












# DoS Protection

using multi-prefix query counting











# Overview

- amplification att.



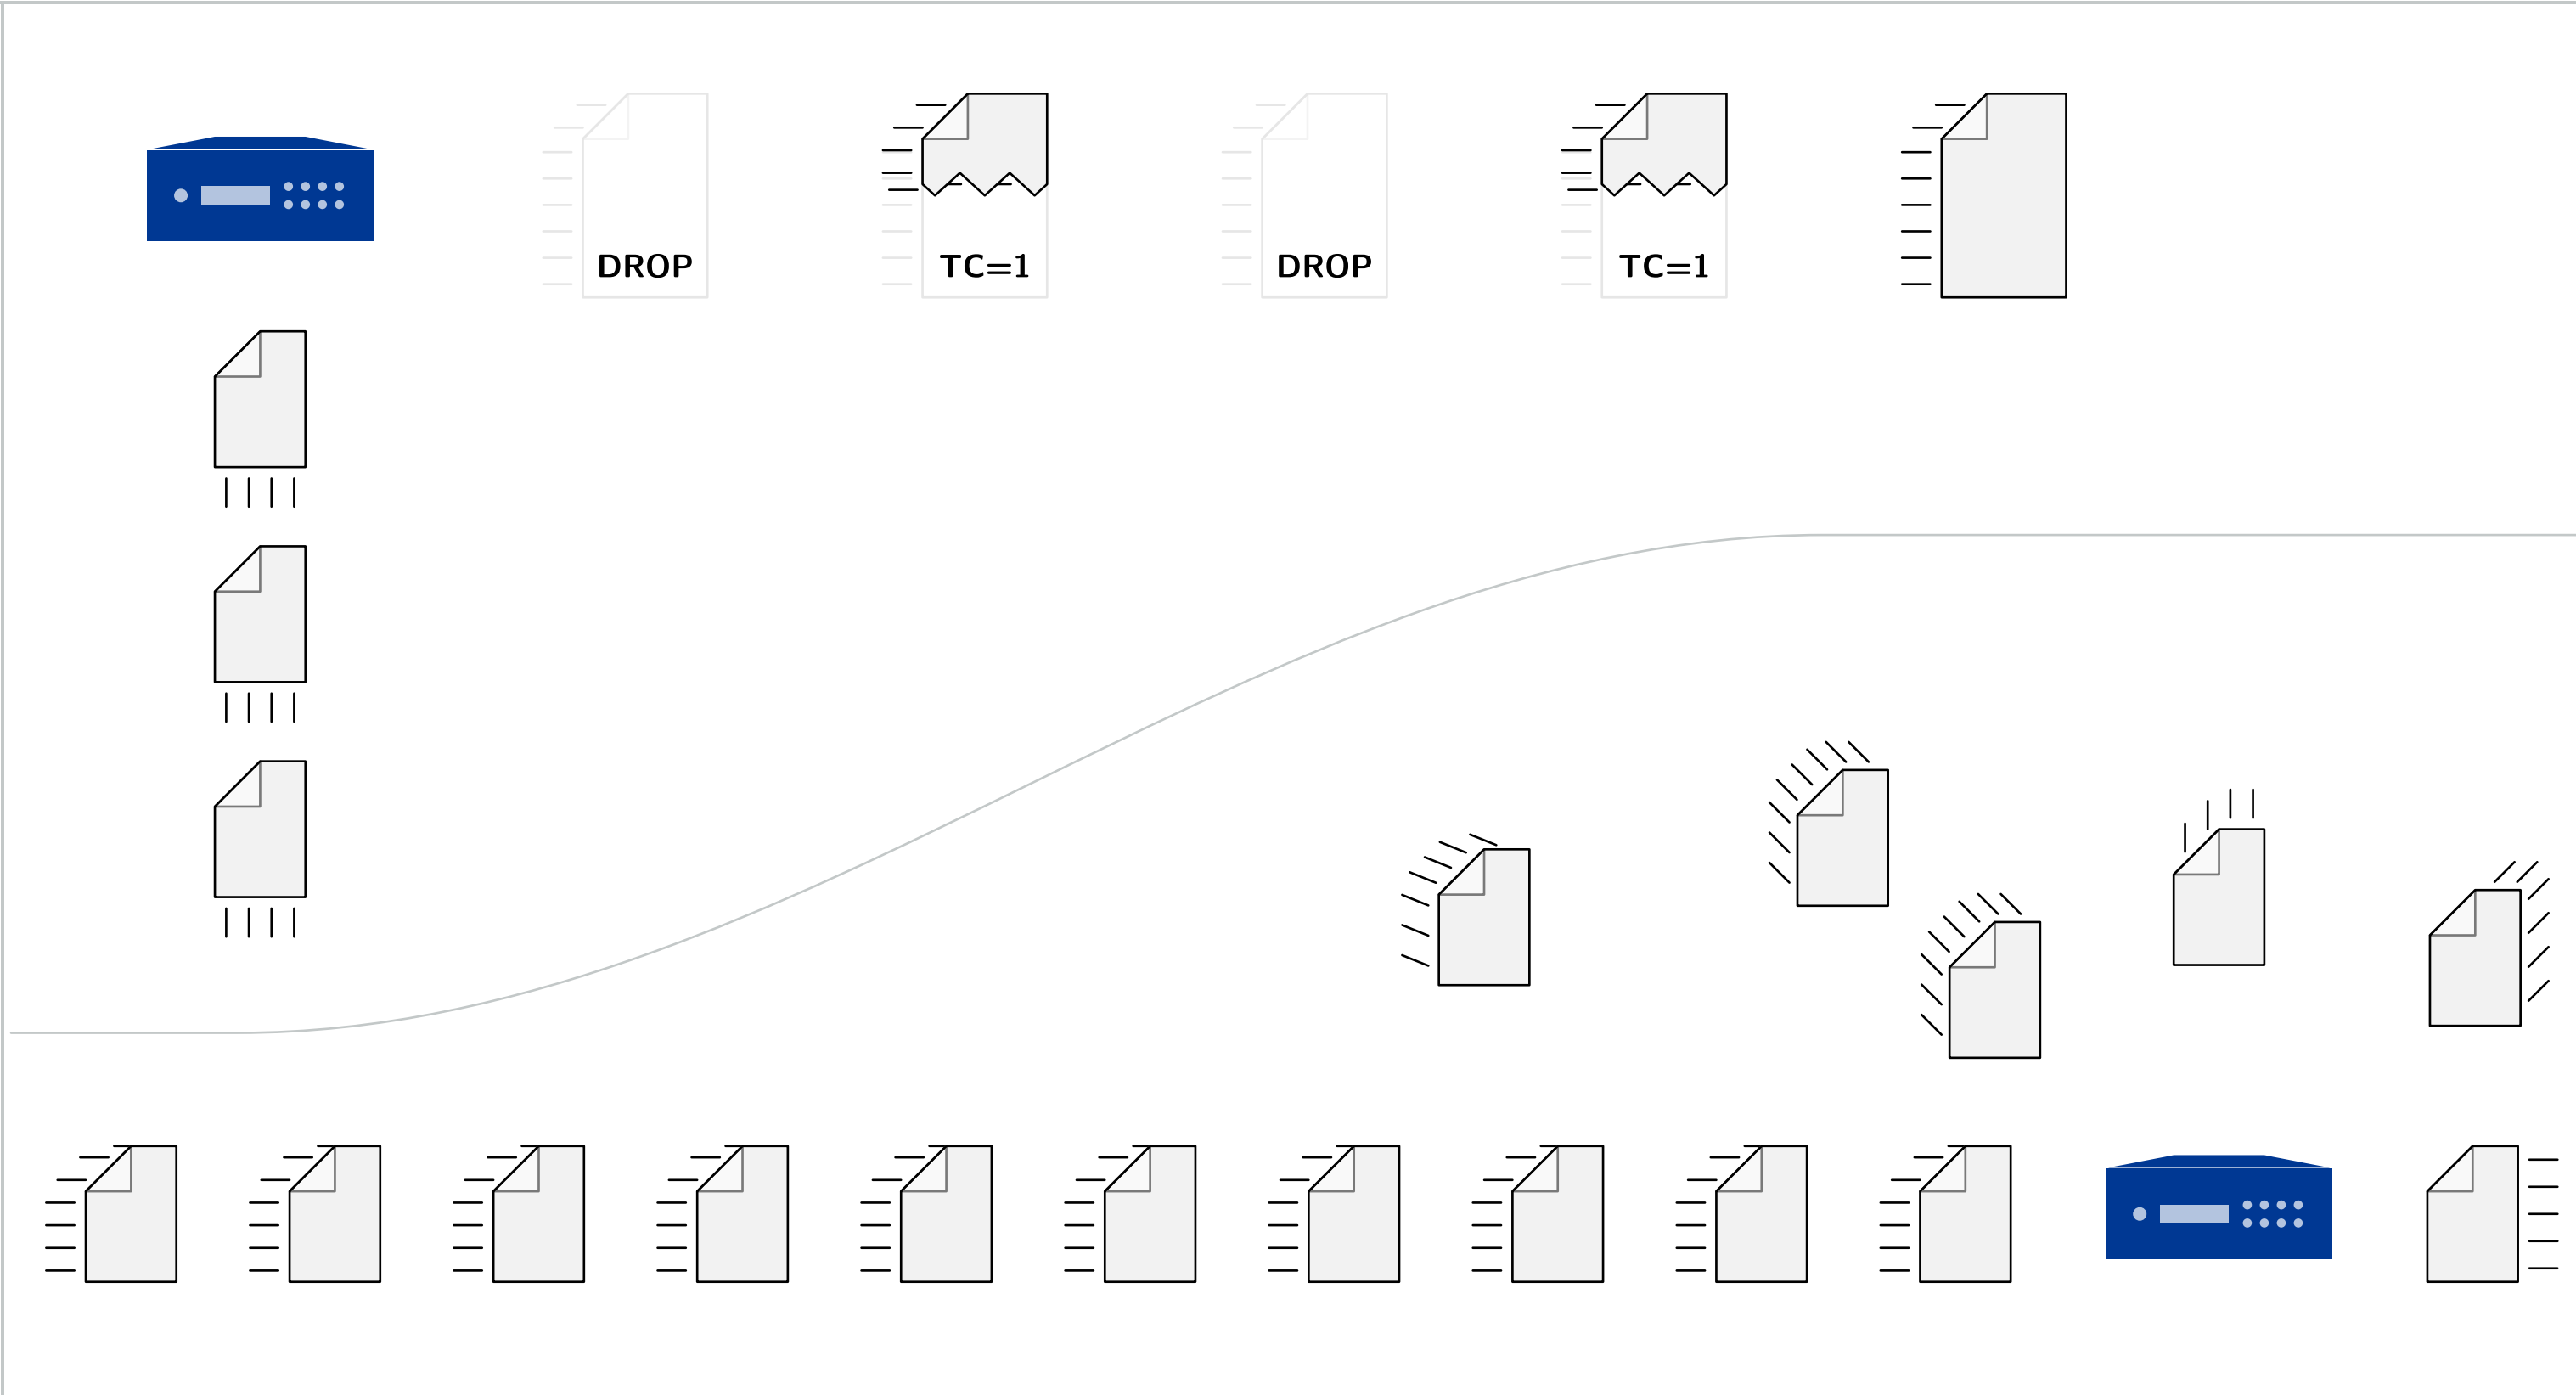
# Overview

- amplification att.
  - rate limiting



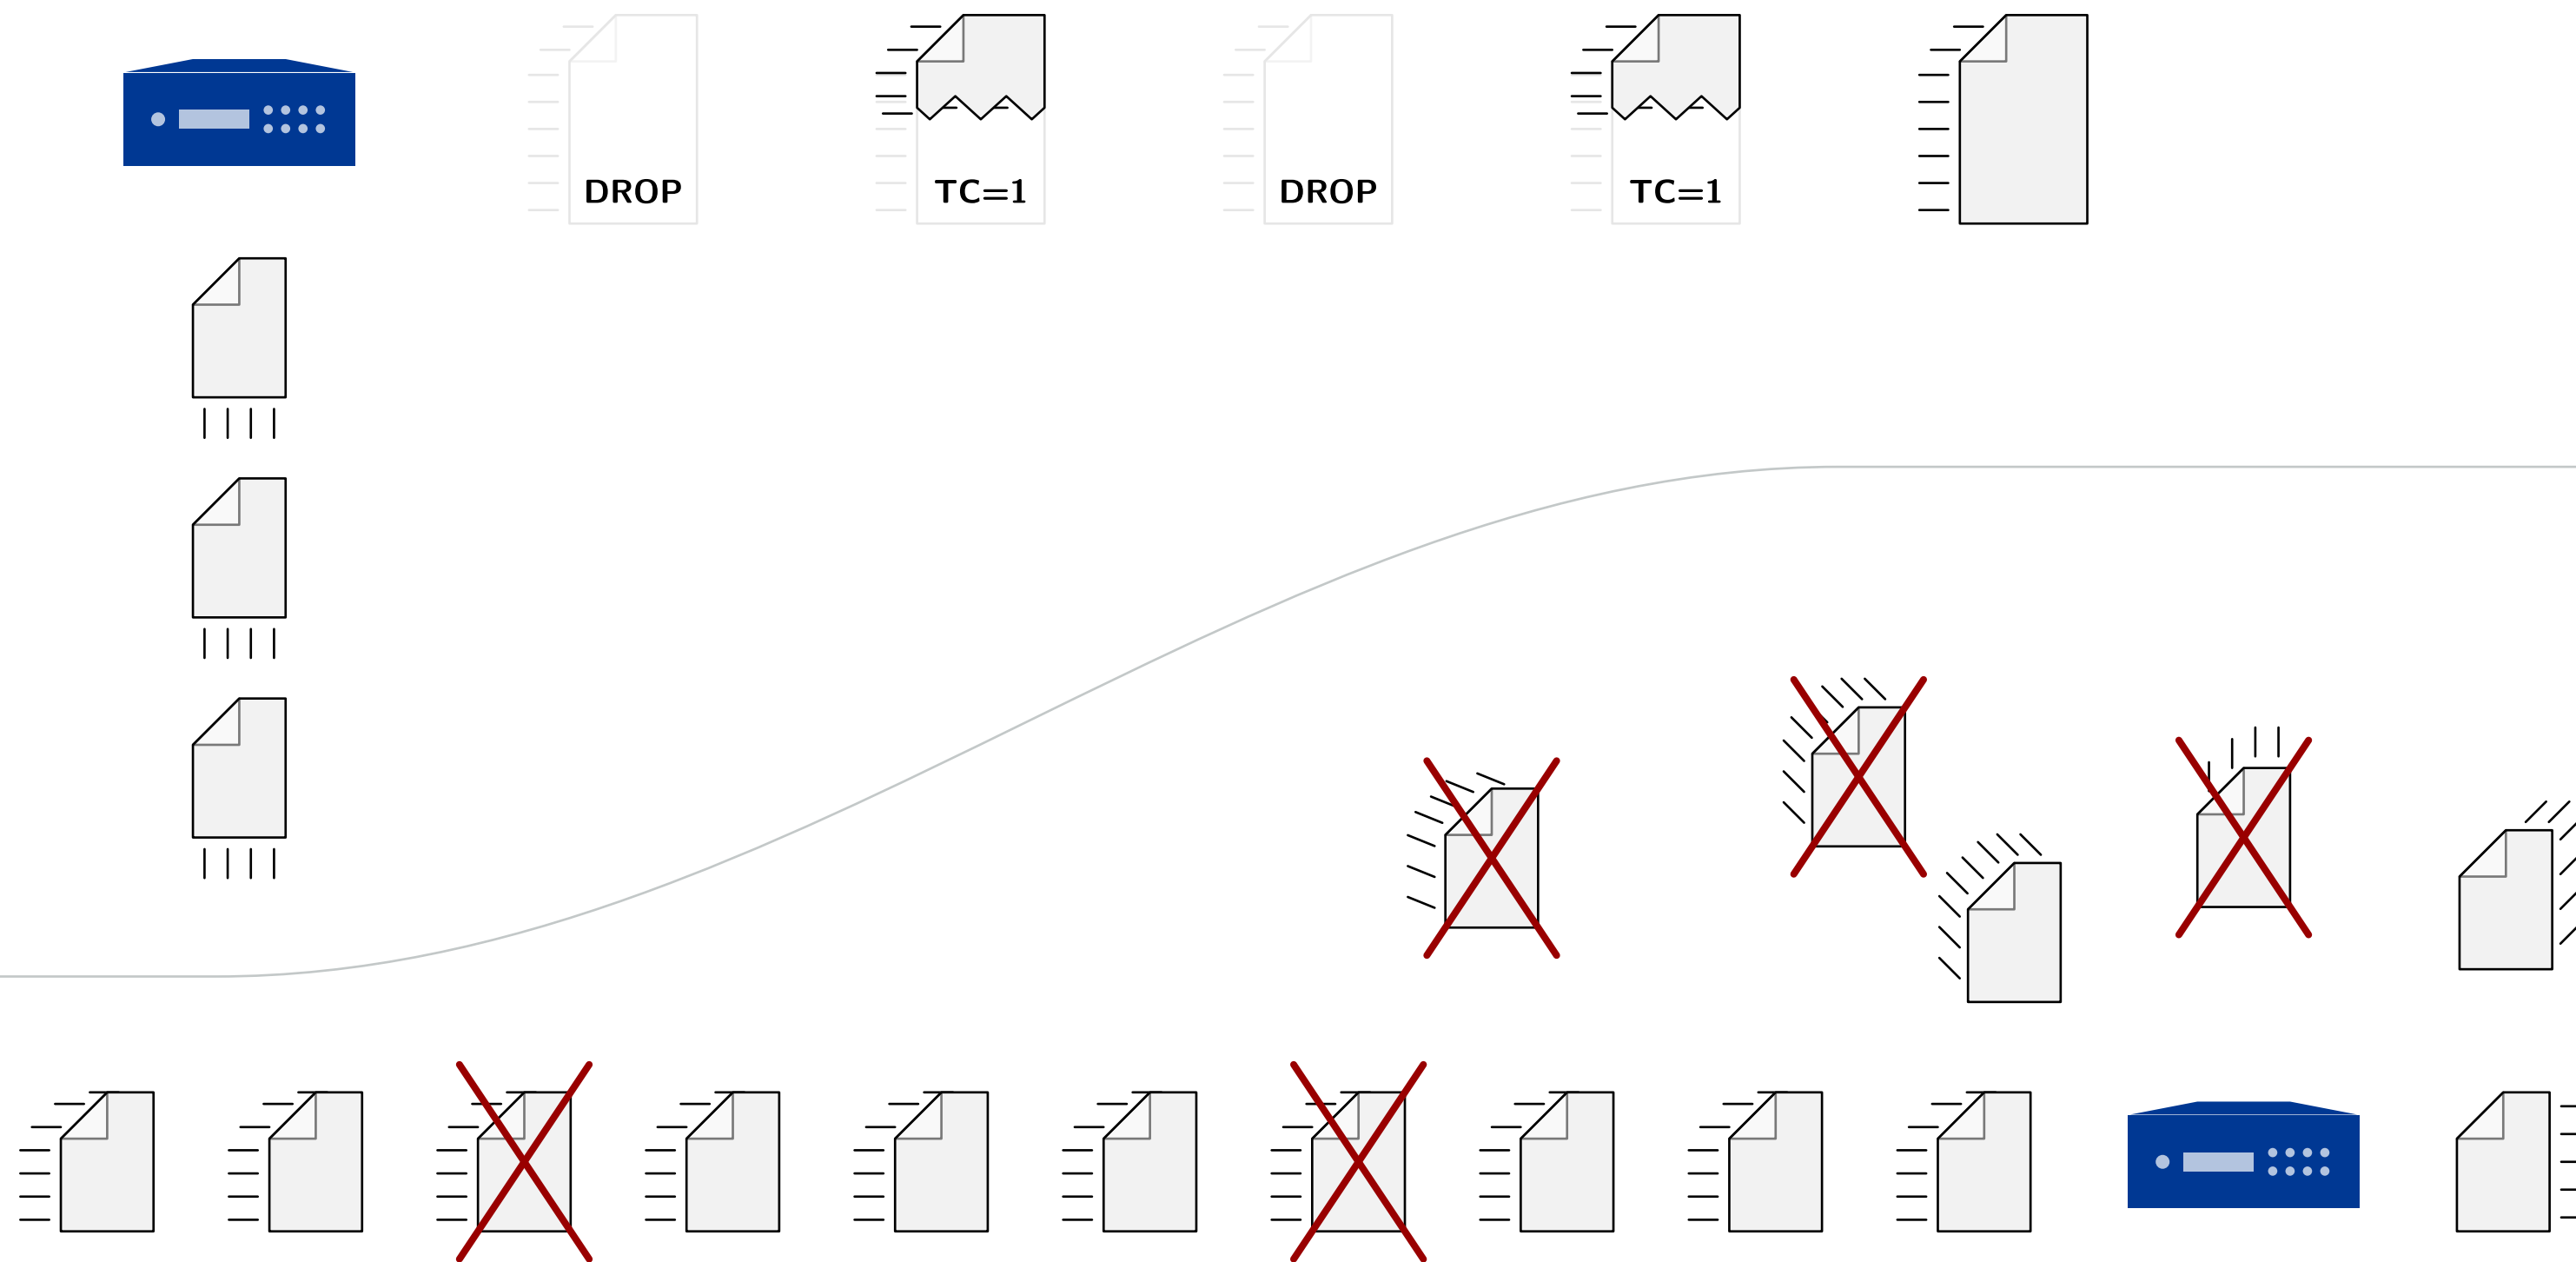
# Overview

- amplification att.
  - rate limiting
- server overload




# Overview

- amplification att.
  - rate limiting
- server overload



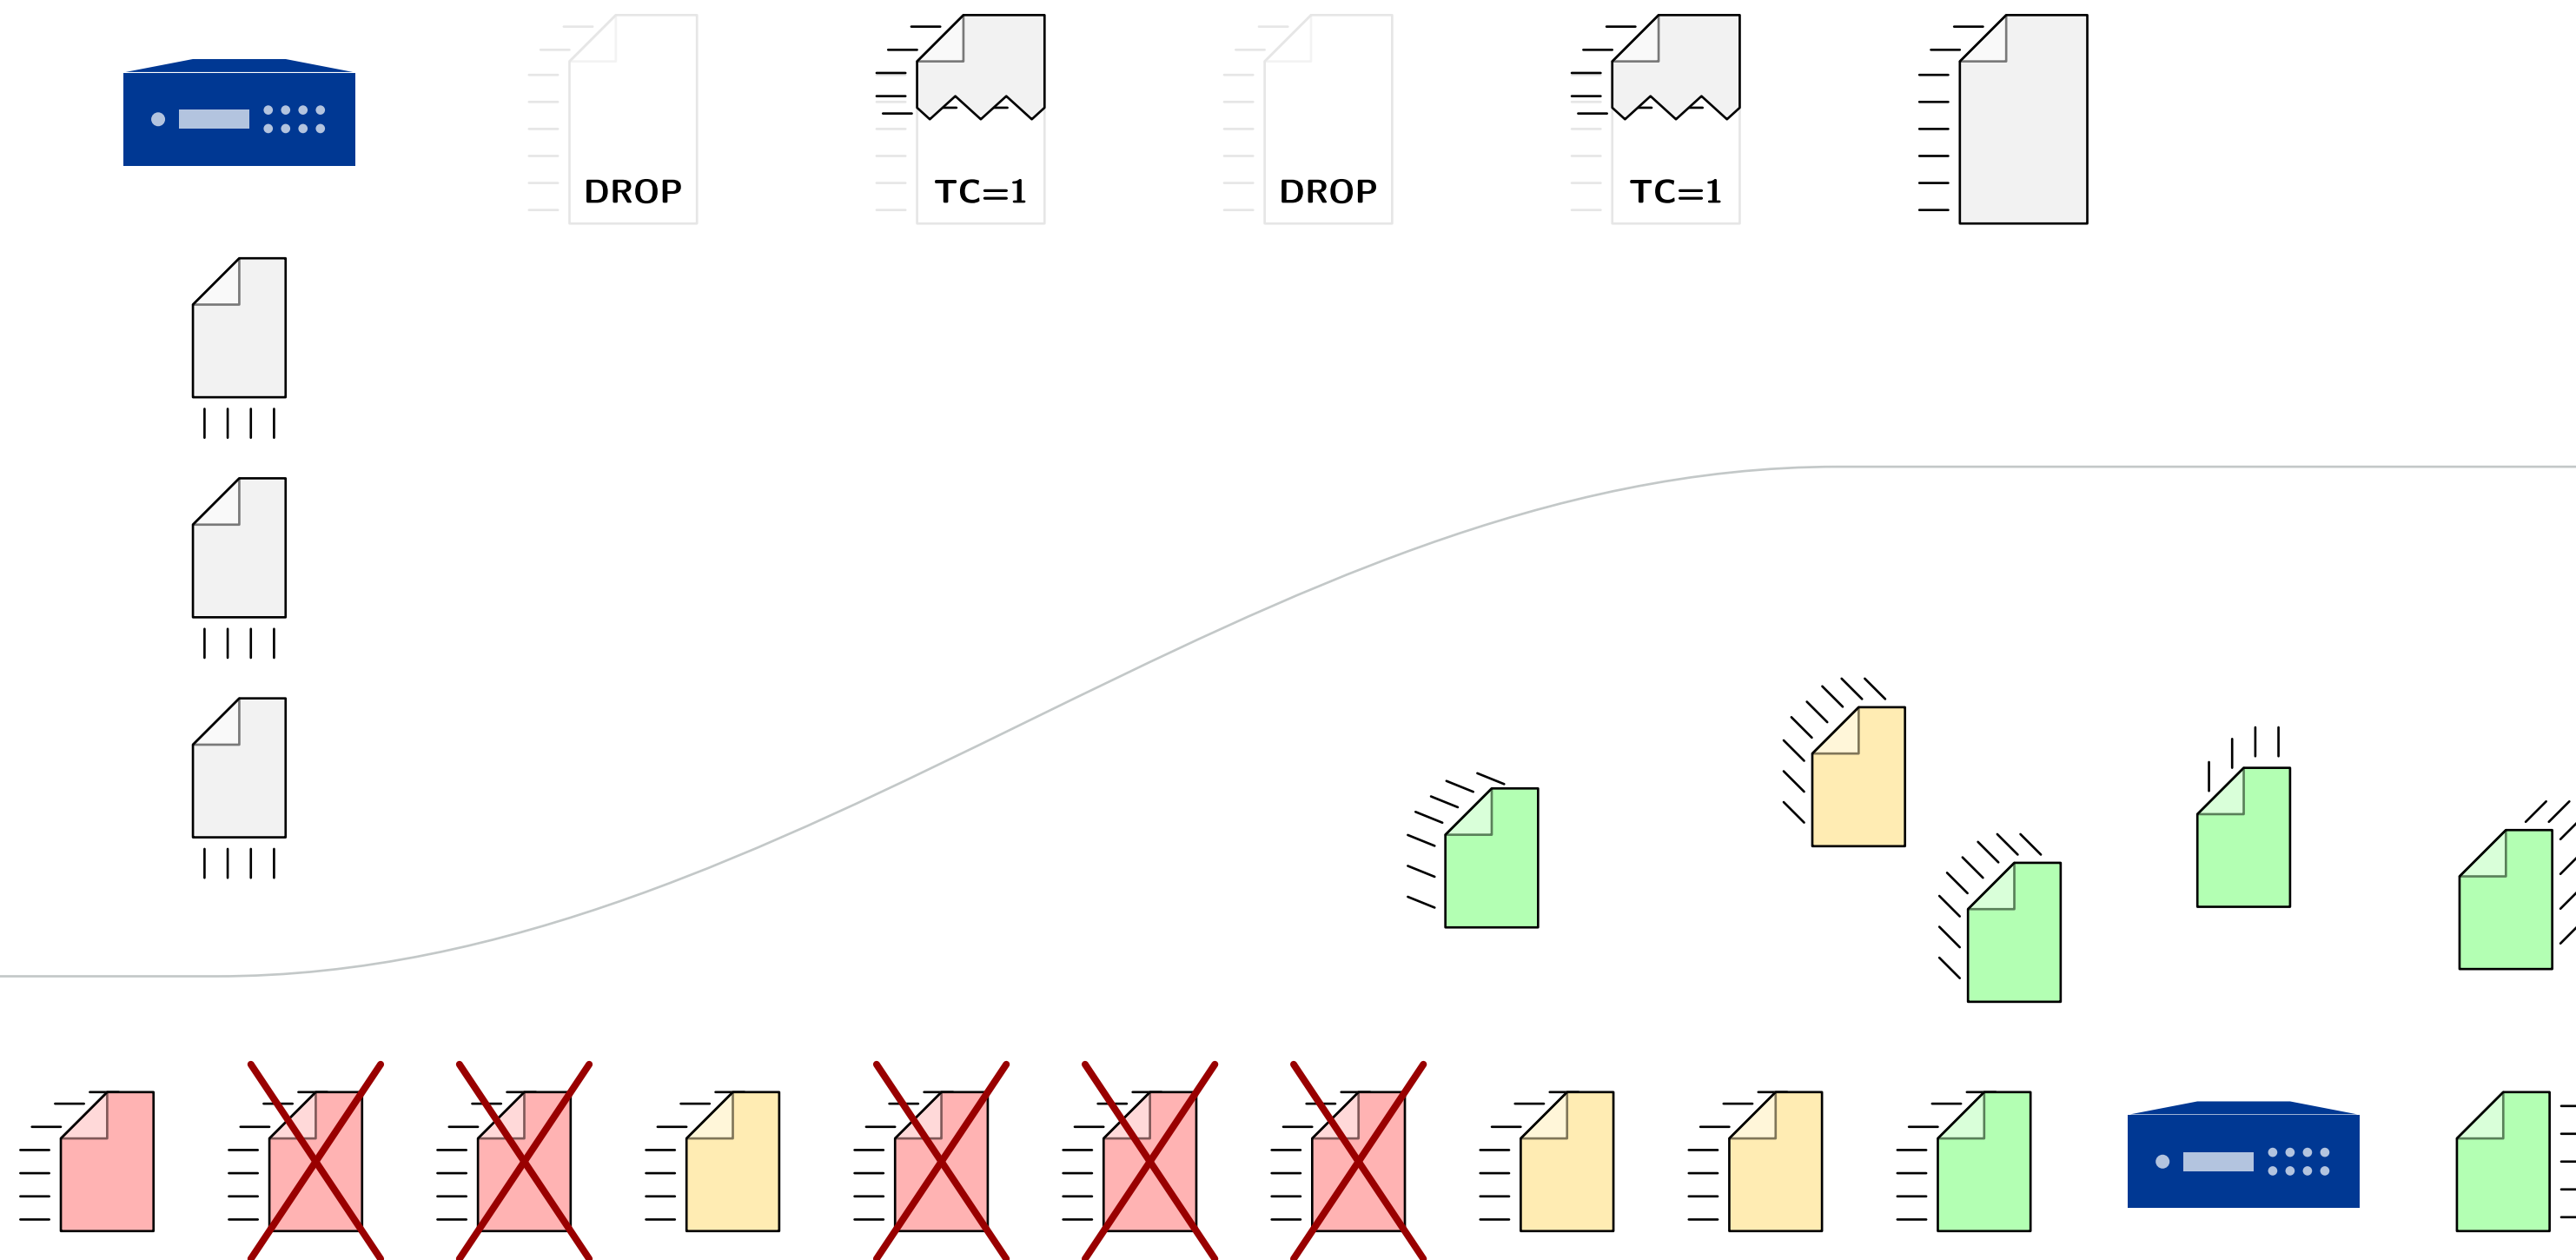
# Overview

- amplification att.
  - rate limiting
- server overload
  - **prioritization**



# Overview

- amplification att.
  - rate limiting
- server overload
  - prioritization



# Limiting individual clients

- counters for addresses
