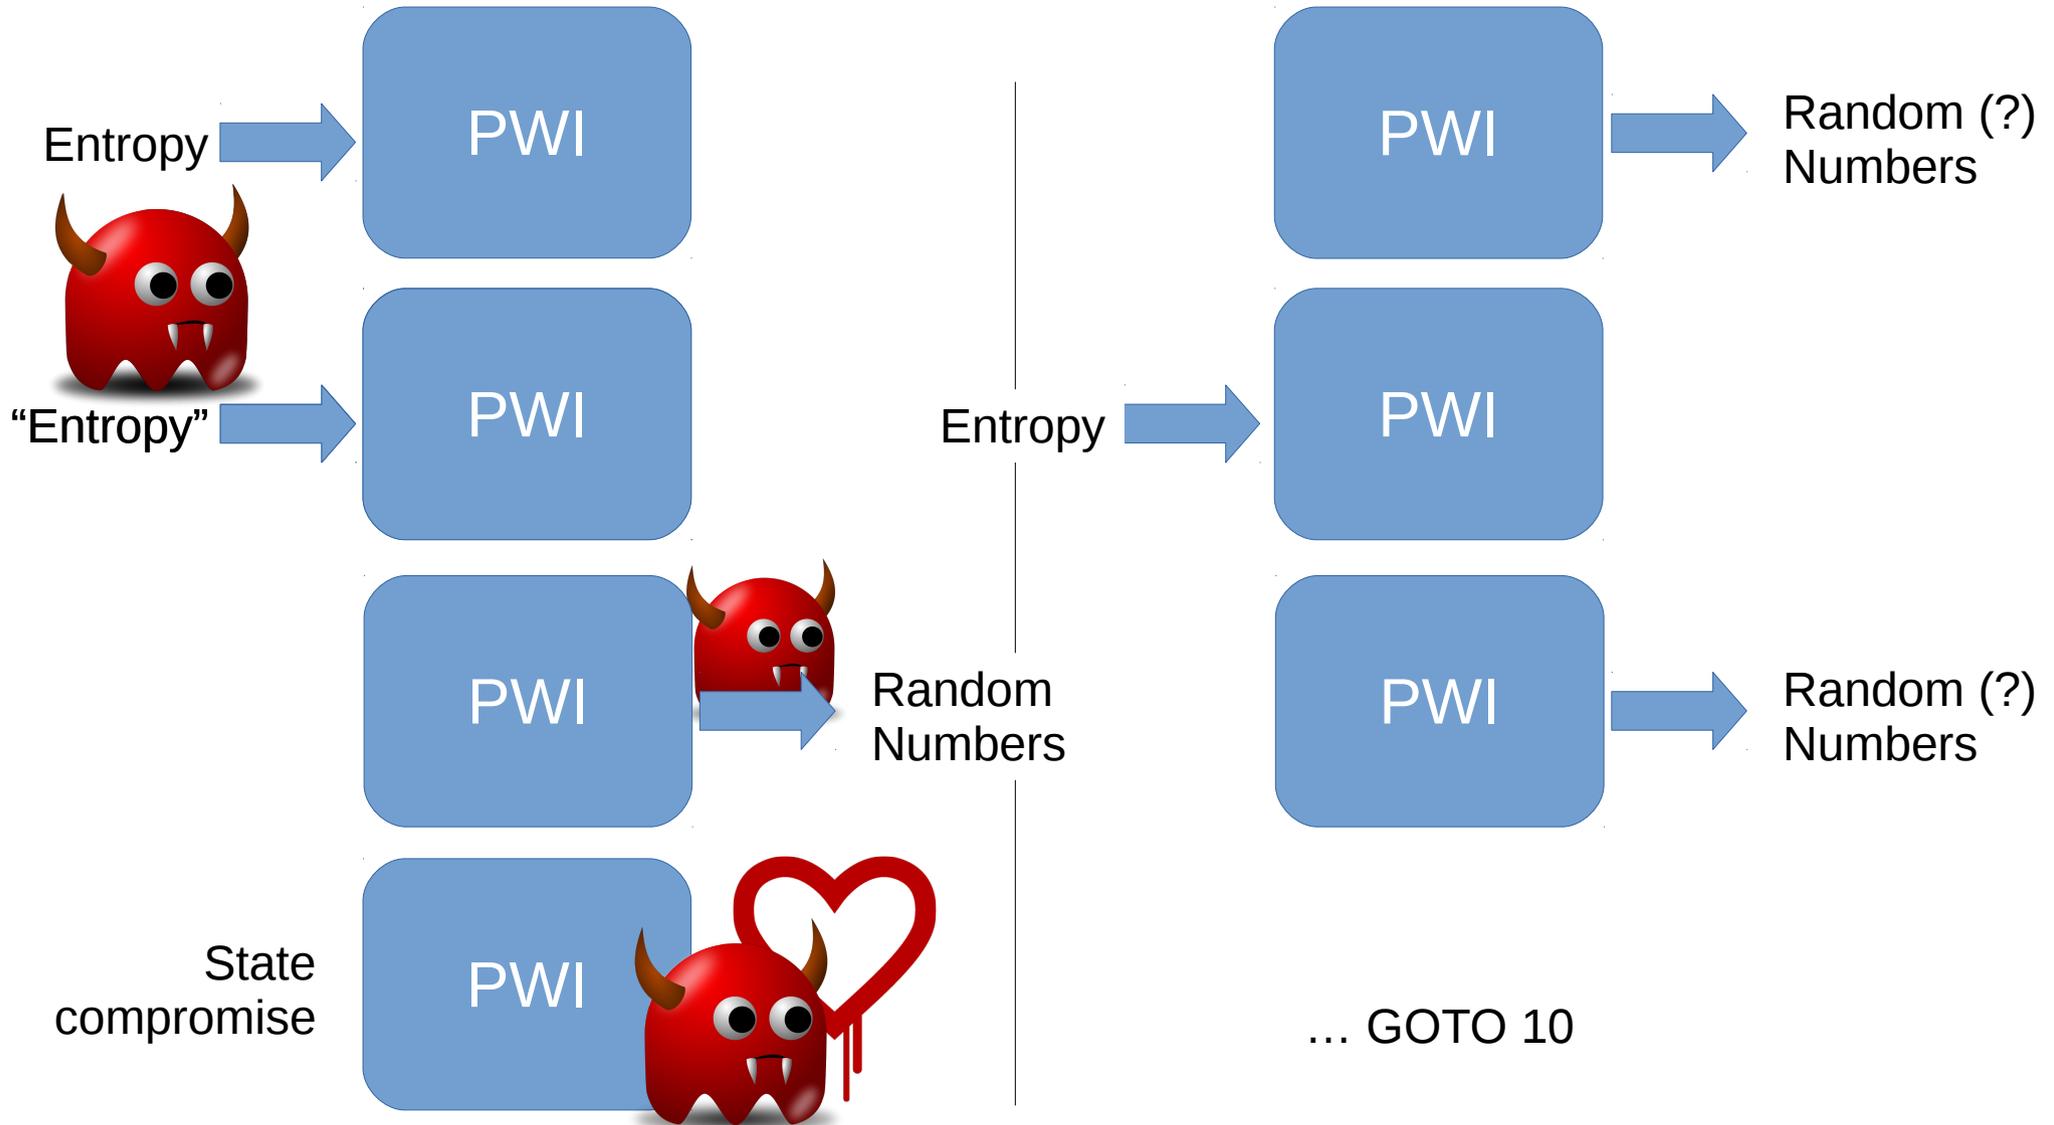
# PRNGs with Input



# PRNGs with Input



# PRNGs with Input



# Types of security



**Resilience:** Basic security. No state compromise, everything looks random



**Forward security:** Random values are safe even if PWI state is compromised in the future.

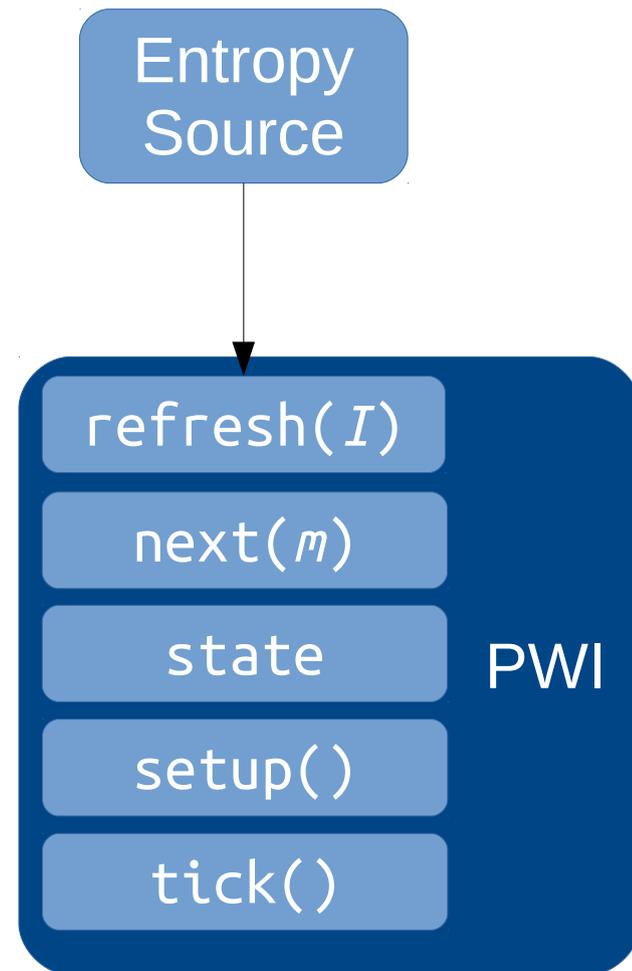


**Backward security:** Random values are safe even if PWI state was compromised in the past (as long as we've harvested enough entropy since then).

**Robustness:** Both forward and backward security, even if the adversary can *tamper* with state.

# PWI Primitive

- Started with PWI model of [DPRVW'13]
- Some simple extensions:
  - Blocking
  - Multiple interfaces
  - Explicit setup
  - Asynchronous behavior





# PWI Oracles

Grabs bits from entropy source, feeds them to PWI. Leaks some side-channel info to attacker.



Entropy Source

ES-Refresh()

get-next()

next-ror()

get-state()

set-state()

wait()

PWI

- refresh( $I$ )
- next( $m$ )
- state
- setup()
- tick()

$b = ?$



# PWI Oracles

Calls next(). Returns the result.



Entropy Source

ES-Refresh()

get-next()

next-ror()

get-state()

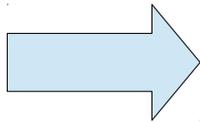
set-state()

wait()

PWI

- refresh( $I$ )
- next( $m$ )
- state
- setup()
- tick()

$b = ?$



# PWI Oracles

Calls next().

- Returns result if  $b = 0$ .
- Returns random string if  $b = 1$ .



Entropy Source

ES-Refresh()

get-next()

next-ror()

get-state()

set-state()

wait()

refresh( $I$ )

next( $m$ )

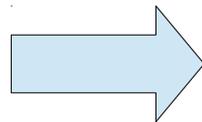
state

setup()

tick()

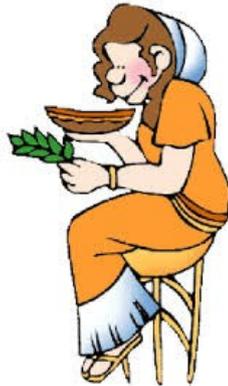
