



# Cautionary Tales on Implementing the Software That People Want

*Be Careful What You Wish For. You Might Get It!!!*

# How Did Paul Get This Way?

- High school class: IBM mainframe & HP Basic (1973-1976)
- University: Computer science & mechanical engineering, business applications (1976-1981)
  - Started supporting self by coding in Summer June 1977
- Contract programming (1981-1985)
- Systems administration and Internet research (1986-1990)
- Concurrent proprietary UNIX (1990-2000)
- Linux kernel concurrency and realtime (2001-present)


# Cautionary Quotes

---

- The first secret of getting what you want is knowing what you want. *Arthur D. Hlavaty*
- If you don't know what you want, you will probably never get it. *Oliver Wendell Holmes, Jr.*
- If you don't know what you want, you end up with a lot you don't. *Chuck Palahniuk*


# 1990: Stochastic Fairness Queueing

# 1990: Stochastic Fairness Queueing




**Advanced Gateway Server (AGS) router, Cisco, US, 1986**  
This Cisco router was able to map one network protocol into another. The software was originally developed by Bill Yeager at Stanford, then licensed and enhanced by Cisco (from "San Francisco") founders Len Bosack and Sandy Lerner.  
*Gift of Alex Boshannak, X1539 98*

# 1990: Queueing Problem




1 Megabit network is *fast*.

# 1990: Fair Queueing



Give each flow its own queue!  
Yeah, you and how many 10MHz CPUs???


# 1990: Fair Queueing



Give each flow its own queue!  
Yeah, you and how many 10MHz CPUs???




# 1990: Stochastic Fair Queueing: Hash



Hash IP-address/Port quadruple for wondrous end-to-end fairness!!!


# 1990: Paul's Internet Vision



Hash IP-address/Port quadruple for wondrous end-to-end fairness!!!

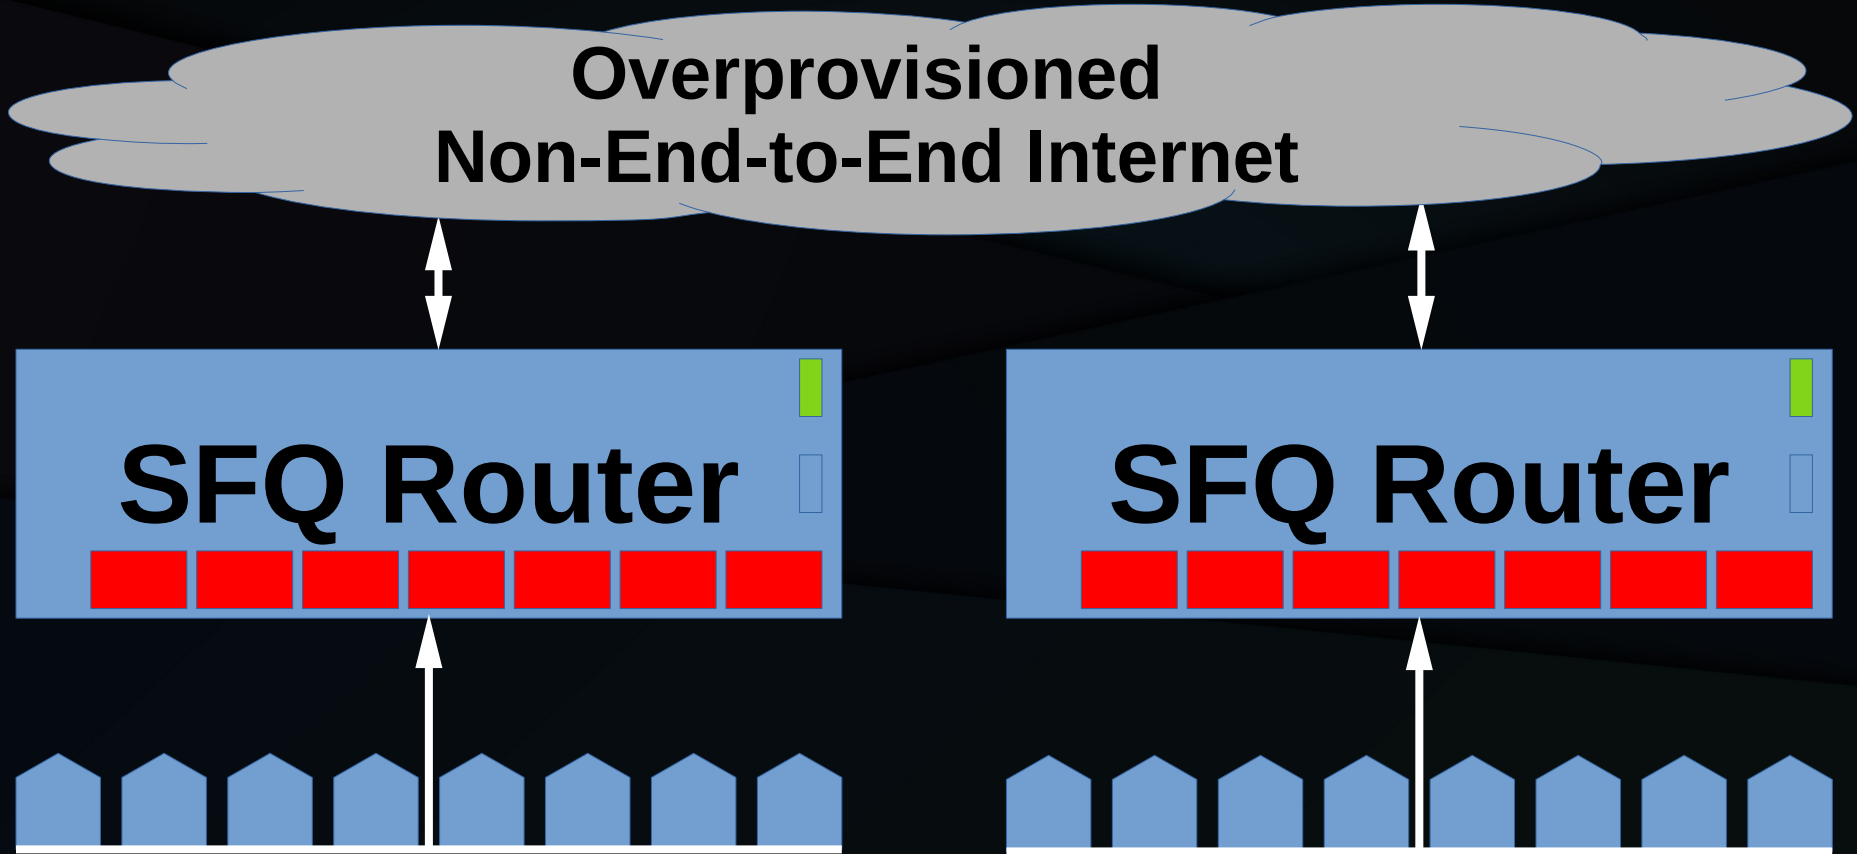
# 1990: Paul's Internet Vision

**Delusion**



Hash IP-address/Port quadruple for wondrous end-to-end fairness!!!

# 1990: What Internet Did Instead



Internet gateways hash Ethernet MAC addresses for approximate real-world fairness.

# 1990 SFQ: What Went Wrong?

- Solved wrong problem: End-to-end fairness
  - Correct problem: Hop-by-hop & endpoint fairness
  - By sheer dumb luck, my algorithm handled both
- Research-quality code: Get the paper out!!!
  - Cisco and Linux-kernel engineers fixed this
- Used heavily until about 2015
  - FQ-CODEL and CAKE now address bufferbloat
  - Dave Taht, Eric Dumazet, Toke Høiland-Jørgensen, ...

# 1990 SFQ: What Went Wrong?

- Solved wrong problem: End-to-end fairness
  - Correct problem: Hop-by-hop & endpoint fairness
  - By sheer dumb luck, my algorithm solved both
- Research-quality code: Got paper out!!!
  - Engineers at Cisco and Linux kernel fixed this
- Used heavily until about 2015
  - FQ-CODES and CAKE now address bufferbloat
  - Dave Taht, Eric Dumazet, Toke Høiland-Jørgensen, ...

**Bad idea badly implemented**

# 1990 SFQ: What Went Wrong?

- Solved wrong problem: End-to-end fairness
  - Correct problem: Hop-by-hop & endpoint fairness
  - By sheer dumb luck, my algorithm worked
- Research-quality code: Got it wrong at!!!
  - Engineers at Cisco and others later fixed this
- Used heavily unorthodox ideas
  - FQ-CODES: no address bufferbloat
  - Dave Dagon, Amir A. Amazez, Toke Høiland-Jørgensen, ...

**Bad idea badly implemented,  
resuscitated by dumb luck**

# 1990 SFQ: What Went Wrong?

- Solved wrong problem: End-to-end fairness
  - Correct problem: Hop-by-hop & endpoint fairness
  - By sheer dumb luck, algorithm worked!!!
- Research-quality implementation
  - Engineers at Cisco fixed this
- Used heavily unproven ideas
  - FQ-CODE
  - Dave ... , Toke Høiland-Jørgensen ...

**Prematurely implemented,  
root of abstraction is the  
bad idea b/c of all evil  
resuscitated, dumb luck**



# 1990 SFQ: What Instead?

- Solved wrong problem: End-to-end fairness
  - Correct problem: Hop-by-hop & endpoint fairness
  - By sheer dumb luck, my algorithm happened to work
- Research-quality code: Get the code out!!!
  - Engineers at Cisco and in the kernel fixed this
- Used heavily until 2015
  - FQ-CODEL now address bufferbloat
  - Dave Taht, Nicolas Dumazet, Toke Høiland-Jørgensen, ...

**Live among your users!!!**

# 1980s: Eight-Bit CRM



# 1980s: Eight-Bit CRM

- CRM application built to spec under contract
- The company loved it!

# 1980s: Eight-Bit CRM

- CRM application built to spec under contract
- The company loved it!
- Their prospective customers, not so much

# 1980s: Eight-Bit CRM

- CRM application built to spec under contract
- The company loved it!
- Their prospective customers, not so much
- Dumb luck: They paid me *before* bankruptcy

# 1980s: Eight-Bit CRM

- CRM application built to spec und
- The company loved it!
- Their prospective customers, not so much
- Dumb luck: The company failed *before* bankruptcy

**Bad idea well implemented**

# 1980s: Eight-Bit CRM

- CRM application built to spec und
- The company loved it!
- Their prospective cust so much
- Dumb luck: The before bankruptcy

**Bad idea well implemented,  
but hey, I got paid???**

# 1980s: Eight-Bit CRM: What Instead?

**Time and Grade: Experience**




# 1980s: Acoustic Navigation




# 1980s: Acoustic Navigation

**But with insane quantities  
of shock mounting for  
shipboard use**



# 1980s: Acoustic Navigation (Pre-GPS)



...then convert time to distance and triangulate!!!


# Acoustic Navigation Complications

- If the ship's position was known when deploying the transponder, there would be no need for acoustic navigation
- Transponders do not fall exactly straight down through four miles of water
- Ocean surface is not perfectly level
- Sound does not travel in a straight line through ocean water
- Sound does not travel at a uniform speed through ocean water
- Dolphins like to play with transponders

# Acoustic Navigation Complications

- If the ship's position was known when deploying the transponder, there would be no need for acoustic navigation
- Transponders do not fall exactly straight down through four miles of water
- Ocean surface is not perfectly level
- Sound does not travel in a straight line through ocean water
- Sound does not travel at a uniform speed through ocean water
- Dolphins like to play with transponders


# Acoustic Navigation Calibration (1/2)



# Acoustic Navigation Complications


- If the ship's position was known when deploying the transponder, there would be no need for acoustic navigation
- Transponders do not fall exactly straight down through four miles of water
- Ocean surface is not perfectly level
- Sound does not travel in a straight line through ocean water
- Sound does not travel at a uniform speed through ocean water
- Dolphins like to play with transponders

# Acoustic Navigation Calibration (2/2)







# Acoustic Navigation Calibration (2/2)




# Acoustic Navigation Calibration (2/2)




# Acoustic Navigation Calibration (2/2)



# Acoustic Navigation Calibration (2/2)



# Acoustic Navigation Calibration (2/2)



Then calculate sound velocity as a function of depth, and finally do ray-tracing.


# Acoustic Navigation Complications

- If ship's position known when deploying transponder, no need for system
- Transponders do not fall exactly straight down through four miles of water
- Ocean surface is not perfectly level
- Sound does not travel in a straight line through ocean water
- Sound does not travel at a uniform speed through ocean water
- Dolphins like to play with transponders
- Error minimization has difficulty with three unknowns per transponder

# Acoustic Navigation Complications


- If ship's position known when deploying transponder, no need for system
- Transponders do not fall exactly straight down through four miles of water
- Ocean surface is not perfectly level
- Sound does not travel in a straight line through ocean water
- Sound does not travel at a uniform speed through ocean water
- Dolphins like to play with transponders
- Error minimization has difficulty with three unknowns per transponder

# Acoustic Navigation: Measure Depth







# Acoustic Navigation: Measure Depth




# Acoustic Navigation: Measure Depth




# Acoustic Navigation: Measure Depth




# Acoustic Navigation: Measure Depth




# Acoustic Navigation: Measure Depth




# Acoustic Navigation: Measure Depth



# Acoustic Navigation: Measure Depth



# Acoustic Navigation: Measure Depth





# Cautionary Quote

---

- "Everyone knows that debugging is twice as hard as writing a program in the first place. So if you're as clever as you can be when you write it, how will you ever debug it?" - *Brian W. Kernighan*, "The Elements of Programming Style", 2nd Edition, Chapter 2.

# Cautionary Quote

- "Everyone knows that debugging is twice as hard as writing a program in the first place. So if you're as clever as you can be when you write it, how will you ever debug it?" - *Brian W. Kernighan*, "The Elements of Programming Style", 2nd Edition, Chapter 2.
- While programming, you are living in blissful ignorance of important requirements. These requirements make themselves known during debugging.
- Which is but one cause of Kernighan's observation.

# Cautionary Quote

- "Everyone knows that debugging is twice as hard as writing a program in the first place. So if you're as clever as you can be when you write it, how will you ever debug it?"  
*Kernighan*, "The Elements of Programming Style", 2nd Edition, Chapter 2.
- While programming, successful ignorance of important requirements make themselves known during development.
- Which is a direct consequence of Kernighan's observation.

**I failed to understand that I was competing with a file cabinet**

# Cautionary Quote

- "Everyone knows that debugging is twice as hard as writing a program in the first place. So if you're as clever as you can be when you write it, you will never debug it."  
*Kernighan, "The Elements of Programming Style", 2nd Edition, Chapter 2.*
- While programming, I was in a state of successful ignorance of important requirements. I should have made myself known during the development phase.
- Which is a direct consequence of Kernighan's observation.

**And the file cabinet that I was  
I failed to undo a file cabinet  
competing with a file cabinet won**

# 1970s: Student Housing System



# 1970s: Student Housing System



# 1970s: Student Housing System




# 1970s: Student Housing System






# Student Housing Temporal Confusion




# Student Housing Temporal Confusion

Student started on Friday and was not amused by the bill.



# Student Housing Temporal Confusion

Student started on Friday and was not amused by the bill. My manager had the usual 1970s earthy suggestion for alternative uses of the money.



# Student Housing Temporal Confusion

Student started on Monday and was not paid by the bill. My mother had the usual earthy suggestion for alternative uses of the money.

**Problem: Months vary in length**

Mon.

Tue.

Thu.

Fri.

Sat.

Sun.

\$

\$

\$

\$\$\$\$\$

\$

\$

# Student Housing Temporal Confusion

Student started on Monday and was not paid for the bill. My usual suggestion for the use of the money.



# Student Housing Temporal Confusion

Student started on Friday and was not able to pay the bill. My manager gave the usual 10% discount as a suggestion for creative uses of the money.




**Good idea implemented poorly**

# 1990s: Clustered Database Servers




# Shared Disks For Availability Win!!!






# Shared Disks For Availability Win!!!




All data is still accessible!!!

