Paul E. McKenney, Meta Platforms Kernel Team

Linux Plumbers Conference Refereed Track, November 14, 2023

Hunting Heisenbugs

Heisenbugs and impressionism: The closer you get, the less you see!

© 2023 Meta Platforms

Overview

- Heisenbugs, Then and Now
- How to Hunt Heisenbugs
- Heisenbugs: The Goal

Overview

- Heisenbugs, Then and Now
- How to Hunt Heisenbugs
- Heisenbugs: The Goal

How to avoid hunting heisenbugs!!!

Heisenbugs, Then and Now

Heisenbugs, Then and Now

- Heisenbugs used to result in horrible life-changing experiences
- But they are increasingly just “Tuesday”

Fleets as Heisenbug Detectors

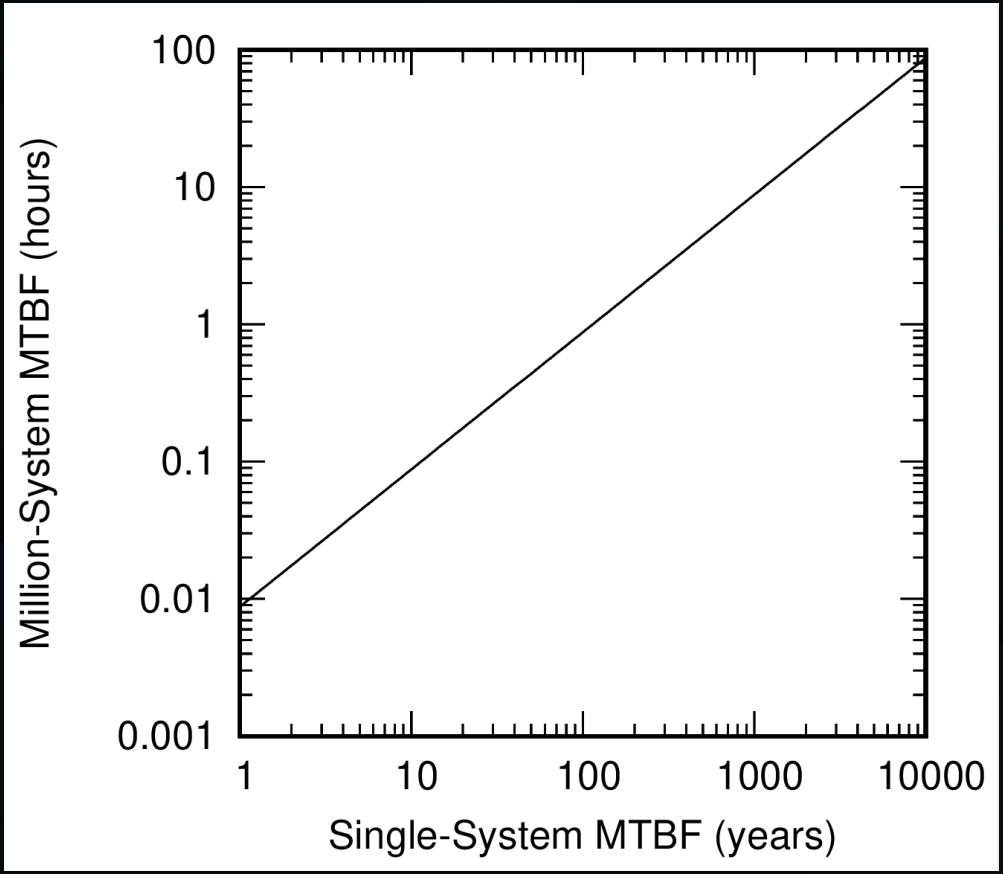

After Heisenbug Detected?

- Debug based on console output
- Collect debug information via BPF
- Collect debug information via kernel patch
- Set up kernel debugger

After Heisenbug Detected?

- **Debug based on console output**
- Collect debug information via BPF
- Collect debug information via kernel patch
- Set up kernel debugger

One failure per 4K systems per week (70-year single-system MTBF)

After Heisenbug Detected?

- Debug based on console output
- Collect debug information via BPF
- Collect debug information via kernel patch
- **Set up kernel debugger**

After Heisenbug Detected?

- Debug based on console output
- Collect debug information via BPF
- Collect debug information [illegible] kernel patch
- Set up kernel deb[illegible]

Except that 1M instances of kgdb is not fun...

After Heisenbug Detected?

- Debug based on console output
- Collect debug information via BPF
- Collect debug information via kernel patch
- Set up kernel debugger

Except that 1M installs of kgdb is not fun...

Nor is waiting for your alleged fix to deploy!!!

After Heisenbug Detected?

- Debug based on console output
- Collect debug information via BPF
- Collect debug information via kernel patch
- Set up kgdb

Nor is waiting for ... of kgdb is not fun...

Need a better way

Except that 1M ins... ...yed fix to deploy!!!

But What If I Don’t Have A Big Fleet?

- I hunted heisenbugs long before “having” a fleet!!!

But What If I Don’t Have A Big Fleet?

- I hunted heisenbugs long before “having” a fleet!!!
- Tens of billions of Linux instances
  - Good to have fewer things going bump in the night
- Potential safety-critical benefit
  - Whether we like it or not, Linux kernel is already increasingly used in safety-critical applications
  - Acceptance test suffices in many cases

Key Heisenbug-Hunting Trick

Key Heisenbug-Hunting Trick

- Why is it a heisenbug?

Key Heisenbug-Hunting Trick

- Why is it a heisenbug?
- Because it occurs only rarely
- Any added debugging changes timing
- Slight changes in timing can reduce incidence

Key Heisenbug-Hunting Trick

- Why is it a heisenbug?
- Because it occurs only r
- Any added debug
- Slight chan

iming

an reduce incidence

Anti-Heisenbug: Reduce MTBF!!!

What If I Do Have a Fleet?

- One-week test on 50 systems to validate to run for one year on 1M systems?

What If I Do Have a Fleet?

- One-week test on 50 systems to validate to run for one year on 1M systems?
  - MTBF of test systems must be ***six orders of magnitude*** shorter than MTBF of fleet systems

What If I Do Have a Fleet?

- One-week test on 50 systems to validate to run for one year on 1M systems?
  - MTBF of test systems must be ***six orders of magnitude*** shorter than MTBF of fleet systems
  - **In many cases, this is eminently doable**

How to Hunt Heisenbugs

Increase Workload Intensity

- Leverage the philosophy of my high-school track and cross-country coach

Increase Workload Intensity

- Leverage the philosophy of m… school track and cross-countr…

Race days were the easy days

Increase Workload Intensity: Kernel

- Most production systems major in userspace execution (if not idle!)
  - 10% kernel utilization is high
  - A few percent kernel utilization is commonplace

Increase Workload Intensity: Kernel

- Most production systems major in userspace execution (if not idle!)
  - 10% kernel utilization is high
  - A few percent kernel utilization is commonplace
- One-to-two orders of magnitude MTBF reduction just by focusing on kernel execution!

Increase Workload Intensity: Kernel

- Most production systems major i erspace execution (if not idle!)
  - 10% kernel utilization is
  - A few percent ker ilization is commonplace
- One-to-two rs of magnitude MTBF reductio st by focusing on kernel execution!

rcutorture majors in this

Increase Workload Intensity: Kernel

- Most produc[illegible] systems major i[illegible]erspace executi[illegible]
  - 10% kerne[illegible]
  - A few percent k[illegible] commonplace
- One-to-two [illegible]s of h[illegible] BF reductio[illegible]t by focusing o[illegible]el execution!

Anti-Heisenbug: Increase Intensity!!!

Increase Workload Intensity: Kernel

- Special case of testing suspicious subsystems in isolation
- Configure application to beat up kernel
- Run kernel portion of workload from traces taken from application
- Increase CPUs, memory, I/O, ...

Increase Workload Intensity: Kernel

- Special case of testing suspicious subsyste
  in isolation
- Configure applicati nel
- Run kernel po kload from traces
  taken fro ation
- I CPUs, memory, I/O, ...

Ask yourself:
What has caused trouble in the past?

Increase Workload Intensity: Kernel

- Special case of testing suspicious subsystem in isolation
- Configure kernel
- Run kernel power traces taken from production
- Increase CPUs, memory, I/O, ...

Anti-Heisenbug: Look for and promote trouble!!!

What has caused trouble in the past?

Workload-Intensity Caution: Latency

- Increasing load normally increases delays
  - Scheduler queueing
  - Lock contention
  - Memory contention