PWI

$b = ?$



# PWI Oracles

Returns the PWI state.



Entropy Source

ES-Refresh()

get-next()

next-ror()



get-state()

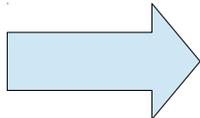
set-state()

wait()

PWI

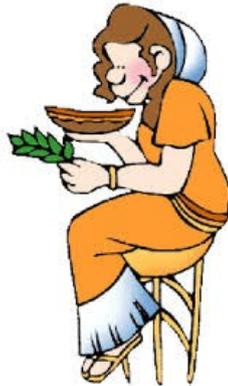
- refresh( $I$ )
- next( $m$ )
- state
- setup()
- tick()

$b = ?$



# PWI Oracles

Sets the PWI state to a value specified by the Adversary.



Entropy Source

ES-Refresh()

get-next()

next-ror()

get-state()

set-state()

wait()

refresh( $I$ )

next( $m$ )

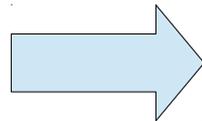
state

setup()

tick()

PWI

$b = ?$



# PWI Oracles

Prompts the PWI to perform next scheduled atomic task. No return value.



Entropy Source

ES-Refresh()

get-next()

next-ror()

get-state()

set-state()

wait()

refresh( $I$ )

next( $m$ )

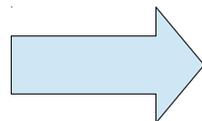
state

setup()

tick()

PWI

$b = ?$



# PWI Oracles

The get-state and set-state oracles make the state “corrupt”.



Entropy Source

ES-Refresh()

get-next()

next-ror()

get-state()

set-state()

wait()

refresh( $I$ )

next( $m$ )

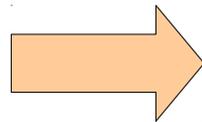
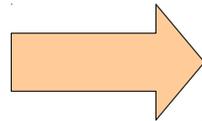
state

setup()

tick()

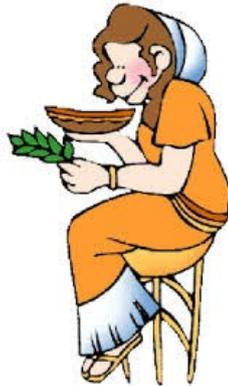
PWI

$b = ?$



# PWI Oracles

While the state is corrupt, the Adversary is cut off from the next-ror oracle.



Entropy Source

ES-Refresh()

get-next()

next-ror()

get-state()

set-state()

wait()

refresh( $I$ )

next( $m$ )

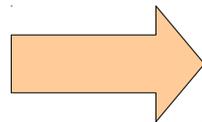
state

setup()

tick()

PWI

$b = ?$



# PWI Oracles

The state remains corrupt until the PWI harvests a specified amount of entropy.



