Cassandra - Intro

Big Data Systems

Dr. Rubi Boim

Cassandra

- (Most popular) NoSQL wide column database

- Open source

Motivation (for this course)

- Main database selected for the course

- Hands on practice

- Platform to learn wide column data modeling

Agenda

- History

- Architecture

- Data model (Original)

- Data model (CQL)

- Examples

Agenda

- **History**

- Architecture

- Data model (Original)

- Data model (CQL)

- Examples

History

- ## Create by Facebook in 2008
  paper: Cassandra - A Decentralized Structured Storage System

  - ## By one of the authors of Dynamo
    Avinash Lakshman was previously at Amazon, then moved to FB

Facebook's motivation

- Design a system to fulfill the storage needs of the **inbox search problem**

  - Billions of writes per day

  - Scale (much more) with the number of users

- A year later Cassandra was used to power much more services within Facebook

Interesting fact

- Facebook released Cassandra as Open source

- Later on Facebook abandoned Cassandra and moved to different technologies
  HBase to power the inbox search

- However Cassandra today is used by endless number of companies
  Apple, Twitter, Netflix, Uber and much more

Initial requirements (2008)

- Access / mange petabytes of data in <u>real time</u>

- Wide applicability

- Highly scalable

- Highly available

- **Completely distributed (P2P)**

- **Tunable consistency**

- **<u>Fast</u>**

Agenda

- History

- **Architecture**

- Data model (Original)

- Data model (CQL)

- Examples

Architecture

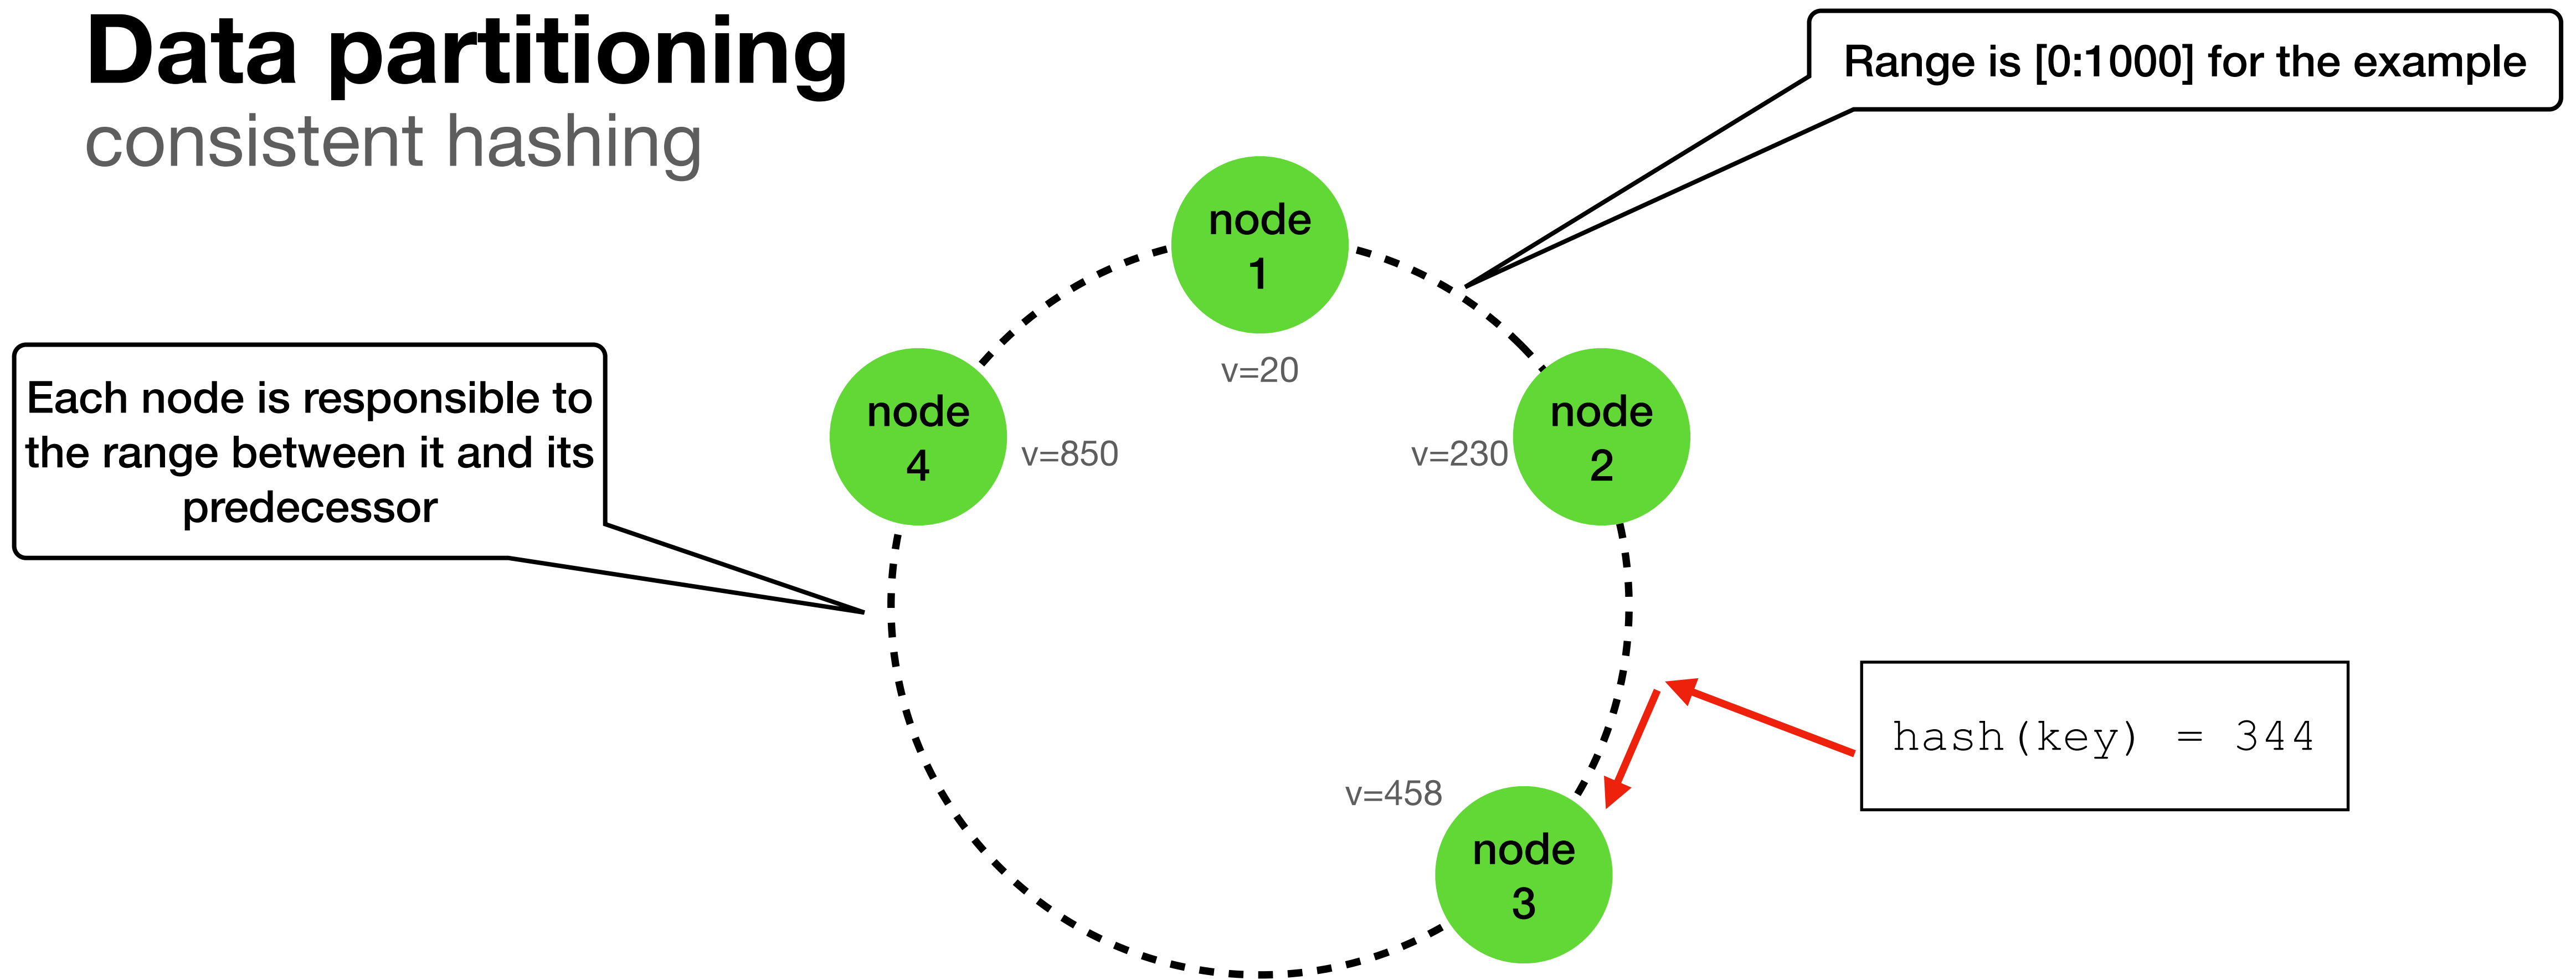**Cassandra: daughter of Dynamo and Bigtable**

- Dynamo
  Distributed storage, replication

- Bigtable
  Data model, storage engine

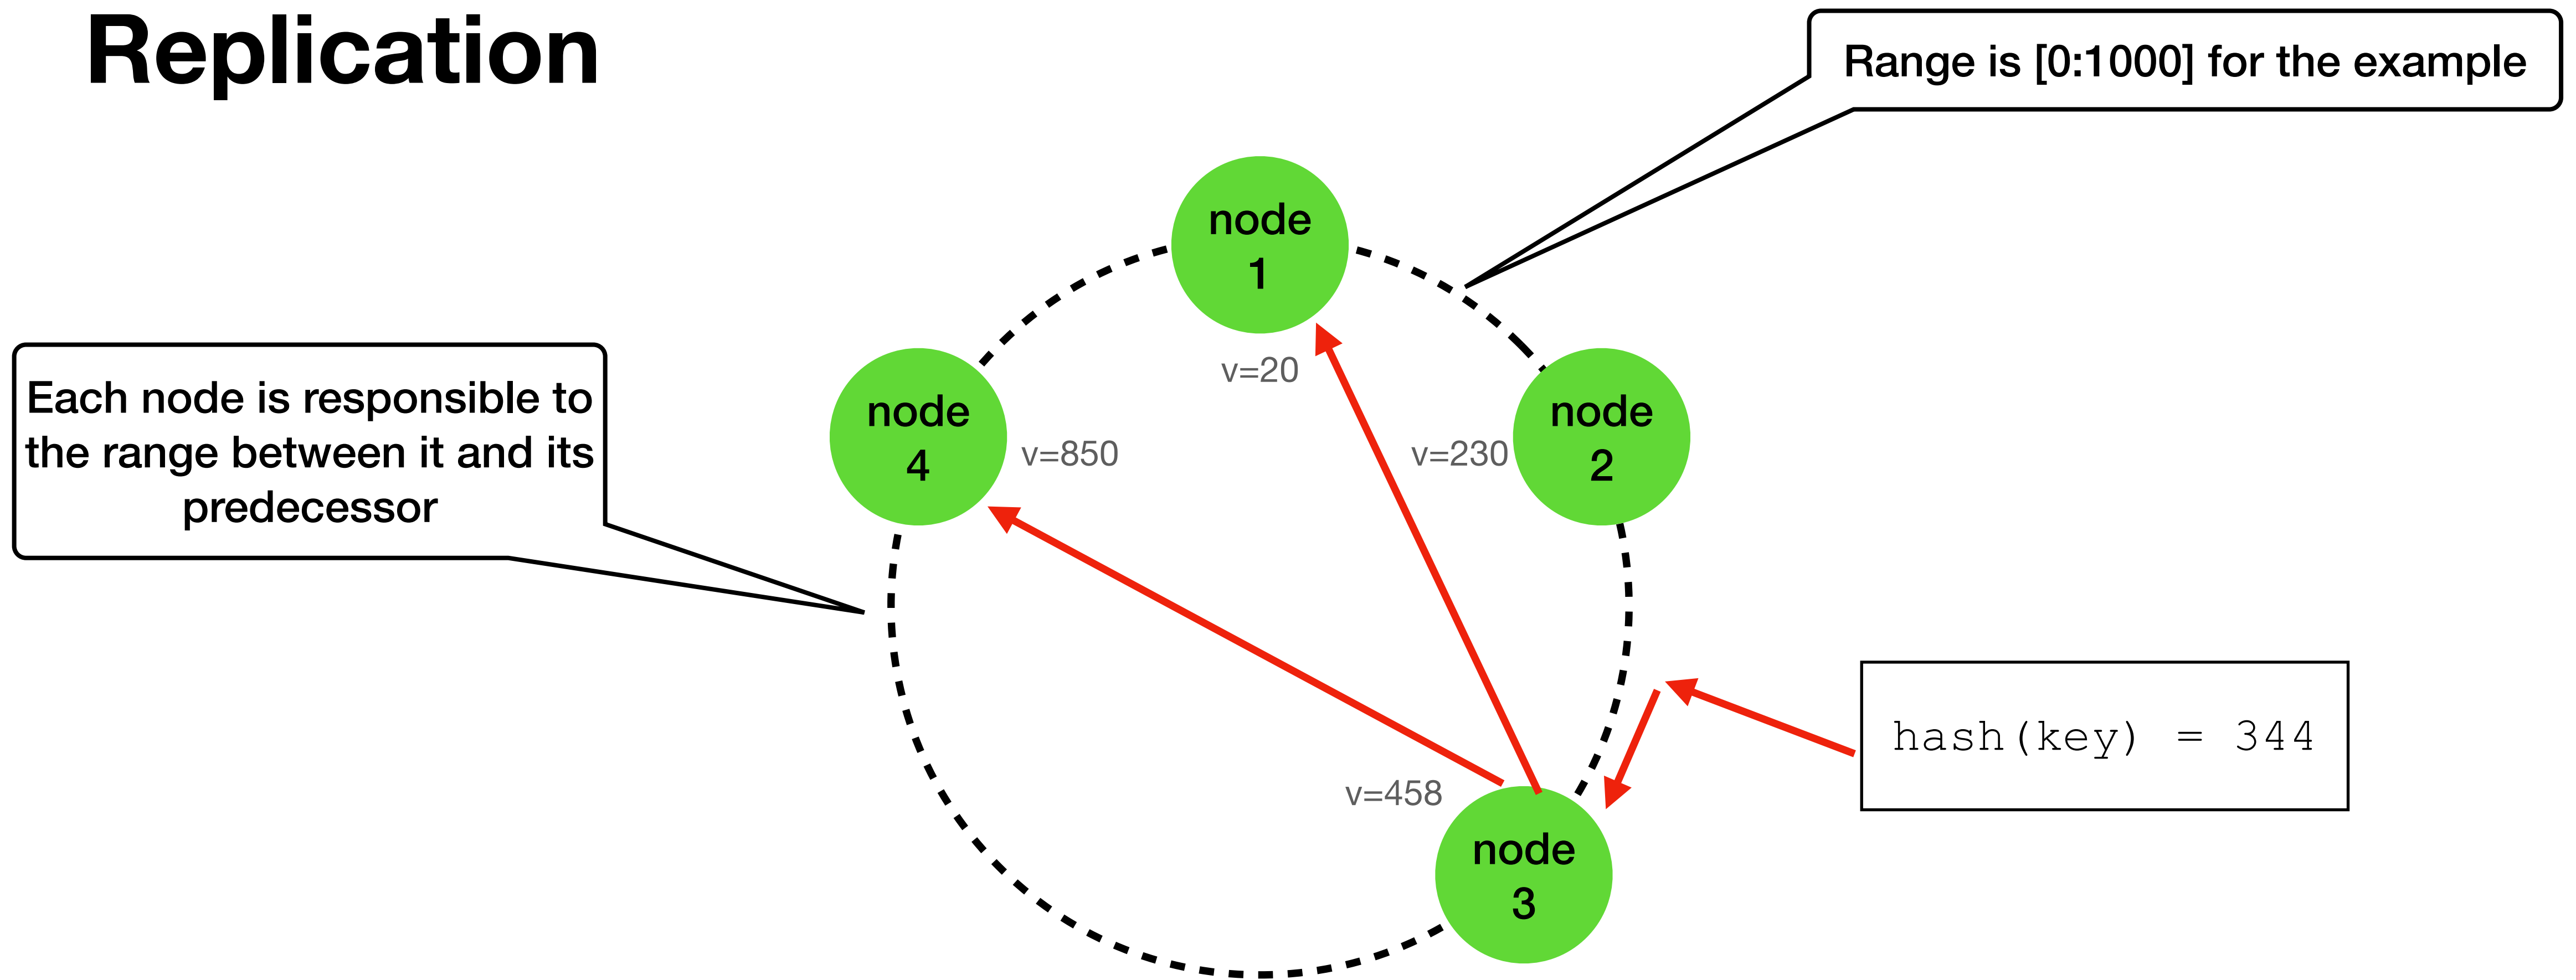Cassandra was designed as a best-in-class combination of both systems

Dynamo —> Cassandra (1)
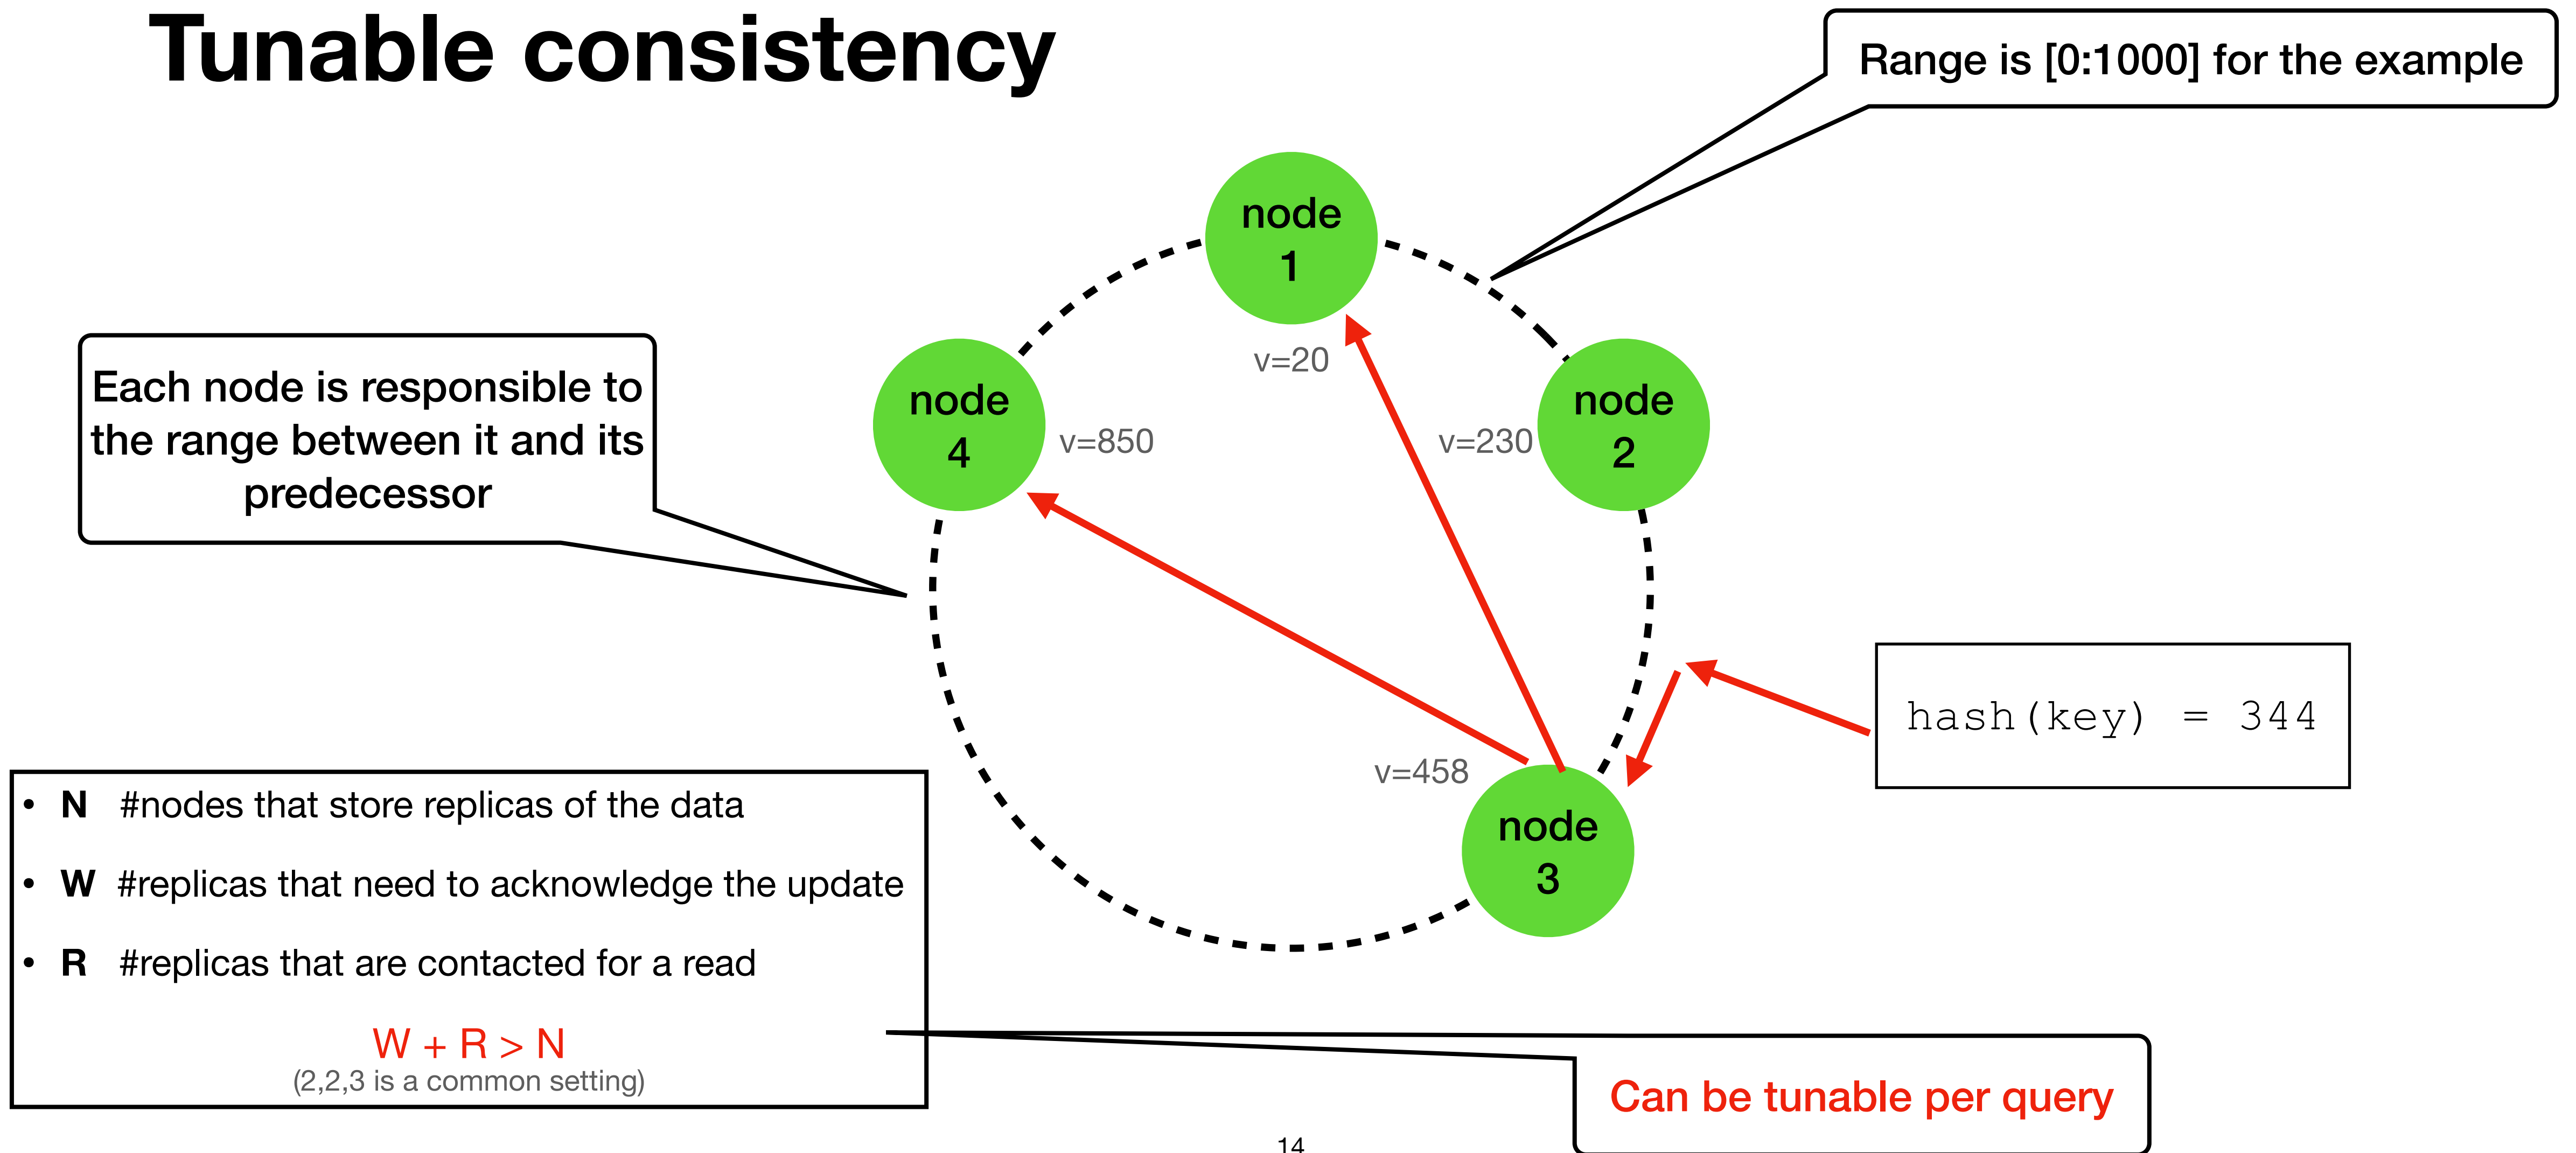
Data partitioning
consistent hashing

Range is [0:1000] for the example

node
1

v=20

node
4

v=850

v=230

node
2

Each node is responsible to the range between it and its predecessor

v=458

hash(key) = 344

node
3

Dynamo —> Cassandra (2)