# 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
- Unable to reproduce in the lab

# Chaos-Monkey Challenges

- Crash dump was a complete disaster area
  - No hints for on-site debugging instrumentation
- Unable to reproduce in the lab
- Eventually, found test case: 5-27-hour MTBF

# Chaos-Monkey Challenges

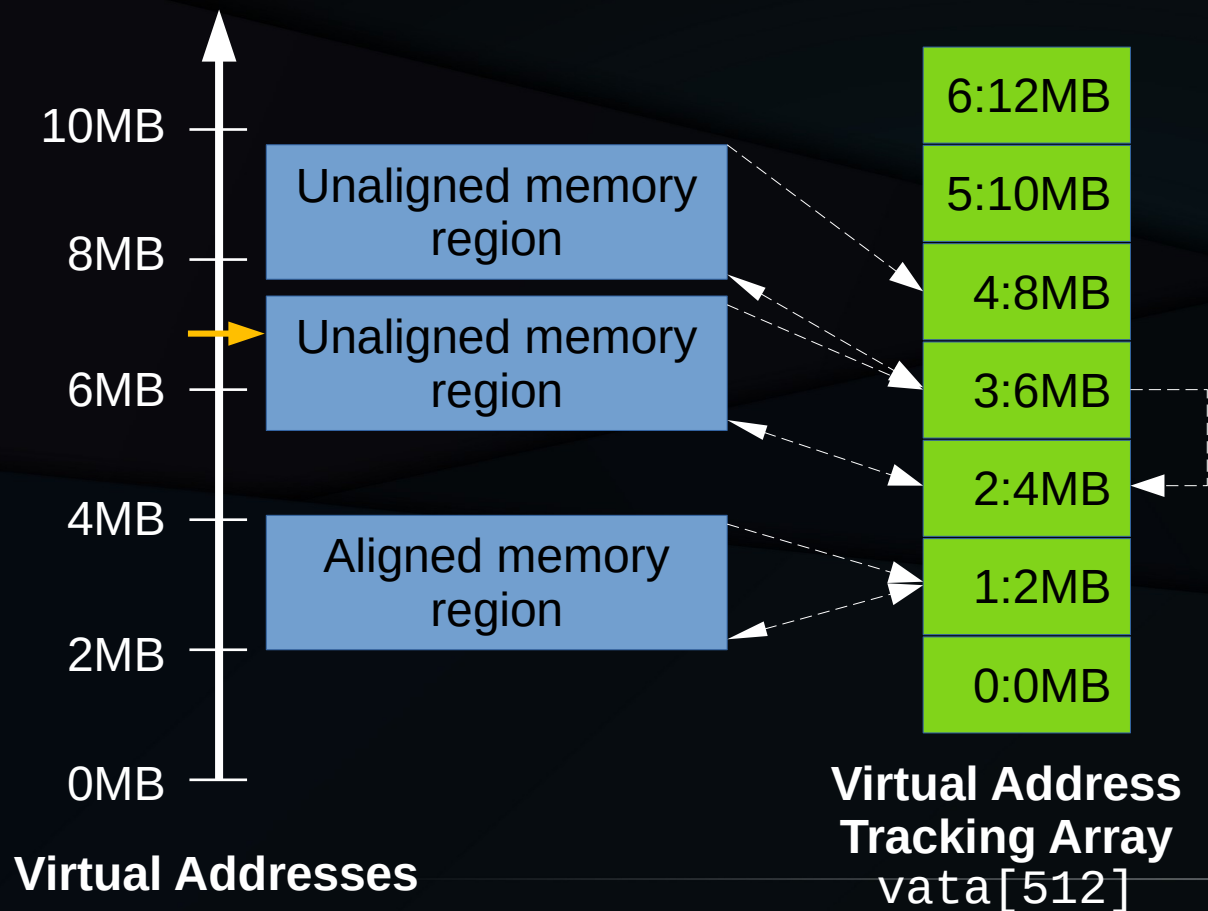
- Crash dump was a complete disaster
  - No hints for on-site debugging or documentation
- Unable to reproduce in lab
- Eventually, found a fix: case: 5-27-hour MTBF

Created rudimentary digital twin?

# Chaos-Monkey Challenges

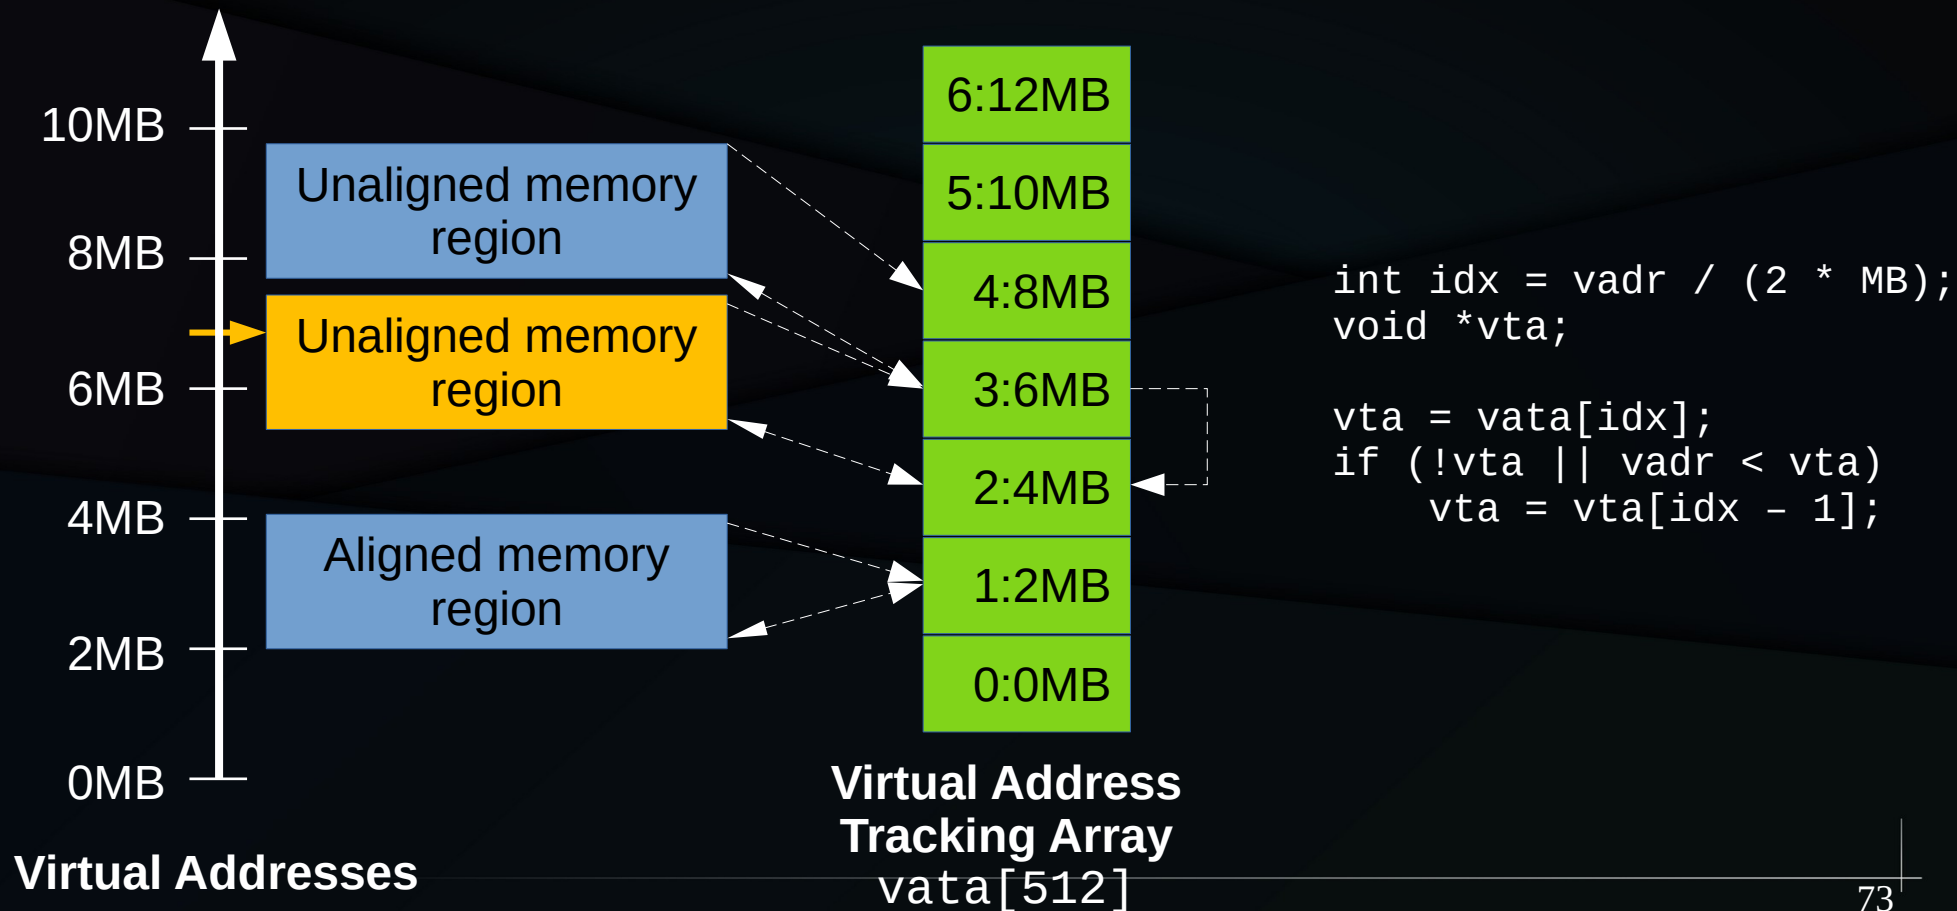
- Crash dump was a complete disaster area
  - No hints for on-site debugging instrumentation
- Unable to reproduce in the lab
- Eventually, found test case: 5-27-hour MTBF
  - But need week-long test for any alleged fix!!!
  - And it was now Memorial Day weekend...