Entropy Source

ES-Refresh()

get-next()

next-ror()

get-state()

set-state()

wait()

PWI

- refresh( $I$ )
- next( $m$ )
- state
- setup()
- tick()

$b = ?$



# Measuring security

Attacker's advantage is:

$$\mathbf{Adv}_{\text{PWI}}^{\text{xxx}}(A) = |\Pr [ A \Rightarrow 1 \mid b = 1 ] - \Pr [ A \Rightarrow 1 \mid b = 0 ]|$$



$$\mathbf{Adv}_{\text{PWI}}^{\text{RES}}(A), \mathbf{Adv}_{\text{PWI}}^{\text{FWD}}(A), \mathbf{Adv}_{\text{PWI}}^{\text{BWD}}(A), \mathbf{Adv}_{\text{PWI}}^{\text{ROB}}(A)$$

# Secure setup()

- setup() should place the PWI in a “good” state
  - Some state is sensitive
  - Other state is not (counters, buffered entropy, etc.)
- Define a **masking function**  $M$  such that  $M(S)$  is a “good version” of  $S$ .

State produced by setup()

$$\mathbf{Adv}_M^{\text{init}}(A) = |\Pr [ A(S_0) \Rightarrow 1 ] - \Pr [ A(M(S_0)) \Rightarrow 1 ]|$$

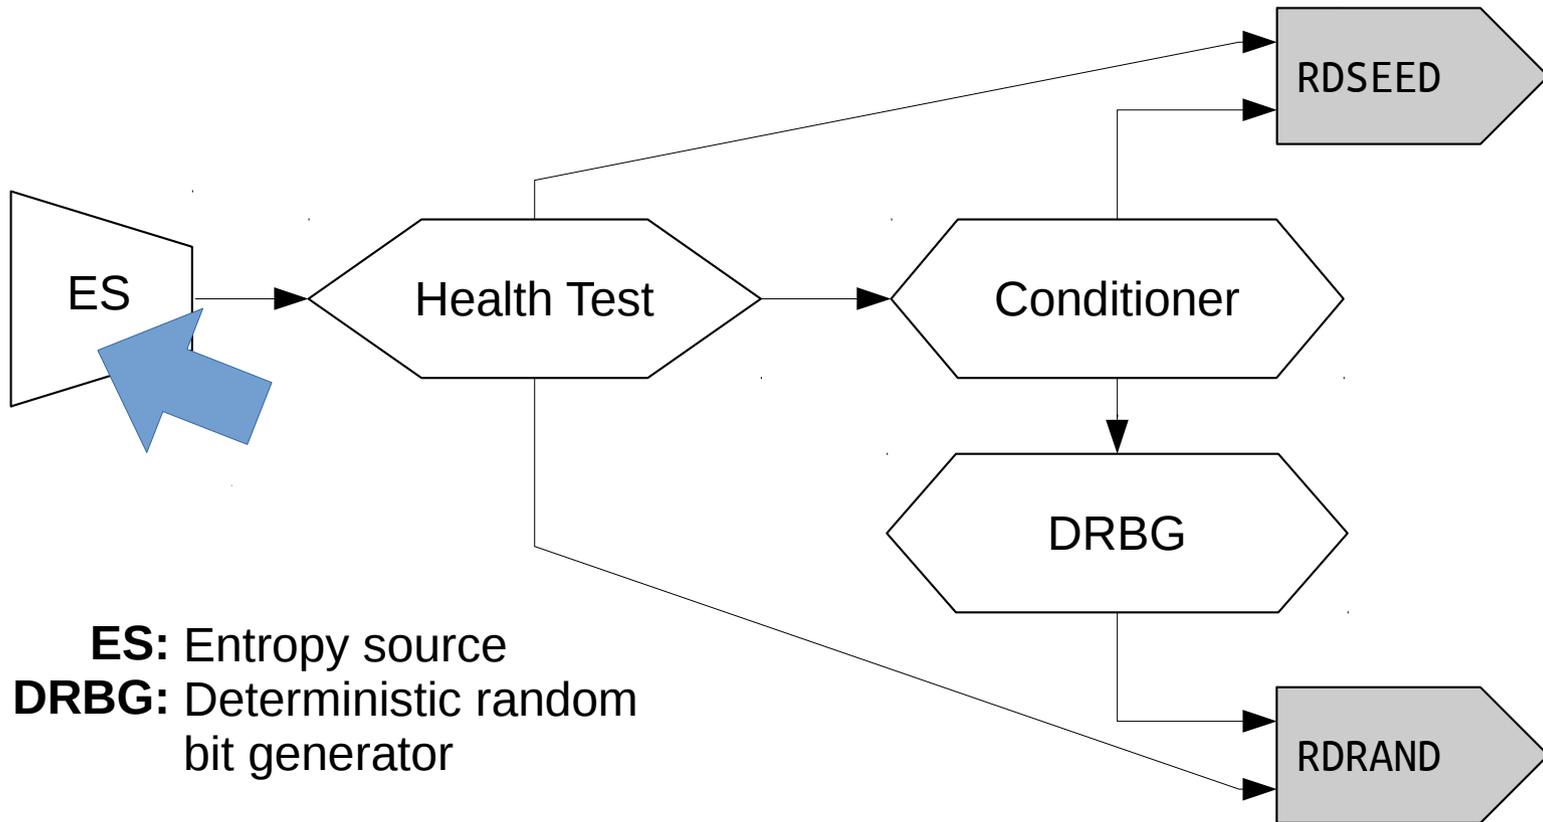
$$\mathbf{Adv}_{\text{PWI}}^{\text{FWD}}(A) \leq \mathbf{Adv}_M^{\text{init}}(B) + \mathbf{Adv}_{\text{PWI}}^{\text{FWD}/M}(A)$$

Forward security when starting from masked state

# Agenda

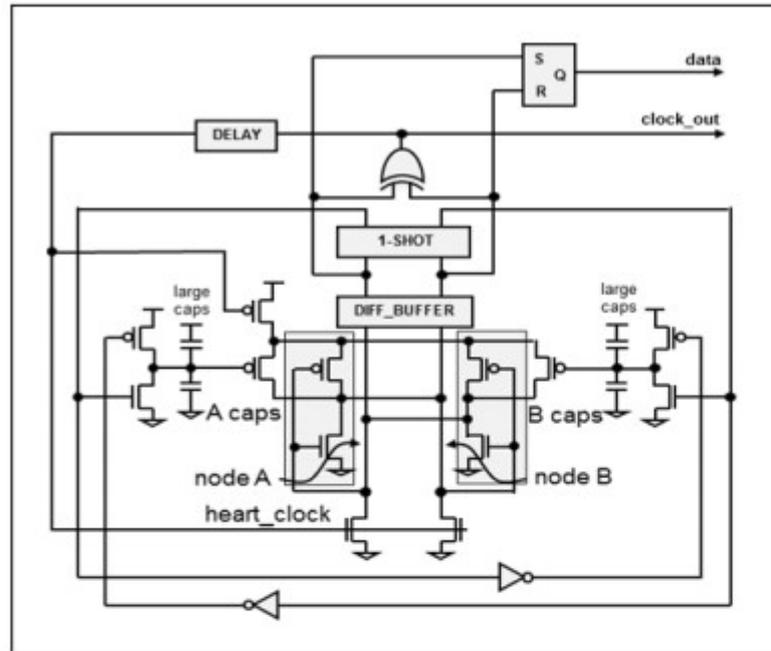
- Intel's Secure-Key RNG design
- The Model : “PRNGs With Input” (PWIs)
- **Analysyis**