Replication

Range is [0:1000] for the example

Each node is responsible to the range between it and its predecessor

hash(key) = 344

node 1
v=20

node 4
v=850

node 2
v=230

node 3
v=458

Dynamo —> Cassandra (3)

Tunable consistency

Range is [0:1000] for the example

Each node is responsible to the range between it and its predecessor

node 1
v=20

node 4
v=850

node 2
v=230

v=458

node 3

hash(key) = 344

- **N** #nodes that store replicas of the data

- **W** #replicas that need to acknowledge the update

- **R** #replicas that are contacted for a read

W + R > N
(2,2,3 is a common setting)
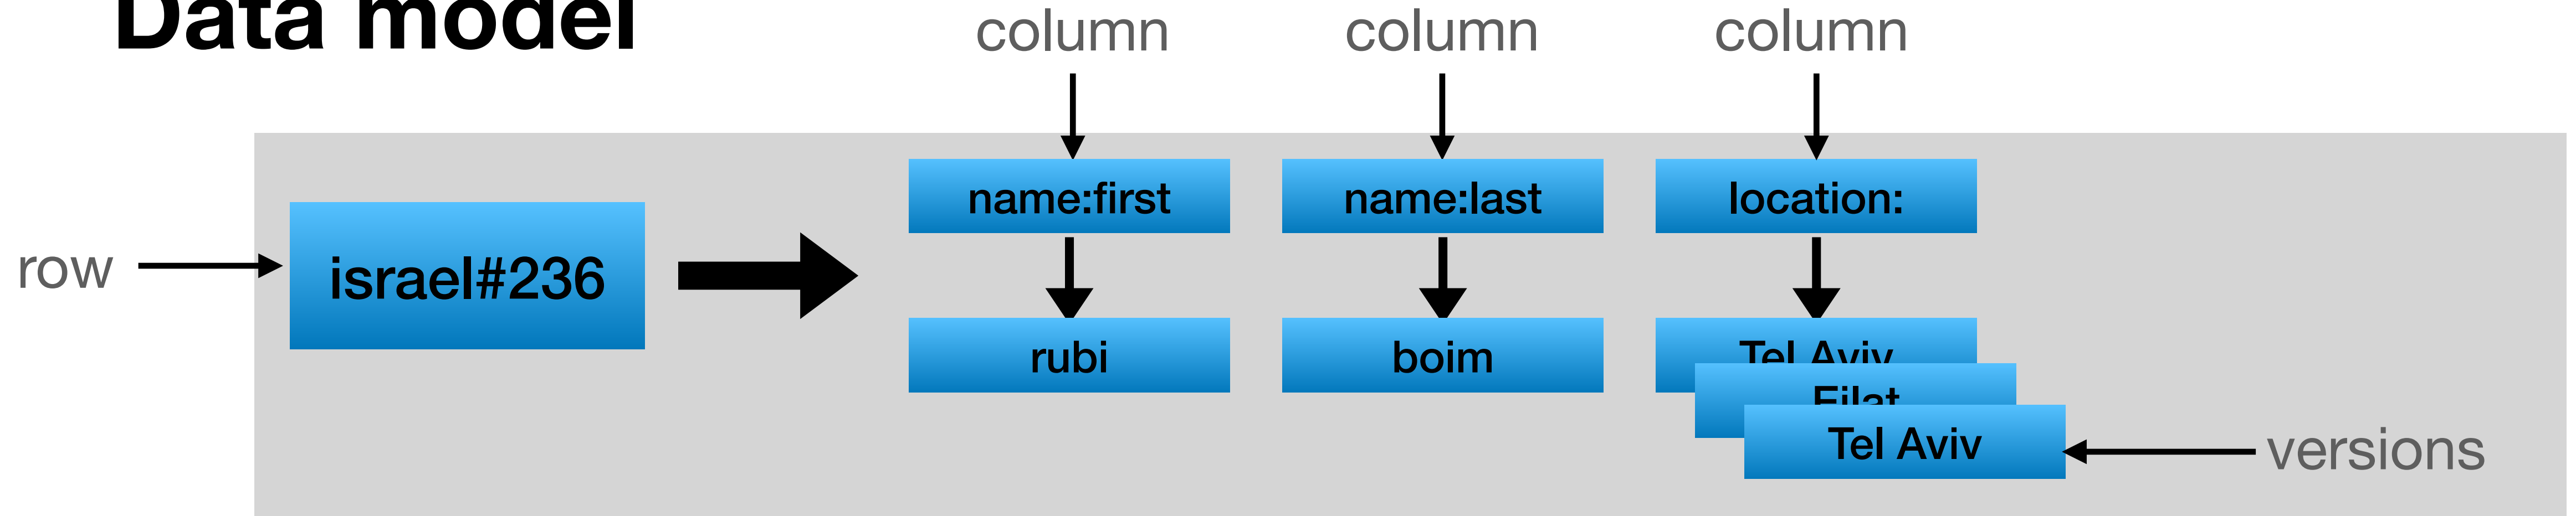
Can be tunable per query

Dynamo —> Cassandra (4)

- ## Distributed cluster membership
  fully distributed P2P, masterless

- ## Failure detection
  gossip

- ## Incremental scale-out on commodity hardware
  different types of hardware size
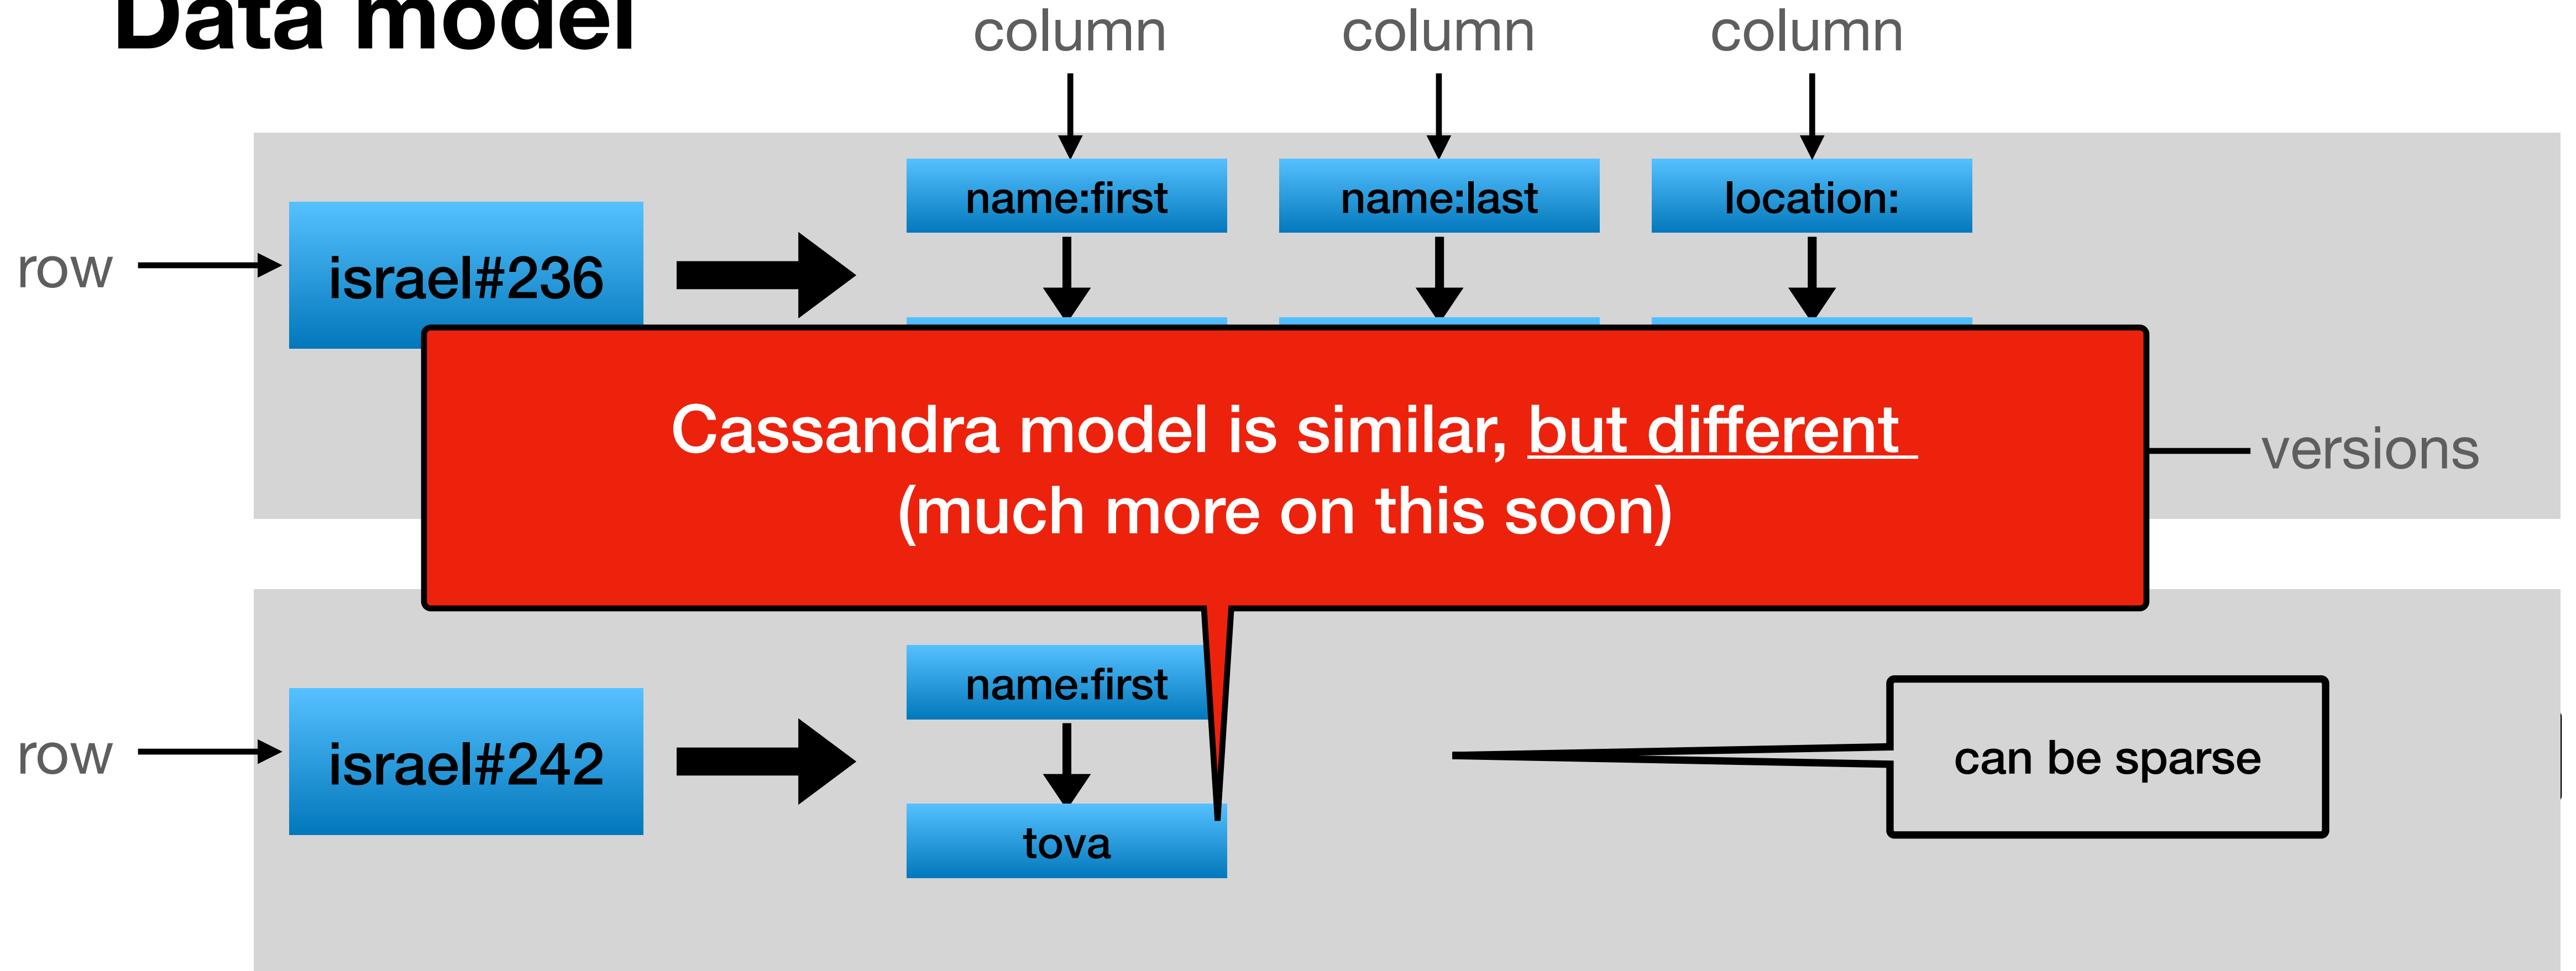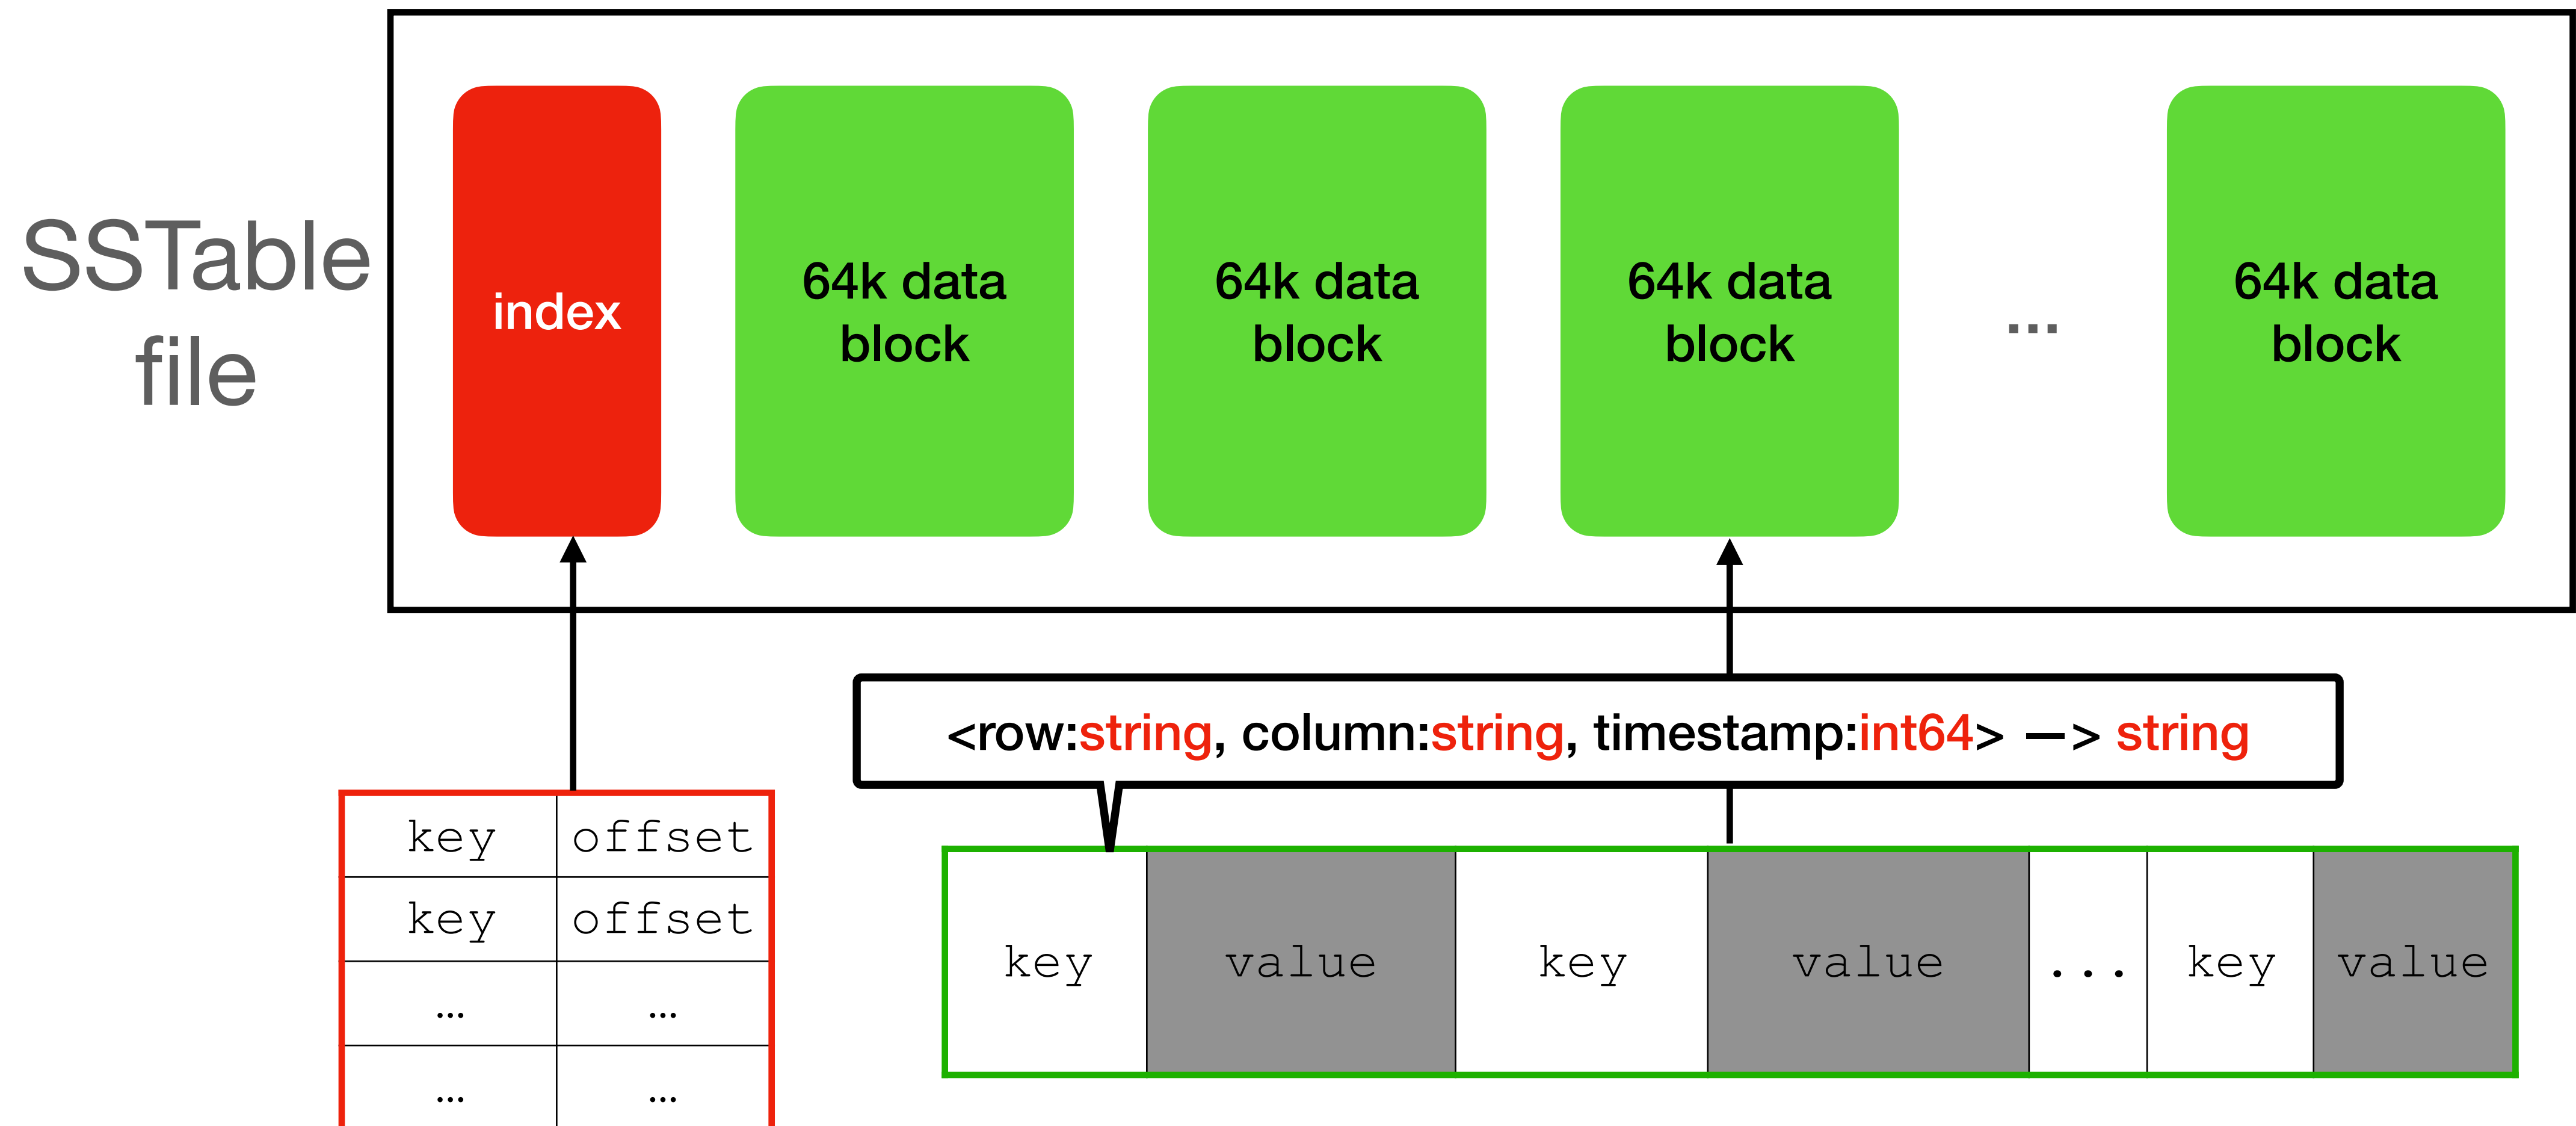
Bigtable —> Cassandra (1)

Data model

column   column   column

row   israel#236

name:first → rubi

name:last → boim

location: → Tel Aviv / Eilat / Tel Aviv   versions

row   israel#242

name:first → tova

can be sparse

<row:string, column:string, timestamp:int64> —> string

Bigtable —> Cassandra (1)

Data model

column   column   column

name:first   name:last   location:

row → israel#236 ➡

Cassandra model is similar, <u>but different</u>
(much more on this soon)

versions

row → israel#242 ➡   name:first
                          ↓
                        tova

can be sparse
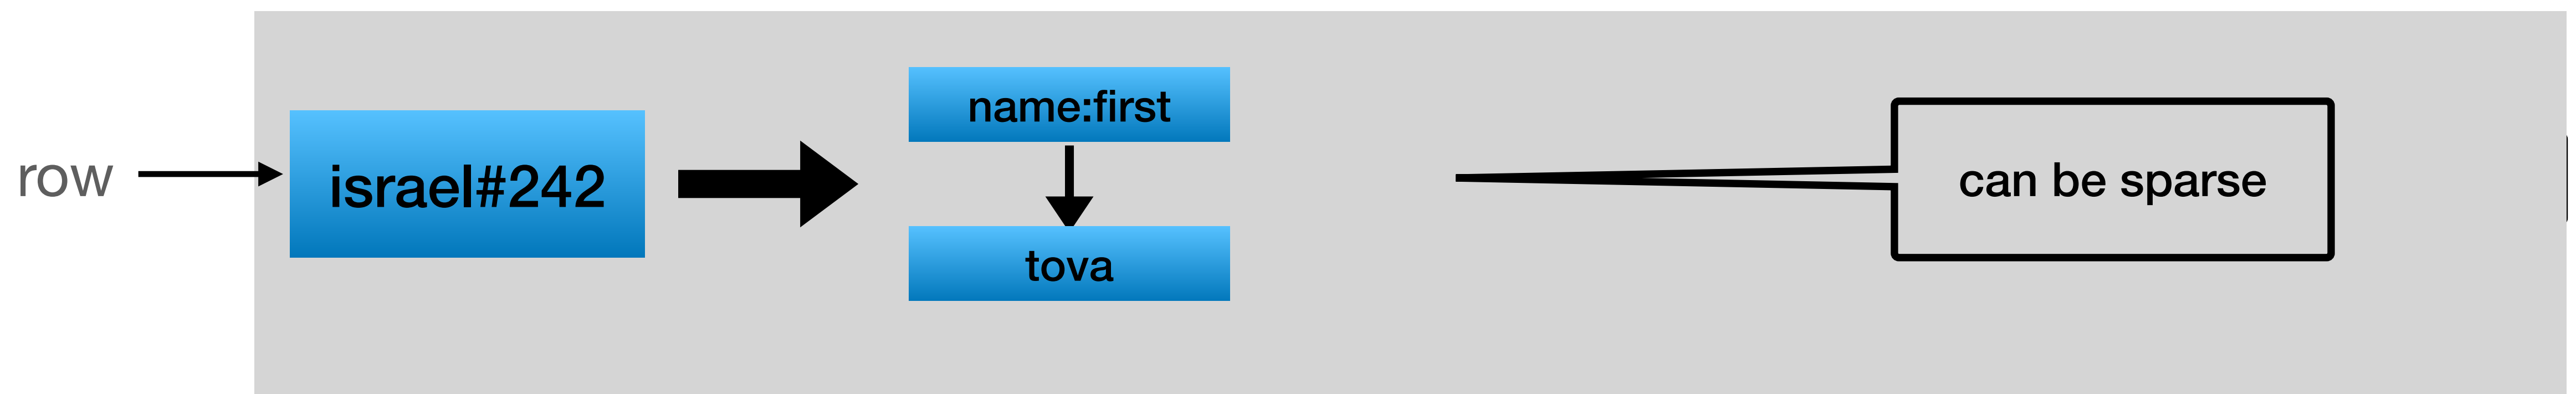
<row:string, column:string, timestamp:int64> —> string

Bigtable —> Cassandra (2)
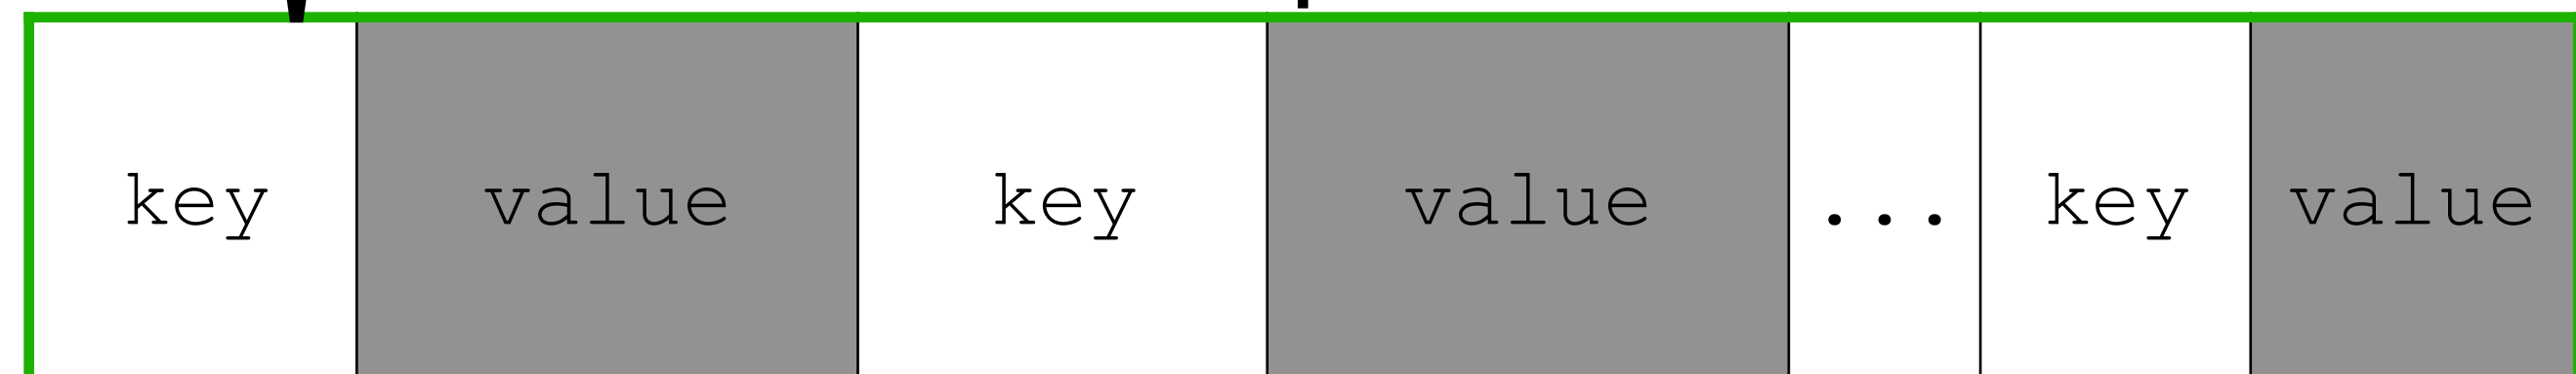
Storage engine - SSTables

Bigtable —> Cassandra (2)

Storag

SSTable
file

Writes are super fast because they do NOT require a read.