  - instant limit  $L_I$

	:
172.16.96.1	count in $[0, L_I)$
	⋮
2001:db8::734	count in $[0, L_I)$
	⋮




# Limiting individual clients

- counters for addresses
  - instant limit  $L_I$

	:
172.16.96.1	count in $[0, L_I)$
	⋮
2001:db8::734	count in $[0, L_I)$
	⋮

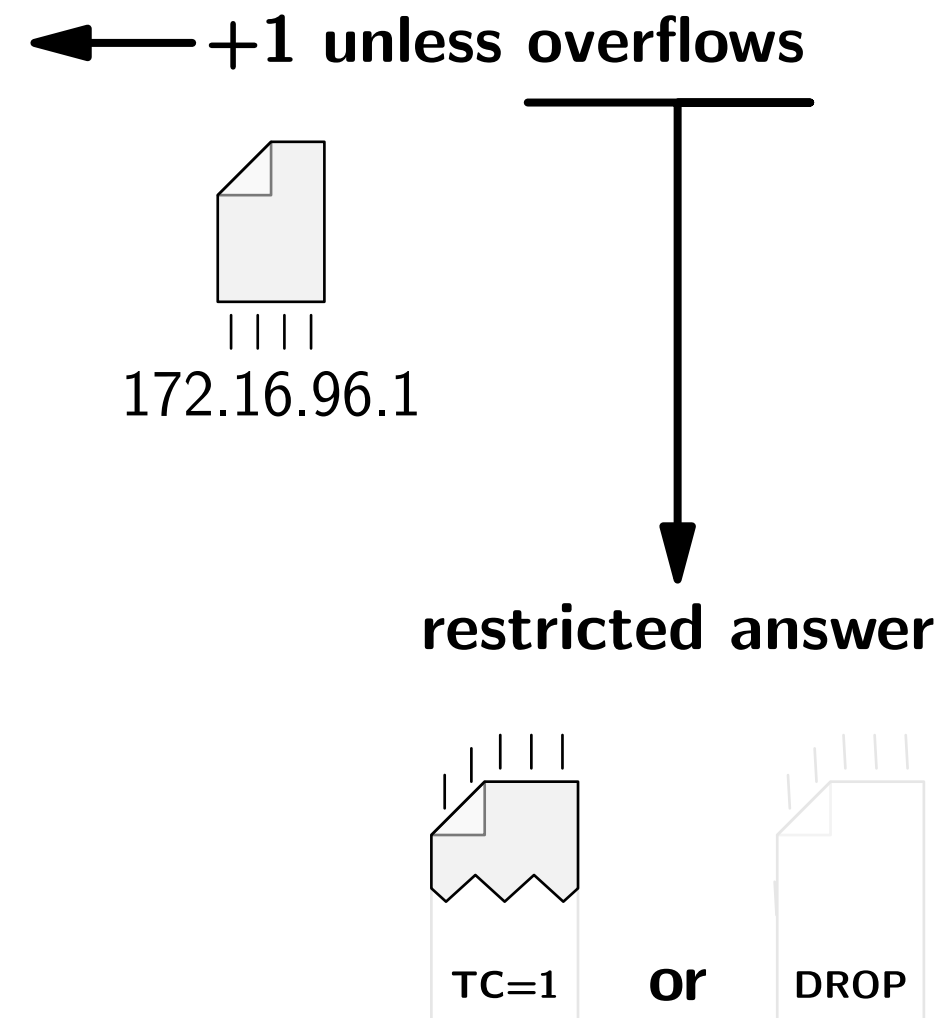
← +1 unless overflows



# Limiting individual clients

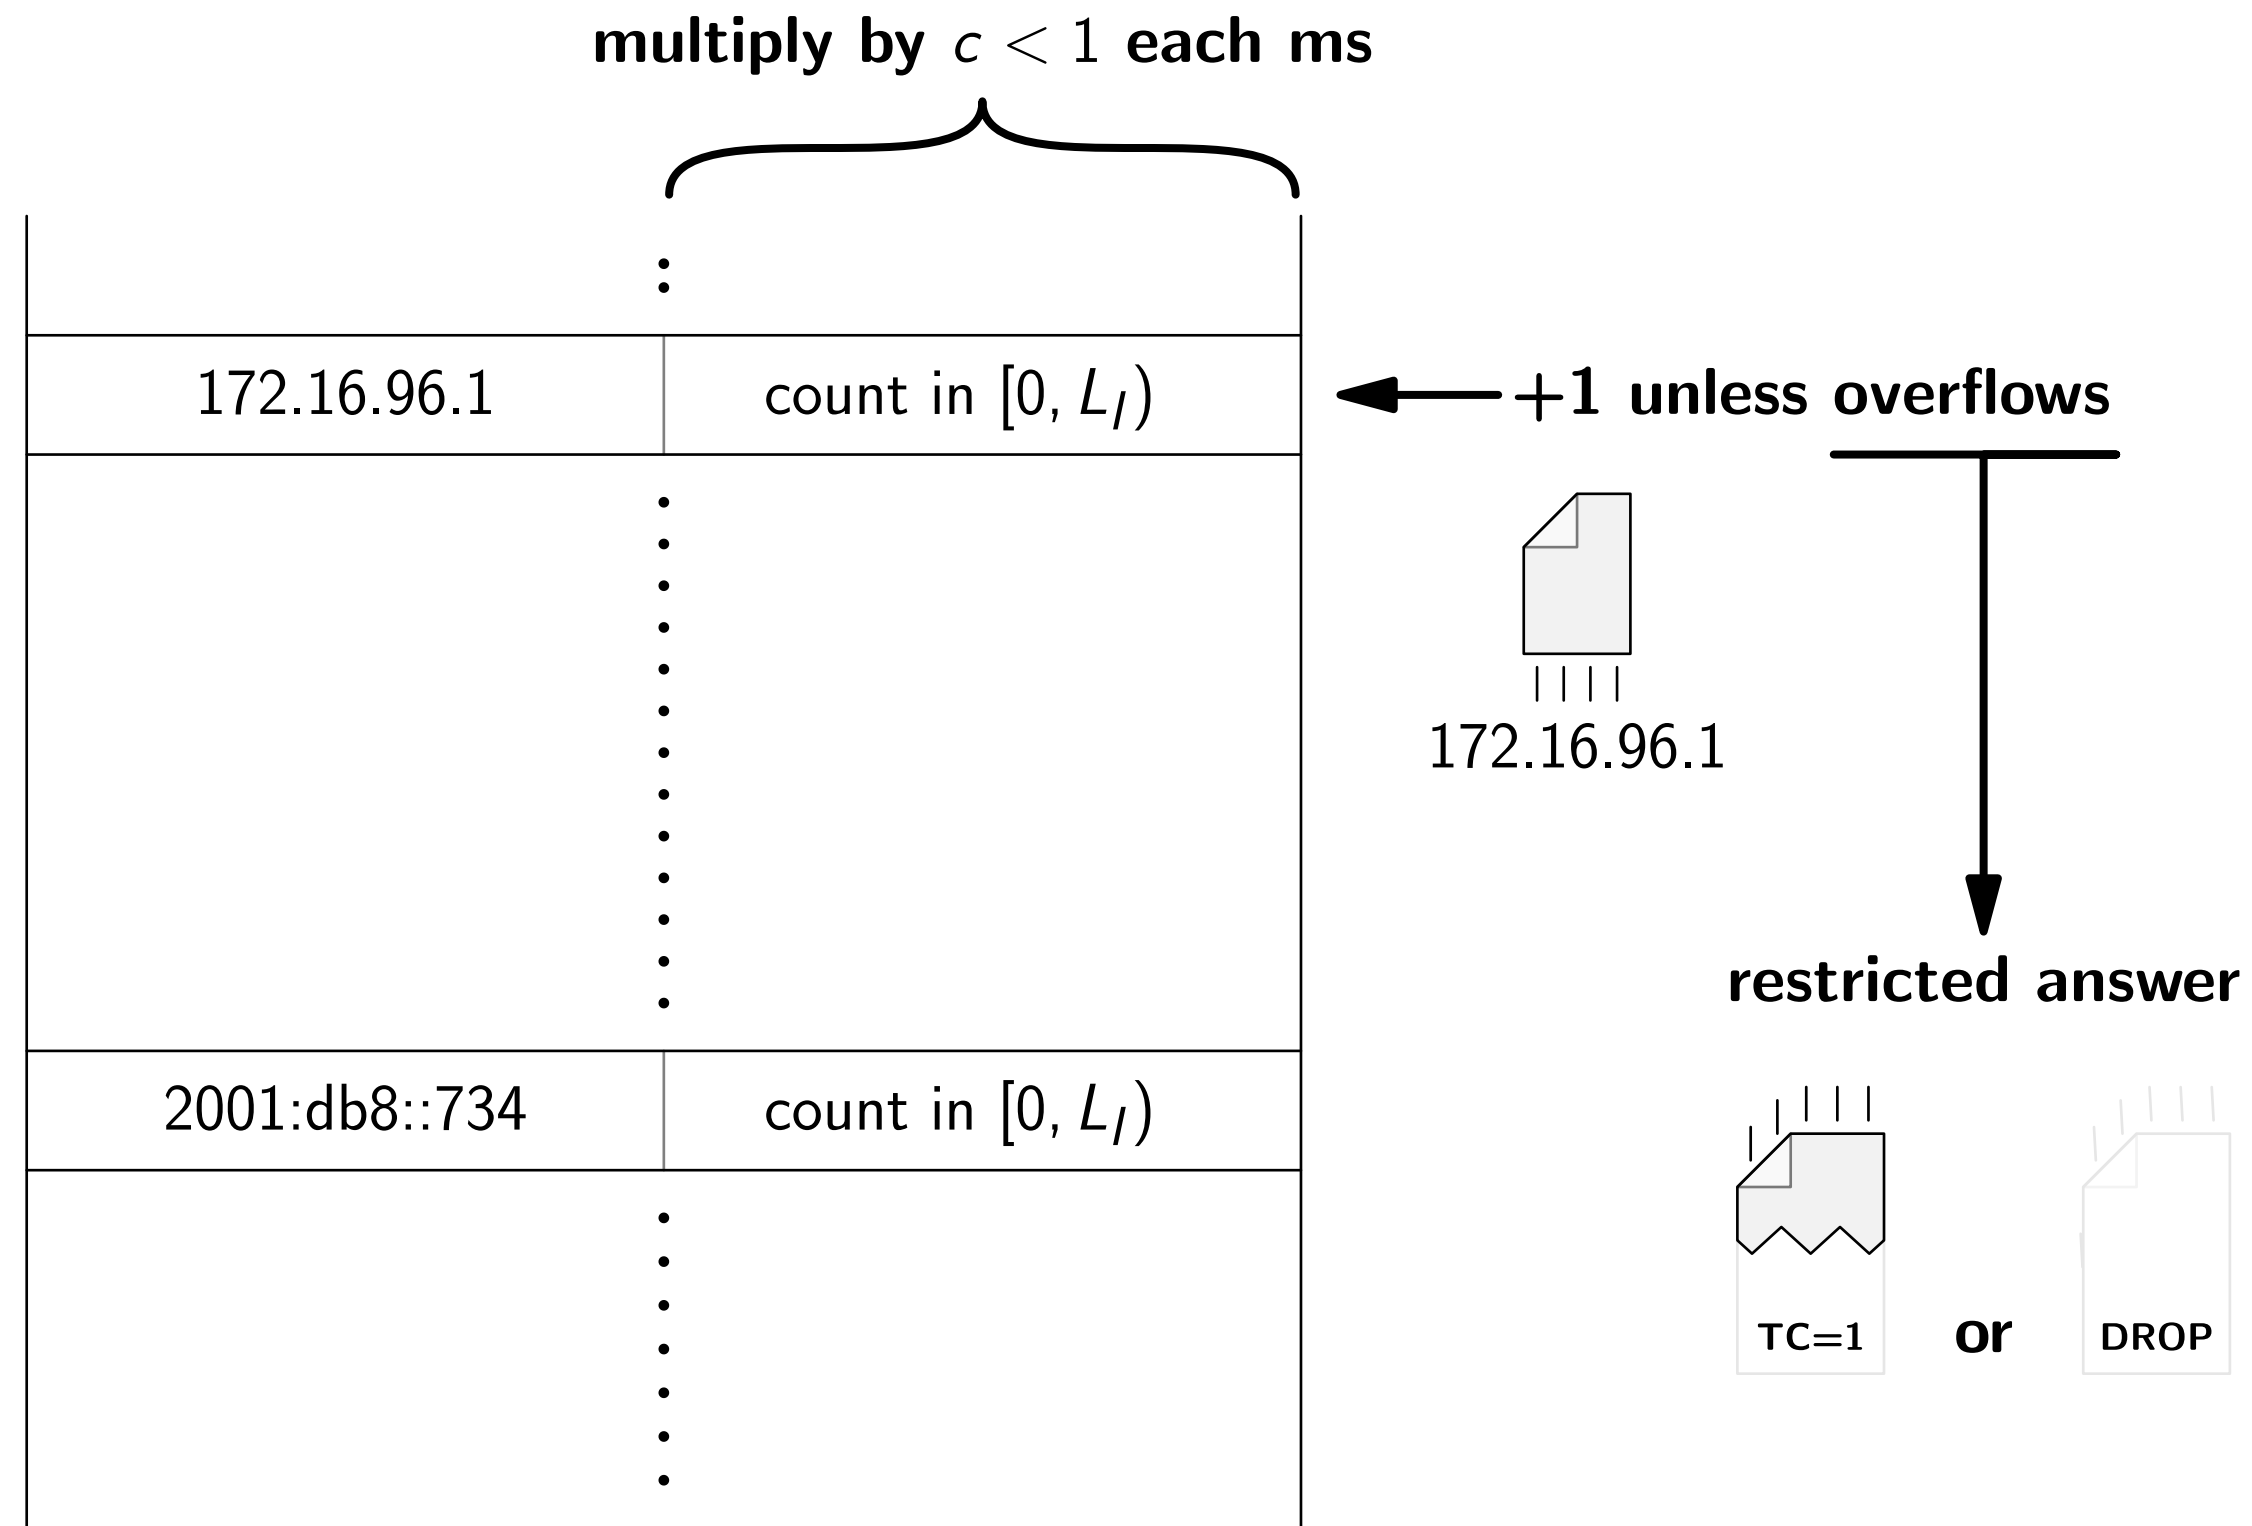
- counters for addresses
  - instant limit  $L_I$

	:
172.16.96.1	count in $[0, L_I)$
	⋮
2001:db8::734	count in $[0, L_I)$
	⋮




# Limiting individual clients

- counters for addresses
  - instant limit  $L_I$
- **exponential decay**
  - **rate limit**