Theoretical Latency: M/M/1 Queue

Theoretical Latency: M/M/1 Queue

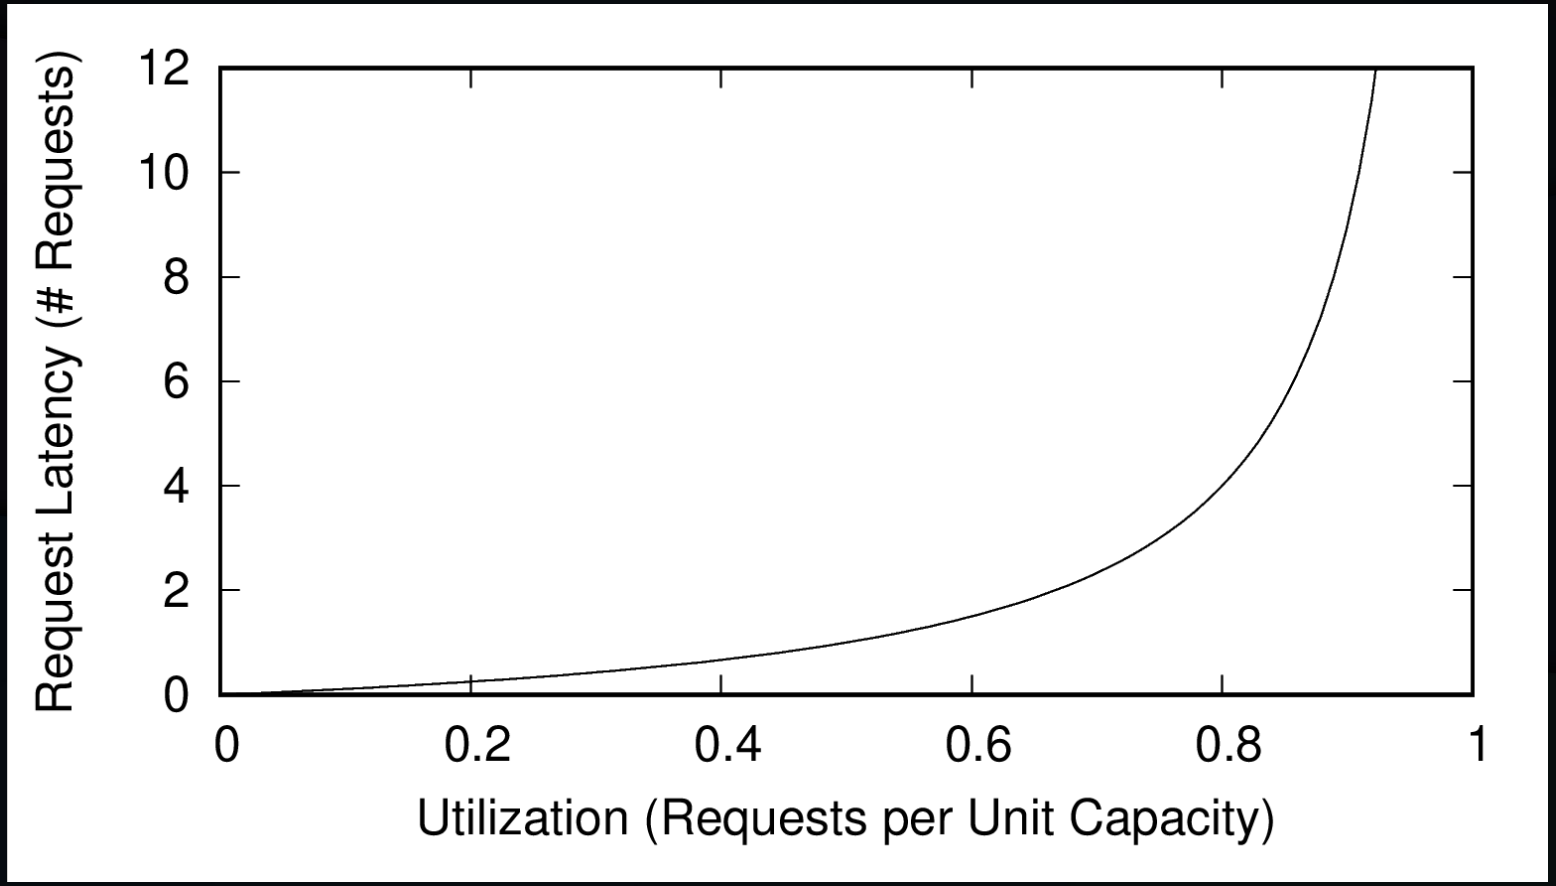

Theoretical Latency: M/M/1 Queue

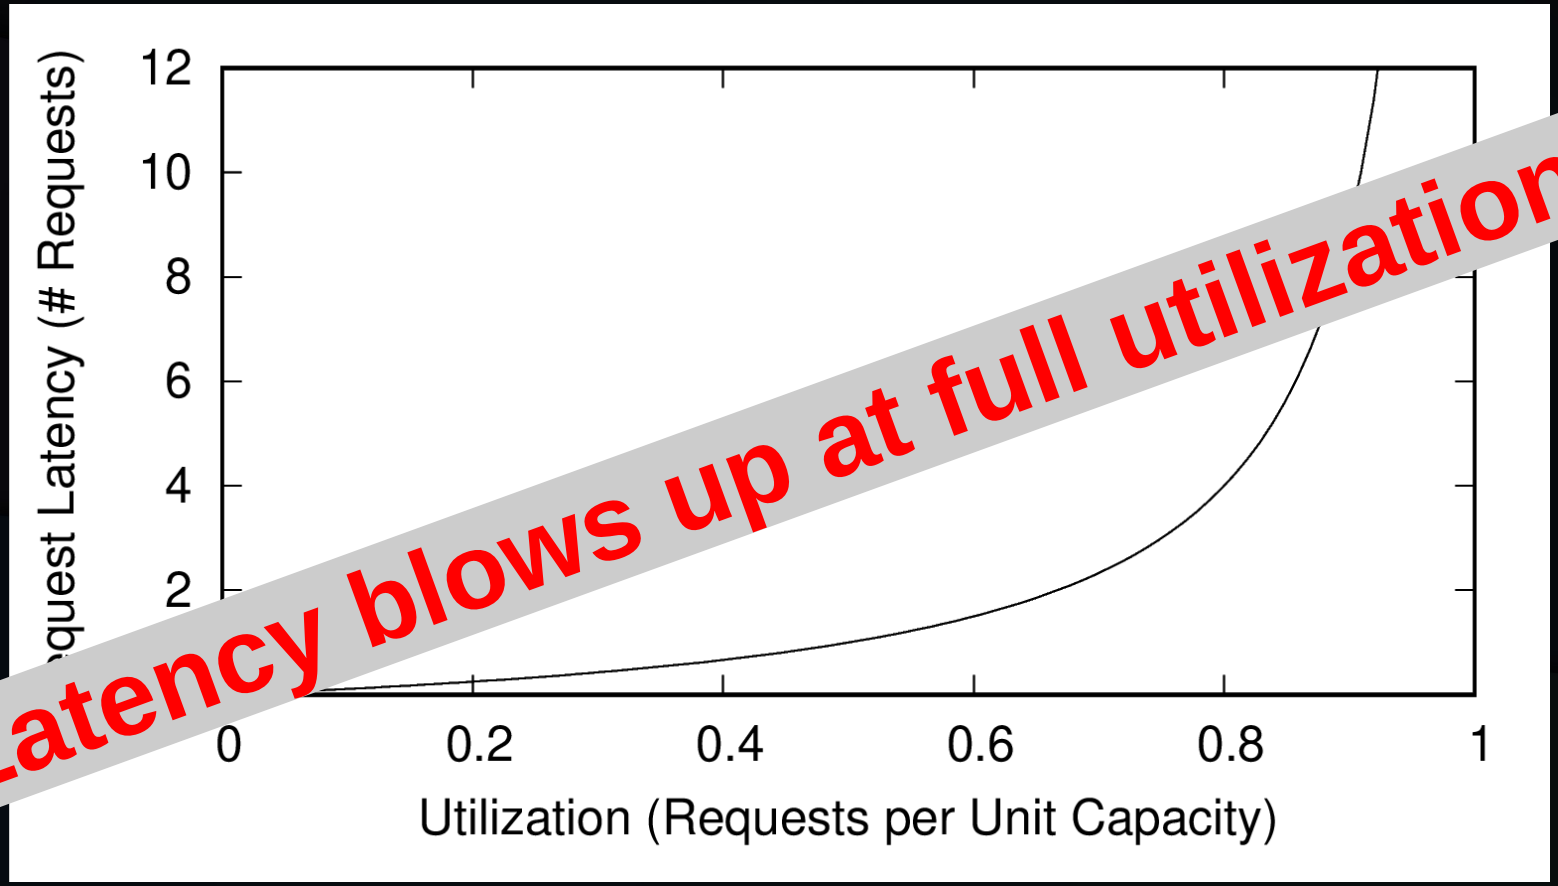

Theoretical Latency: M/M/1 Queue

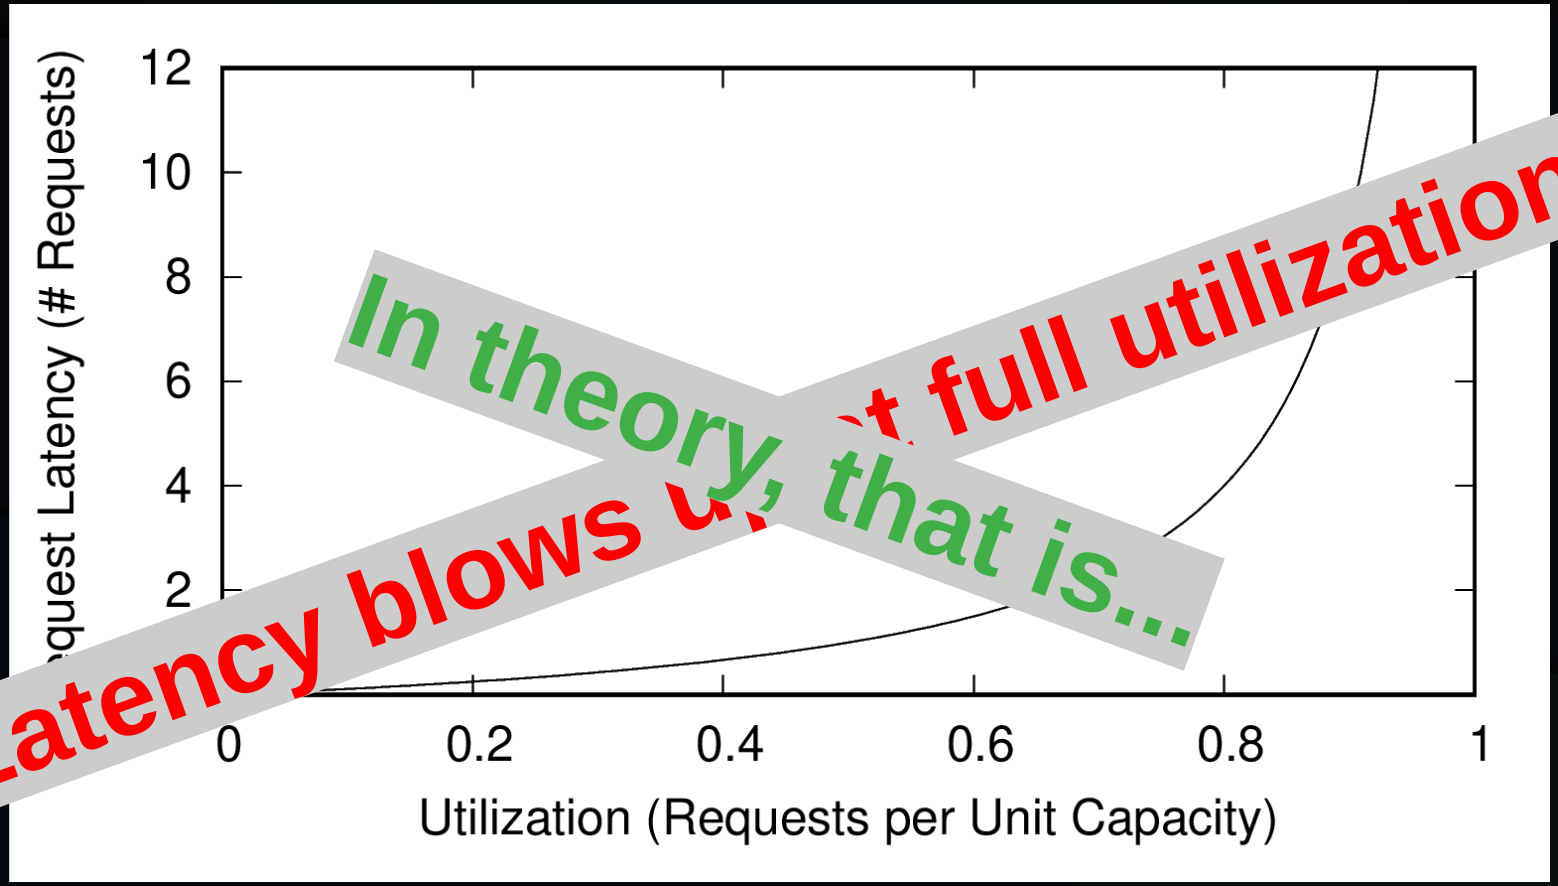

Queues Are Finite!!! M/M/1/k Queue

Latency From Finite Queueing

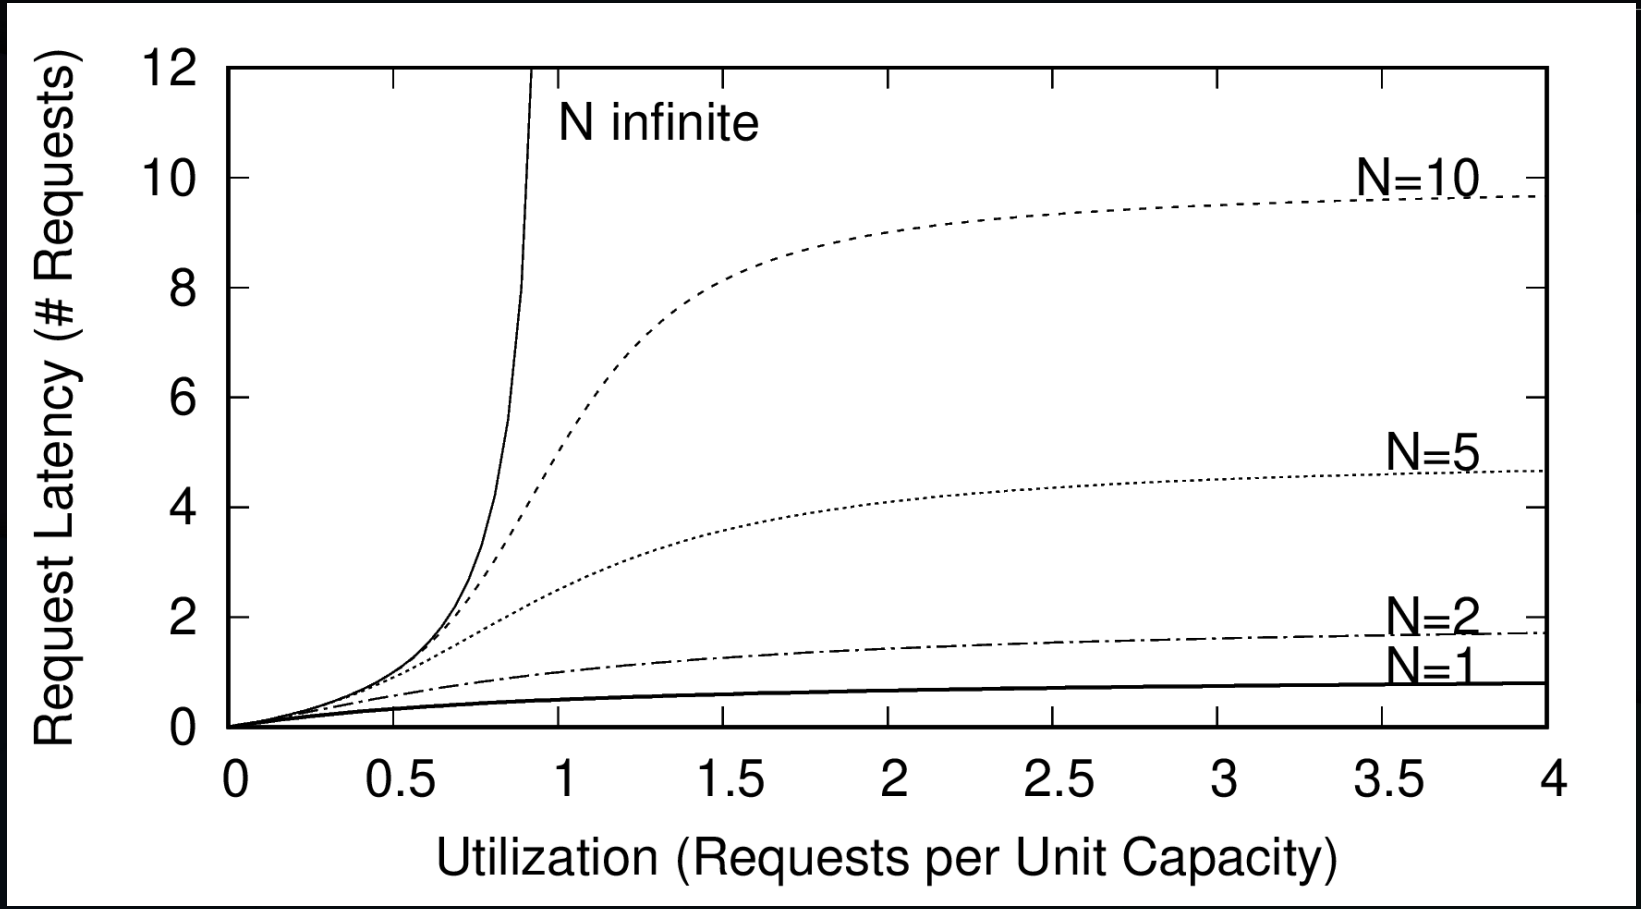

Latency From Finite Queueing

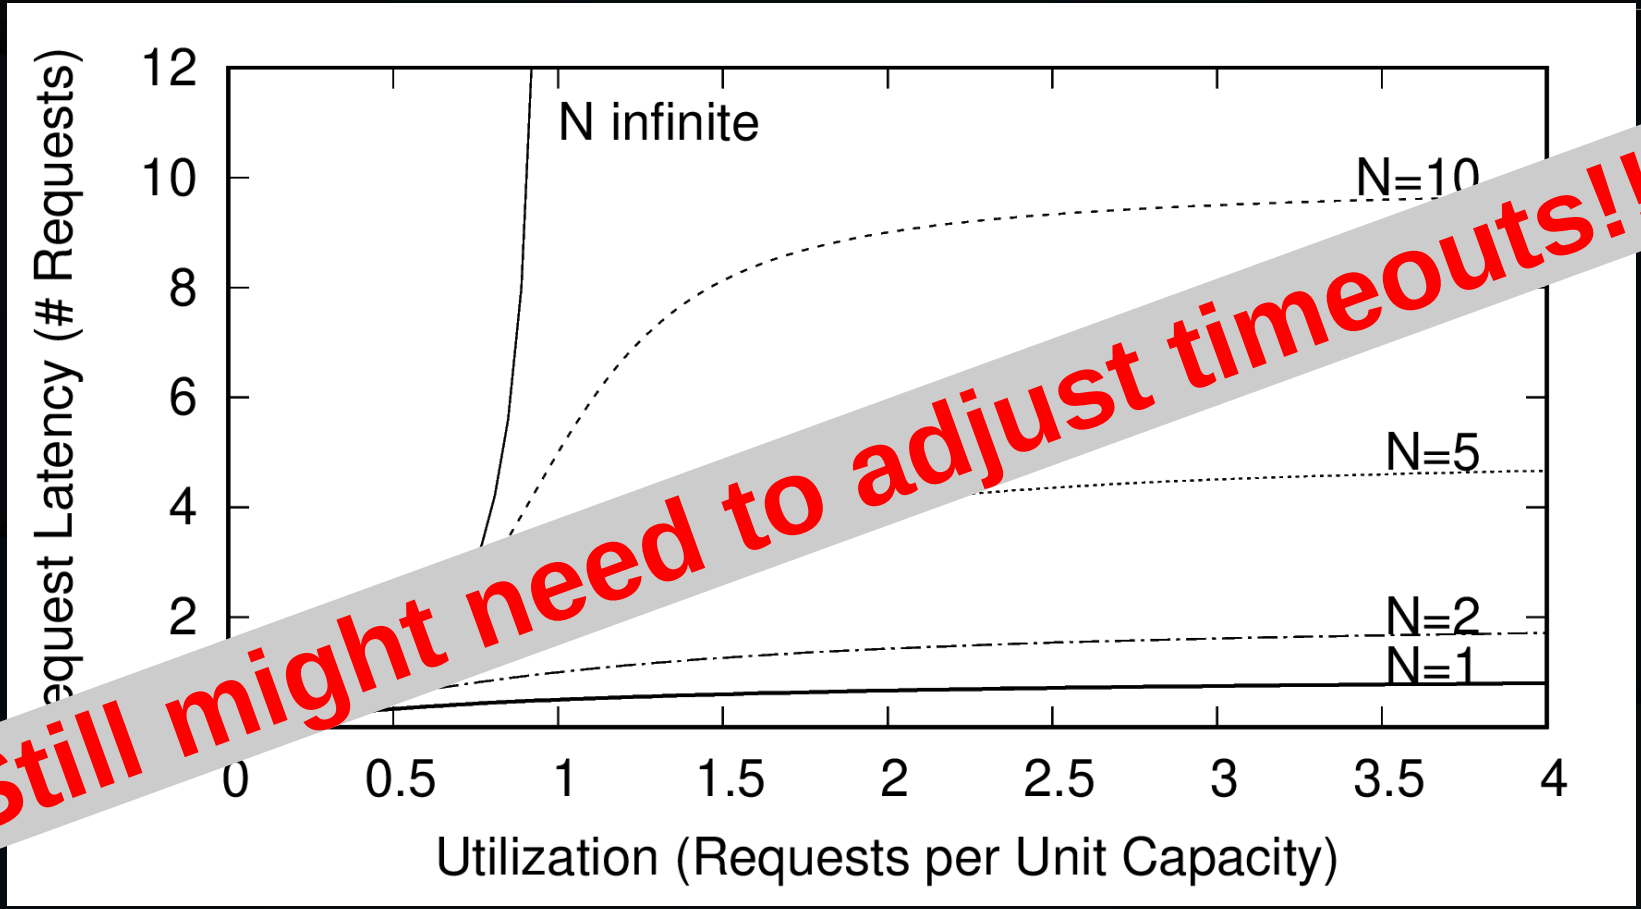

Still might need to adjust timeouts!!!

Inject Strategic Delays

Inject Strategic Delays

- Intensify one part of the workload by de-intensifying another
- Examples:
  - Running on multi-socket systems injects cache-miss latencies *
  - rcutorture injects delays during grace-period initialization to promote races with CPU-hotplug operations
  - Old days: Run CPUs at different speeds

* https://paulmck.livejournal.com/62071.html

RCU CPU-Hotplug Strategic Delays

- RCU need not wait on offline CPUs
  - Nor on CPUs that online after grace period start
    - Though it is OK to wait on them
  - But RCU does need to be clear on whether or not it needs to wait on a given CPU
  - And RCU does need to “keep its own books” on which CPUs are online

RCU CPU-Hotplug Strategic Delays

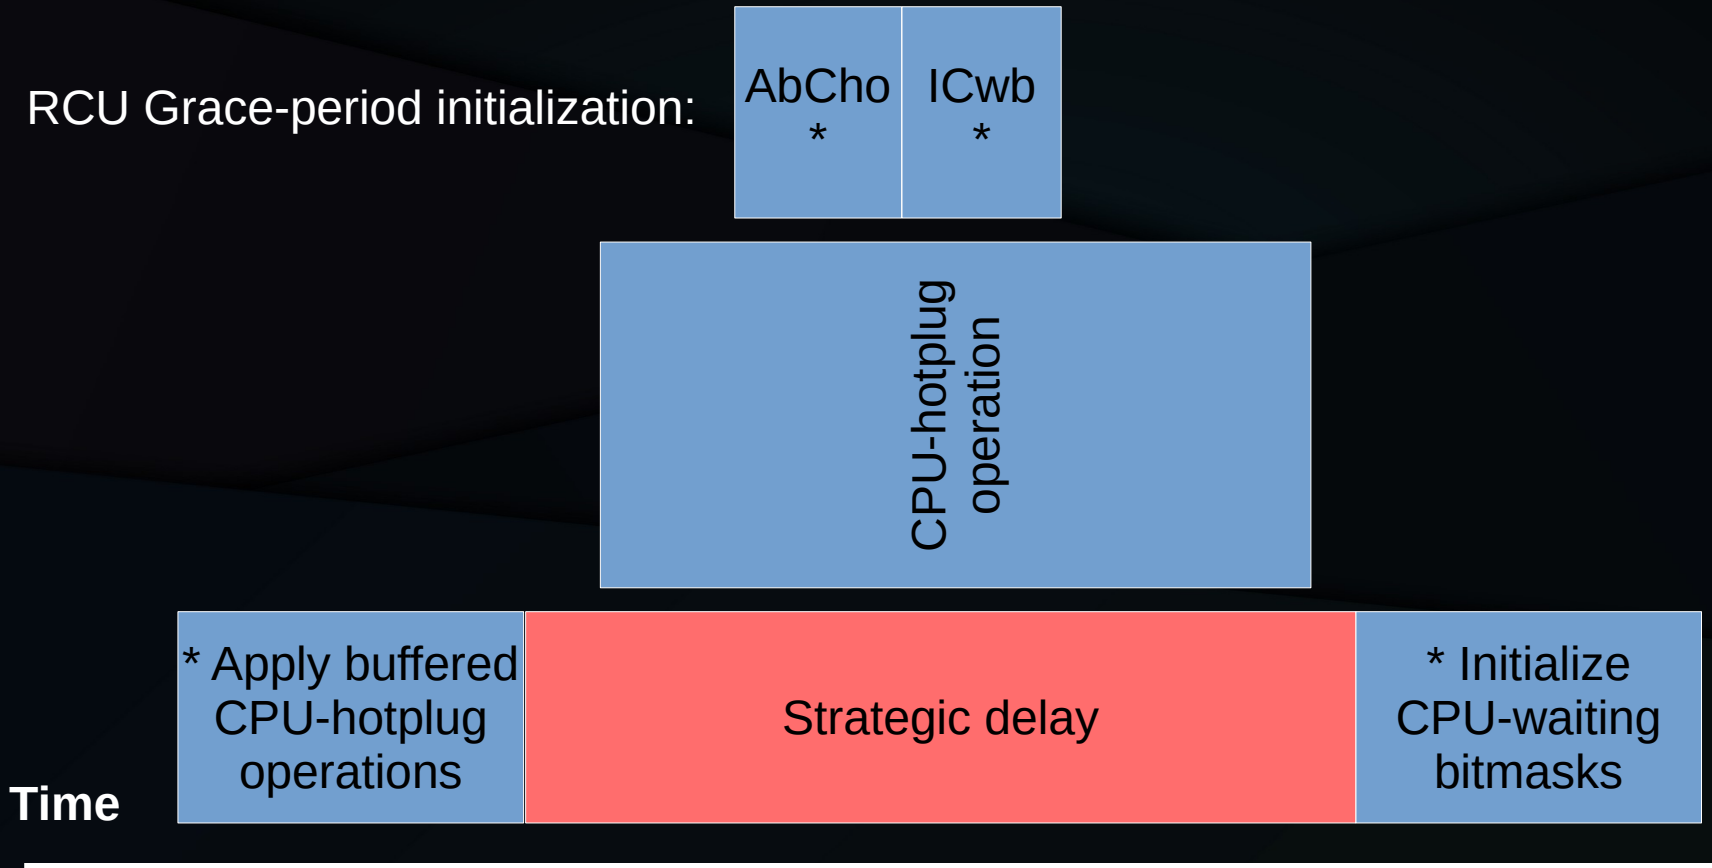

Count Near Misses

Count Near Misses

- USA FAA requires reporting of near misses *
  - Higher probability of near miss than of collision
- Near misses can help hunting heisenbugs
  - More quickly evaluate commits, configurations, and effectiveness of other anti-heisenbugs
  - Especially helpful when verifying fixes

* https://www.faa.gov/air_traffic/publications/atpubs/aim_html/chap7_section_7.html

