

# SafeFetch:

Practical Double-Fetch Protection with  
Kernel-Fetch Caching

Victor Duta

Mitchel Alonserij   Cristiano Giuffrida



# Teaser

- Kernel **double-fetch** bugs
- State-of-the-art relies on COW semantics
- SafeFetch idea => per-syscall caching of user data
- Low performance impact:
  - Less than **5%** geomean overhead

# Double-Fetch Anatomy



Thread @A

User  
context

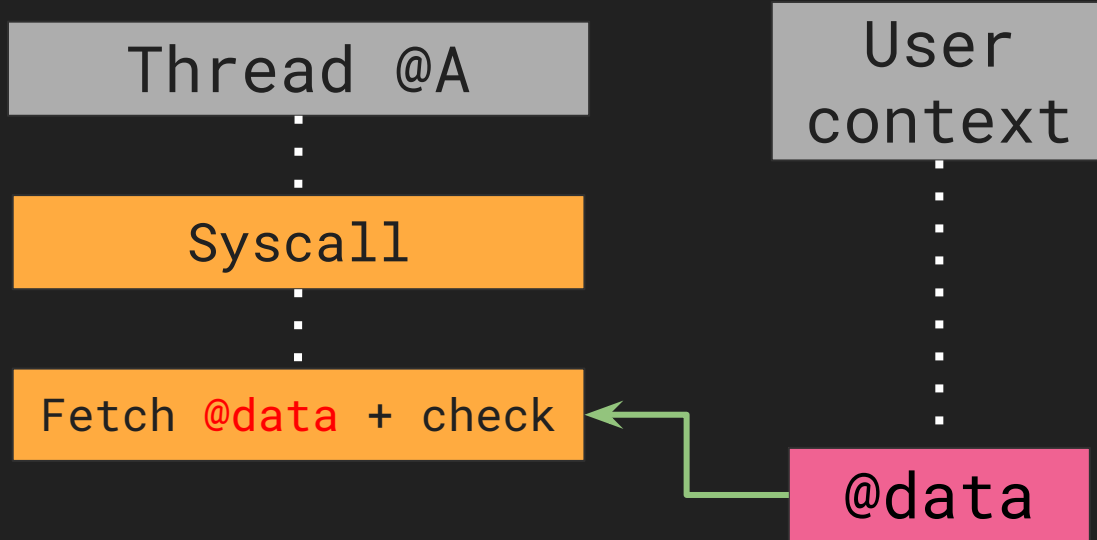


@data

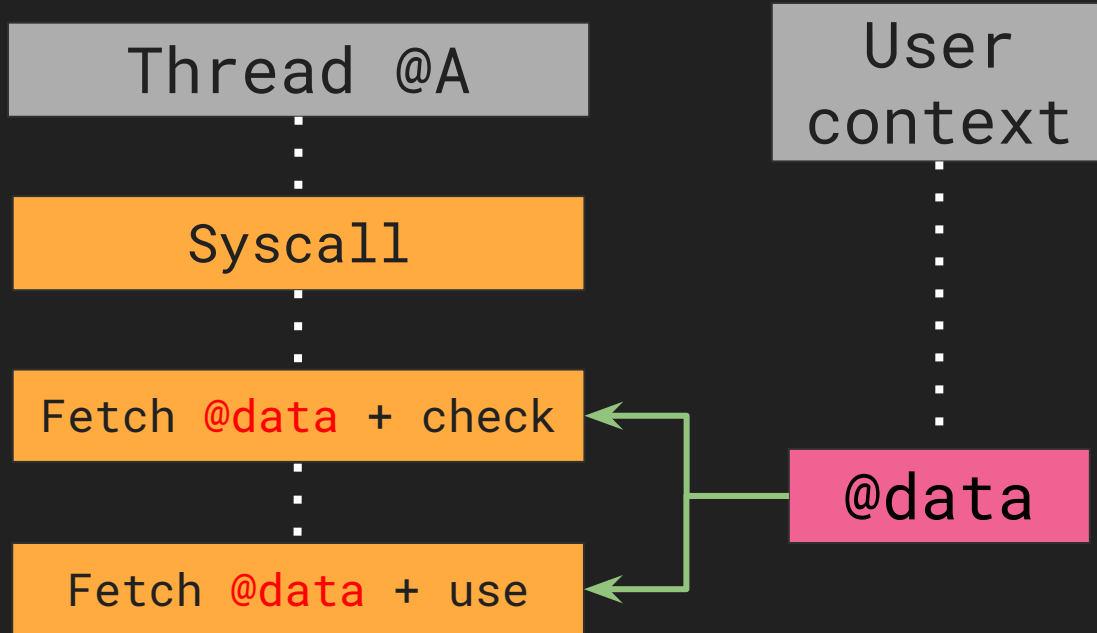
# Double-Fetch Anatomy



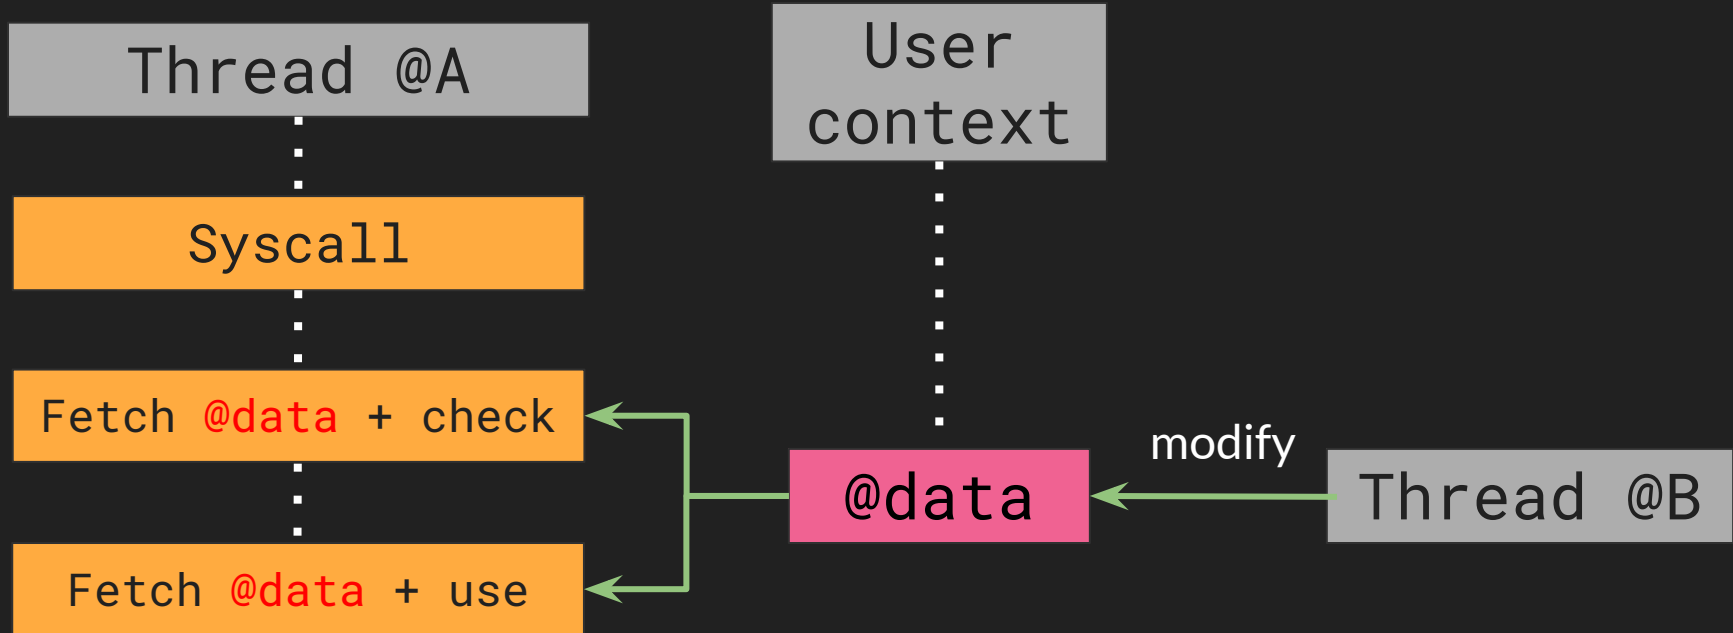
# Double-Fetch Anatomy



# Double-Fetch Anatomy



# Double-Fetch Anatomy



# Midas

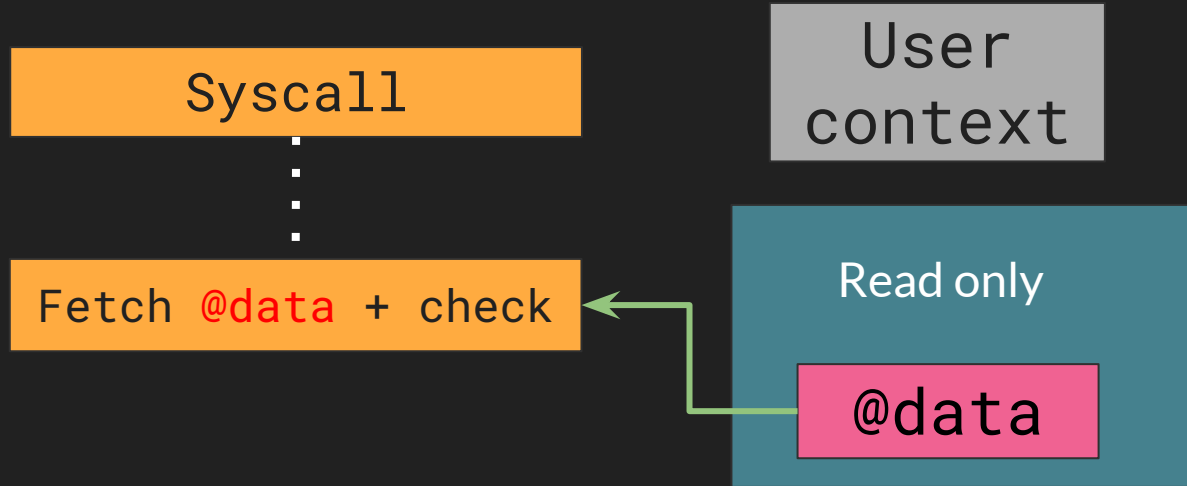