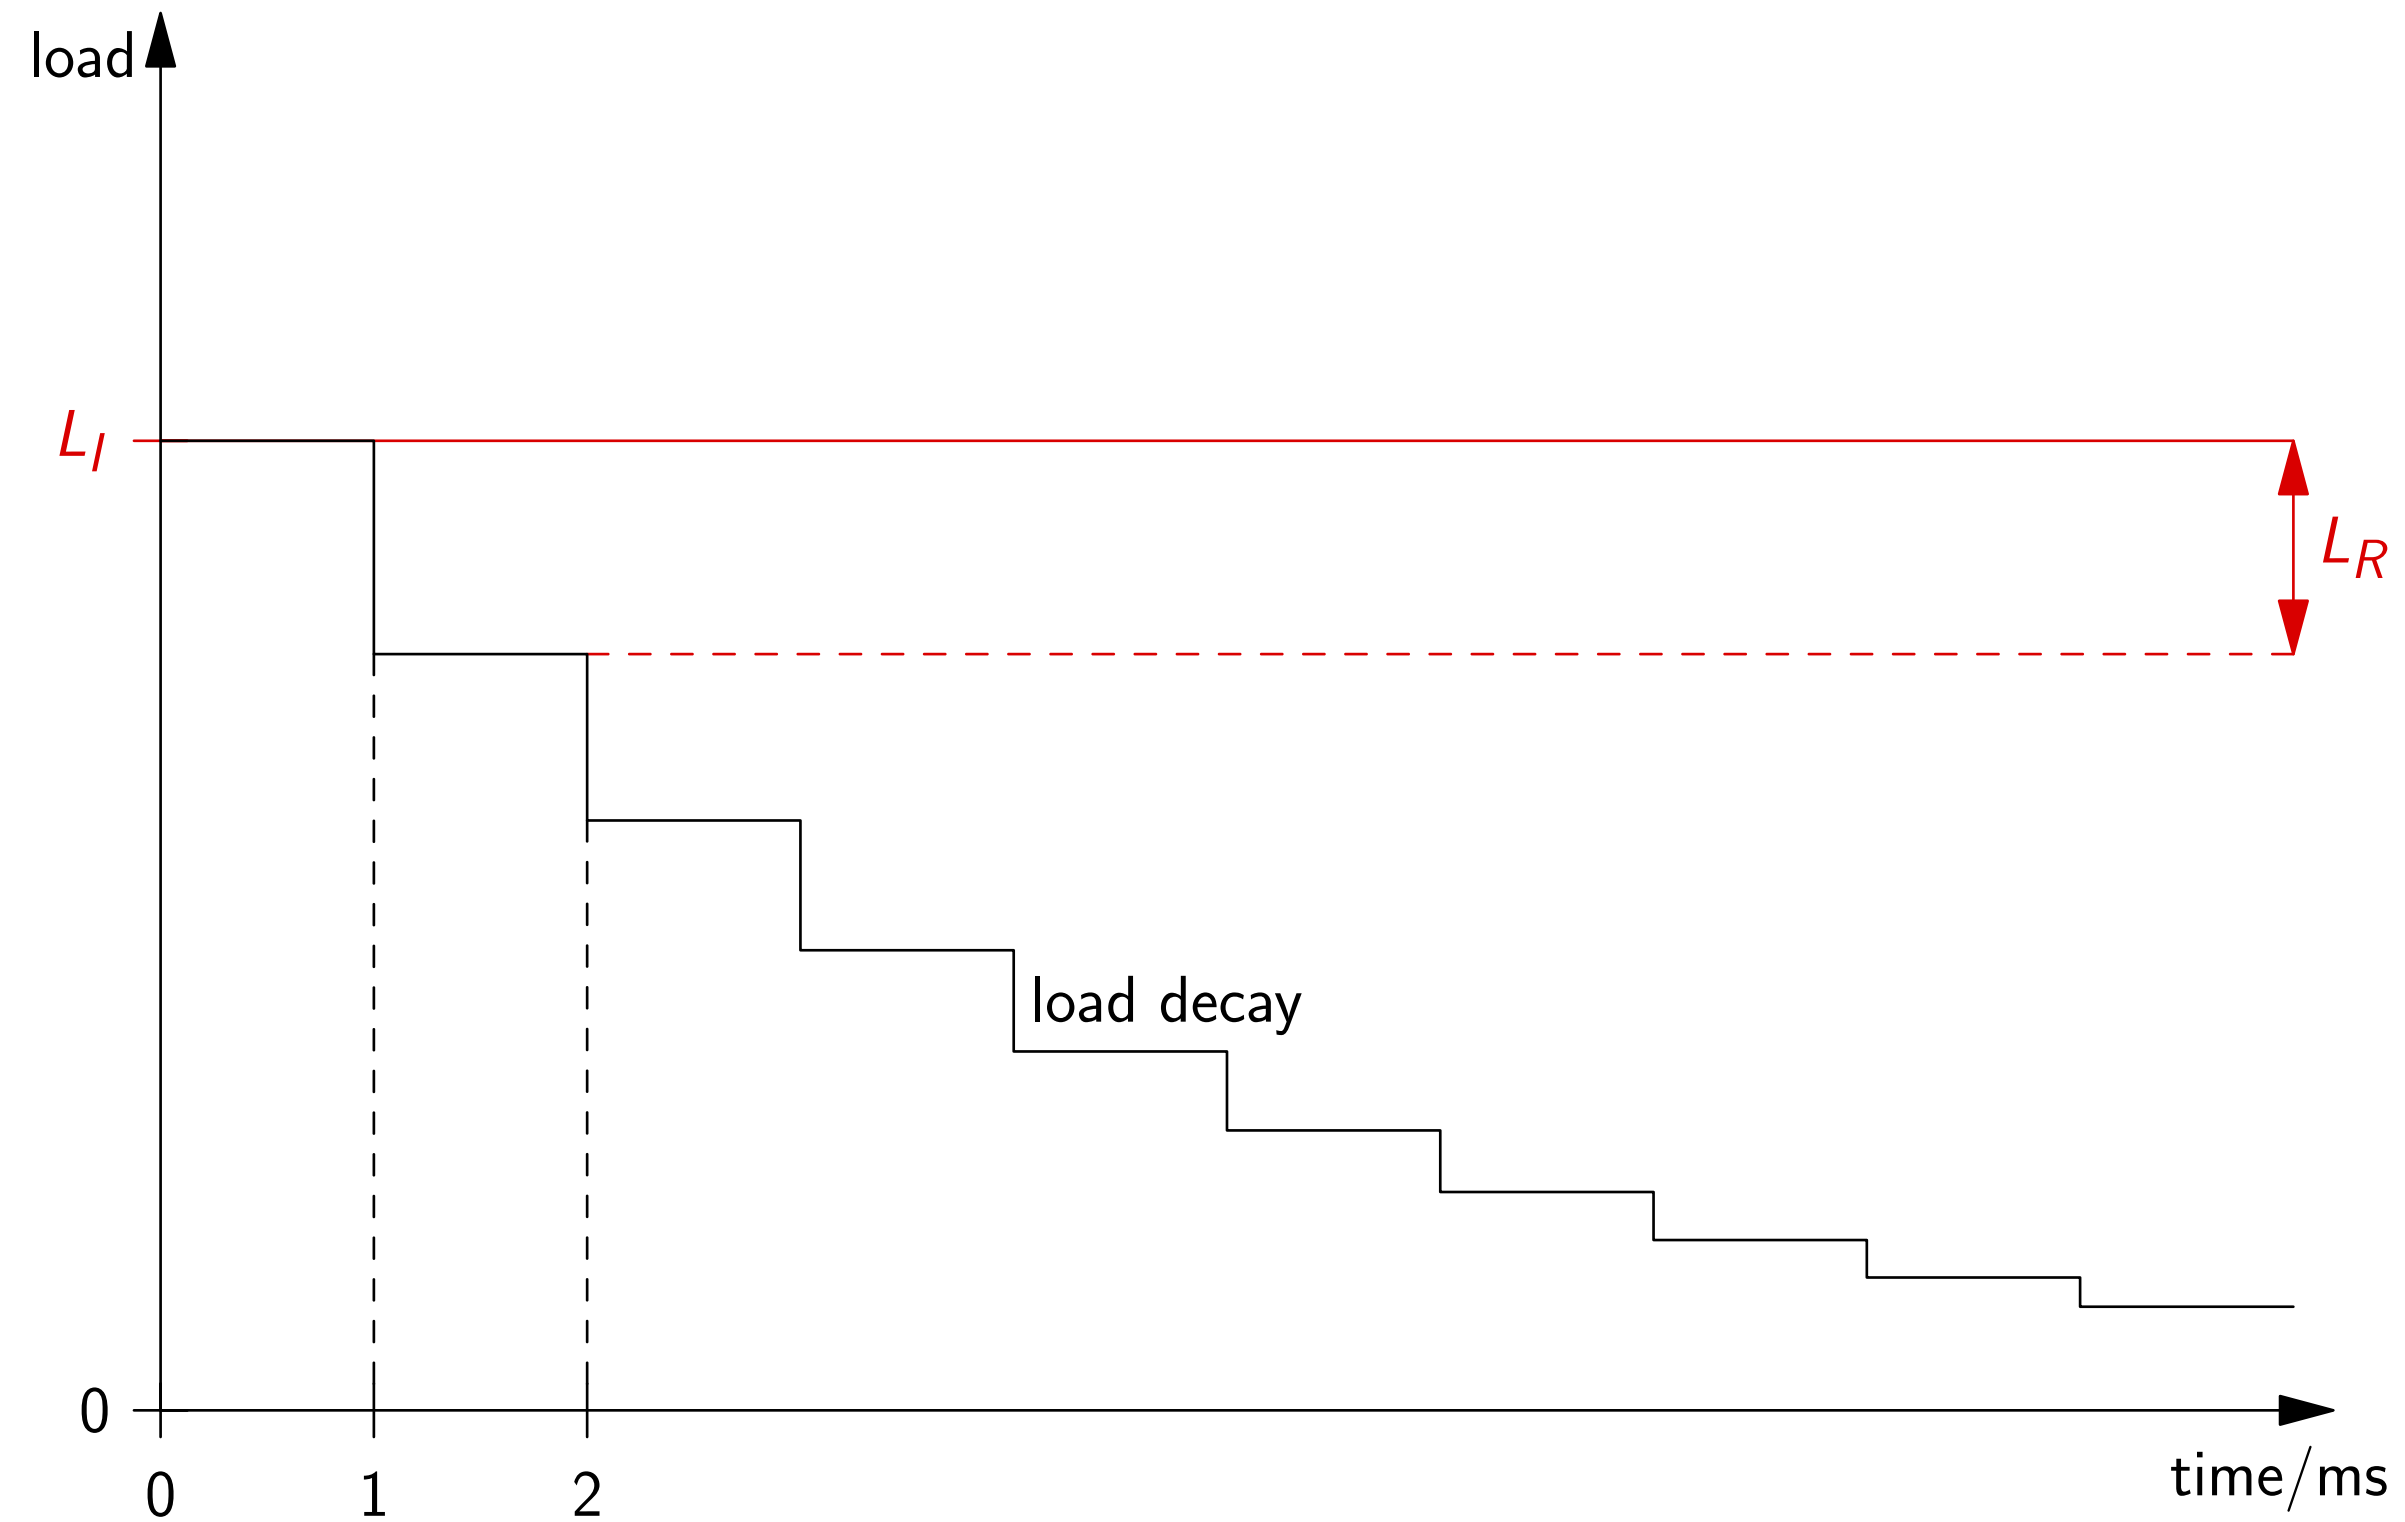
# Exponential decay

- instant limit  $L_I$



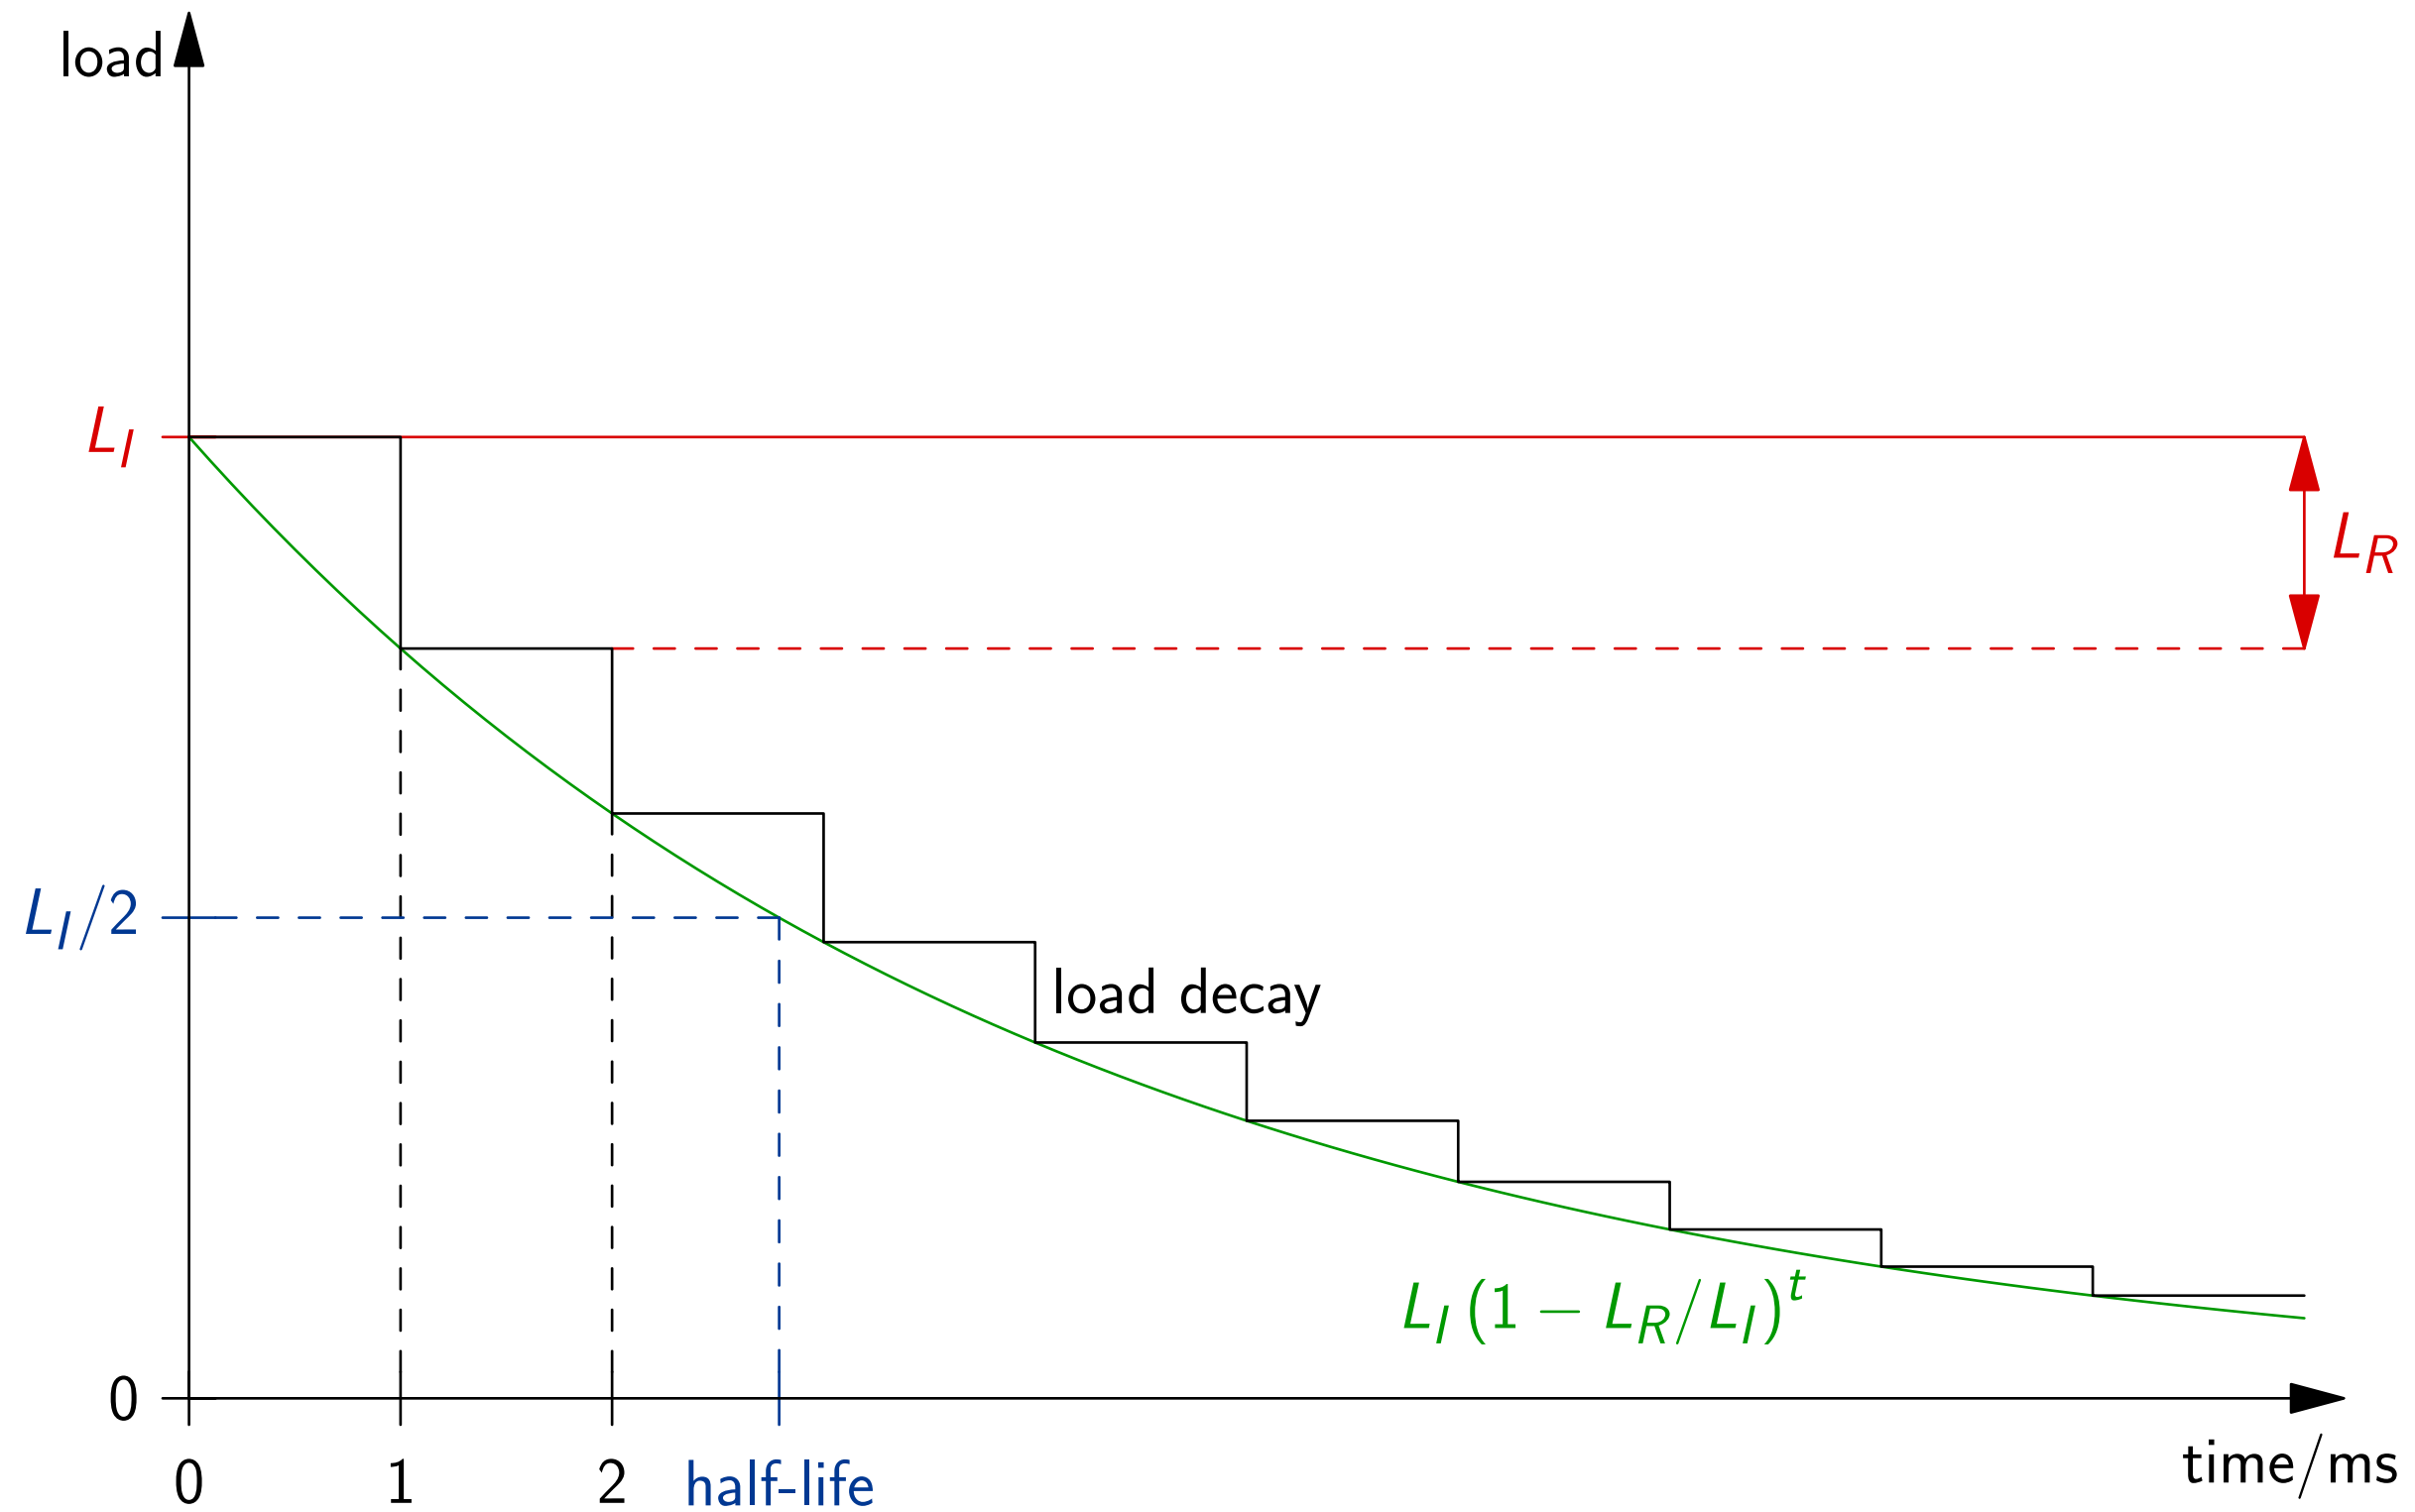
# Exponential decay

- instant limit  $L_I$
- rate limit  $L_R$ 
  - per ms




# Exponential decay

- instant limit  $L_I$
- rate limit  $L_R$ 
  - per ms
- **half-life**




# Exponential decay

- instant limit  $L_I$
- rate limit  $L_R$ 
  - per ms
- half-life



# Exponential decay

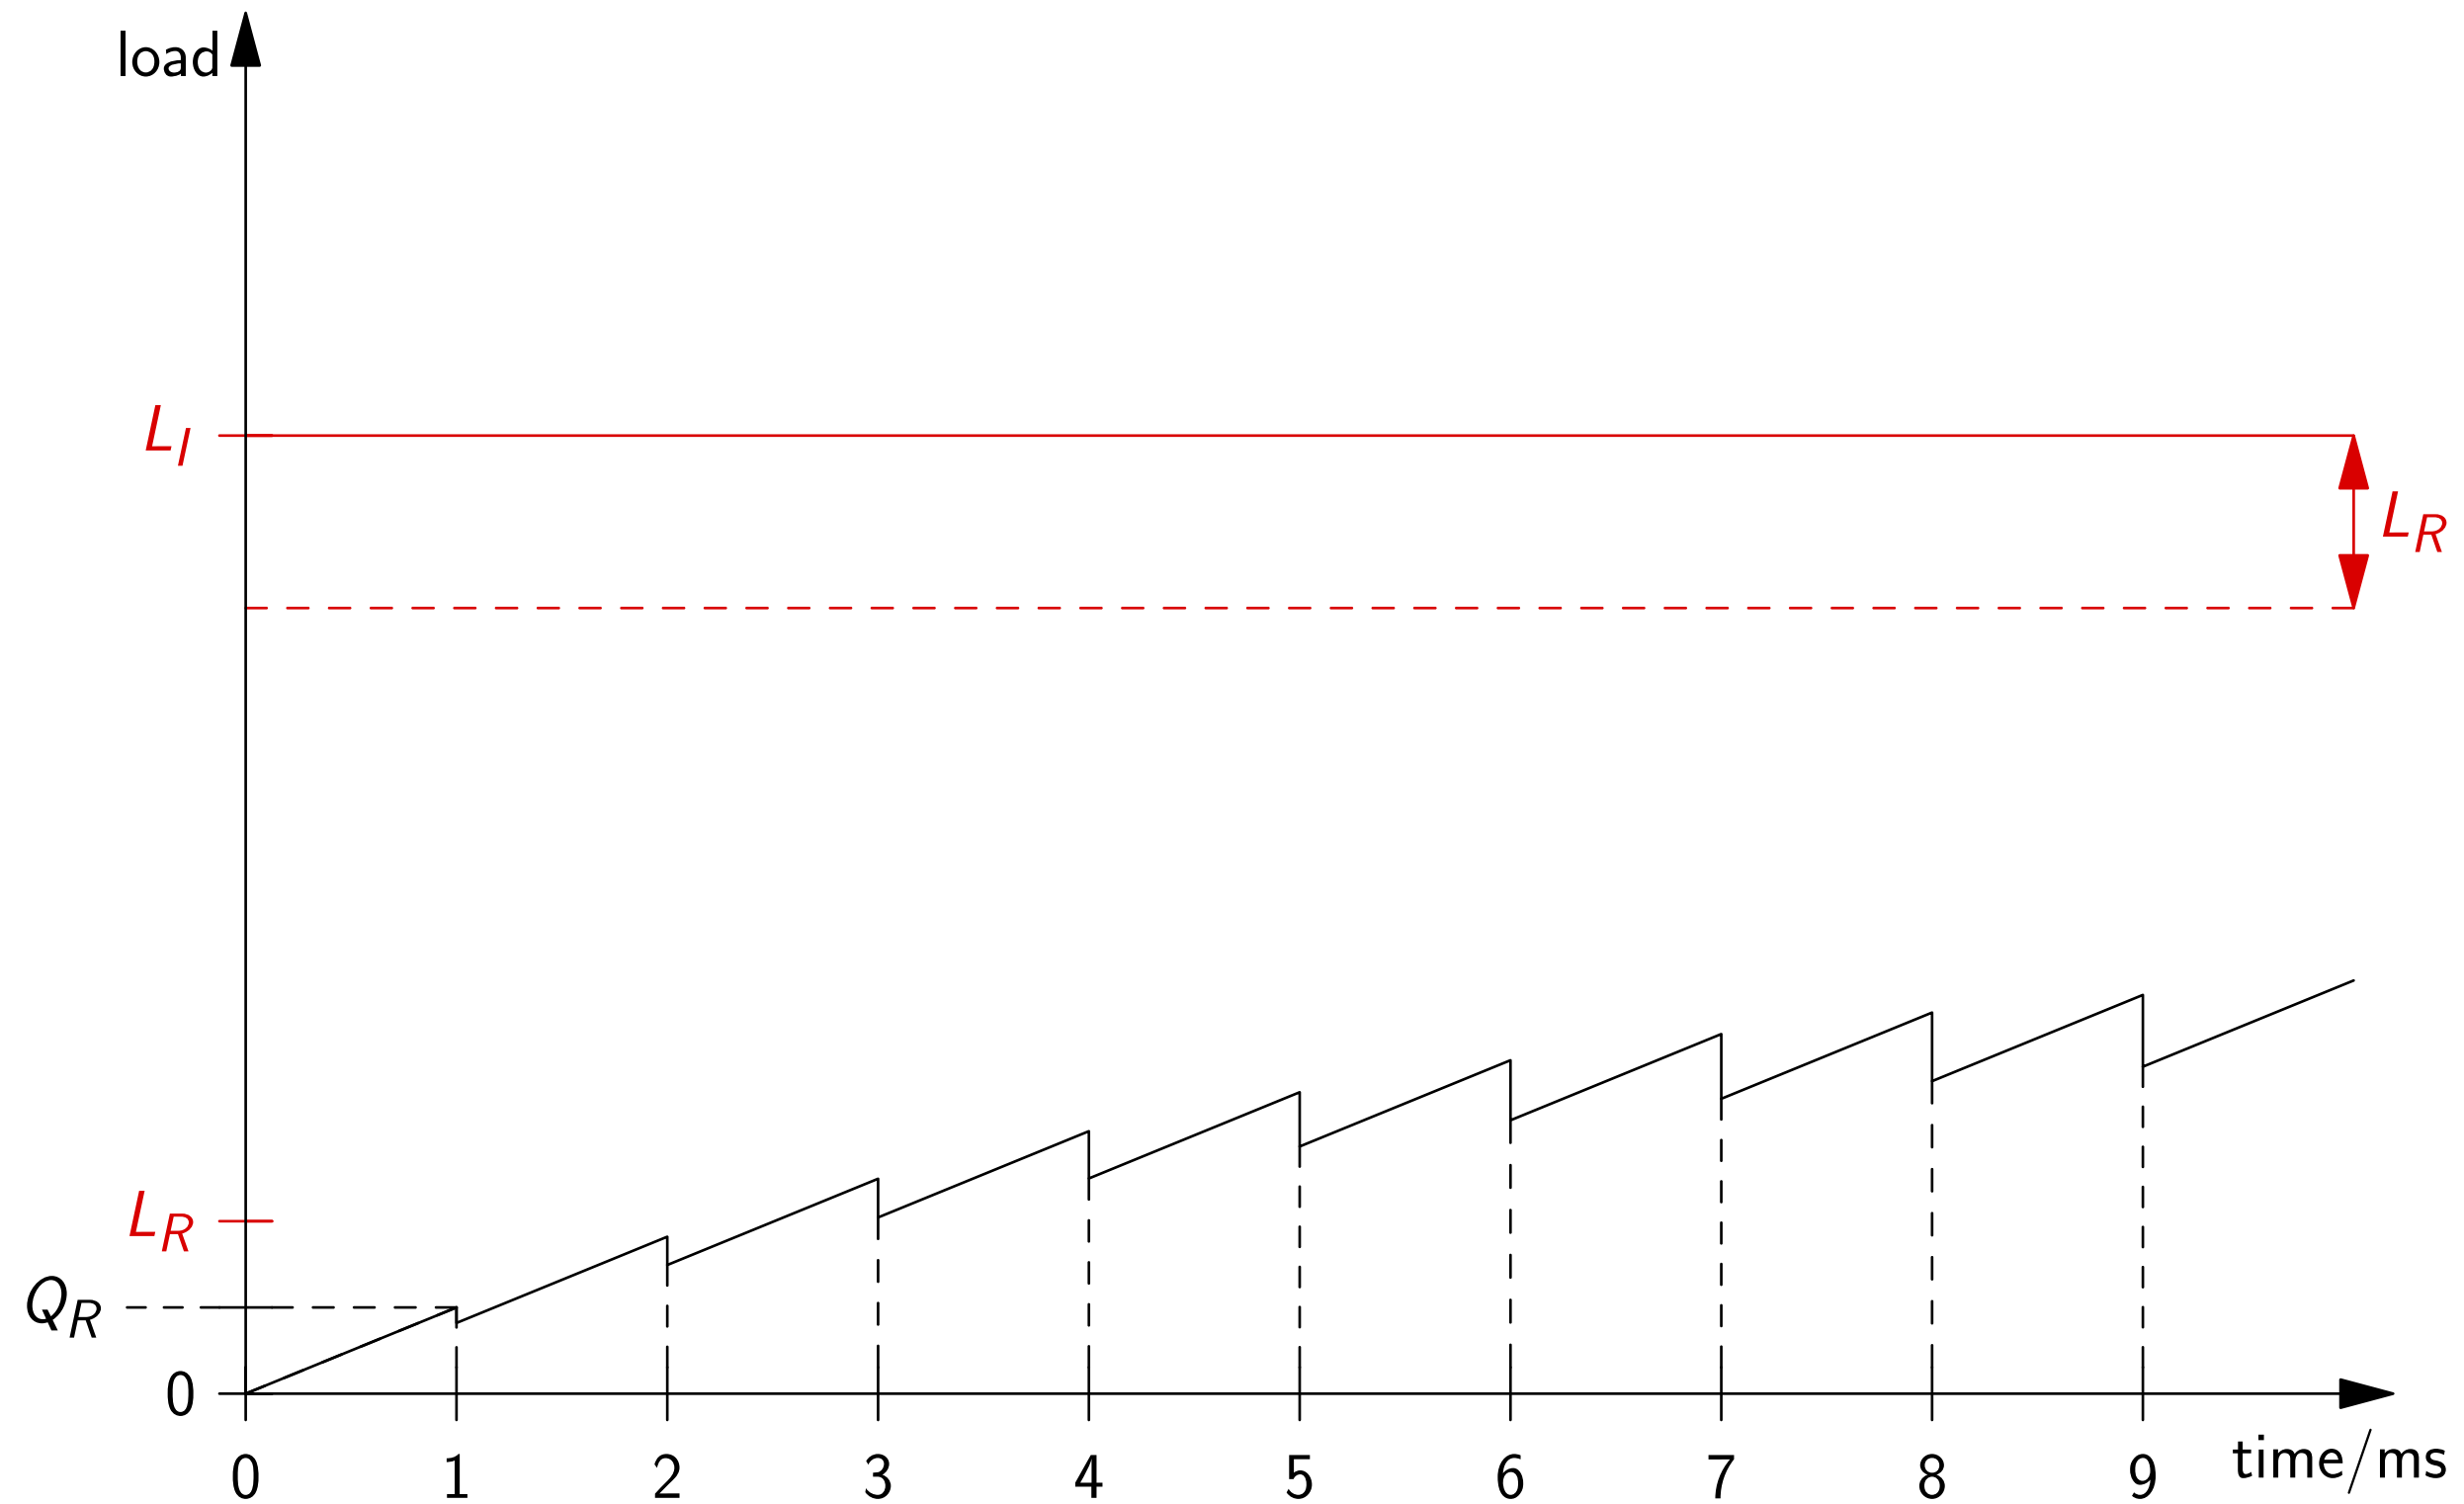
- instant limit  $L_I$
- rate limit  $L_R$ 
  - per ms
- half-life