Even If only a single column is added/deleted

key	offset
key	offset
…	…
…	…

<row:string, column:string, timestamp:int64> —> string

key	value	key	value	…	key	value

Bigtable —> Cassandra (3)

Storage engine - the rest of the gang

- CommitLog

- Memtable

- SSTable (and compactions)

- Bloom filters

# Cassandra ring

ApacheCassandra™= NoSQL Distributed Database

NODE

NODE

NODE

NODE

NODE

NODE

NODE

NODE

cassandra

1 Installation = 1 NODE
✓ Capacity + 1TB
✓ Throughput: 3000 Tx/sec/core

DataCenter | Ring

Communication:
✓ Gossiping

Cassandra - partitions
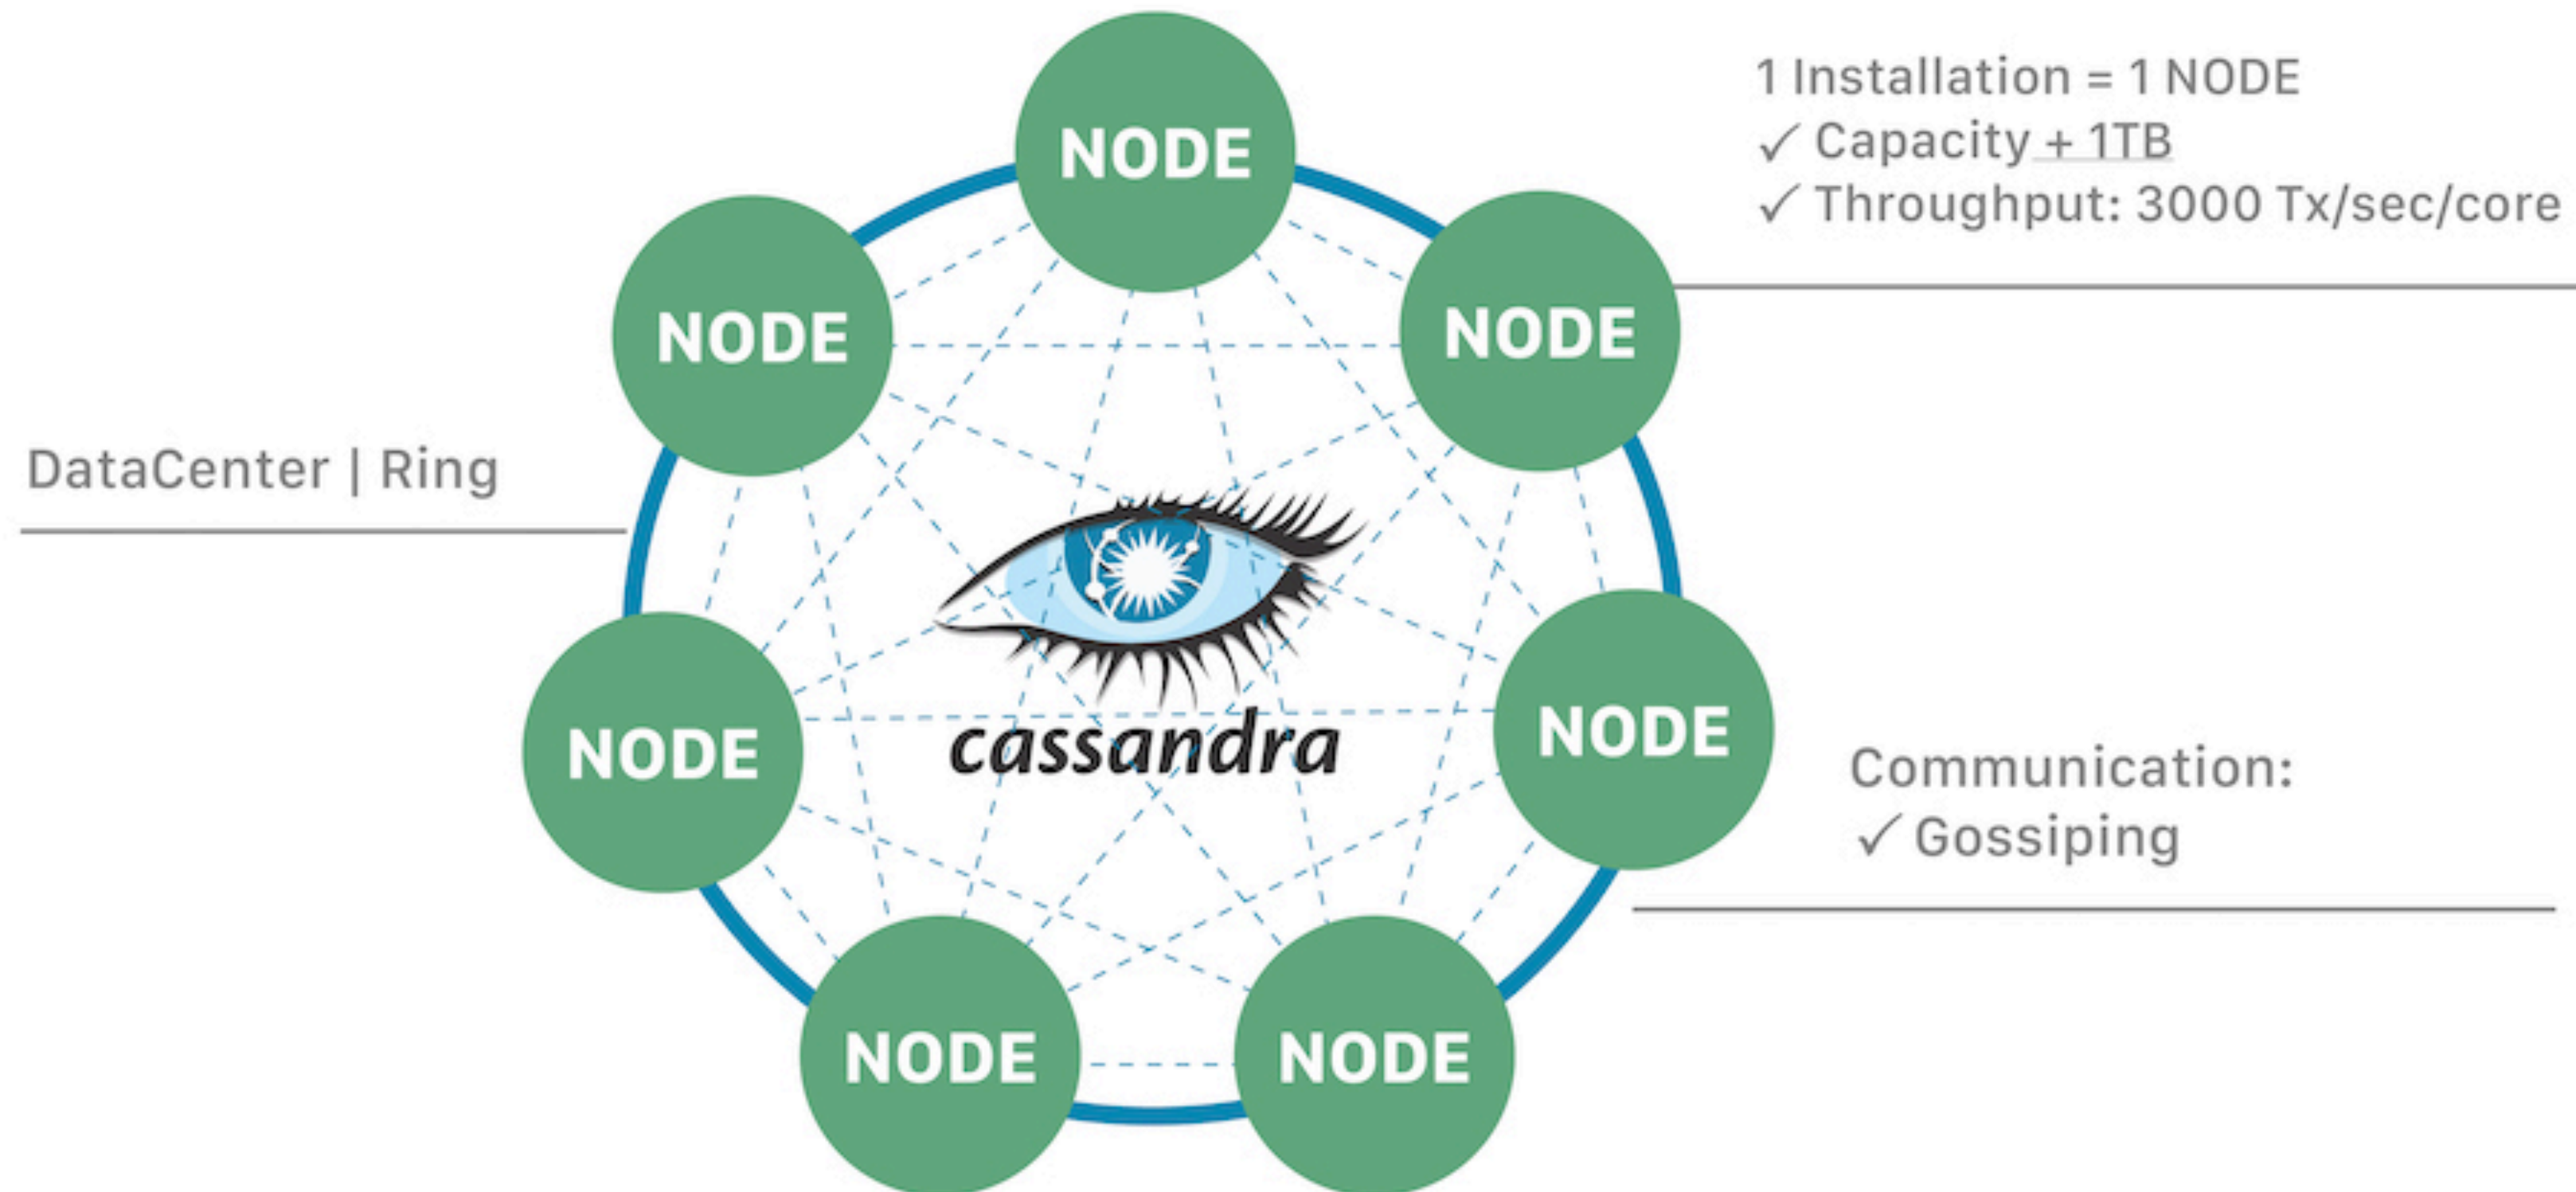
Partitions are similar to ROWs in Bigtable
(In a few slides…)

COUNTRY	CITY	POPULATION
USA	New York	8.000.000
USA	Los Angeles	4.000.000
FR	Paris	2.230.000
DE	Berlin	3.350.000
UK	London	9.200.000
AU	Sydney	4.900.000
DE	Nuremberg	500.000
CA	Toronto	6.200.000
CA	Montreal	4.200.000
FR	Toulouse	1.100.000
JP	Tokyo	37.430.000
IN	Mumbai	20.200.000

Partition Key

USA	New York	8.000.000
USA	Los Angeles	4.000.000

DE	Berlin	3.350.000
DE	Nuremberg	500.000

FR	Paris	2.230.000
FR	Toulouse	1.100.000

UK	London	9.200.000

JP	Tokyo	37.430.000

AU	Sydney	4.900.000
IN	Mumbai	20.200.000

CA	Toronto	6.200.000
CA	Montreal	4.200.000

22

Cassandra - partitions
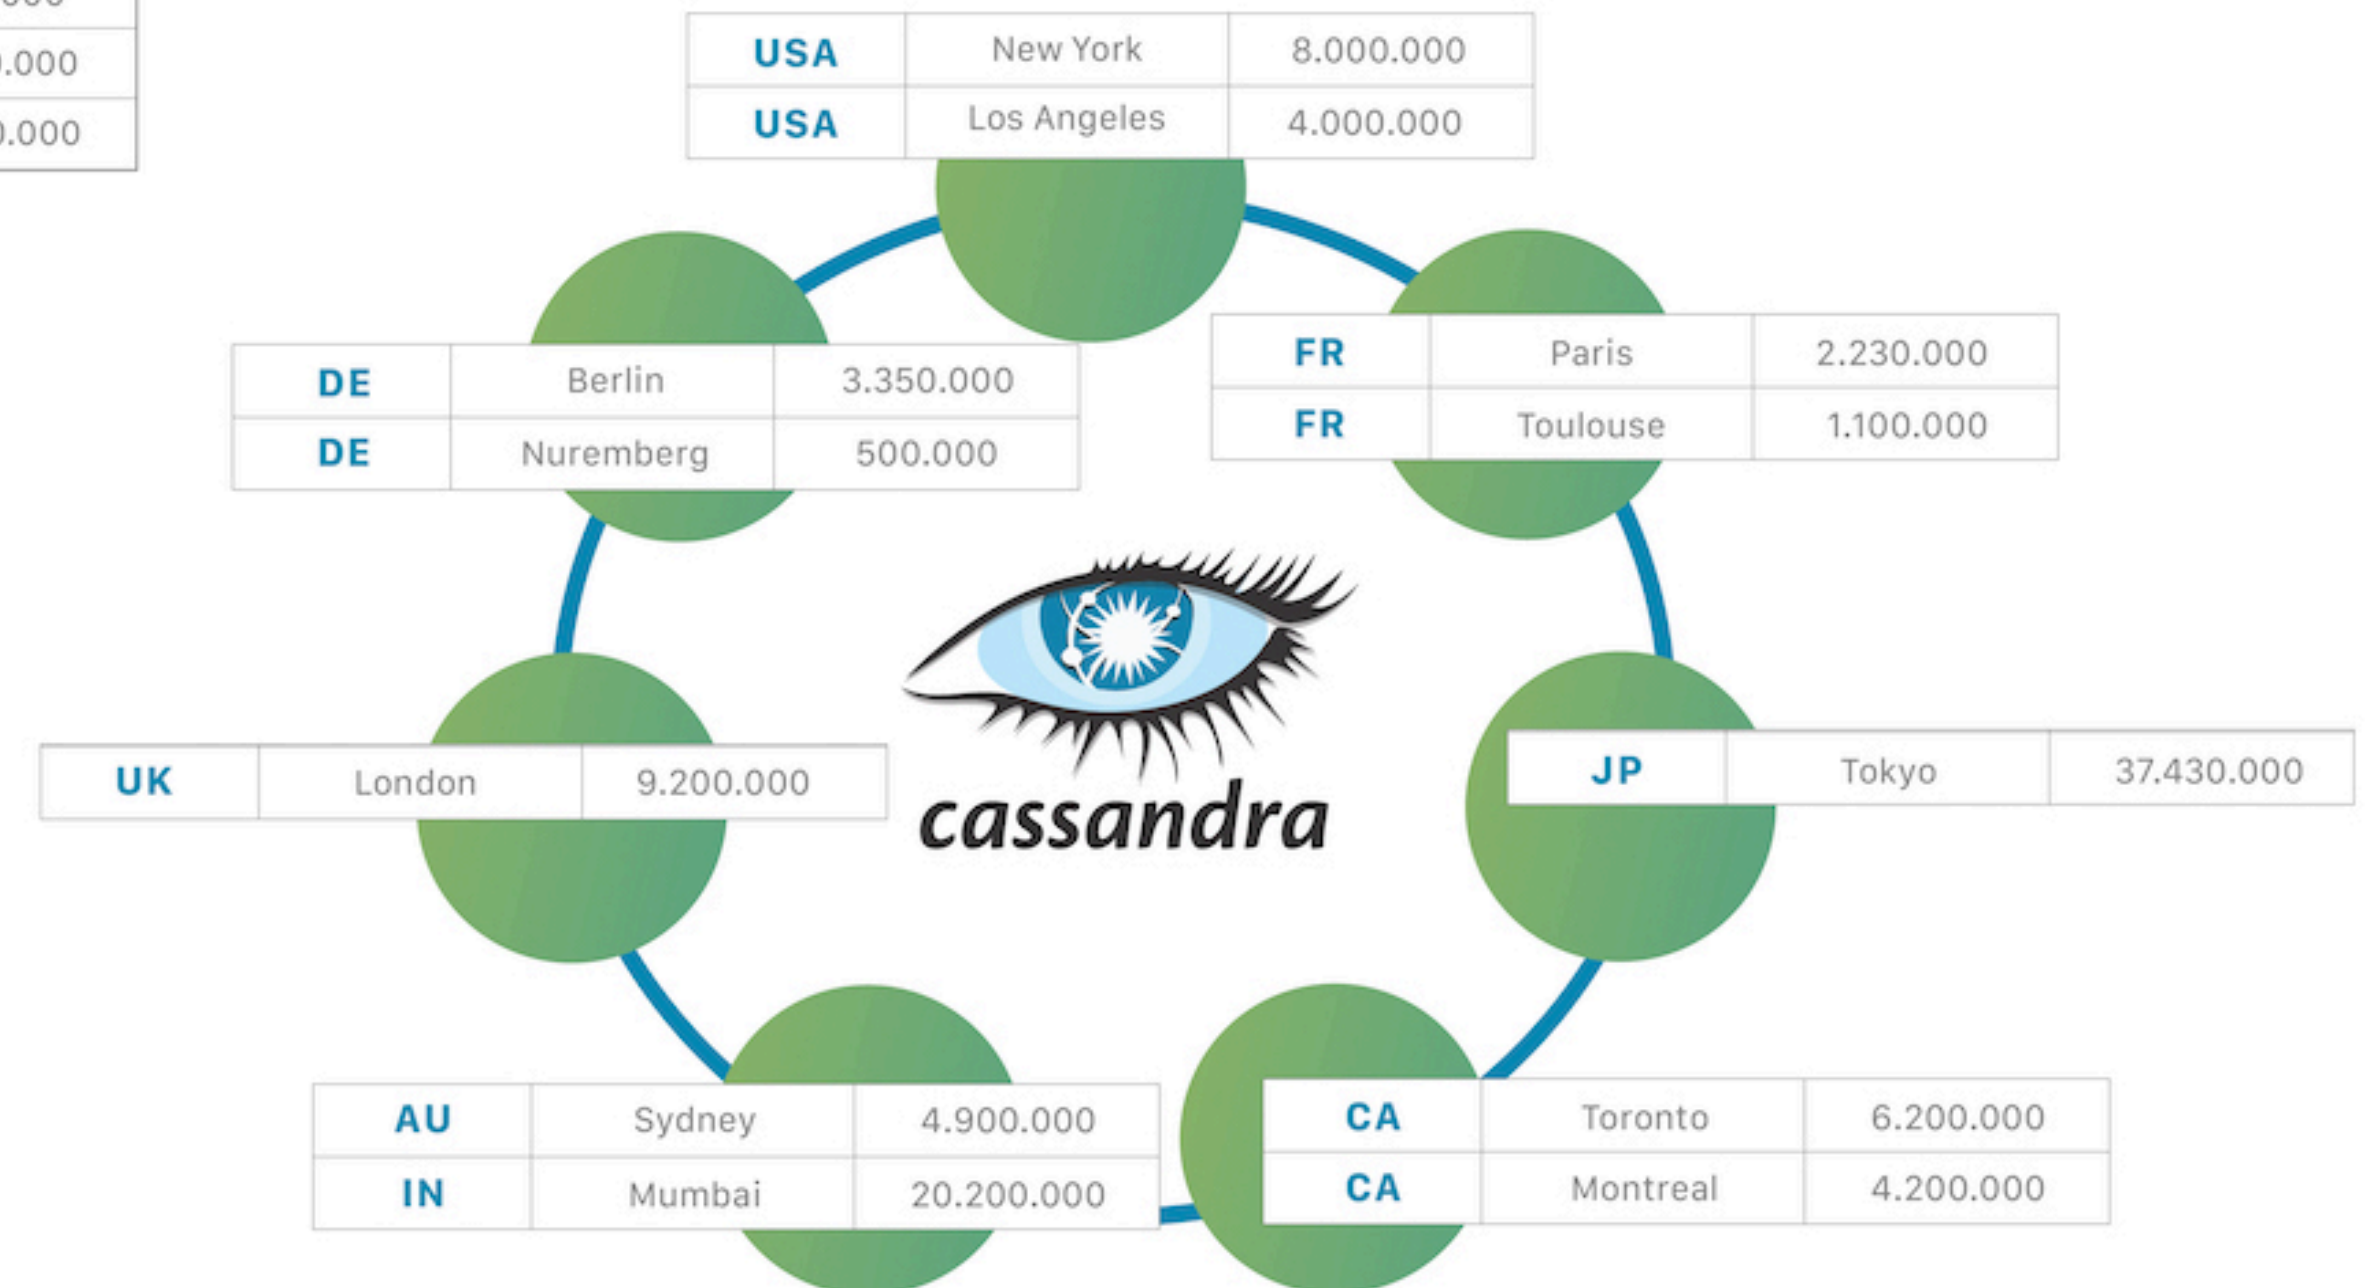
Cassandra - globally distributed
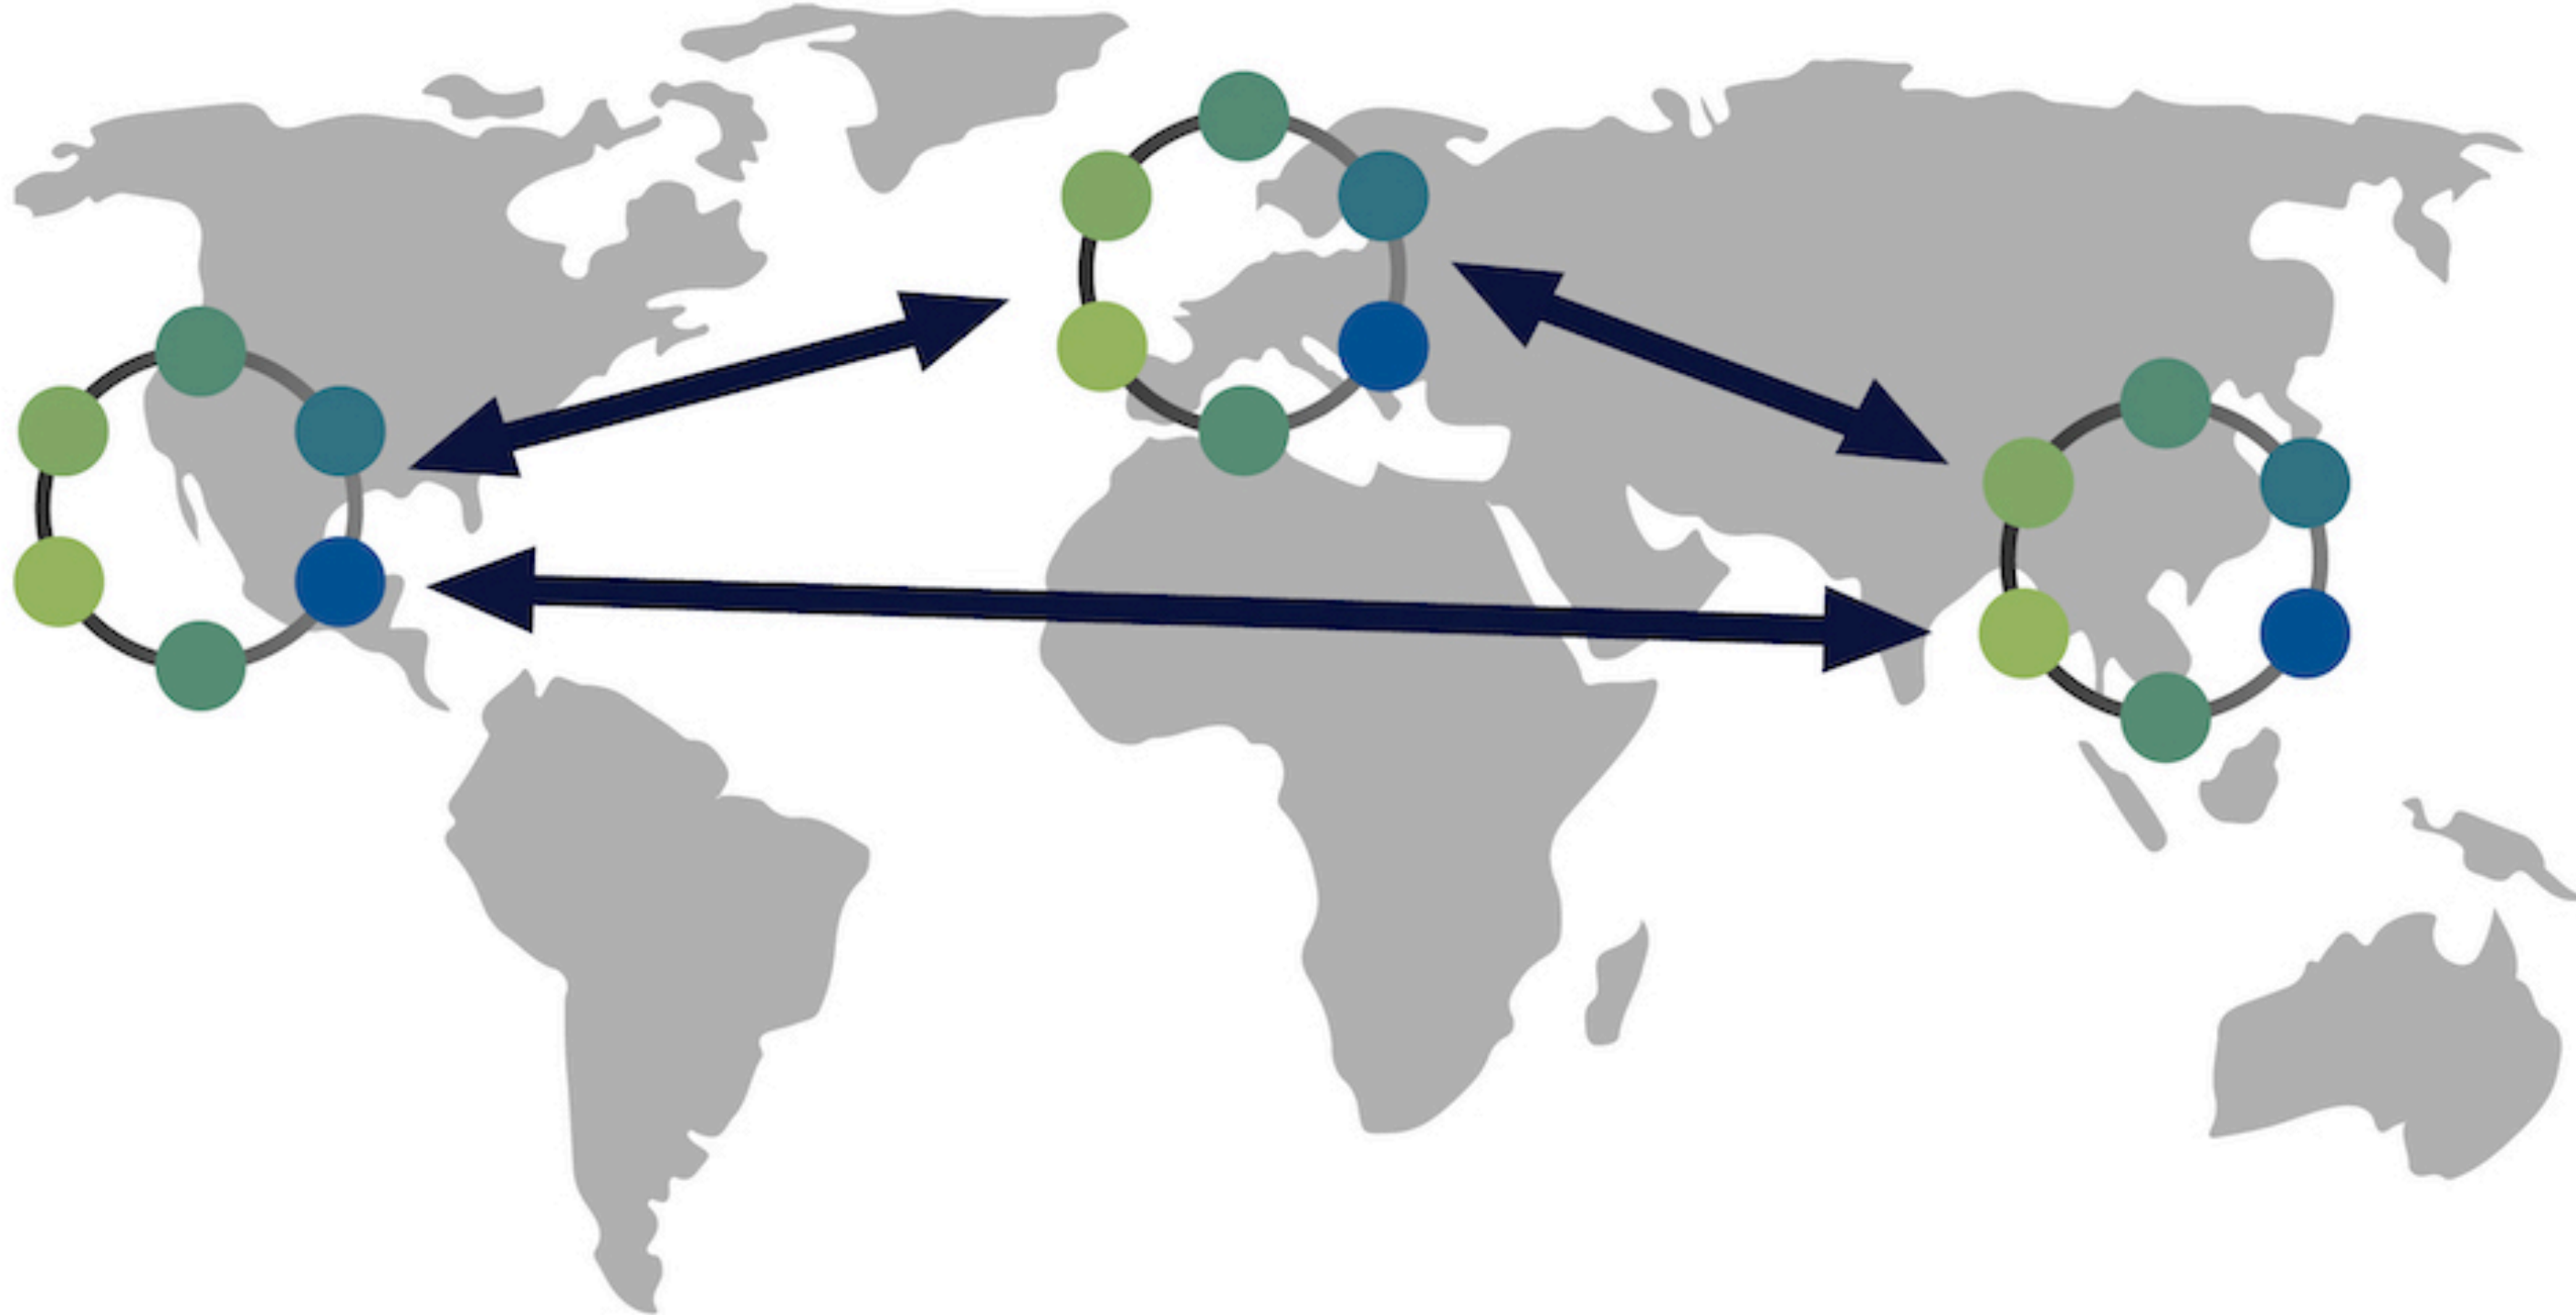
Cassandra - tunable consistency