**ES:** Entropy source  
**DRBG:** Deterministic random bit generator

# Entropy Source

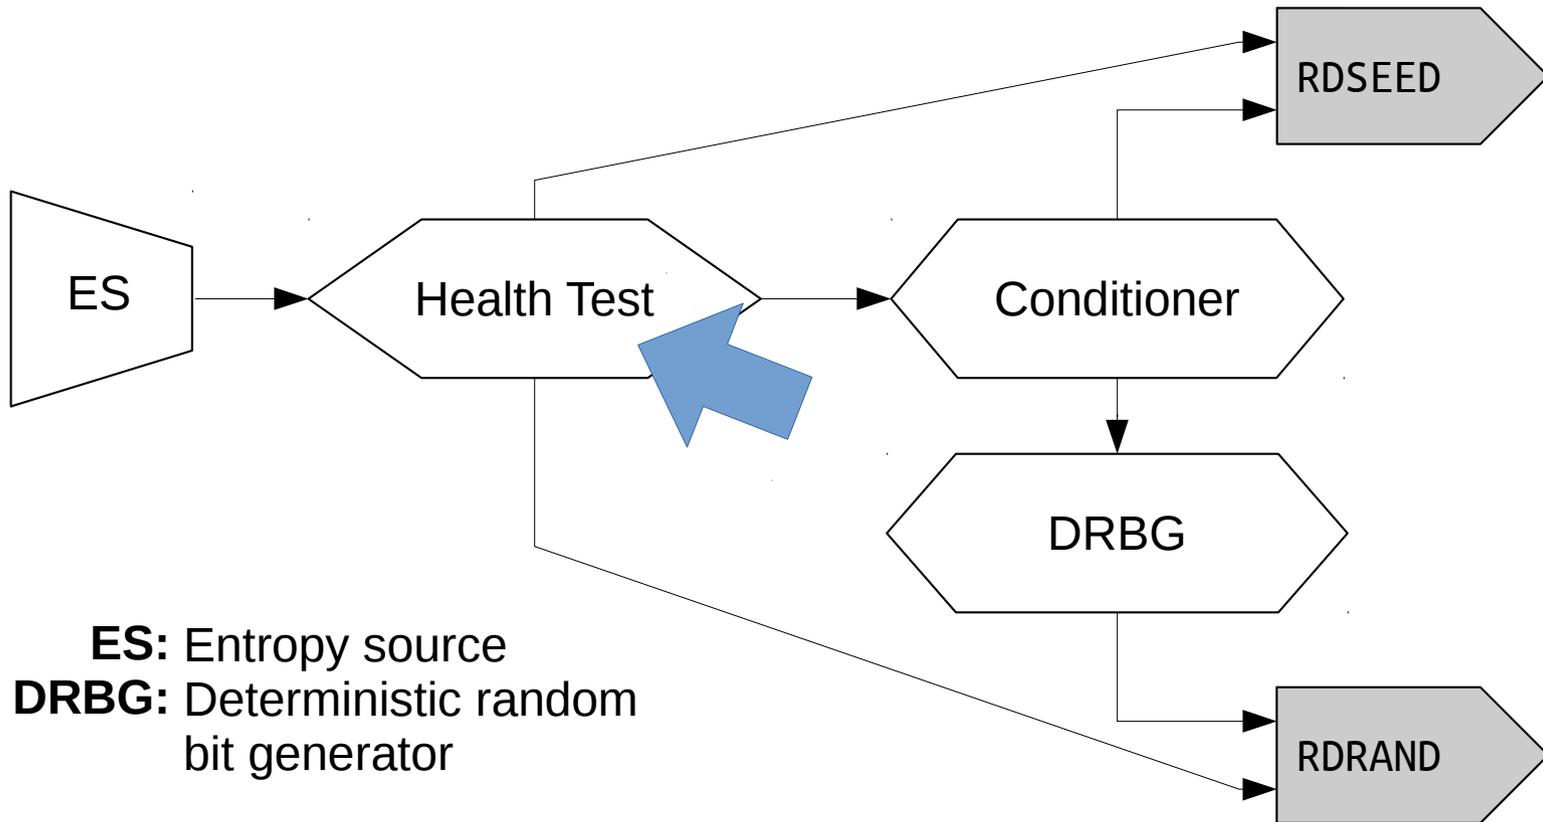


**Figure 2: Entropy source for the Intel RNG (from [8])**

The ES is a dual differential jamb latch with feedback. It is a latch formed by two cross-coupled inverters (nodes A and B). The circuit is self-clocking (heart\_clock), and designed such that when the clock is running, the circuit enters a metastable state. The circuit then resolves to one of two possible states, determined randomly by thermal noise in the system. The settling of the circuit is biased by the differential in the charges on the capacitors (A caps and B caps). The state to which the latch resolves is the random bit of output of the ES.

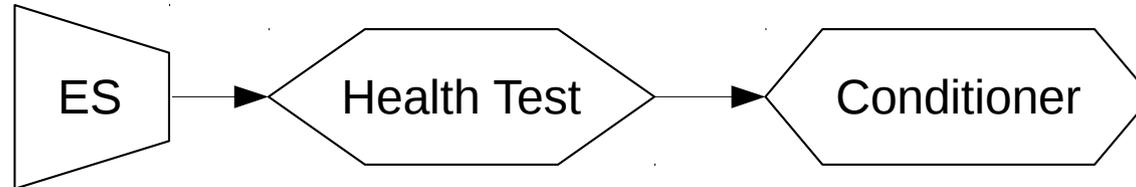
The circuit is also designed with feedback to seek out its metastable region. Based on how the latch resolves, a fixed amount of charge is drained from one capacitor and added