# Constant query rate example

- instant limit  $L_I$
- rate limit  $L_R$ 
  - per ms
- query rate  $Q_R$ 
  - per ms




# Constant query rate example

- instant limit  $L_I$
- rate limit  $L_R$ 
  - per ms
- query rate  $Q_R$ 
  - per ms



# Constant query rate example

- instant limit  $L_I$
- rate limit  $L_R$ 
  - per ms
- query rate  $Q_R$ 
  - per ms



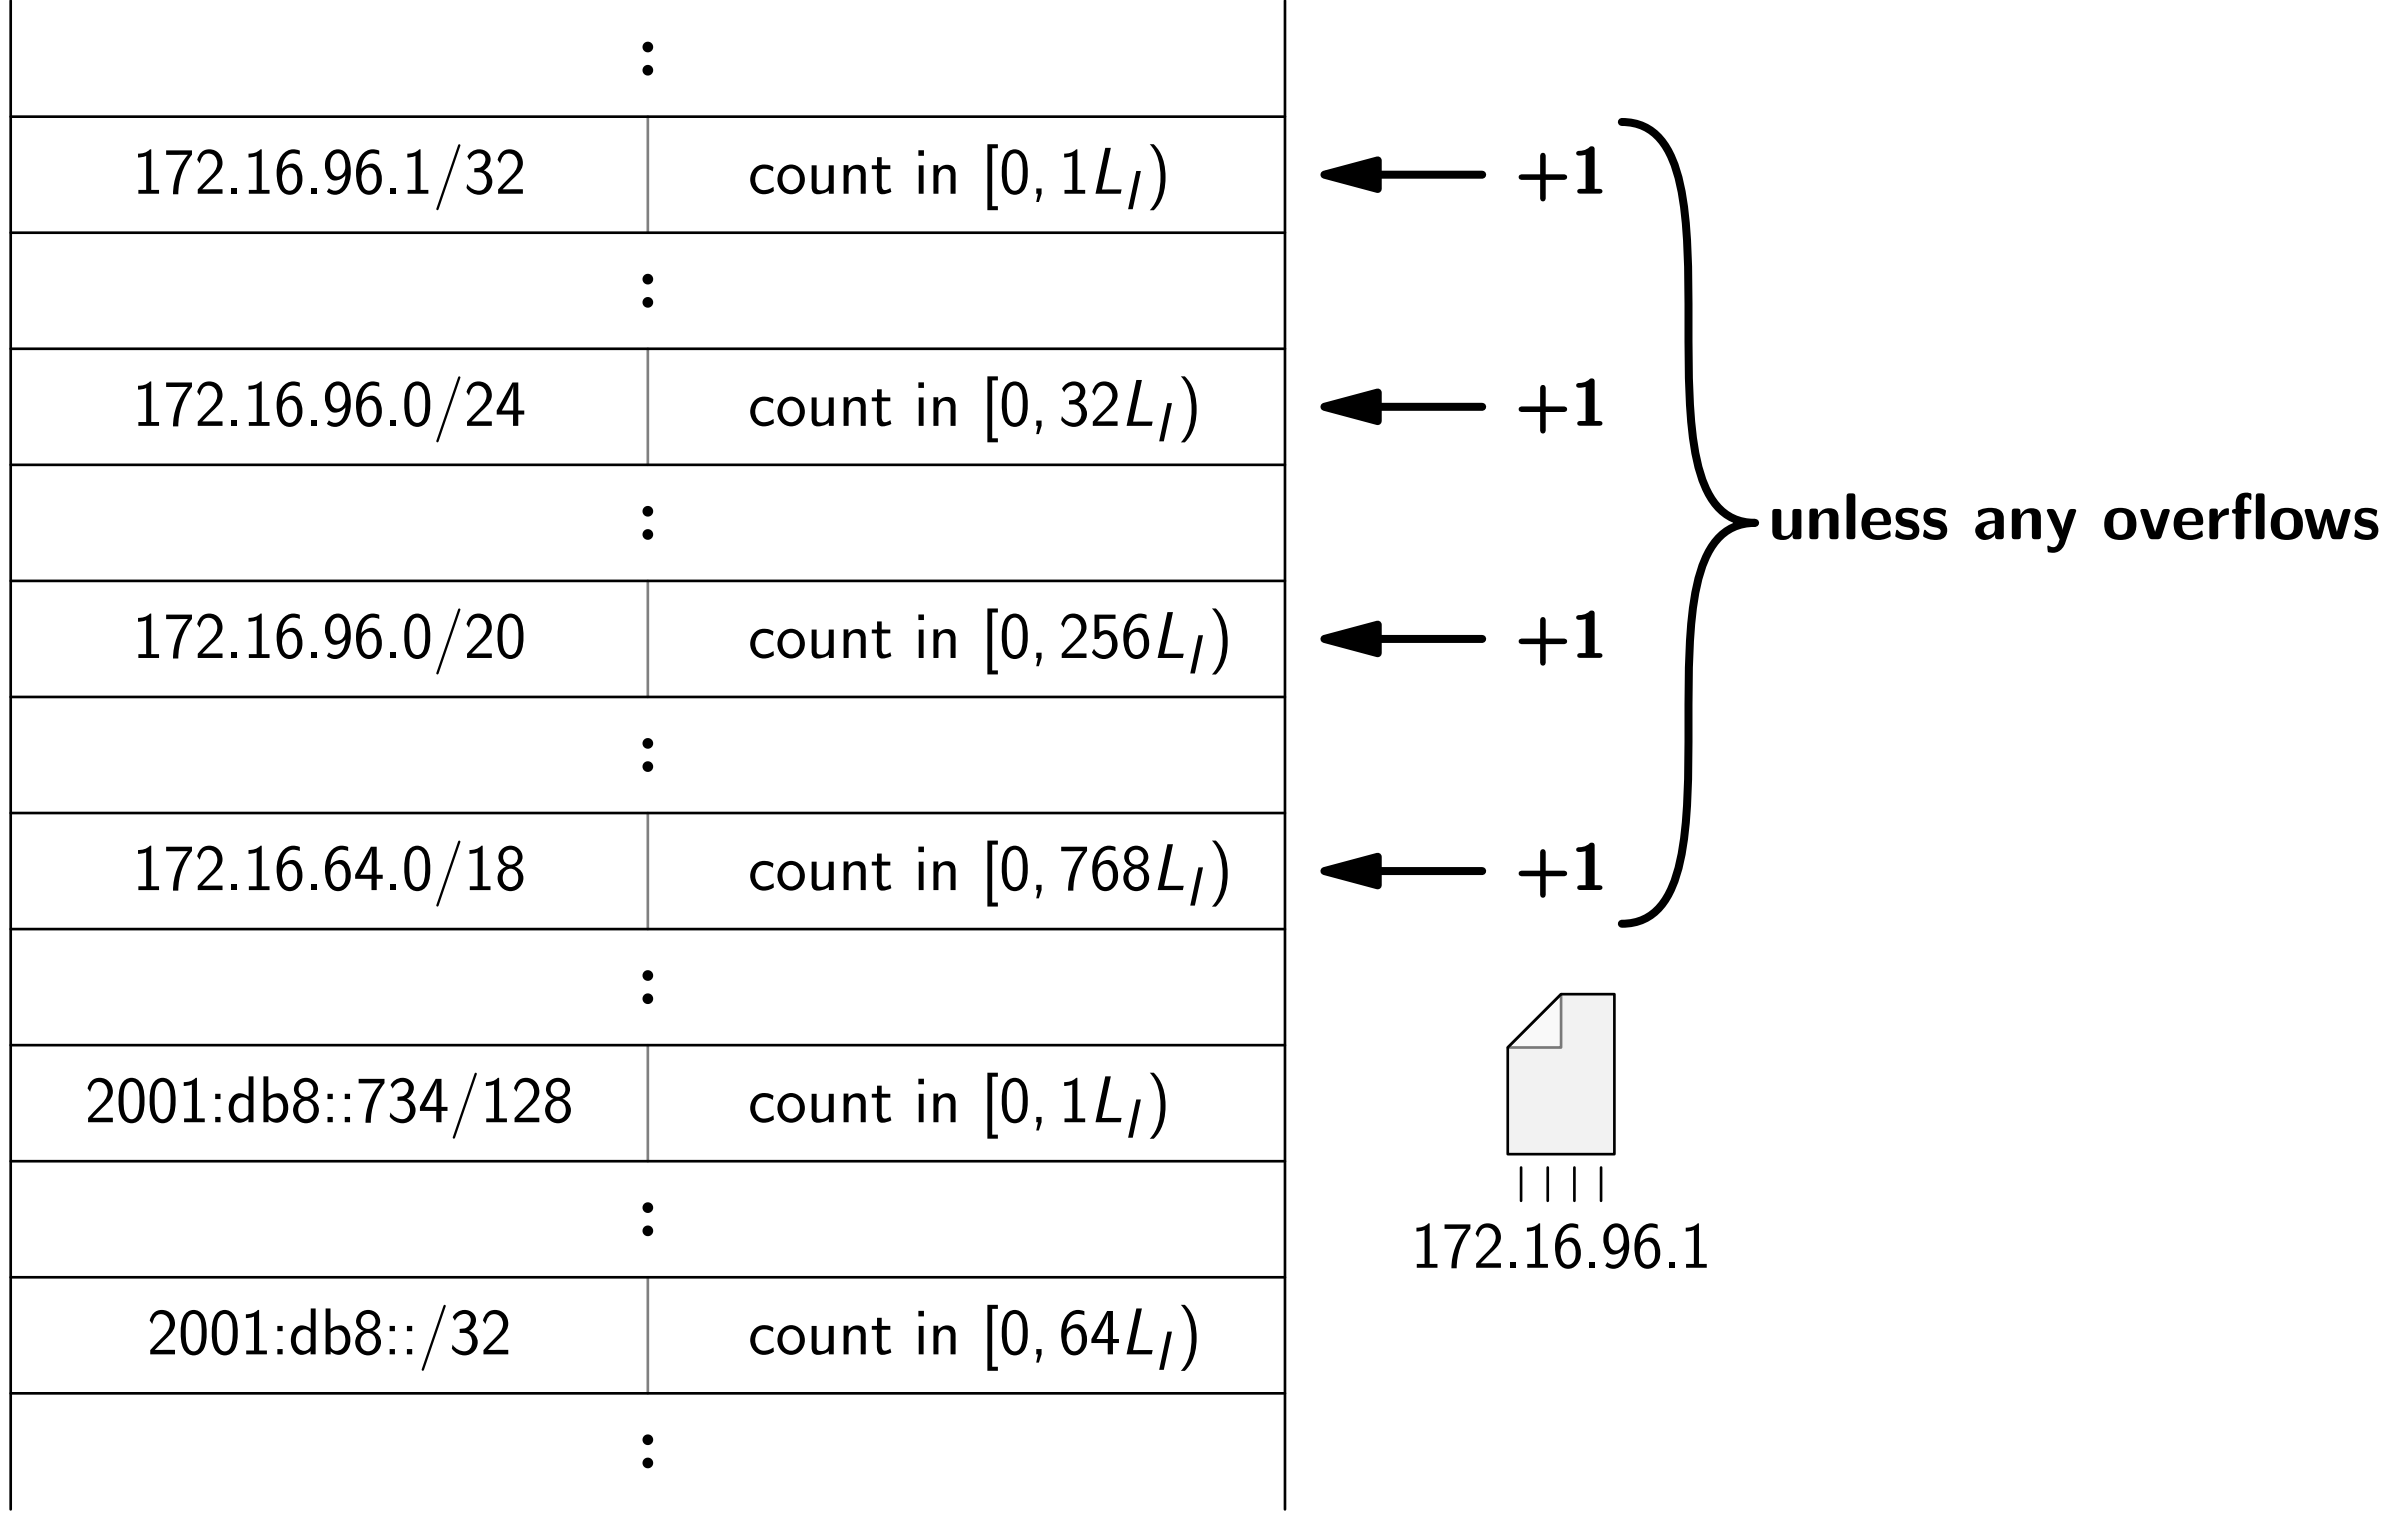
# Limiting networks

- IPv4
  - /32: 1
  - /24: 32
  - /20: 256
  - /18: 768
- IPv6
  - /128: 1
  - /64: 2
  - /56: 3
  - /48: 4
  - /32: 64

:	
172.16.96.1/32	count in $[0, 1L_I)$
:	
172.16.96.0/24	count in $[0, 32L_I)$
:	
172.16.96.0/20	count in $[0, 256L_I)$
:	
172.16.64.0/18	count in $[0, 768L_I)$
:	
2001:db8::734/128	count in $[0, 1L_I)$
:	
2001:db8::/32	count in $[0, 64L_I)$
:	

# Limiting networks

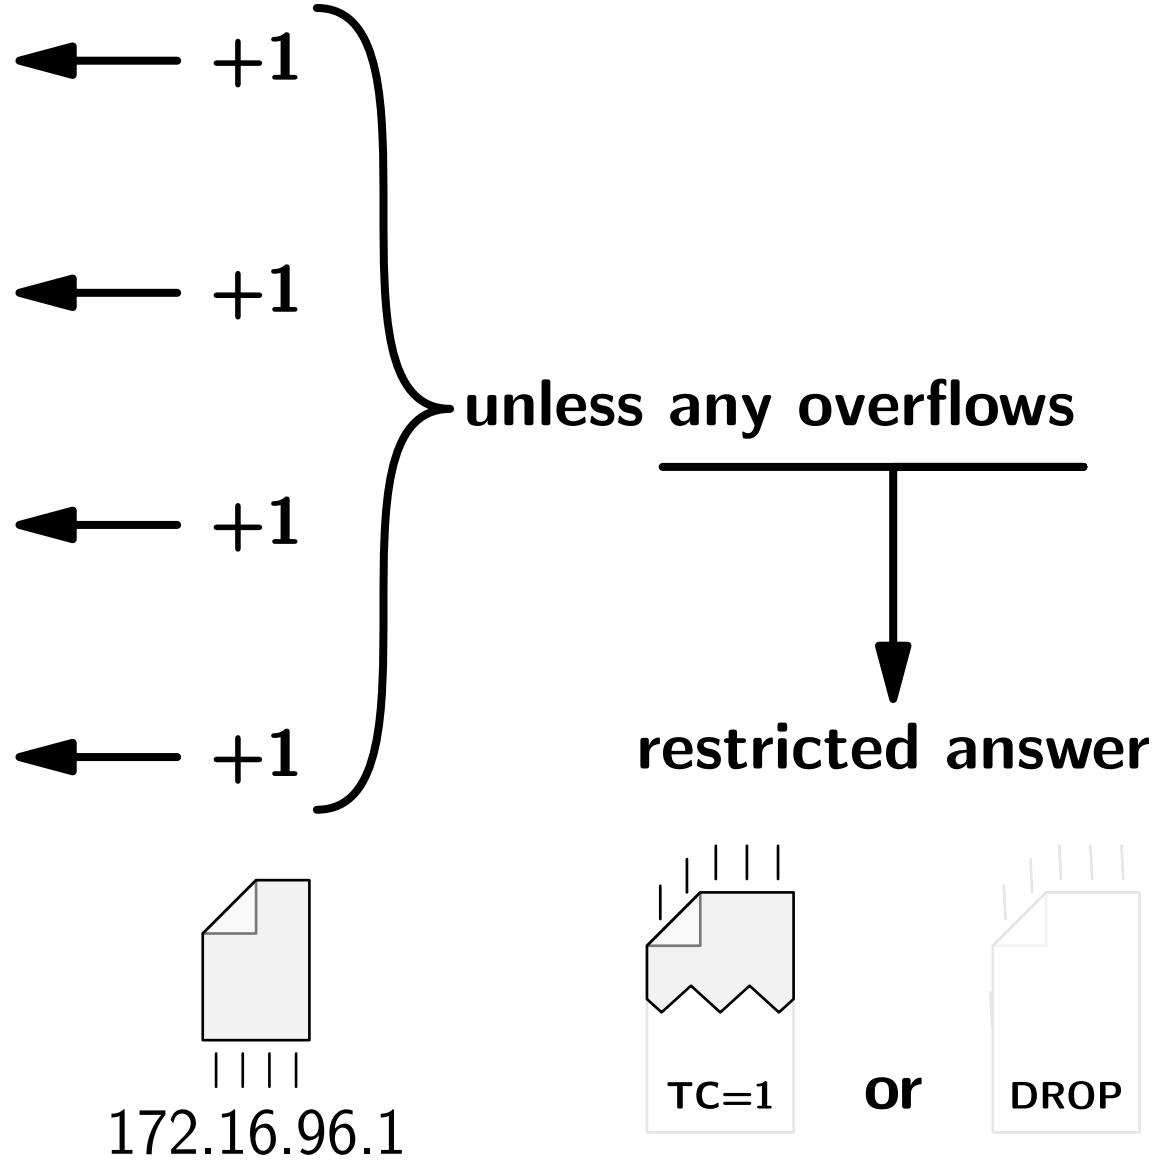
- IPv4
  - /32: 1
  - /24: 32
  - /20: 256
  - /18: 768
- IPv6
  - /128: 1
  - /64: 2
  - /56: 3
  - /48: 4
  - /32: 64



# Limiting networks

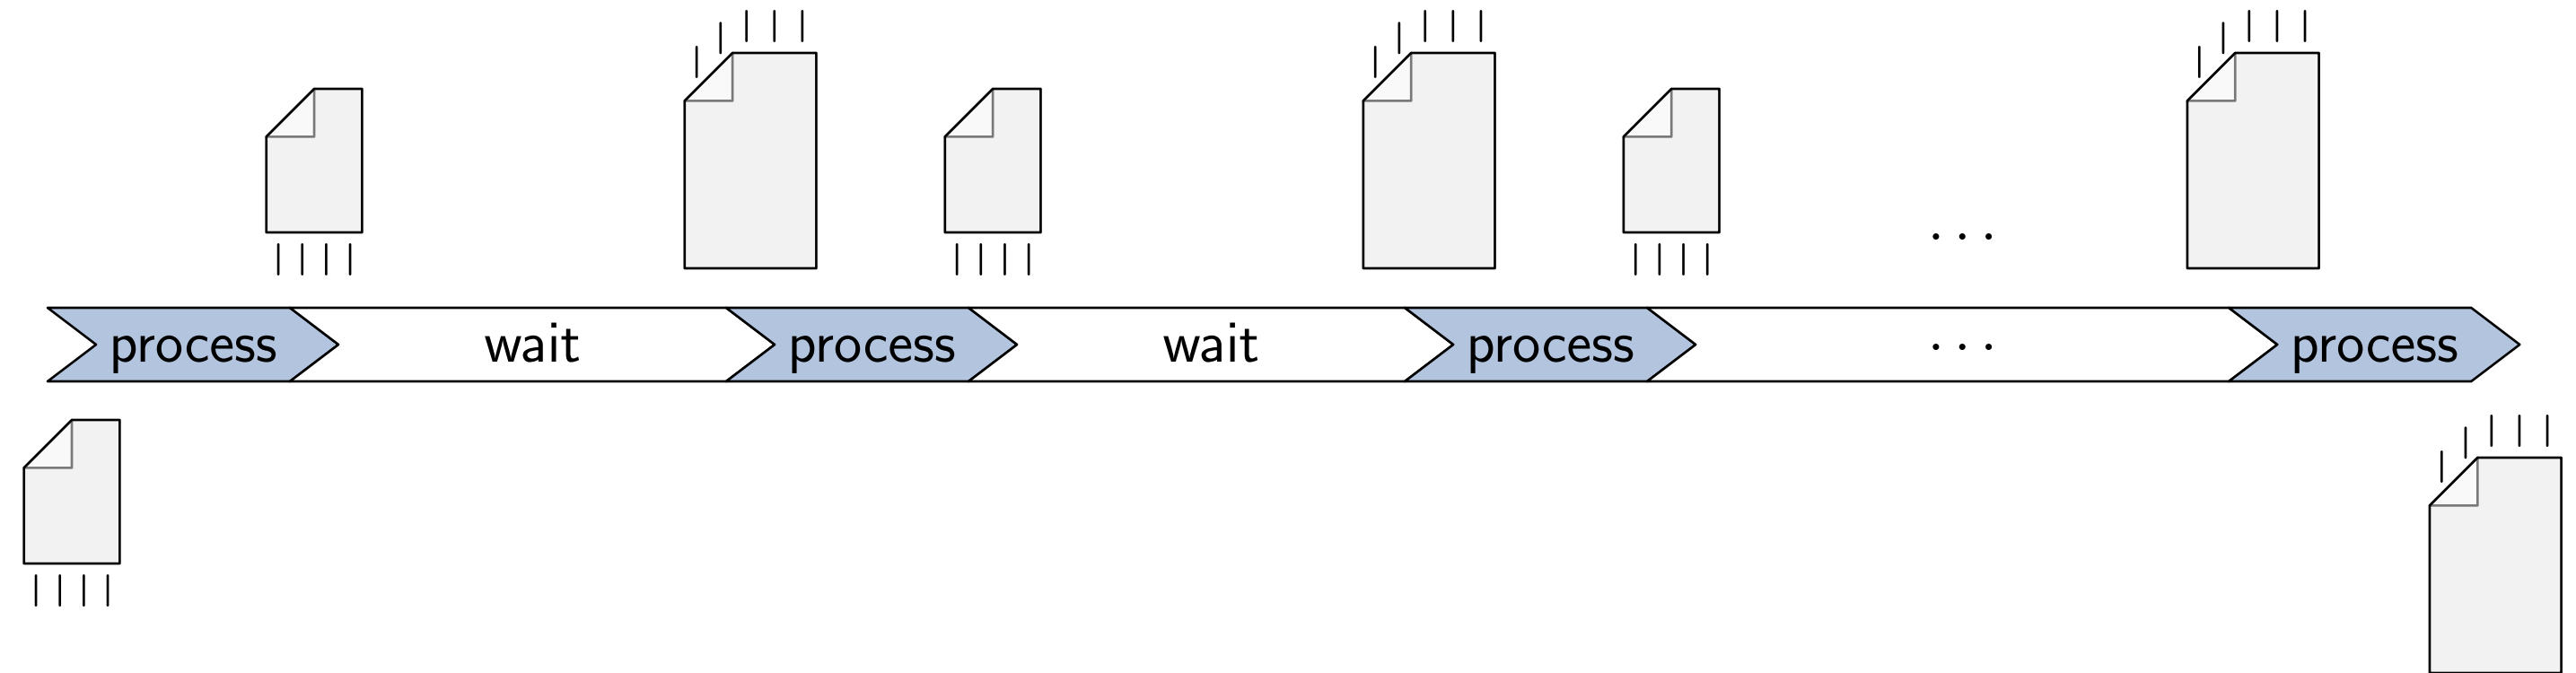
- IPv4
  - /32: 1
  - /24: 32
  - /20: 256
  - /18: 768
- IPv6
  - /128: 1
  - /64: 2
  - /56: 3
  - /48: 4
  - /32: 64

:	
172.16.96.1/32	count in $[0, 1L_I)$
:	
172.16.96.0/24	count in $[0, 32L_I)$
:	
172.16.96.0/20	count in $[0, 256L_I)$
:	
172.16.64.0/18	count in $[0, 768L_I)$
:	
2001:db8::734/128	count in $[0, 1L_I)$
:	
2001:db8::/32	count in $[0, 64L_I)$
:	