Consistency
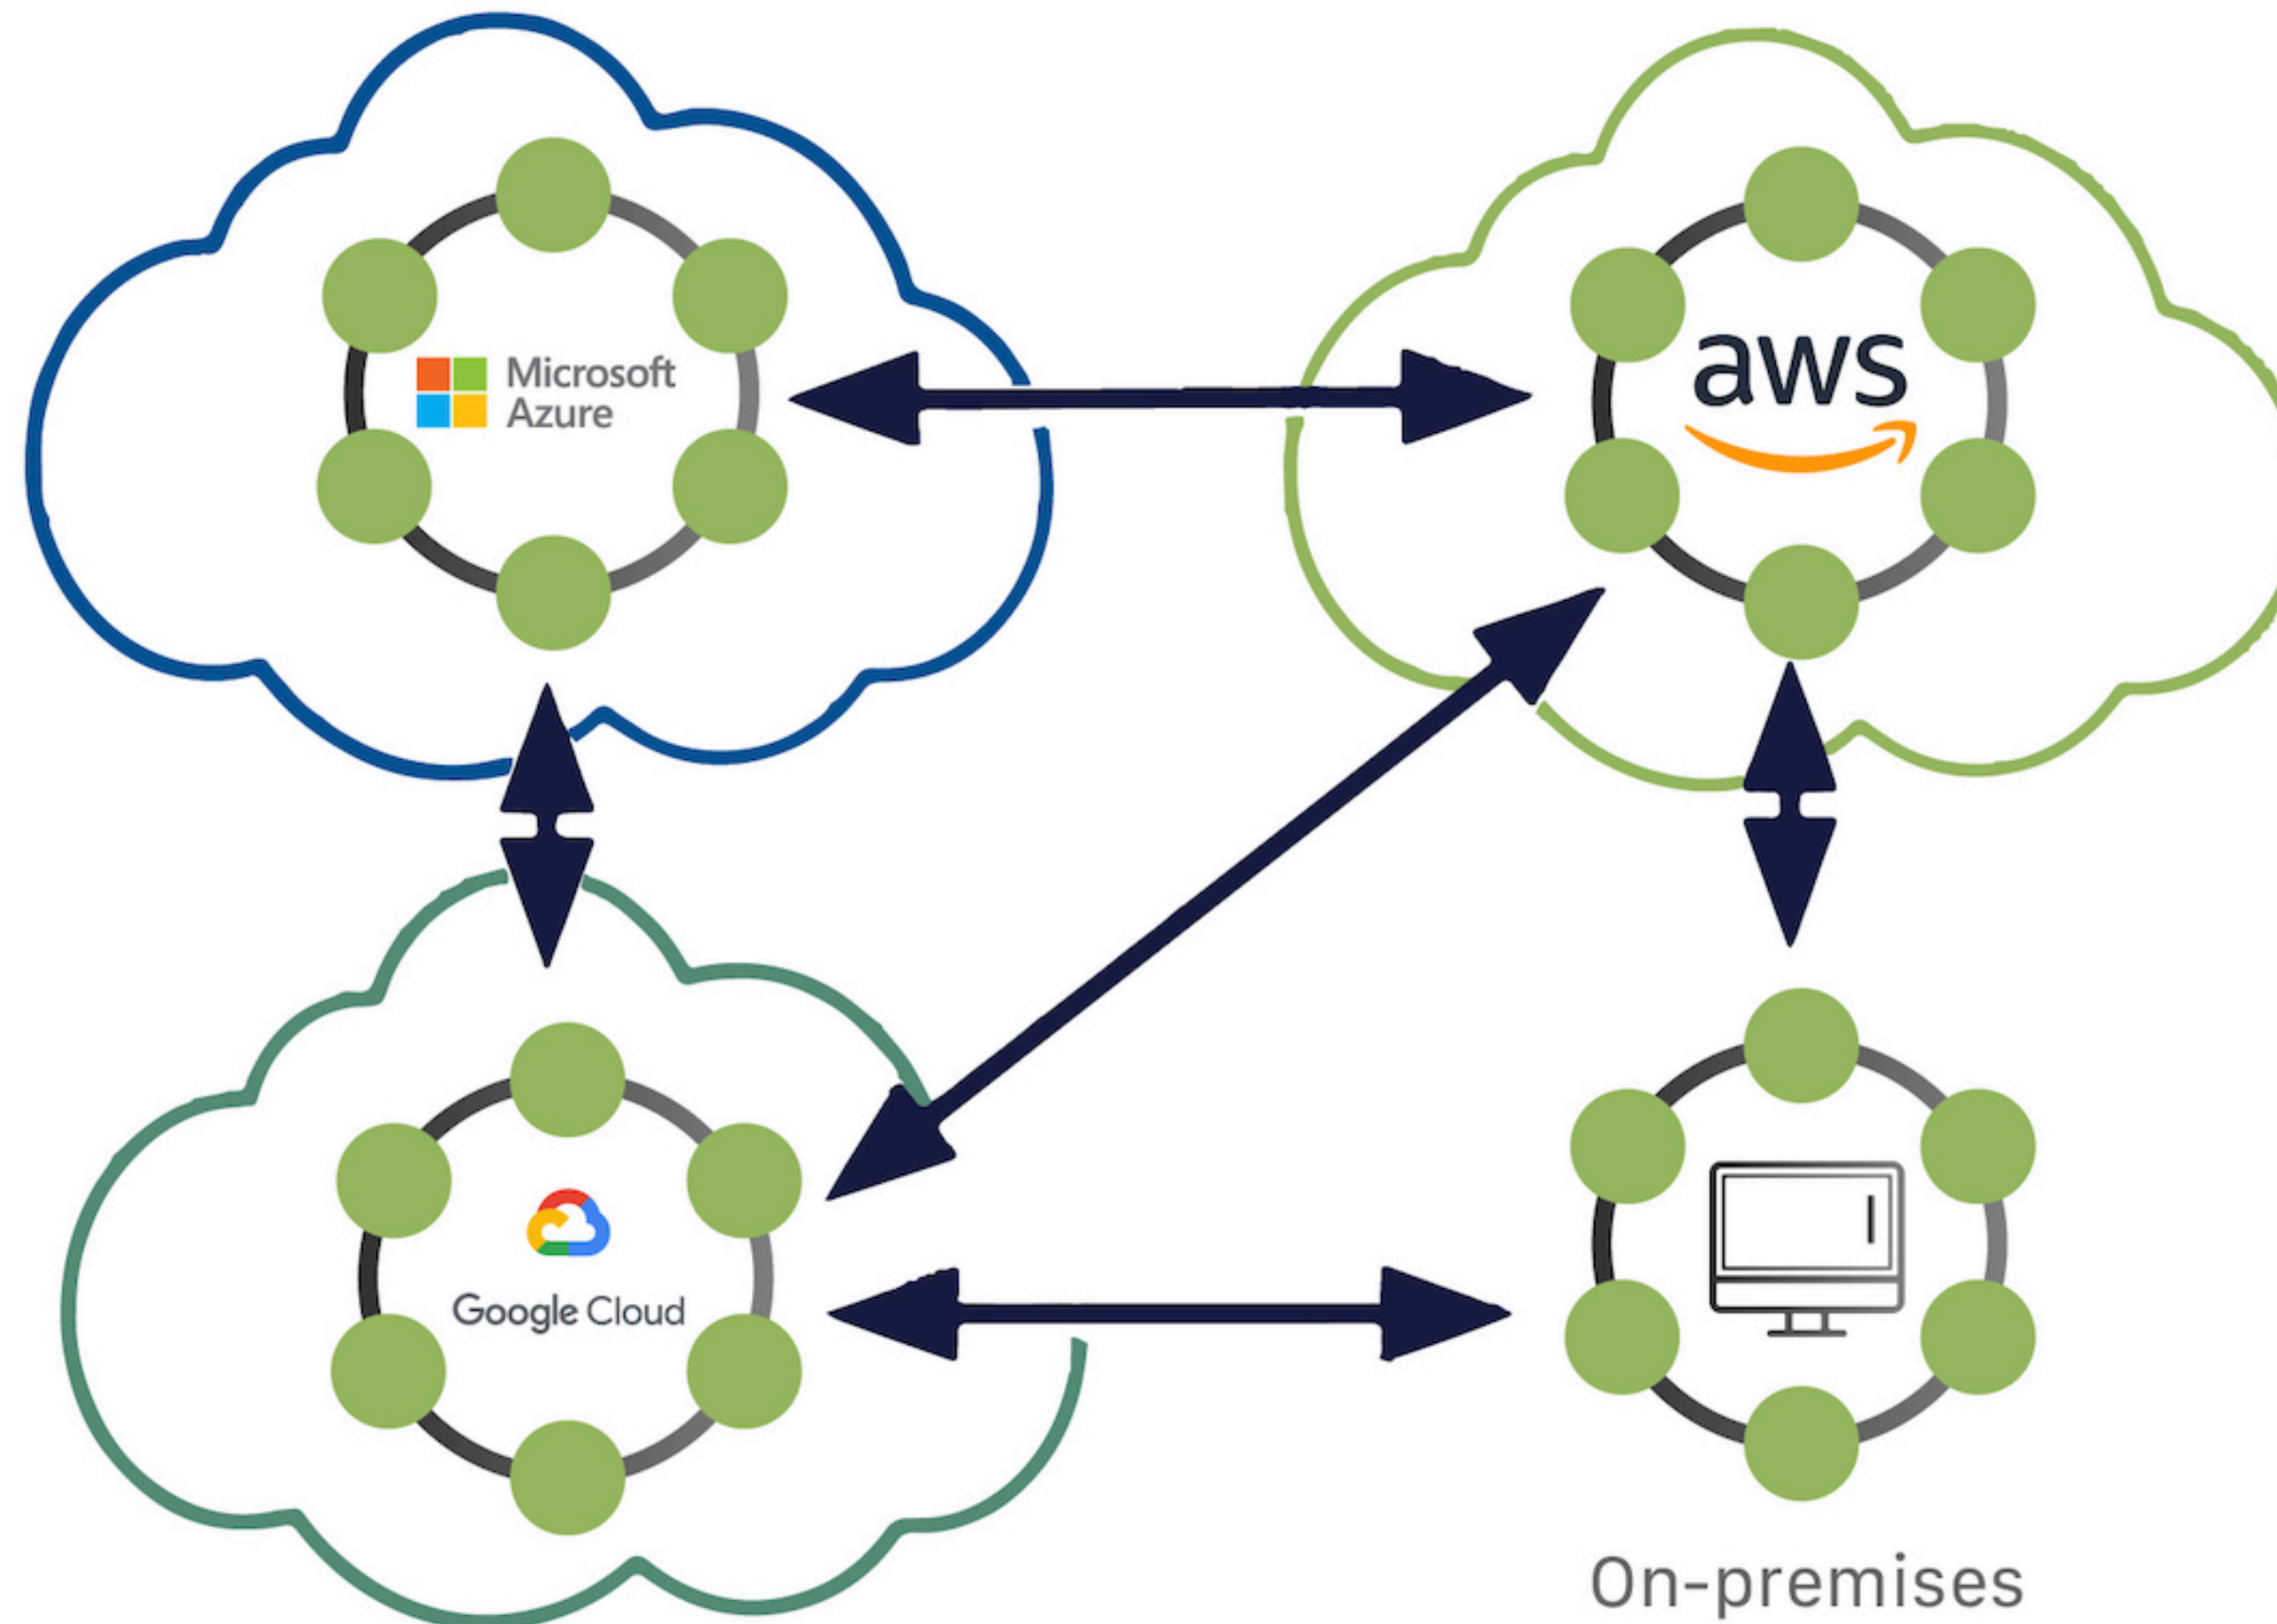
Availability

Can configure

C* default

cassandra

Partition
Tolerance
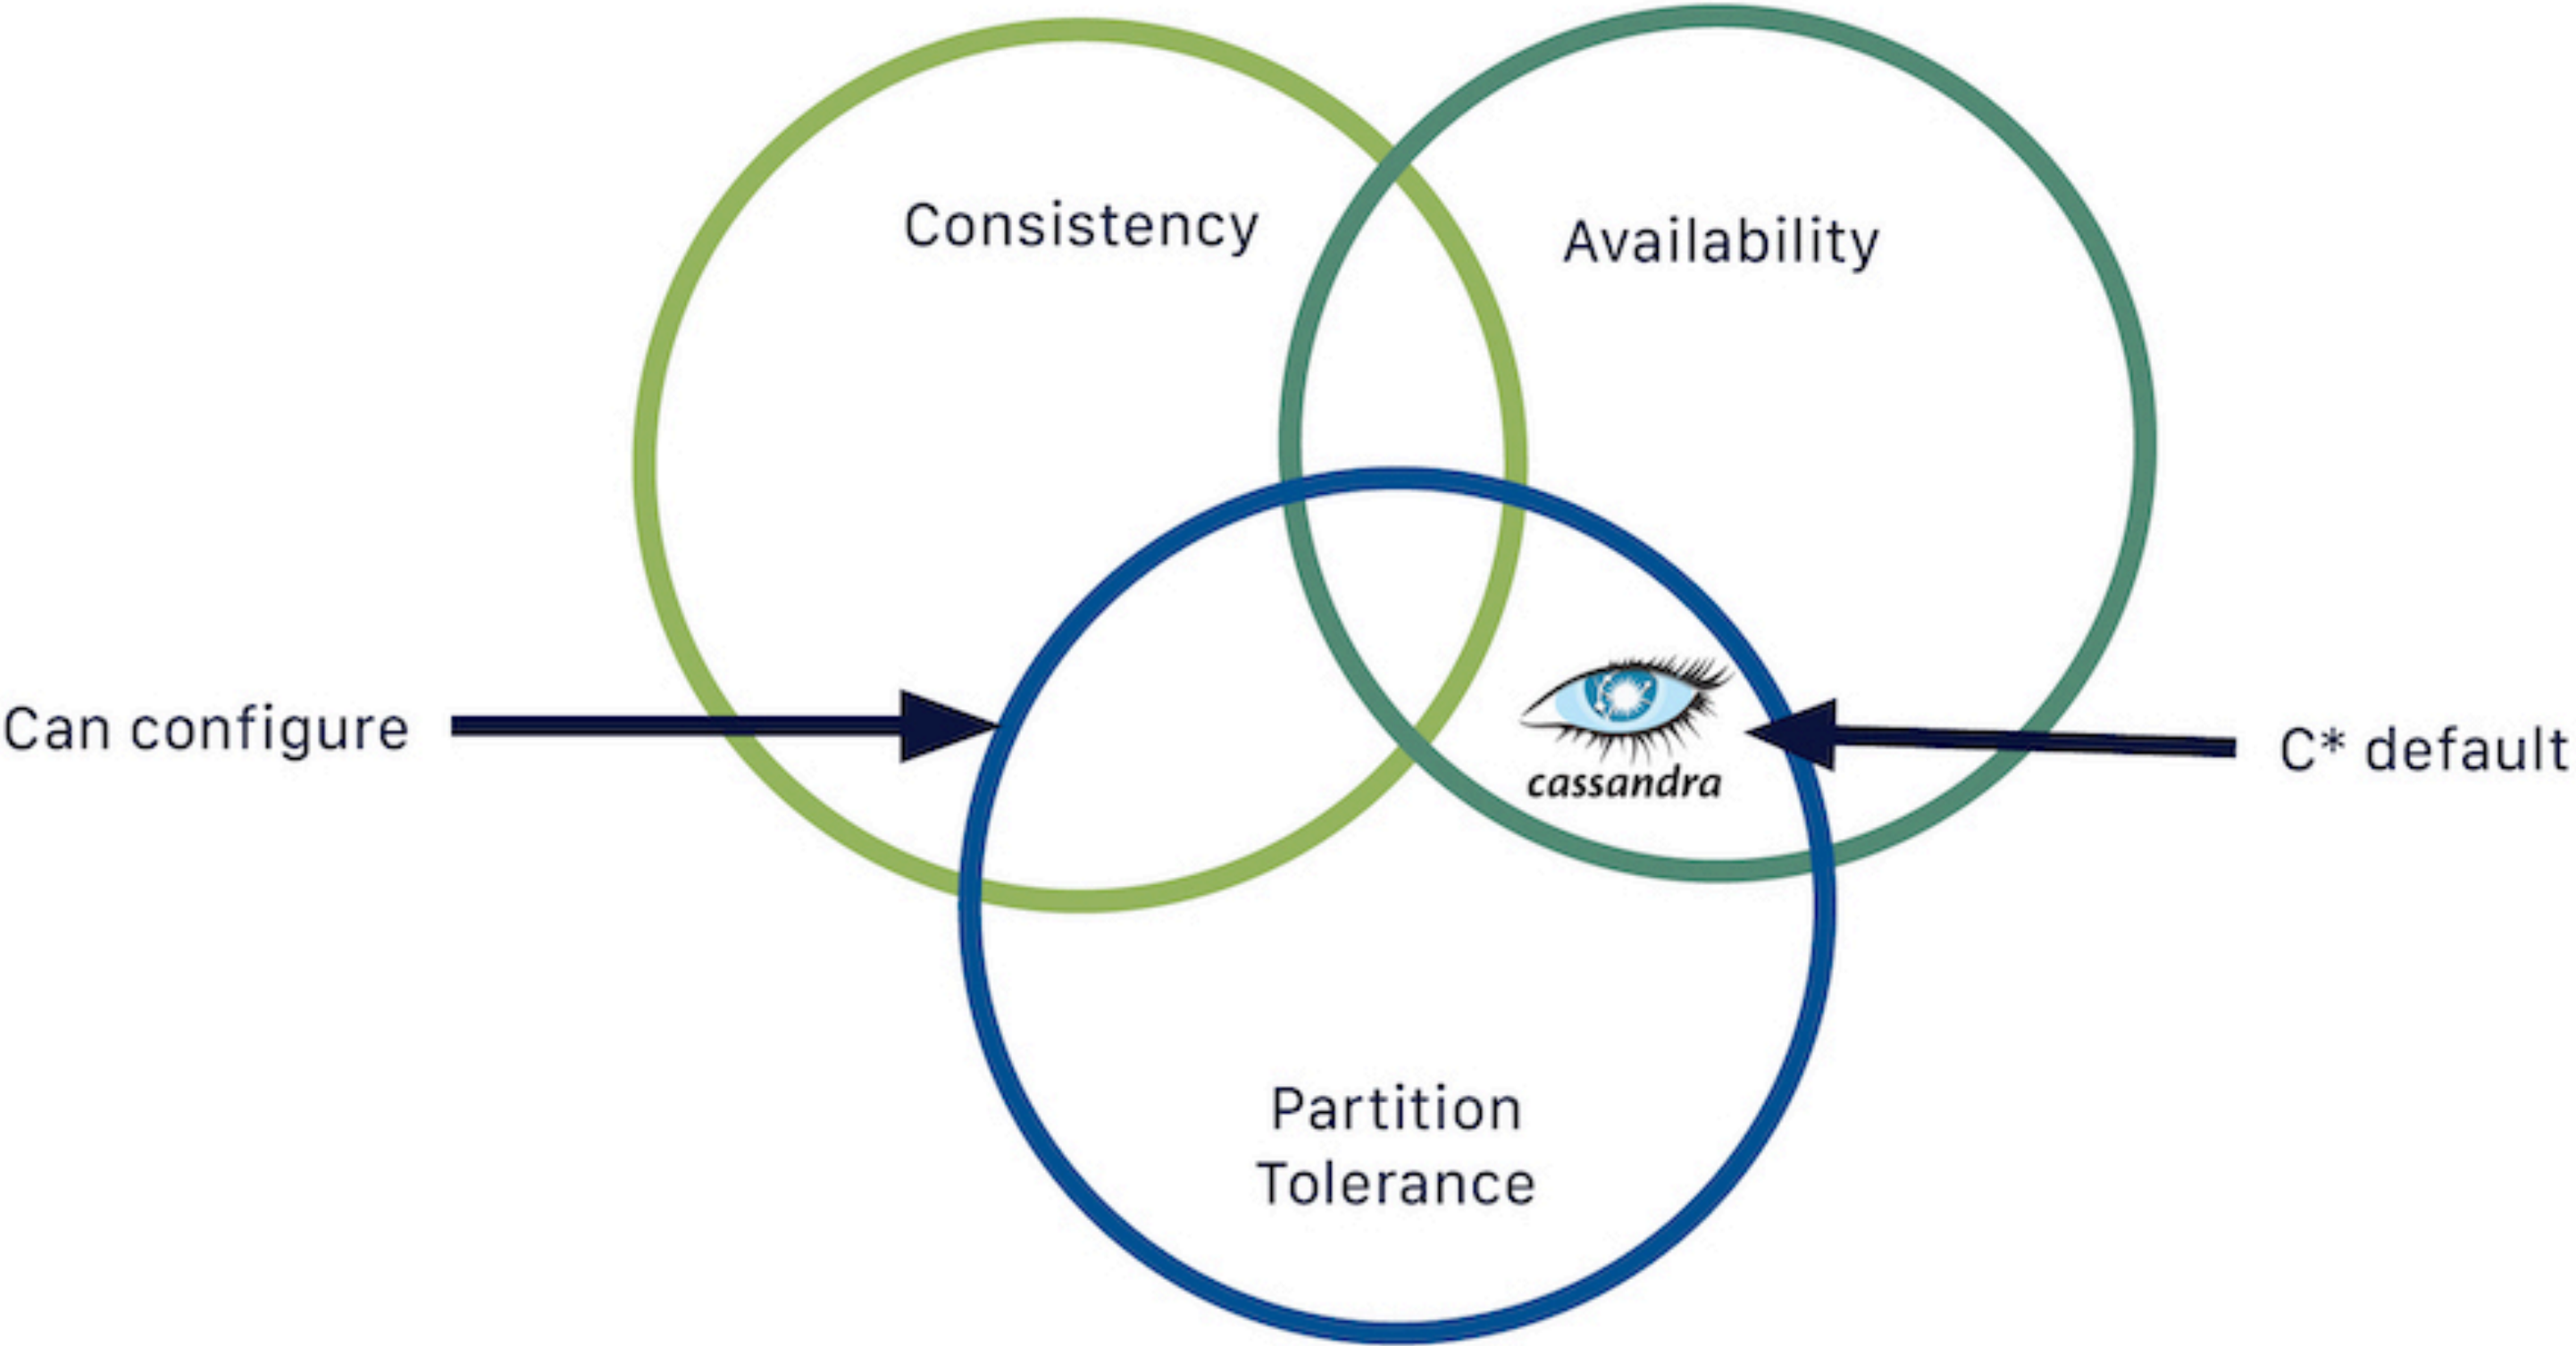
Cassandra - flexibility