# Hint From Stack Trace

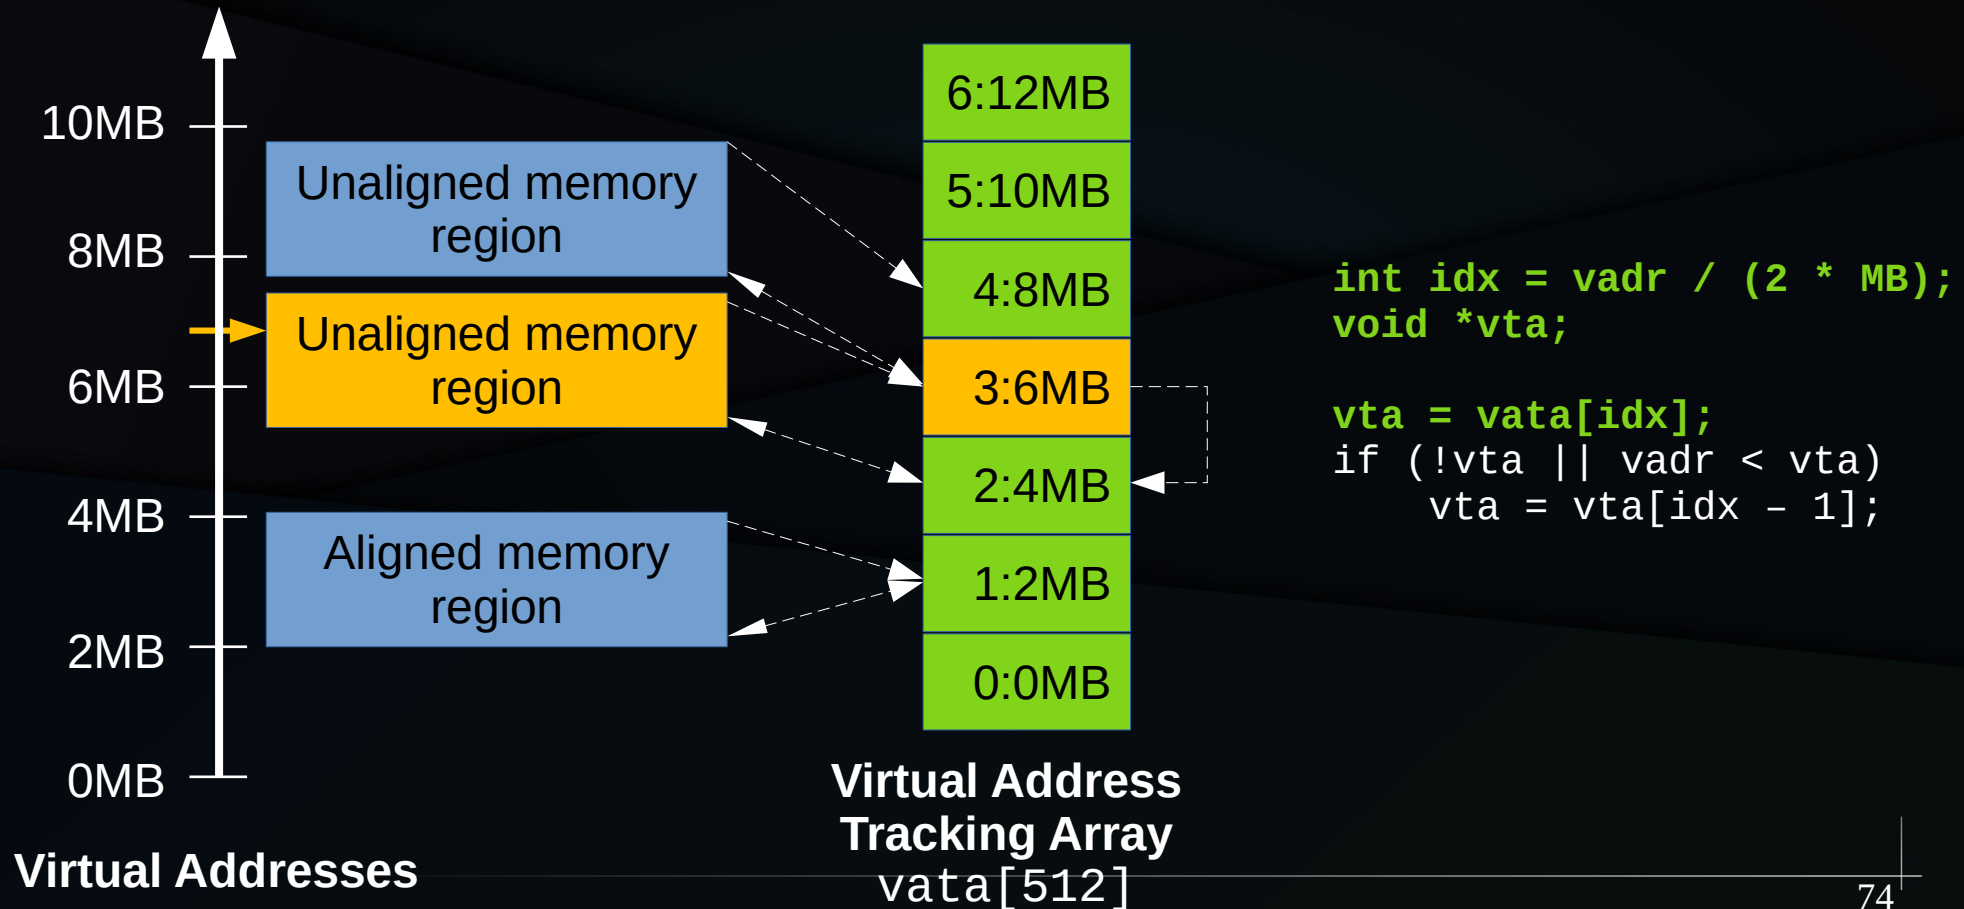




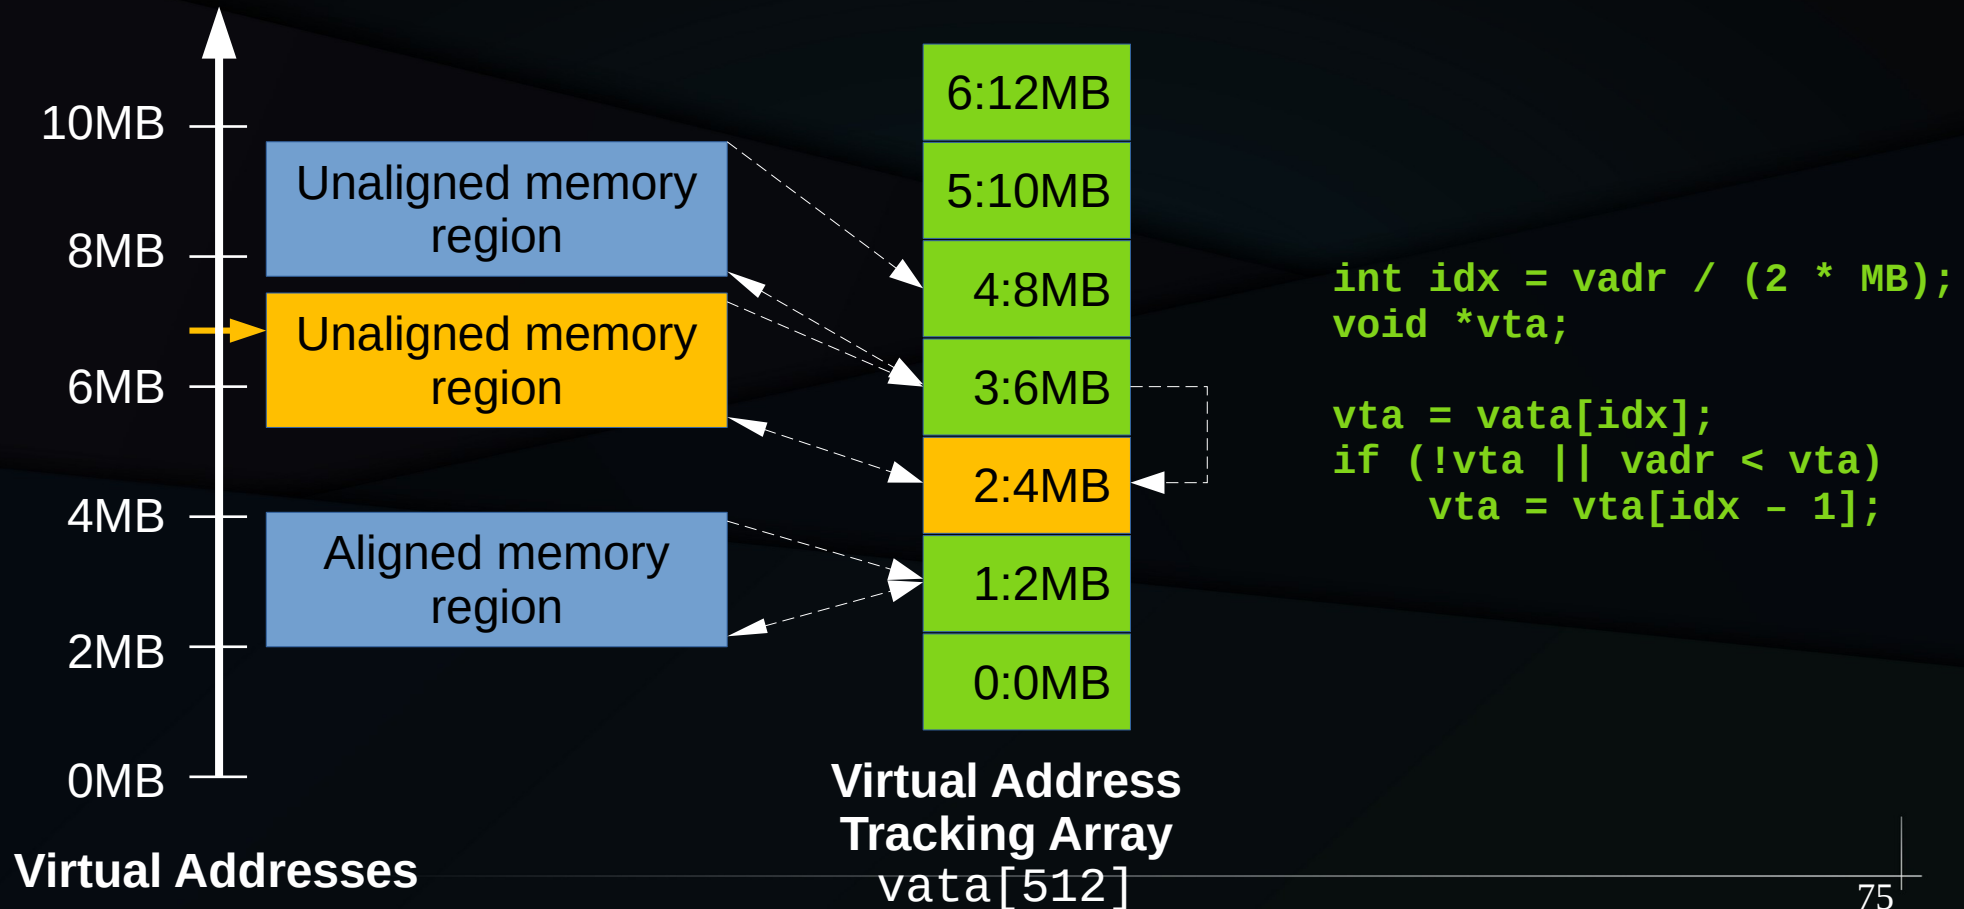
# Hint From Stack Trace



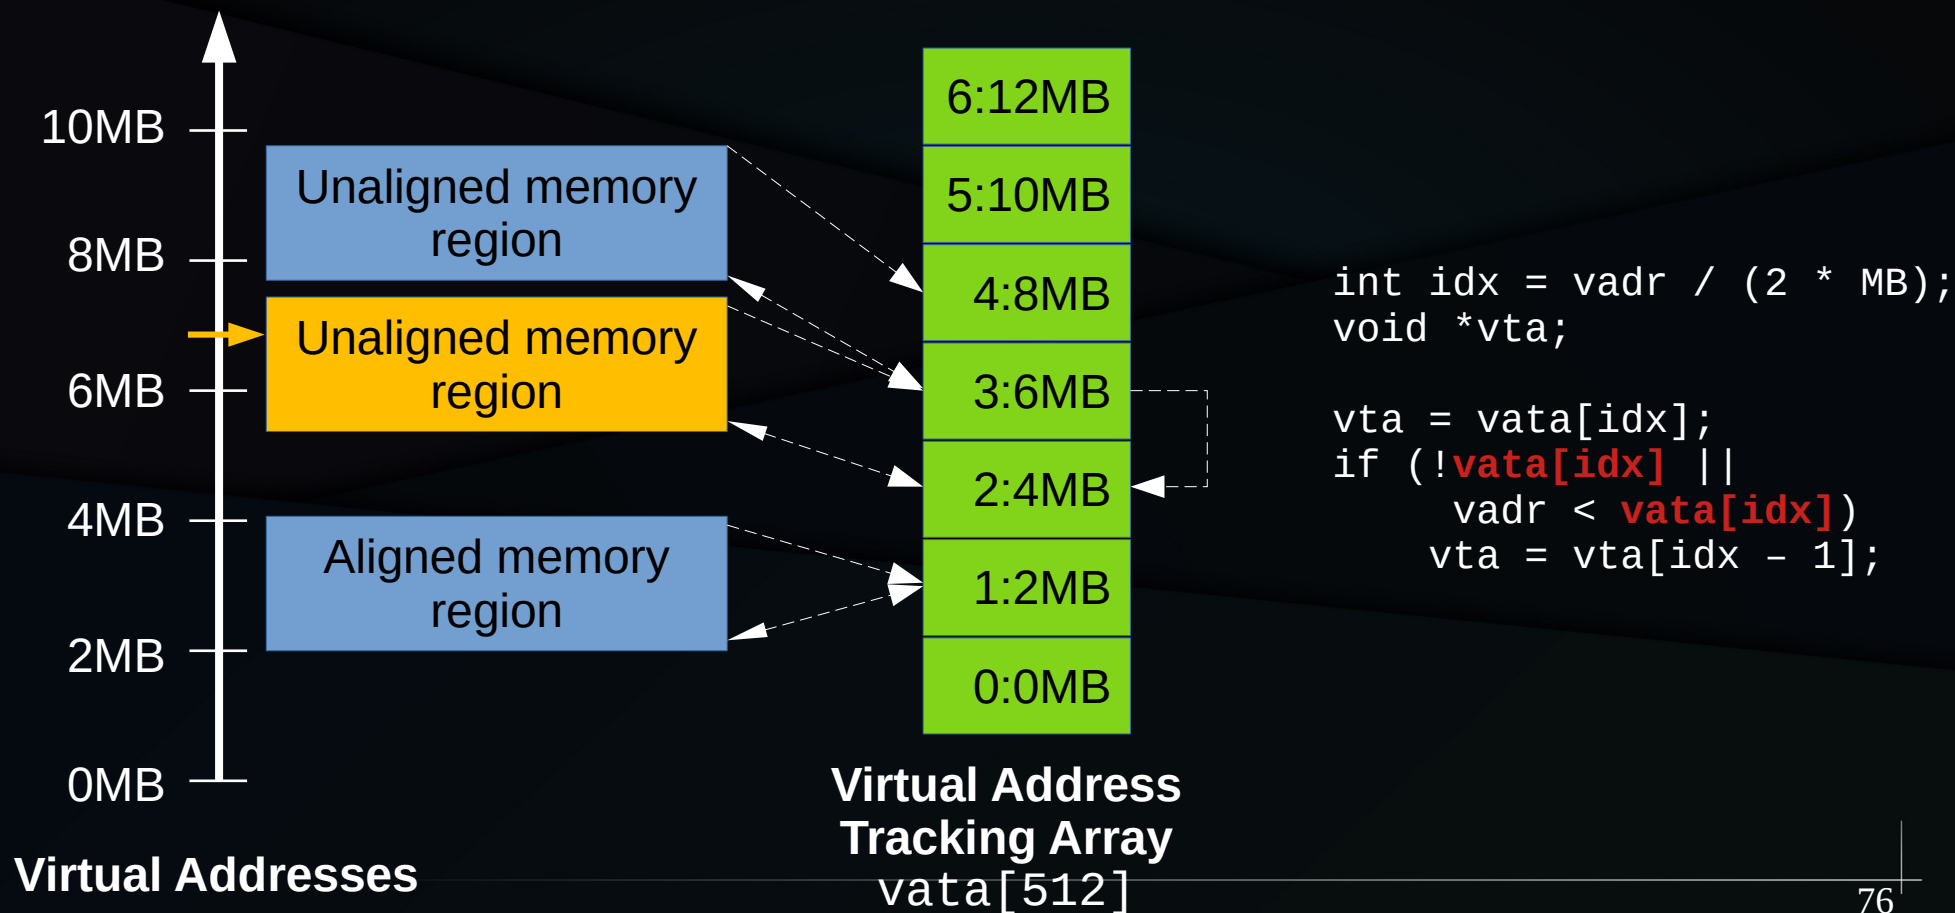
# Hint From Stack Trace



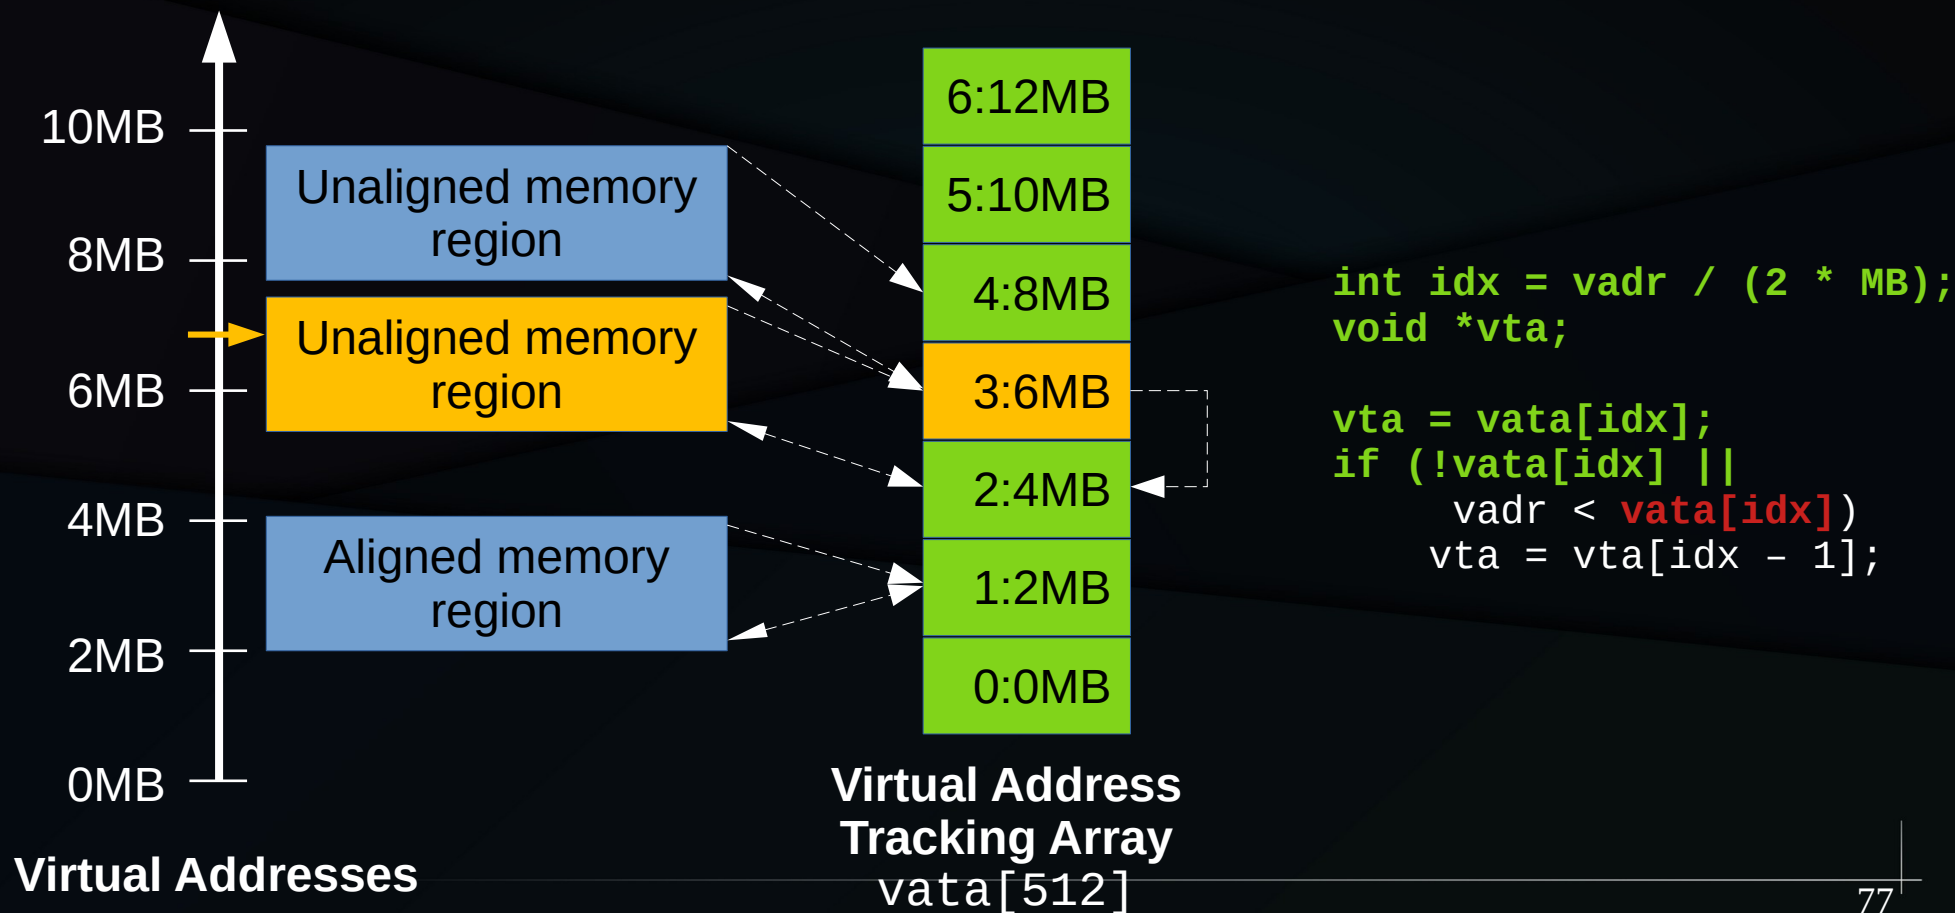
# Hint From Stack Trace



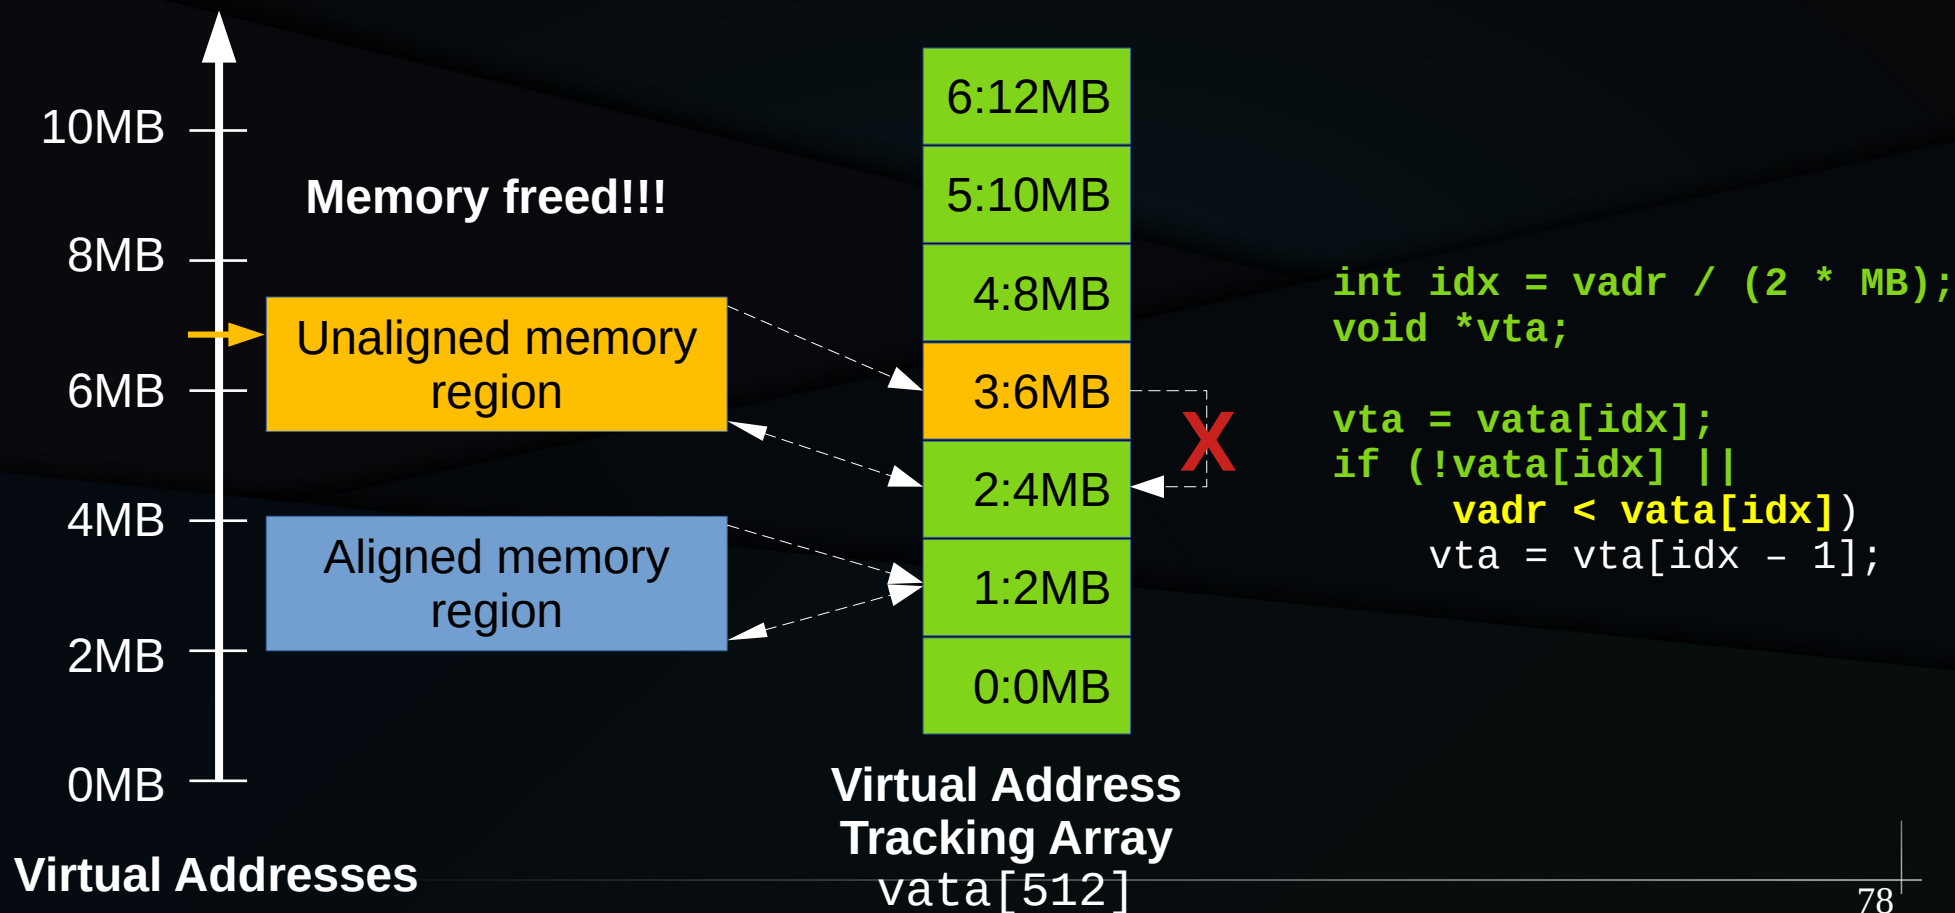