Count Near Misses: RCU Example

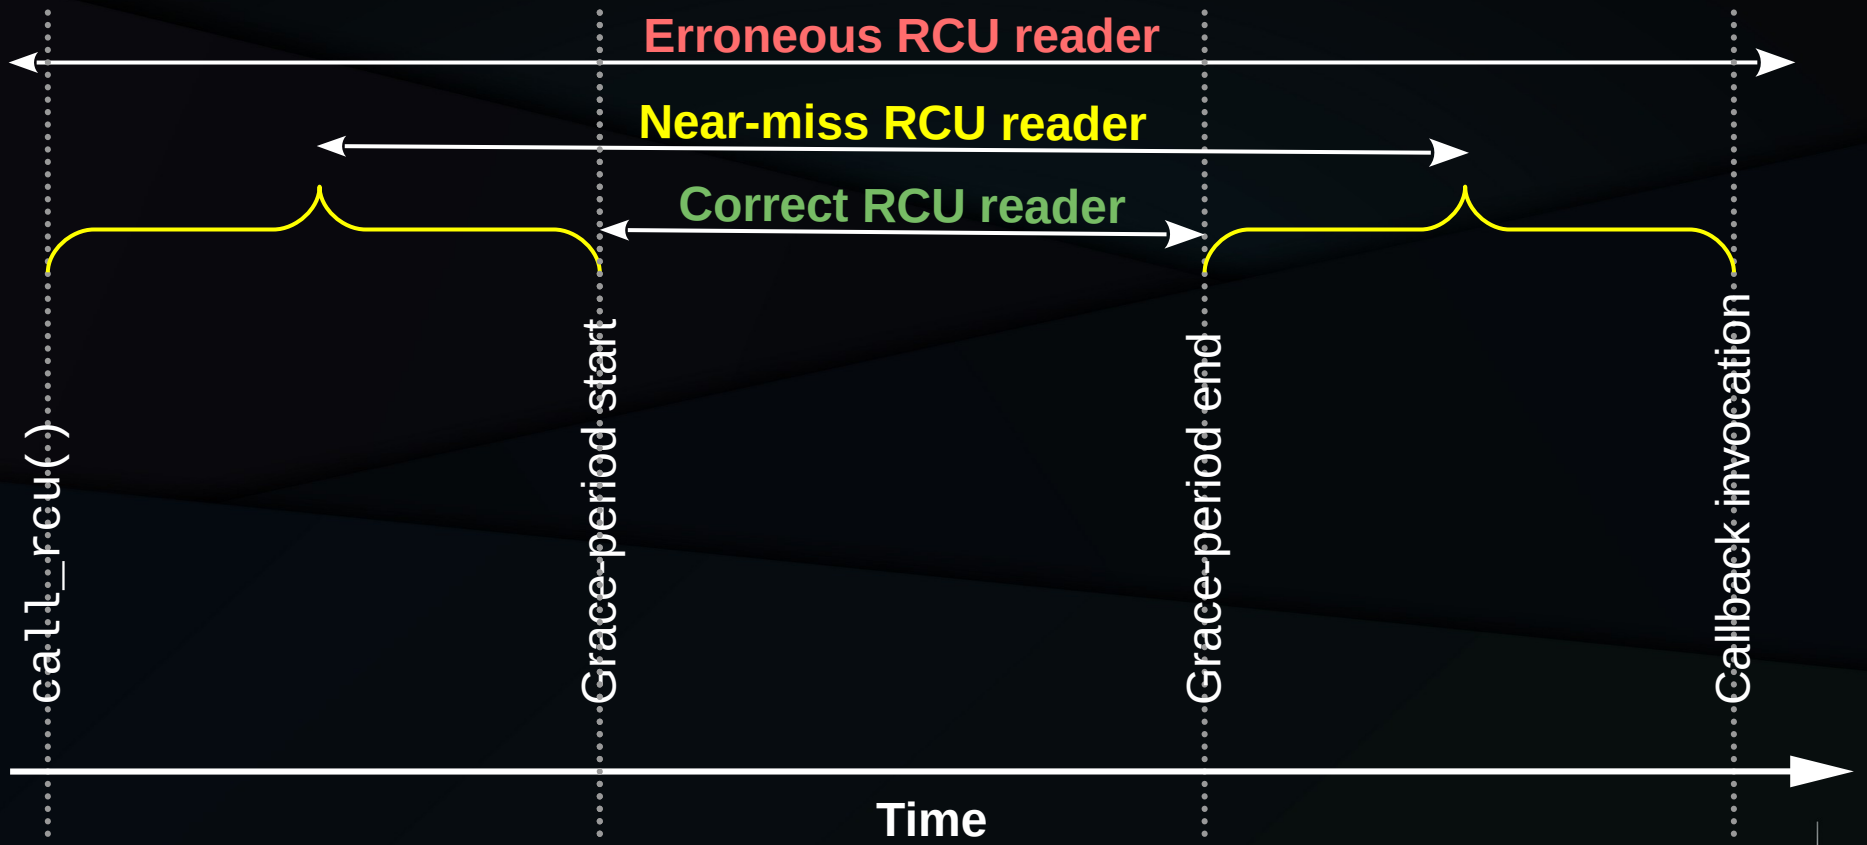

Count Near Misses: RCU Example

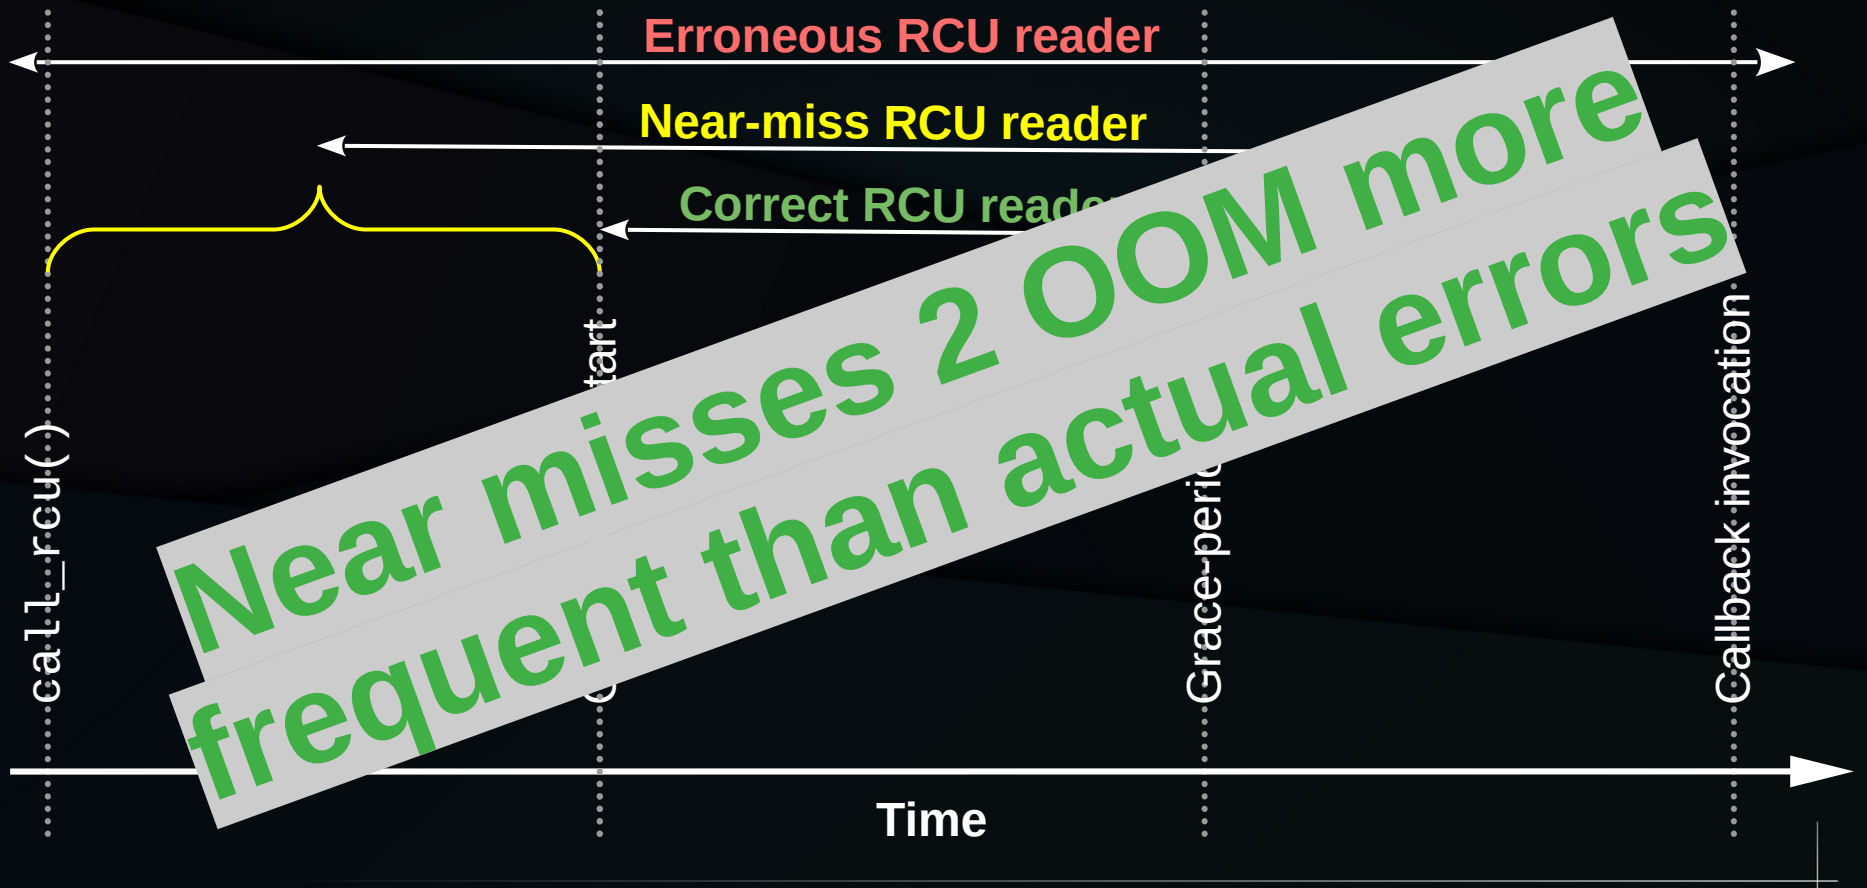

Count Near Misses: RCU Example

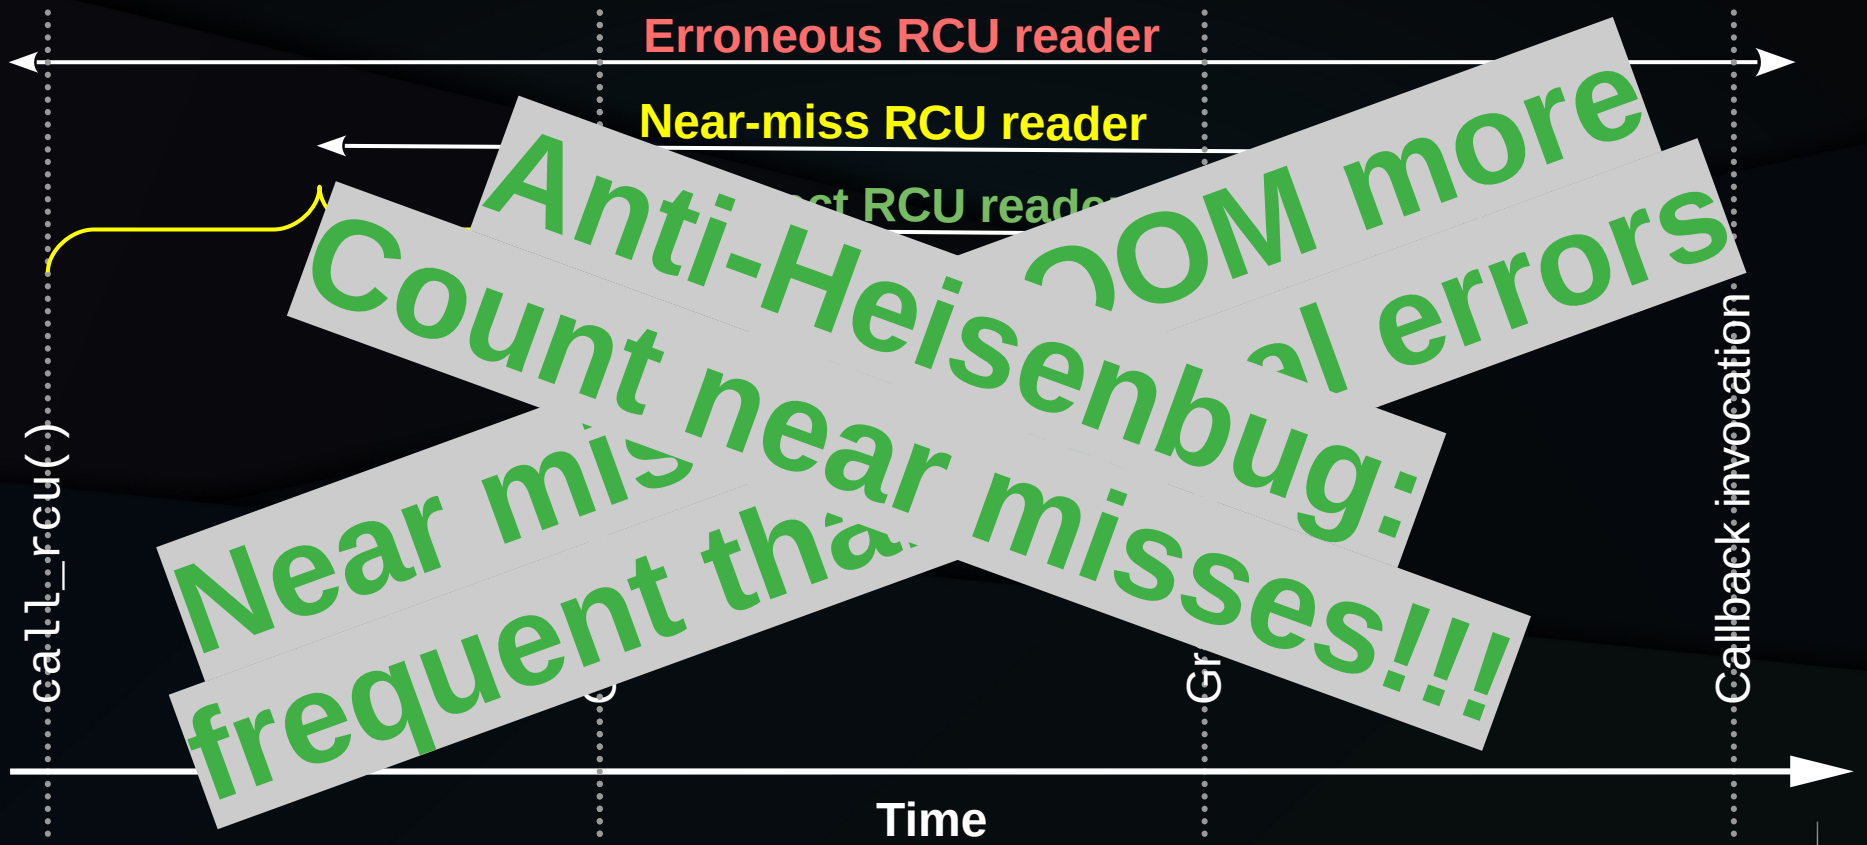

Make Rare Events Happen Frequently

Make Rare Events Happen Frequently

- Utilization, redux
- Force rare error conditions
- Force rare slowpath execution
- Add delays to race-prone code
- “The nuclear option”