Open source

Agenda

- History

- Architecture

- **Data model (Original)**
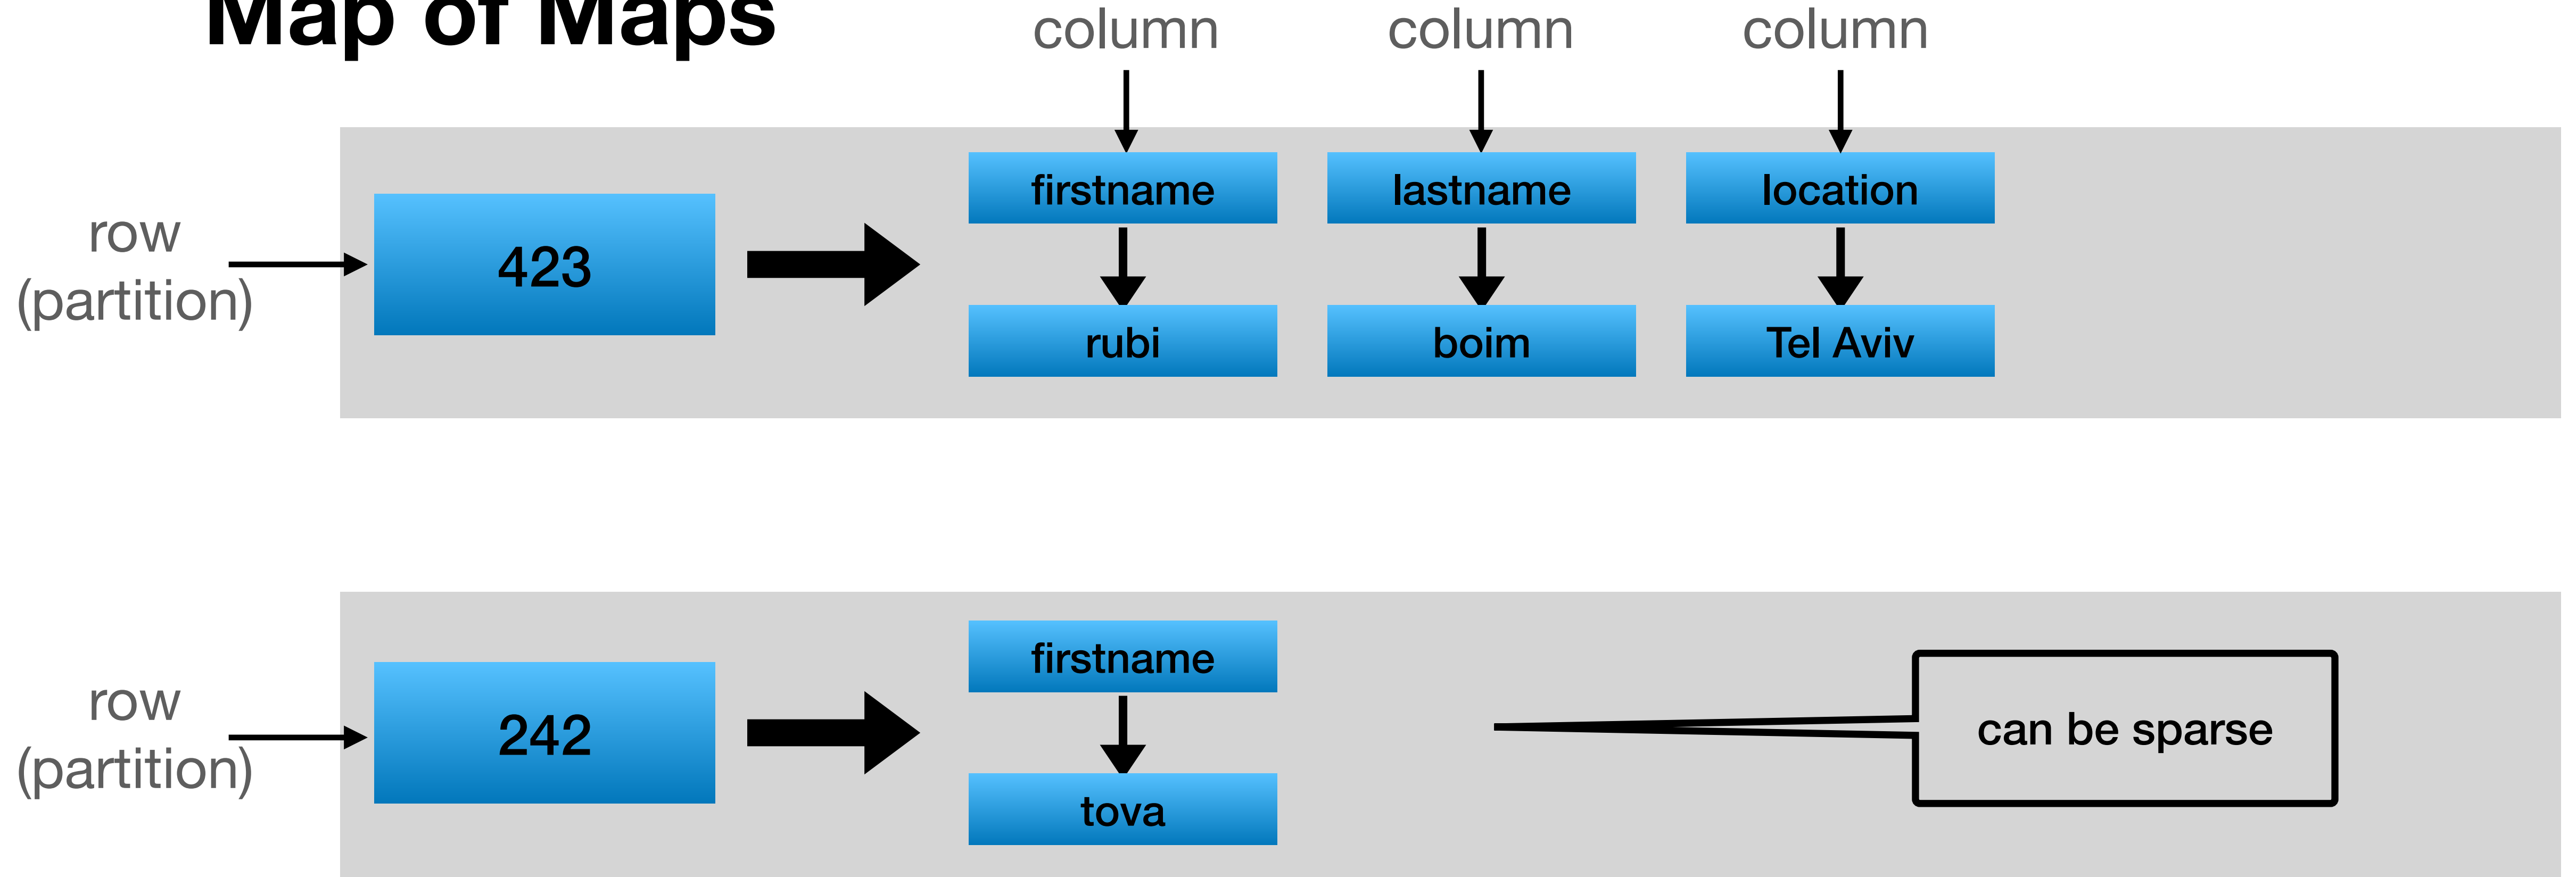
- Data model (CQL)

- Examples

Data model (original)

Map of Maps

- Wide column

- Similar to Bigtable

  - Without the versions

  - Rows are NOT ordered

Data model (original)

Map of Maps

column   column   column

row
(partition) → 423 ➡️ firstname → rubi   lastname → boim   location → Tel Aviv

row
(partition) → 242 ➡️ firstname → tova   can be sparse

<row:string, column:string> —> string

Data model (original)

API

- Thrift API

  - Standard / Open source

- Straightforward API
  "put / get / delete"

Data model (original)

Schema less

- When you create a table, it has no columns

  - Originally a table was called "column family"
    this will change in a few slides

  Note: a "column family" in Bigtable is something else

- Each row can add / remove columns

- Data is saved as raw bytes

Data model (original)

S

-

Think about it, if you look at the schema of the database
you only see something like:

```
table users
table movies
```
…
No info (id, name, title, rating…) is available <u>at the schema level</u>.

<u>You would need to "explore" the data to understand it</u>

-

- Data is saved as raw bytes

Data model (original)

Schema less - main problem

- Hard to understand the logic

  - Need to read code and explore data

- Data errors (no strict forcing for reading/writing data)
  for example, you save `long` values but read `int`

Data model - Update

After a few years a MAJOR update to Cassandra

- The storage had not changed at all

- Same wide column

- But the data modeling / API is completely different

"New" data model - CQL

Main goal

- Add "schema" to tables ("similar" to relational DB)

- Use an API "similar" to SQL

video_id	title	year	duration
1	Bad Boys	1995	119
2	Top Gun	1986	110
3	American Pie	1999	96

- Use the same storage/model behind the scenes

Agenda

- History

- Architecture

- Data model (Original)

- **Data model (CQL)**

- Examples

Data model (CQL) - TLDR

- A table in Cassandra has

  - Partition keys

  - Clustering columns

  - Columns
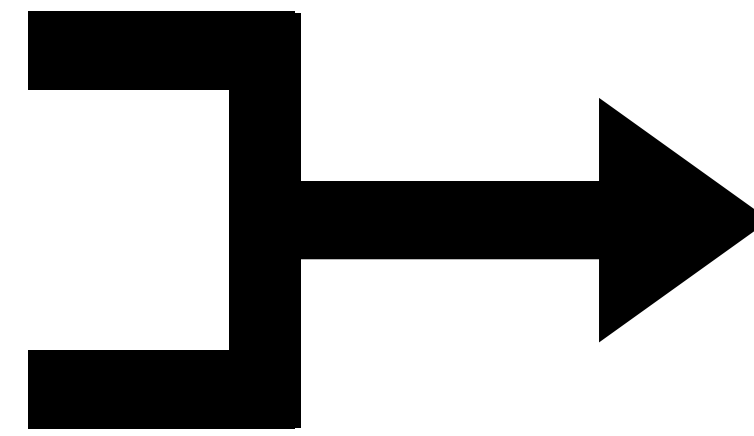
Primary key

Partition Key

- Partition key: <u>the first part of the primary key</u>

- `hash(partition key)` defines the <u>partition</u>

- Dynamo will map between partitions and nodes
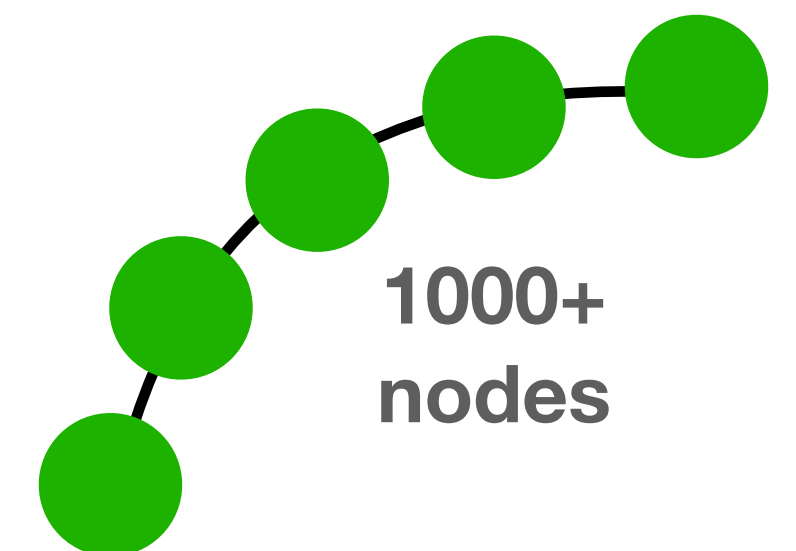 partition range: [-2^63 to +2^63]

1000+ nodes

Partition Key

- Partition key: <u>the first part of the primary key</u>

- `hash(p`——————————————`on`

A partition is similar to a ROW in Bigtable

- Dynamo will map between partitions and nodes
  partition range: [-2^63 to +2^63]

  1000+ nodes

Partition Key

- Partition key: the first part of the pri[mary]

- `hash(partition key)` defines the partition

- Dynamo will map between partitions and nodes
  partition range: [-2^63 to +2^63]

Throughout the rest of the course, whenever you would see the "1000 nodes ring" image think how the slide's information would be applied on large scale

1000+ nodes
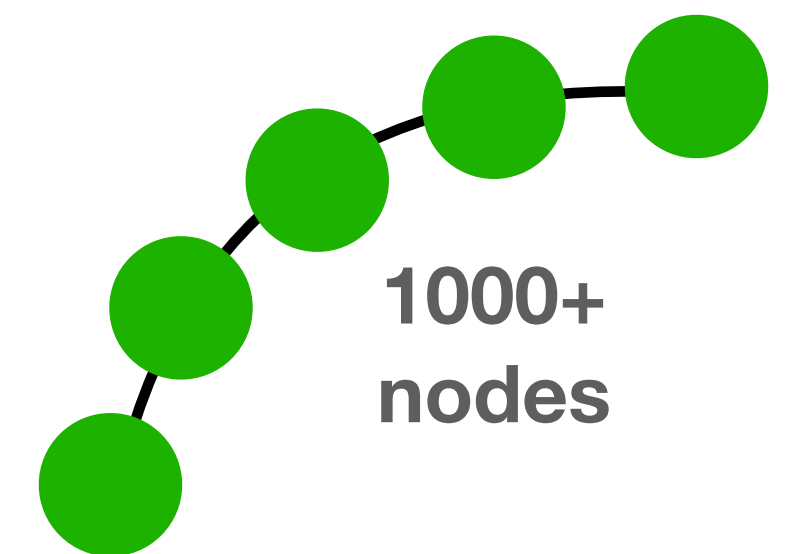
Classic (relational) view

```
CREATE TABLE movies (
    video_id      BIGINT,
    title         TEXT
    year          INT,
    duration      INT,
    PRIMARY KEY (video_id)
);
```
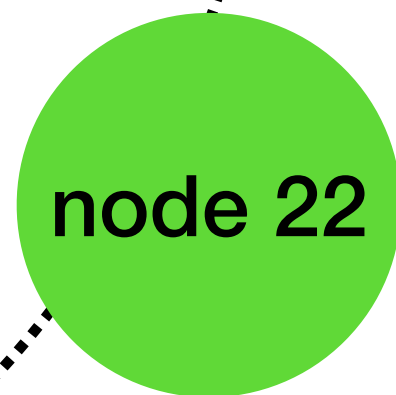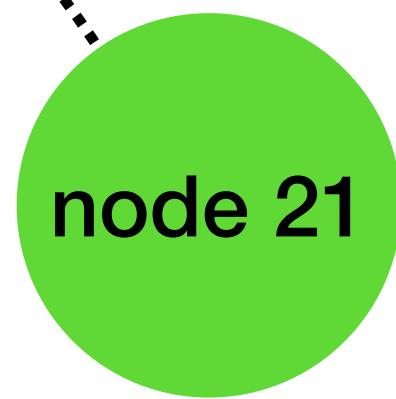
movies	
video_id	K
title	
year	
duration	

What is the partition key of Top Gun?

movies

video_id	title	year	duration
1	Bad Boys	1995	119
2	Top Gun	1986	110
3	American Pie	1999	96

Relational —> wide columns

video_id	title	year	duration
1	Bad Boys	1995	119
2	Top Gun	1986	110
3	American Pie	1999	96

node 21

node 22

```
CREATE TABLE movies (
    video_id        BIGINT,
    title           TEXT
    year            INT,
    duration        INT,
    PRIMARY KEY (video_id)
);
```

video_id	title	year	duration
1	Bad Boys	1995	119
2	Top Gun	1986	110
3	American Pie	1999	96

node 21

node 22
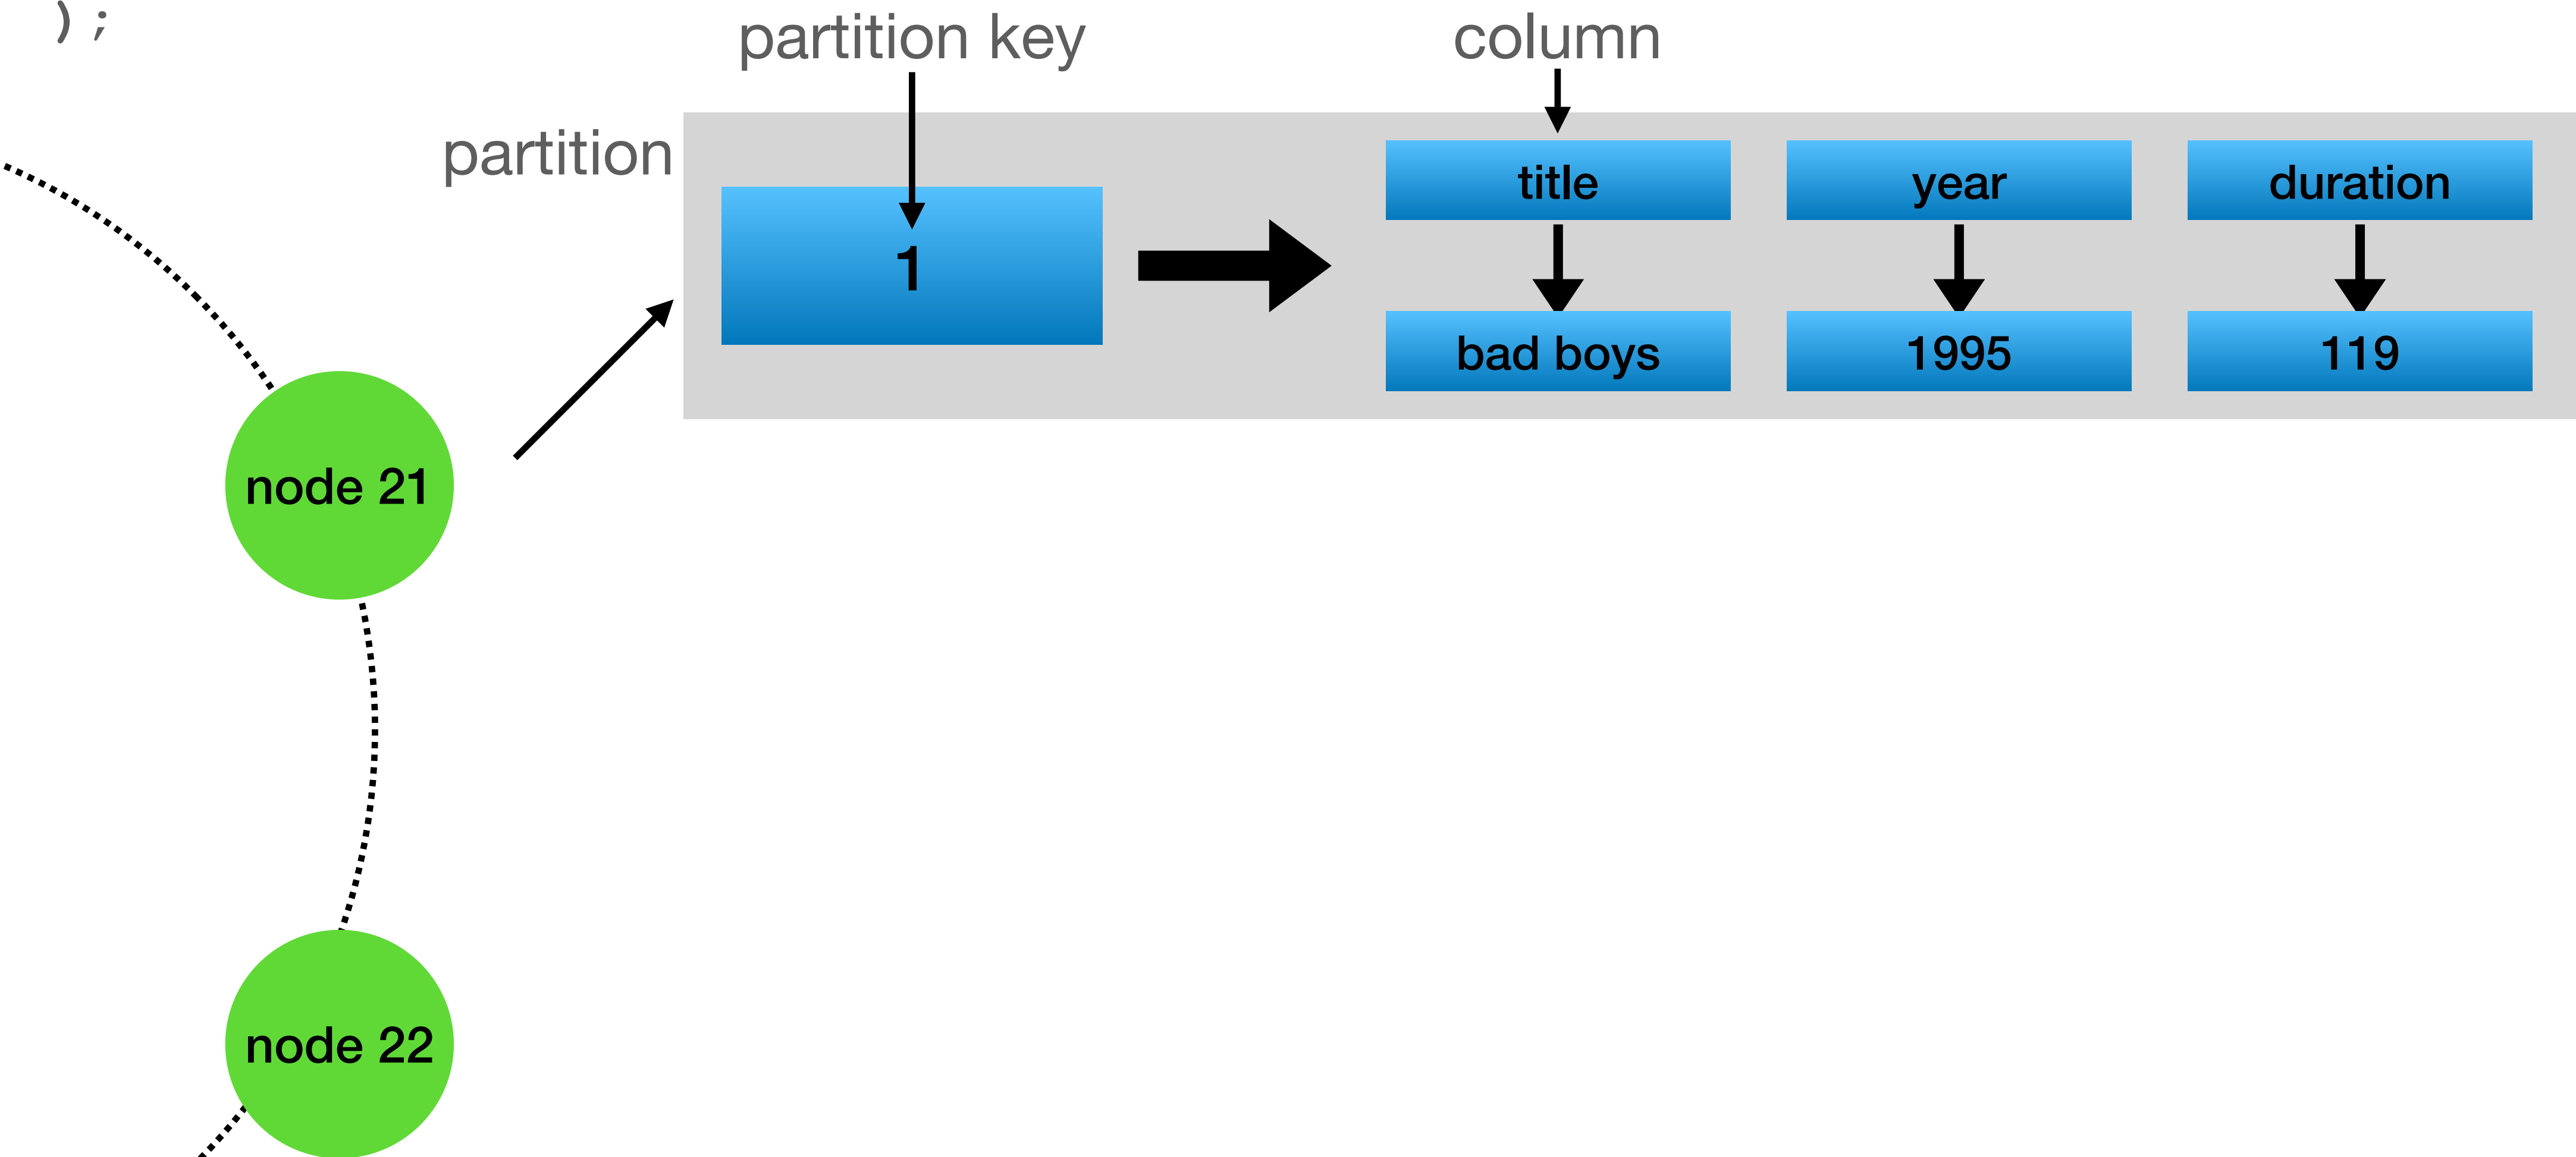
```
CREATE TABLE movies (
    video_id        BIGINT,
    title           TEXT
    year            INT,
    duration        INT,
    PRIMARY KEY (video_id)
);
```