# Hint From Stack Trace: Compiler Fun



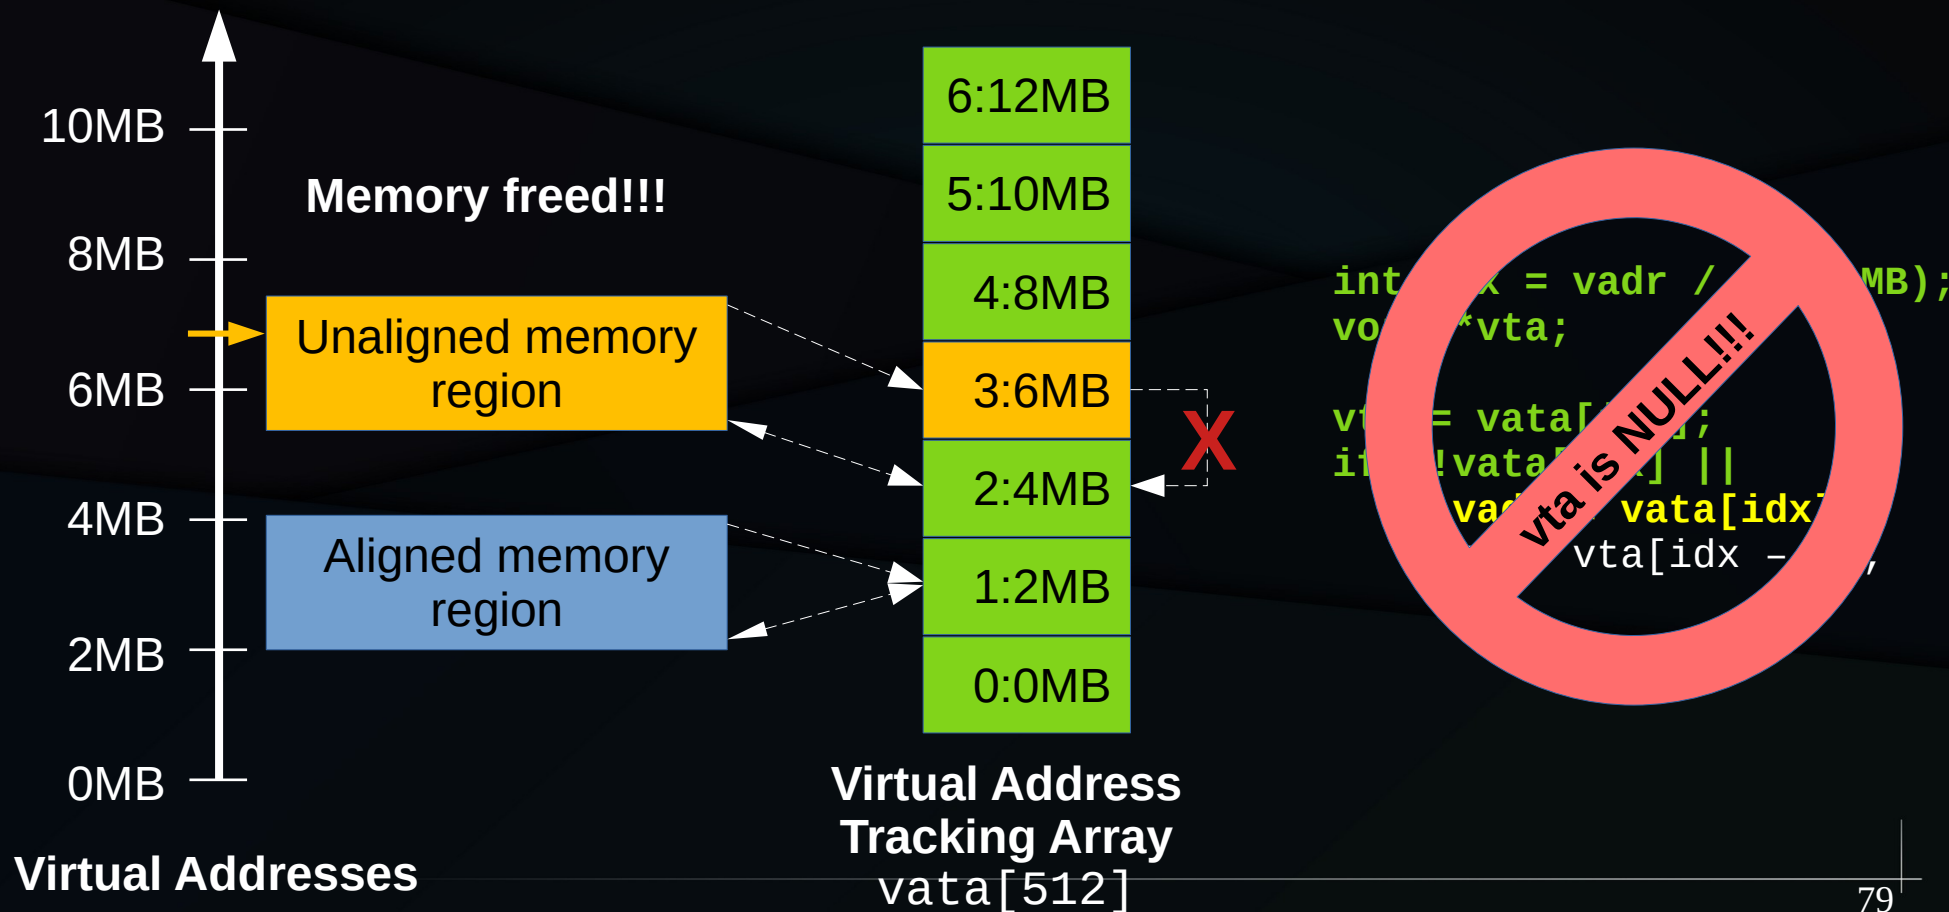
# Compiler Fun In Failure Case



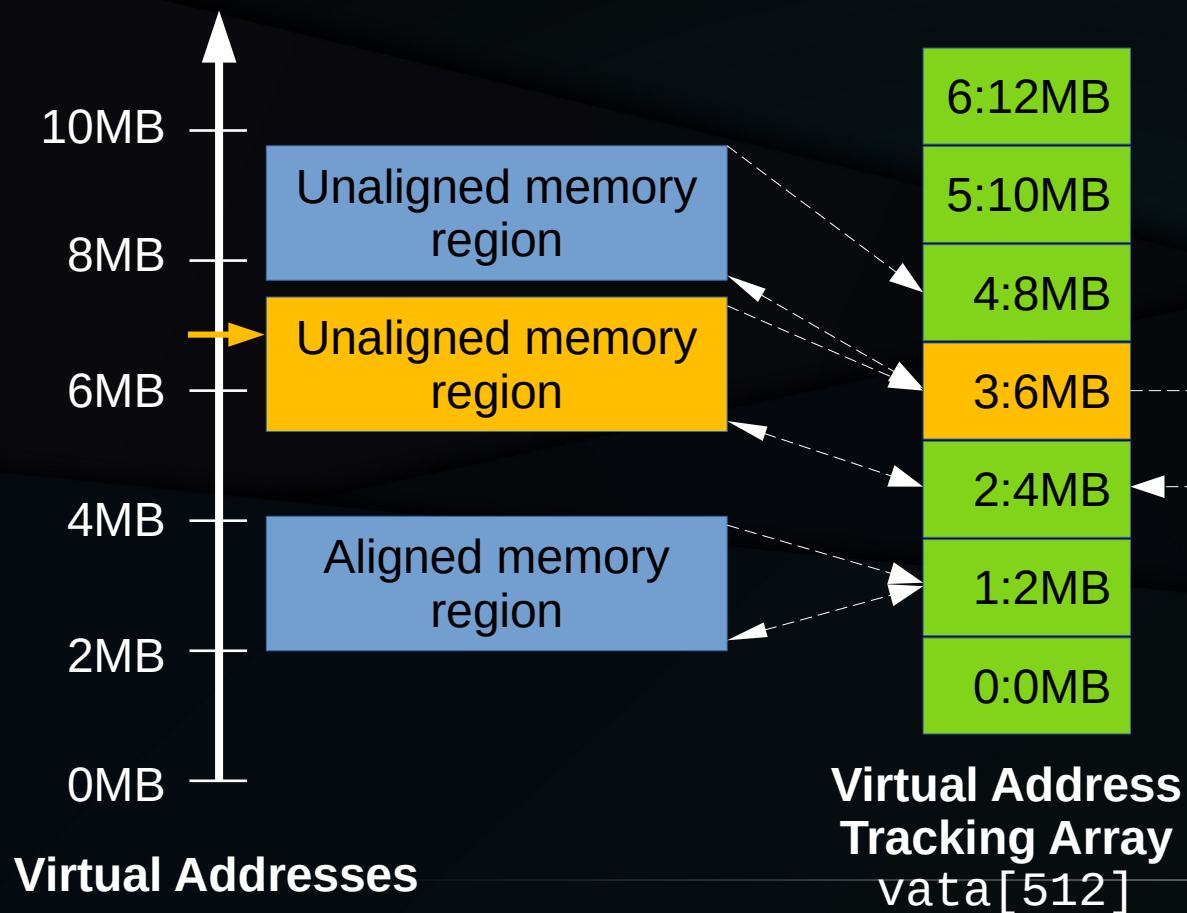
# Compiler Fun In Failure Case: Update



# Compiler Fun In Failure Case: Update



# Thwarting Compiler Fun

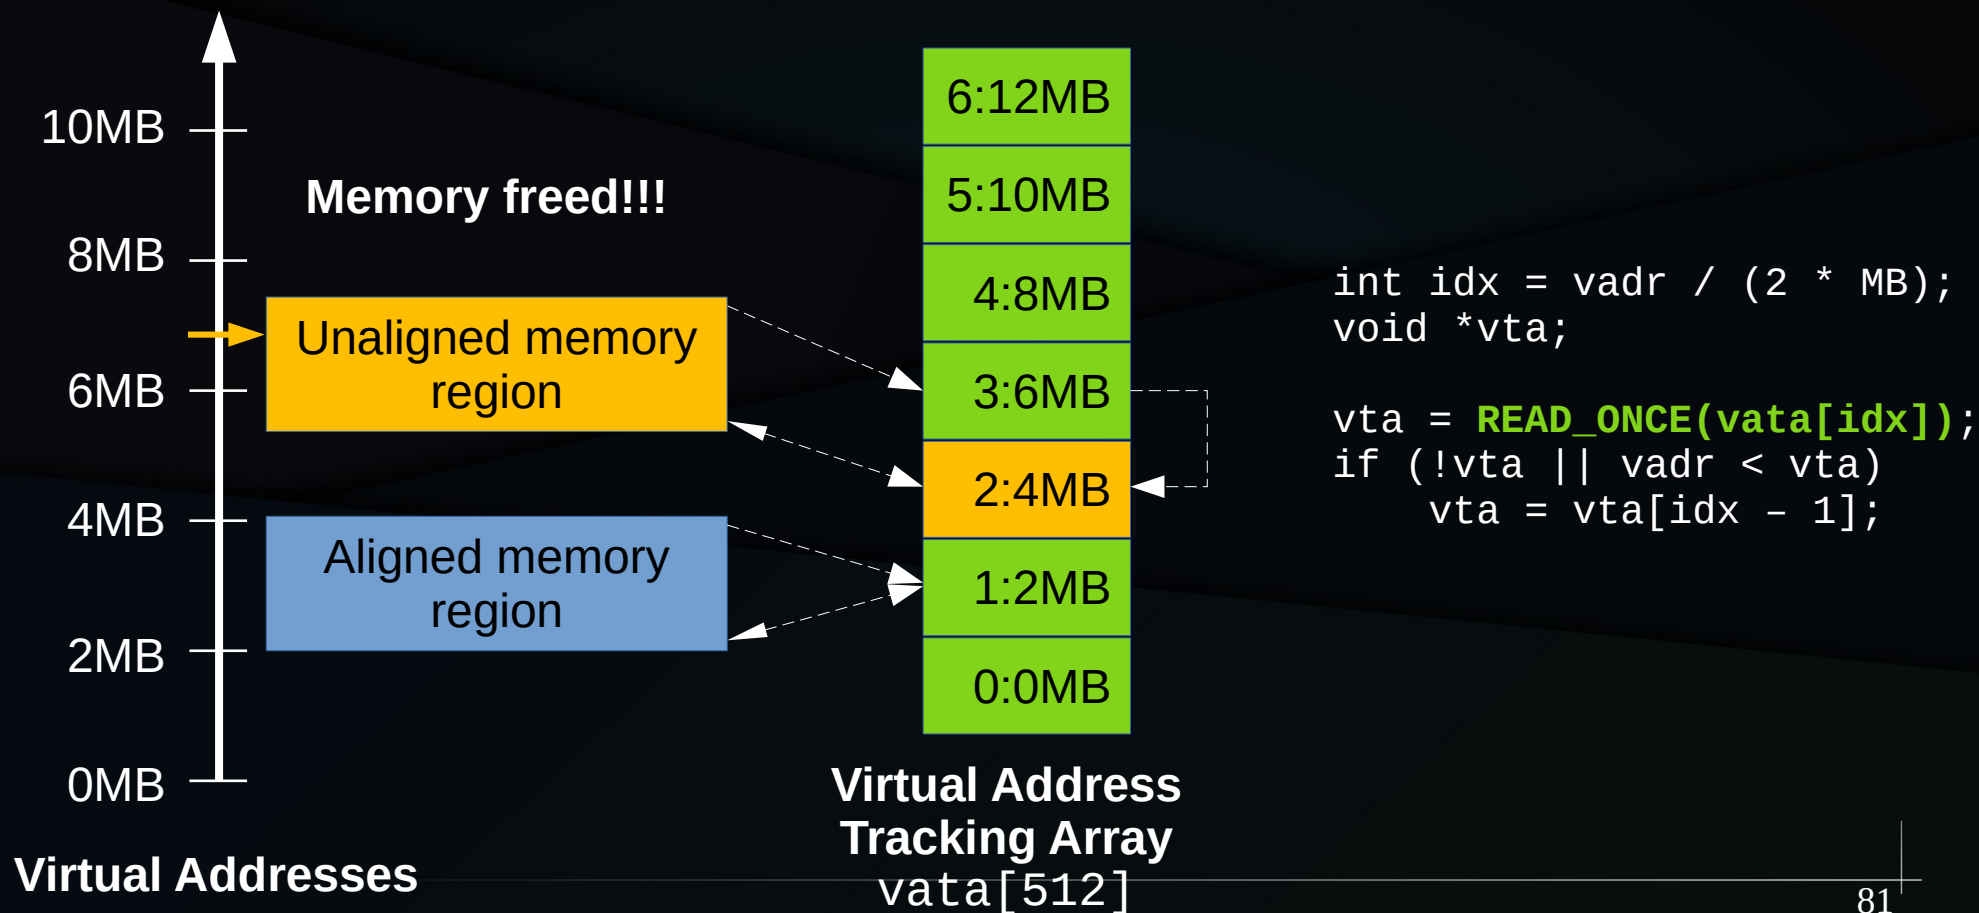


```
int idx = vadr / (2 * MB);  
void *vta;
```