Utilization and Rare Events

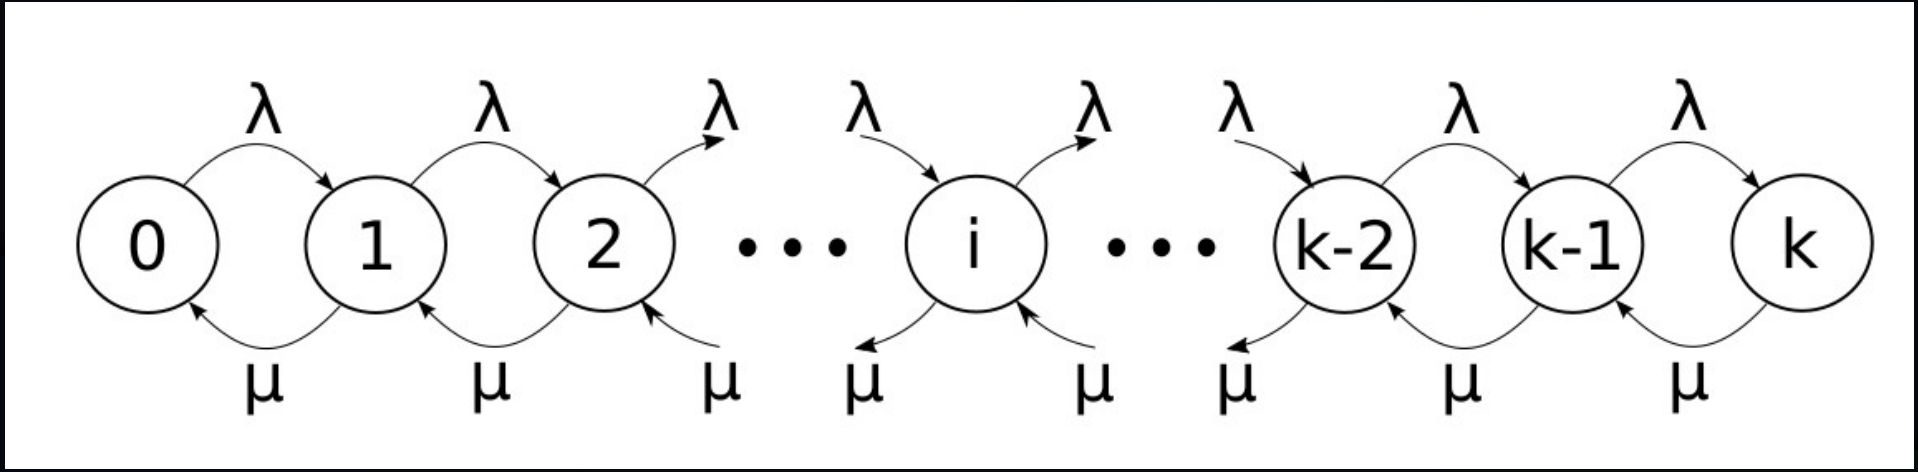

Utilization and Rare Events

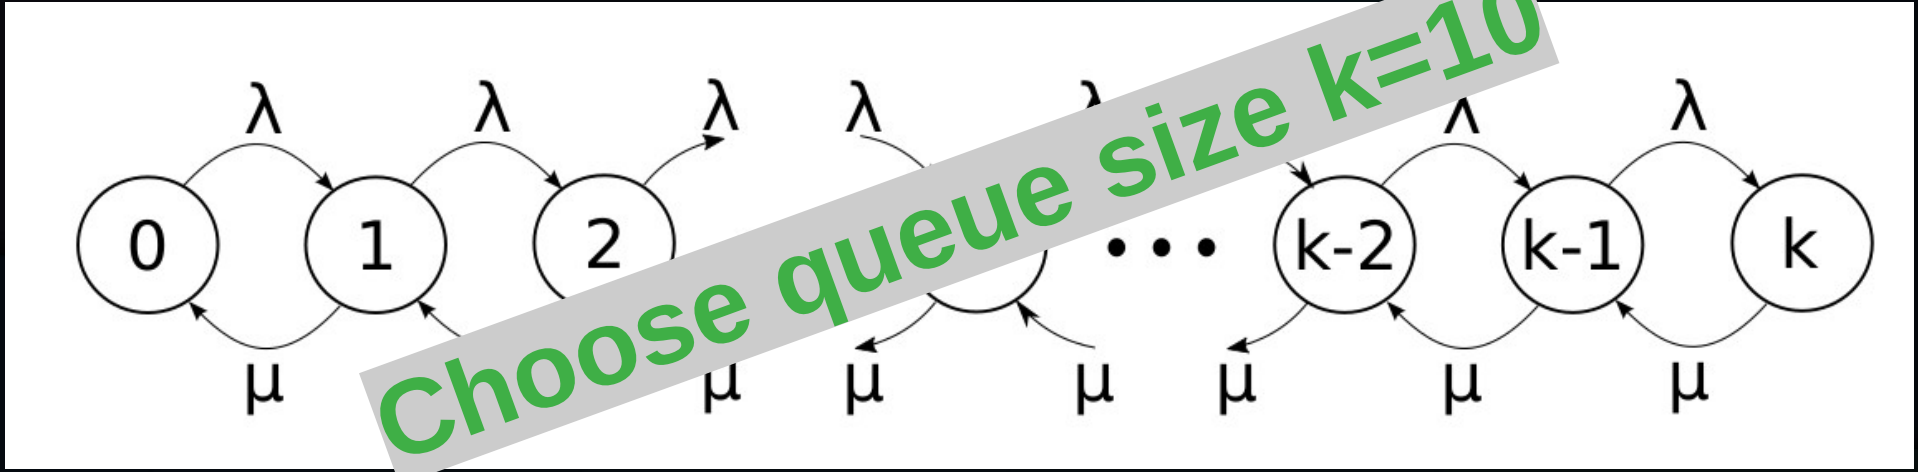

Choose queue size k=10

Utilization and Rarity of Events

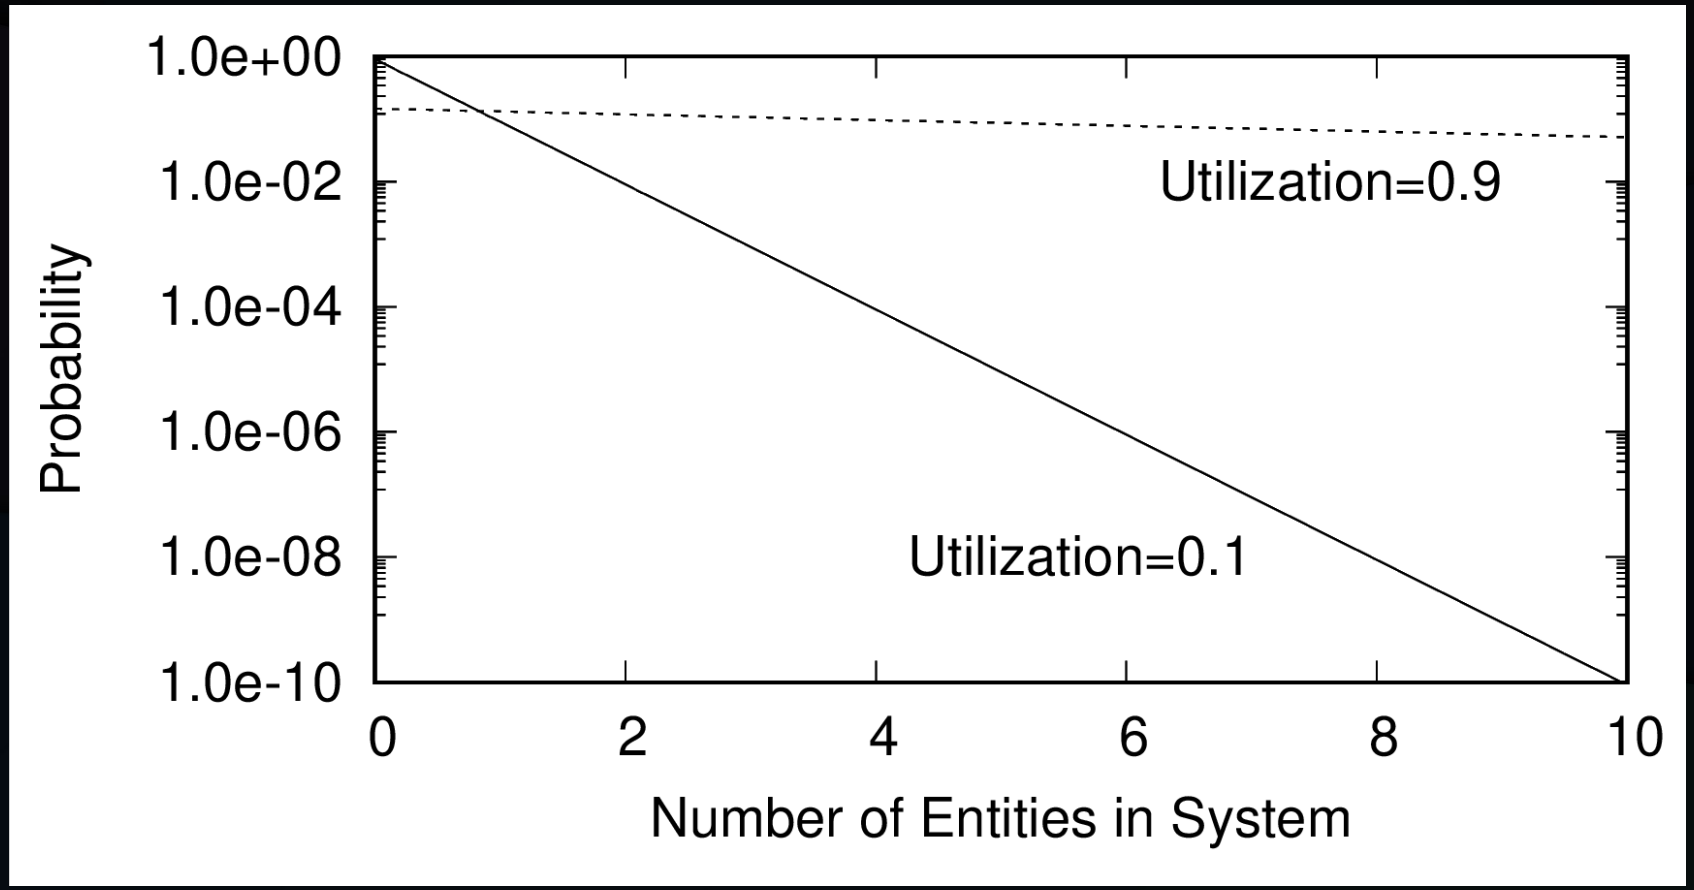

Utilization and Rarity of Events

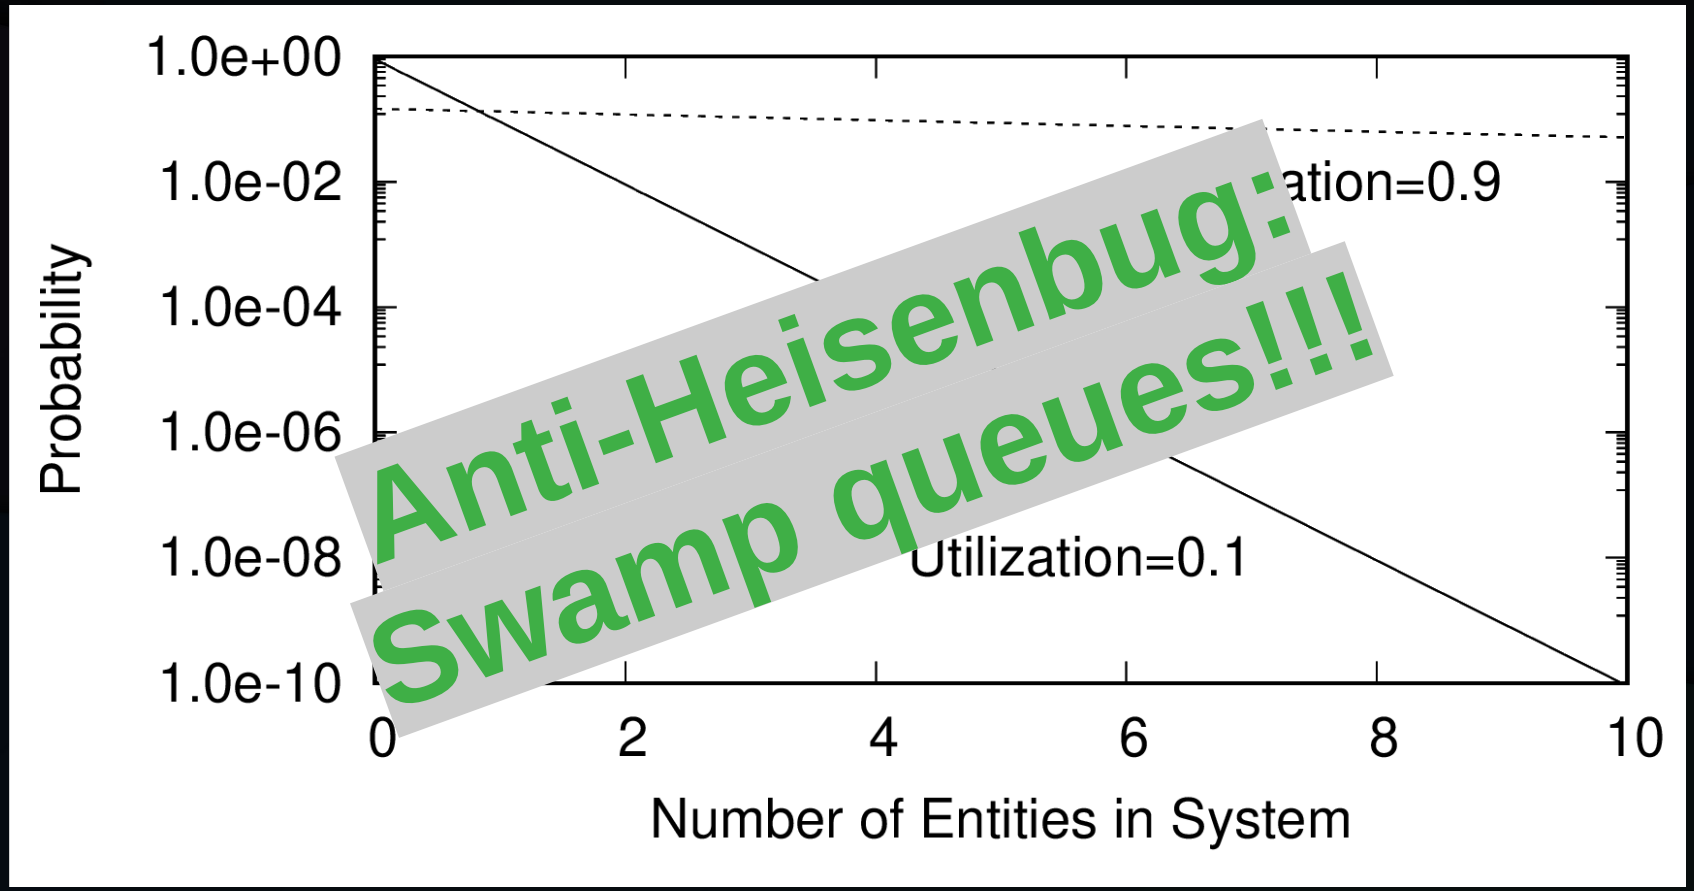

Utilization and Rarity of Events

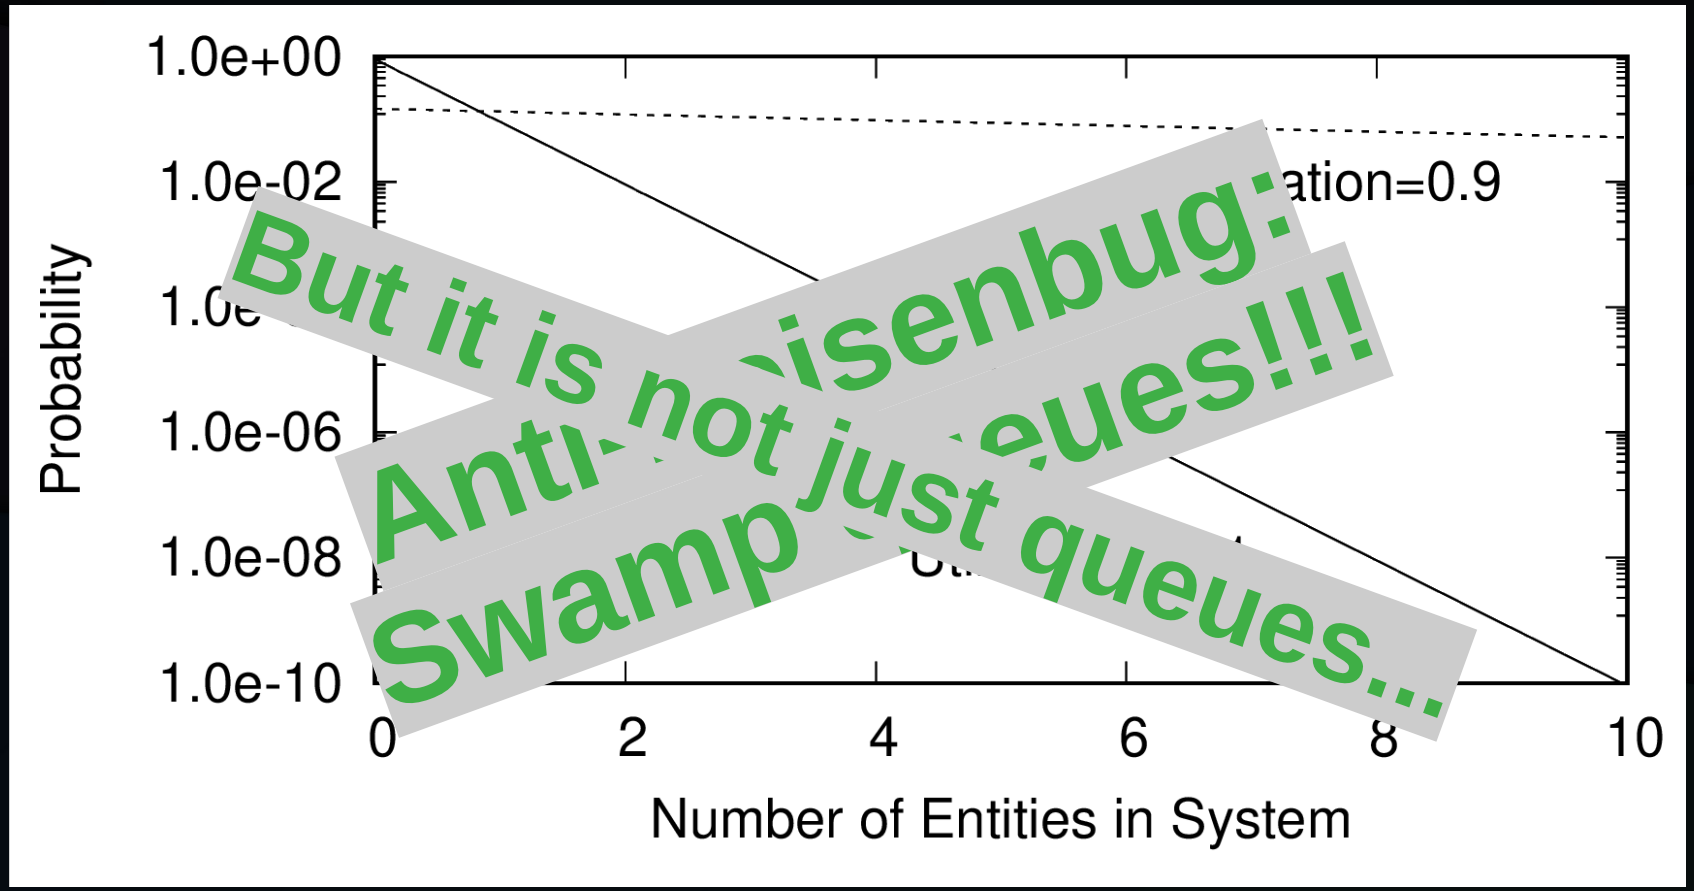

DYNIX/ptx Memory Allocator ca. 1993

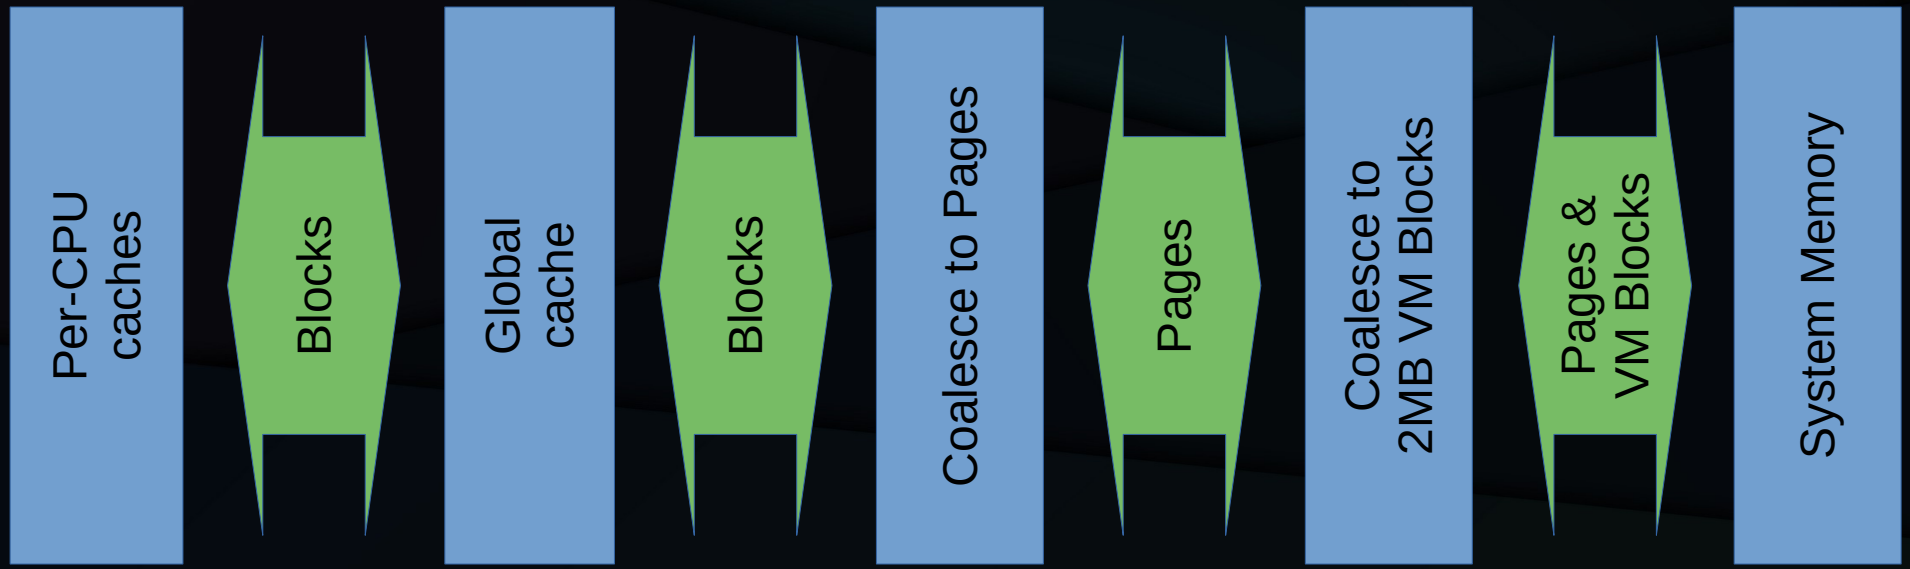