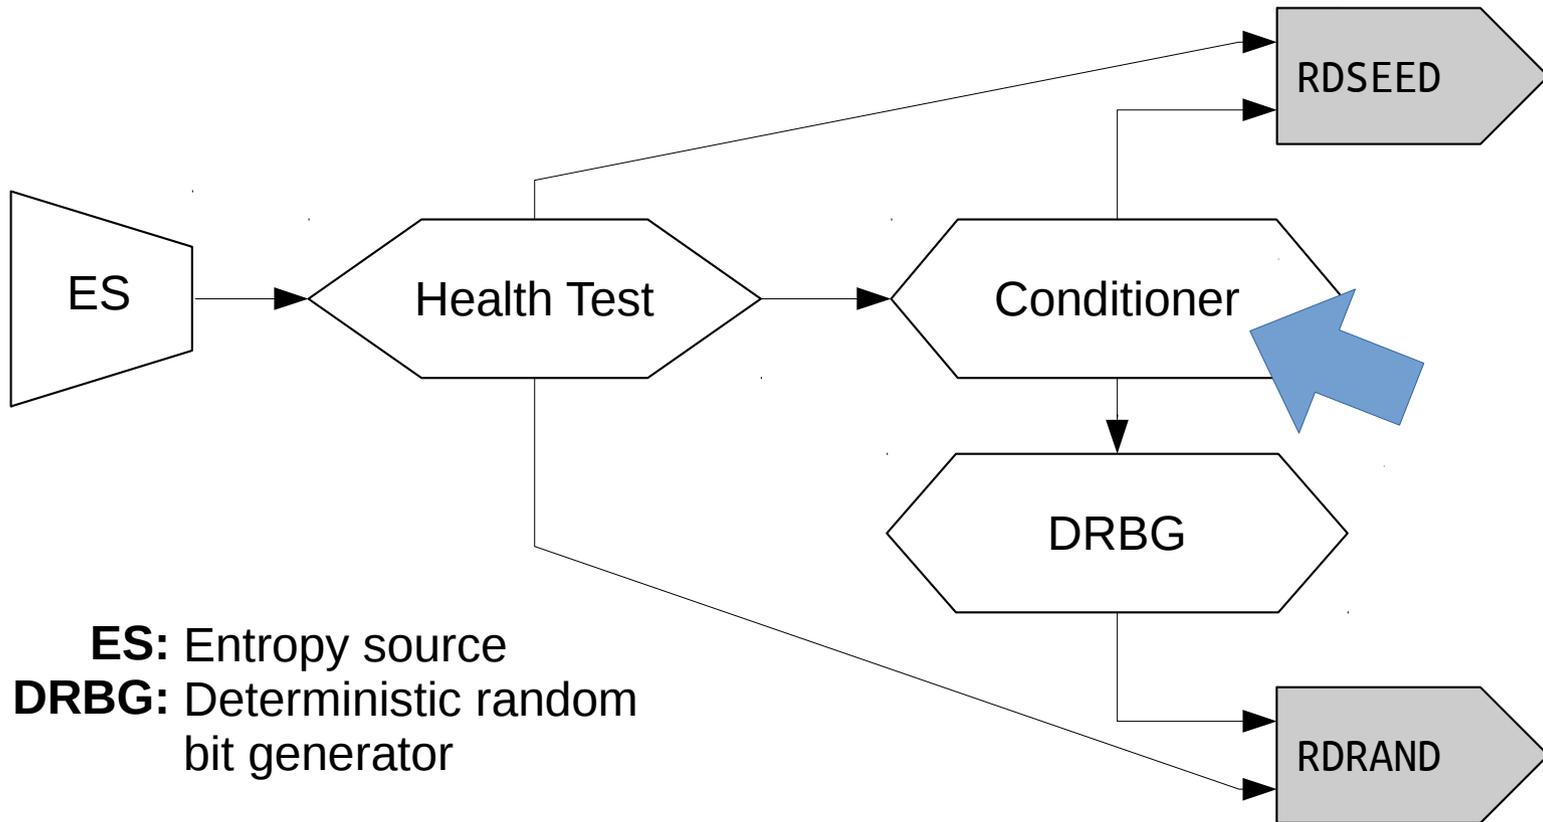


**ES:** Entropy source  
**DRBG:** Deterministic random  
bit generator

# Entropy source assumptions

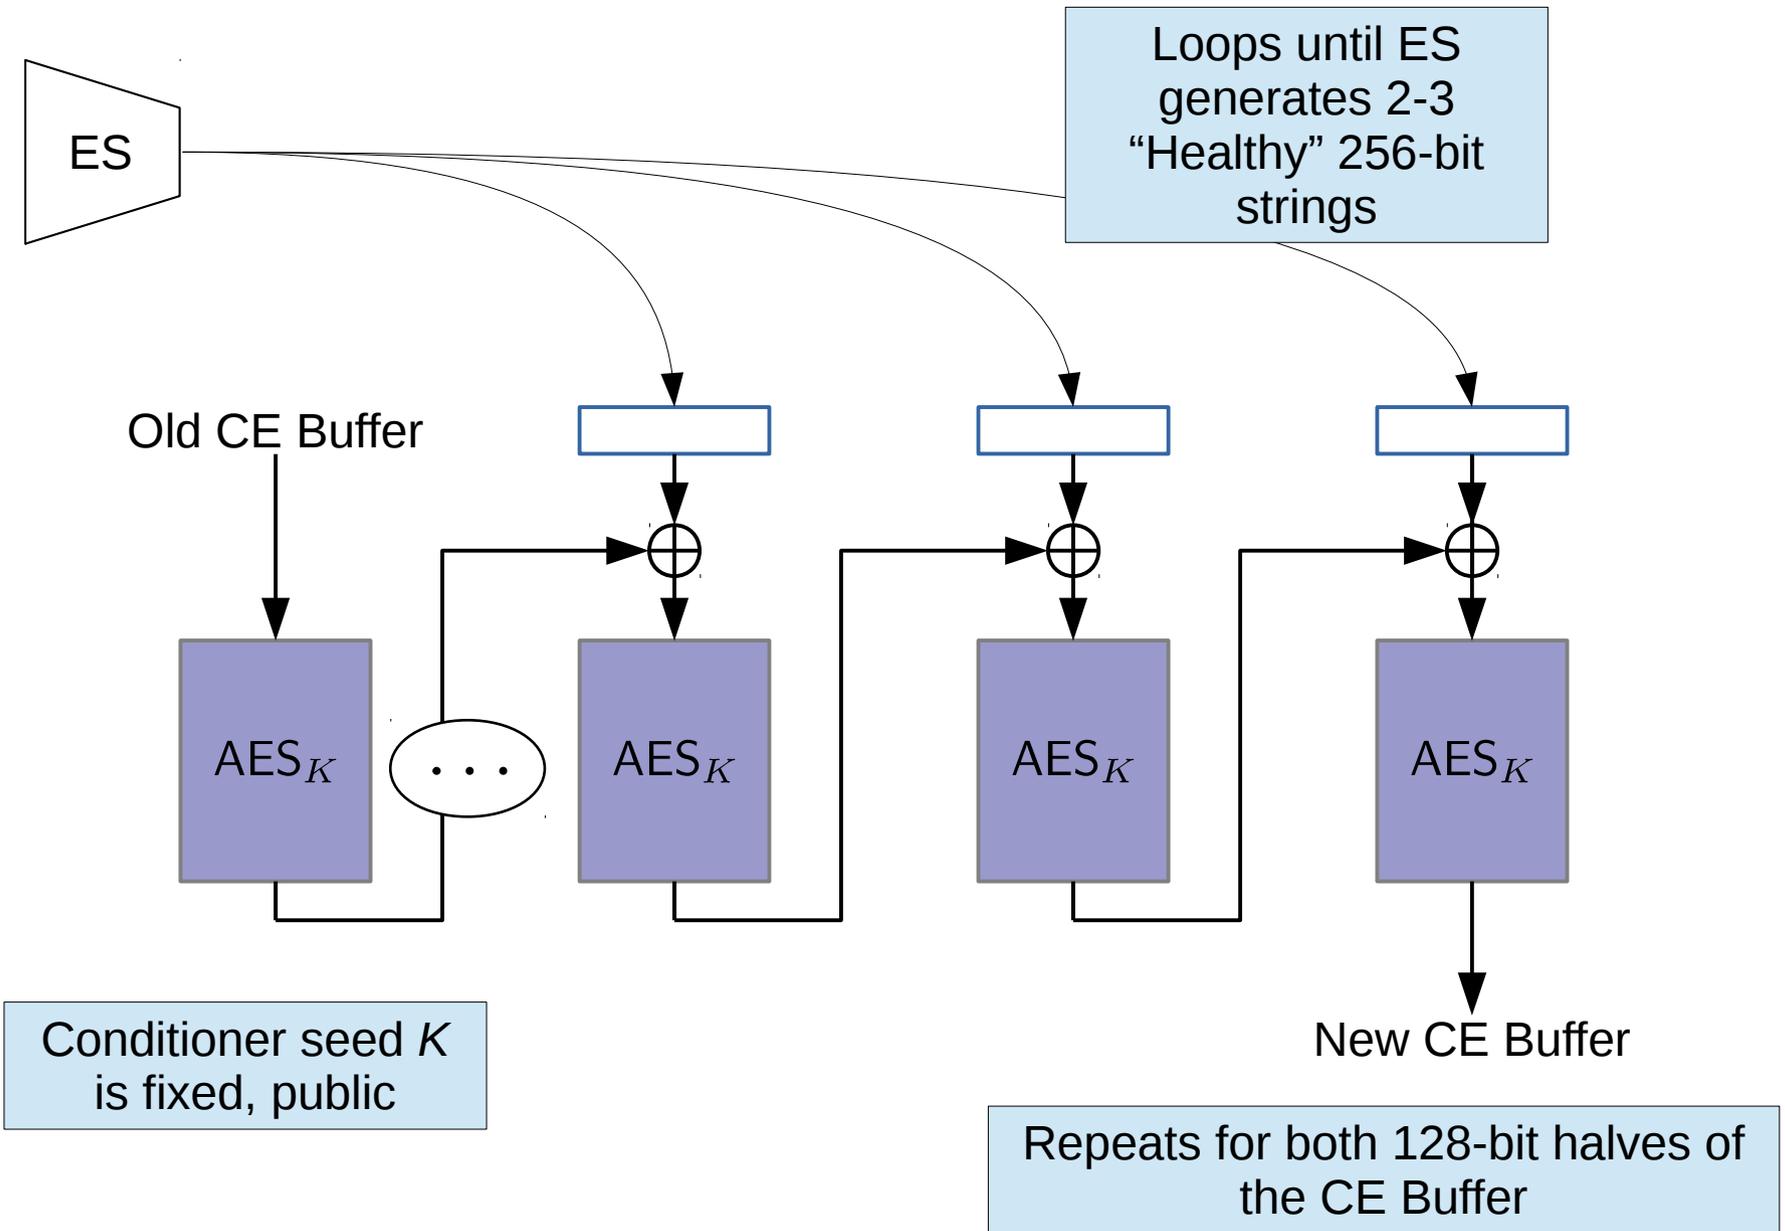


- Assume “healthy” samples have min-entropy  $\gamma$ 
  - Will estimate value from CRI analysis
  - (Conservatively) assume no entropy from “unhealthy” samples
- Assume ES will eventually produce a healthy sample
  - Perfect entropy source = 1% of samples unhealthy
  - Say any 256-bit sample is healthy with probability  $\geq \beta$