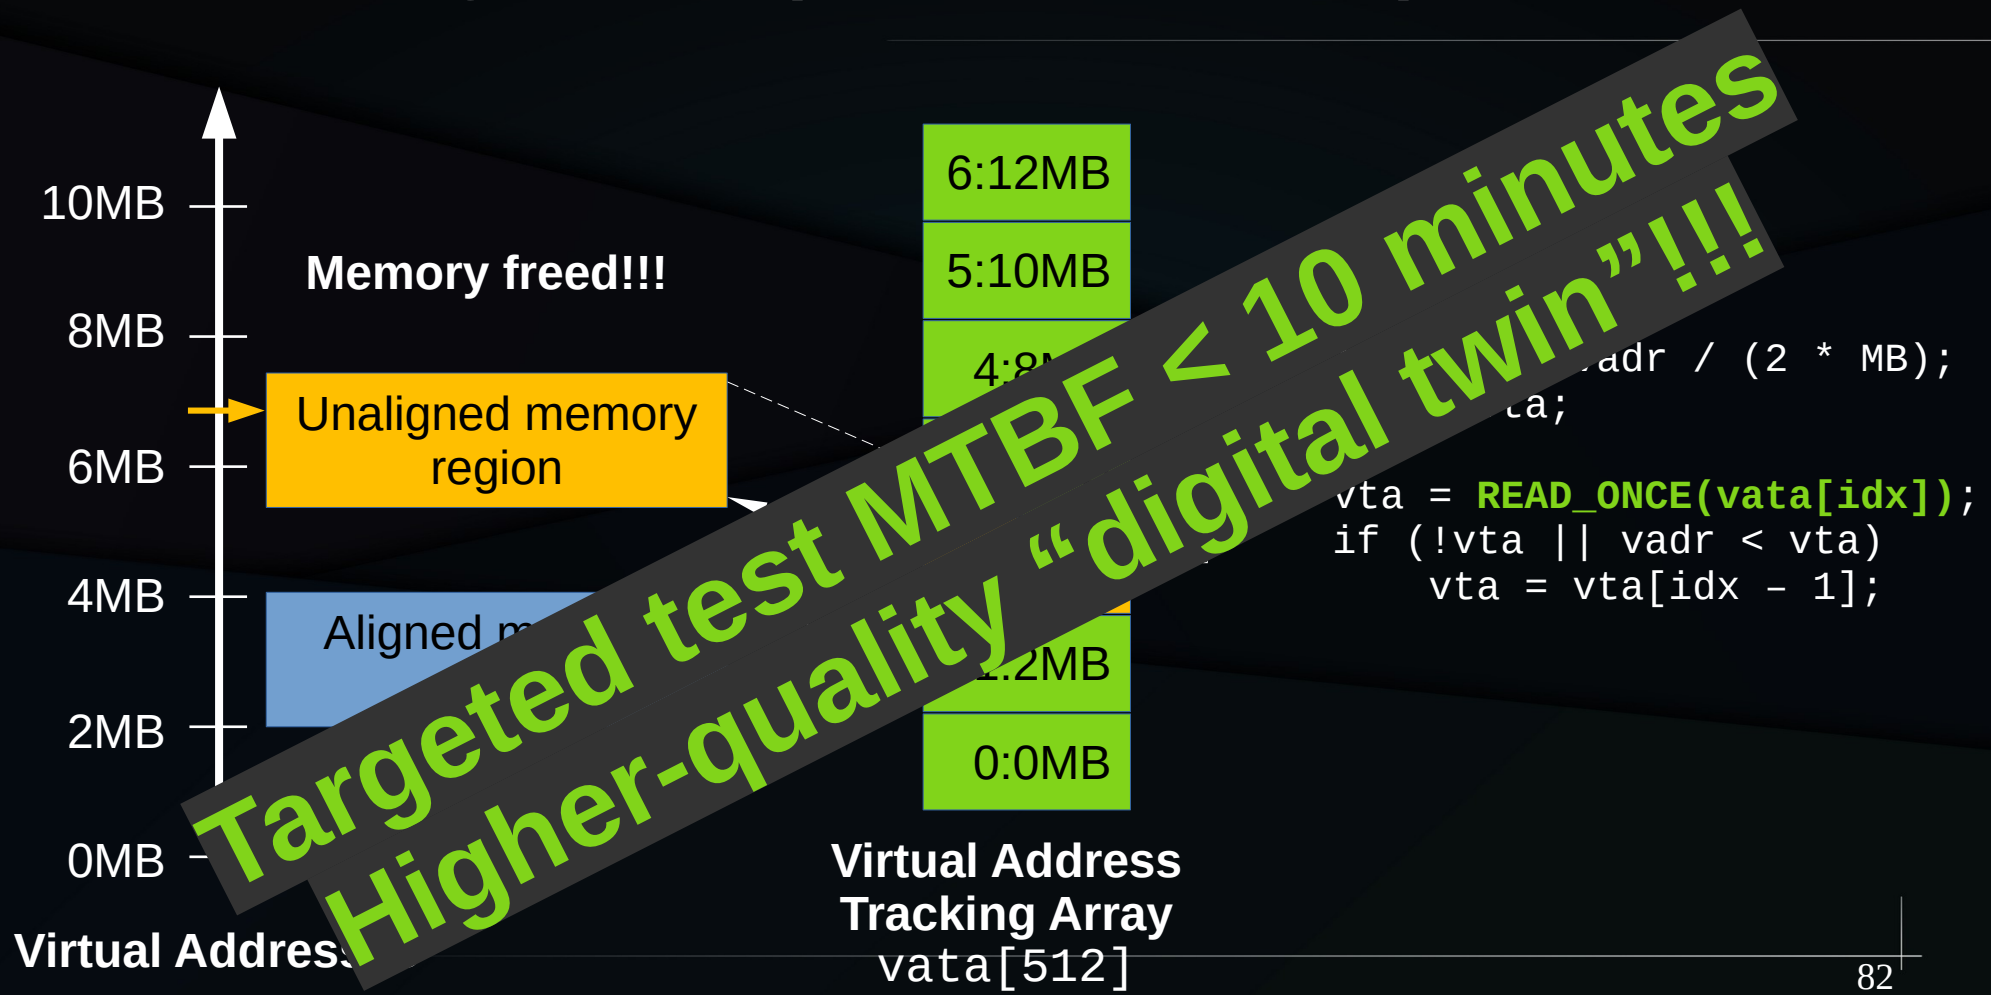
```
vta = READ_ONCE(vata[idx]);  
if (!vta || vadr < vta)  
    vta = vata[idx - 1];
```




# Thwarting Compiler Fun: Update OK




# Thwarting Compiler Fun: Update OK



# Shared Disks For Availability Win!!!



# Shared Disks For Availability Win!!!



# Shared Disks For Availability Win!!!



# 1970s: My First Professional Project



# 1970s: My First Professional Project

- Pro-bono computer dating program for National Honor Society fundraiser during my senior year in high school
- Questions from Home Economics teacher
- Simple Hamming-distance matching with expected 1970s constraints on matches
- Students' paper questionnaires transcribed to paper tape, then read into program
- Simple, effective, worked great!!!

# One Unsatisfied Customer

- Senior girl matched only with freshmen boys
  - And she really did check the seniors-only box
- Program looked to be correct
- Turned out to be data-entry error
- Correct program is not enough
  - Environment and processes matter!!!



# One Unsatisfied Customer

- Senior girl matched only with fresh
  - And she really did check the only box
- Program looked to be correct
- Turned out to be a data-entry error
- Correct program is not enough
- Environment and processes matter!!!

**Good idea implemented properly,**

# One Unsatisfied Customer

- Senior girl matched only with fresh
  - And she really did check the
- Program looked to be
- Turned out to be a error
- Corrected but not enough
- and processes matter!!!

**Good idea implemented properly,  
but I was also overall project lead!**

# Cautionary Quote

---

- A lot of success in life and business comes from knowing what you want to avoid. - *Charlie Munger*

# 2004: Real-Time Linux

# 2004: Real-Time Linux


- Early 2000s: Many requests for real-time Linux
  - But “enterprise-grade real-time Linux”
- Except that no such thing existed
- And my employer had strict rules for contracts calling for mythical creatures

# 2004: Real-Time Linux

- Early 2000s: Many requests for real-time Linux
  - But “enterprise-grade real-time Linux”
- Except that no such thing existed
- And my employer had strict rules for contracts calling for mythical creatures


**No Bid**

# 2004: Dawn of Multicore Embedded



# Multicore Embedded for Real Time!!!

Non-realtime code here

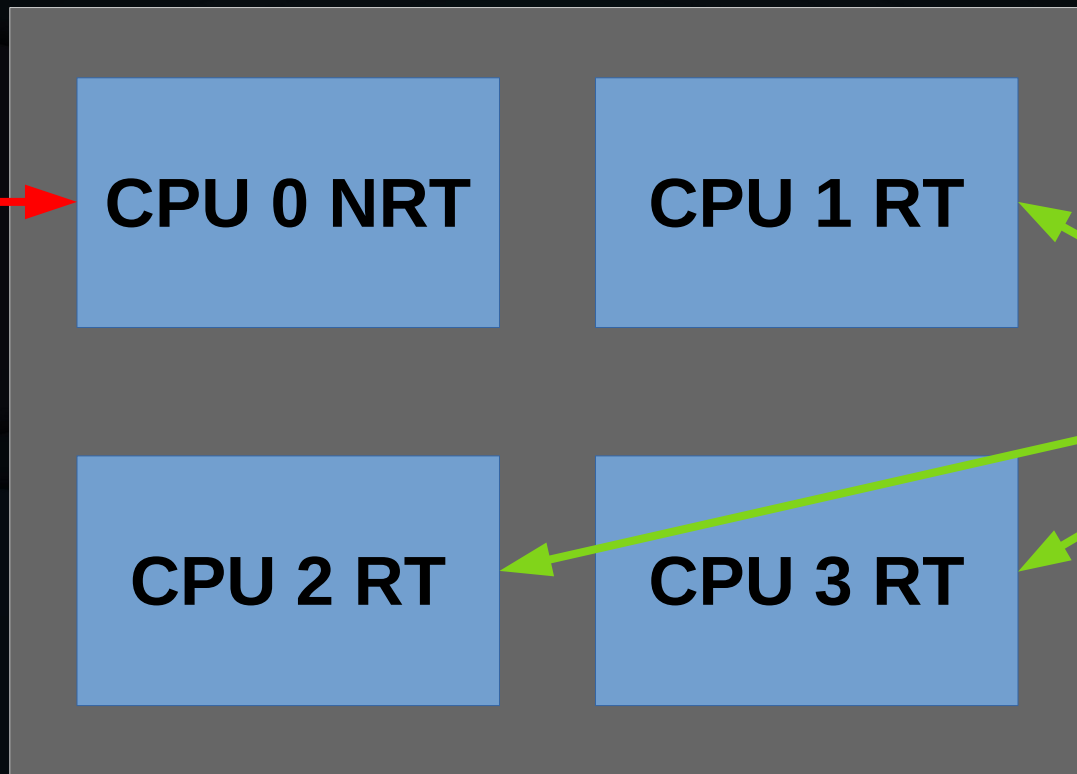


Realtime code here



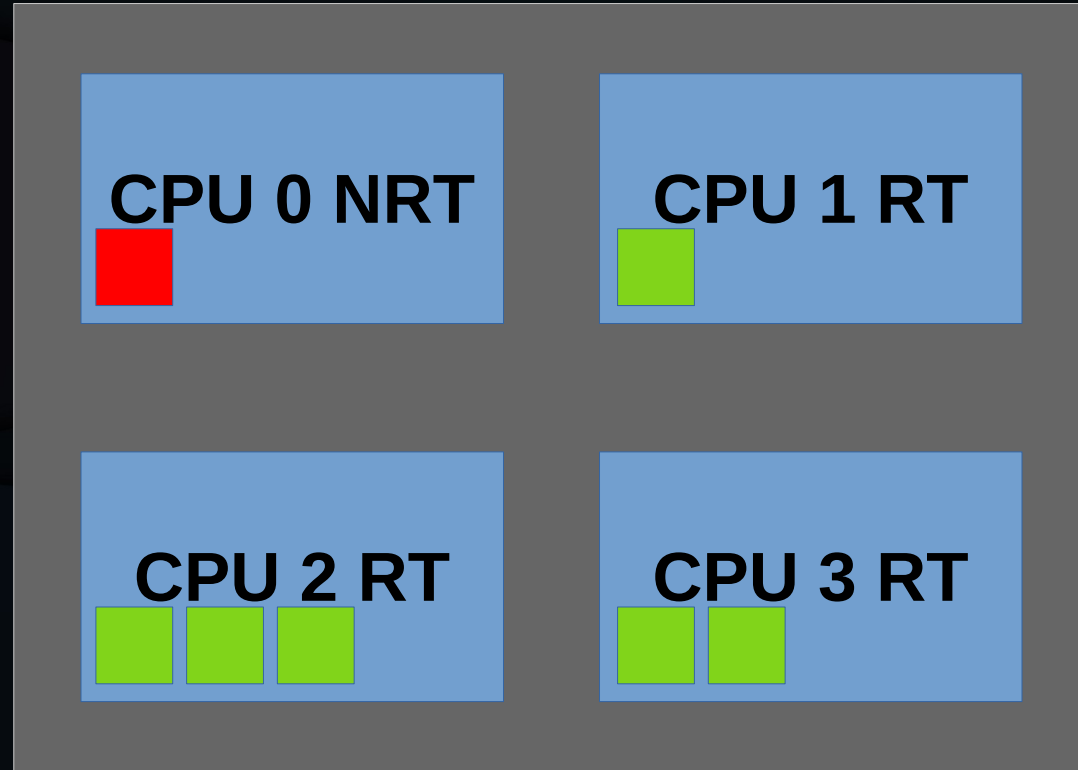
# Multicore Embedded for Real Time!!!