video_id	title	year	duration
1	Bad Boys	1995	119
2	Top Gun	1986	110
3	American Pie	1999	96

partition key   column

partition

1   →   title   year   duration

        bad boys   1995   119

node 21

node 22
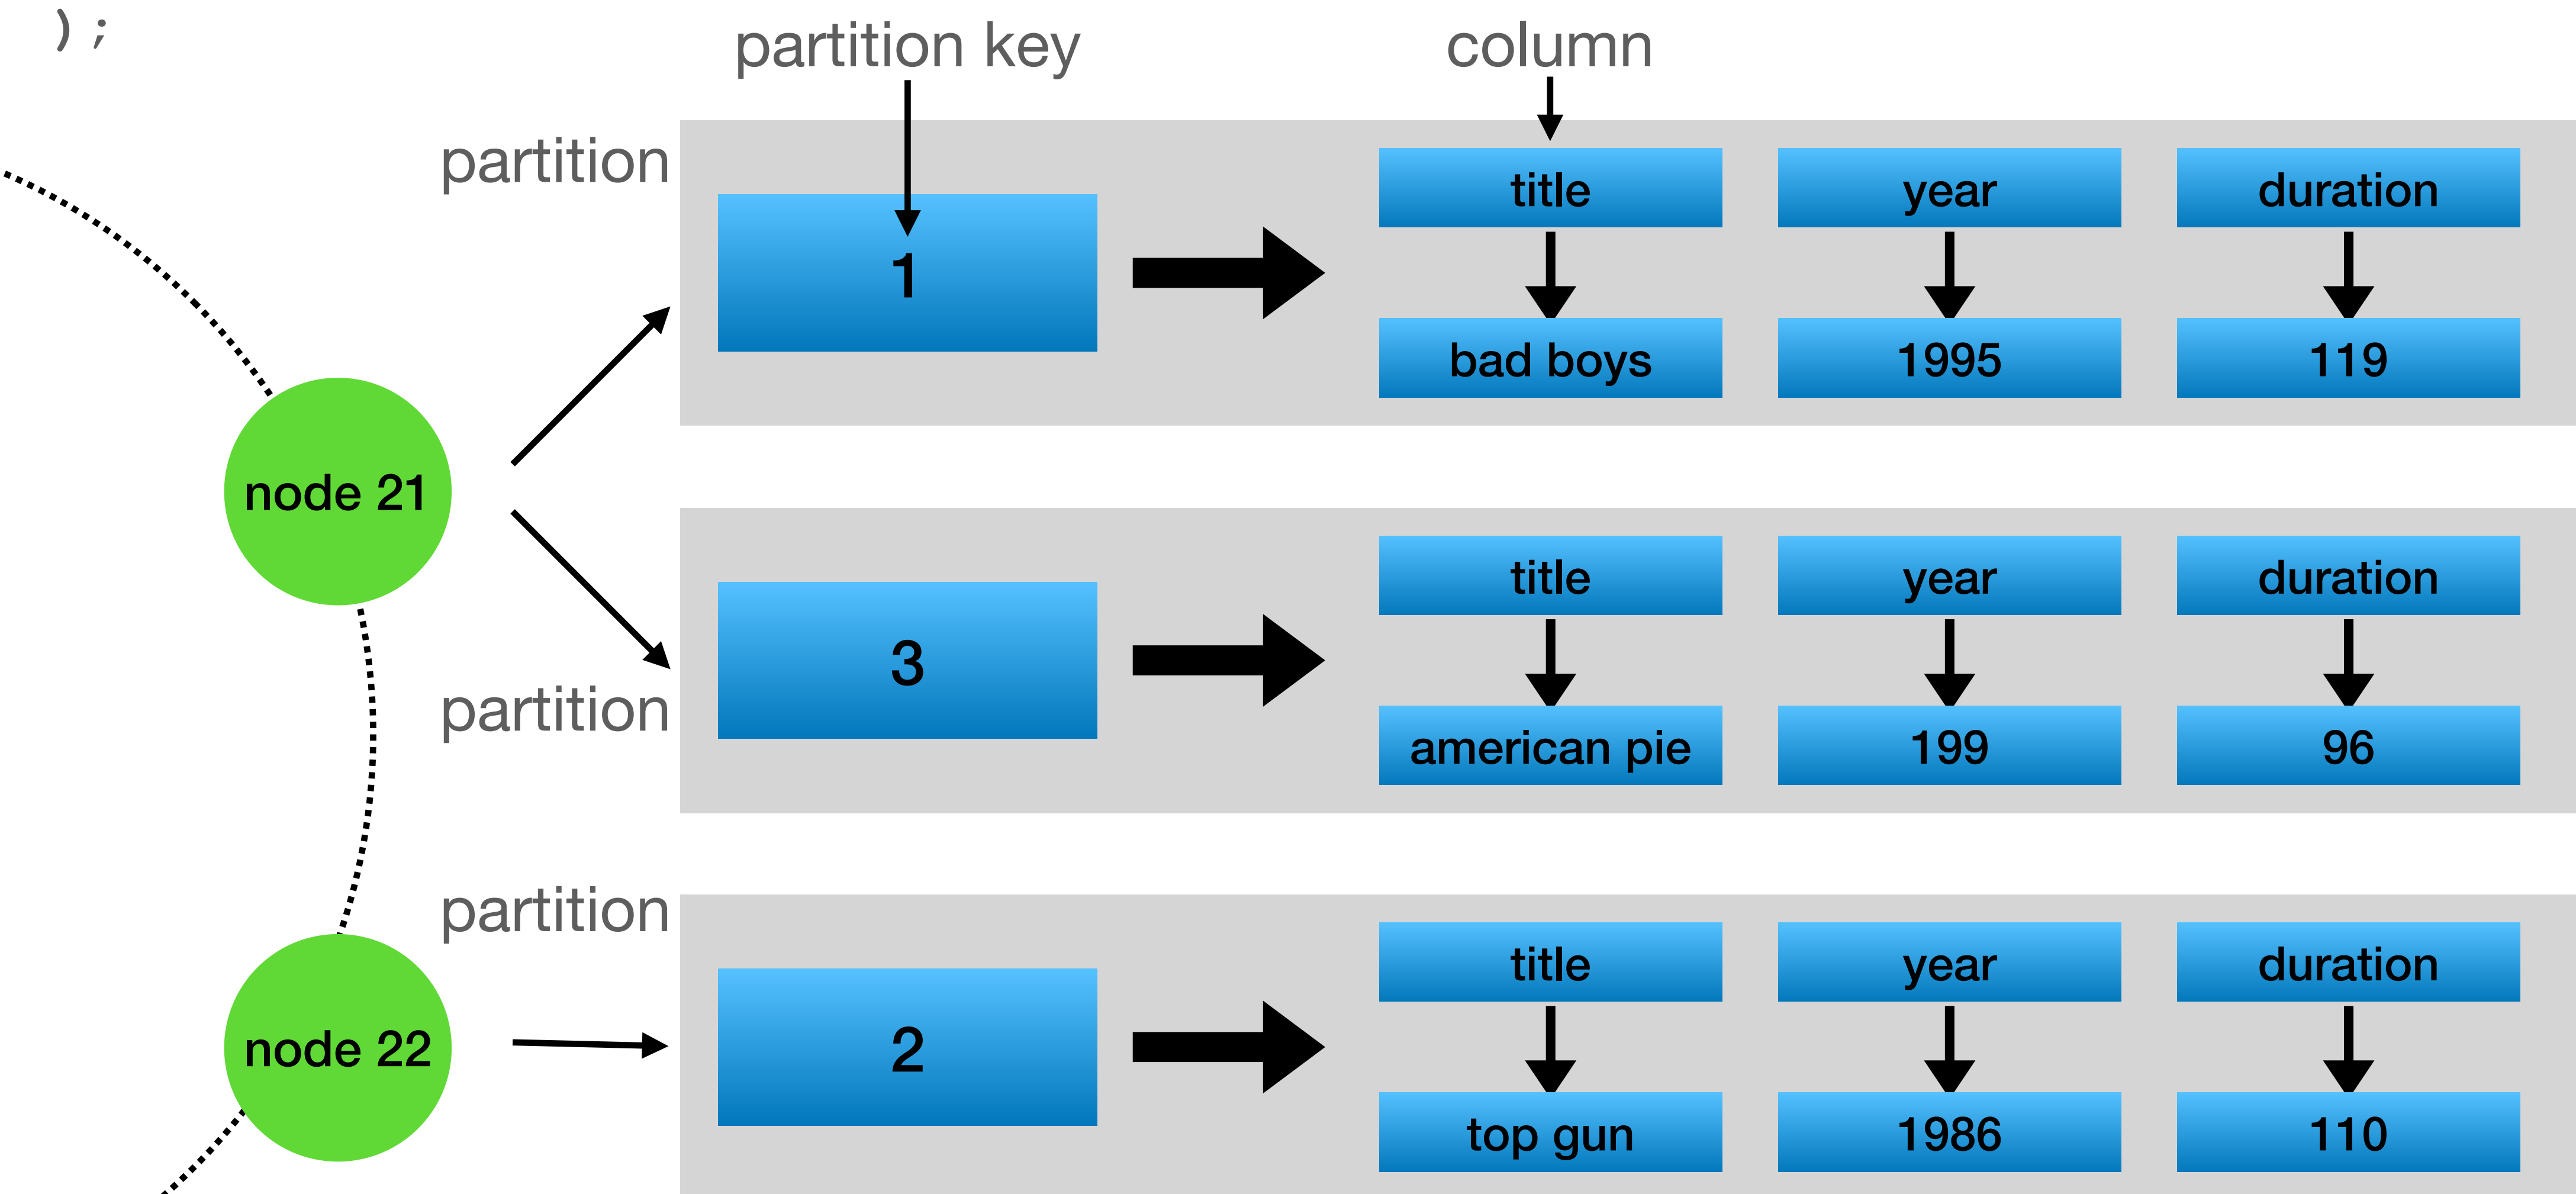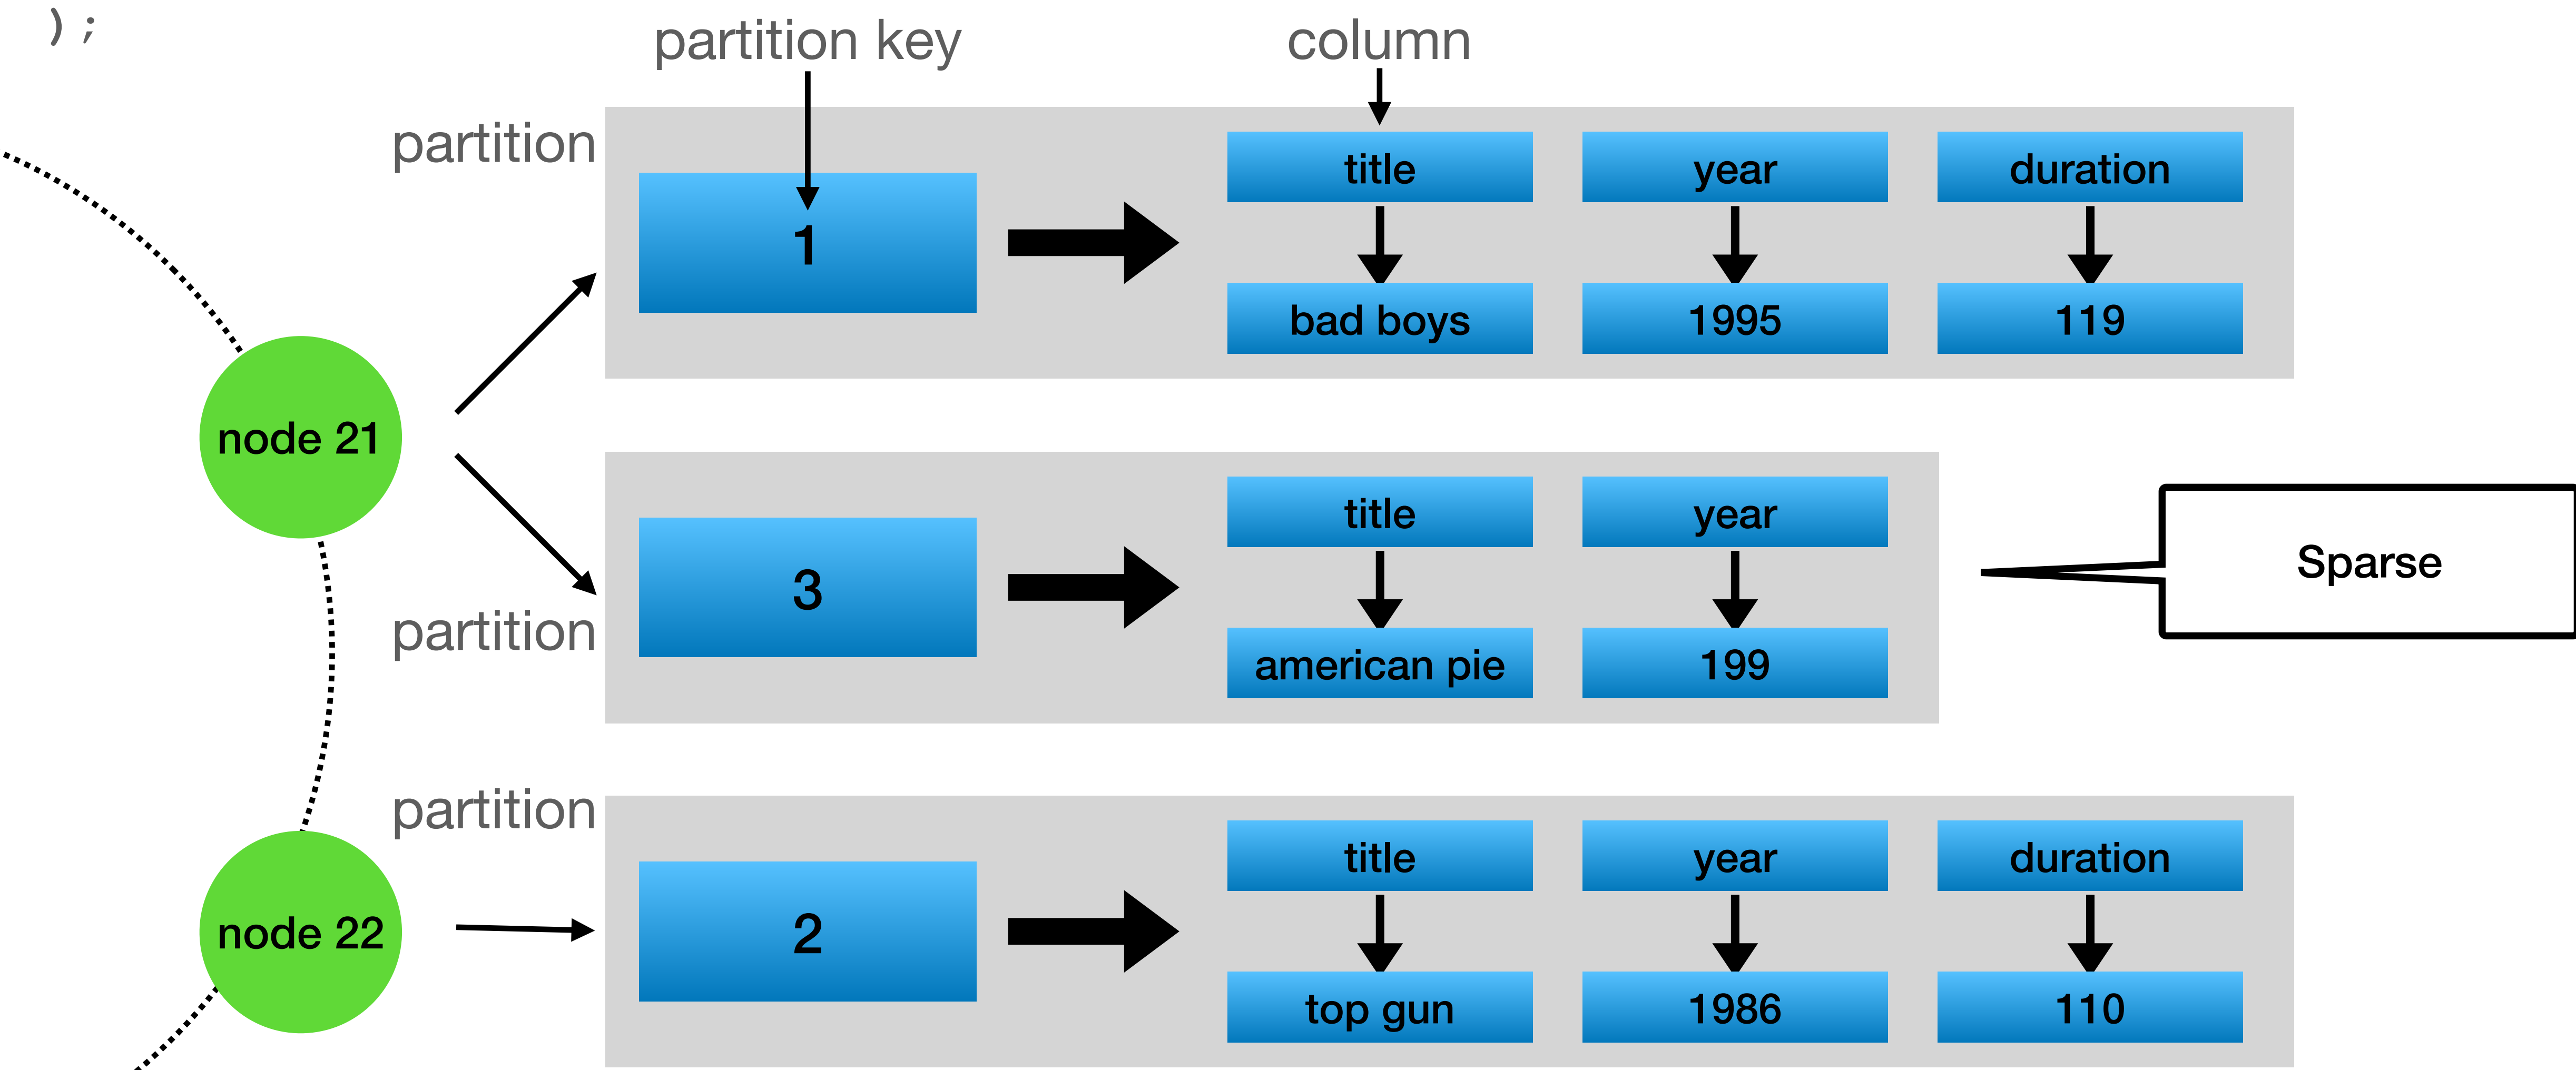
```
CREATE TABLE movies (
    video_id        BIGINT,
    title           TEXT
    year            INT,
    duration        INT,
    PRIMARY KEY (video_id)
);
```

video_id	title	year	duration
1	Bad Boys	1995	119
2	Top Gun	1986	110
3	American Pie	1999	96

partition key   column

partition

node 21

1 → title → bad boys | year → 1995 | duration → 119

3 → title → american pie | year → 199 | duration → 96

partition

node 22

partition

2 → title → top gun | year → 1986 | duration → 110

45

```
CREATE TABLE movies (
    video_id        BIGINT,
    title           TEXT
    year            INT,
    duration        INT,
    PRIMARY KEY (video_id)
);
```

video_id	title	year	duration
1	Bad Boys	1995	119
2	Top Gun	1986	110
3	American Pie	1999	

partition key   column

partition

node 21

1 → title → bad boys | year → 1995 | duration → 119

3 → title → american pie | year → 199

Sparse

partition

node 22

partition

2 → title → top gun | year → 1986 | duration → 110

46

SELECTing data efficiently

- Cassandra needs to know where the data is located —> **We need to provide the partition key**

```
SELECT * FROM movies WHERE video_id = 1
```

Partition key

1000+ nodes

SELECTing data efficiently

- Cassandra needs to know where the data is located —> **We need to provide the partition key**

```
SELECT * FROM movies WHERE video_id = 1
```