**ES:** Entropy source  
**DRBG:** Deterministic random  
bit generator

# Conditioner: CBC-MAC



# Does this work?

Theorem from [DGHKR Crypto '04] says CBC-MAC (over a random permutation) works as an entropy extractor **but**

- Intel RNG recycles state
- Bound degrades quickly with input length

# Does this work?

Theorem from [DGHKR Crypto '04] says CBC-MAC (over a random permutation) works as an entropy extractor **but**

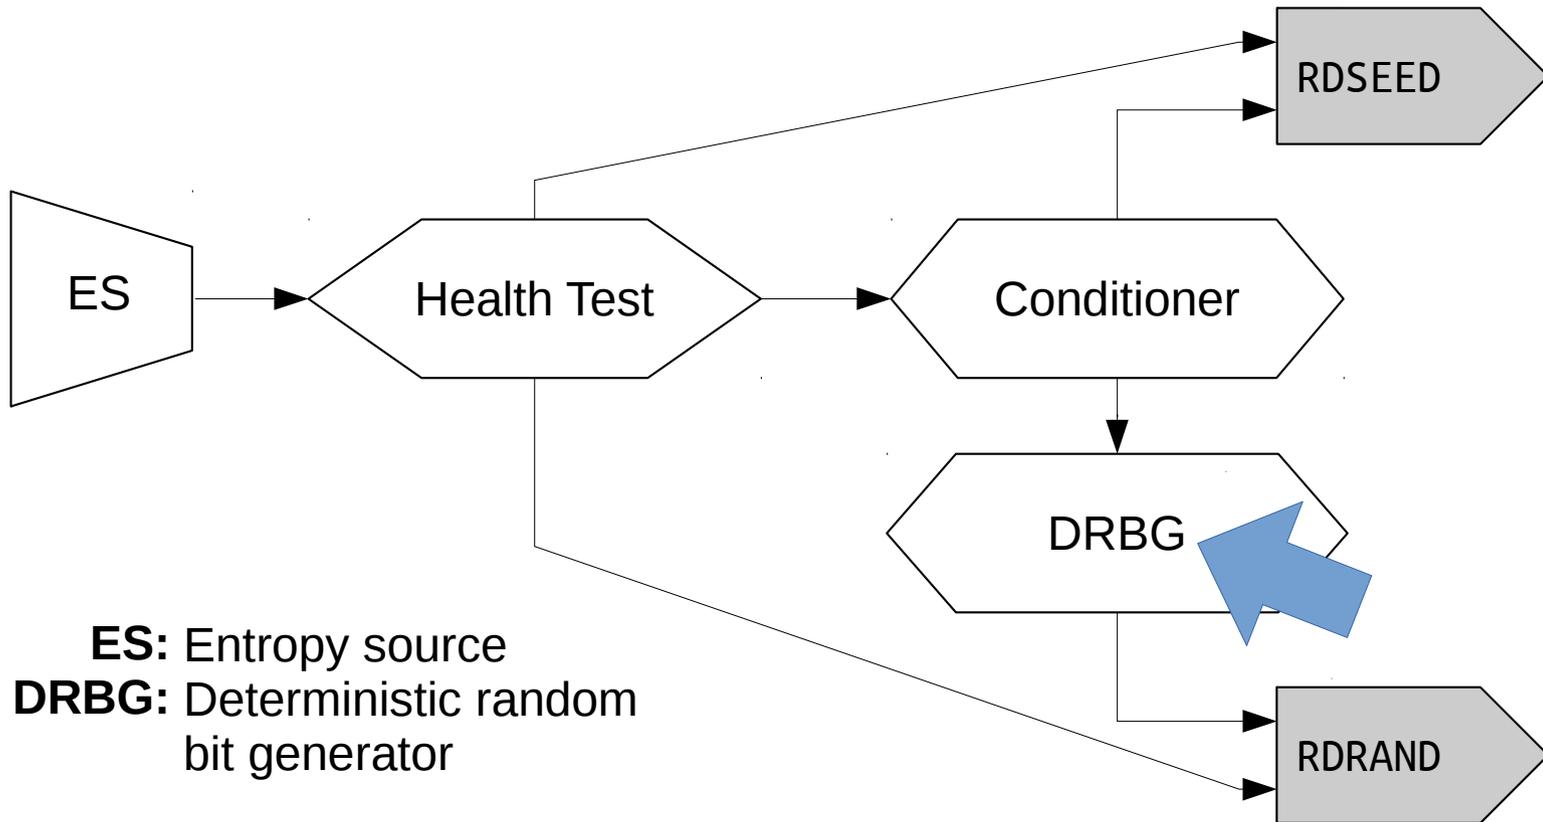
- Intel RNG recycles state Not too hard to patch
- Bound degrades quickly with input length