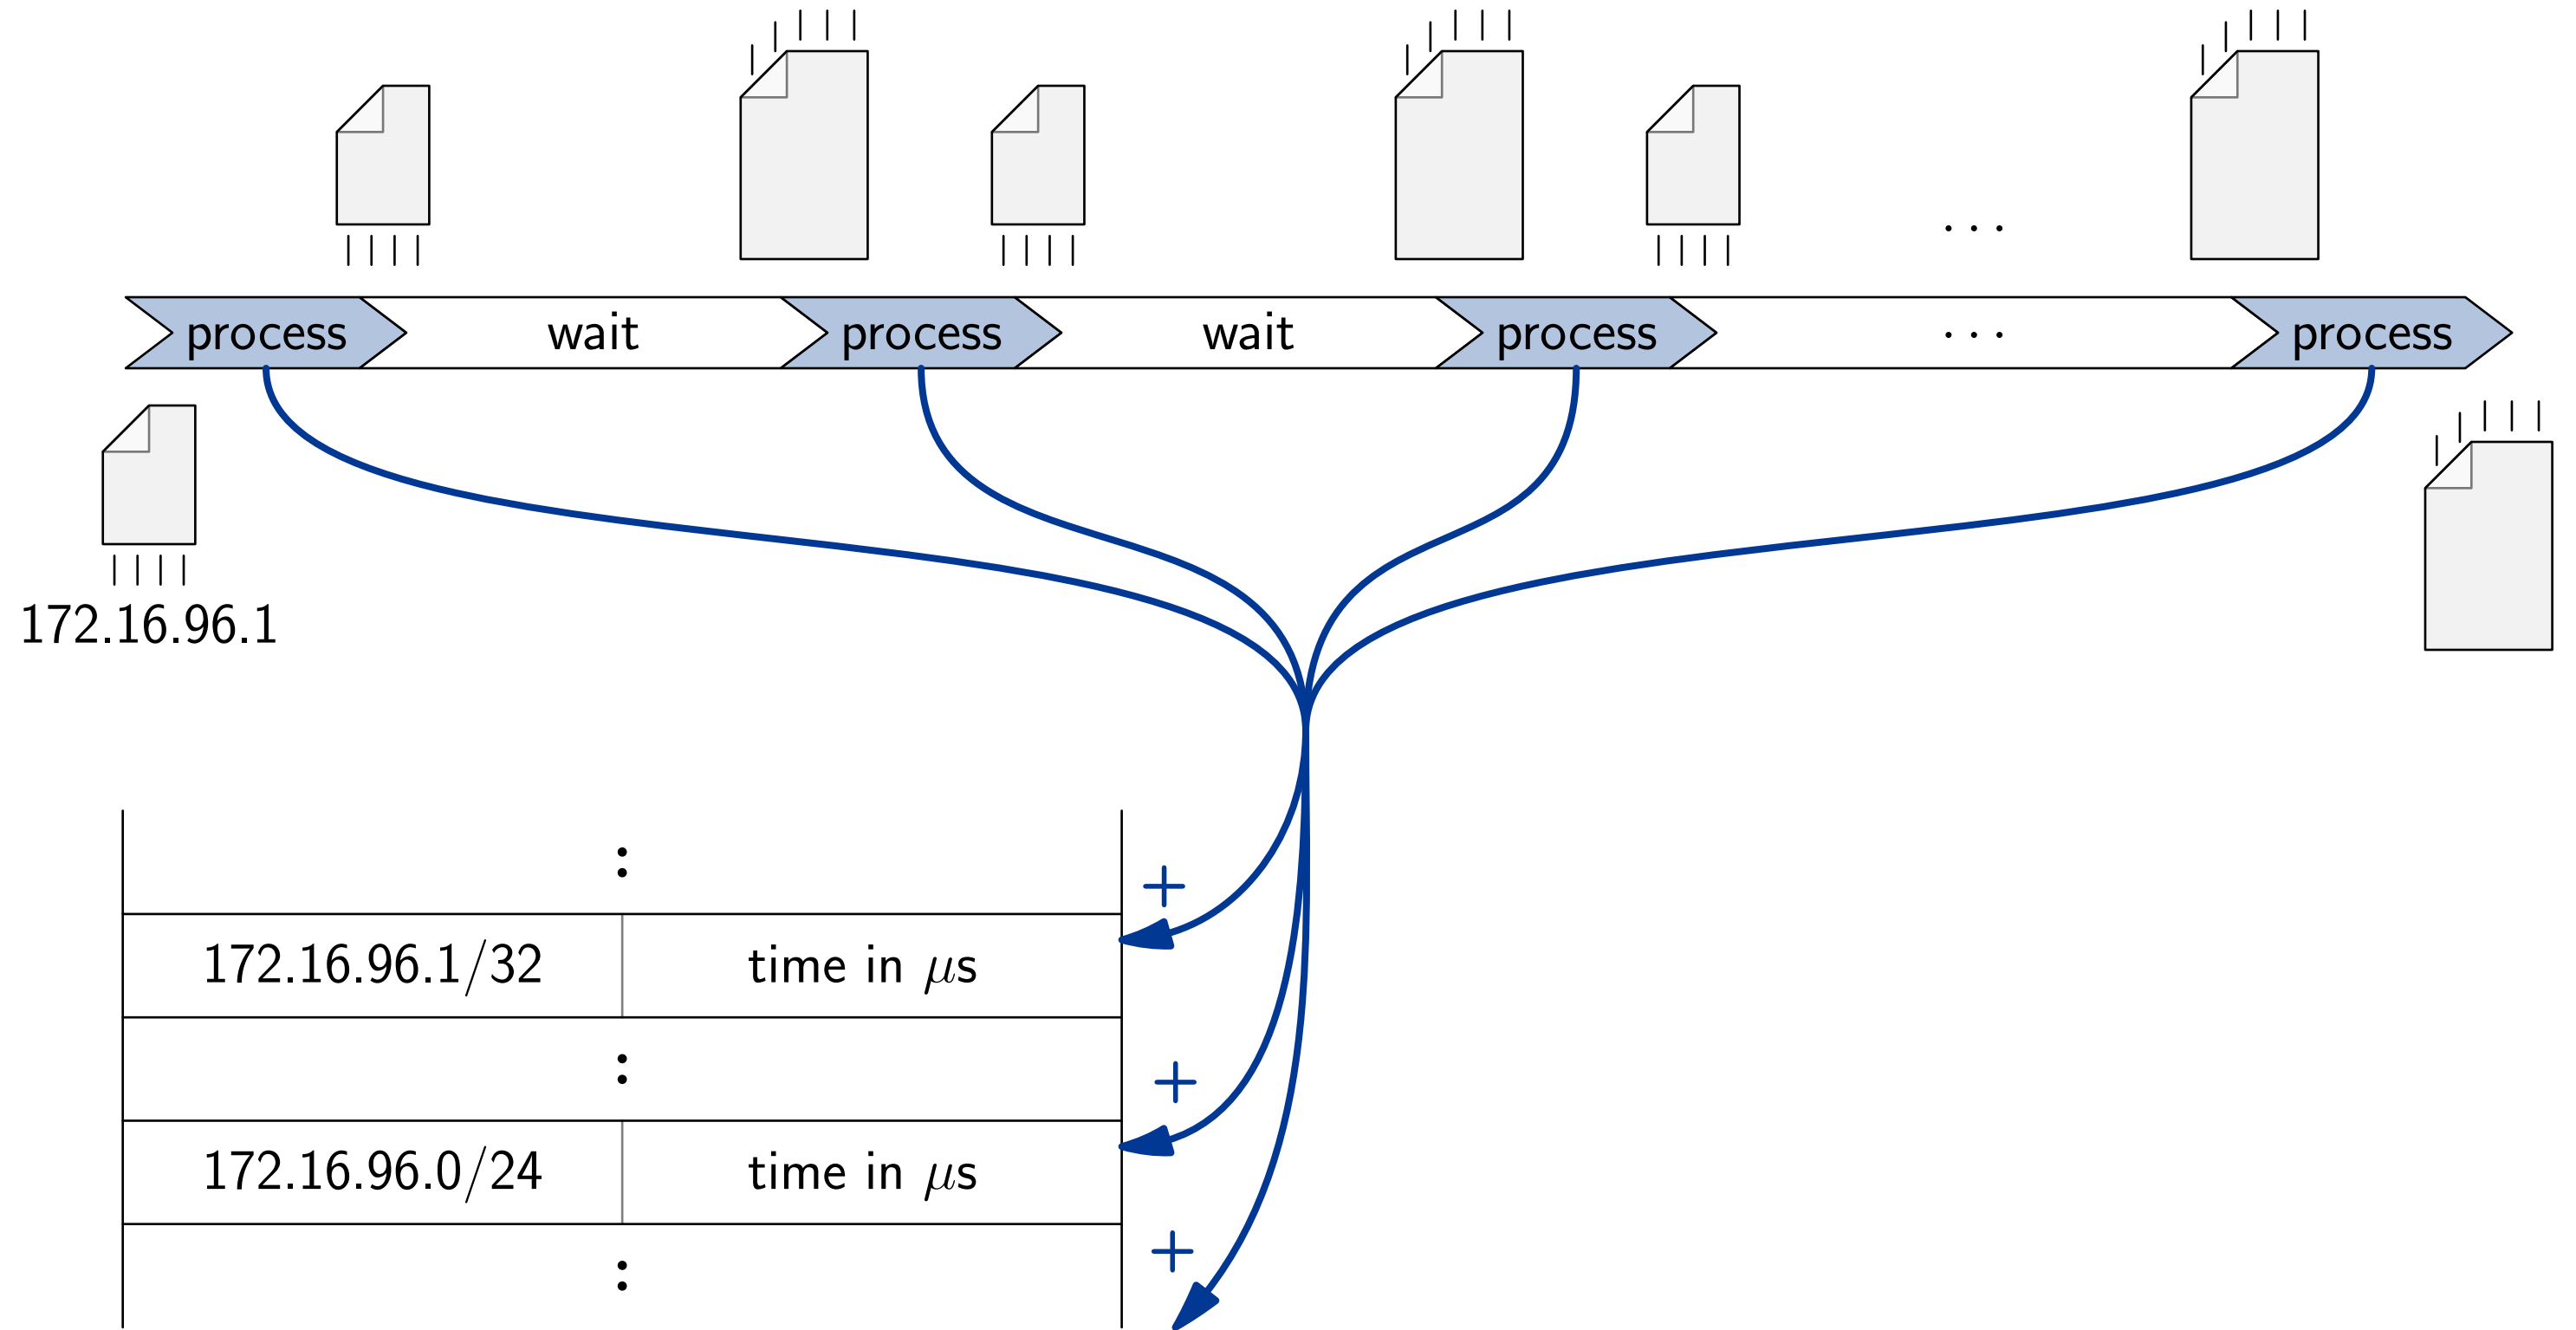
# Prioritization

- **measuring time**
  - **only cpu, no wait**



# Prioritization


- measuring time
  - only cpu, no wait
  - **add to table values**





# Prioritization

- measuring time
  - only cpu, no wait
  - add to table values
- defer to queues
  - based on values



# Implementation

- hashing


hash(172.16.96.1/32) = 10111101100001101010110000110101001110100111010011101101001000001011111

:	:
172.16.96.1/32	3 ∈ [0, 1L <sub>I</sub> )
:	:
172.16.96.0/24	15.34 ∈ [0, 32L <sub>I</sub> )
:	:
172.16.96.0/20	123 ∈ [0, 256L <sub>I</sub> )
:	:
2001:db8::734/128	7.569 ∈ [0, 1L <sub>I</sub> )
:	:
2001:db8::/32	33.21 ∈ [0, 64L <sub>I</sub> )
:	:

# Implementation

- hashing
  - buckets

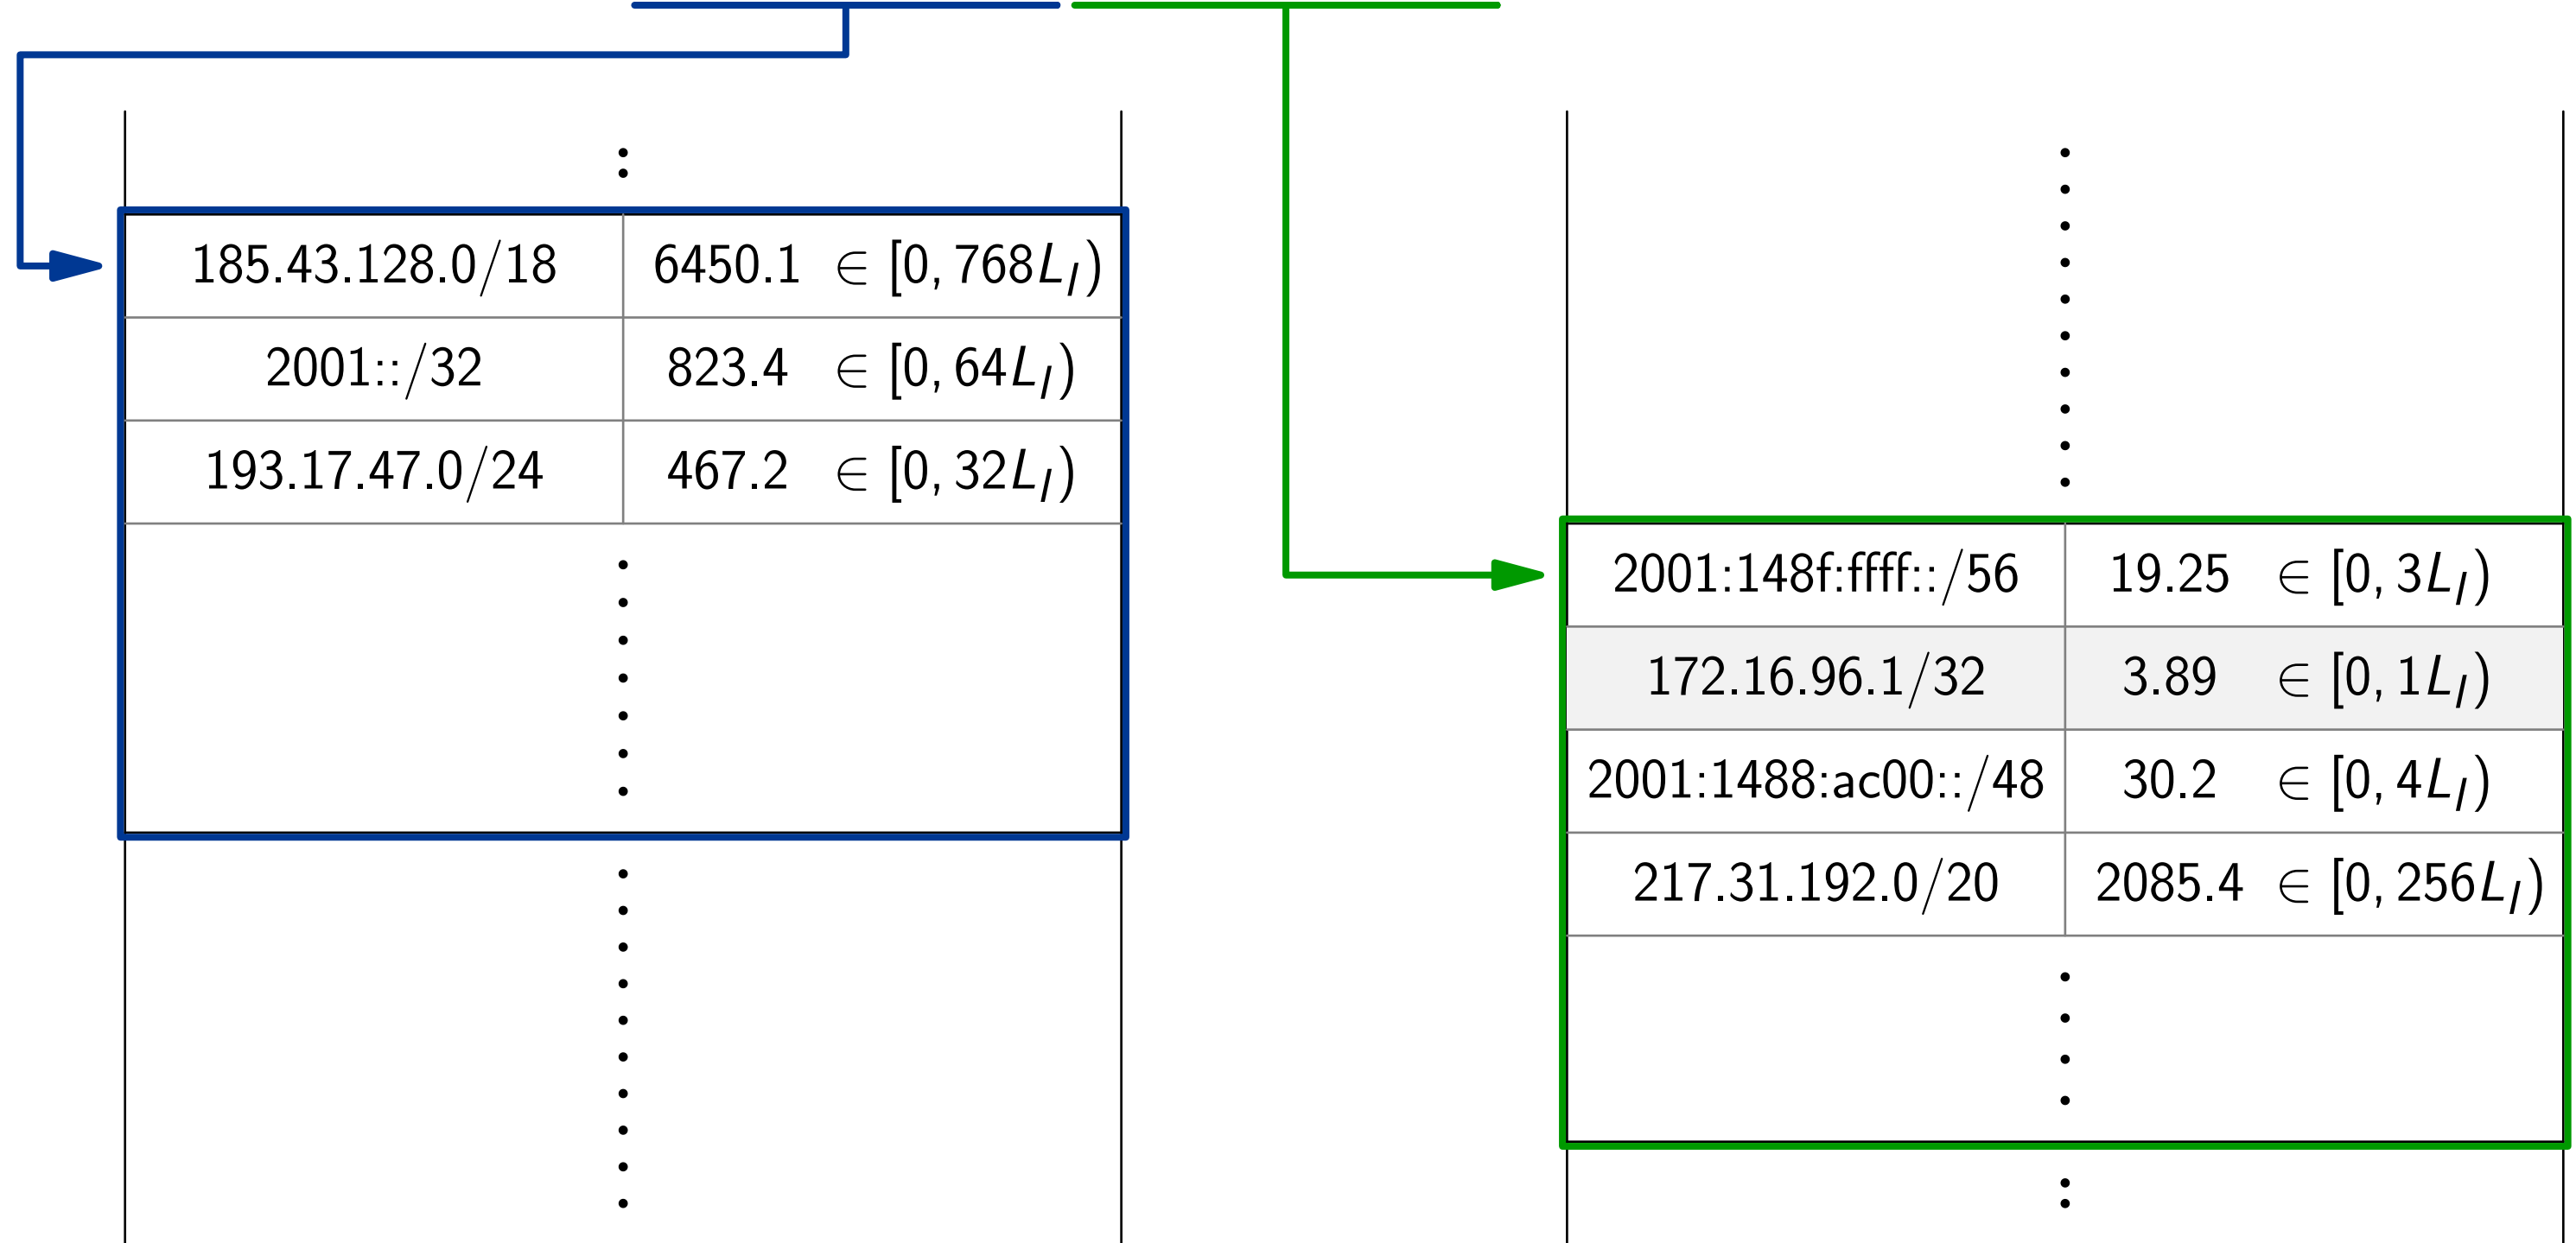
hash(172.16.96.1/32) = 1011110110000110101011000011010100111010011101101001000001011111



# Implementation

- hashing
  - buckets
  - **two tables**

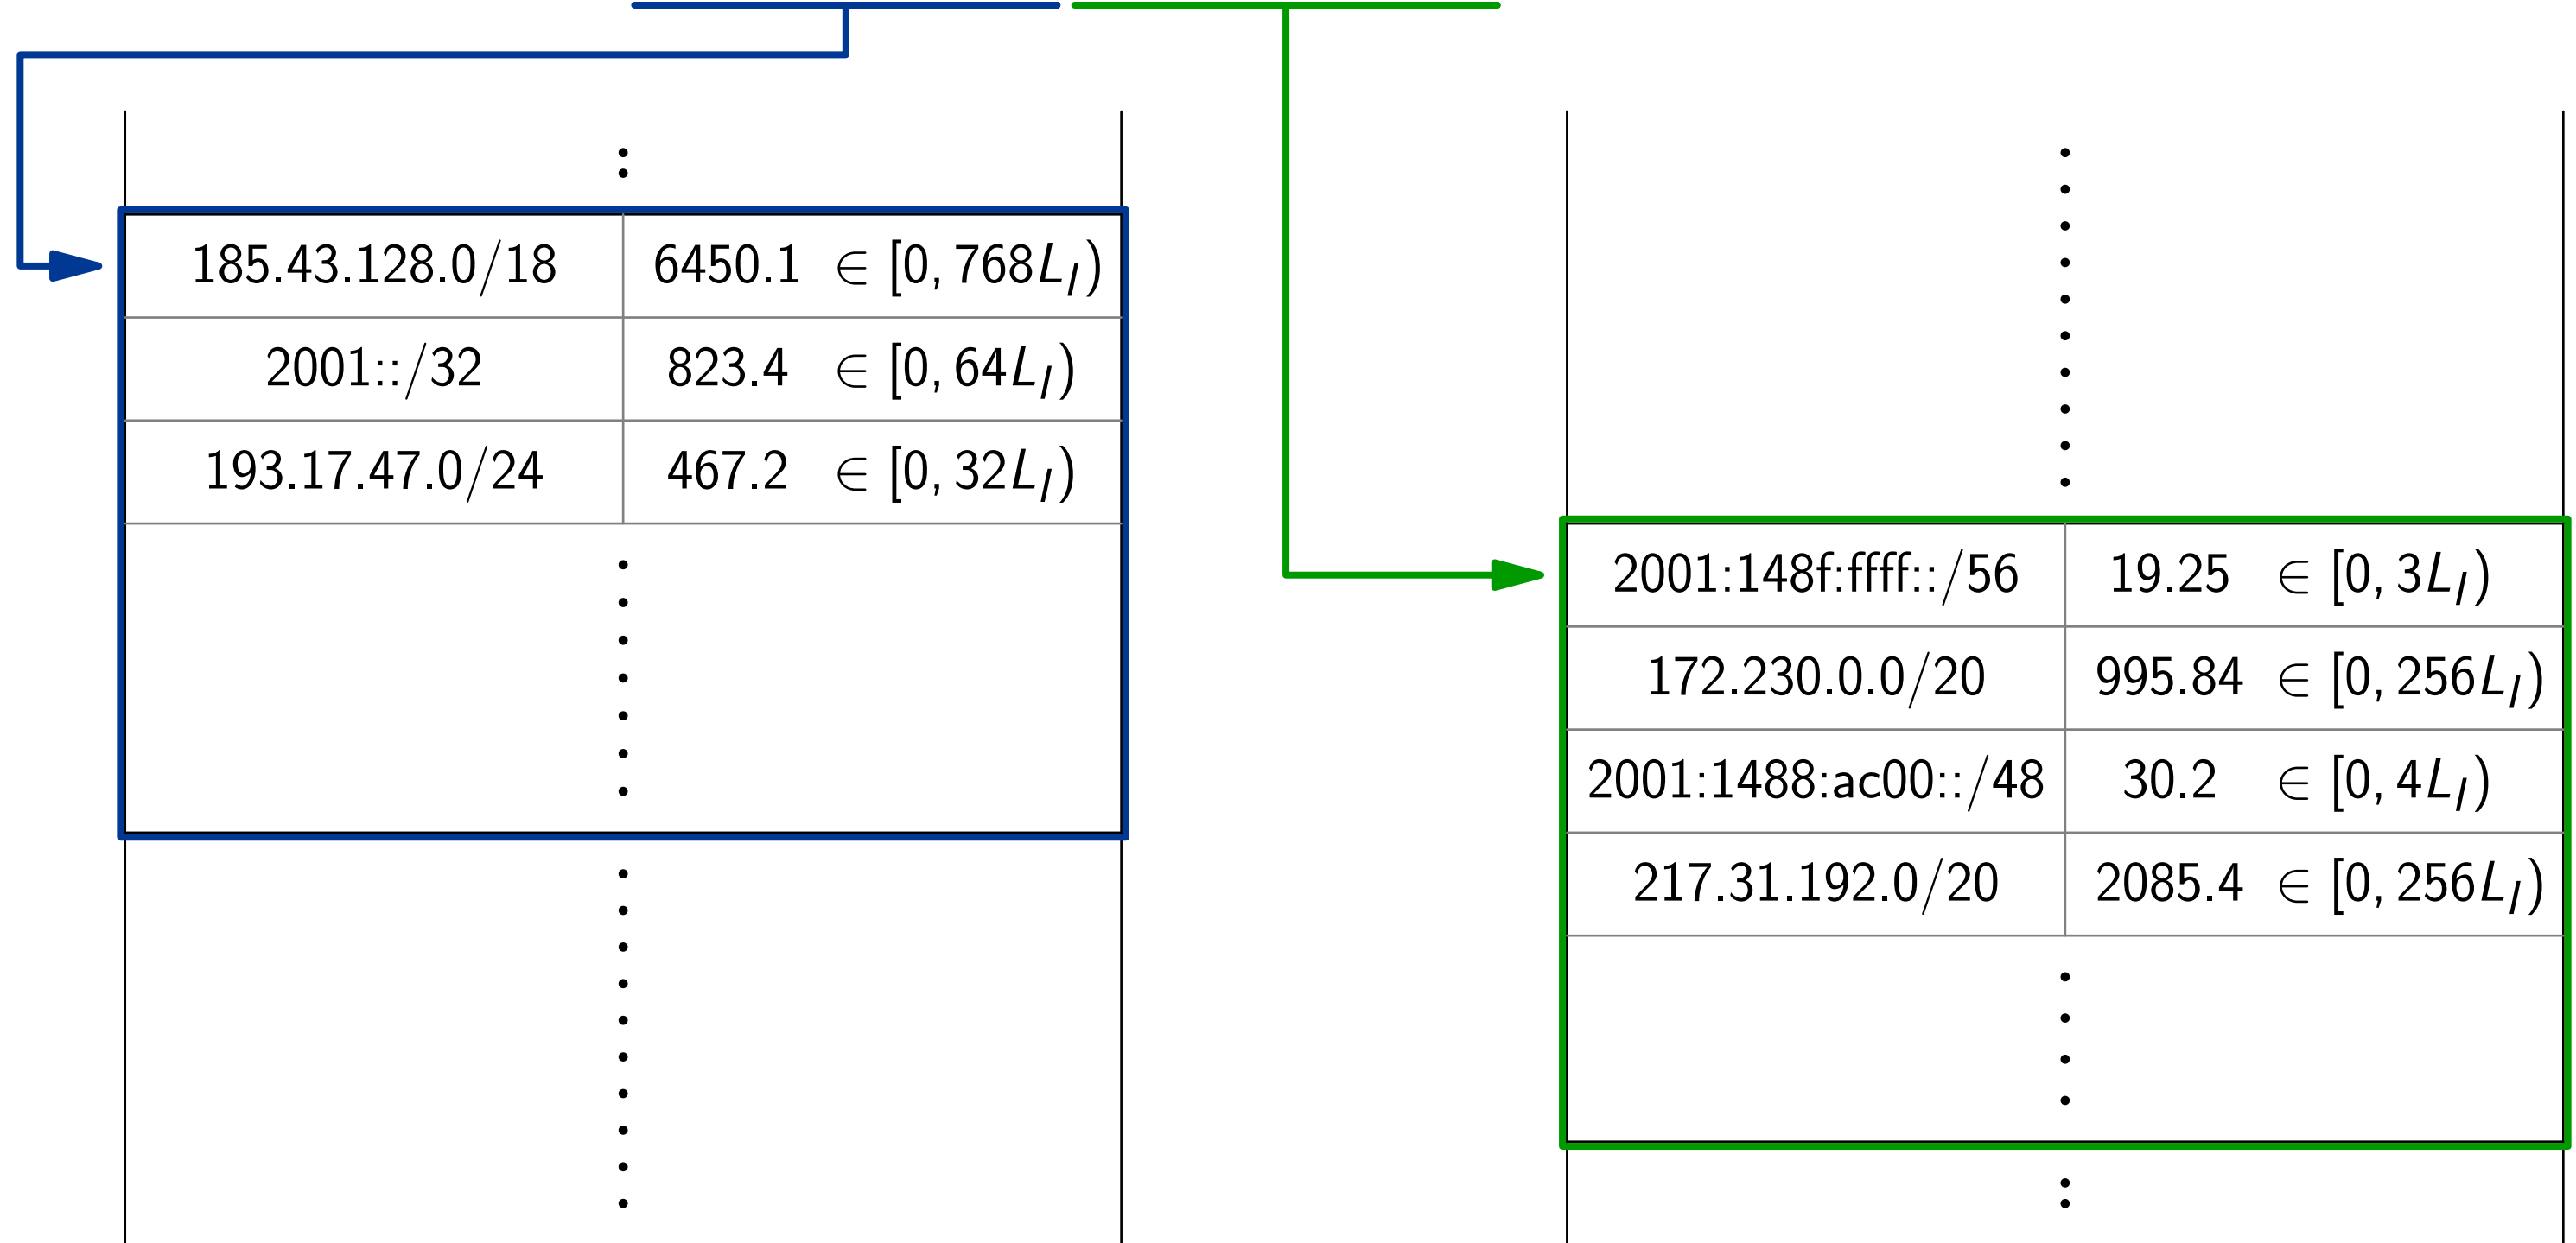
hash(172.16.96.1/32) = 1011110110000110101011000011010100111010011101101001000001011111