Partition key

partition key = 1
—>
token = 12312
—>
data is on node 22
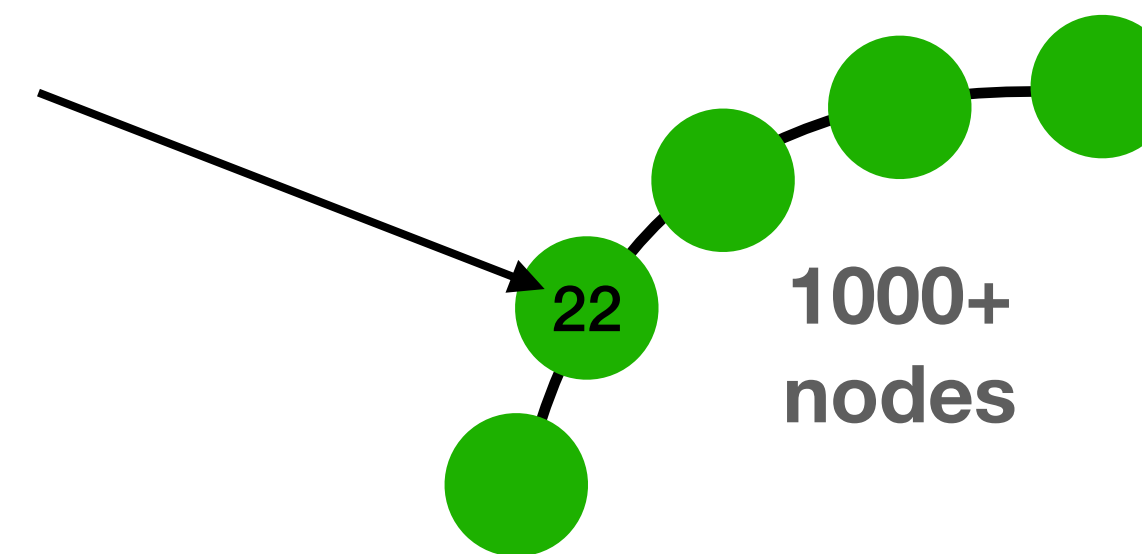
22   1000+
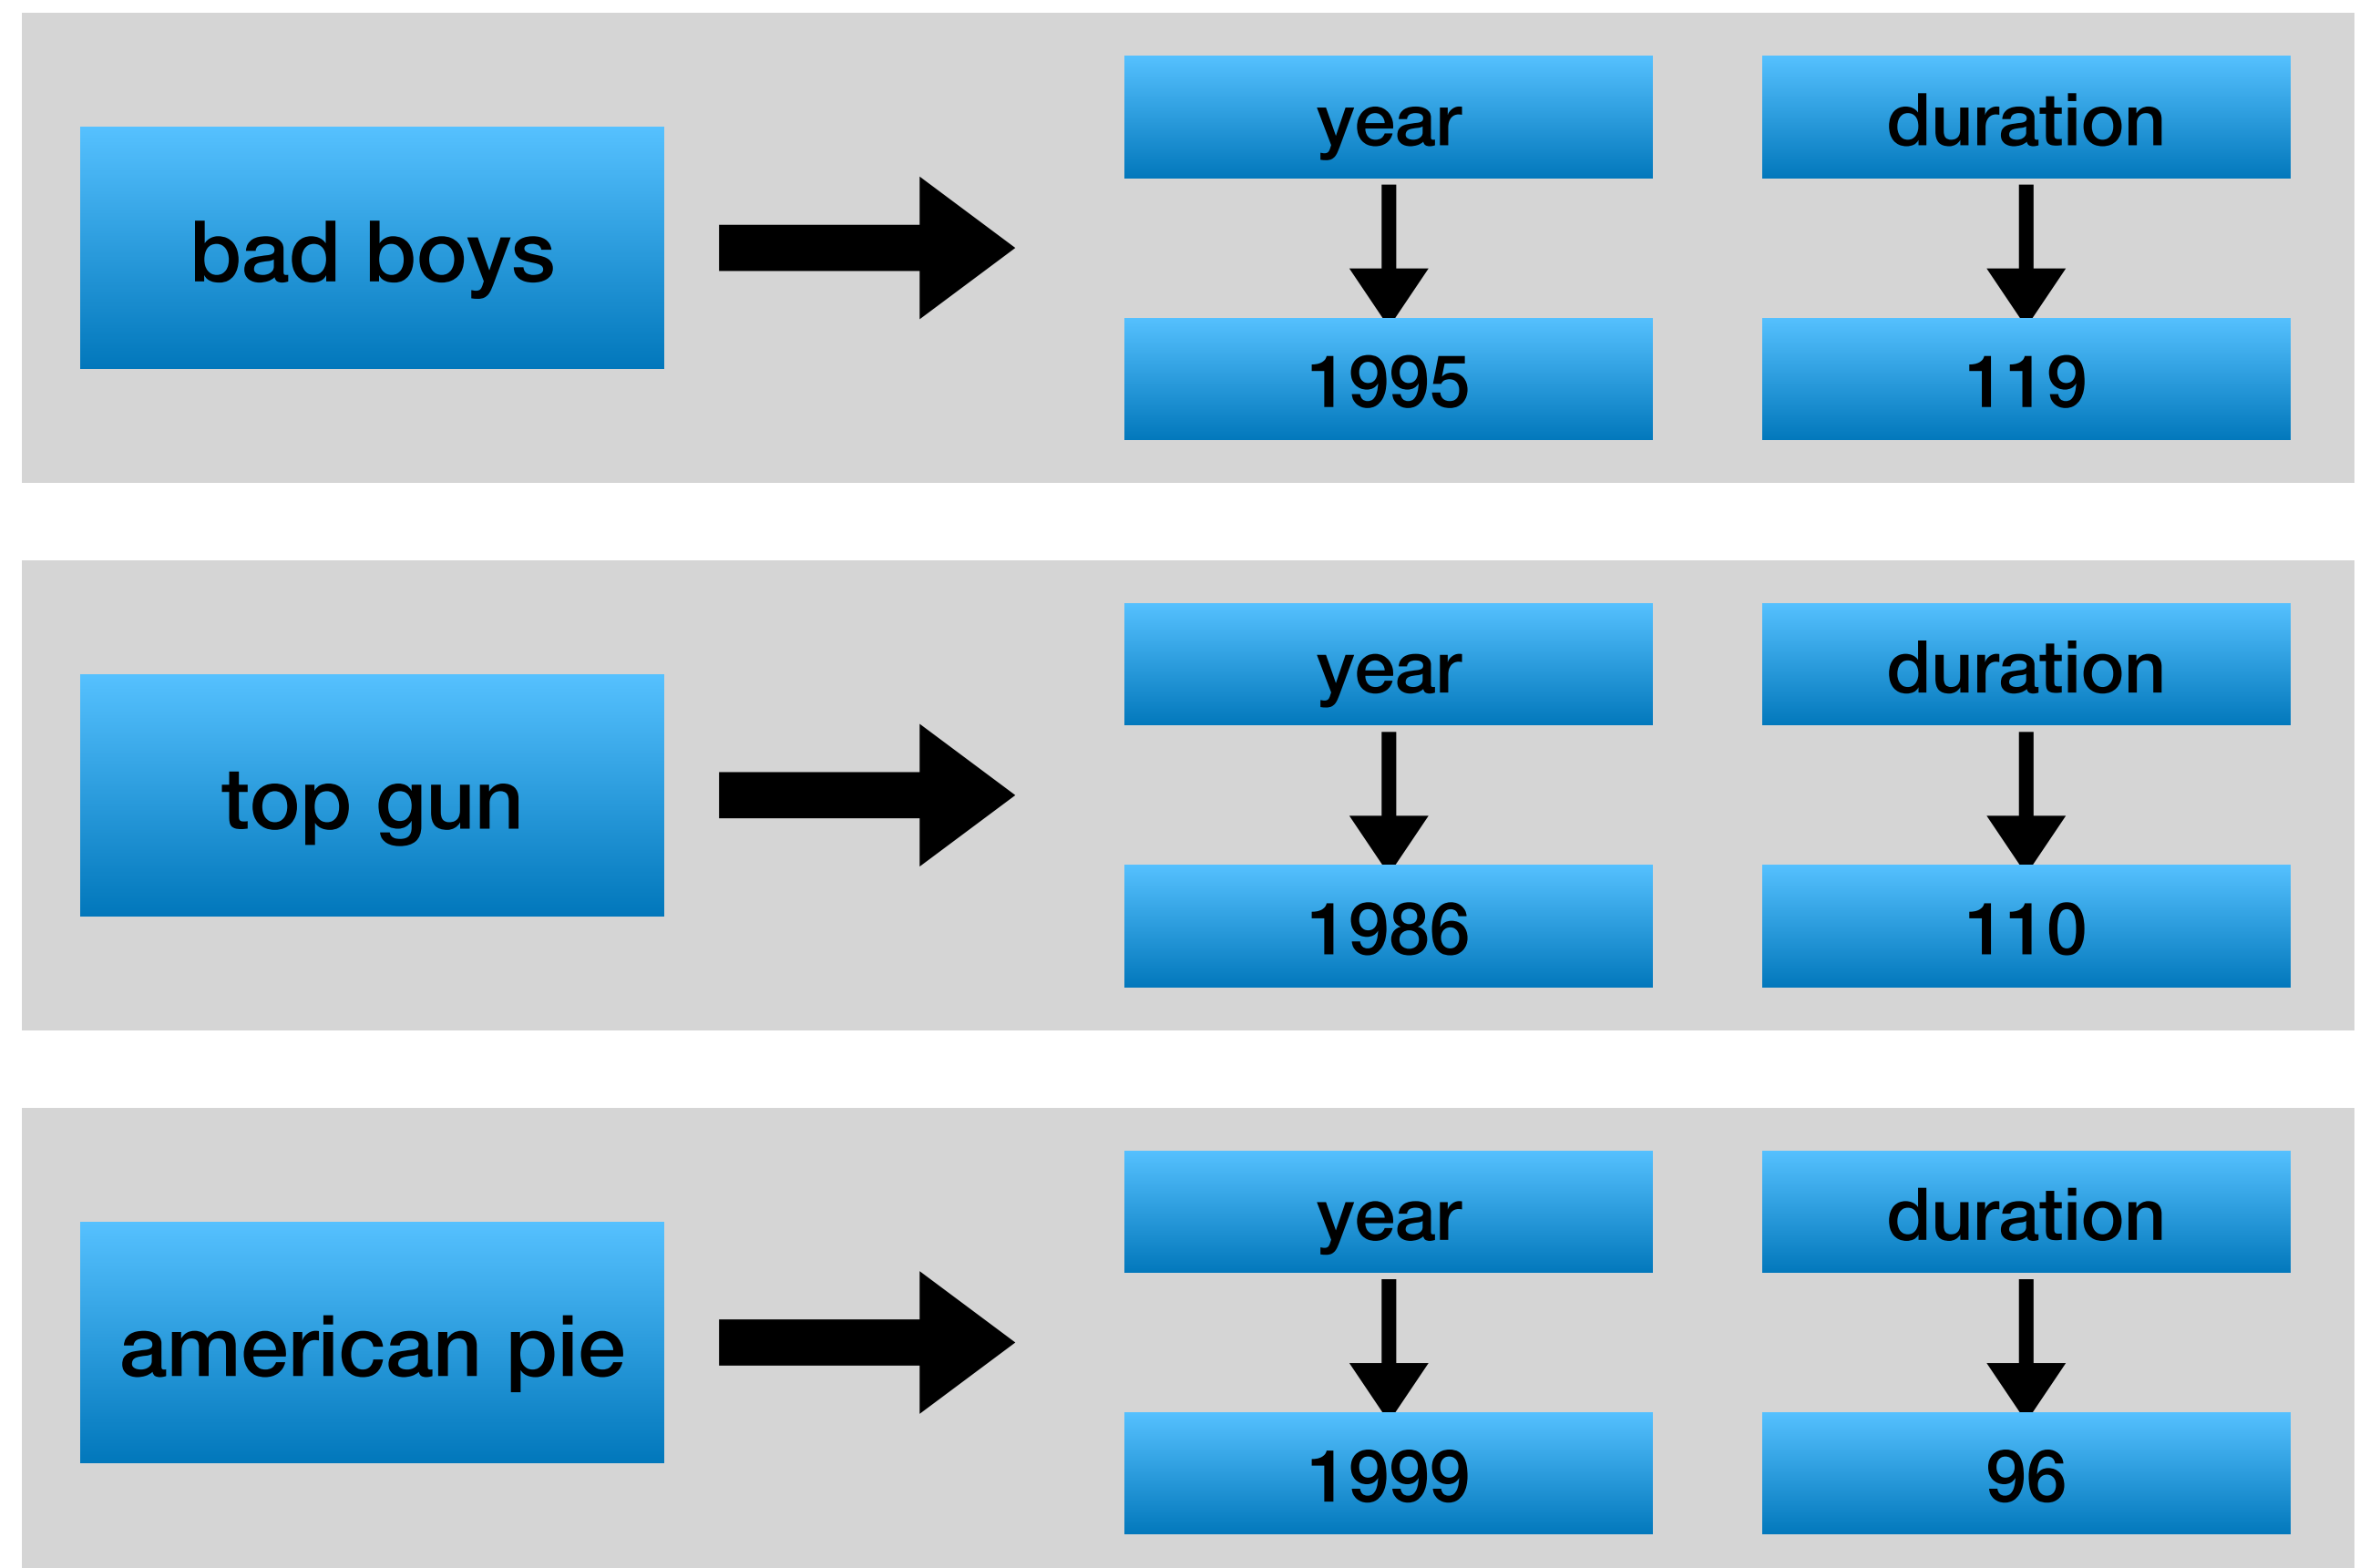nodes

Simple partition key

- Single column (similar to previous example)

```
CREATE TABLE "movies" (
    title       TEXT
    year        INT,
    duration    INT,
    PRIMARY KEY (title)
);
```

The title is the key / partition key

bad boys	→	year	duration
		1995	119

top gun	→	year	duration
		1986	110

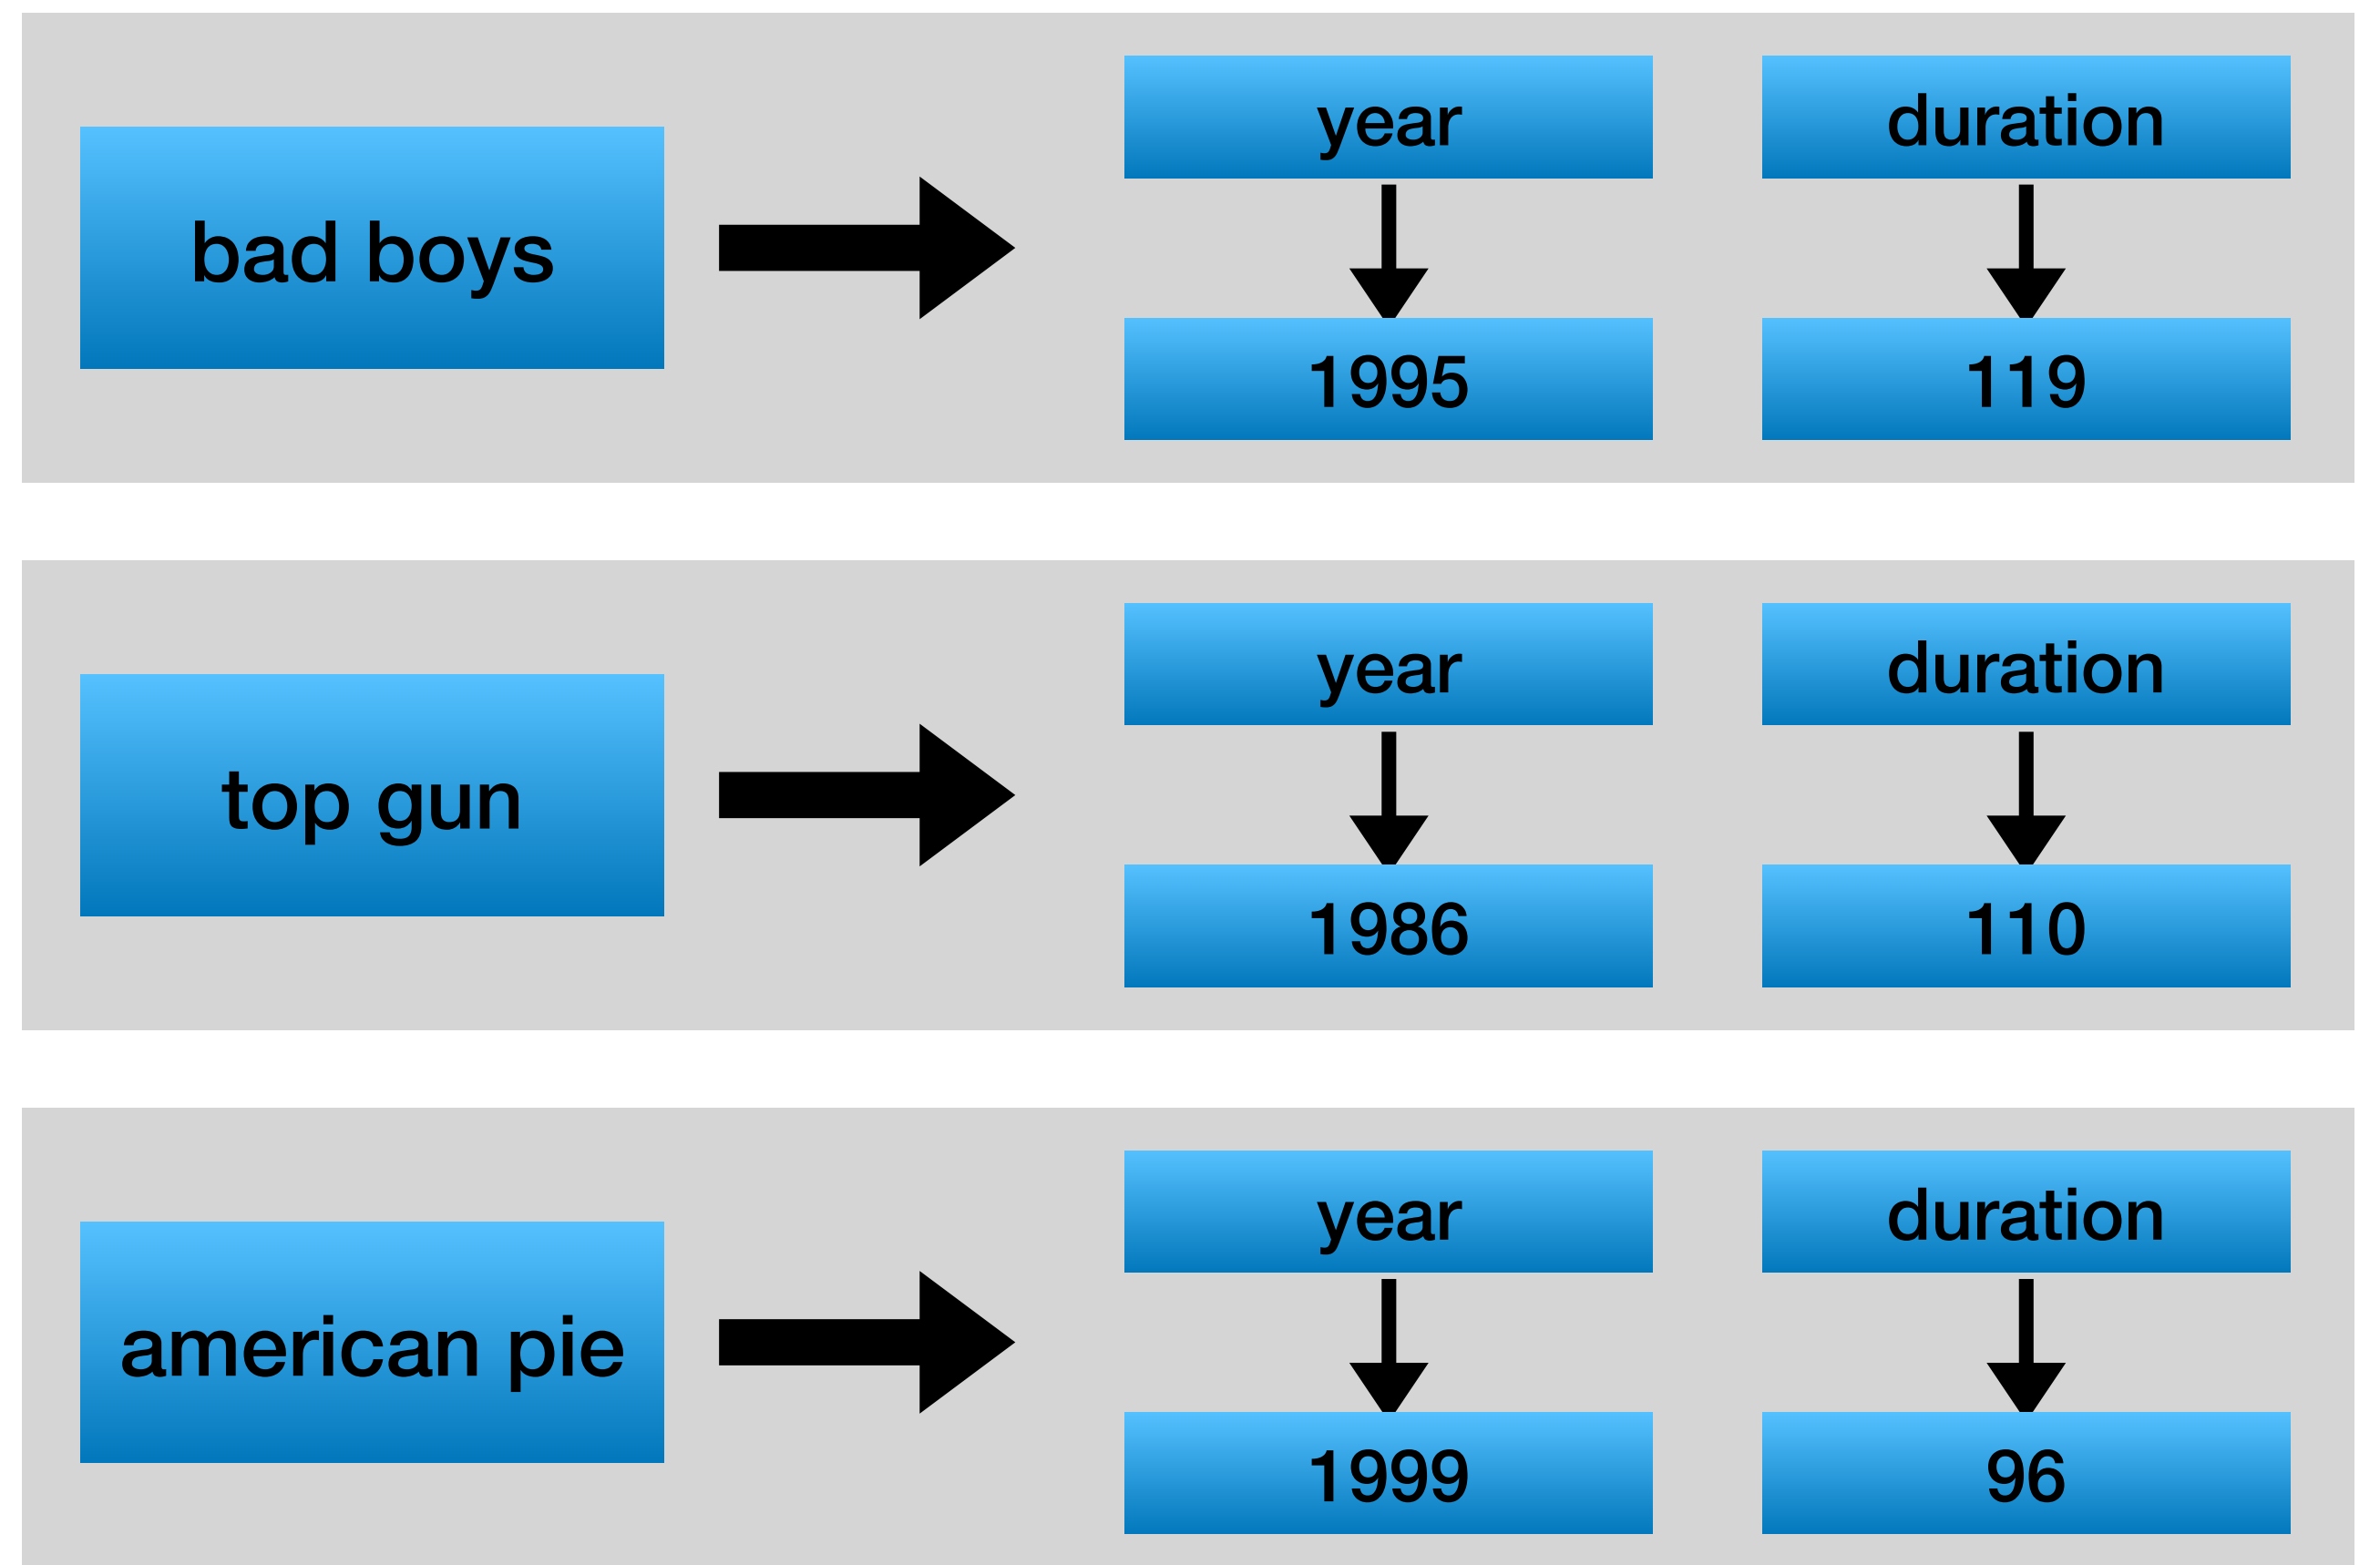| american pie	→	year	duration
		1999	96

Simple partition key

- Single column (similar to previous example)

```
CREATE TABLE "movies" (
    title       TEXT
    year        INT,
    duration    INT,
    PRIMARY KEY (title)
);
```

The title is the key / partition key
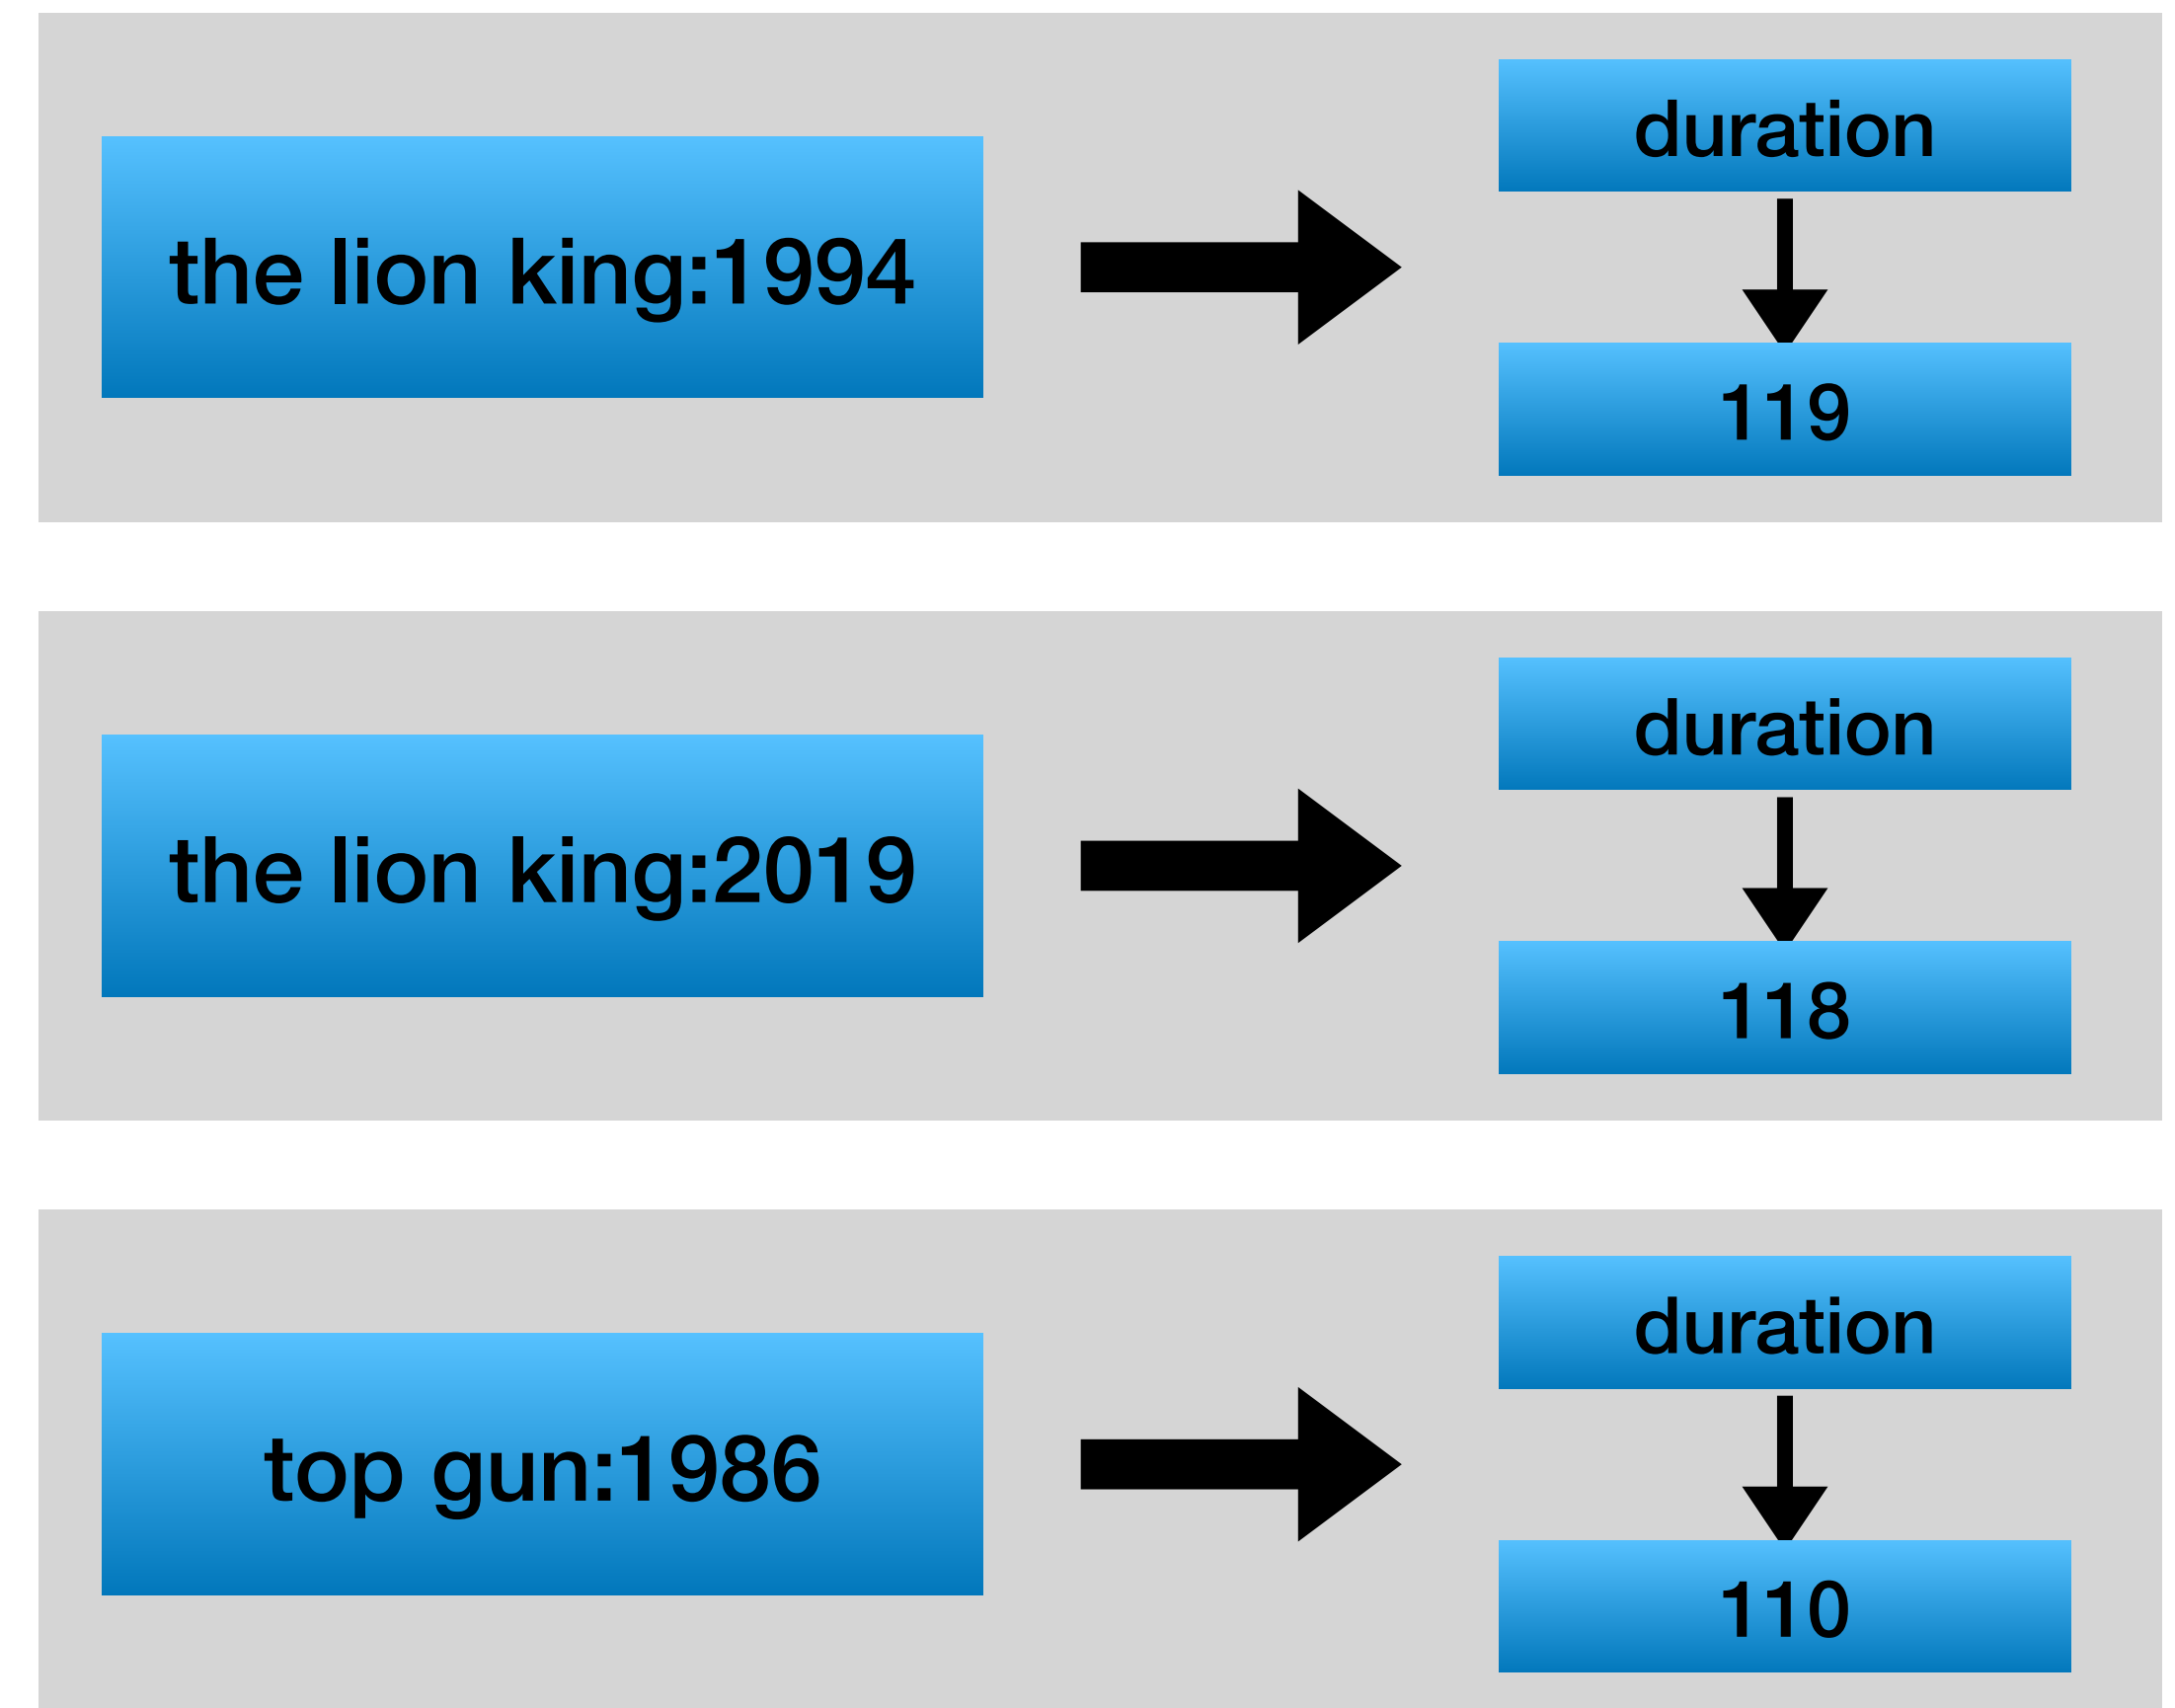
What happens if we have a remake?