Non-realtime code here

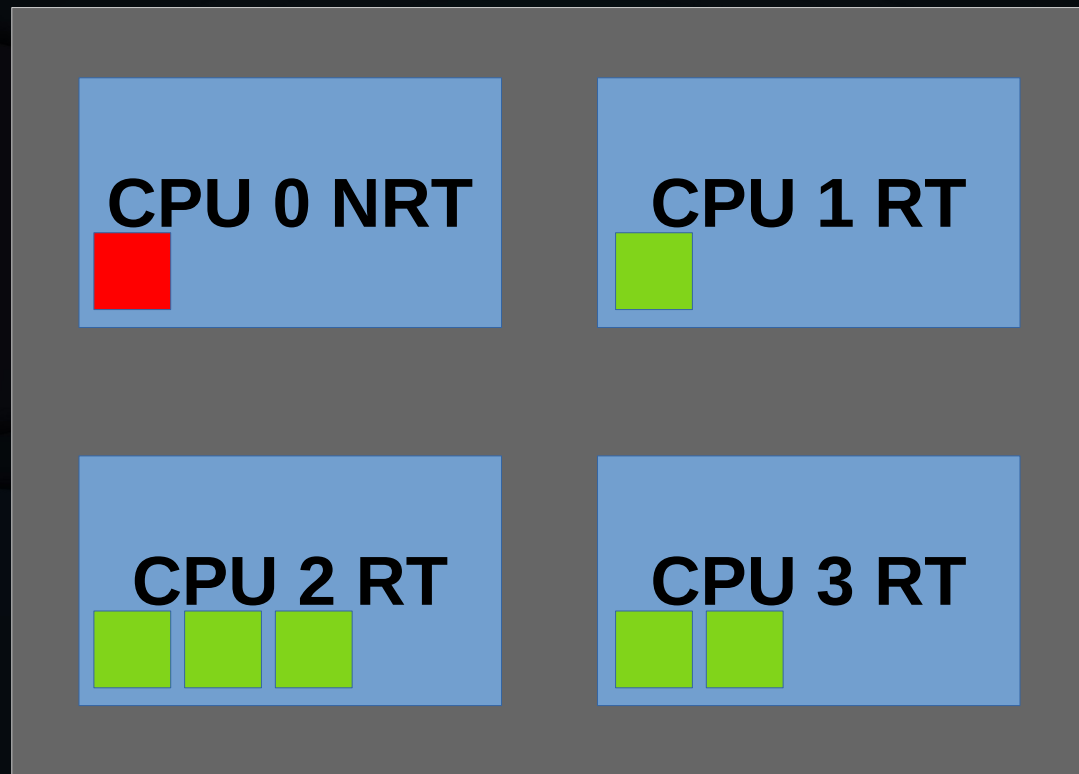


Realtime code here

# Multicore Embedded for Real Time!!!

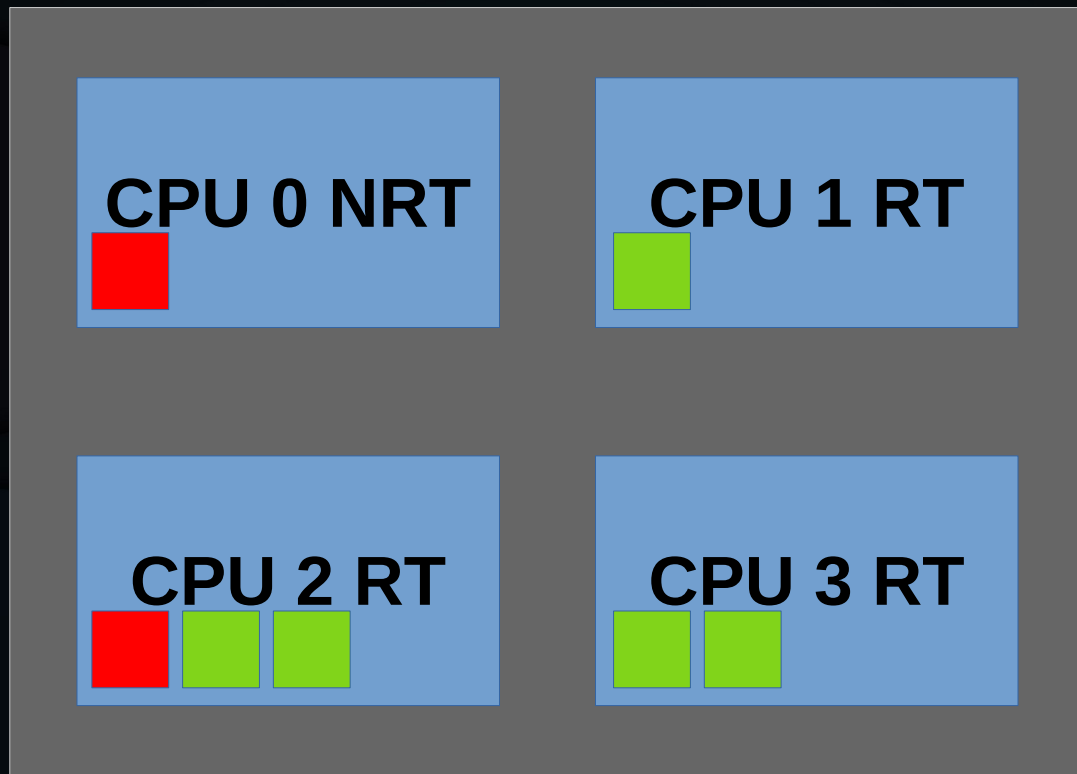


# Multicore Embedded for Real Time!!!



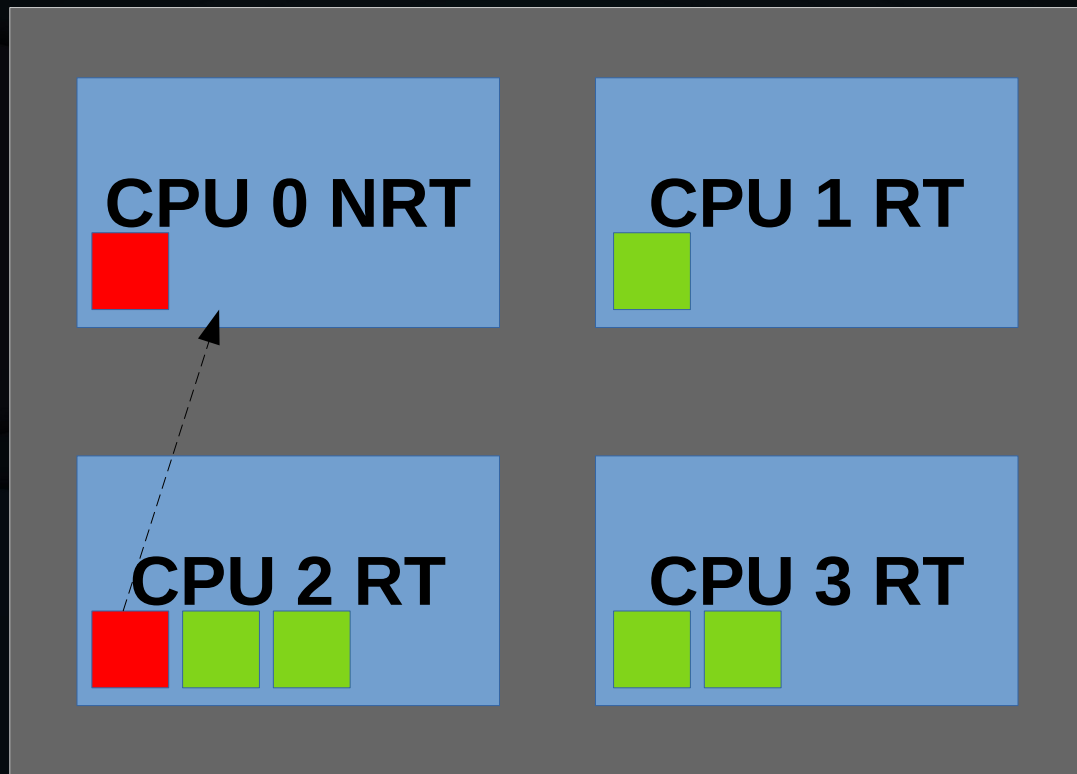
Respond to non-real-time activity by migrating.

# Multicore Embedded for Real Time!!!



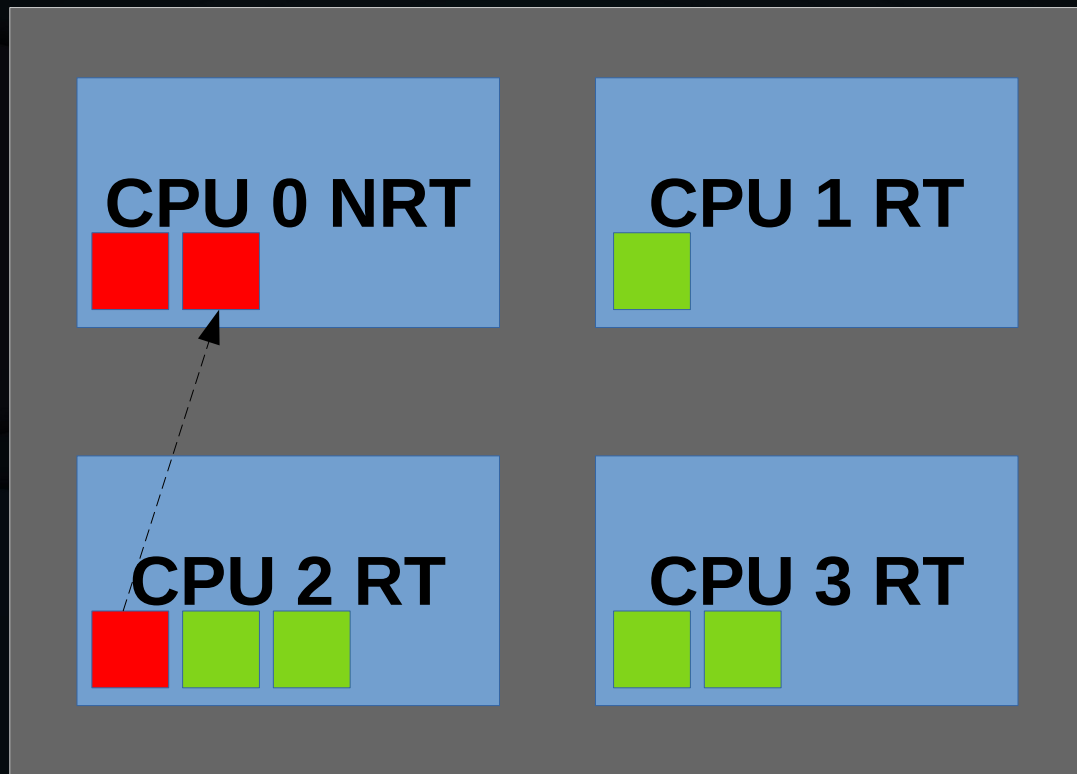
Respond to non-real-time activity by migrating.

# Multicore Embedded for Real Time!!!



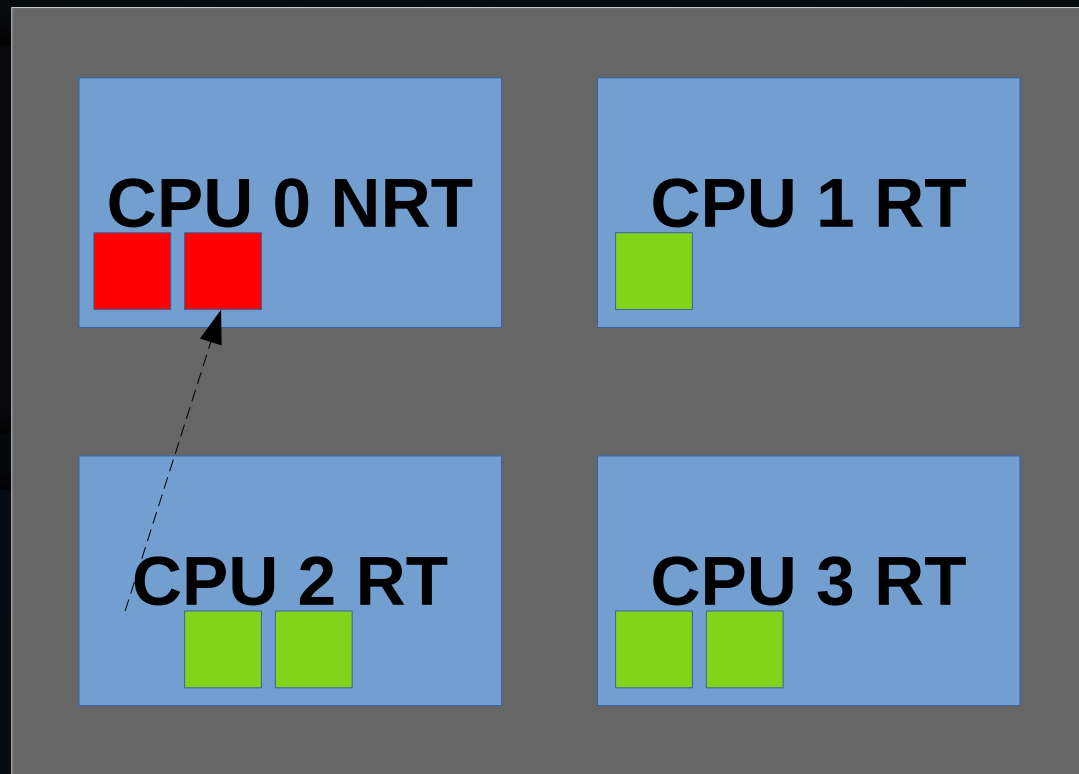
Respond to non-real-time activity by migrating.

# Multicore Embedded for Real Time!!!



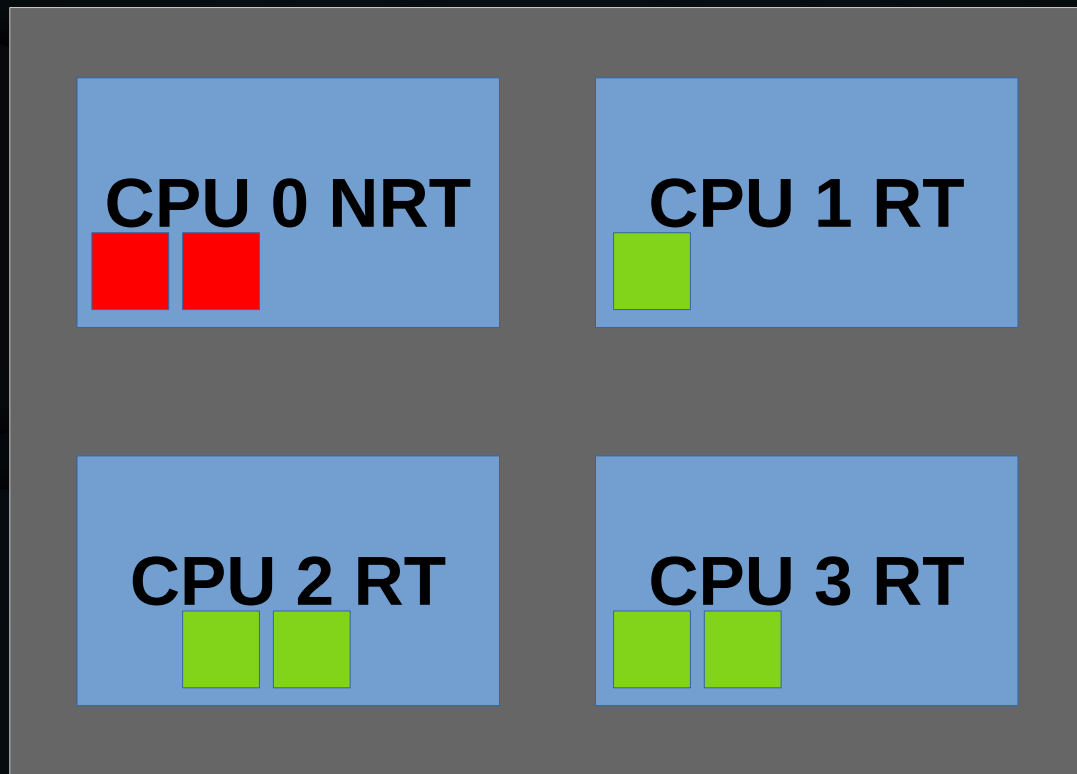
Respond to non-real-time activity by migrating.

# Multicore Embedded for Real Time!!!



Respond to non-real-time activity by migrating.

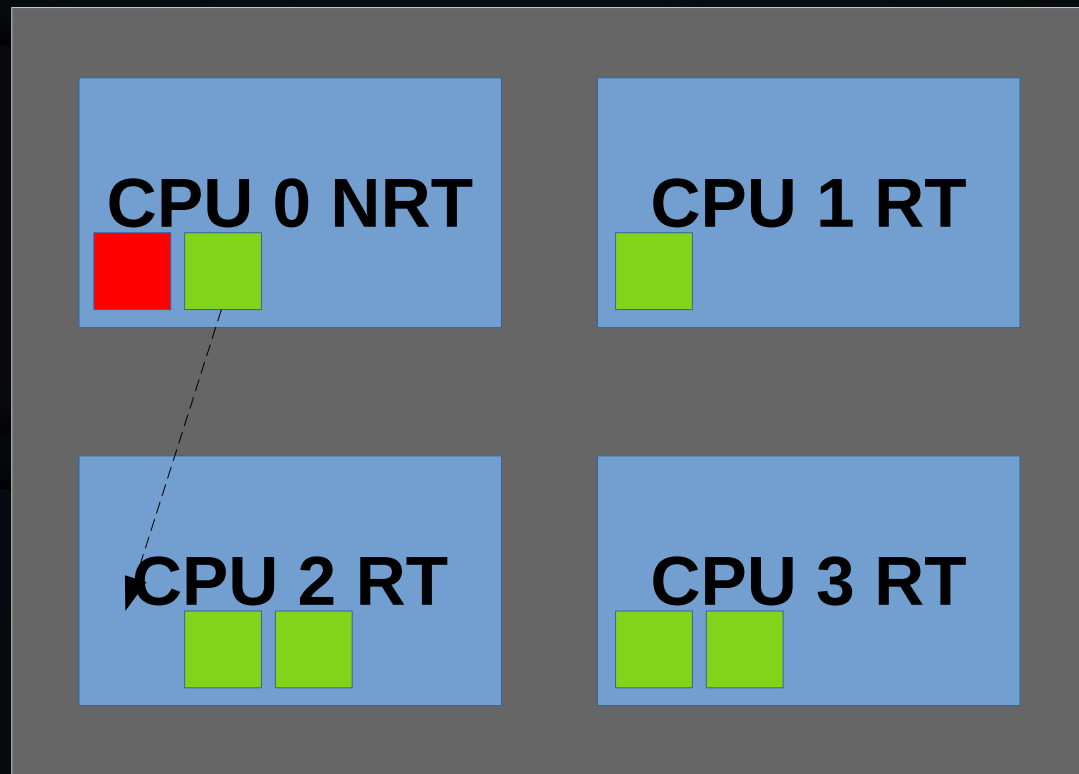
# Multicore Embedded for Real Time!!!



Respond to non-real-time activity by migrating.