syscall

User  
context

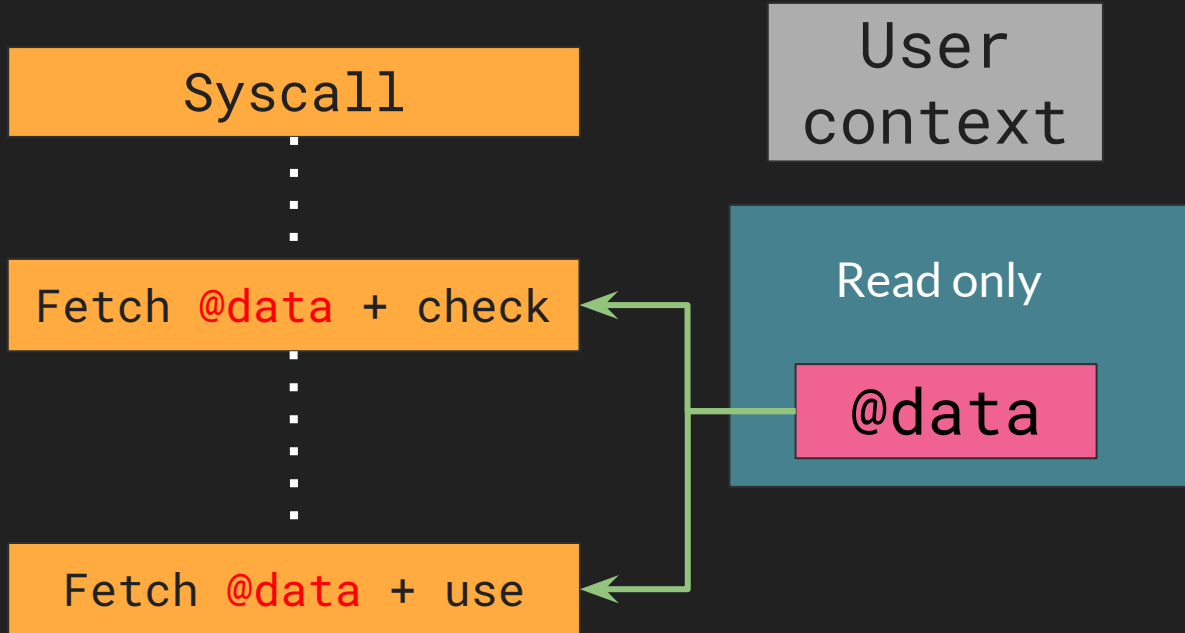
@data



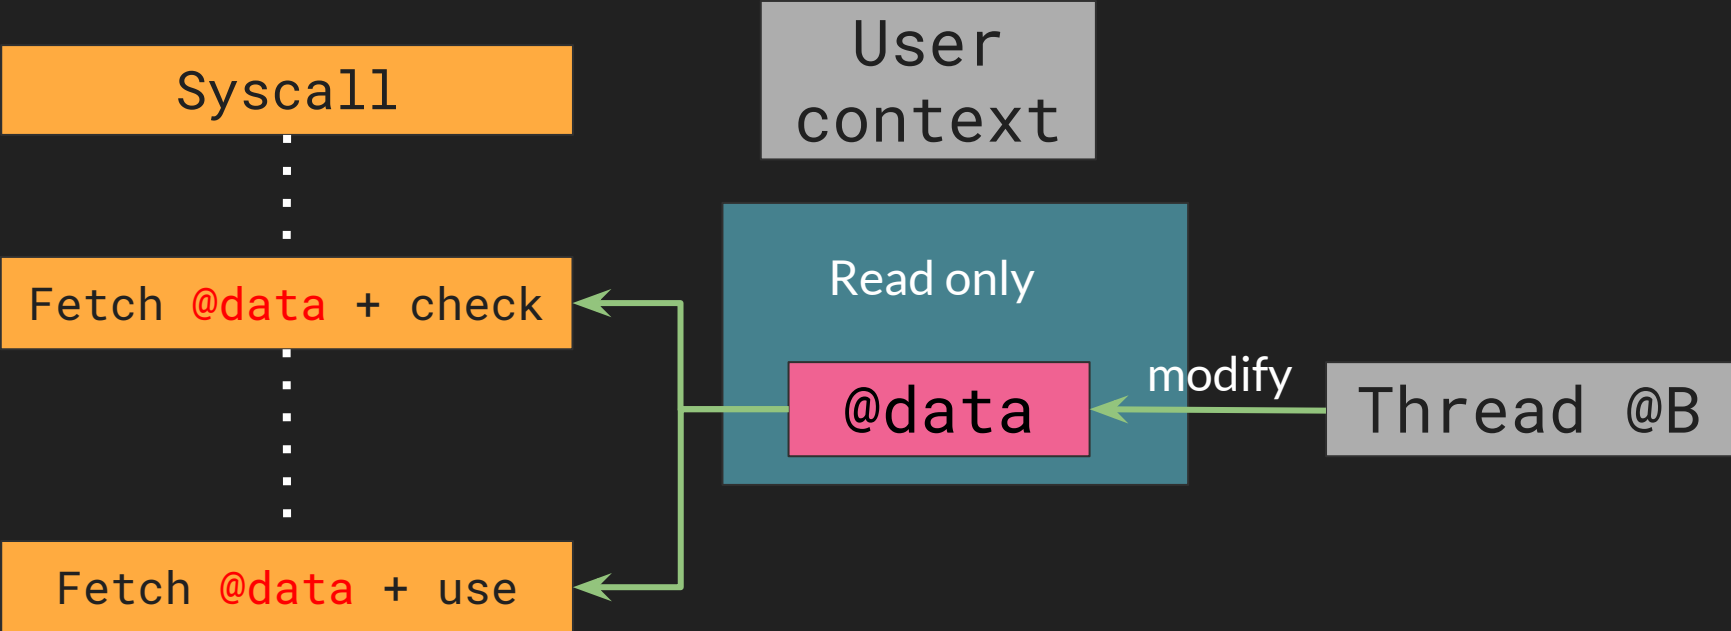
# Midas



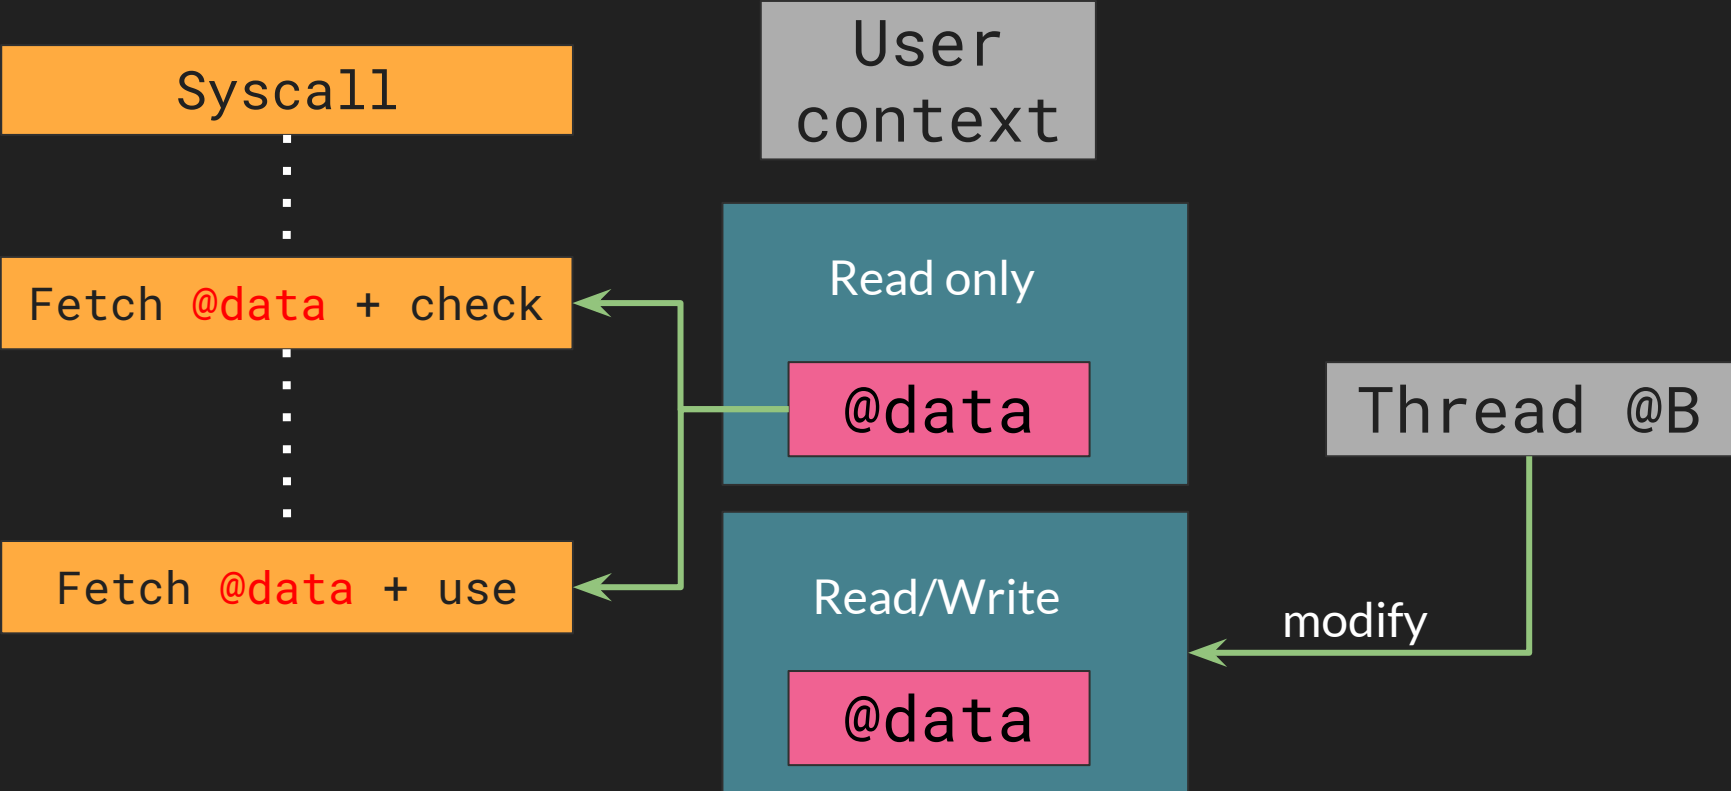
# Midas



# Midas



# Midas



# Key Insights

# Key Insights

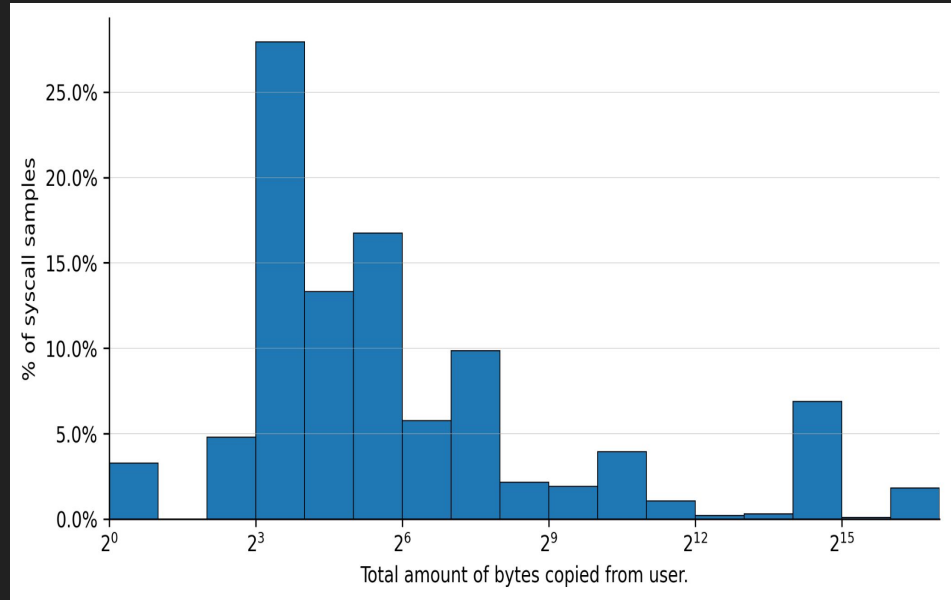
- 60% of syscall executions do not fetch **user** data

# Key Insights

- 60% of syscall executions do not fetch **user** data
- On average syscalls fetch far less than a **page** of data

# Key Insights

- 60% of syscall executions do not fetch user data
- On average syscalls fetch far less than a page of data

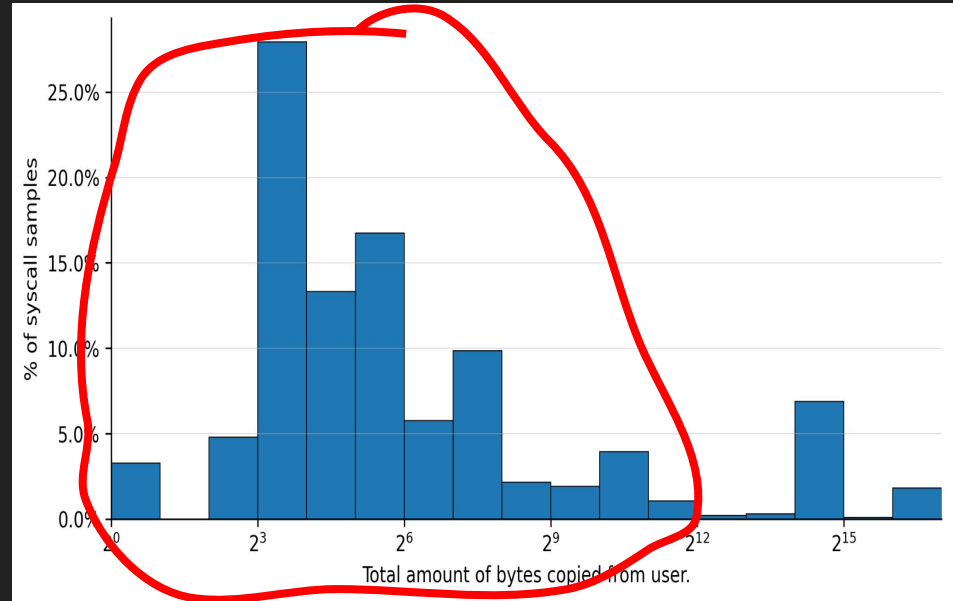




# Key Insights

- 60% of syscall executions do not fetch user data