Later on added per-NUMA-node caches

DYNIX/ptx Memory Allocator ca. 1993

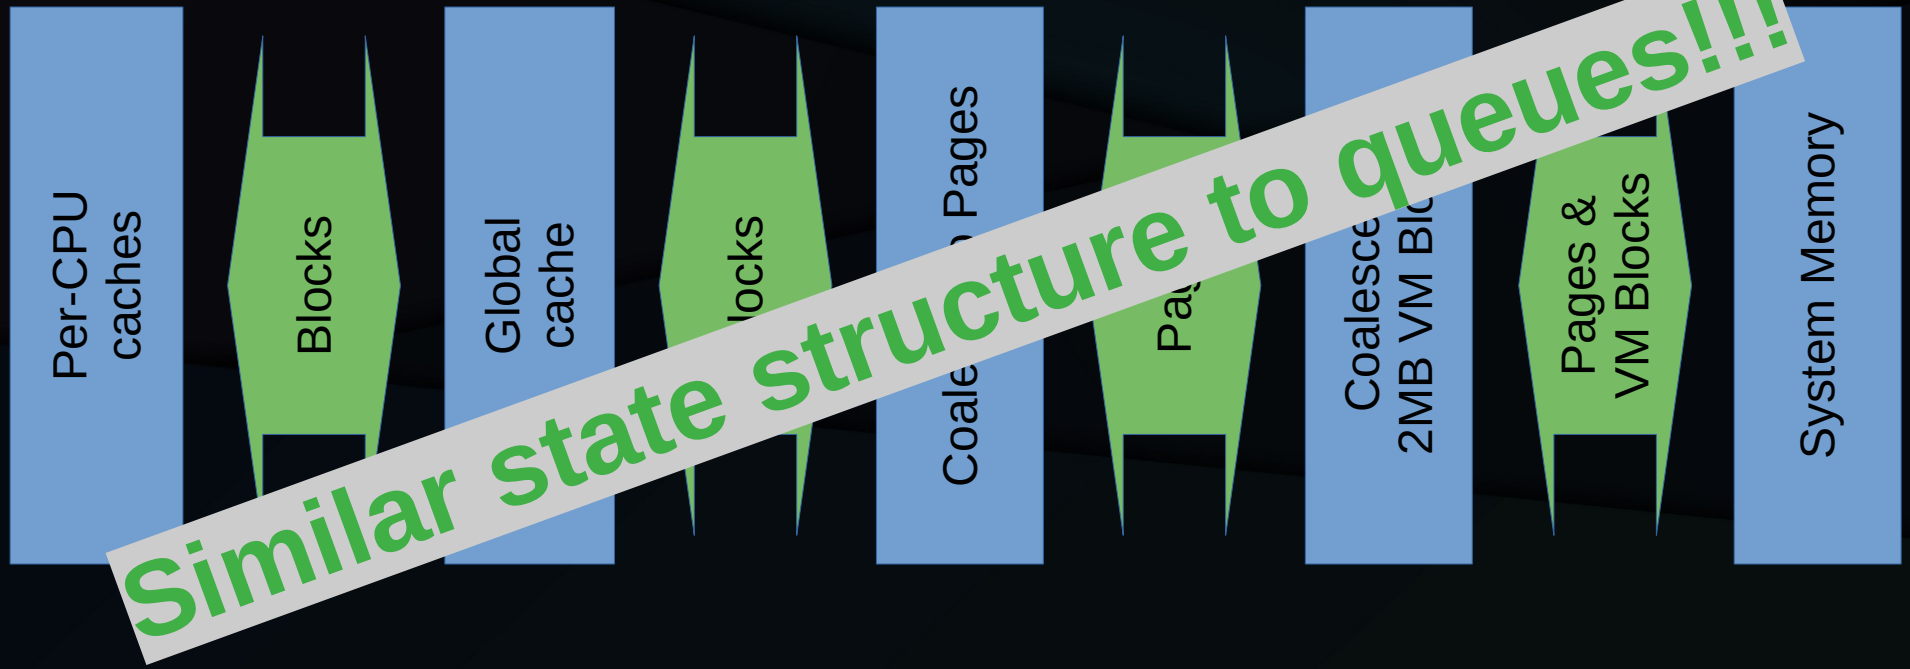

Later on added per-NUMA-node caches

Shared Disks For Availability Win!!!

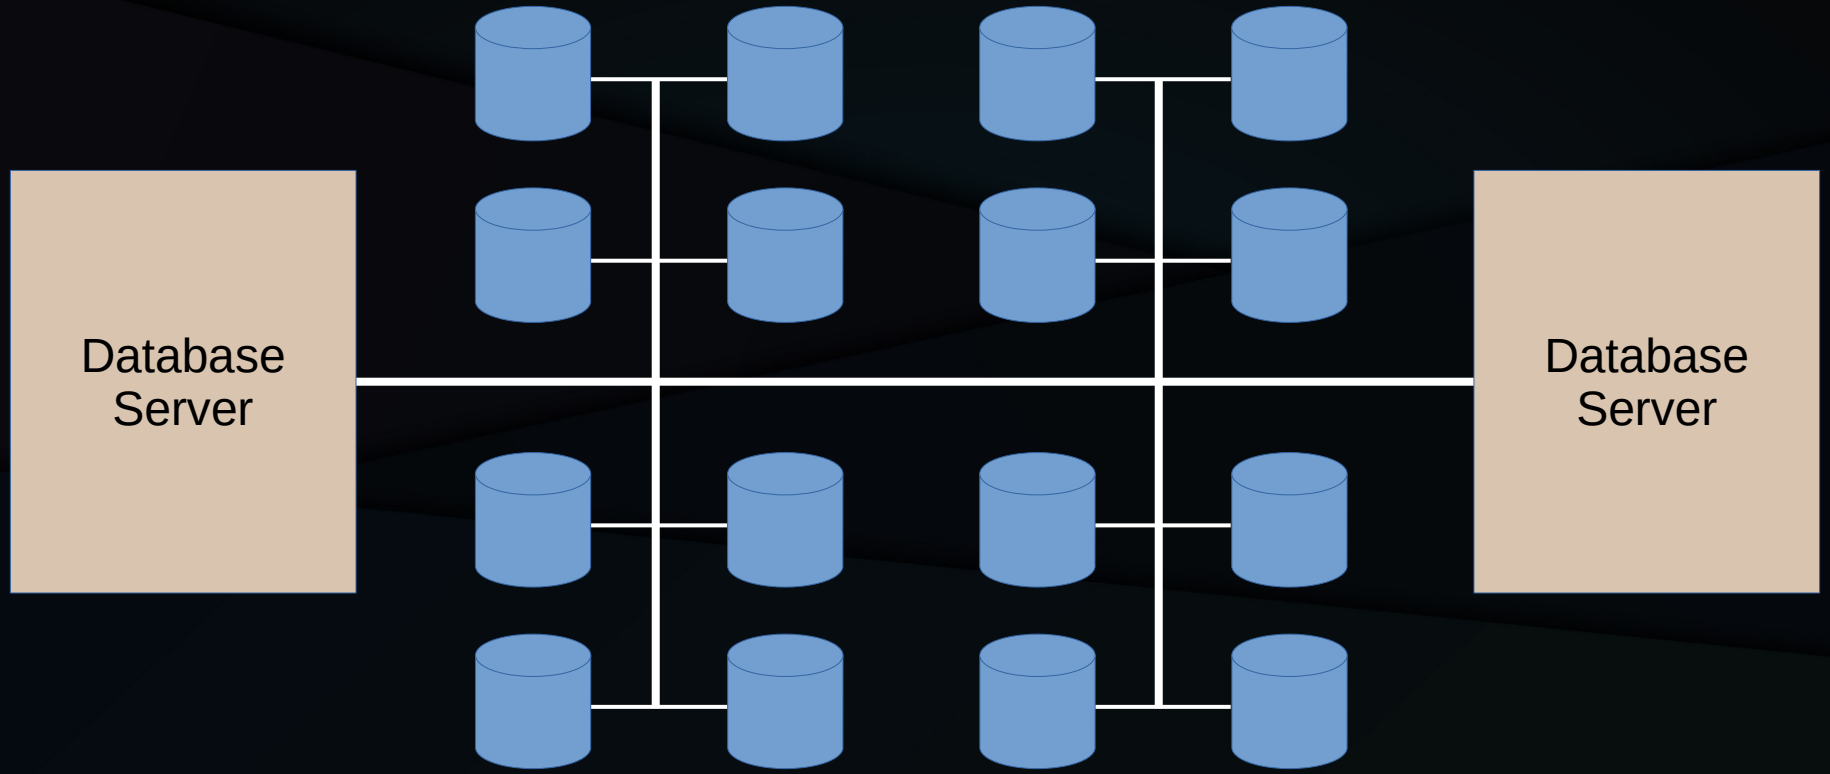

Shared Disks For Availability Win!!!

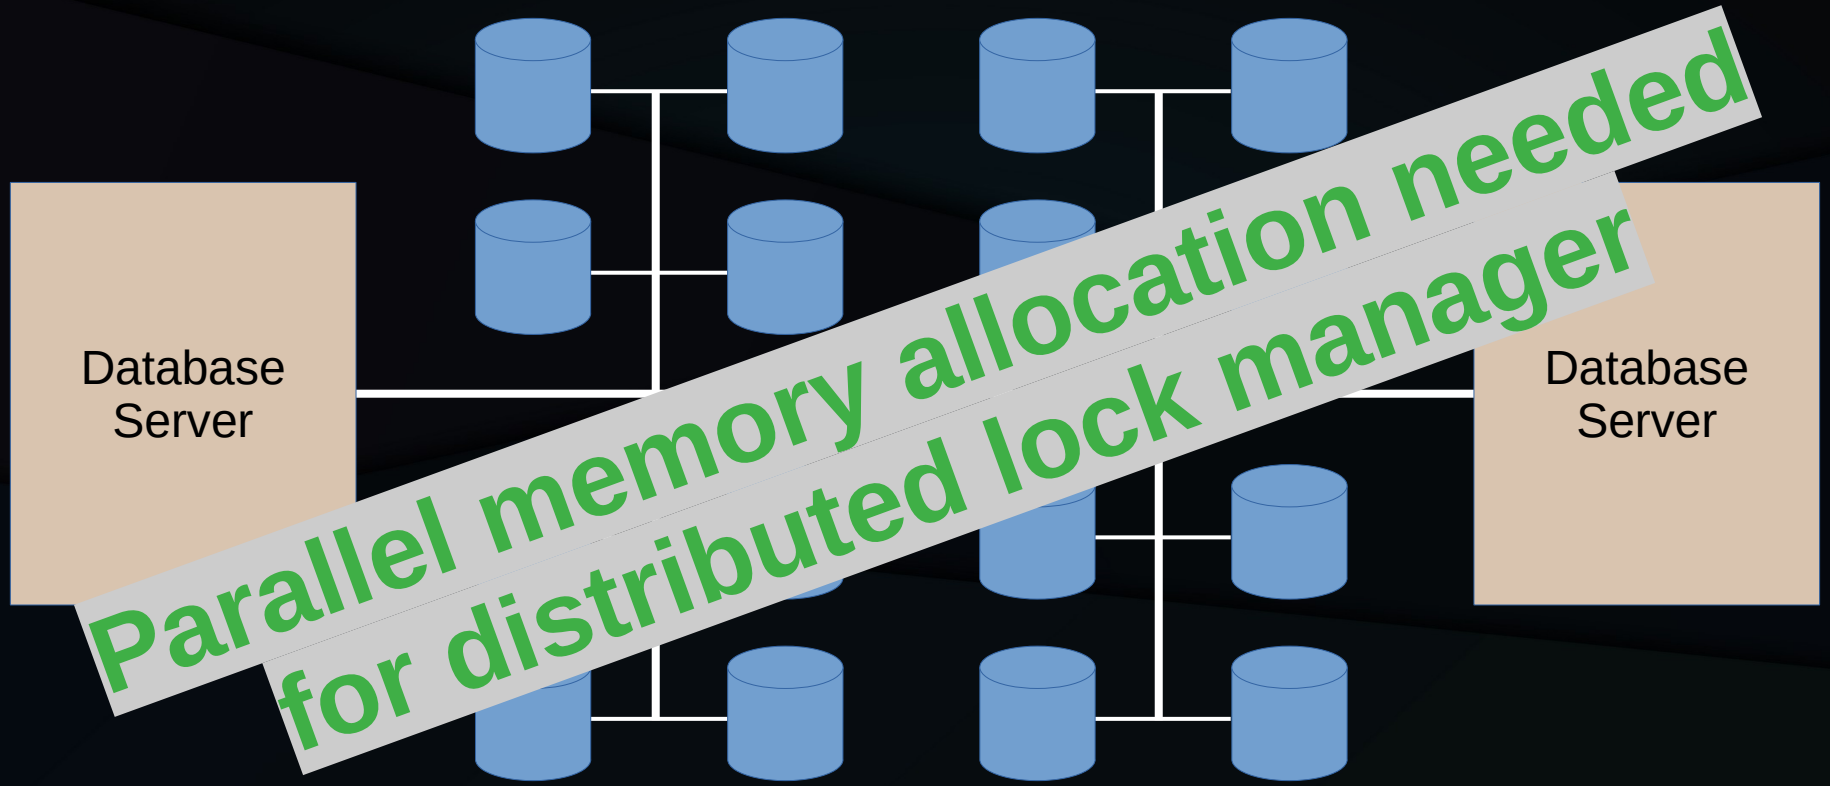

Shared Disks For Availability Win!!!

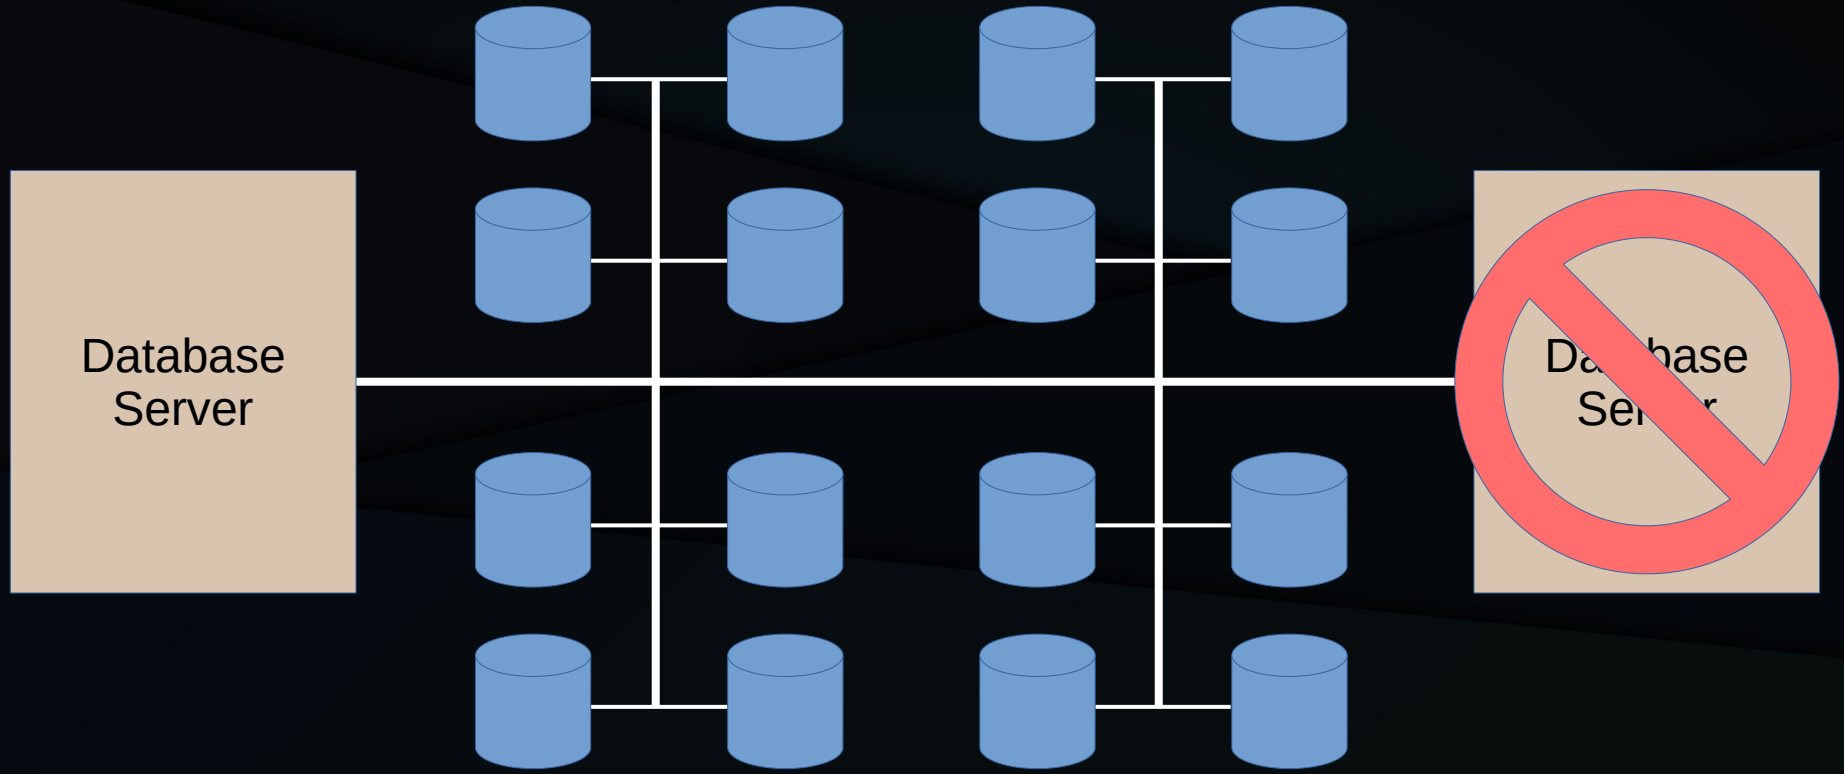

All data is still accessible!!! Of course, sites should test this frequently...