# Implementation

- hashing
  - buckets
  - two tables
- evicting

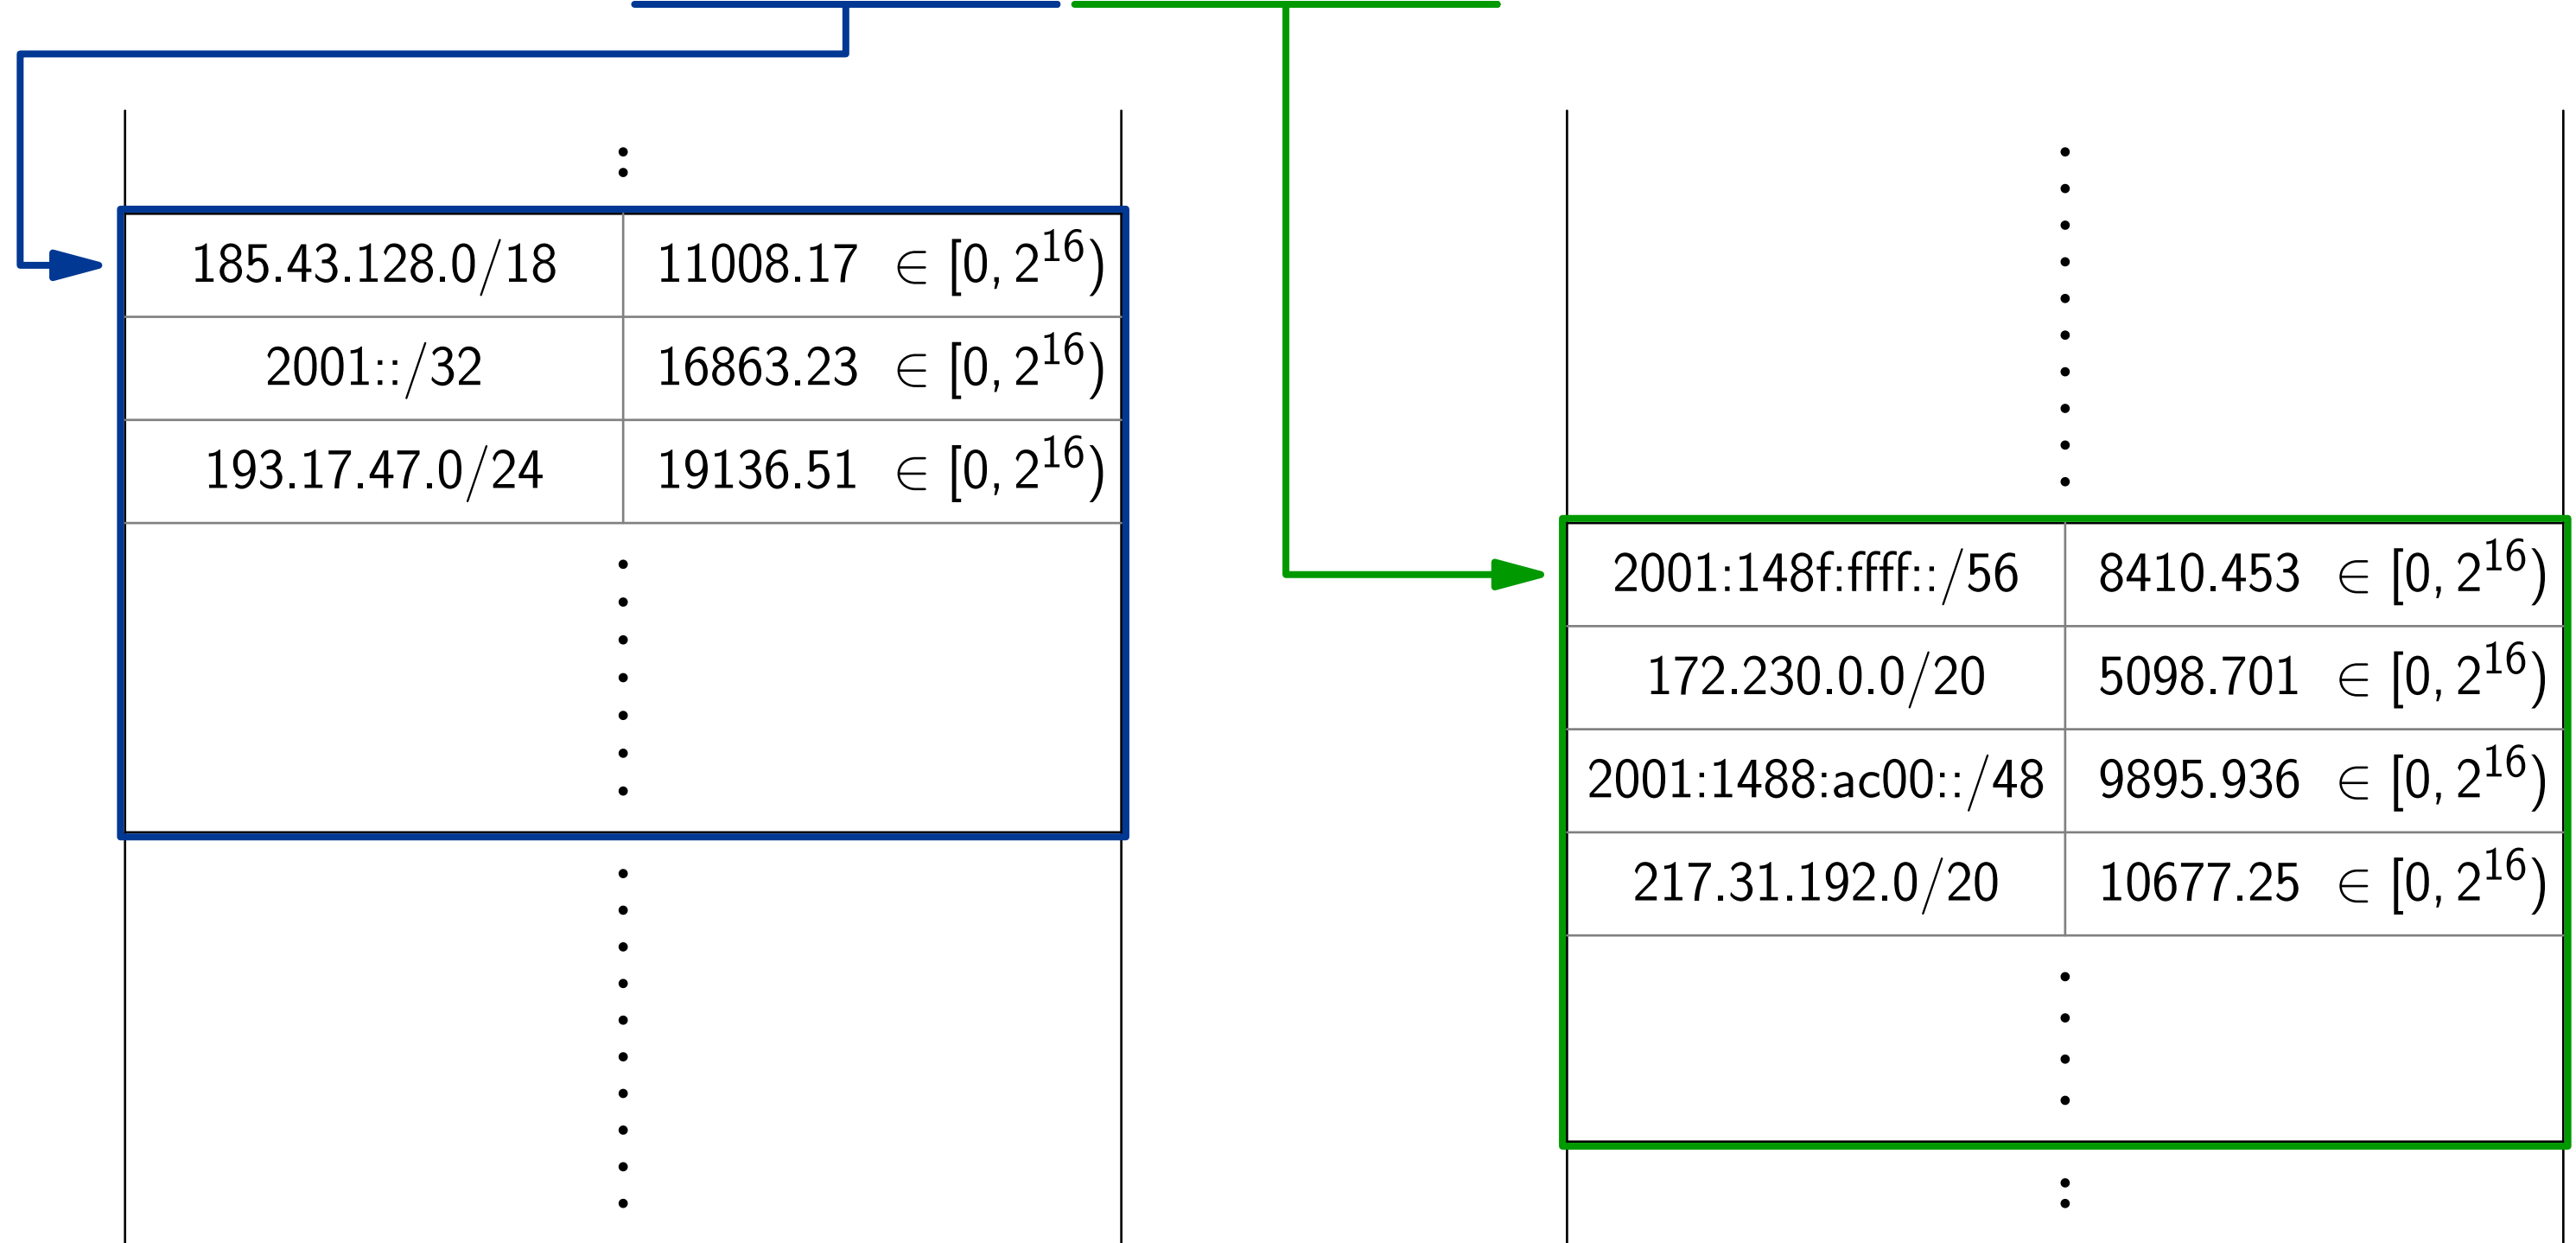
hash(172.16.96.1/32) = 1011110110000110101011000011010100111010011101101001000001011111



# Implementation

- hashing
  - buckets
  - two tables
- evicting
  - **normalized limits**

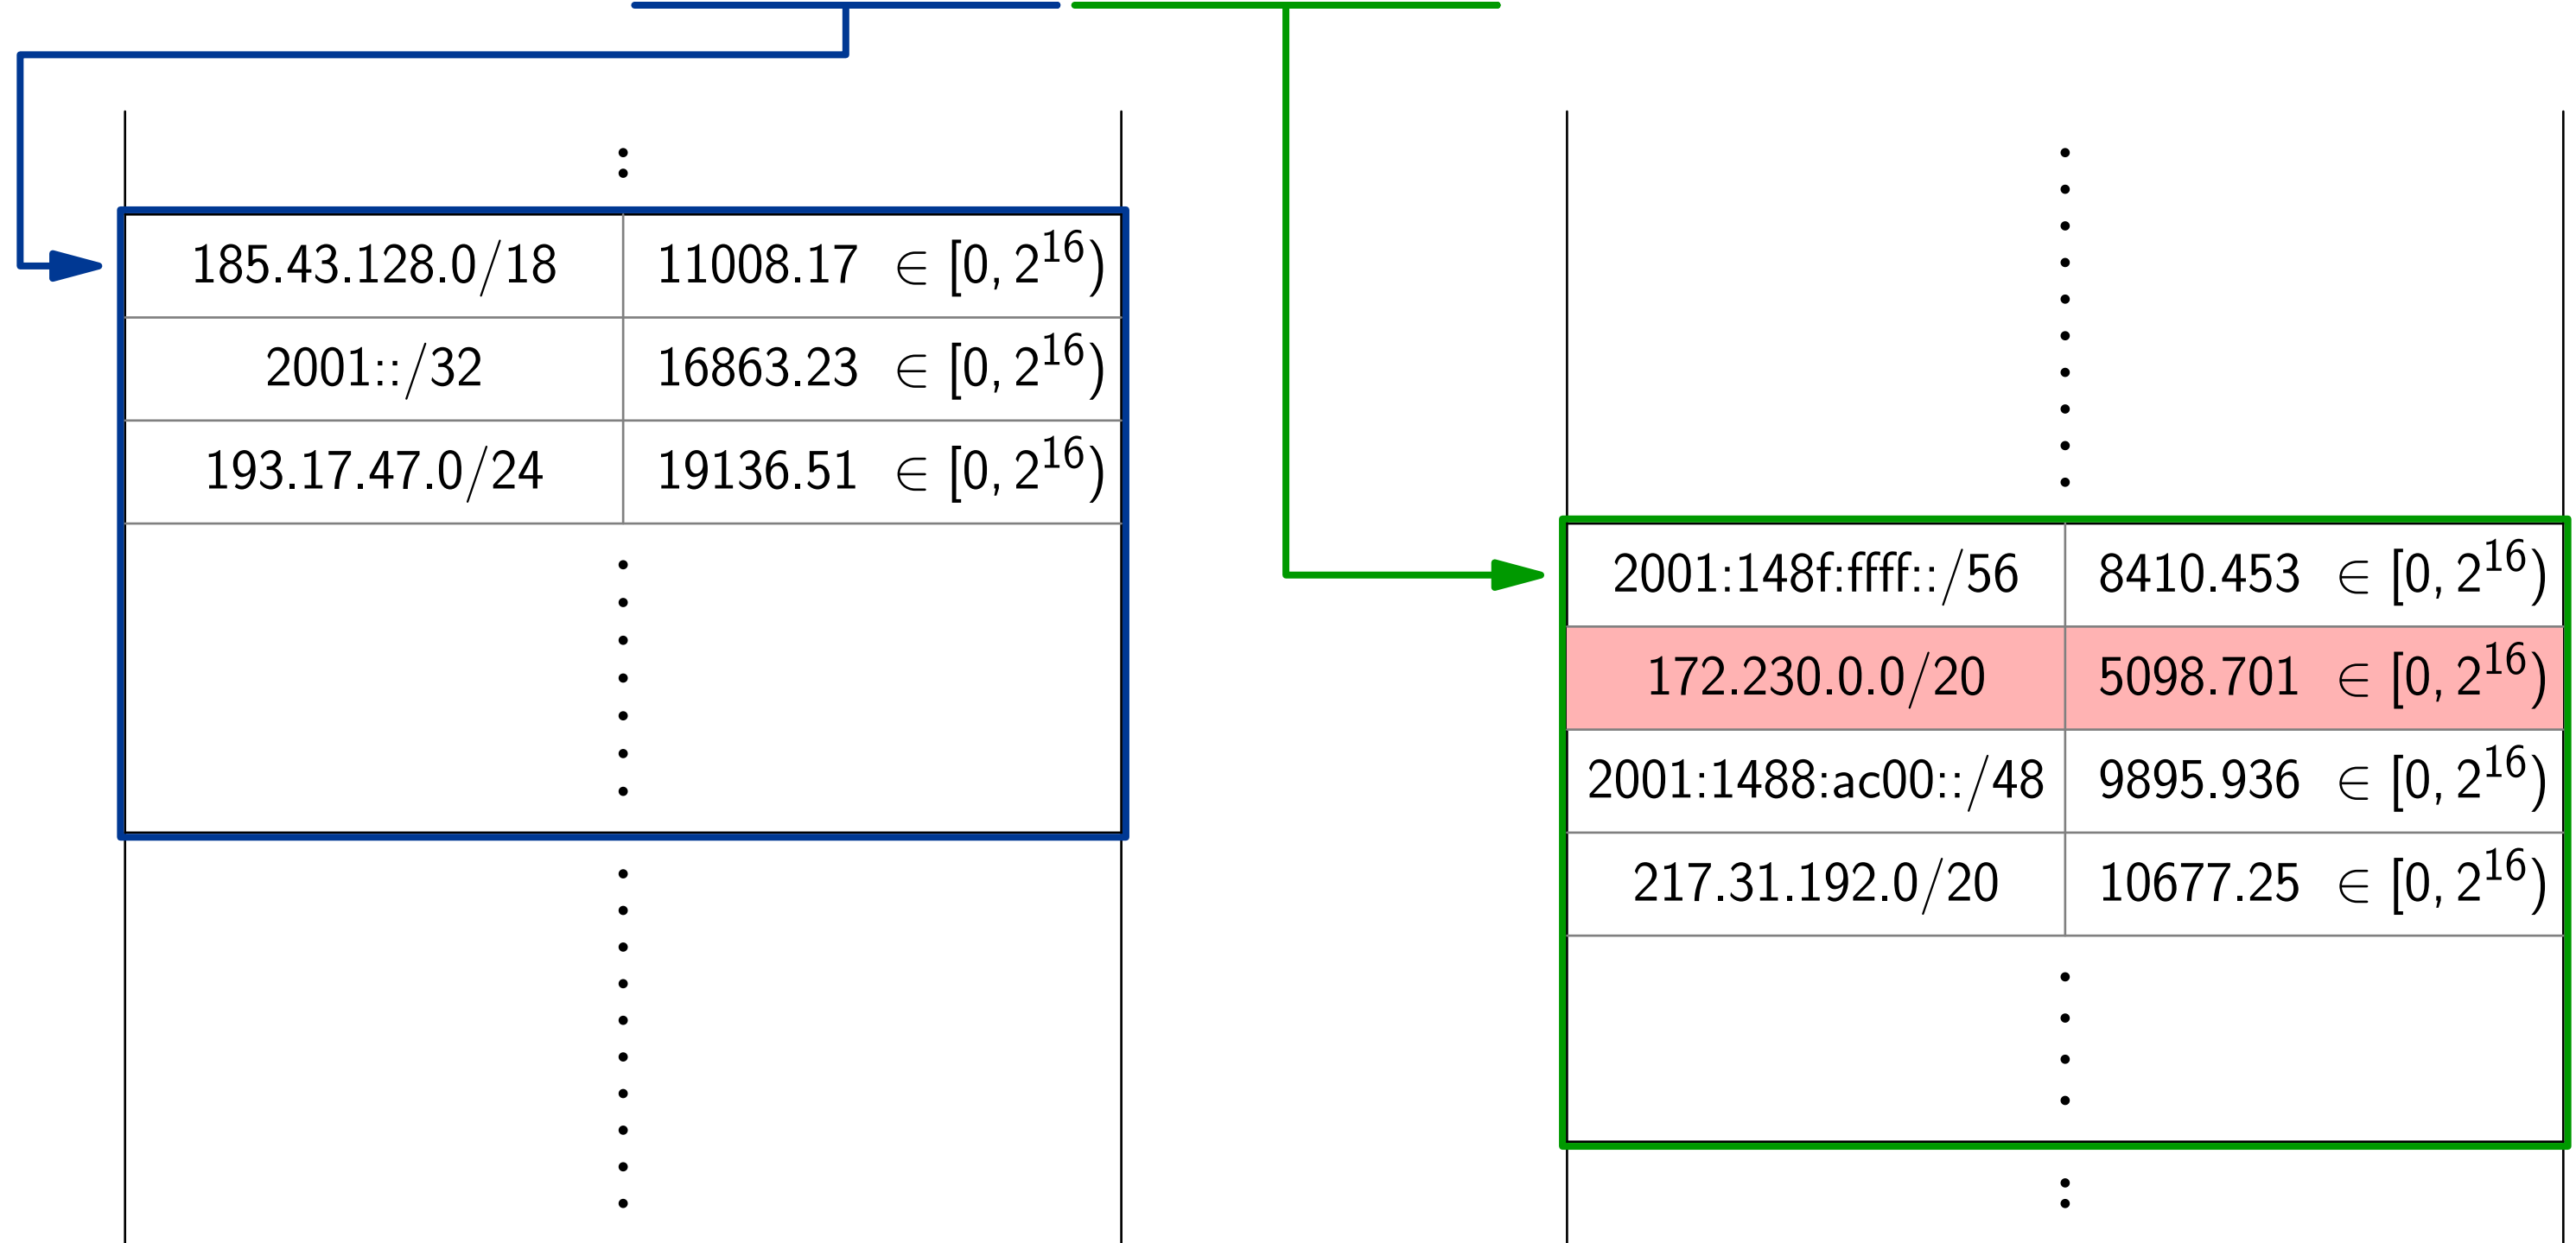
hash(172.16.96.1/32) = 1011110110000110101011000011010100111010011101101001000001011111