- On average syscalls fetch far less than a page of data

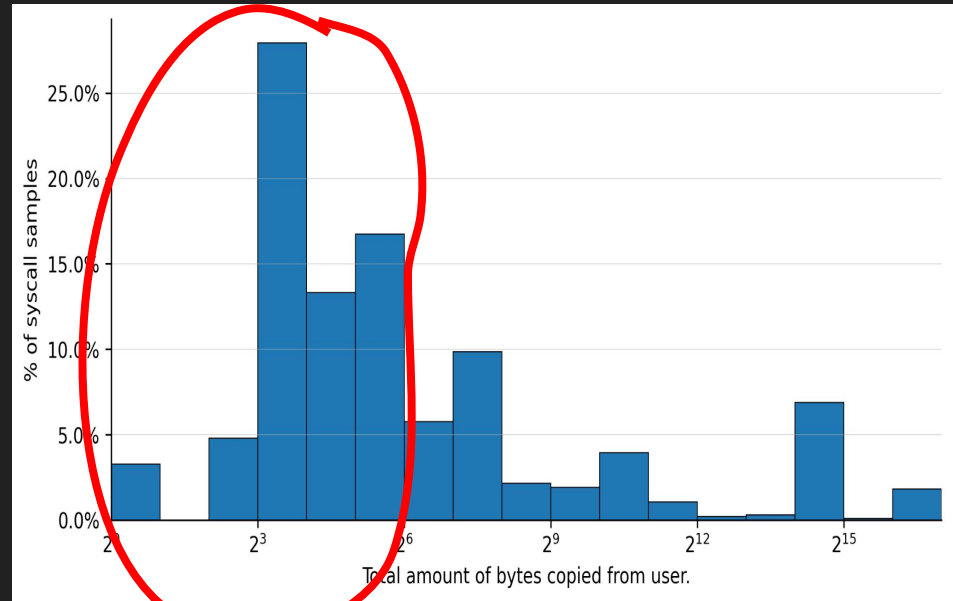
90% of syscalls fetch  $\leq 1$  Page of user data



# Key Insights

- 60% of syscall executions do not fetch user data
- On average syscalls fetch far less than a page of data

90% of syscalls fetch  $\leq 1$  Page of user data

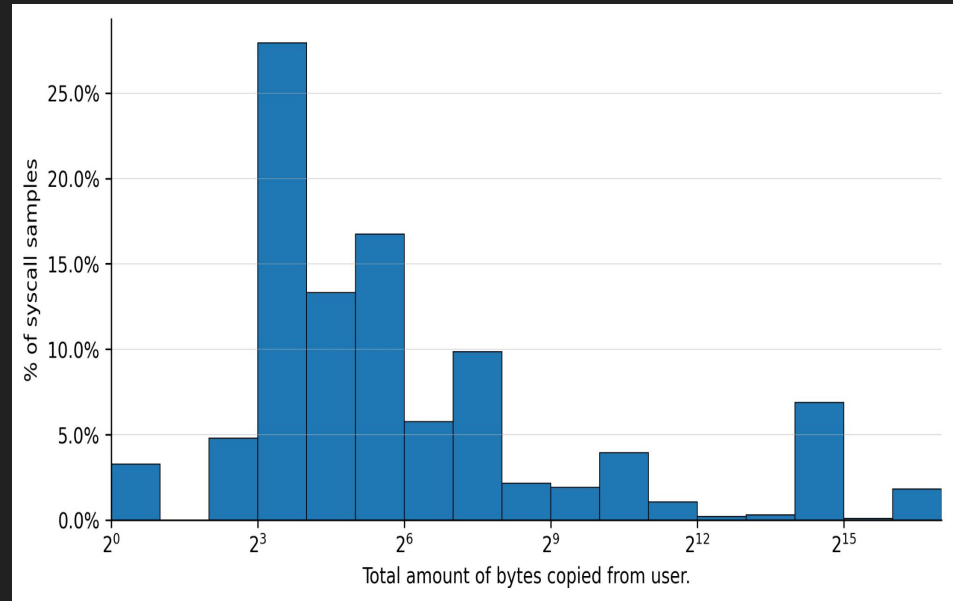


66% of syscalls fetch  $< 64$  bytes of data

# Key Insights

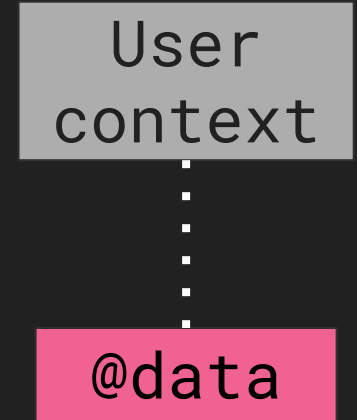

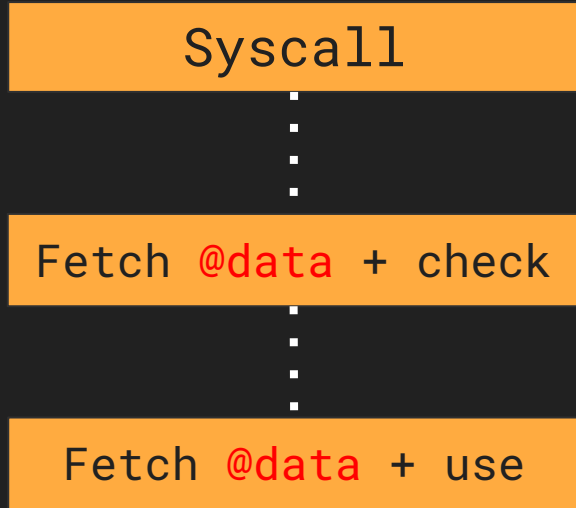
- 60% of syscall executions do not fetch **user** data
- On average syscalls fetch far less than a **page** of data
- Caching is more efficient