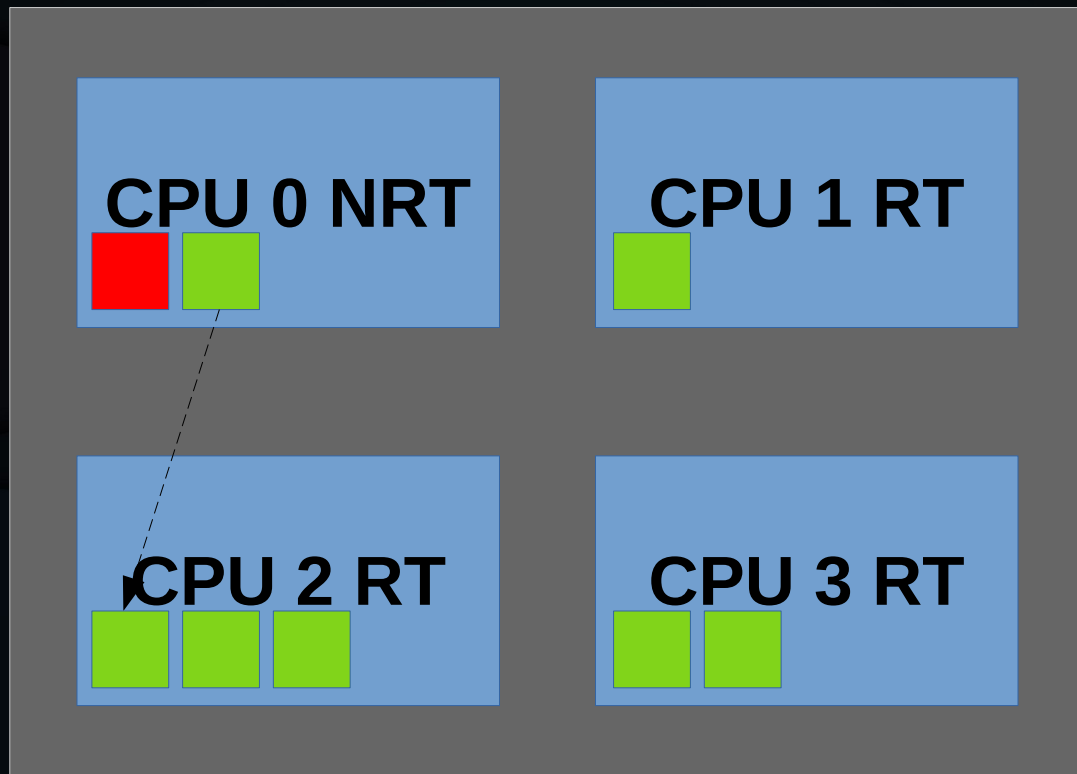


# Multicore Embedded for Real Time!!!



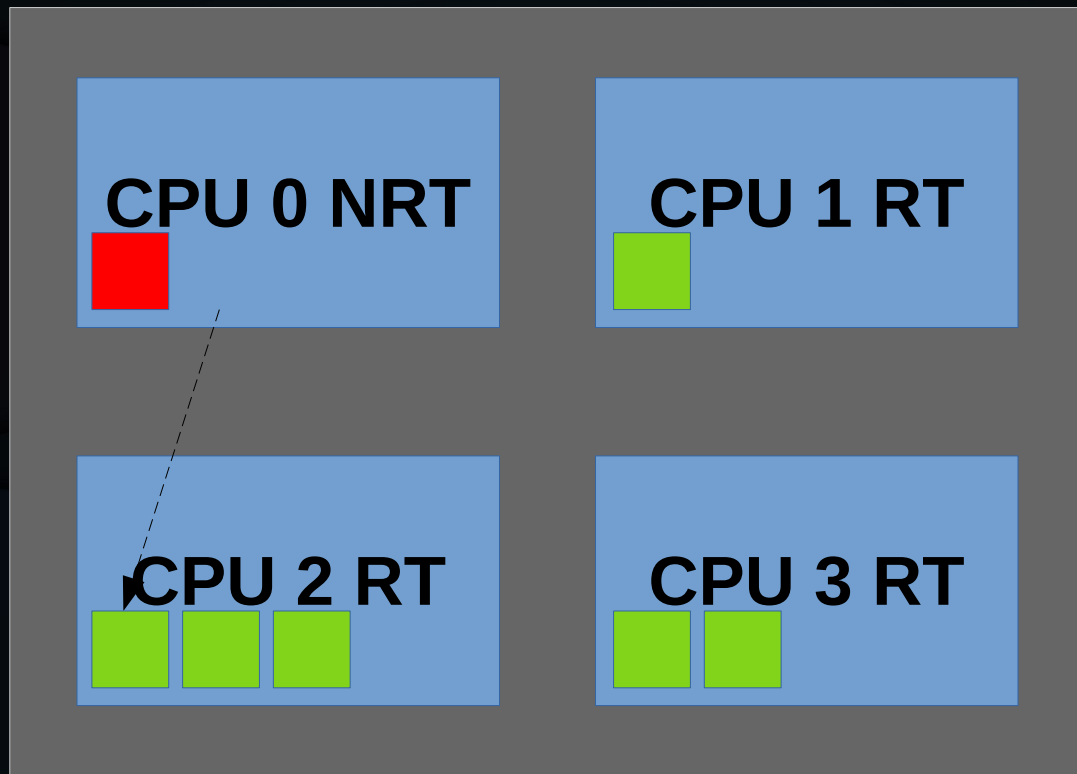
Respond to non-real-time activity by migrating and back when done (system call).

# Multicore Embedded for Real Time!!!



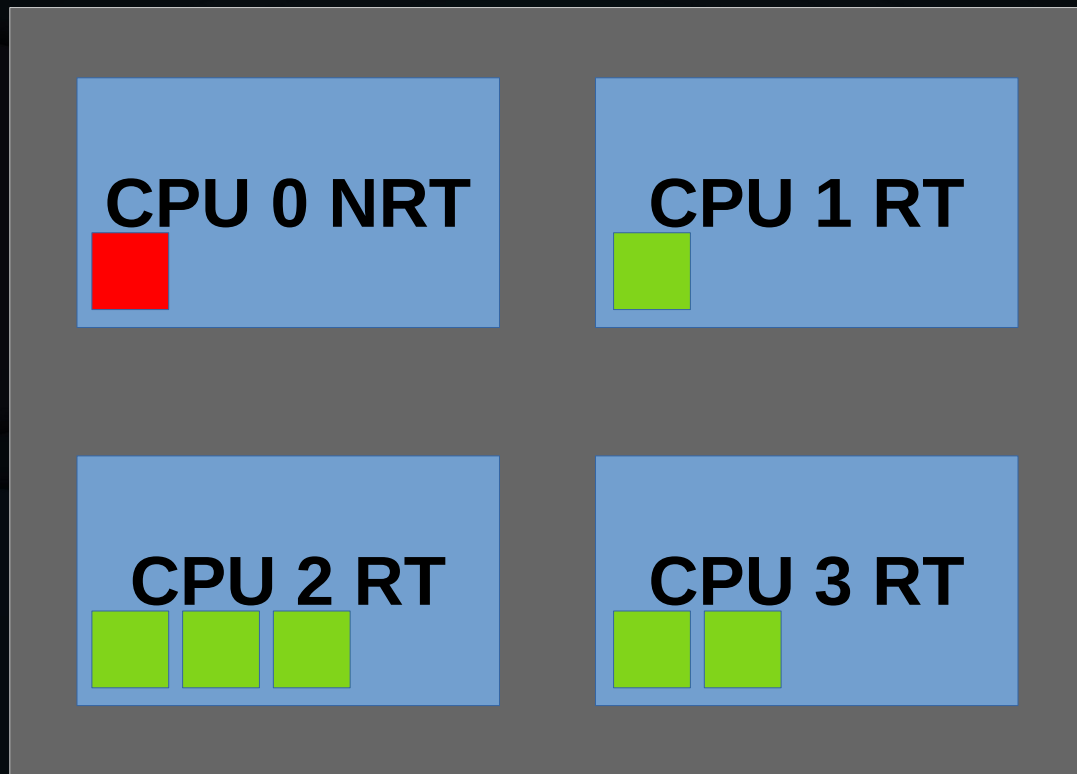
Respond to non-real-time activity by migrating and back when done (system call).

# Multicore Embedded for Real Time!!!



Respond to non-real-time activity by migrating and back when done (system call).

# Multicore Embedded for Real Time!!!



Respond to non-real-time activity by migrating and back when done (system call).

# Multicore Real Time Linux Actions

- Produce patch implementing syscall migration

# Multicore Real Time Linux Actions

- Produce patch implementing syscall migration
- Test it out, works great!

# Multicore Real Time Linux Actions

- Produce patch implementing syscalls
- Test it out, works great!

**There is a real-time effort spinning up.  
But they are rewriting the kernel.  
Pragmatism for the win!!!**

# Multicore Real Time Linux Actions

- Produce patch implementing syscall migration
- Test it out, works great!
- Inform executives real-time Linux is real!!!
- No more need for no-bid!!!



# Multicore Real Time Linux Actions

- Produce patch implementing syscall migration
- Test it out, works great!
- Inform executives real-time Linux is real!!!
- No more need for no-bid!!!
- And we win a large contract!!!

# Multicore Real Time Linux Actions

- Produce patch implementing sysctl migration
- Test it out, works great!
- Inform executives real time Linux is real!!!
- No more need for a no-bid!!!
- And we will win a large contract!!!

**My idea is rejected!!!**

# Multicore Real Time Linux Actions

- Why was my brilliant idea rejected?

# Multicore Real Time Linux Actions

- Why was my brilliant idea rejected?



# Multicore Real Time Linux Actions

- Produce patch implementation and call migration
- Test it out, works
- Inform executives Linux is real!!!
- No more no
- And we will sign a contract!!!

**Rejected!!!  
Except that we have  
contractual commitments  
to meet...**

# Multicore Real Time Linux Actions

- Patch implementation and migration
  - Test it out,
  - Inform executives Linux is real!!!
  - No more no
  - And we will
- Remember that we have**  
**Except we have**  
**rewrite-the-kernel effort?**  
**contractual commitments**  
**to meet contract!!!**

# Multicore Real Time Linux Actions

- Patch implementation and migration
- Test
- Inform executives Linux is real!!!
- No more no
- And we will sign contract!!!

**Remember that we have**  
**Well, I helped them with RCU**  
**Exceeded commitments**  
**contract to 10**  
**rewrite-the-kernel effort?**

# Multicore Real Time Linux Actions

**Remember that we have**

**Well, I helped rewrite the kernel effort?**

**Three from-scratch implementations**

**contract**

• patch implementation and all migration

• Test Linux is real!!!

• No more

• And we will



# 2004: Real-Time Linux



# 2004: Real-Time Linux





# 2004: Real-Time Linux



# Formal Verification

# Formal Verification: Why Bother?

# Installed Base

Million-Year Bug? Once In a Million Years!!!



# Installed Base

**Million-Year Bug? Once In a Million Years!!!**

**Murphy is a nice guy: Everything that can happen, will...**

1

1975  
NHS

# Installed Base

**Million-Year Bug? Once In a Million Years!!!**  
**Murphy is a nice guy: Everything that can happen, will...**  
**...maybe in geologic time**


1

1975  
NHS



# Installed Base

Million-Year Bug? Once in Ten Millennia




# Installed Base

Million-Year Bug? Once per Century



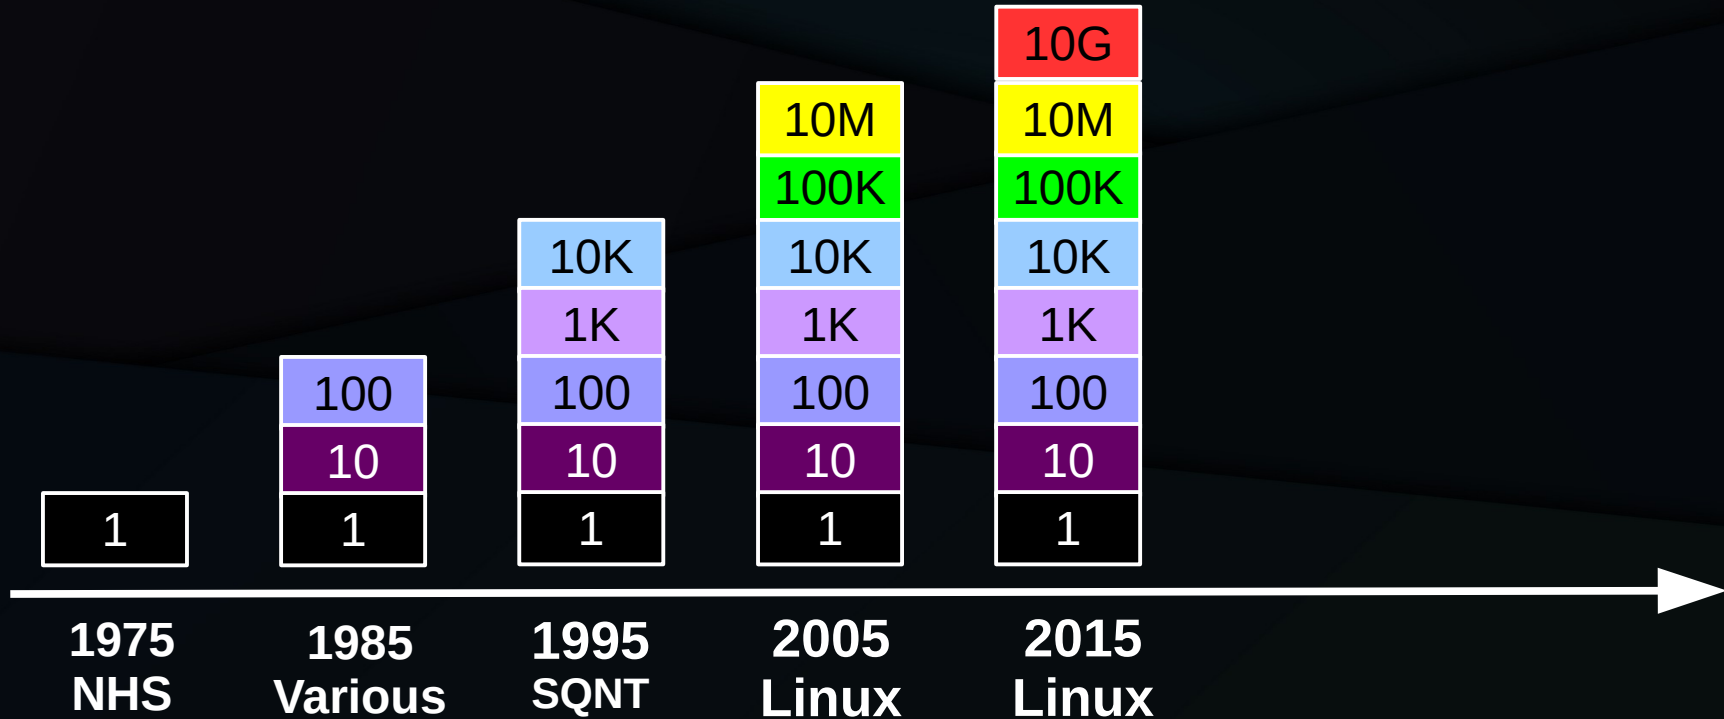
# Installed Base

Million-Year Bug? Once a Month




# Installed Base

Million-Year Bug? Several Times per *Day*




# Installed Base

Million-Year Bug? Several Times per *Hour*




# Installed Base

Million-Year Bug? You don't want to know...



# Installed Base

Million-Year Bug? You don't want to know...  
But has Murphy transitioned  
from a nice guy into a  
homicidal maniac?



# Formal Verification: Why Bother?

- 2017: 20 billion instances of the Linux kernel
  - Million-year MTBF bug fails >50 times per day
  - New kernel version every 2-3 months
- Testing really is feasible for low-duty-cycle devices
  - But not for the ~80 million servers!!!
- Plus Linux is used in safety-critical applications!!!
- Full state-space search is quite attractive



# Formal Verification Experience

- 1993: Promela/spin election algorithm
- 2007: “Quick” RCU (QRCU) verification
- 2008: RCU idle-detection for energy efficiency
- 2012: Verify userspace RCU
- 2014: Verify RCU idle detection for NMIs
- 2018-on: Heavy use of herd7 and LKMM

# Formal Verification Experience

- 1993: Promela/spin election algorithm
- 2007: “Quick” RCU (QRCU) verification
- 2008: RCU idle-detection energy efficiency
- 2012: Verify userspace RCU
- 2014: Verify idle detection for NMIs
- 2018: Heavy use of herd7 and LKMM

**Verifying design, not regression testing**

# Formal Verification Experience

- 1993: Promela/spin election algorithm
- 2007: “Quick” RCU (QRCU) vs. efficiency
- 2008: RCU idle-detection vs. efficiency
- 2012: Verify userspace
- 2014: Verify kernel section for NMIs
- 2018: Verify code of herd7 and LKMM

**Verifying design, not regression testing**  
**Verification valid after bug fix???**

# Formal Verification is *Expensive*

- At best, exponential; in general, undecidable
- “Macho” verification requires full specification
  - Which is large, thus containing lots of bugs!
- Successful formal verification highly restricted:
  - Small programs, simple properties of large programs, or execution-guided verification

# Formal Verification is *Expensive*

- At best, exponential; in general undecidable
- “Macho” verification requires full specification
  - Which is large, thus containing lots of bugs!
- Successful formal verification highly restricted:
  - Small programs, simple properties of large programs, or execution-guided verification

**Powerful when used properly**

# Formal Verification is *Expensive*

- At best, exponential; in general, intractable
- “Macho” verification requires complete specification
  - Which is large, thus a source of bugs!
- Successful formal verification highly restricted:
  - Small number of simple properties of large programs
  - Execution-guided verification

**Powerful when used properly**  
**How to verify the verification?**

# Formal Verification is Heavily Used

- Several test projects on the Linux kernel
- Many proprietary projects verify each commit
- But...
  - Formal verification in the small
  - Check for undesirable properties
    - File bug reports as appropriate

# Formal Verification is Heavily Used

- Several test projects on the Linux kernel
- Many proprietary projects verify each commit
- But...
  - Formal verification is small
  - Check for desirable properties
    - File bug reports as appropriate

De-risk via one-way bet



# Cautionary Quote (Redux)

- A lot of success in life and business comes from knowing what you want to avoid. - *Charlie Munger*

# Cautionary Quotes



- Sometimes you don't even know what you want until you find out you can't have it. - *Meghan O'Rourke*
- Sometimes we don't know what we want until we don't get it. - *Sloane Crosley*
- We don't know what we want, but we are ready to bite somebody to get it - *Will Rogers*