# Implementation

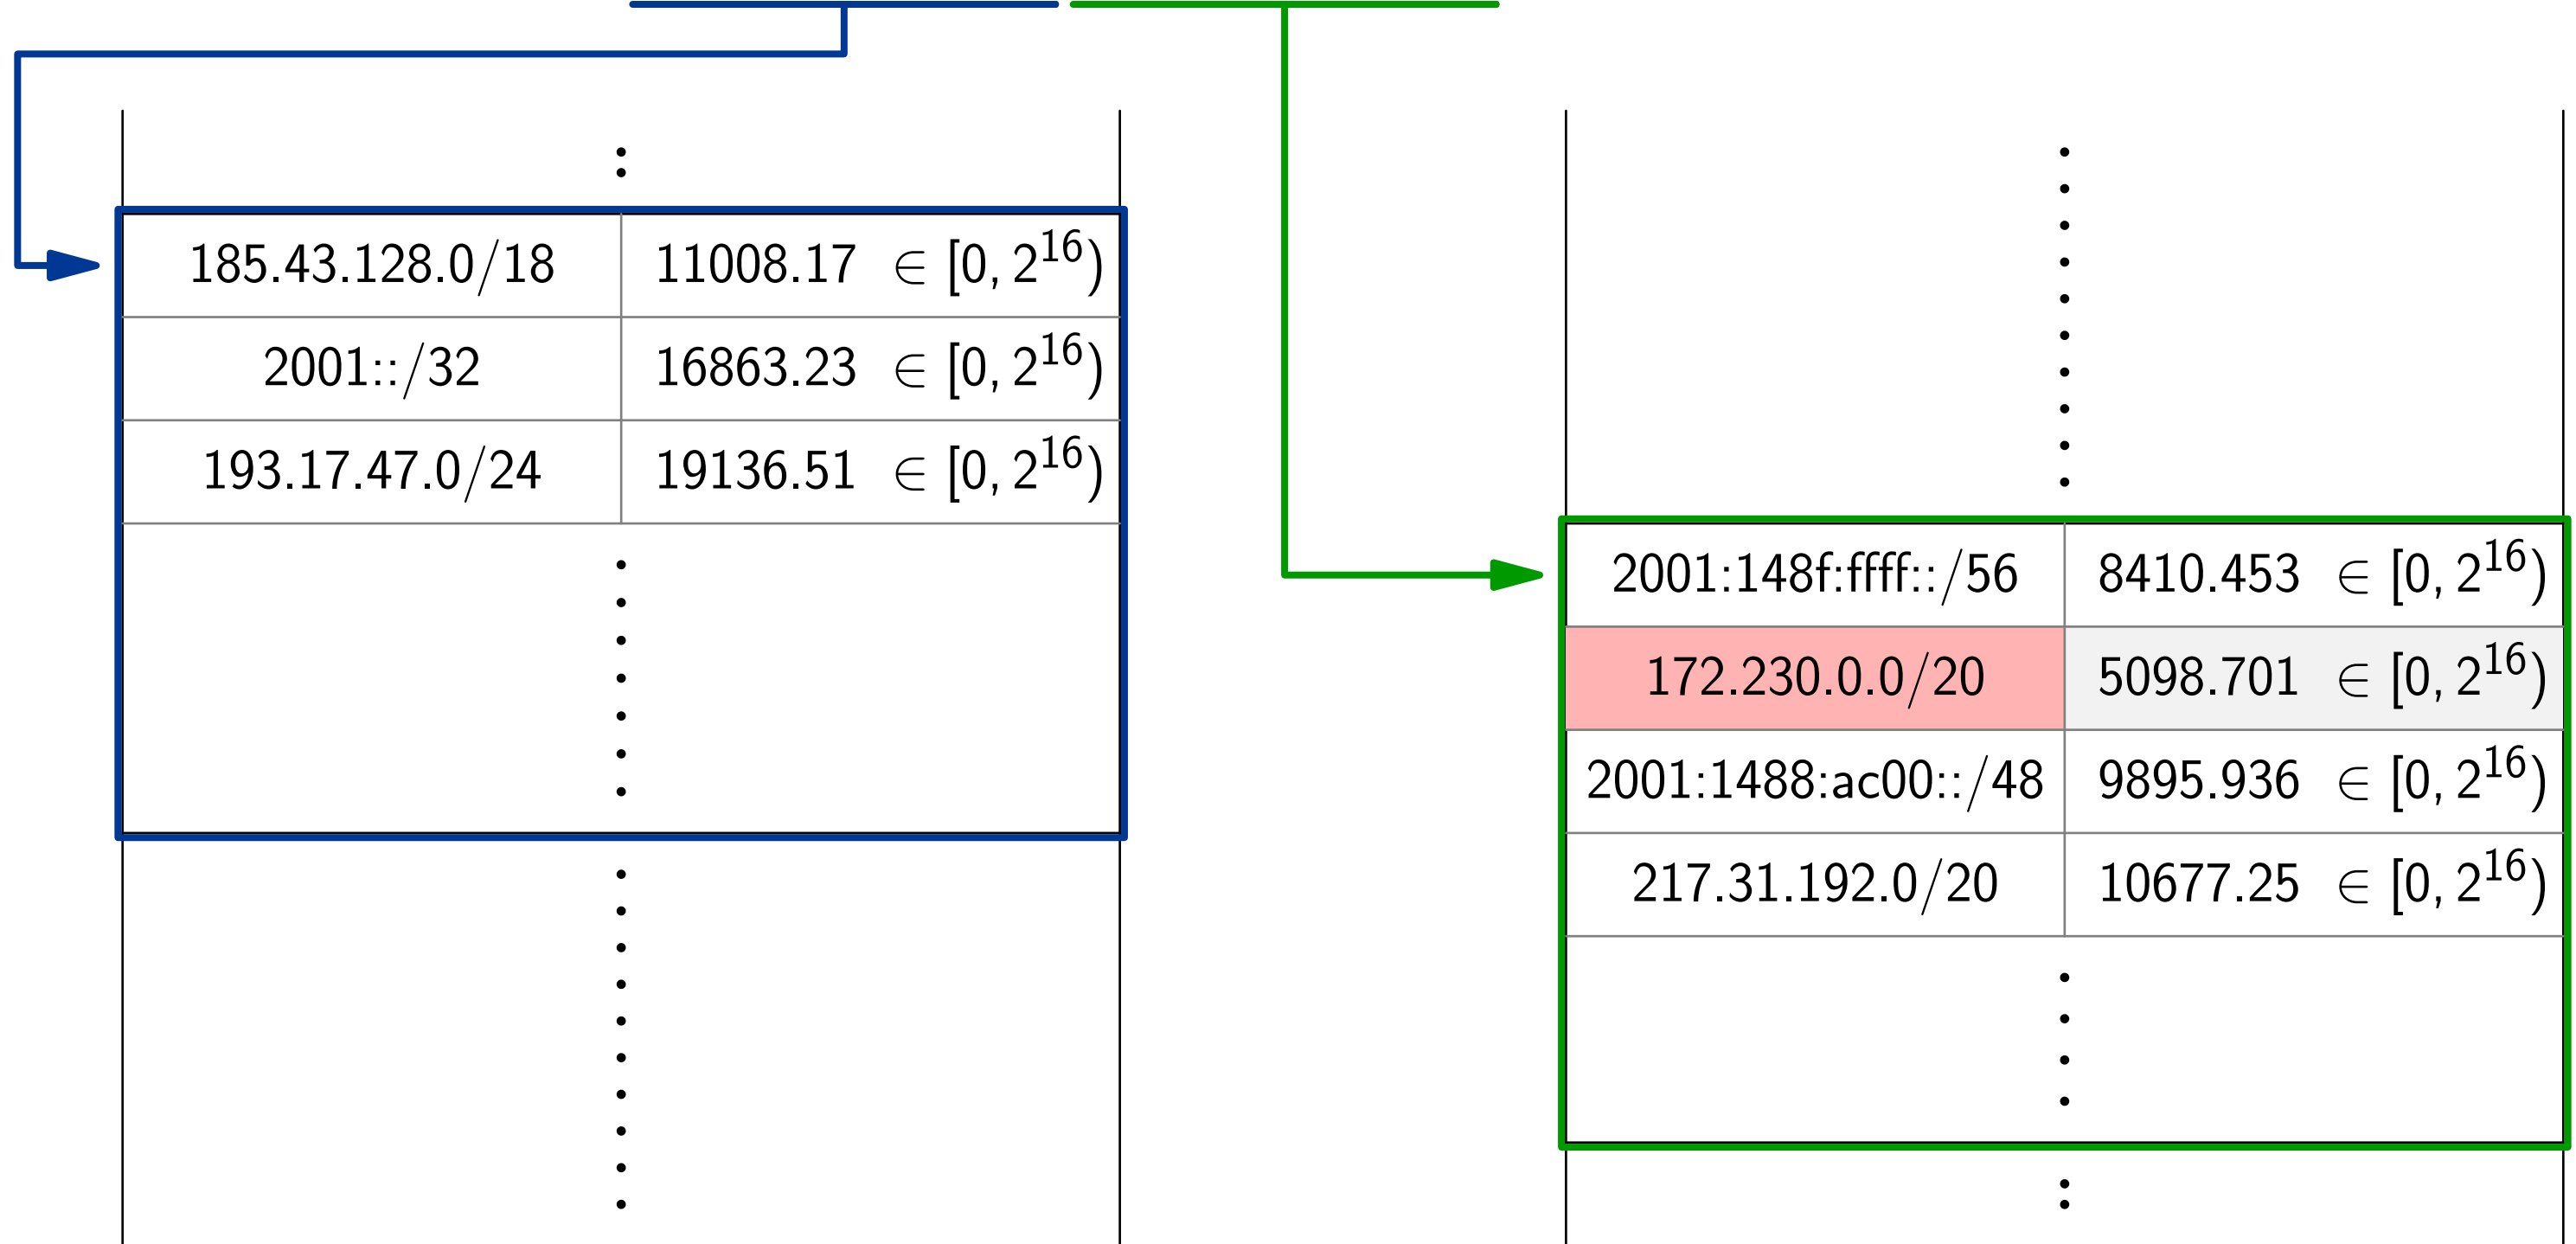
- hashing
  - buckets
  - two tables
- evicting
  - normalized limits
  - **choosing minimal**

hash(172.16.96.1/32) = 10111101100001101010110000110101001110100111010011101101001000001011111



# Implementation

hash(172.16.96.1/32) = 1011110110000110101011000011010100111010011101101001000001011111

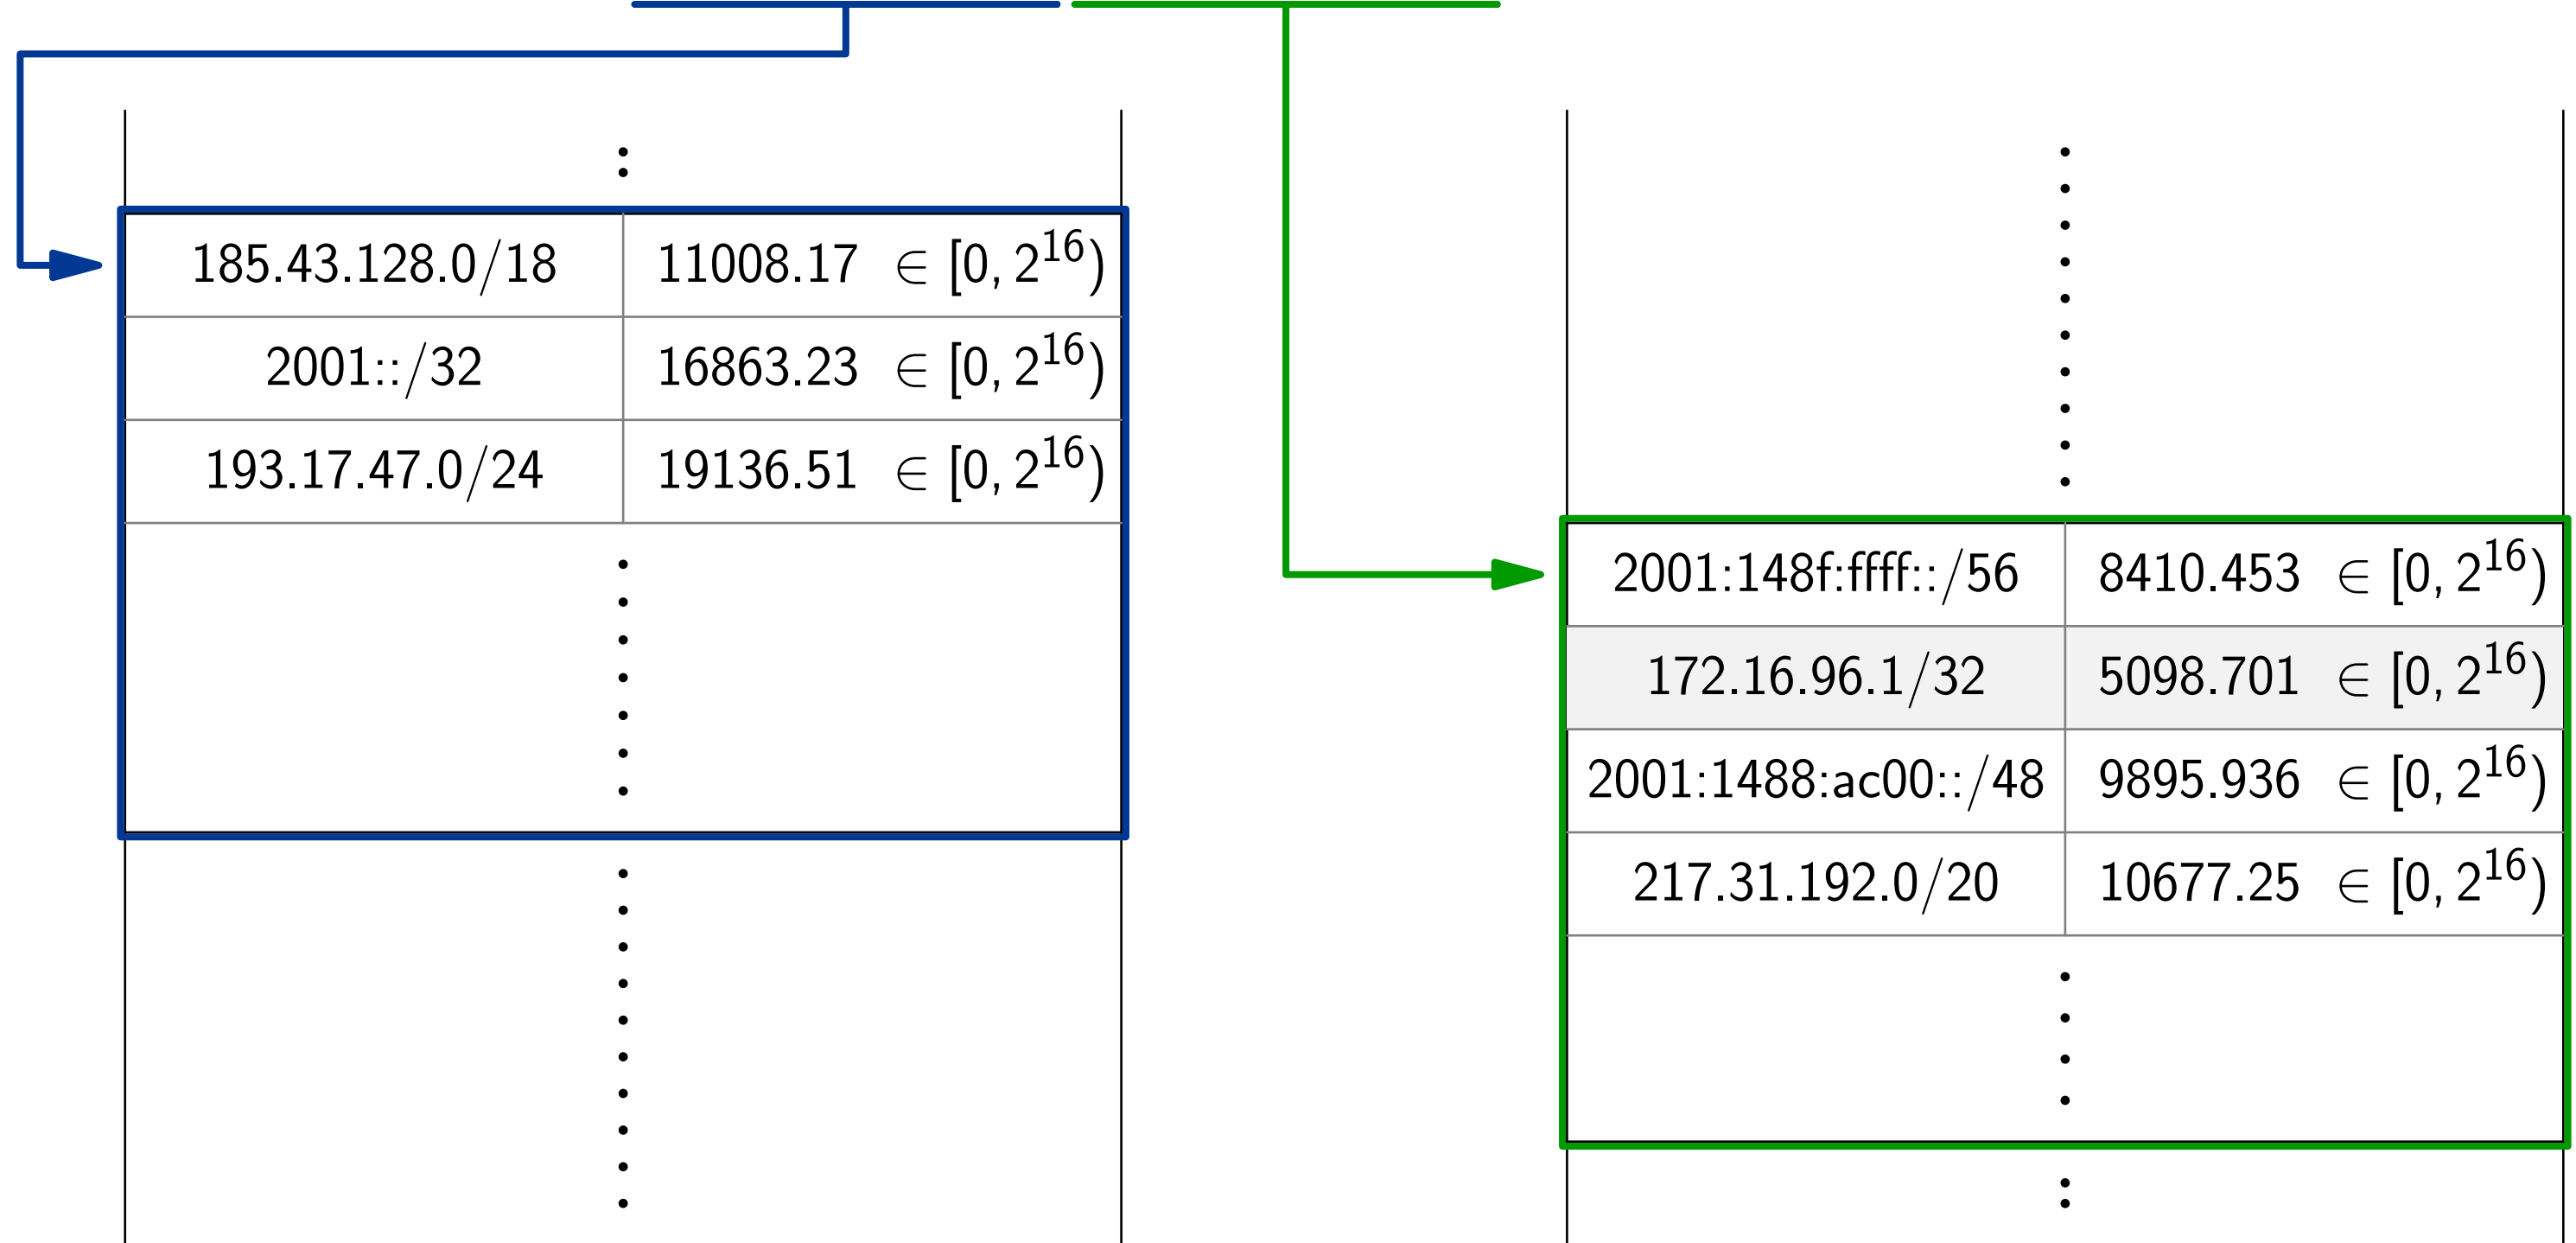


- hashing
  - buckets
  - two tables
- evicting
  - normalized limits
  - choosing minimal
  - **keeping value**



# Implementation


hash(172.16.96.1/32) = 1011110110000110101011000011010100111010011101101001000001011111



- hashing
  - buckets
  - two tables
- evicting
  - normalized limits
  - choosing minimal
  - keeping value
- lazy decay

# Implementation

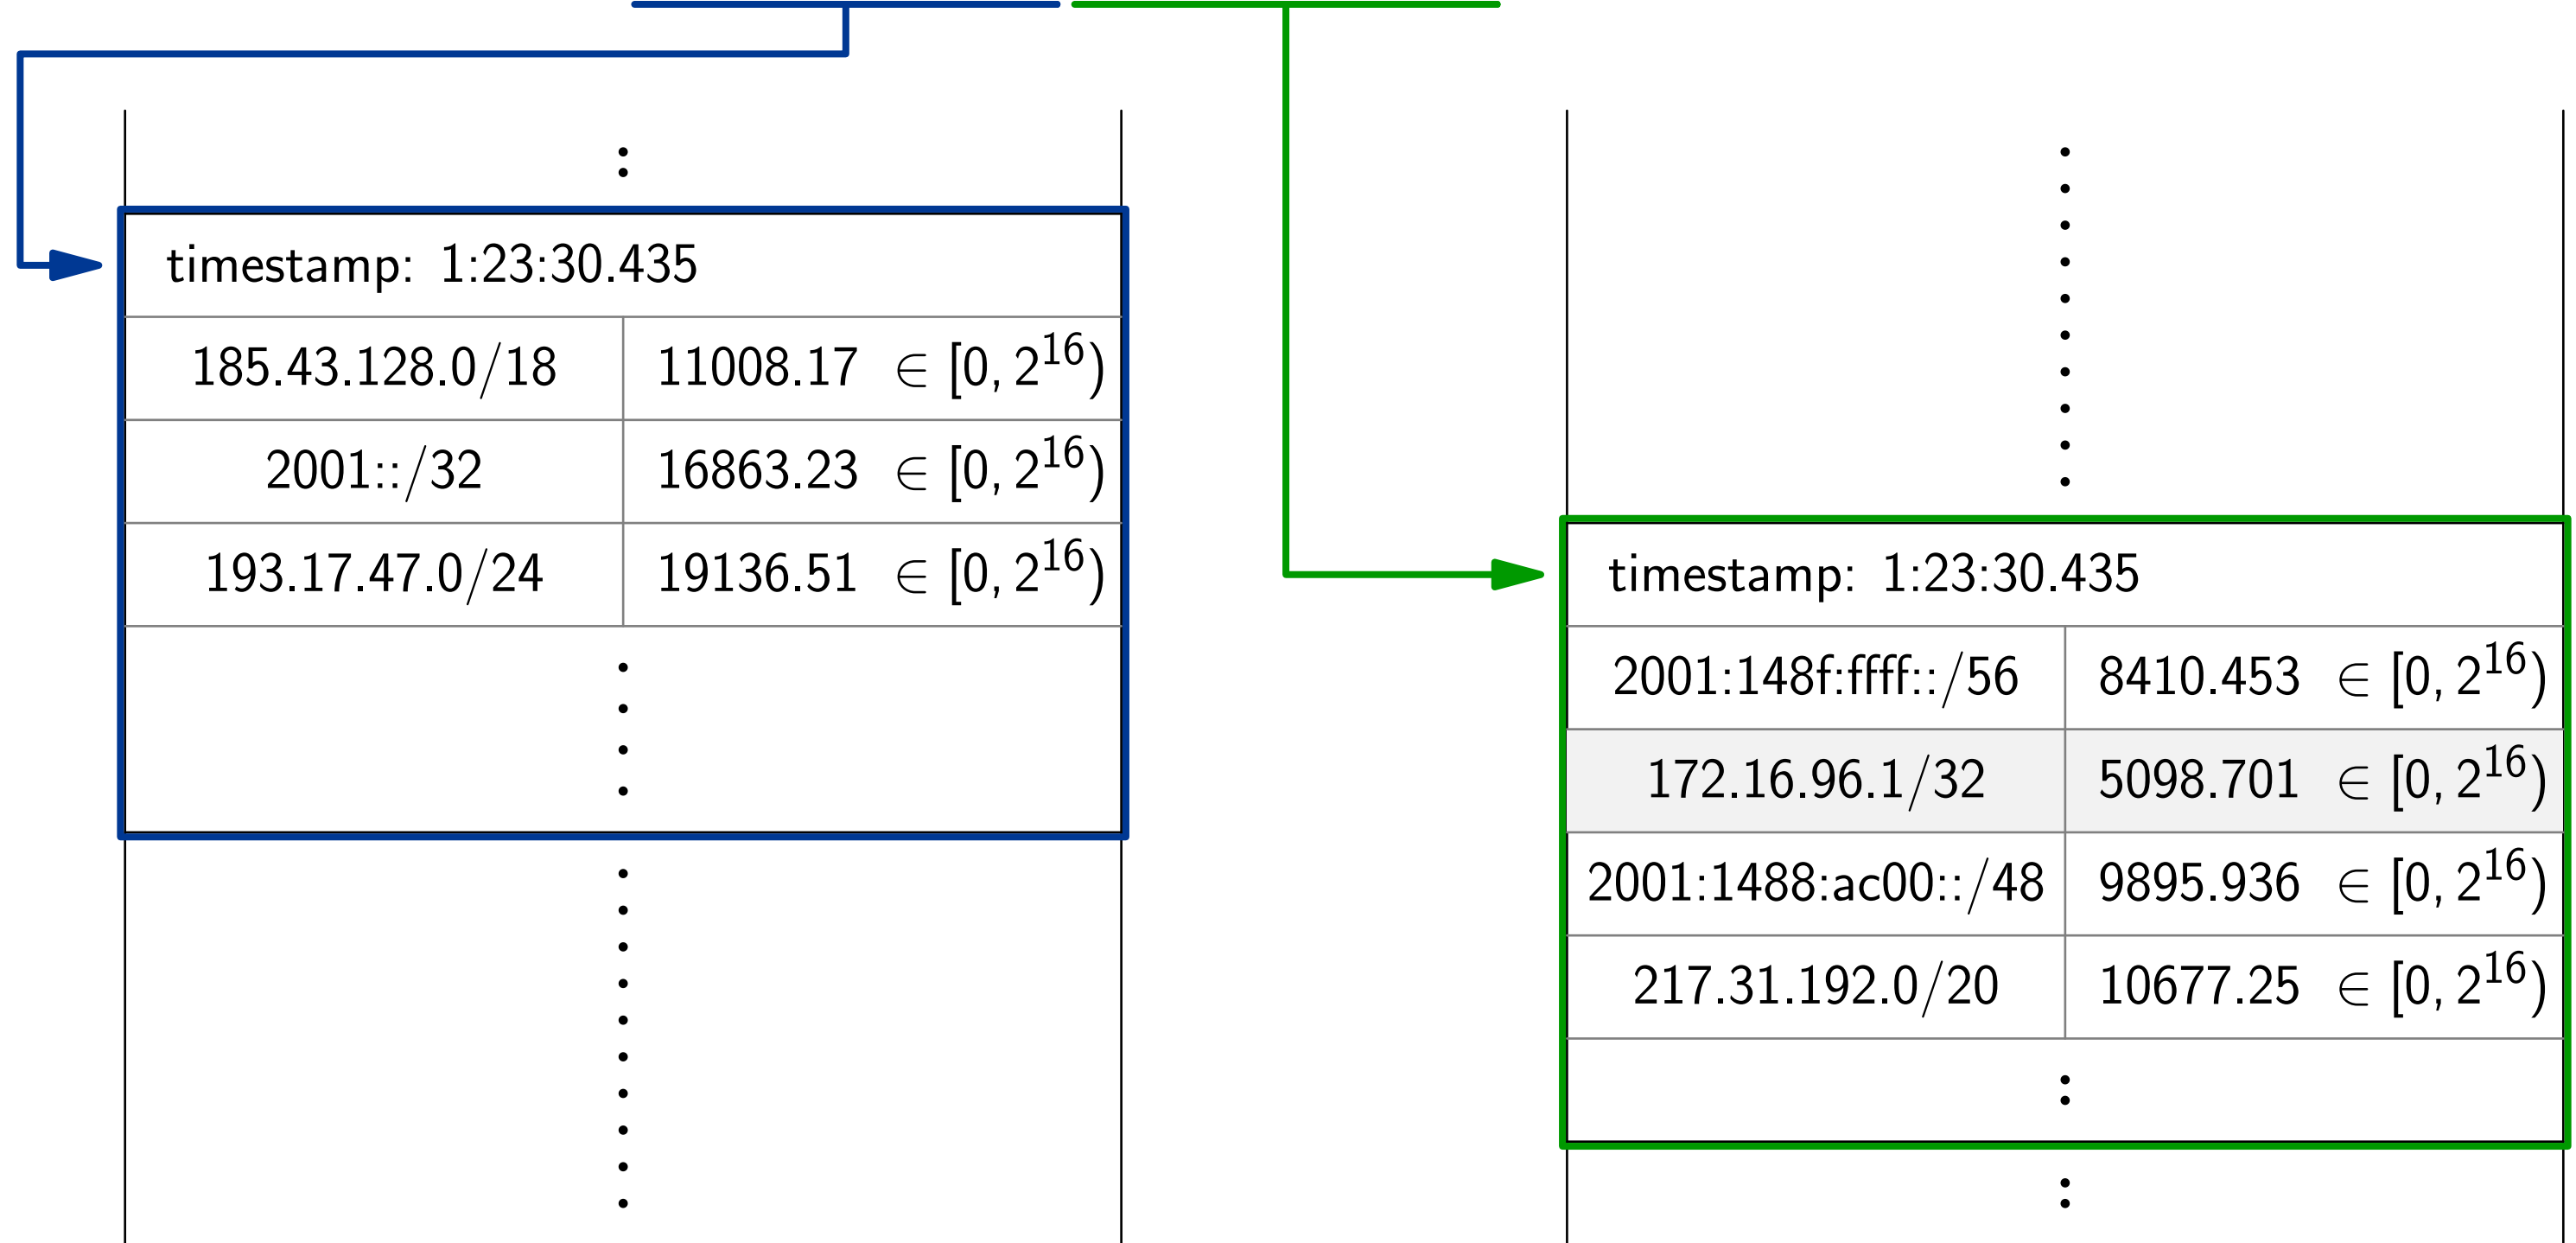
hash(172.16.96.1/32) = 1011110110000110101011000011010100111010011101101001000001011111