90% of syscalls fetch  $\leq 1$  Page of user data

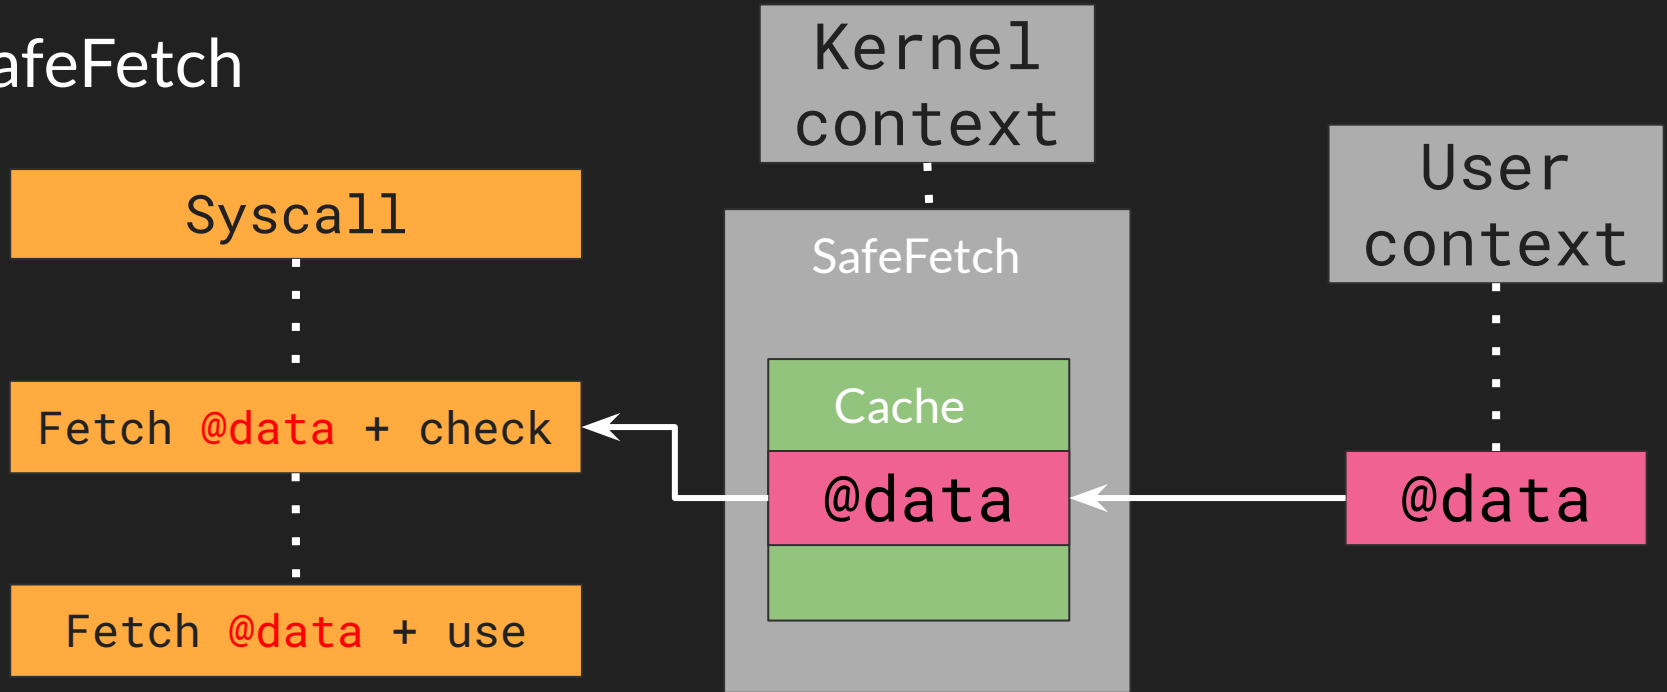


66% of syscalls fetch  $< 64$  bytes of data

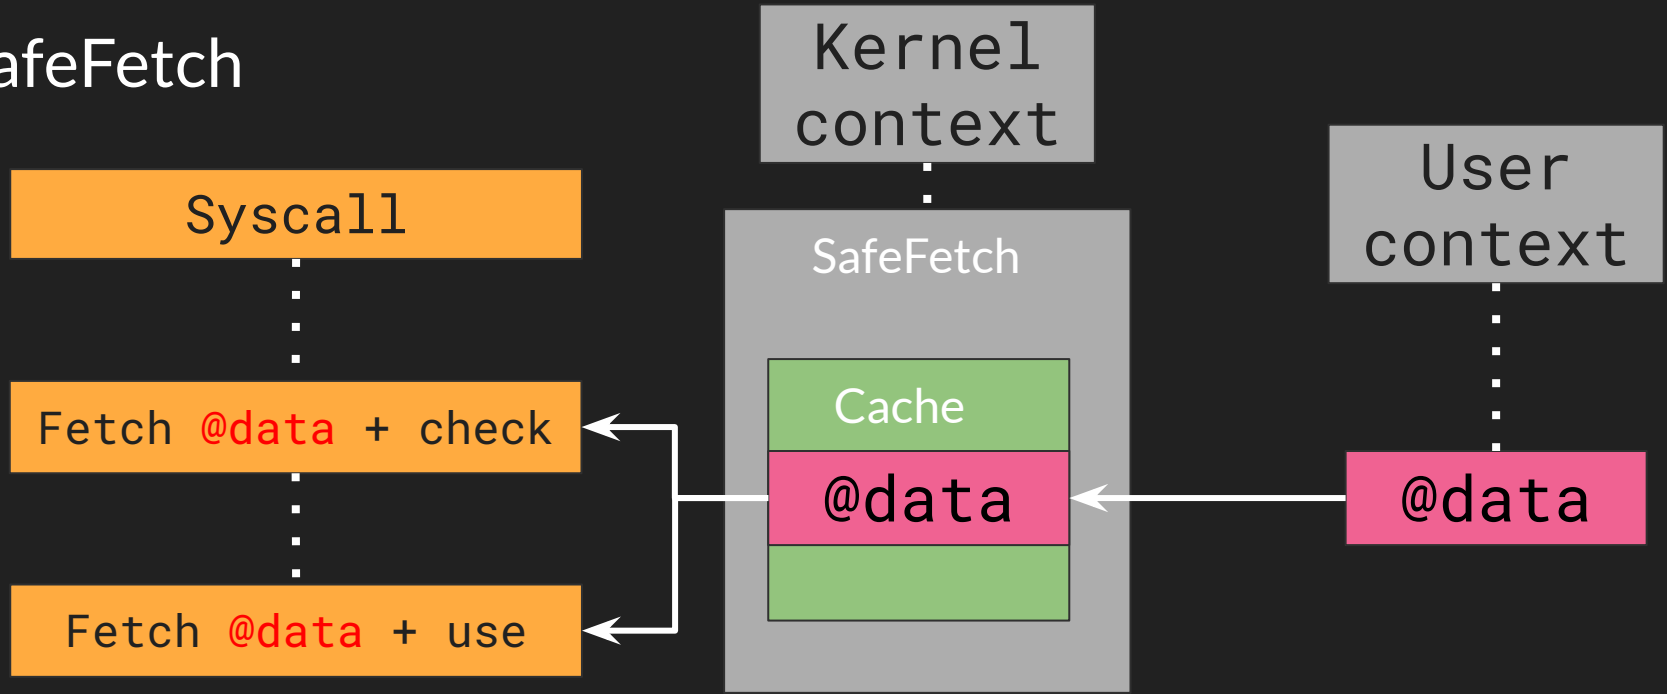
# SafeFetch



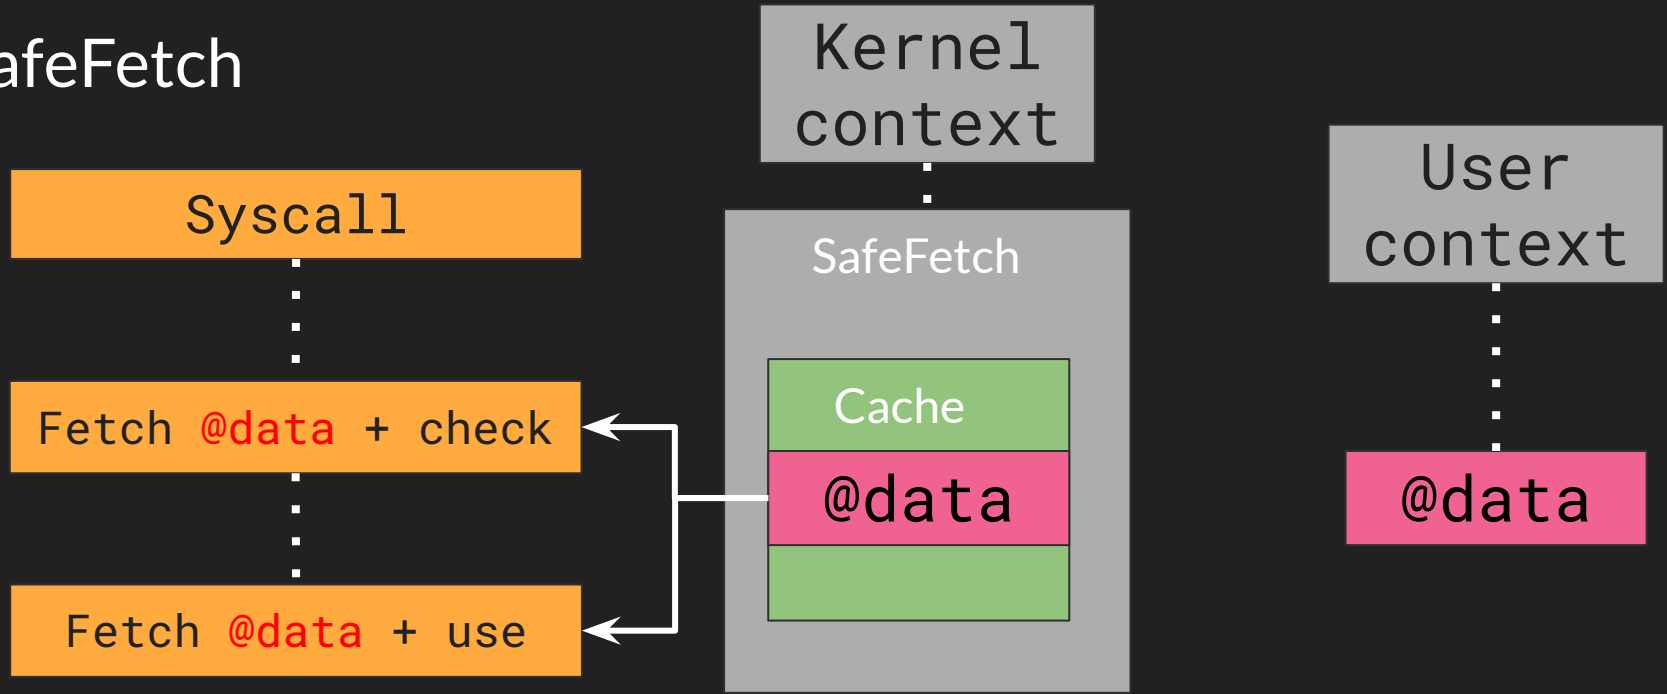