Shared Disks For Availability Win!!!

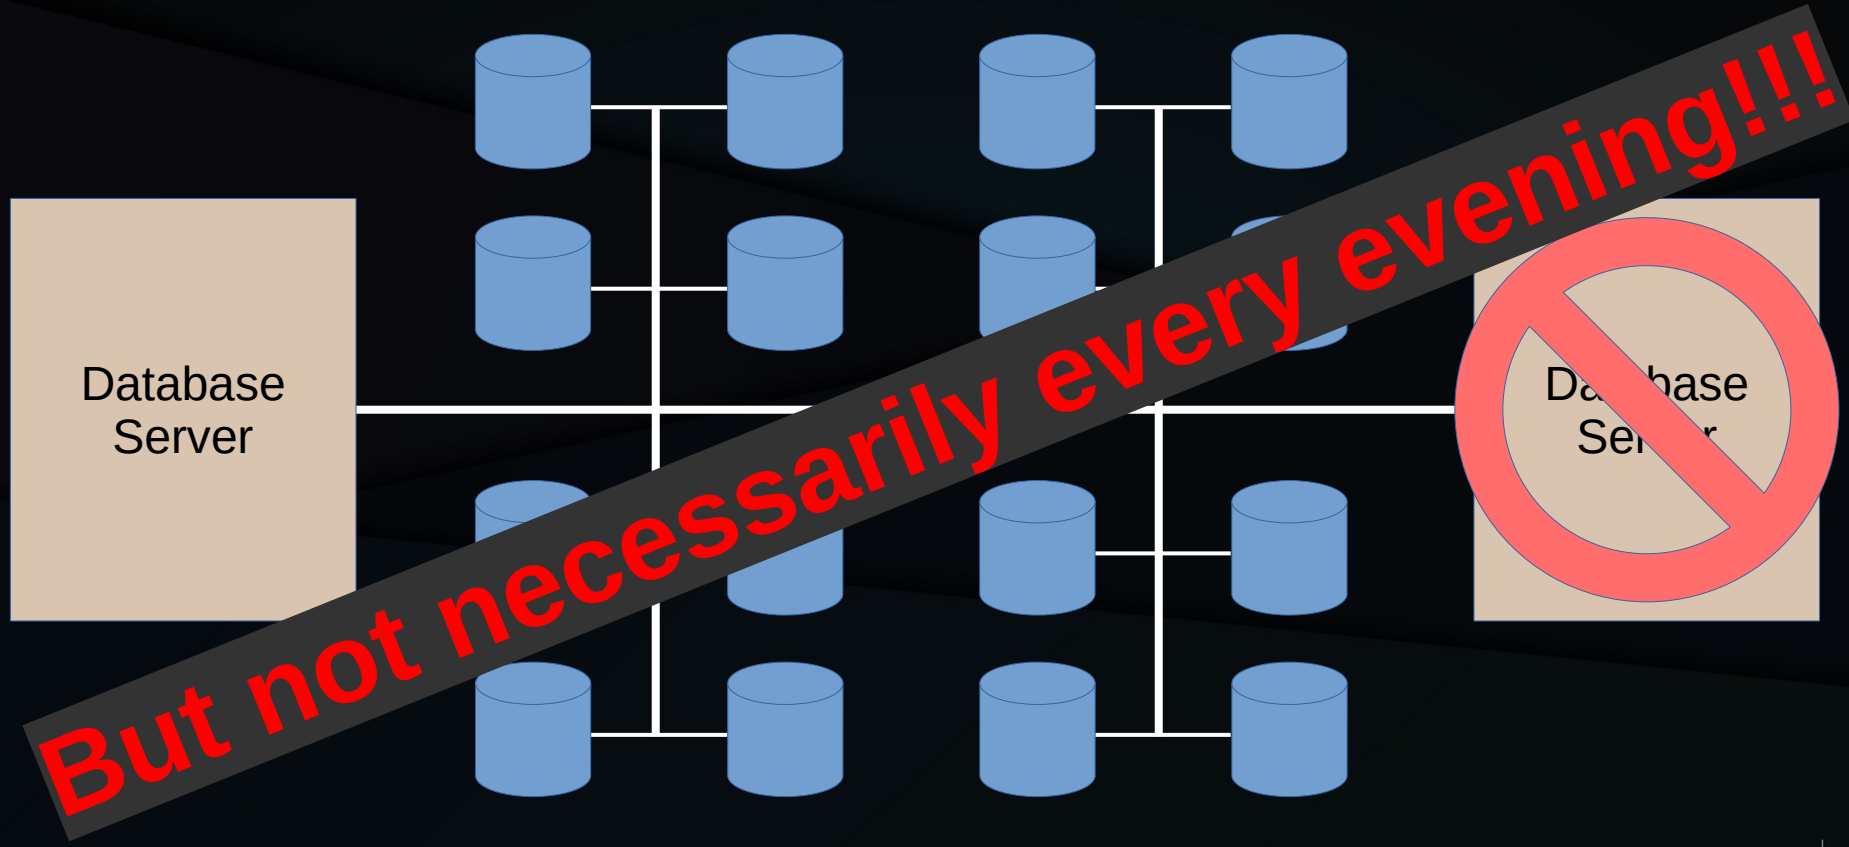

All data is still accessible!!! Of course, sites should test this frequently...

Chaos-Monkey Challenges

- Crash dump was a complete disaster area
  - No hints for on-site debugging instrumentation
- Eventually found test case: 5-27-hour MTBF
  - But need week-long test for any alleged fix!!!

Hint From Stack Trace

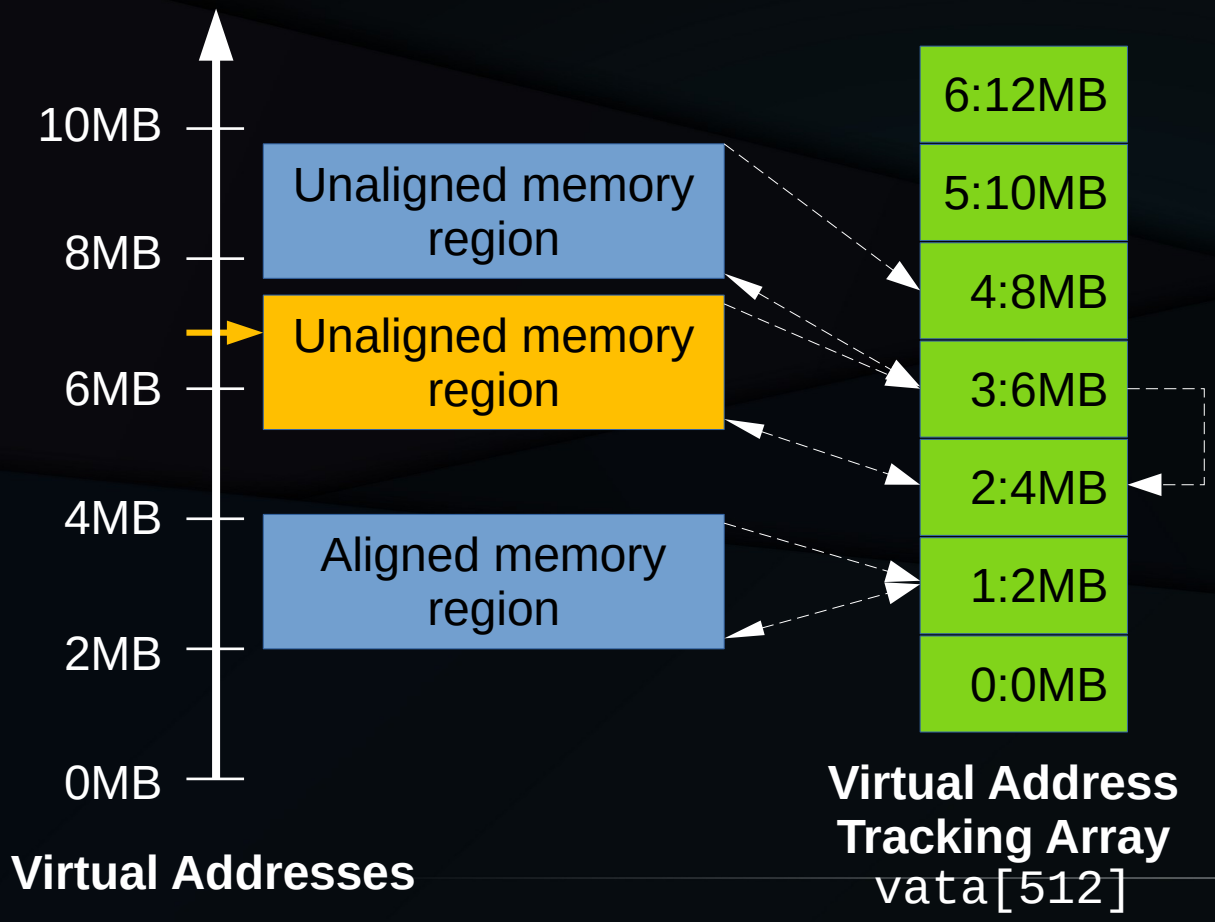


```
int idx = vadr / (2 * MB);
void *vta;

vta = vata[idx];
if (!vta || vadr < vta)
    vta = vta[idx – 1];
```

Hint From Stack Trace: Compiler Fun

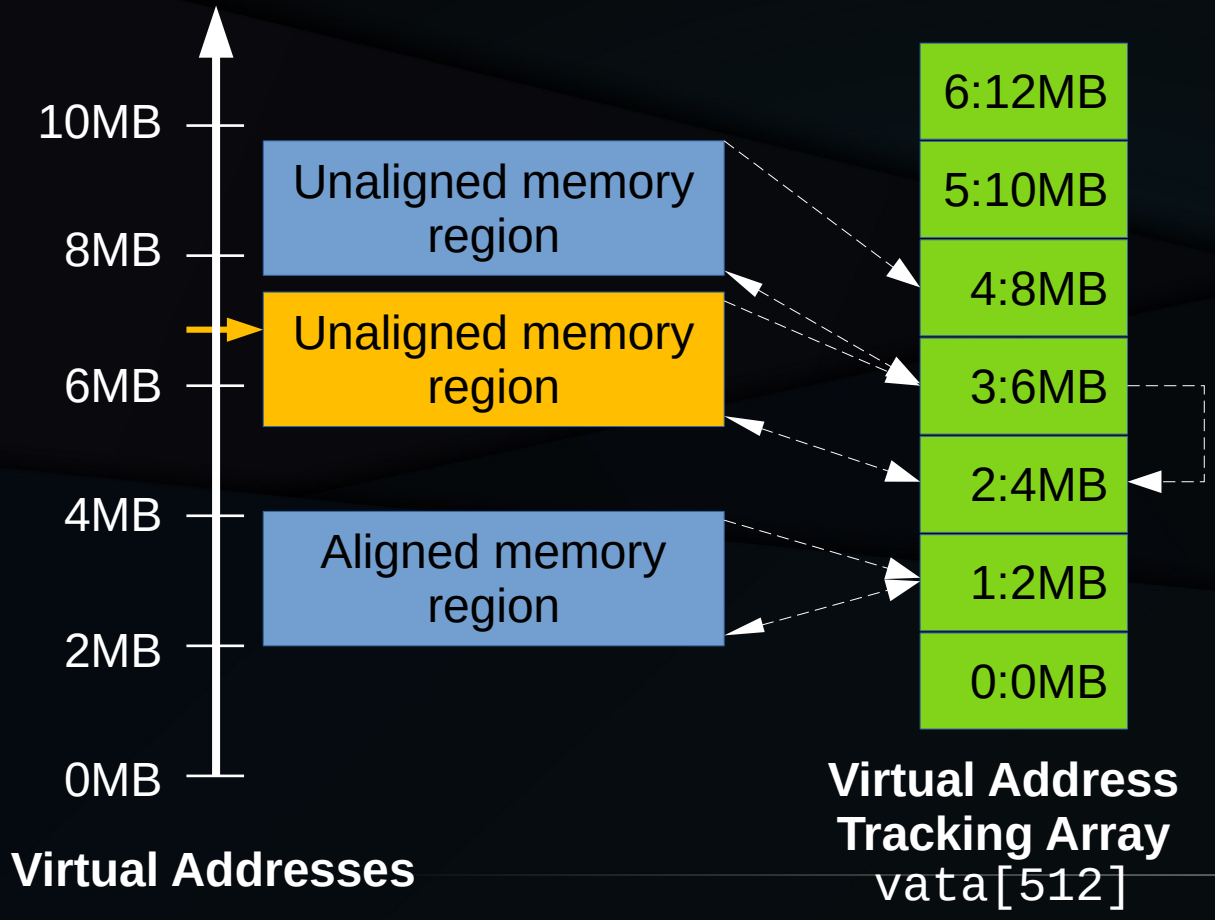


```
int idx = vadr / (2 * MB);
void *vta;

vta = vata[idx];
if (!vata[idx] ||
      vadr < vata[idx])
    vta = vta[idx – 1];
```

Hint From Stack Trace: Compiler Fun

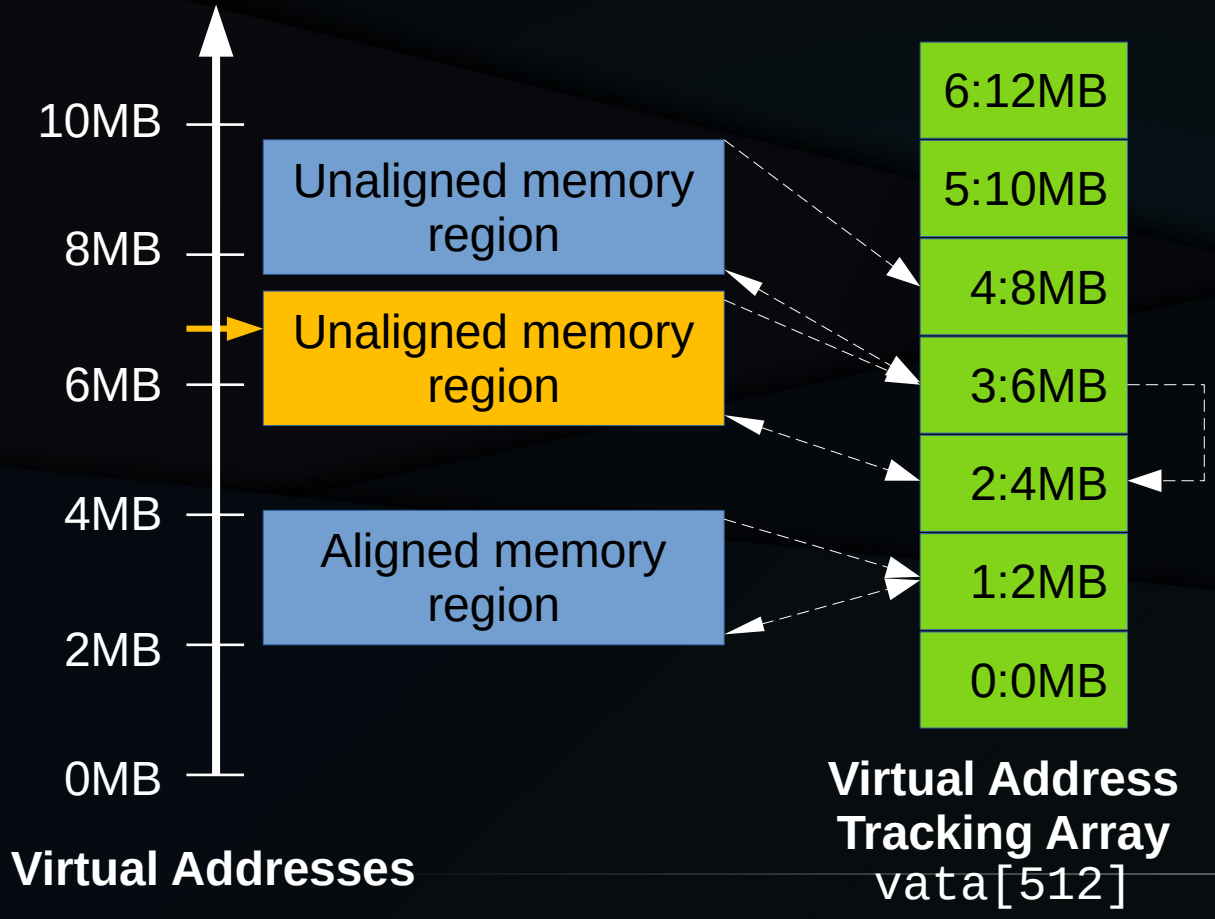


```
int idx = vadr / (2 * MB);
void *vta;

vta = READ_ONCE(vata[idx]);
if (!vta || vadr < vta)
    vta = vta[idx – 1];
```