- hashing
  - buckets
  - two tables
- evicting
  - normalized limits
  - choosing minimal
  - keeping value
- lazy decay

# Implementation

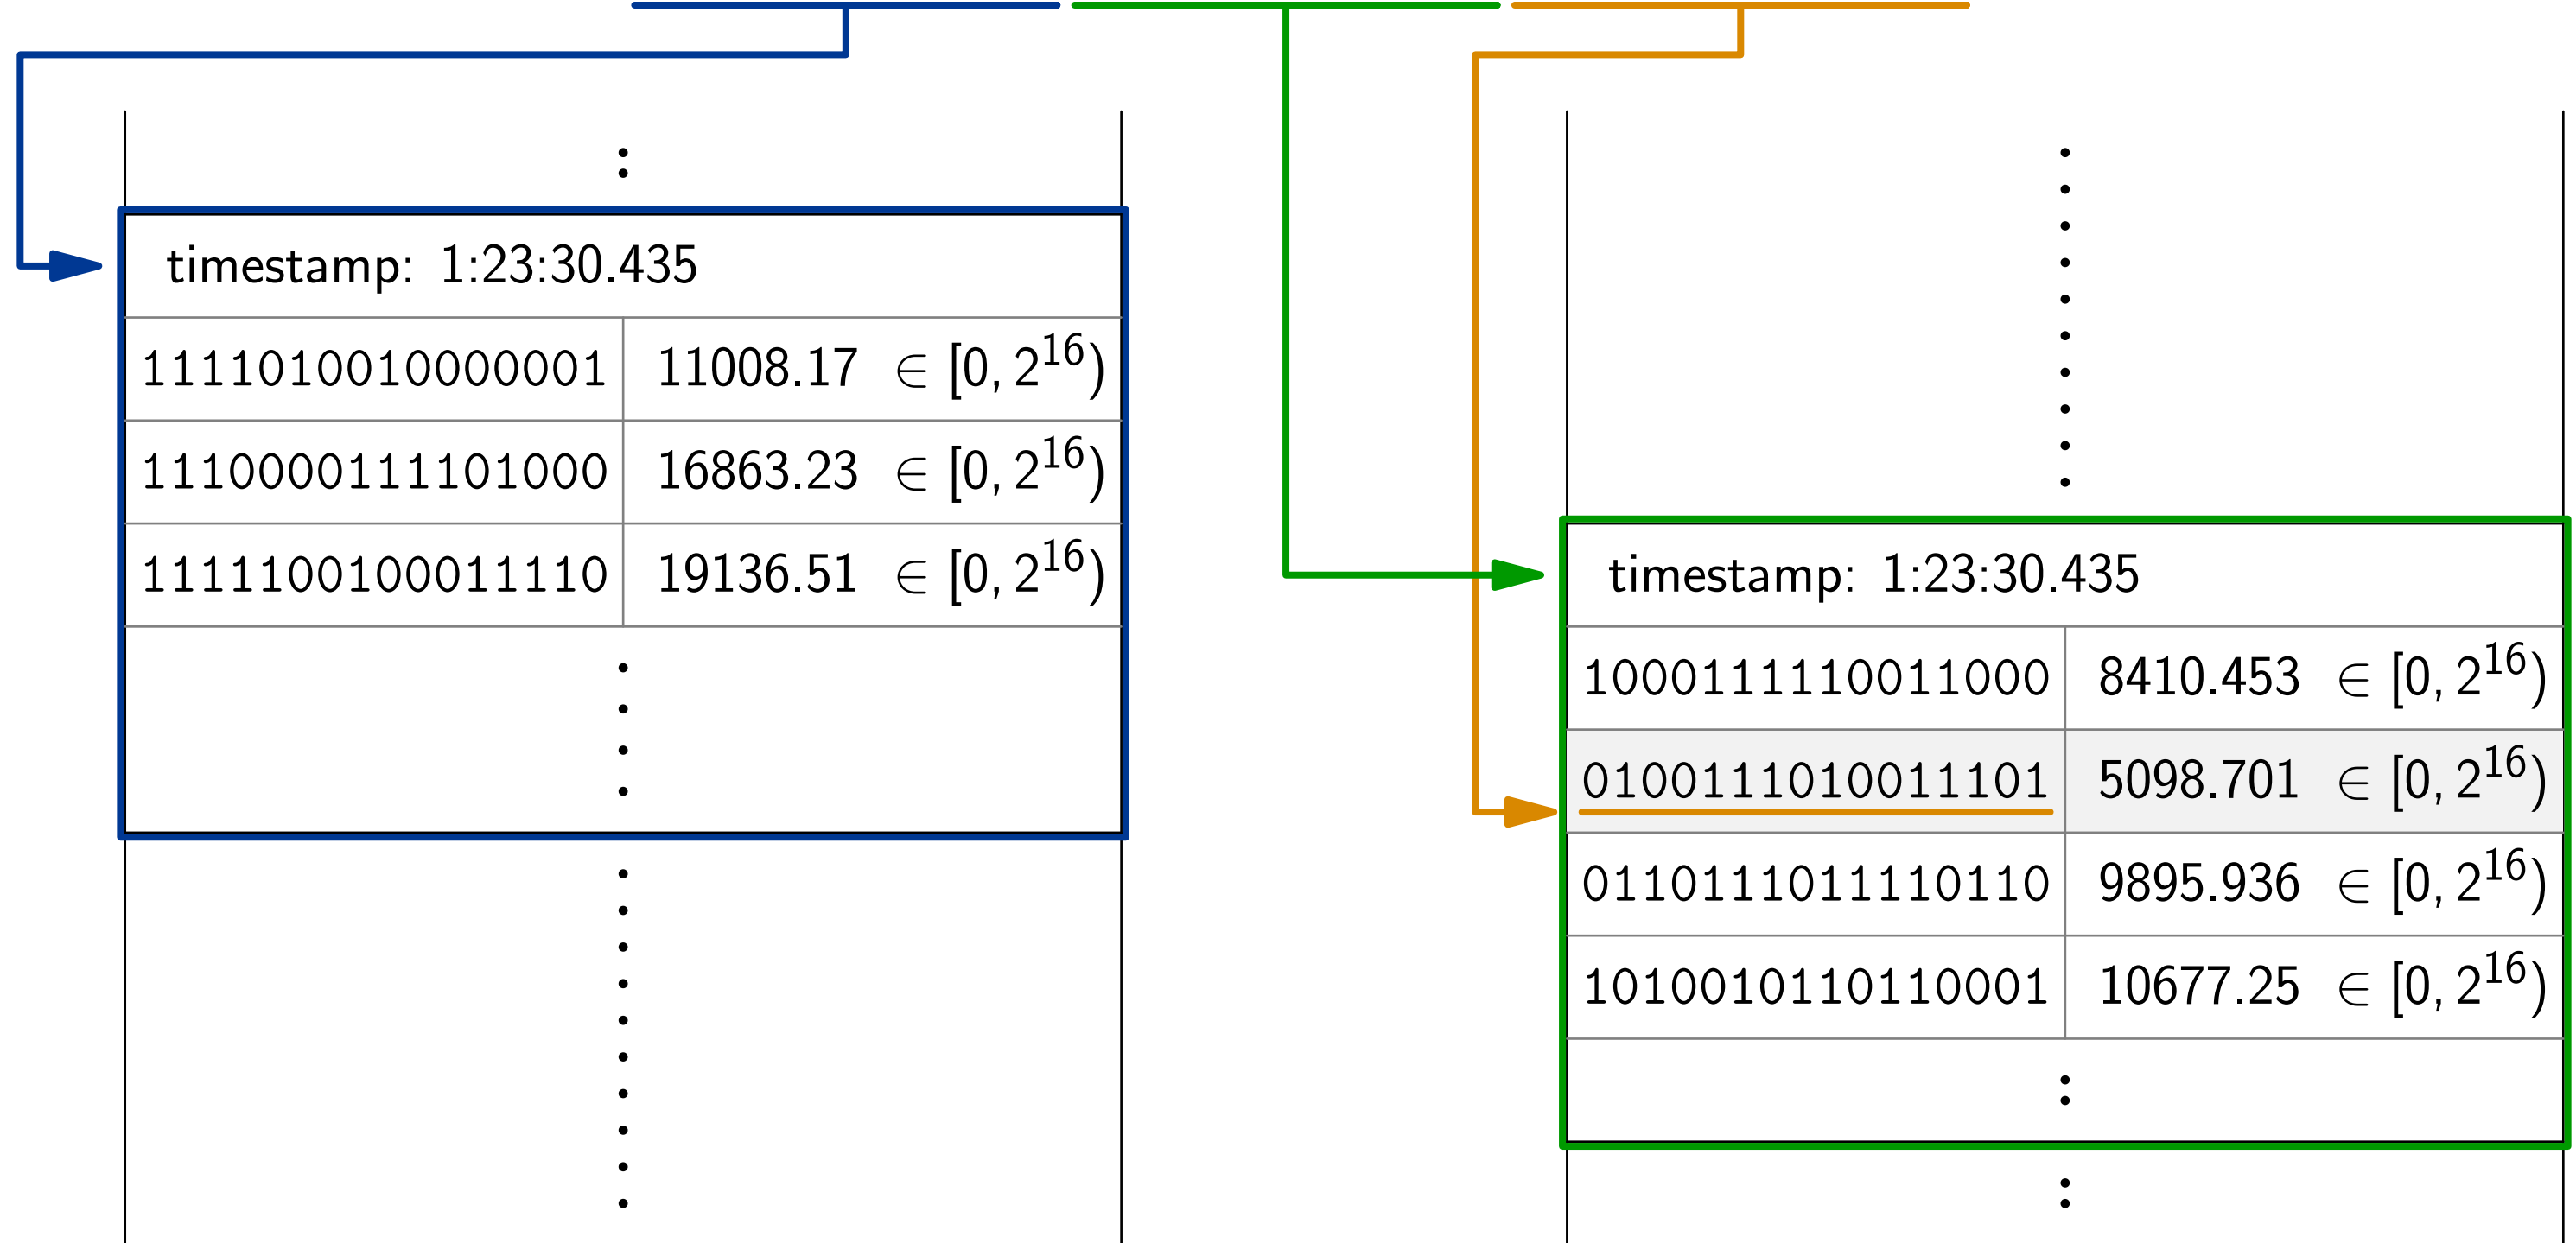
hash(172.16.96.1/32) = 1011110110000110101011000011010100111010011101101001000001011111



- hashing
  - buckets
  - two tables
- evicting
  - normalized limits
  - choosing minimal
  - keeping value
- lazy decay
- memory layout

# Implementation

hash(172.16.96.1/32) = 1011110110000110101011000011010100111010011101101001000001011111



- hashing
  - buckets
  - two tables
- evicting
  - normalized limits
  - choosing minimal
  - keeping value
- lazy decay
- memory layout
  - hashed labels

# Implementation

- hashing
  - buckets
  - two tables
- evicting
  - normalized limits
  - choosing minimal
  - keeping value
- lazy decay
- memory layout
  - hashed labels
  - **prob. rounding**

hash(172.16.96.1/32) = 101111011000011010101100001101010011

:	
timestamp: 1:23:30.435	
1111010010000001	11008.17 ∈ [0, 2 <sup>16</sup> )
1110000111101000	16863.23 ∈ [0, 2 <sup>16</sup> )
1111100100011110	19136.51 ∈ [0, 2 <sup>16</sup> )
:	:
:	:
:	:
:	:
:	:
:	:
:	:

rounding: 5098.7008

16-bit base value      16-bit probability of rounding up (70.08 %)

timestamp: 1:23:30.435	
1000111110011000	8410.453 ∈ [0, 2 <sup>16</sup> )
<u>0100111010011101</u>	5098.701 ∈ [0, 2 <sup>16</sup> )
0110111011110110	9895.936 ∈ [0, 2 <sup>16</sup> )
1010010110110001	10677.25 ∈ [0, 2 <sup>16</sup> )
:	:
:	:

# Implementation

- hashing
  - buckets
  - two tables
- evicting
  - normalized limits
  - choosing minimal
  - keeping value
- lazy decay
- memory layout
  - hashed labels
  - **prob. rounding**

hash(172.16.96.1/32) = 101111011000011010101100001101010011

:		
timestamp: 1:23:30.435		
1111010010000001	[11008.17]	< 2 <sup>16</sup>
1110000111101000	[16863.23]	< 2 <sup>16</sup>
1111100100011110	[19136.51]	< 2 <sup>16</sup>
:		
:		
:		
:		
:		
:		
:		


rounding: 5098.7008

16-bit base value      16-bit probability of rounding up (70.08 %)

timestamp: 1:23:30.435		
1000111110011000	[8410.453]	< 2 <sup>16</sup>
<u>0100111010011101</u>	[5098.701]	< 2 <sup>16</sup>
0110111011110110	[9895.936]	< 2 <sup>16</sup>
1010010110110001	[10677.25]	< 2 <sup>16</sup>
:		
:		

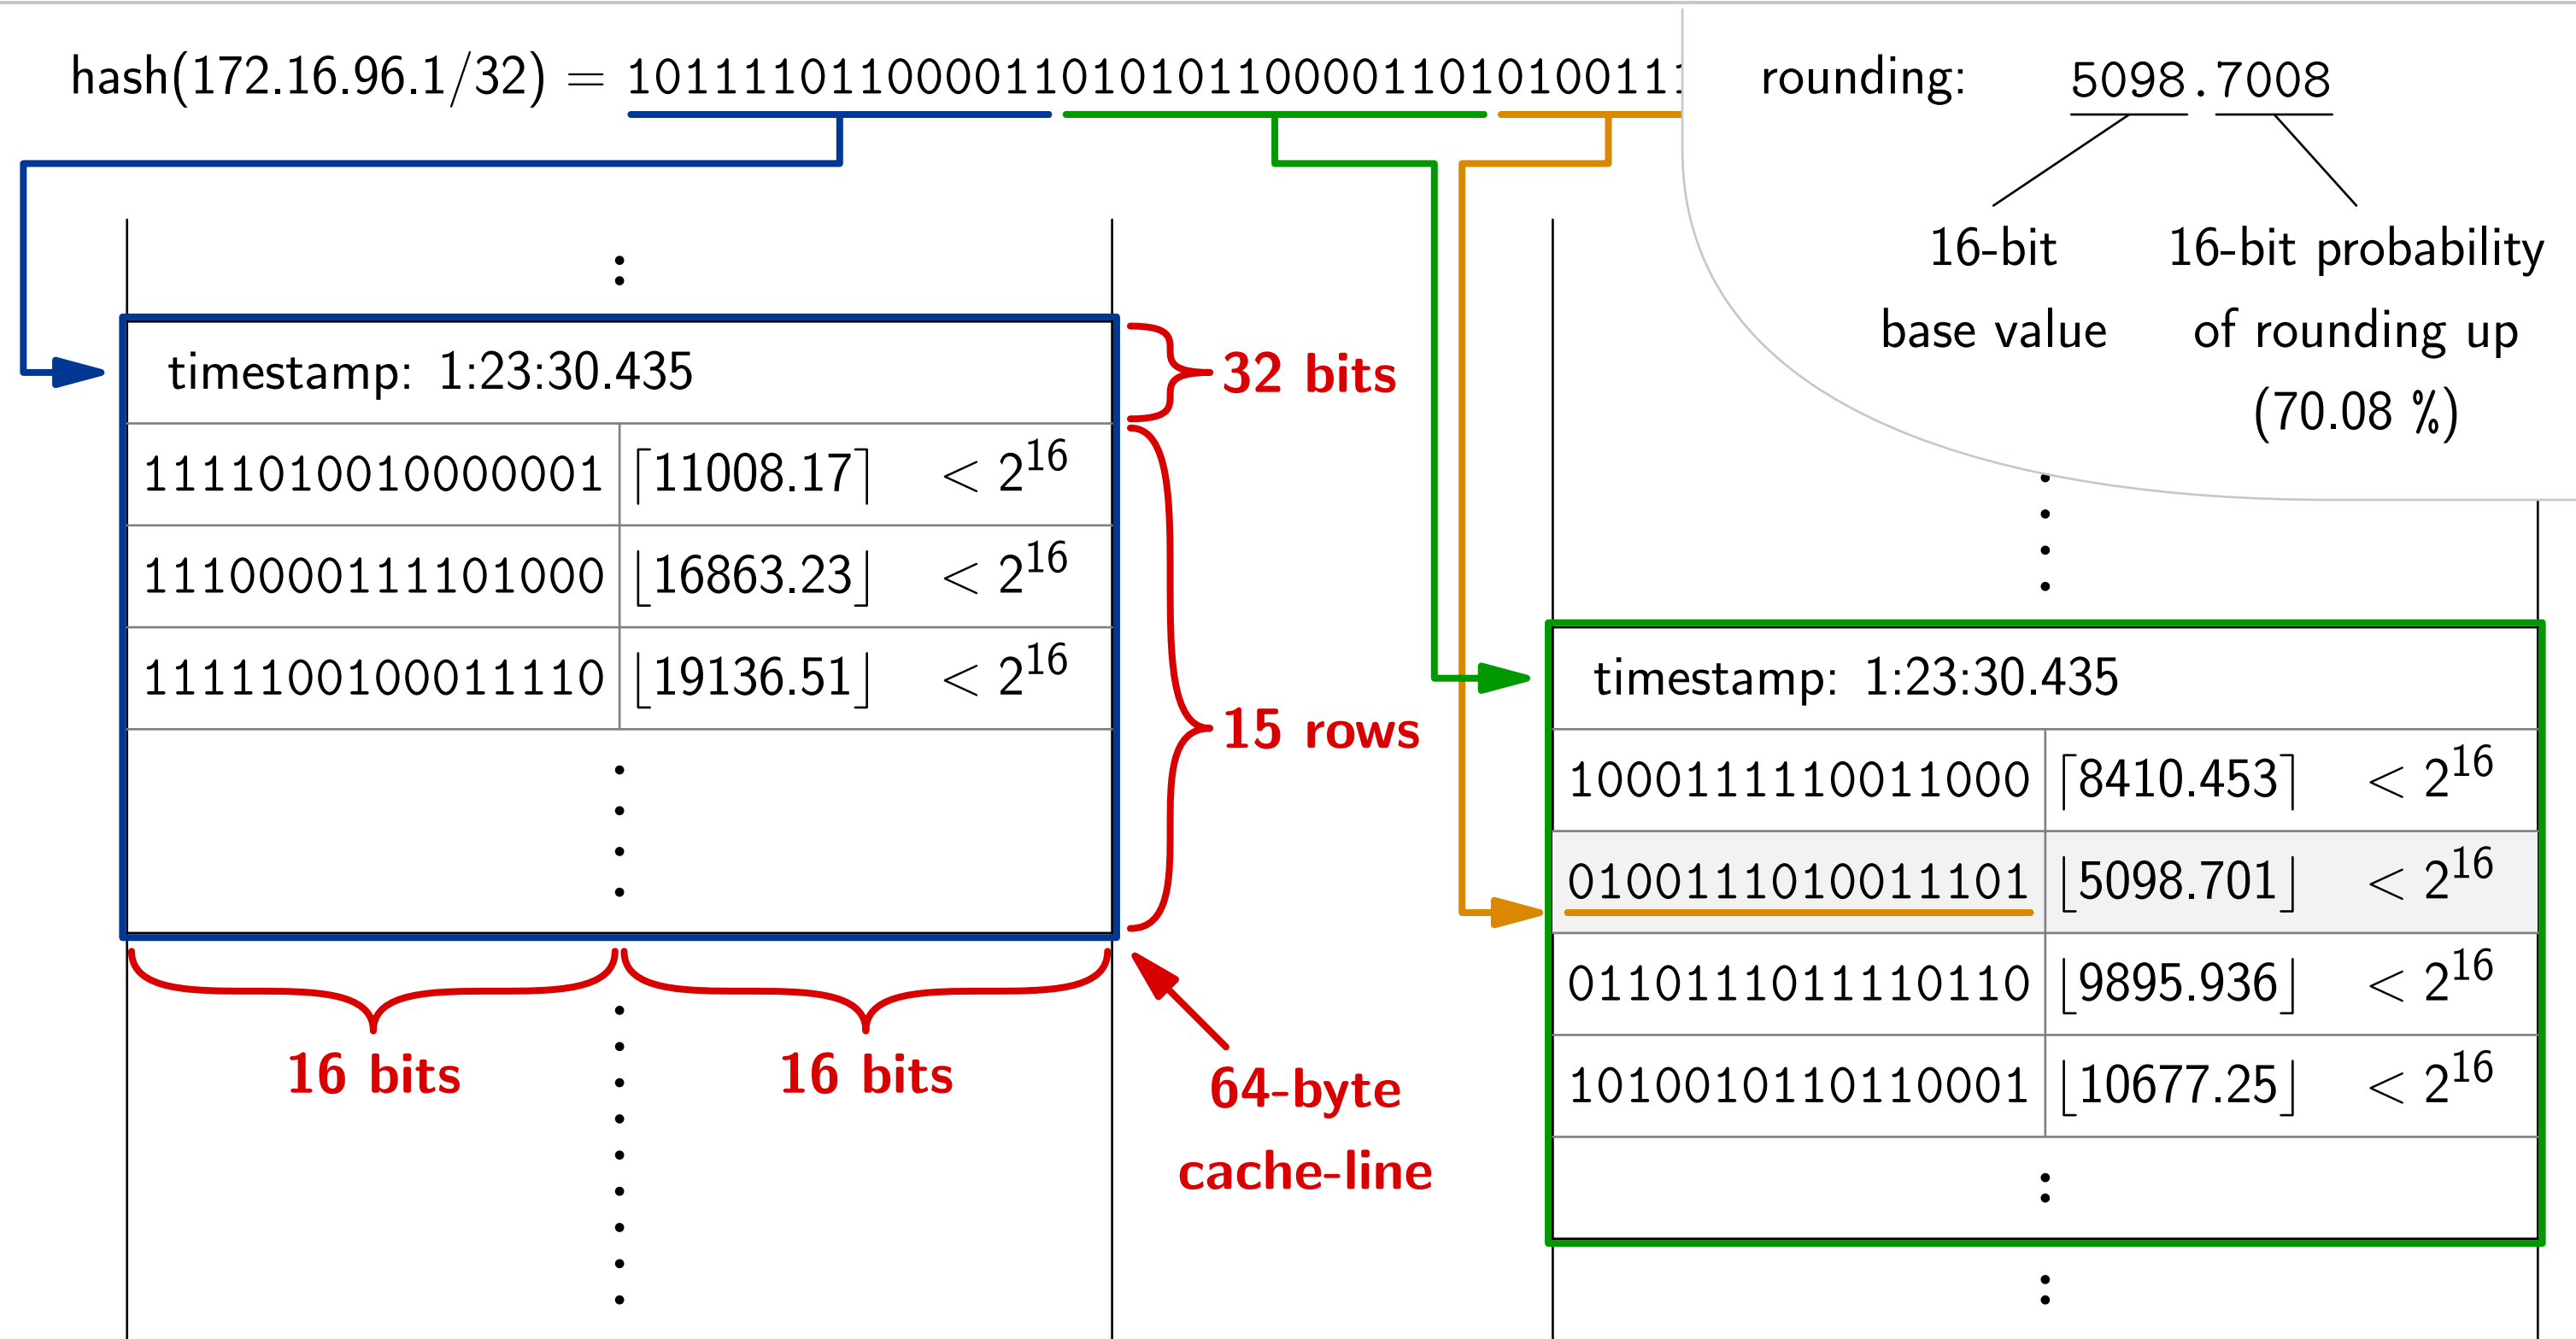
# Implementation

- hashing
  - buckets
  - two tables
- evicting
  - normalized limits
  - choosing minimal
  - keeping value
- lazy decay
- memory layout
  - hashed labels
  - prob. rounding
  - **fit in cache-line**




# Implementation

- hashing
  - buckets
  - two tables
- evicting
  - normalized limits
  - choosing minimal
  - keeping value
- lazy decay
- memory layout
  - hashed labels
  - prob. rounding
  - fit in cache-line
- optimizations
  - prefetching
  - lock-free
  - vectorization





# Summary



- rate-limiting
  - counting UDP queries
  - truncating or dropping
- prioritization
  - measuring time
  - reordering
- counters
  - instant/rate limit
  - exponential decay
  - higher limits for shorter prefixes
- implementation ⇒

hash(172.16.96.1/32) = 101111011000011010101100001101010011...

rounding:  $\frac{5098.7008}{16\text{-bit base value}}$  (16-bit probability of rounding up (70.08 %))

timestamp: 1:23:30.435		
1111010010000001	[11008.17]	$< 2^{16}$
1110000111101000	[16863.23]	$< 2^{16}$
1111100100011110	[19136.51]	$< 2^{16}$
⋮	⋮	⋮

32 bits (for the first row)

15 rows (for the first three rows)

16 bits (for the first half of the row)

16 bits (for the second half of the row)

64-byte cache-line

timestamp: 1:23:30.435		
1000111110011000	[8410.453]	$< 2^{16}$
<u>0100111010011101</u>	[5098.701]	$< 2^{16}$
0110111011110110	[9895.936]	$< 2^{16}$
1010010110110001	[10677.25]	$< 2^{16}$
⋮	⋮	⋮