# SafeFetch



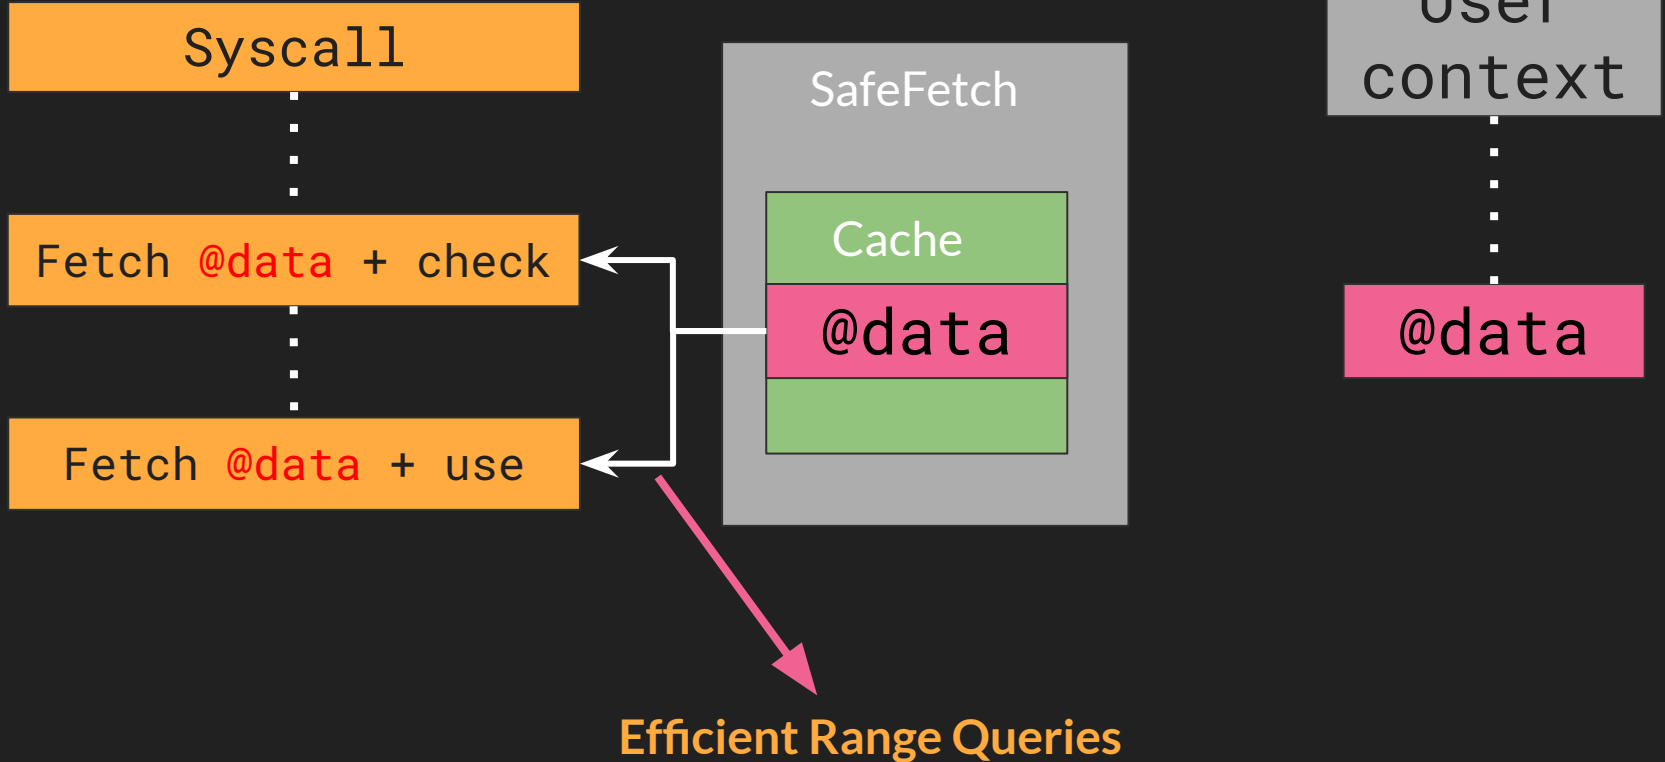
# SafeFetch



# SafeFetch

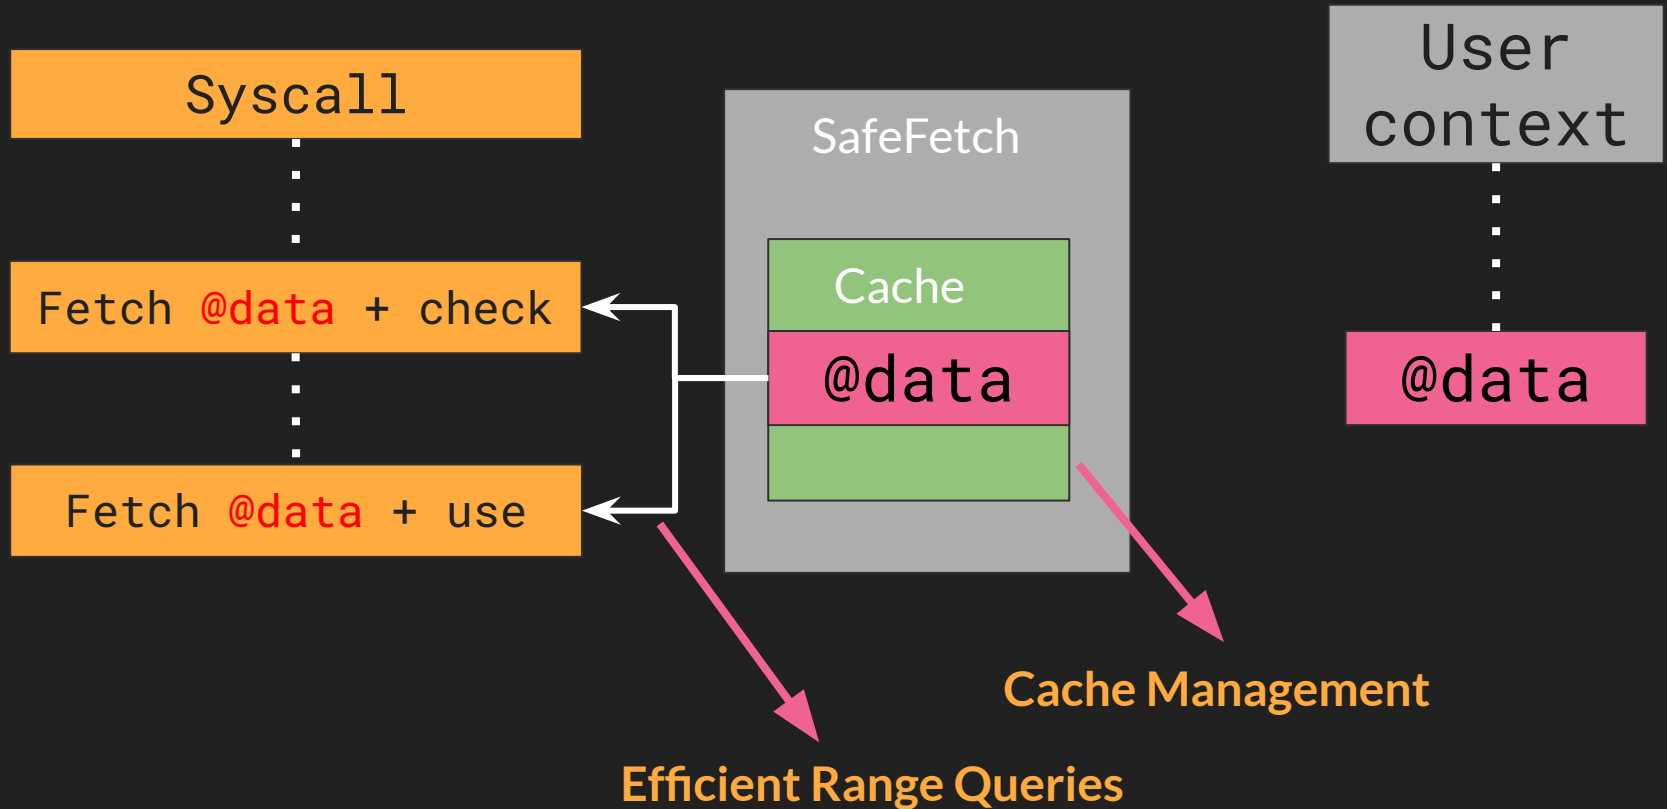


# SafeFetch






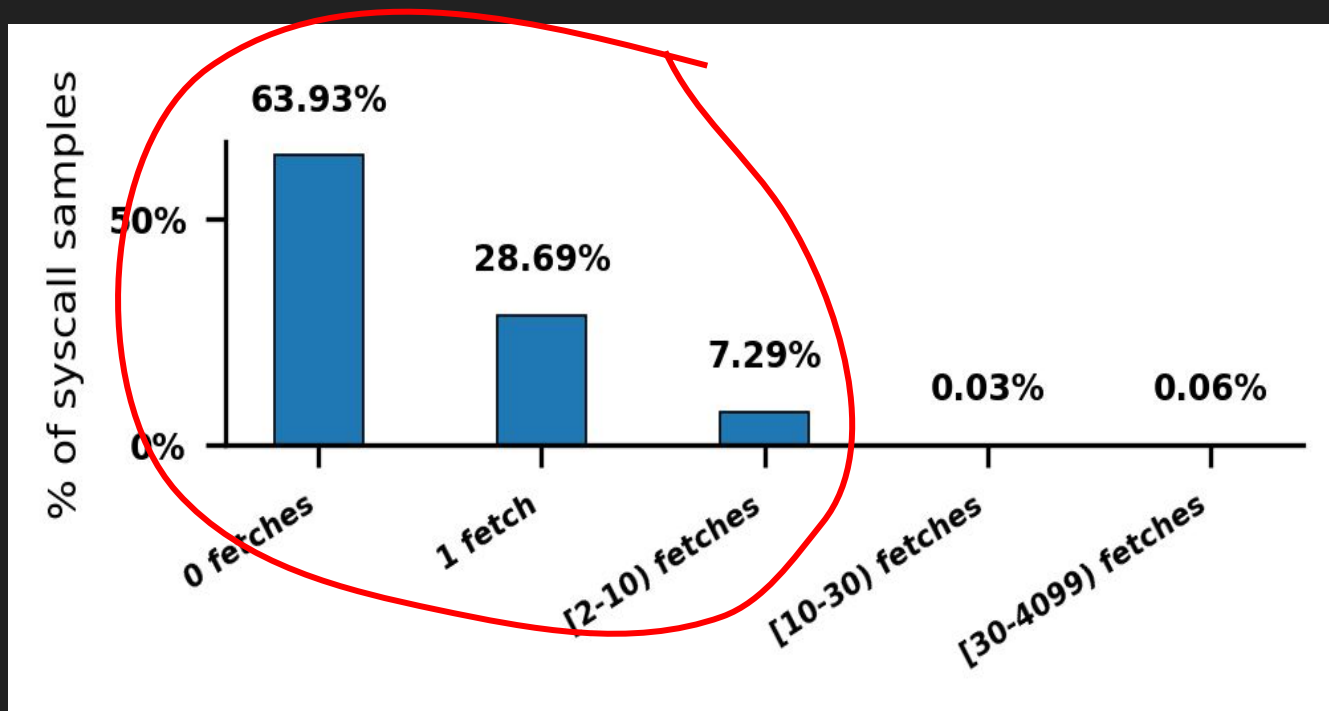
# SafeFetch




# Efficient queries



# Efficient queries

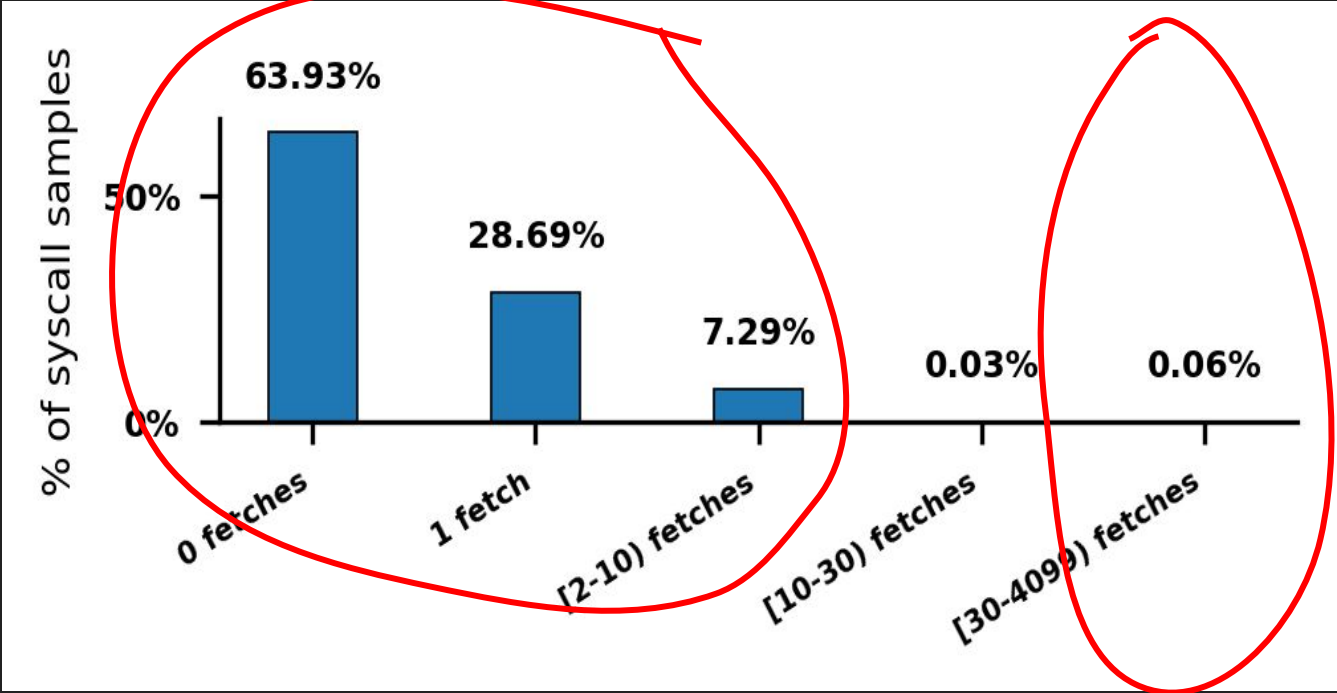