DYNIX/ptx Memory Allocator ca. 1993

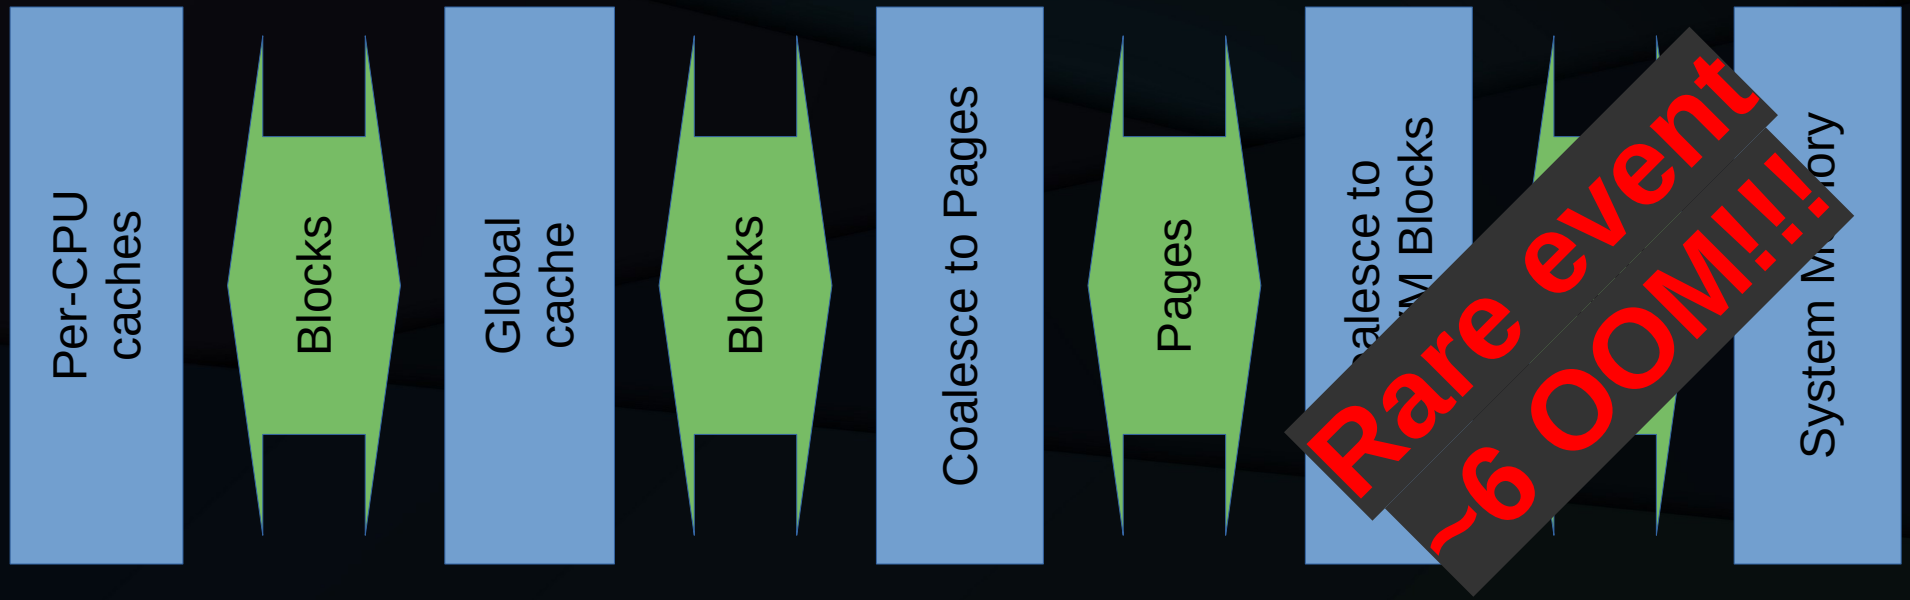

DYNIX/ptx Memory Allocator ca. 1993

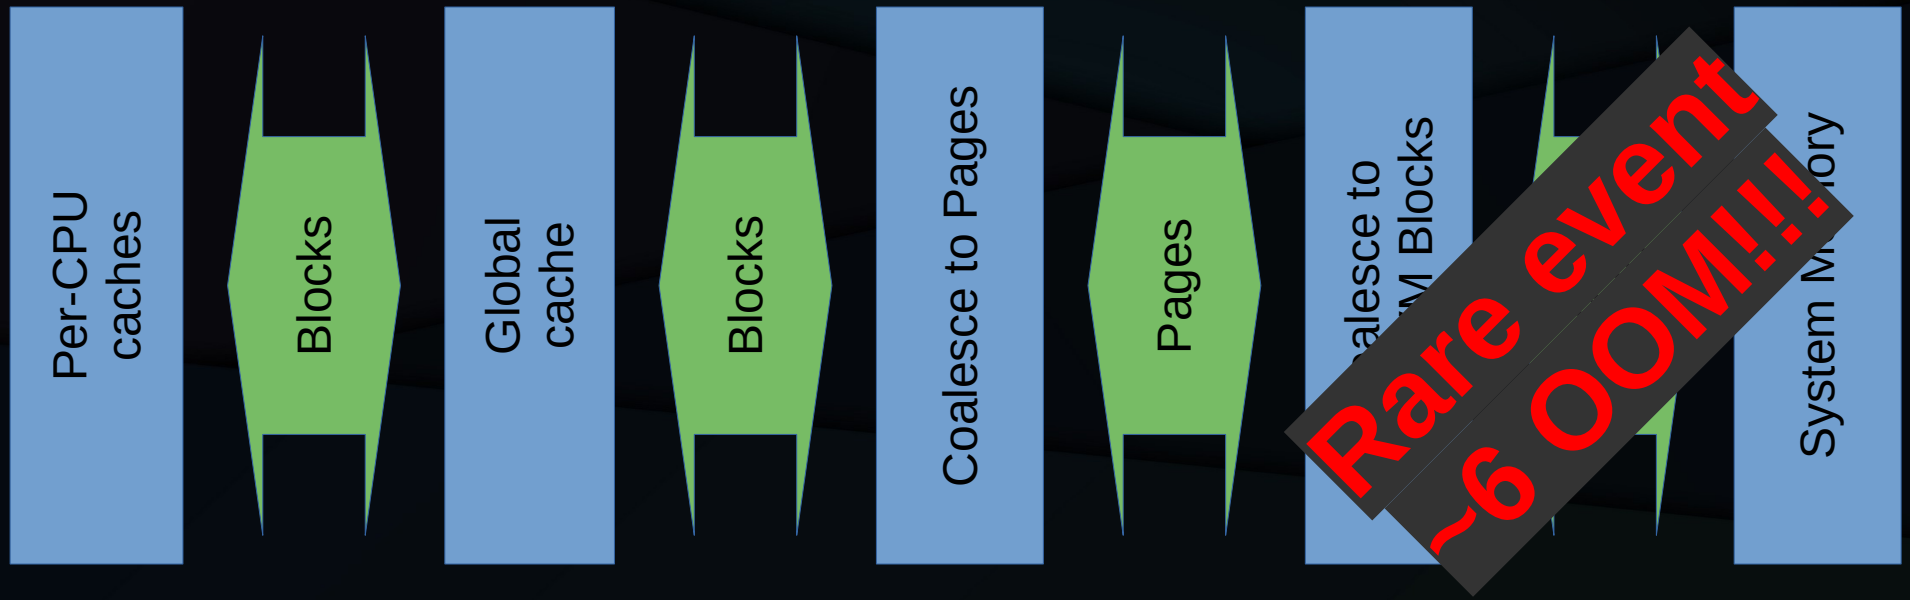

Focused test reduced MTBF to 12 minutes, 1-2 OOM better than stress test

Focused Test vs. Stress Test

	Focused Test	Stress test	Existing testing
MTBF	12 minutes	5 to 27 hours	Infinite?
Basis	Exact bug	Customer workload	Past experience
Hardware	Minimal	A few large systems	Many systems
Development	Day or two	Few person-weeks	Large over years
Applicability	Narrow *	Modest	Wide
Impact	Profound contention	Heavy load	Wide variation

* No I/O, few tasks, modest stress on scheduler, almost no userspace

DYNIX/ptx Memory Allocator ca. 1993

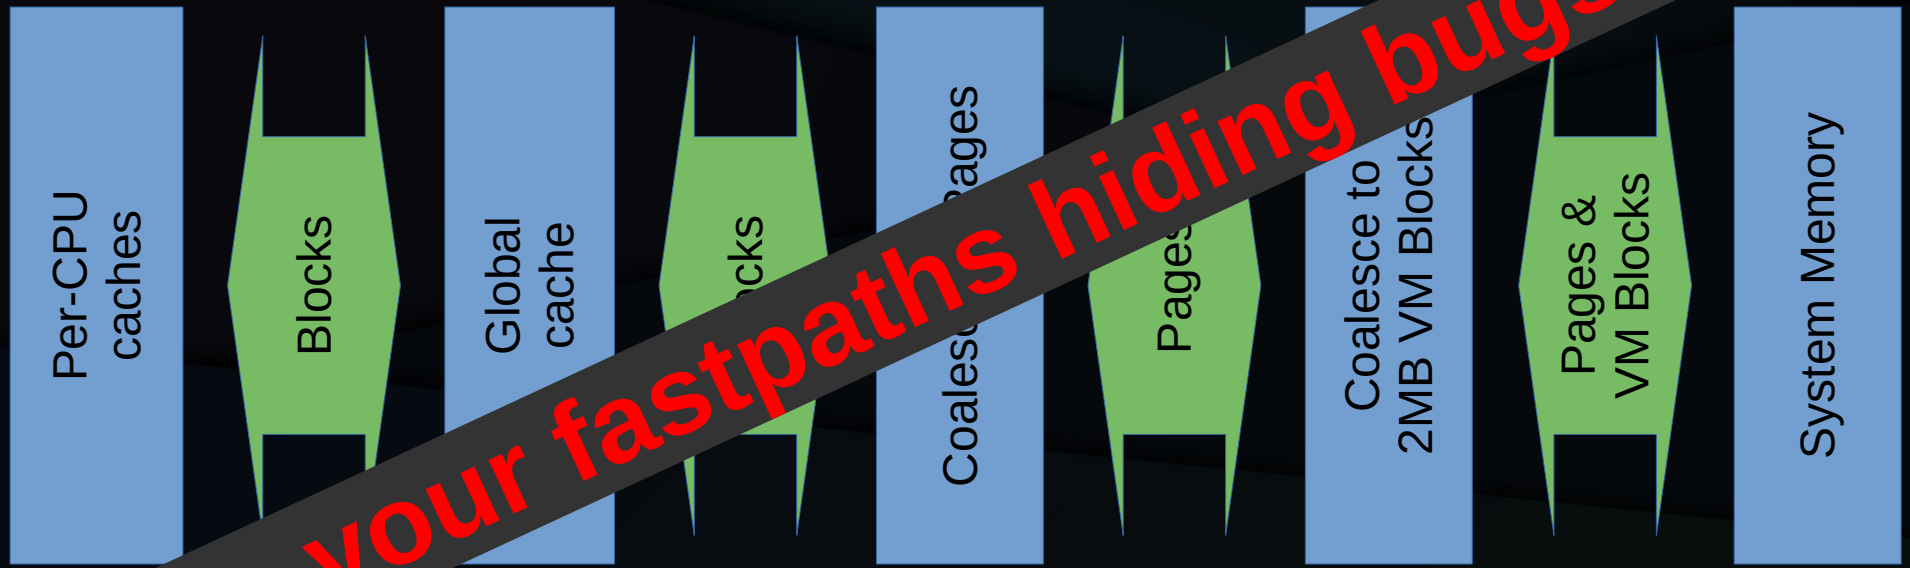

Focused test reduced MTBF to 12 minutes, 1-2 OOM better than stress test

DYNIX/ptx Memory Allocator ca. 1993

Focused test reduced MTBF to 12 minutes, 1-2 OOM better than stress test

DYNIX/ptx Memory Allocator ca. 1993

Focused test reduced MTBF to 12 minutes, 1-2 OOM better than stress test

Safely Disabling Fastpaths: Options

- Run on small systems
  - Four-CPU guest OSes for the win!
- Accept massive contention
- Run code developed for old systems on newer highly integrated systems

Hardware Latency Trends

Year	Sockets	CPUs	CAS Latency (ns)
2008	4	16	95.9
2017	1	56	101.9
2017	4	224	442.9
2022	2	224	147.0

Hardware Latency Trends

Year	Sockets		Latency (ns)
2008	4	6	95.9
2017	1		101.9
201		224	442.9
		224	147.0

Newer systems handle memory contention better

Fewer fastpaths? Larger systems? Decisions, decisions...

Overlapping RCU Readers

```
rcu_read_lock();
preempt_disable();
rcu_read_unlock();
local_irq_disable();
preempt_enable();
local_bh_disable();
local_irq_enable();
local_bh_enable();
```

Overlapping RCU Readers

```
rcu_read_lock();
preempt_disable();
rcu_read_unlock();
local_irq_disable();
preempt_enable();
local_bh
loc
local_bh_enable();
```

Rare combination unless you are running rcutorture

Other Rare Events

- Transitions to and from RCU idle
- CPU hotplug operations (boot and suspend)
- RCU callback flooding
- Memory near-exhaustion
- Transparent hugepage split/coalescing
- And many many more...

Other Rare Events

- Transitions to and from RCU idle
- CPU hotplug operations (…uspend)
- RCU callback fl…
- Memory …
- Transp…nit hu… …alescing
- And many ma…

Combine rare events: multiplicative decreases in MTBF

Other Rare Events

- Transitions to and from RCU idle
- CPU hotplug (suspend)
- RCU callbacks
- Memory pressure
- Transparent hugepages
- And many more

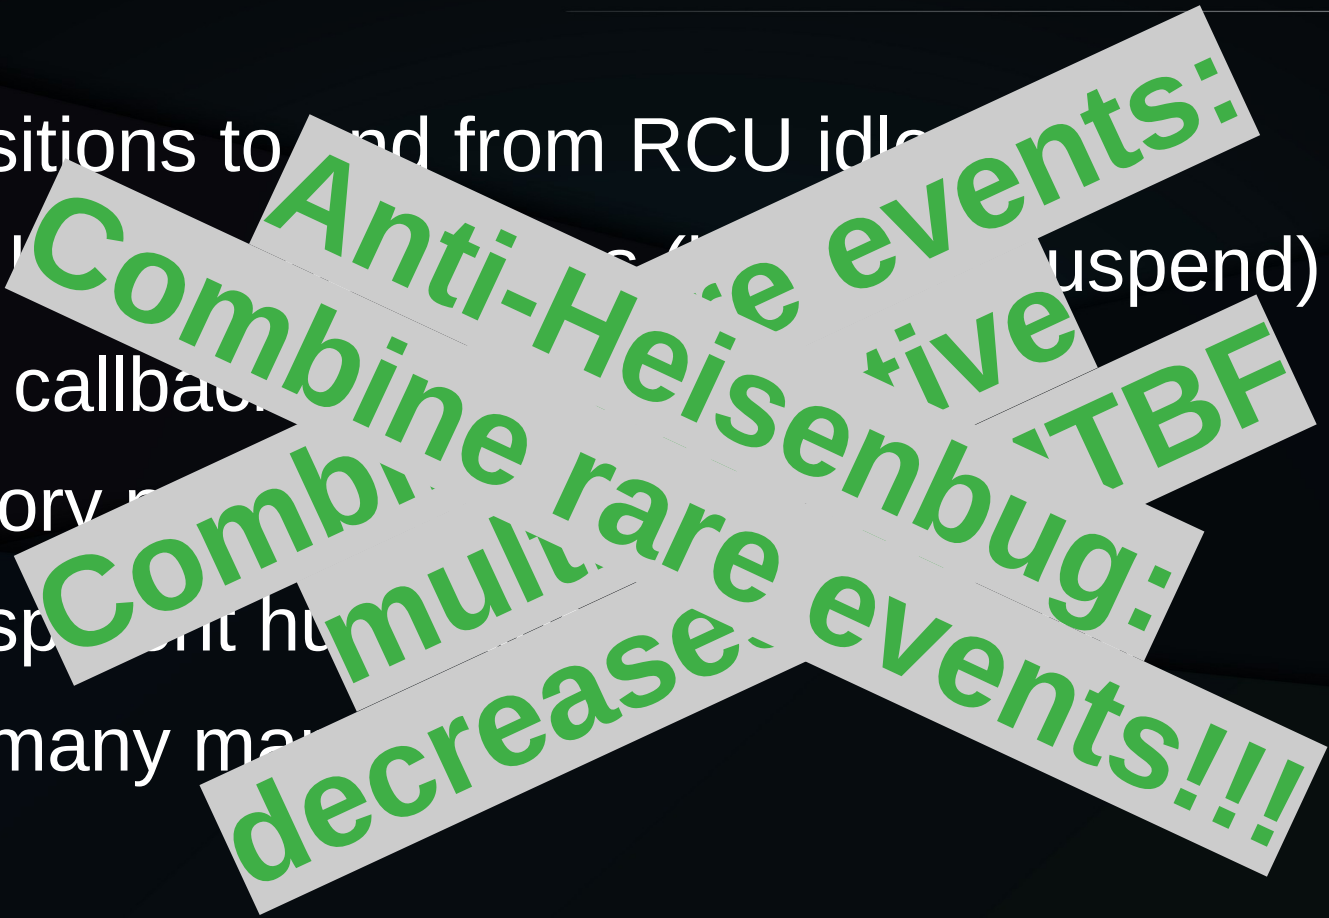

Other Rare Events

- Transitions to and from RCU idle
- CPU ... uspend)
- Transp... hu...
- And many ma...

Anti-... e events:

decrease events!!!