# Natural Selection

# Natural Selection

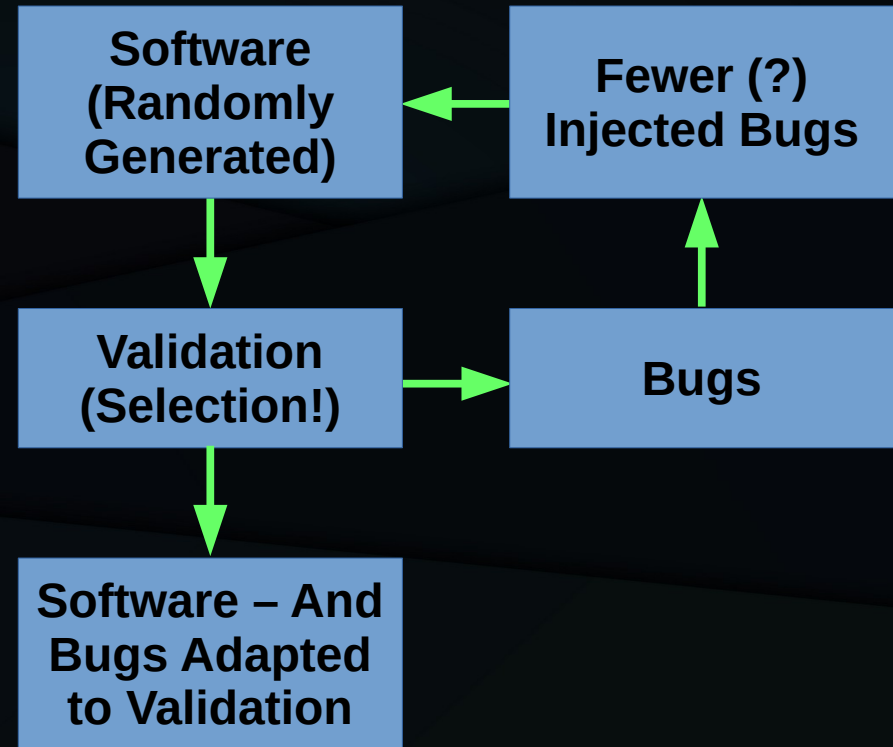


# Natural Selection: Not Just Lifeforms



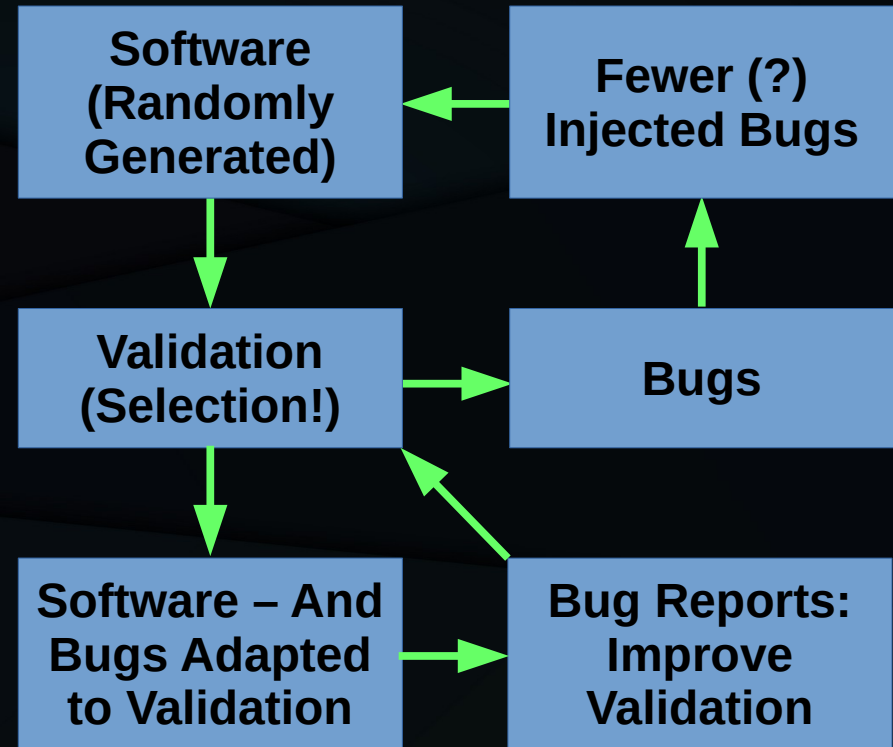
Agile methods attempt to push this methodology back to the specification

# Natural Selection: Bugs are Software!





# Natural Selection: Bugs are Software!



# Validate Only Intended Use Cases




**Current Validated  
Use Cases**




# Major Development Generates Bug




# After Validation and Bug Fixing



# After Another Round of Development

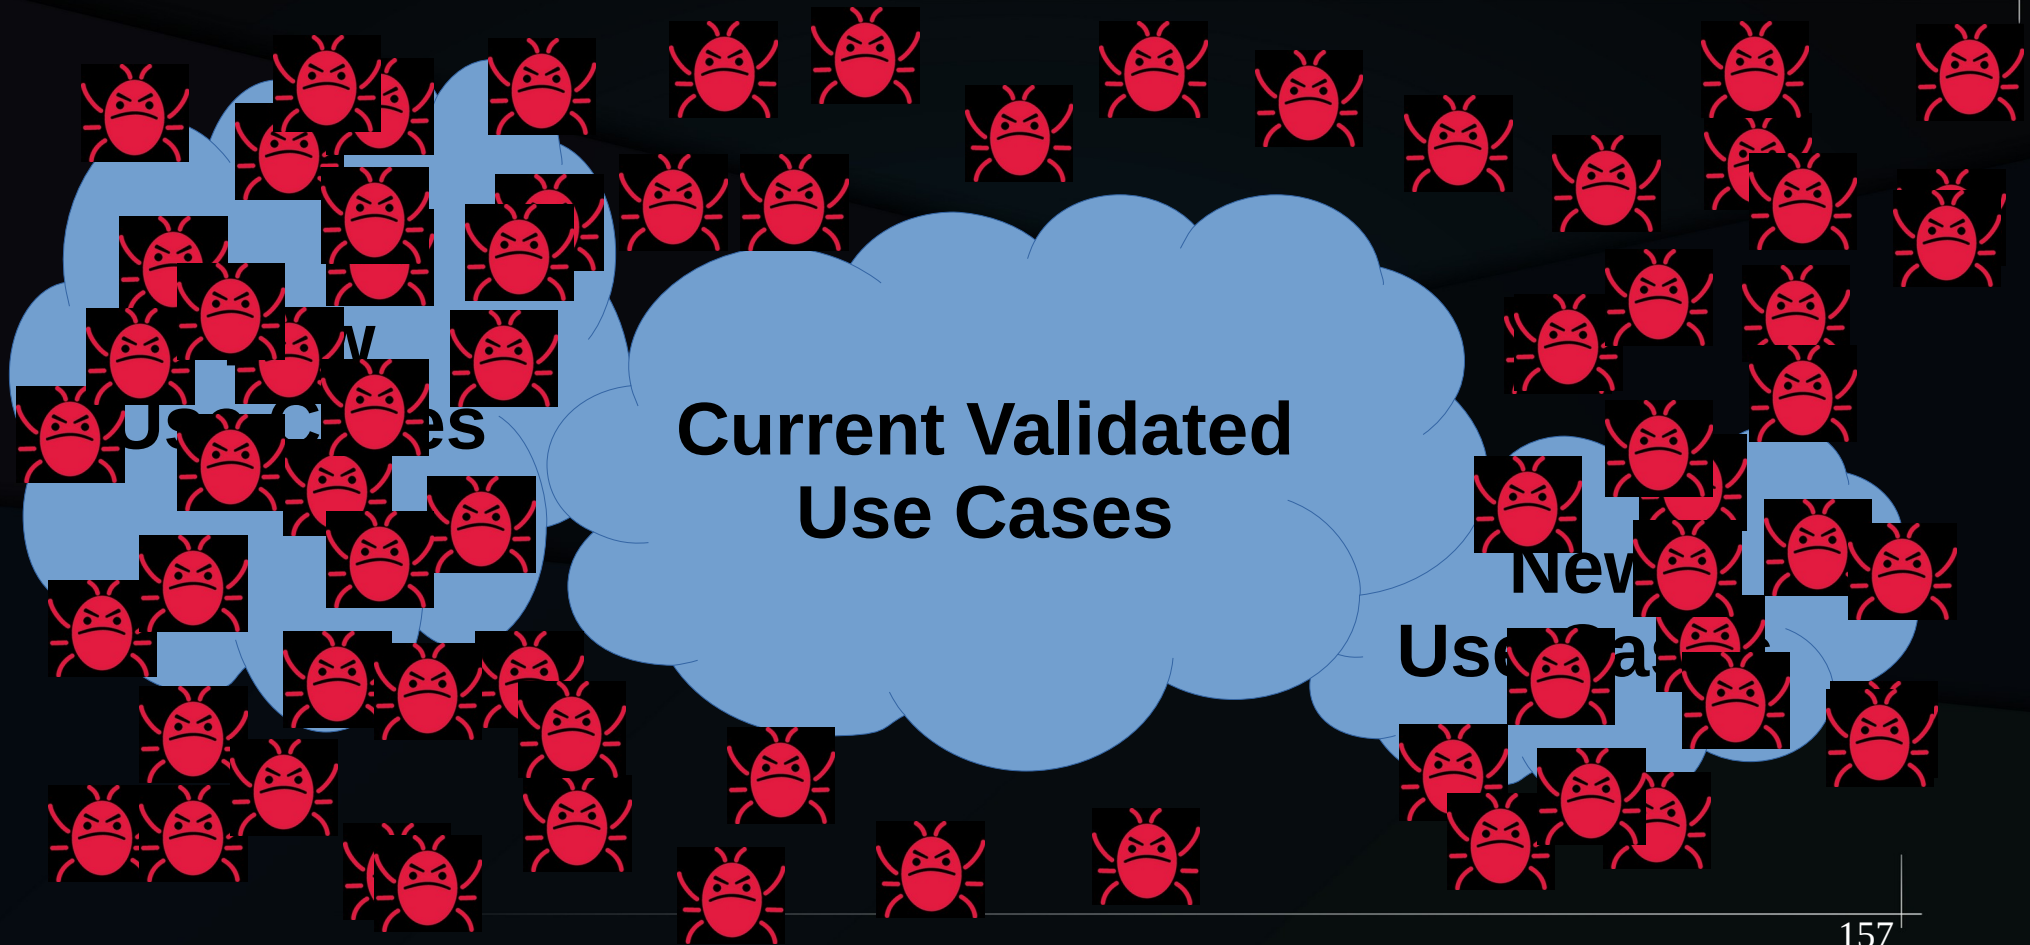


# More Validation and Bug Fixing






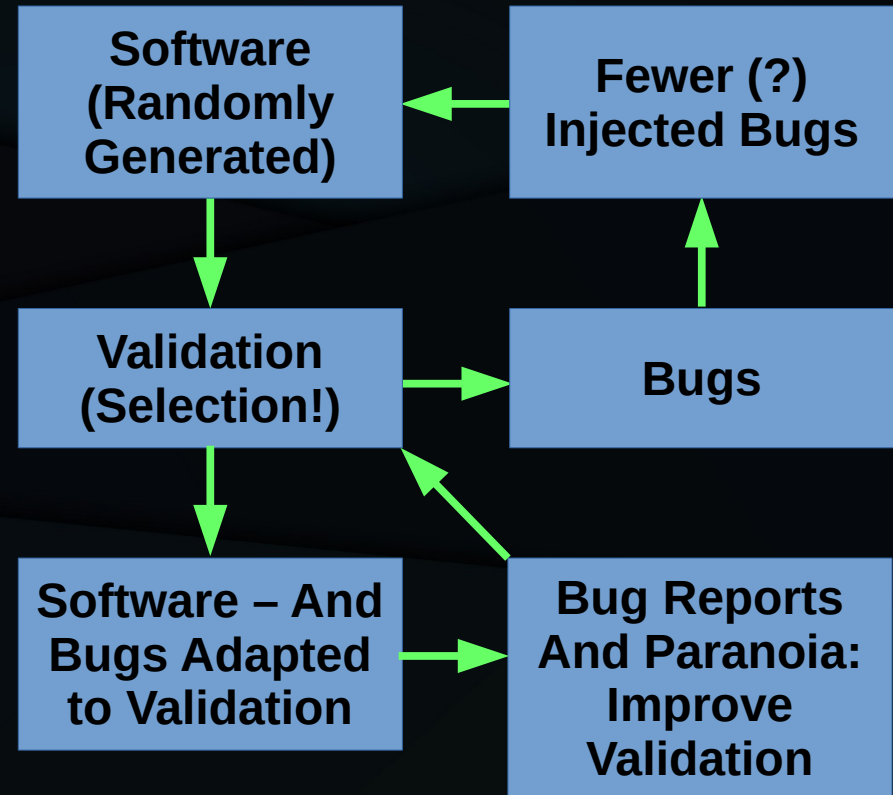

# New Use Cases: Walls of Bugs!!!



# New Use Cases: Walls of Bugs!!!



# Natural Selection: Bugs are Software!



# “Natural Selection” is a Euphemism



# **“Natural Selection” is a Euphemism**

**If your tests are not failing, they are not  
improving your software**

# **“Natural Selection” is a Euphemism**

**If your tests are not failing, they are not  
improving your software**

**If your users are not complaining, they  
are not improving your software**

# Why Would Users Fail to Complain?

- They don't know who to complain to
- The last N times they complained:
  - Nothing useful happened
  - They were yelled at or otherwise belittled
- Your software is successful
  - And has thus “faded into the woodwork”

# Cautionary Quotes

---

- Customers don't know what they want until we've shown them. - *Steve Jobs*

# Cautionary Quotes

- Customers don't know what they want until we've shown them. - *Steve Jobs*
- If there is any one secret of success, it lies in the ability to get the other person's point of view and see things from that person's angle as well as from your own. - *Henry Ford*

# Cautionary Quotes

- Customers don't know what they want until we've shown them. - Steve Jobs
- If there is any one secret of success, it lies in the ability to get into other person's point of view and see things from that person's angle as well as your own. - Henry Ford

**You must live among your users**

# Cautionary Quotes

- Customers don't know what they want until we've shown them. - Steve Jobs
- If there is any one secret of success, it lies in the ability to get the most from what you have and see it from other people's point of view and see it from other people's angle as well as your own. - Henry Ford

**You must live among your users**  
**You must complain on their behalf**

# Summary



# Summary

---

- People don't know what they want

# Summary

---

- People don't know what they want
- But for software developers, this is no excuse

# Summary

---

- People don't know what they want
- But for software developers, this is no excuse
  - You have only failed if you have given up...until then it's called learning. - *Unknown*

# Summary

---

- People don't know what they want
- But for software developers, this is no excuse
  - You have only failed if you have given up...until then it's called learning. - *Unknown*
  - You are not a failure until you start blaming others for your mistakes. - *John Wooden*

# But I Had It Easy!

- High school class: IBM mainframe & HP Basic (1973-1976)
- University: Computer science & mechanical engineering, business applications (1976-1981)
  - Do the assigned work
- Contract programming (1981-1985)
- Systems administration and Internet research (1986-1990)
- Concurrent proprietary UNIX (1990-2000)
- Linux kernel concurrency and realtime (2001-present)

# But I Had It Easy!

- High school class: IBM mainframe & HP Basic (1973-1976)
- University: Computer science & mechanical engineering, business applications (1976-1981)
- **Contract programming (1981-1985)**
  - **Keep them out of court (or even out of jail)**
- Systems administration and Internet research (1986-1990)
- Concurrent proprietary UNIX (1990-2000)
- Linux kernel concurrency and realtime (2001-present)

# But I Had It Easy!

- High school class: IBM mainframe & HP Basic (1973-1976)
- University: Computer science & mechanical engineering, business applications (1976-1981)
- Contract programming (1981-1985)
- **Systems administration and Internet research (1986-1990)**
  - **Keep users happy and fulfill terms of research contracts**
- Concurrent proprietary UNIX (1990-2000)
- Linux kernel concurrency and realtime (2001-present)

# But I Had It Easy!

- High school class: IBM mainframe & HP Basic (1973-1976)
- University: Computer science & mechanical engineering, business applications (1976-1981)
- Contract programming (1981-1985)
- Systems administration and Internet research (1986-1990)
- **Concurrent proprietary UNIX (1990-2000)**
  - **Make it fast and scalable (up to 64 CPUs)**
- Linux kernel concurrency and realtime (2001-present)



# But I Had It Easy!

- High school class: IBM mainframe & HP Basic (1973-1976)
- University: Computer science & mechanical engineering, business applications (1976-1981)
- Contract programming (1981-1985)
- Systems administration and Internet research (1986-1990)
- Concurrent proprietary UNIX (1990-2000)
- **Linux kernel concurrency and realtime (2001-present)**
  - **Make it many things...**

# But I Had It Easy!

- Linux kernel concurrency and realtime (2001-present):
  - Fast and scalable (up to 4096 CPUs)
  - Real-time response (sub-20-microsecond latencies)
  - Energy efficiency
  - Near-bare-metal efficiency to usermode applications
  - Robustness (20 billion instances)
  - Ease of use driven by security
  - Ease of administration (large data centers)

# But I Had It Easy!

- Linux kernel concurrency and realtime (2001)
  - Fast and scalable (up to 4096 CPUs)
  - Real-time response (sub-20ms)
  - Energy efficiency
  - Near-barren hardware
  - Ease of use
  - Ease of integration (large data centers)

Yes, I am proud of my accomplishments, but modern systems are far more complex and user-centric

# But I Had It Easy!

- Linux kernel concurrency and realtime (2001)

- Fast and (up to 4096 CPUs)

- Real-time response

- Energy efficiency

- Near-bar

- Ease

- Ease of use (large data centers)

**My job is to provide accomplishments, not to provide infrastructure**

**Yes, I am proud of my accomplishments, but modern systems are far more complex and more reliable**

# Questions?

# Questions?