# Efficient queries




# Efficient queries

linked list




# Efficient queries

linked list



rb-tree



# Cache Management

# Cache Management

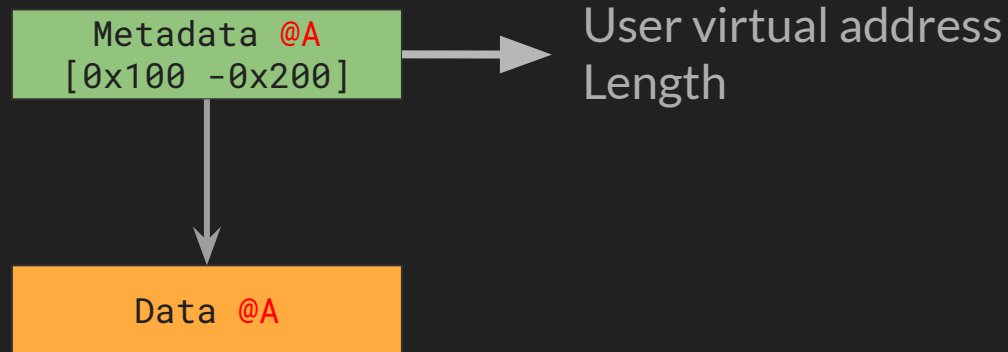
Metadata @A  
[0x100 - 0x200]