If you must choose, choose the events causing the most trouble

Detect, ***Then*** Instrument

Detect, ***Then*** Instrument

- In theory, code executed after the heisenbug occurs does not affect MTBF
  - In practice, code-size changes can affect MTBF, but this is relatively rare, at least until you count on it
- Two tales of timers…

Cartoon of RCU Grace Periods

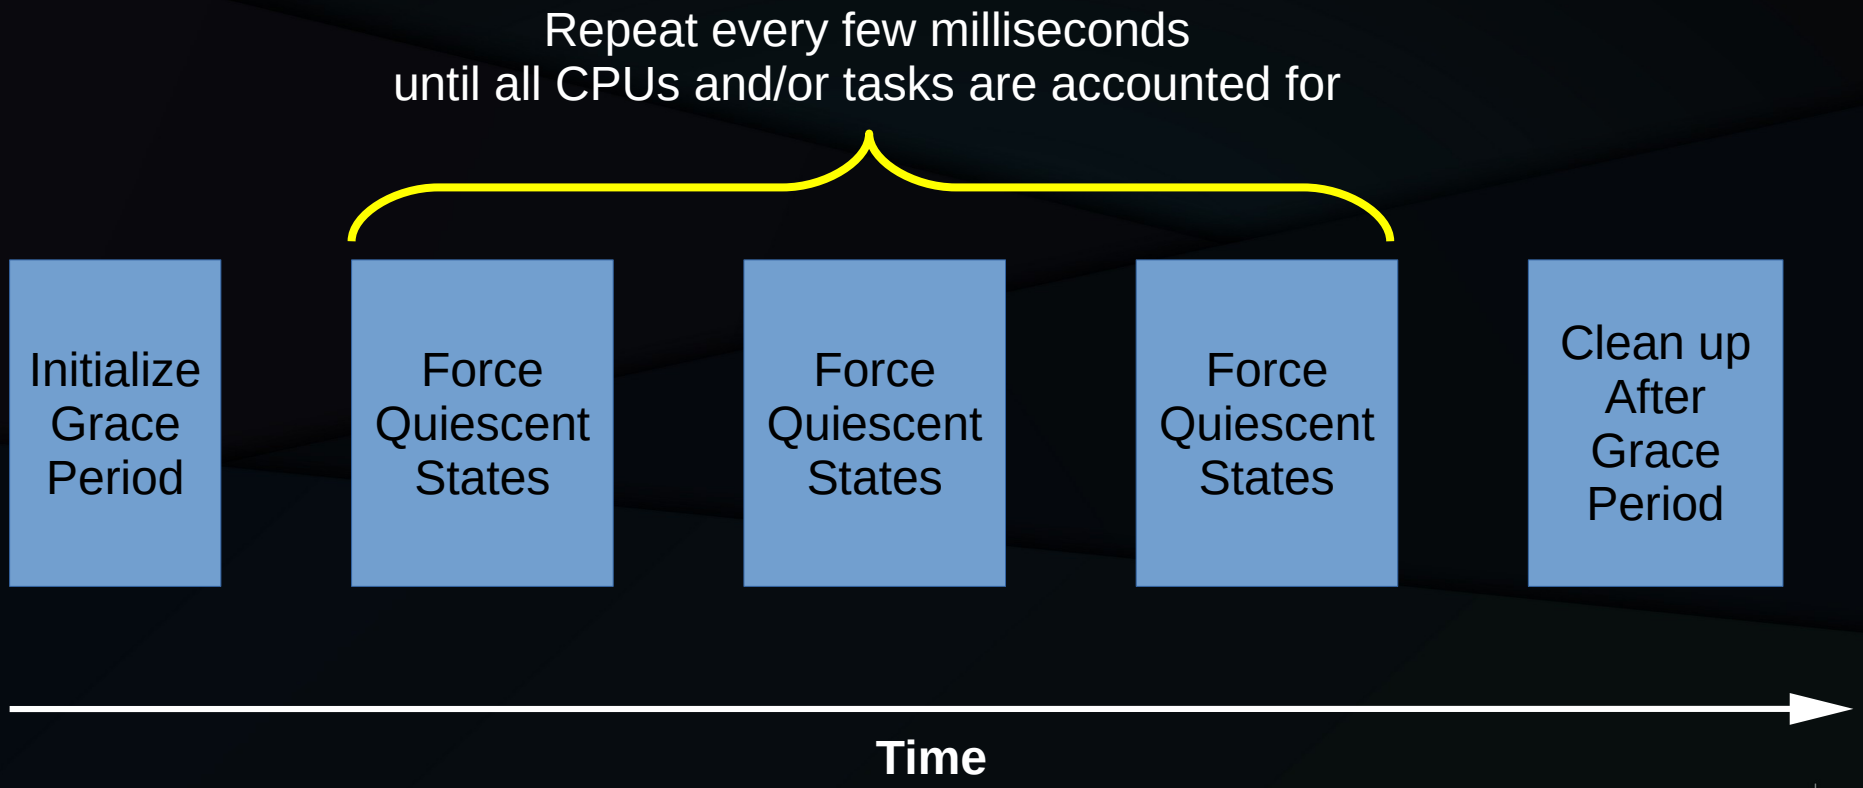

Cartoon of RCU Grace Periods

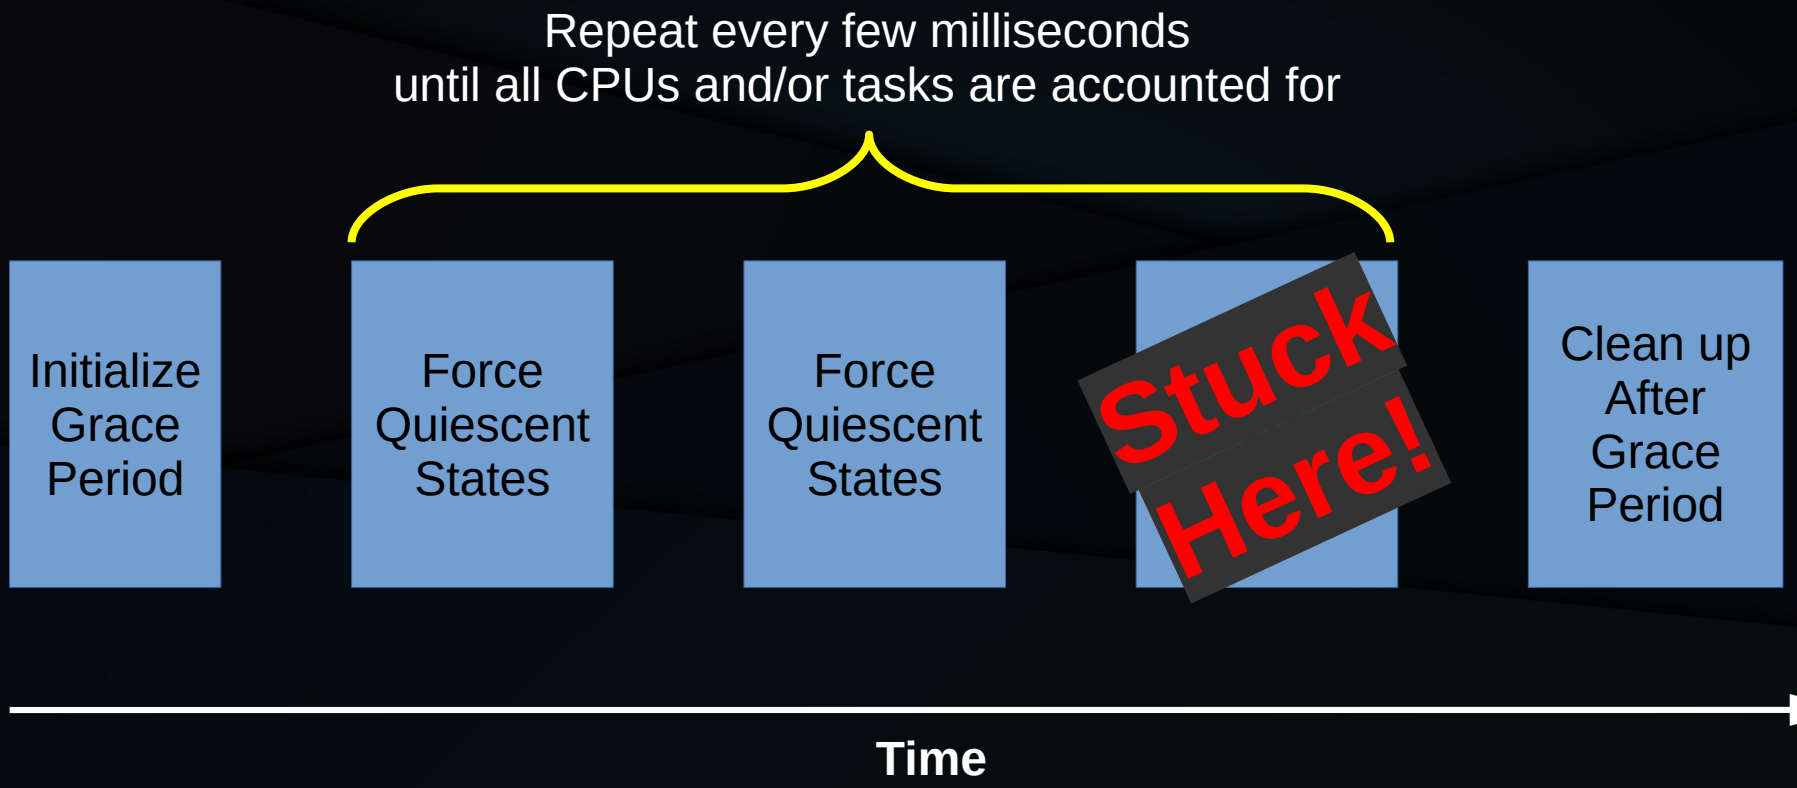

Cartoon of RCU Grace Periods

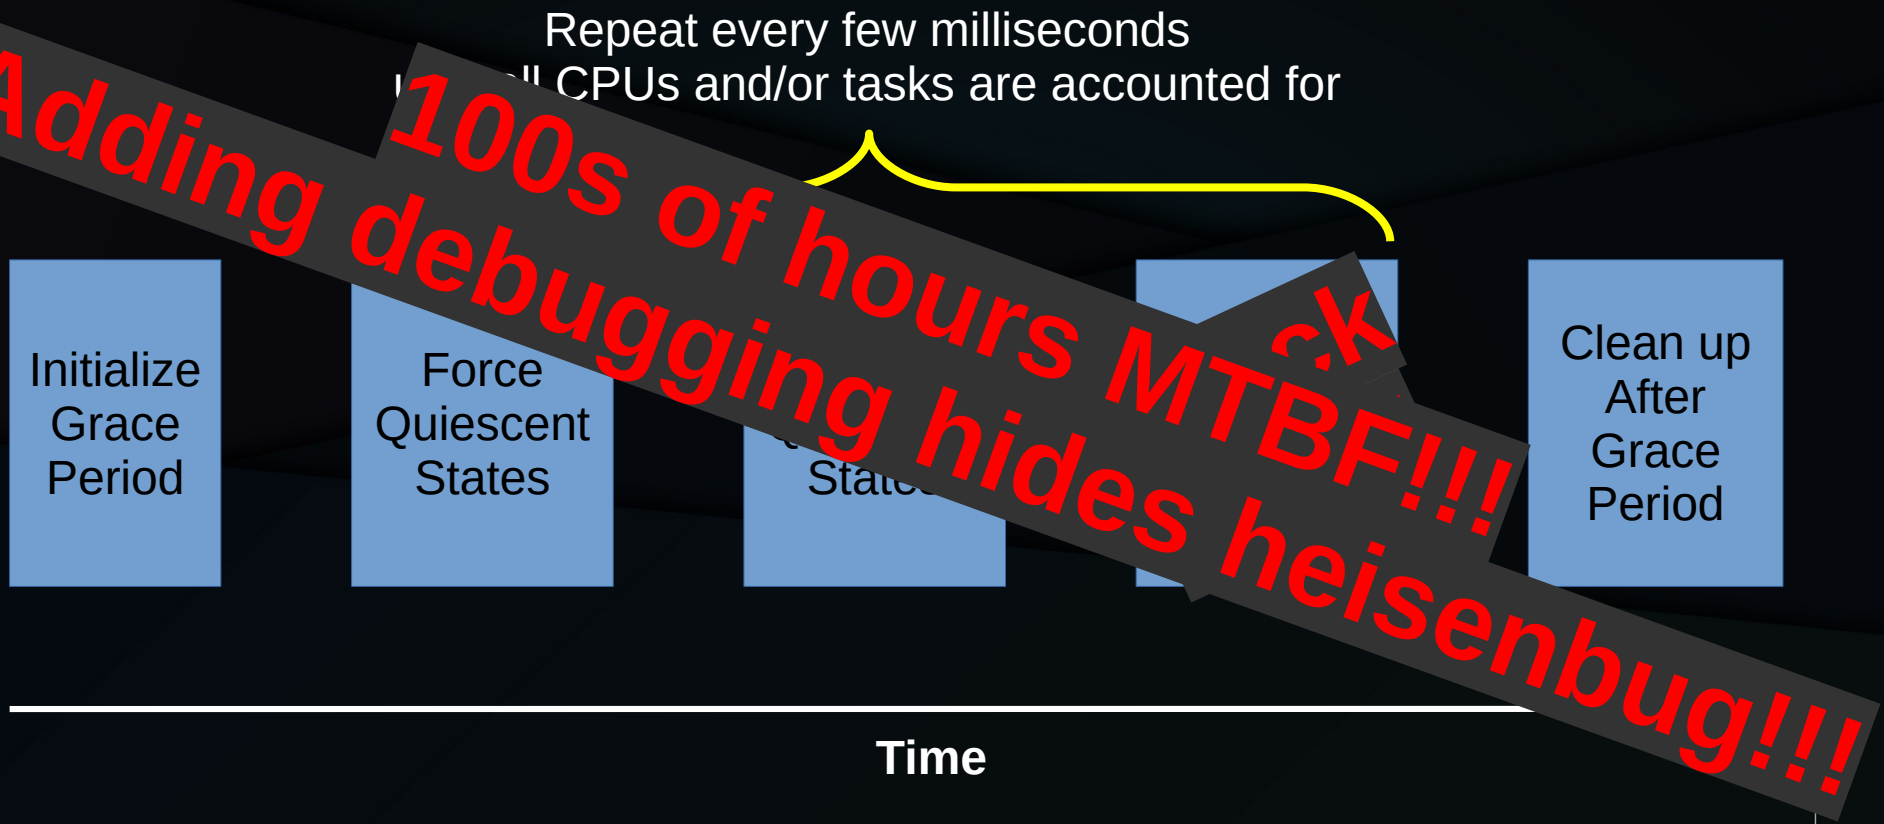

Hunting Stuck-GP Heisenbug

- About a year to reduce MTBF to ~300 hours
  - Choose .config to increase MTBF
  - Increase rate of CPU-hotplug operations
  - Debug still hides heisenbug
  - RCU CPU stall warning restores forward progress

Hunting Stuck-GP Heisenbug

- About a year to reduce MTBF to ~3[illegible]urs
  - Choose .config to increase MT[illegible]
  - Increase rate of CPU-h[illegible] operations
  - Debug still hid[illegible]nbug
  - RCU CB[illegible] warning restores forward progress

Add debug after detection!!!

Adding Debug After Detection

- If timer took more than eight seconds and more than three times as long as was requested, dump debugging information
  - Heuristic, but good enough in this case
  - Bug was due to an interaction between timers, CPU hotplug, and RCU

Adding Debug After Detection

- If timer took more than eight seconds and more than three times as long as was requested, dump debugging information
  - Heuristic, but good enough in this case
  - Bug was due to an interaction between timers, CPU hotplug, and RCU
  - Recent bug between workqueues and RCU?

Adding Debug After Detection

- If timer took more than eight seconds and more than three times as long as was requested, dump debugging information
  - Heuristic, but good for this case
  - Bug was due to interaction between timers, CPU hotplug, and RCU
- Prevent bug between workqueues and RCU?

Post-facto debug did not increase MTBF!!!

Adding Debug After Detection

- If tim… more than eight secon… more than thr… as long as w… quested, dump de…
  - Heuristic, but … this case
  - Bug was … timers, CPU hotp… RCU
  - …ent bug between workqueues … RCU

Anti-heisenbug: Add debugging after bug is detected!!!

Post-facto det… M… not increase

Taking This One Step Further...

- When rare combination of events takes system to a legal but vulnerable state, start the system in that state
  - The nuclear option: White-box testing
  - Exhaustive state testing of `rcu_segcblist`
    - Done in userspace

Taking This One Step Further...

- When rare combination of events [illegible] system to a legal but vulnerable stat[illegible] the system in that state
  - The nuclear opti[illegible]esting
  - Exhausti[illegible] of `rcu_segcblist`
    - [illegible]e

Anti-heisenbug: Force rare risky legal states!!!

Heisenbugs: The Goal

Heisenbugs: The Goal

- What is better than being proficient at hunting heisenbugs?

Heisenbugs: The Goal

- What is better than being proficient at hunting heisenbugs?
- **Not having heisenbugs in the first place!!!**

How to Avoid Hunting Heisenbugs

How to Avoid Hunting Heisenbugs

- No easy way out, but:
  - Careful concurrency-first design
  - Thorough unit testing, including stress testing
  - Thorough integration testing
  - Stringent code-review process
  - Verification, if applicable

How to Avoid Hunting Heisenbugs

- No easy way out, but:
  - Careful concurrency-first desi
  - Thorough unit testi sting
  - Thorough i
  - S process
  - Verific , if applicable

Anti-heisenbug: Don't randomly hack concurrent code!!!

Summary

Summary: How to Hunt Heisenbugs

- Create anti-heisenbugs
  - Reduce MTBF
  - Increase workload intensity
  - Look for and promote trouble
  - Inject strategic delays
  - Count near misses
  - Swamp queues
  - De-emphasize fastpaths
  - Combine rare events
  - Add debugging ***after*** bug is detected
  - Force rare risky legal states (whitebox)

- Avoid (many) heisenbugs
  - Careful concurrency-first design
  - Thorough unit testing, including stress testing
  - Thorough integration testing
  - Stringent code-review process
  - Verification, if applicable
  - Don’t randomly hack concurrent code!!!

Summary: How to Hunt Heisenbugs

- Create anti-heisenbugs
  - Reduce MTBF
  - Increase workload intensity
  - Look for and promote trouble
  - Inject strategic delays
  - Count near misses
  - Swamp queues
  - De-emph
  - Combi
  - Add debugging ***after*** bu tected
  - Force rare risky legal states (whitebox)

- Avoid (many) h ugs
  - Caref y-first design
  - it testing, including
  - integration testing
  - Stringent code-review process
  - Verification, if applicable
  - Don't randomly hack concurrent code!!!

No "silver bullet", but many useful techniques

Questions?

Backup

Finite Requests into Finite Queue

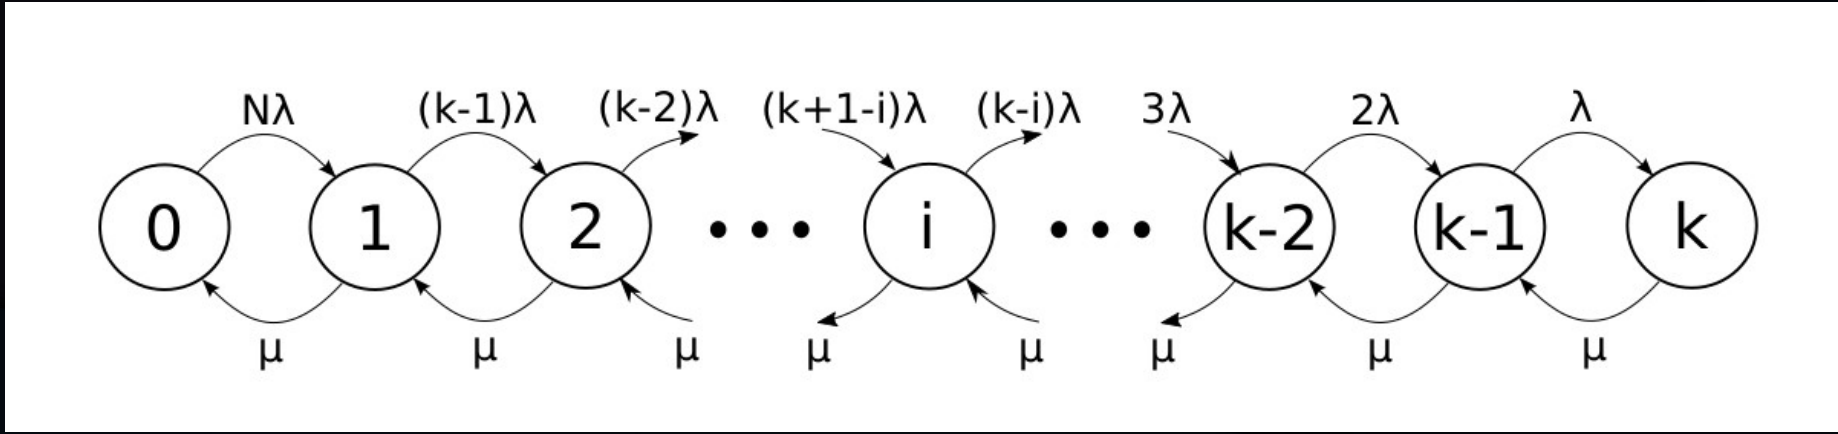

Finite Requests into Finite Queue