# Does this work?

Theorem from [DGHKR Crypto '04] says CBC-MAC (over a random permutation) works as an entropy extractor **but**

- Intel RNG recycles state Not too hard to patch
- Bound degrades quickly with input length

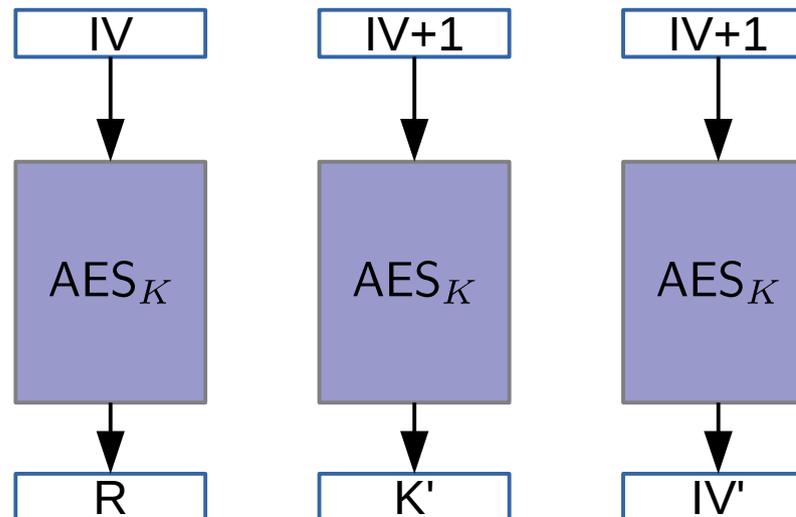
We can impose a fixed “cut-off” point --- don't count entropy or input length past this point.



**ES:** Entropy source  
**DRBG:** Deterministic random bit generator

# DRBG: CTR-AES

- PWI state includes a CTR key and IV
- DRBG operation:
  - Compute  $R, K', IV'$  as below
  - Reassign  $(K, IV) \leftarrow (K', IV')$  Helps with forward security.
  - Return  $R$  as the DRBG output



$$\epsilon(L_m) = \mathcal{O}(L_m + 1)/2^{k/2}$$

$$\hat{\epsilon}(L_m) = \sum_{i=0}^{m-1} \binom{L_m}{i} \beta^i (1 - \beta)^{L_m - i}$$

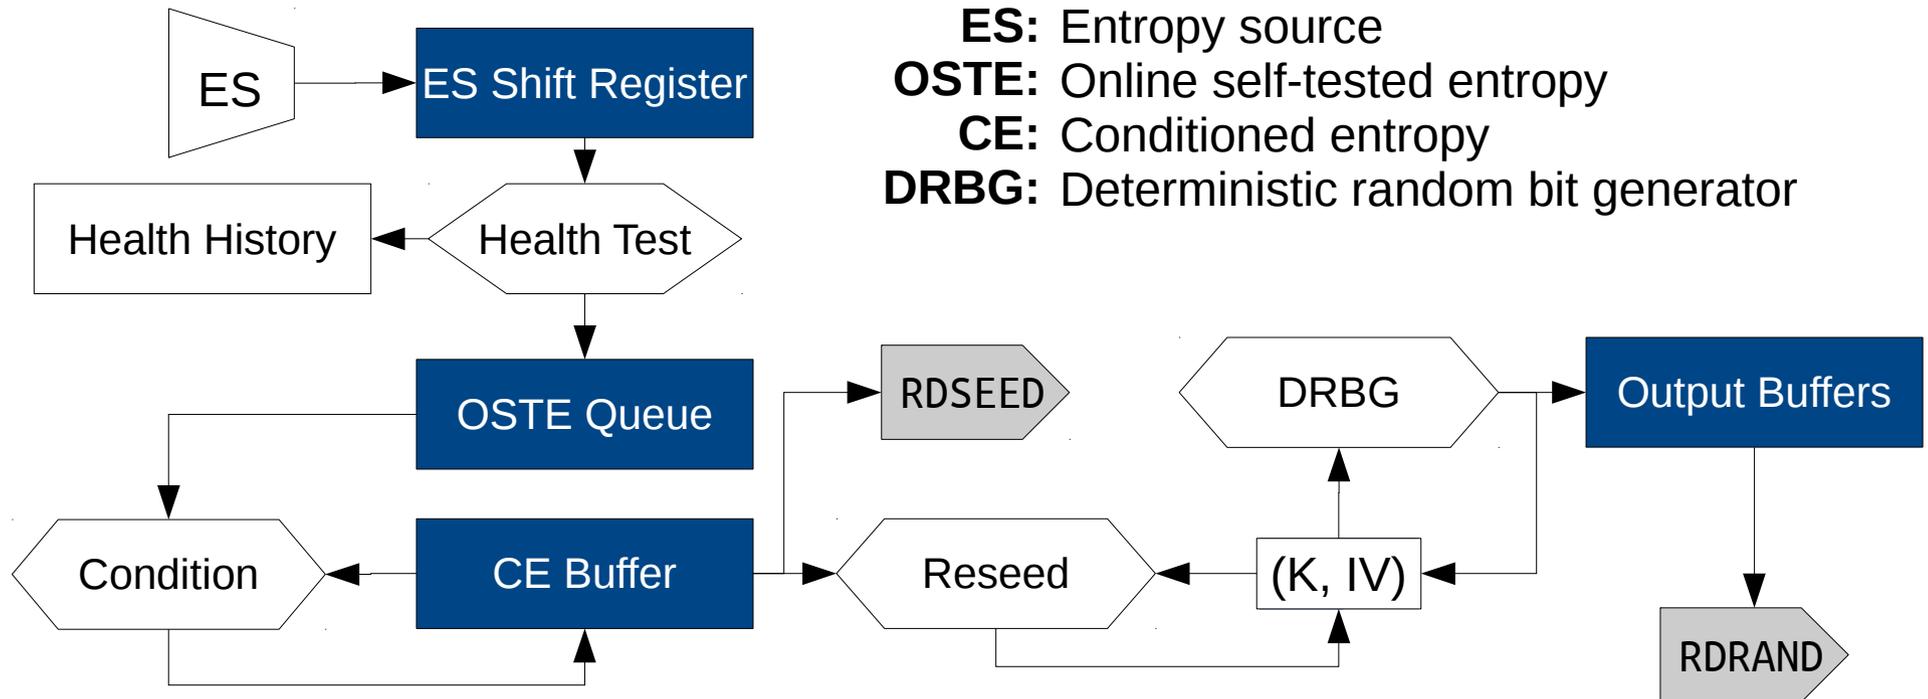
## Results

$$\mathbf{Adv}_{\text{ISK-RNG}, M}^{\text{INIT}}(A) \leq 2^{(k-m\gamma)/2+2} + 4\epsilon(L_m) + 8\hat{\epsilon}(L_m) + 5 \left( \mathbf{Adv}_{\text{AES}}^{\text{prp}}(B) + \frac{3}{2^k} \right)$$