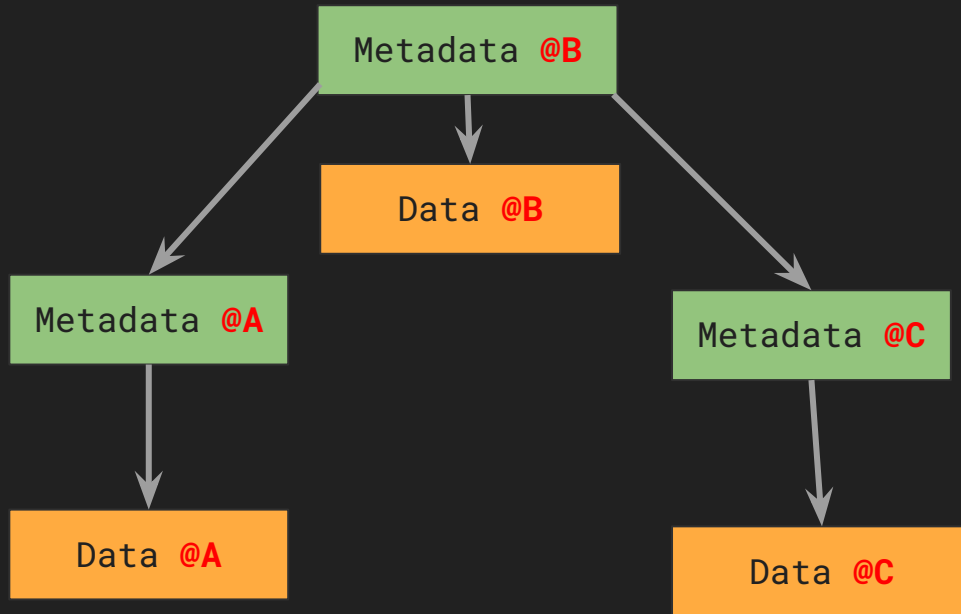
# Cache Management



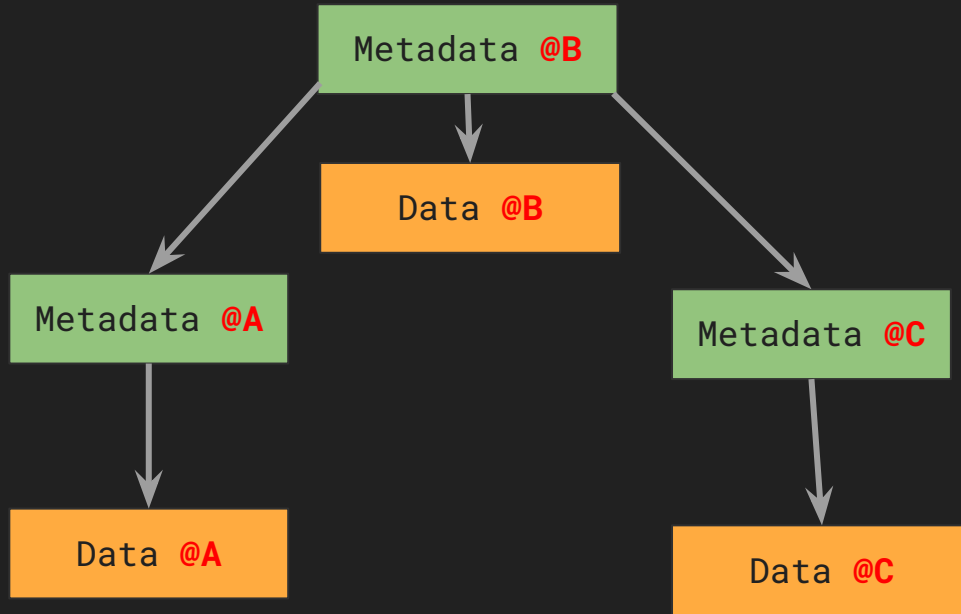
# Cache Management



# Cache management




# Cache management



Region  
allocator


# Cache management



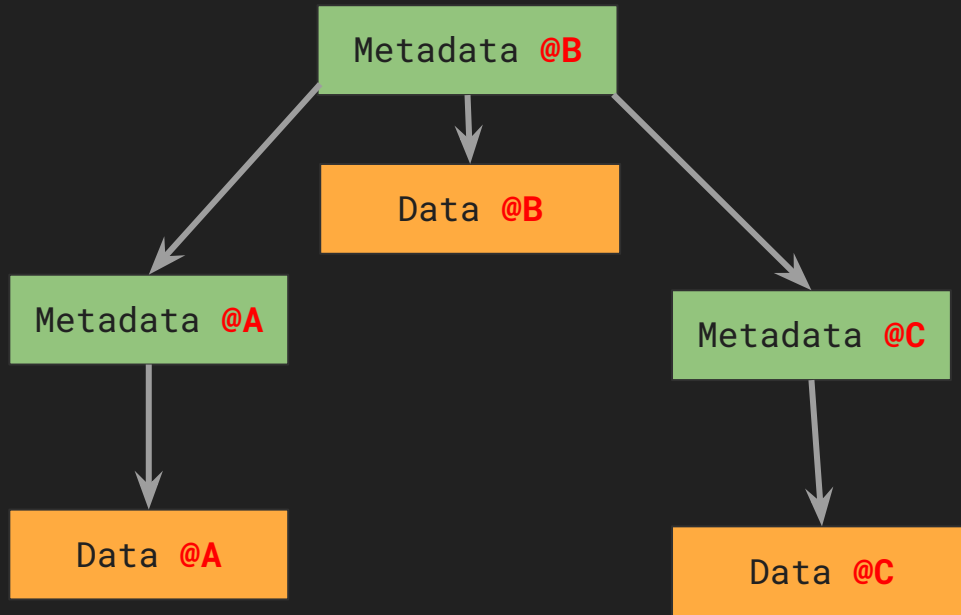
Region  
allocator

Metadata

Data




# Cache management



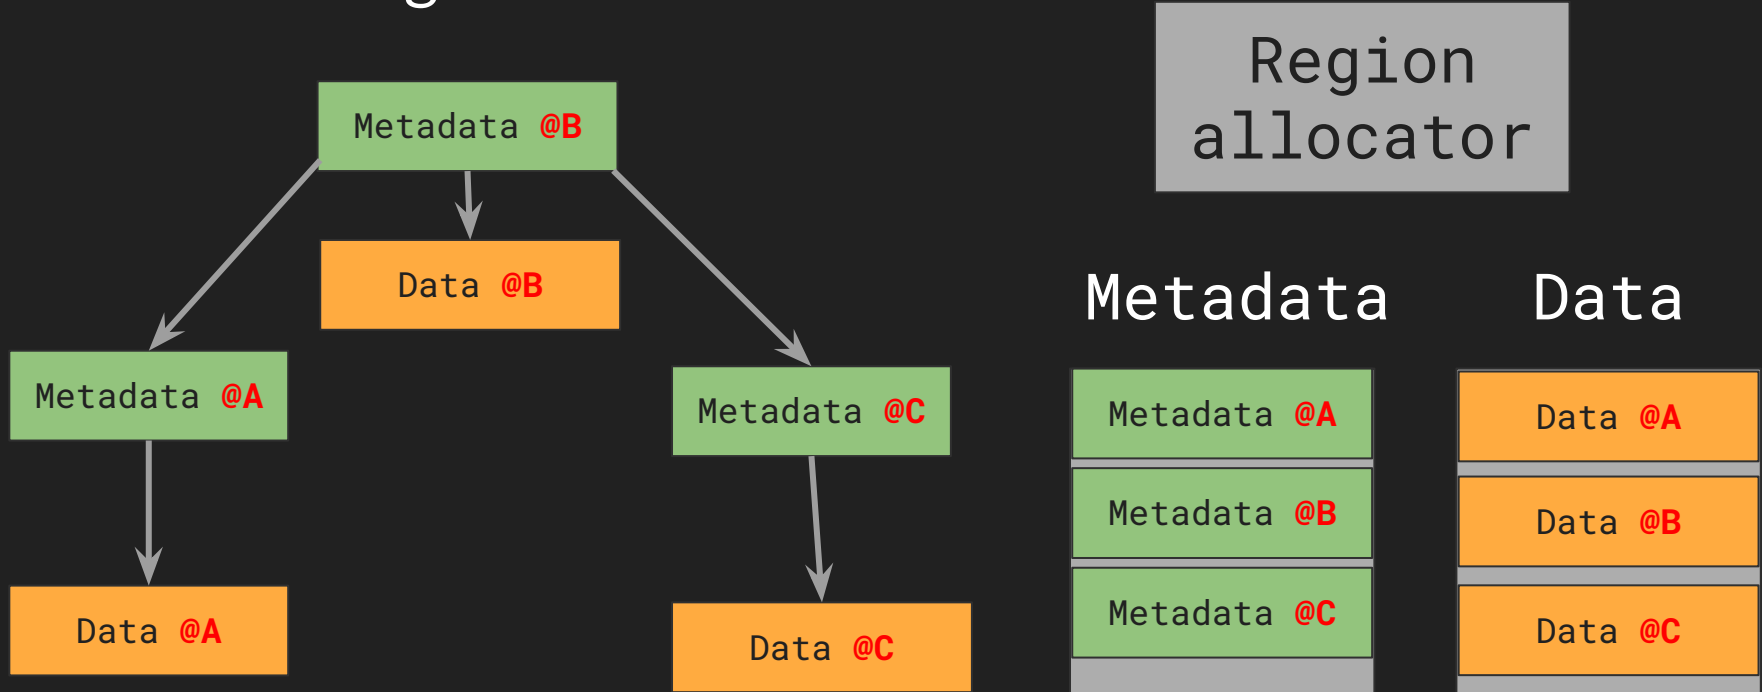
Region  
allocator

Metadata

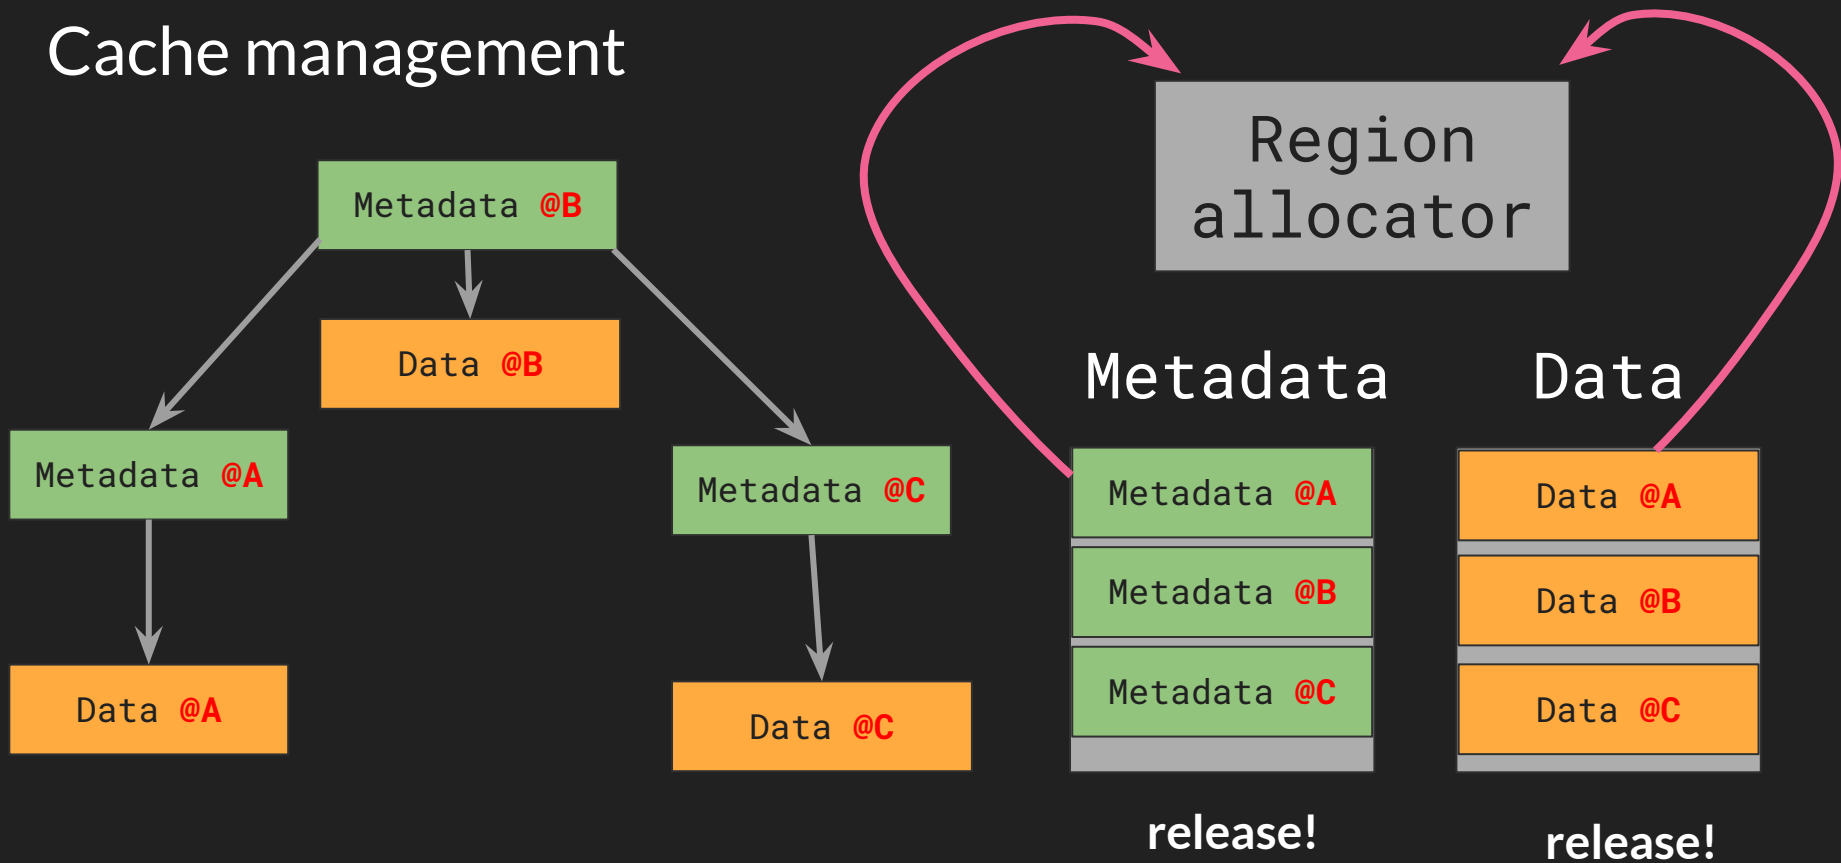
Data



# Cache management



# Cache management







# Results




# Results



# Results



# Results



# Conclusion

- Comprehensive protection against **double-fetch** bugs
- Less than **5%** geometric overhead across a variety of benchmarks
- CVE-2016-6516

Questions?



Github Repository

# Memory utilization

Table 3: SafeFetch-induced memory utilization.

Benchmark	SafeFetch raw memory utilization					
	cache provisioning			zero-copy pins		
	avg.	99 <sup>th</sup>	peak	avg.	99 <sup>th</sup>	peak
LMBench	8.02 KB	8 KB	236 KB	0.65 KB	0 KB	9.6 MB
OSBench	8.04 KB	12 KB	12 KB	0 KB	0 KB	0 KB
Phoronix	8.04 KB	8 KB	44 KB	5.82 KB	16 KB	128 KB

# (Double-)Fetch Rate

Table 5: Statistics for (double) fetch rates.

<b>Benchmark</b>	<b>Statistic</b>	<b>Fetches</b>	<b>Double Fetches</b>
LMBench	<b>Rate</b>	<b>1/2</b>	<b>1/273457</b>
	<b>Syscalls</b>	<b>38</b>	<b>7</b>
	<b>Min/Avg/Max</b>	<b>1/1/459</b>	<b>1/50/67</b>
OSBench	<b>Rate</b>	<b>1/3</b>	<b>1/80</b>
	<b>Syscalls</b>	<b>17</b>	<b>5</b>
	<b>Min/Avg/Max</b>	<b>1/6/134</b>	<b>1/42/43</b>
Phoronix	<b>Rate</b>	<b>1/4</b>	<b>1/5993</b>
	<b>Syscalls</b>	<b>47</b>	<b>21</b>
	<b>Min/Avg/Max</b>	<b>1/1/661</b>	<b>1/1/218</b>
Background	<b>Rate</b>	<b>1/2</b>	<b>1/644</b>
	<b>Syscalls</b>	<b>93</b>	<b>20</b>
	<b>Min/Avg/Max</b>	<b>1/2/4099</b>	<b>1/130/467</b>



# Optimization benefits

Table 4: SafeFetch-induced throughput degradation.


Benchmark	SafeFetch	
	/wo zero-copy (%)	/w zero-copy (%)
AF_UNIX	39.1%	8.4%
Pipe	16.1%	5.1%
Nginx	5.9%	1.3%
Apache	9.3%	1.8%

# Cache management


Metadata




Data



# Cache management



# Cache management



2 page allocations + 2 page releases + extra writes



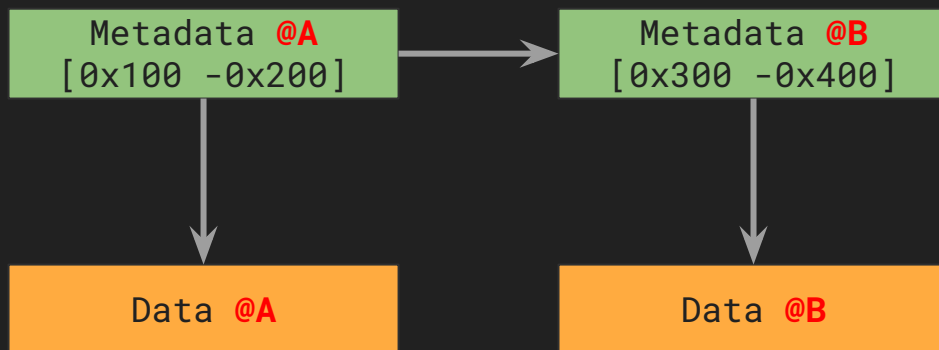


Figure 13: Average normalized time for each fetch vs. number of cached ranges.

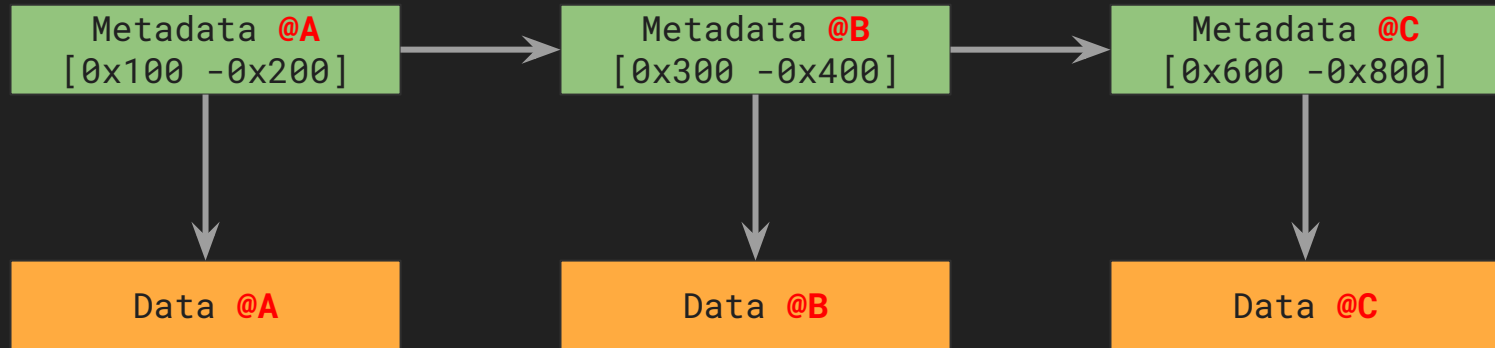
# Results



# Efficient queries

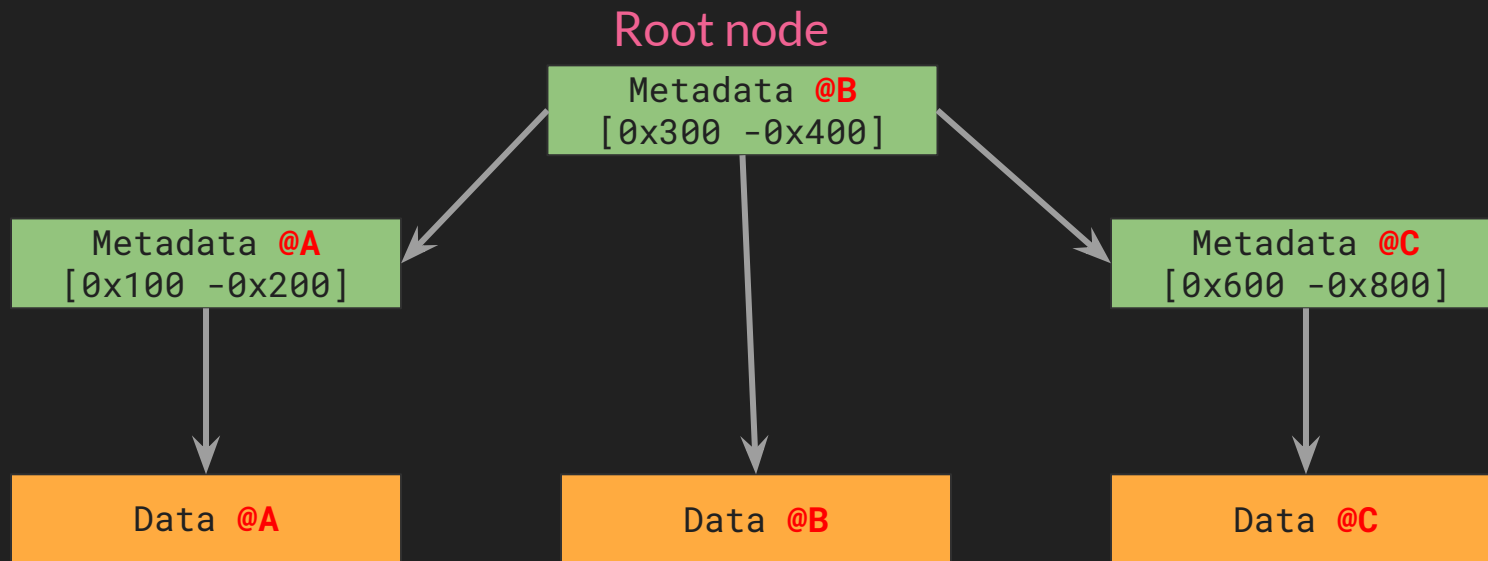


# Efficient queries






# Efficient queries



# Cache management



Region allocator

Metadata

Data