$$\mathbf{Adv}_{\text{RDRAND}}^{\text{FWD}/M}(A) \leq 2(q+4) \left( \mathbf{Adv}_{\text{AES}}^{\text{prp}}(B) + \frac{3}{2^k} \right) + \mathbf{Adv}_{\text{RDRAND}}^{\text{backdoor}}(\text{NSA})$$

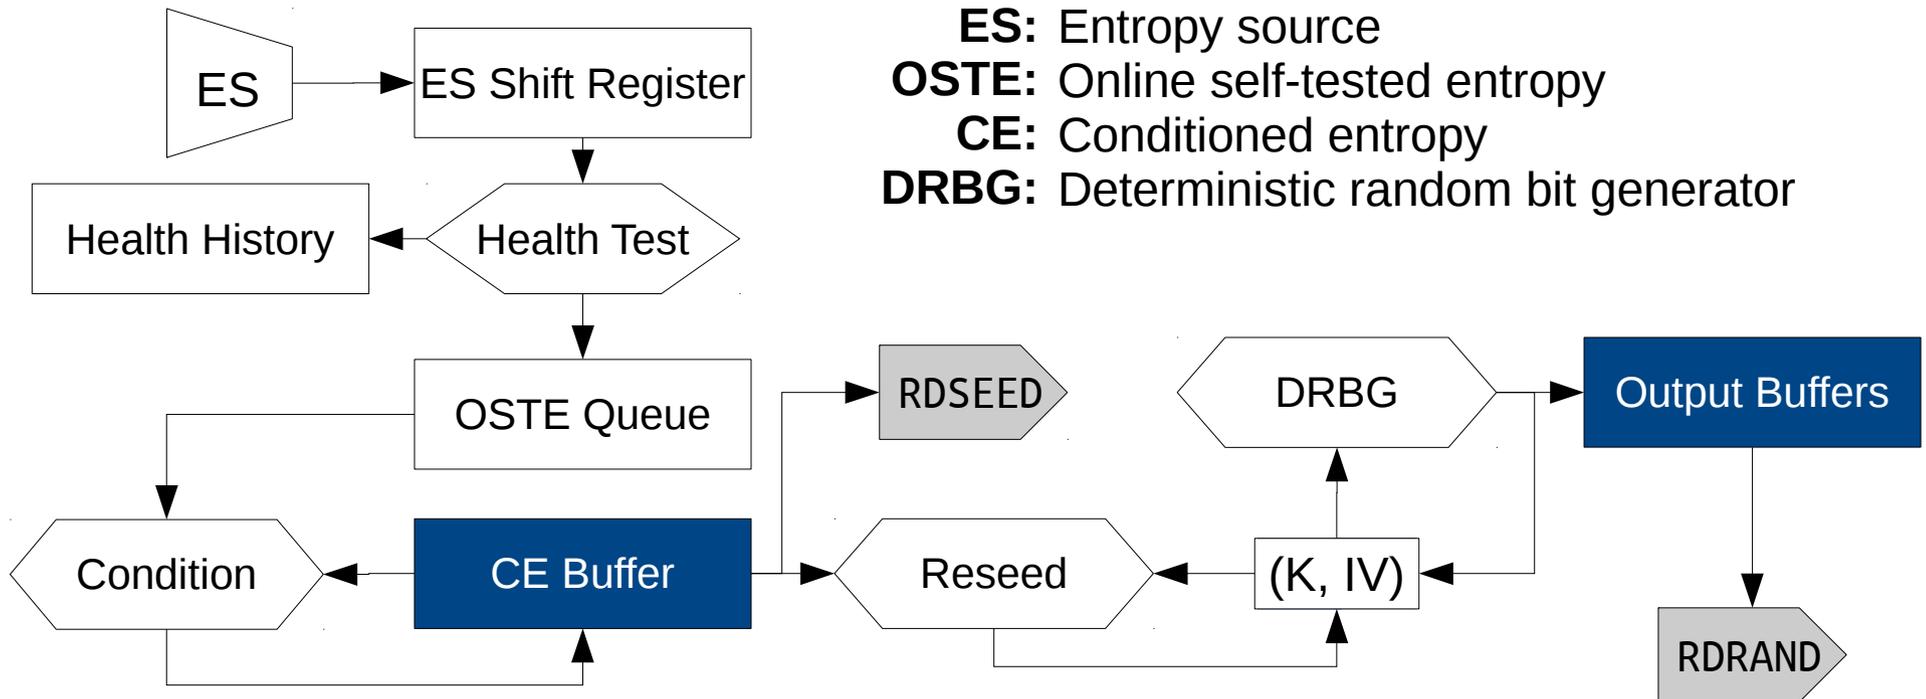
$$\begin{aligned} \mathbf{Adv}_{\text{ISK-RNG}}^{\text{FWD}/M} A &\leq (q+1) \left( 2^{(k-m\gamma)/2} + \epsilon(L_m) + 2(L_m) \right) \\ &\quad + 2(q+4) \left( \mathbf{Adv}_{\text{AES}}^{\text{prp}}(B) + \frac{3}{2^k} \right) \end{aligned}$$

# Limited Backwards Security



- Future outputs, DRBG seeds linger in buffers, no entropy collected when buffers full
- **Only a concern for hardware attacks**
- Security regained after this pipeline is flushed

# Technically no Forward Security



- Old outputs linger in buffers until overwritten
  - Barring ES failure, this takes < 1 microsecond
- **Only a concern for hardware attacks**

# Resilience / Limited Forward Security

Need to start plugging in numbers

- CRI report estimates each 256 bit sample has  $(0.65)^{256}$  bits of entropy. Let's use 128 bits.
- No empirical data for  $\beta$ , but somewhere between 0.5 and 0.99 seems reasonable

Let's to find max number of RDRAND/RDSEED queries that limits advantage to  $2^{-40}$ .

# Resilience / Limited Forward Security

To keep Adversary advantage below  $2^{-40}$ :

- **RDSEED**: Limit ~64MB (!)
- **RDRAND**: Hope Adversary can't bruteforce AES (dominant non-computational term is  $\sim 6q/2^{128}$ ).

**RDSEED** bound is against **computationally unbounded** attackers, doesn't reflect a real weakness...

...we think.

# Bottom line

- RDRAND design seems sound
- RDSEED security bounds problematic, but probably okay in practice