bad boys	→	year	duration
		↓	↓
		1995	119

top gun	→	year	duration
		↓	↓
		1986	110

american pie	→	year	duration
		↓	↓
		1999	96

Composite partition key

• Multi value primary key —> single partition key

```
CREATE TABLE "movies" (
    title        TEXT
    year         INT,
    duration     INT,
    PRIMARY KEY ((title, year))
);
```

Note the double parentheses (more on that later)

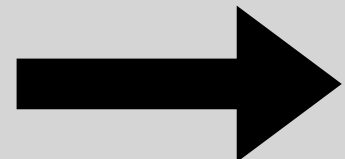# Comparison

```
CREATE TABLE "movies" (
    title           TEXT
    year            INT,
    duration        INT,
    PRIMARY KEY (title)
);
```

```
CREATE TABLE "movies" (
    title           TEXT
    year            INT,
    duration        INT,
    PRIMARY KEY ((title, year))
);
```
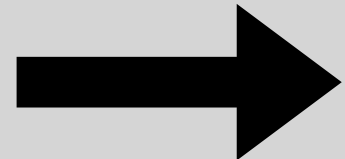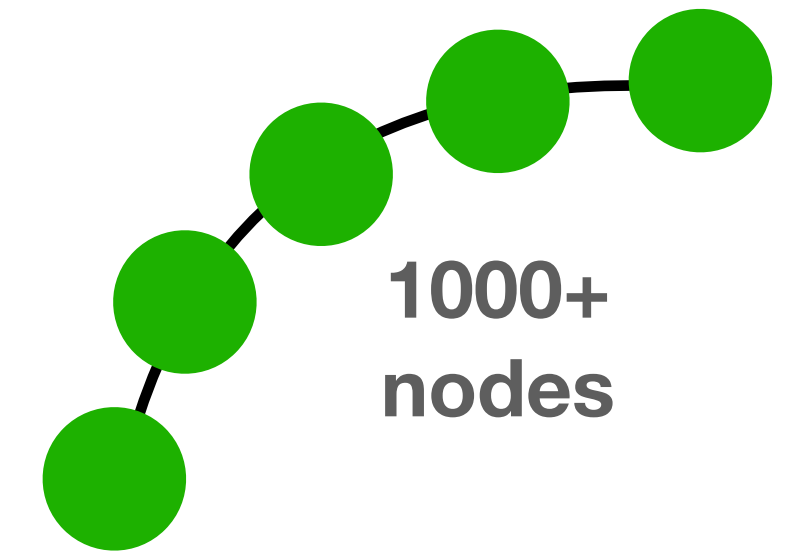
Comparison

```
SELECT * FROM movies
WHERE title = "top gun"

What happens in each case? Why?
```


1000+ nodes

```
CREATE TABLE "movies" (
    title           TEXT
    year            INT,
    duration        INT,
    PRIMARY KEY (title)
);
```

```
CREATE TABLE "movies" (
    title           TEXT
    year            INT,
    duration        INT,
    PRIMARY KEY ((title, year))
);
```

Comparison

```
SELECT * FROM movies
WHERE title = "top gun"

What happens in each case? Why?
```

1000+ nodes

```
CREATE TABLE "movie
    title          TEX
    year           INT
    duration       INT
    PRIMARY KEY (title)
);
```

```
CREATE TABLE "movies" (
    title          TEXT
    year           INT,
    duration       INT,
    PRIMARY KEY ((title, year))
);
```
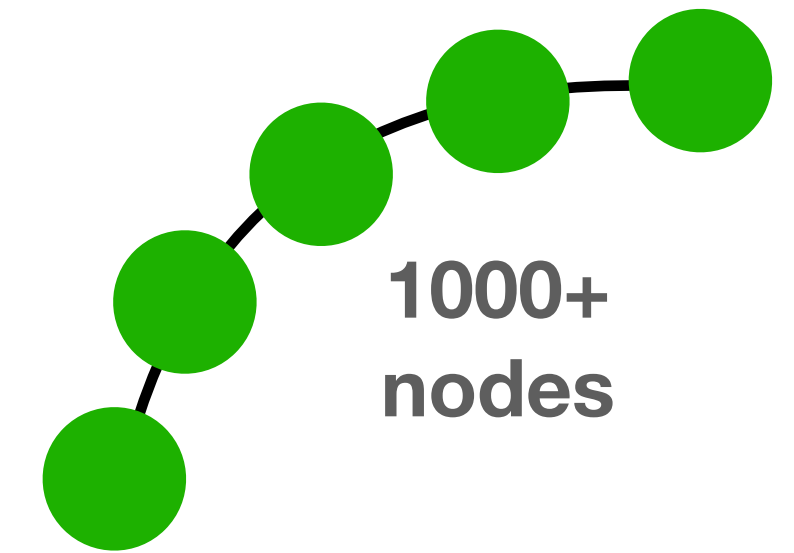
query partition key = "top gun"
—>
token = 12312
—>
data is on node 22 - GREAT

| the lion king | → | year → 2019 | duration → 118 |

| the lion king:2019 | → | duration → 118 |

| top gun | → | year → 1986 | duration → 110 |

| top gun:1986 | → | duration → 110 |

54

Comparison

```
SELECT * FROM movies
WHERE title = "top gun"

What happens in each case? Why?
```

1000+ nodes

```
CREATE TABLE "movies" (
    title        TEXT
    year         INT,
    duration     INT,
    PRIMARY KEY (title
);
```

query partition key = "top gun"
—>
error - cannot calculate token
(hash function expects <u>title + year</u>)

```
CREATE TABLE "movies" (
    title        TEXT
    year         INT,
    duration     INT,
    PRIMARY KEY ((title, year))
```

| the lion king | → | year → 2019 | duration → 118 |

| the lion king:2019 | → | duration → 118 |

| top gun | → | year → 1986 | duration → 110 |

| top gun:1986 | → | duration → 110 |

Discussion

- Can you create a statement that generates a "wide row"?

```
CREATE TABLE "movies" (
    title         TEXT
    year          INT,
    duration      INT,
    PRIMARY KEY (title)
);
```
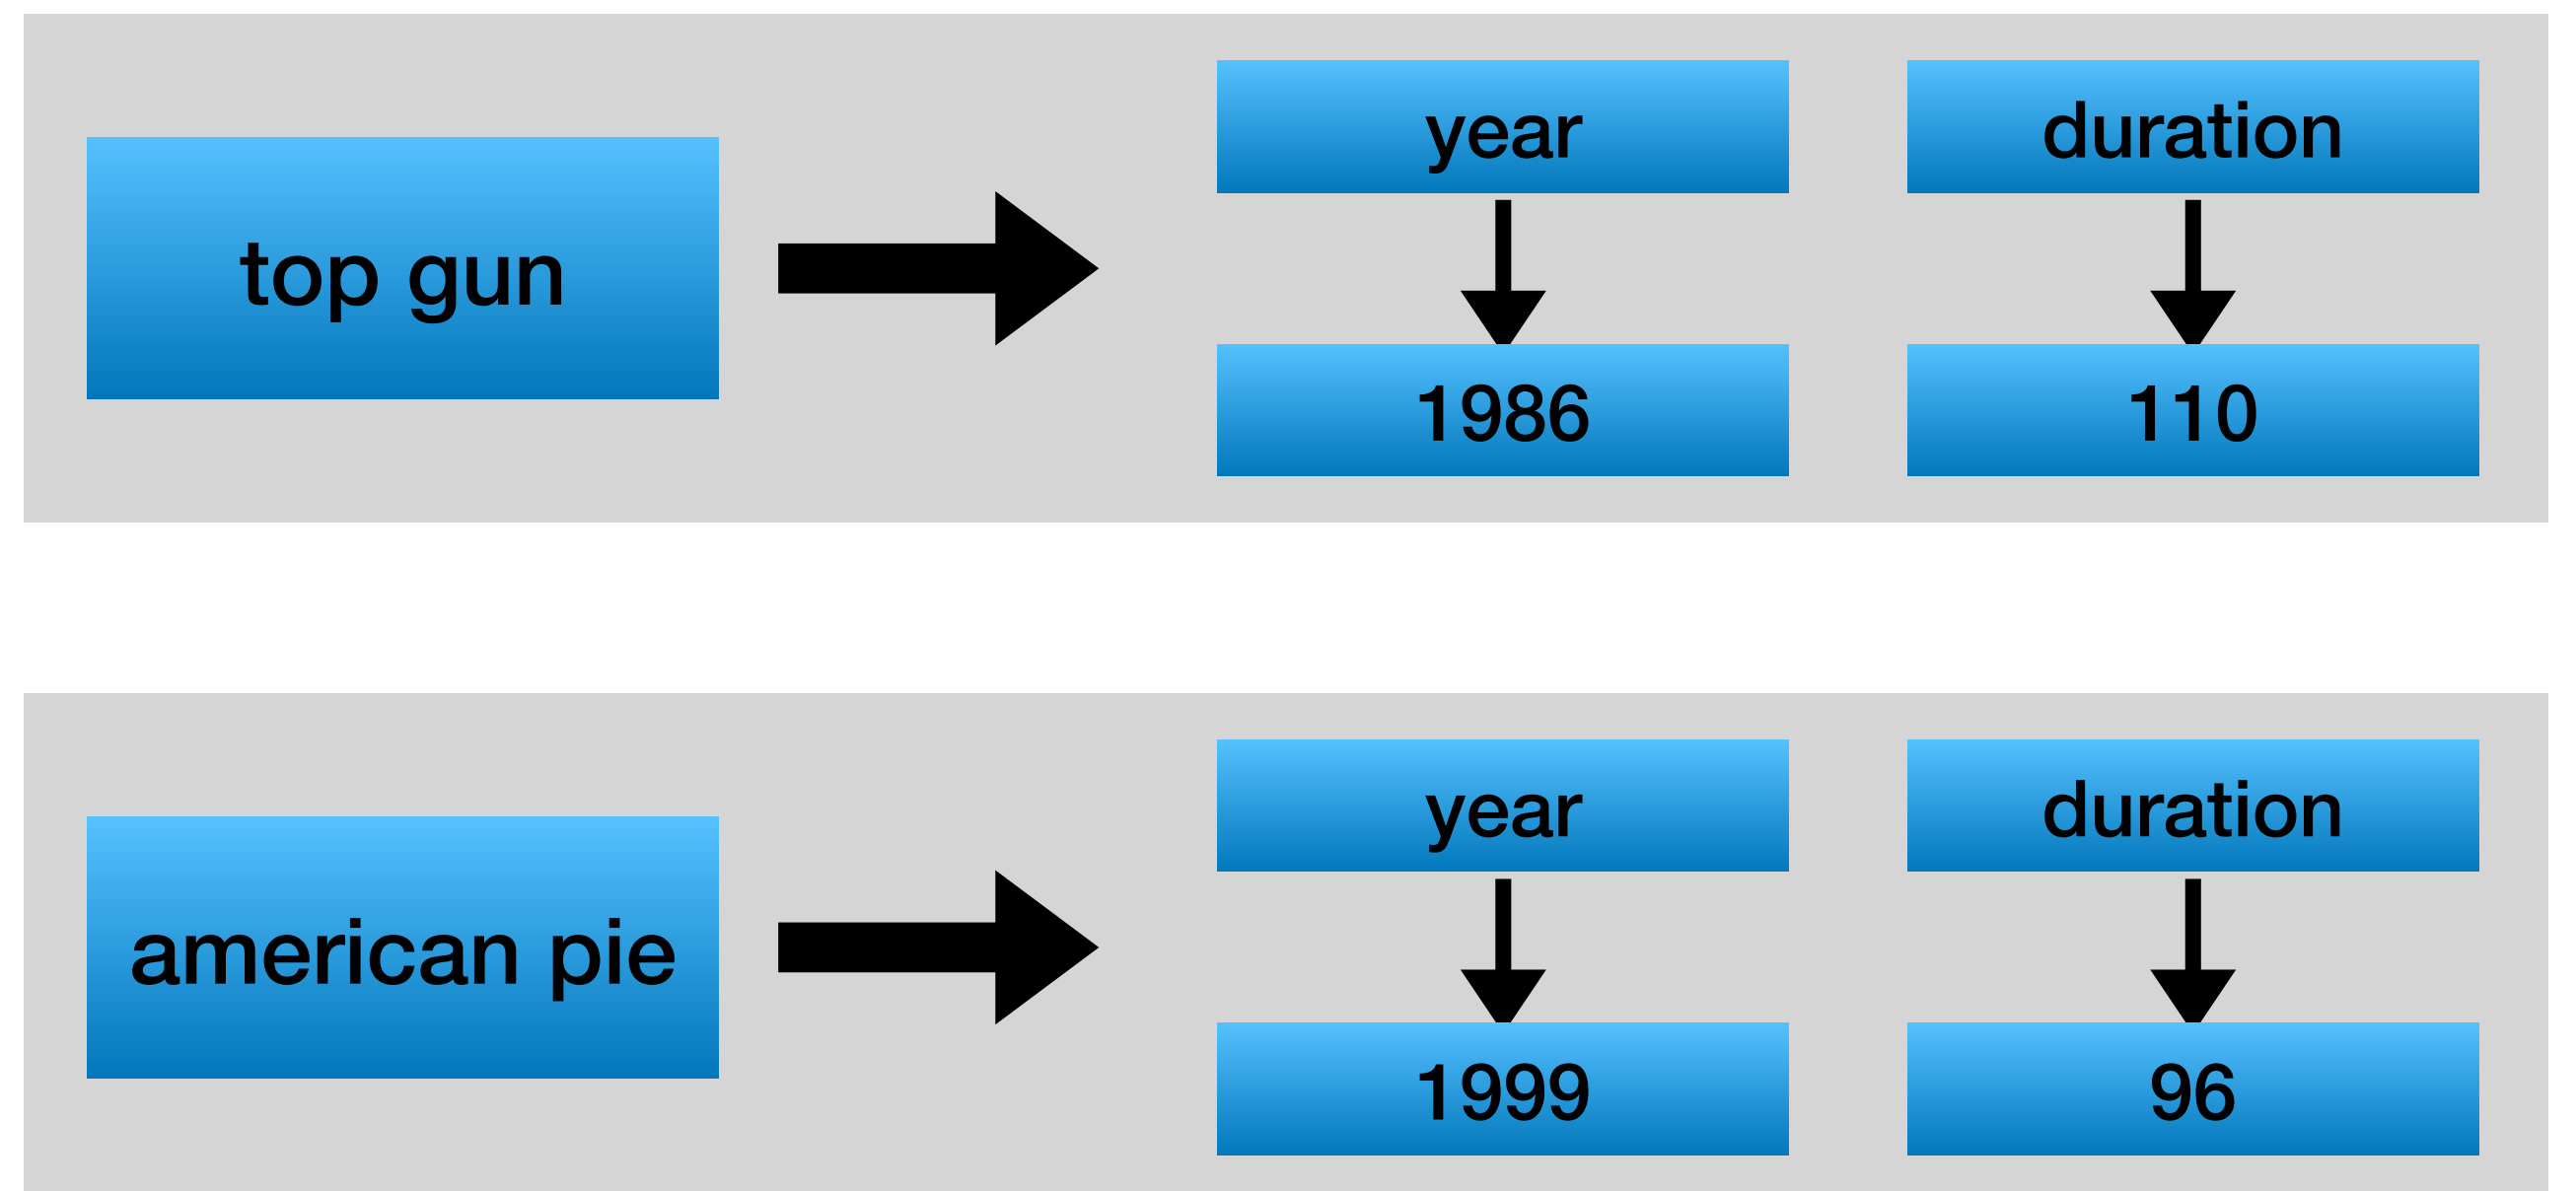
Reminder

- A table in Cassandra has

  - Partition keys

  - Clustering columns
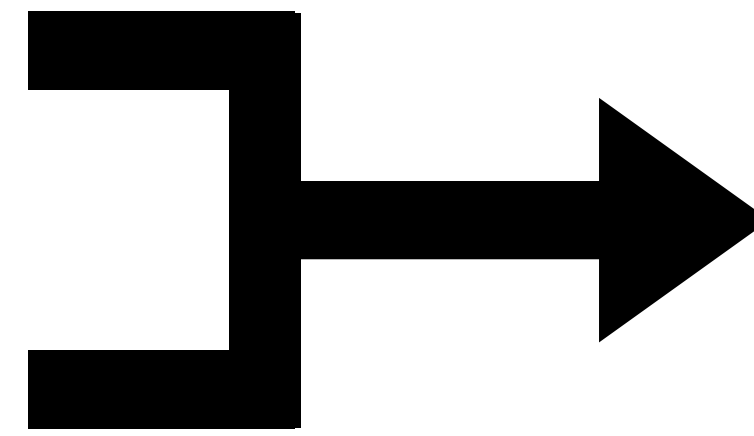
  - Columns

  Primary key

Clustering columns

- Defines the <u>order of the data</u> stored within a partition

- Part of the <u>primary key</u>

```
CREATE TABLE "movies" (
    year            INT,
    title           TEXT,
    duration        INT,
    rating          DOUBLE,
    PRIMARY KEY ((year), title)
);
```

year	title	duration	rating
1995	Bad Boys	119	6.9
1999	American Pie	96	7
1999	The Matrix	136	8.7

**double parentheses defines the partition keys,
the rest are the clustering columns**

Clustering columns

- Defines the <u>order of the data</u> stored within a partition

- Part of the <u>primary key</u>

```
CREATE TABLE "movies" (
    year            INT,
    title           TEXT,
    duration        INT,
    rating          DOUBLE,
    PRIMARY KEY ((year), title)
);
```

**double parentheses defines the partition keys,
the rest are the clustering columns**

year	title	duration	rating
1995	Bad Boys	119	6.9
1999	American Pie	96	7
1999	The Matrix	136	8.7

Sorted within a **single** partition

Clustering columns - format on disk

year	title	duration	rating
1995	Bad Boys	119	6.9
1999	American Pie	96	7
1999	The Matrix	136	8.7

```
CREATE TABLE "movies" (
    year                INT,
    title               TEXT,
    duration            INT,
    rating              DOUBLE,
    PRIMARY KEY ((year), title)
);
```

Clustering columns - format on disk

year	title	duration	rating
1995	Bad Boys	119	6.9
1999	American Pie	96	7
1999	The Matrix	136	8.7

```
CREATE TABLE "movies" (
    year              INT,
    title             TEXT,
    duration          INT,
    rating            DOUBLE,
    PRIMARY KEY ((year), title)
);
```

1995 → Bad Boys:duration → 119   Bad Boys:rating → 6.9

Clustering columns - format on disk

year	title	duration	rating
1995	Bad Boys	119	6.9
1999	American Pie	96	7
1999	The Matrix	136	8.7

```
CREATE TABLE "movies" (
    year            INT,
    title           TEXT,
    duration        INT,
    rating          DOUBLE,
    PRIMARY KEY ((year), title)
);
```

1995 →

Bad Boys:duration → 119

Bad Boys:rating → 6.9

The clustering columns are the prefix of the columns

Clustering columns - format on disk

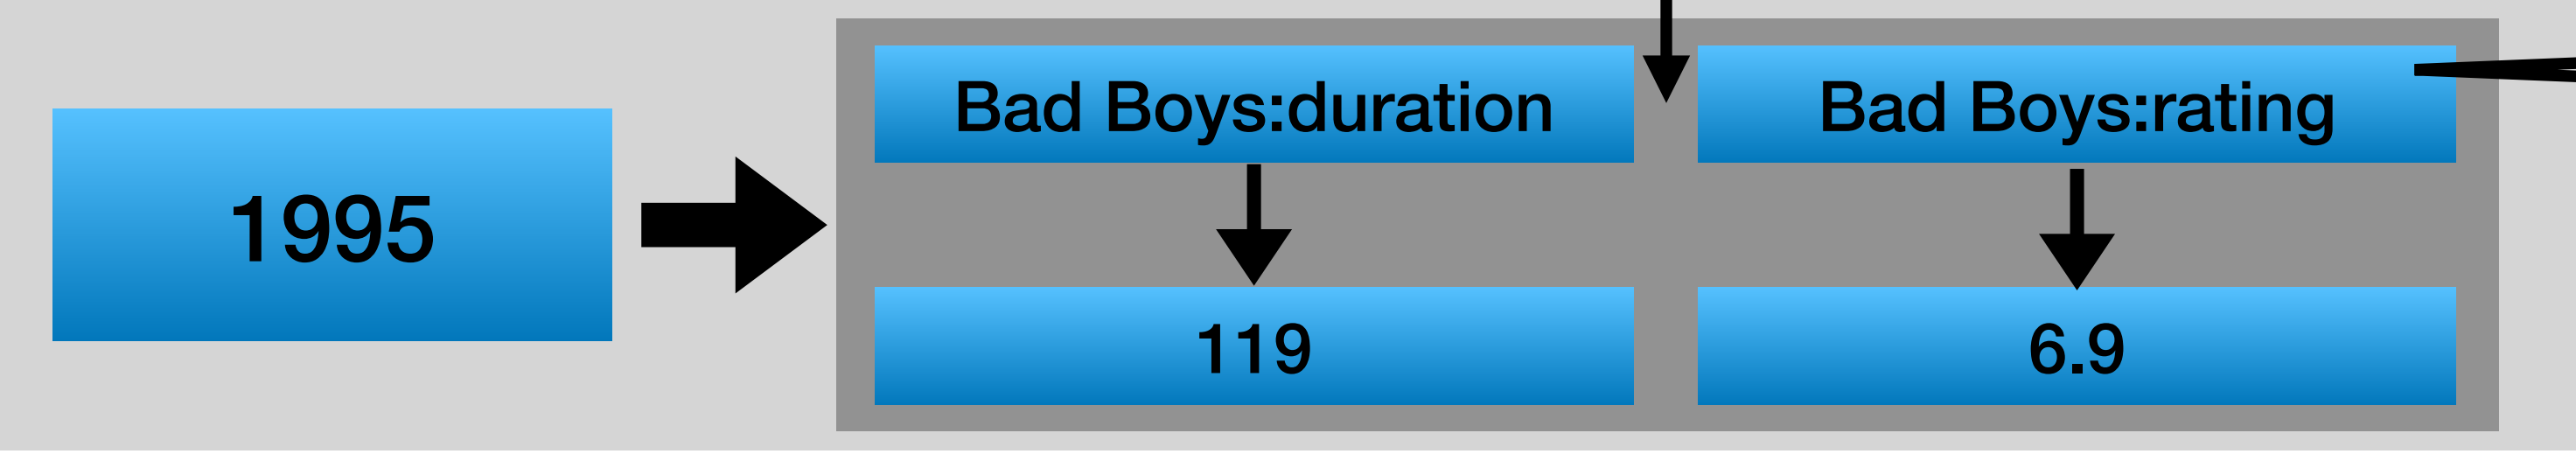| year	title	duration	rating
1995	Bad Boys	119	6.9
1999	American Pie	96	7
1999	The Matrix	136	8.7

```
CREATE TABLE "movies" (
    year            INT,
    title           TEXT,
    duration        INT,
    rating          DOUBLE,
    PRIMARY KEY ((year), title)
);
```

Clustering columns - format on disk

year	title	duration	rating
1995	Bad Boys	119	6.9
1999	American Pie	96	7
1999	The Matrix	136	8.7

```
CREATE TABLE "movies" (
    year            INT,
    title           TEXT,
    duration        INT,
    rating          DOUBLE,
    PRIMARY KEY ((year), title)
);
```
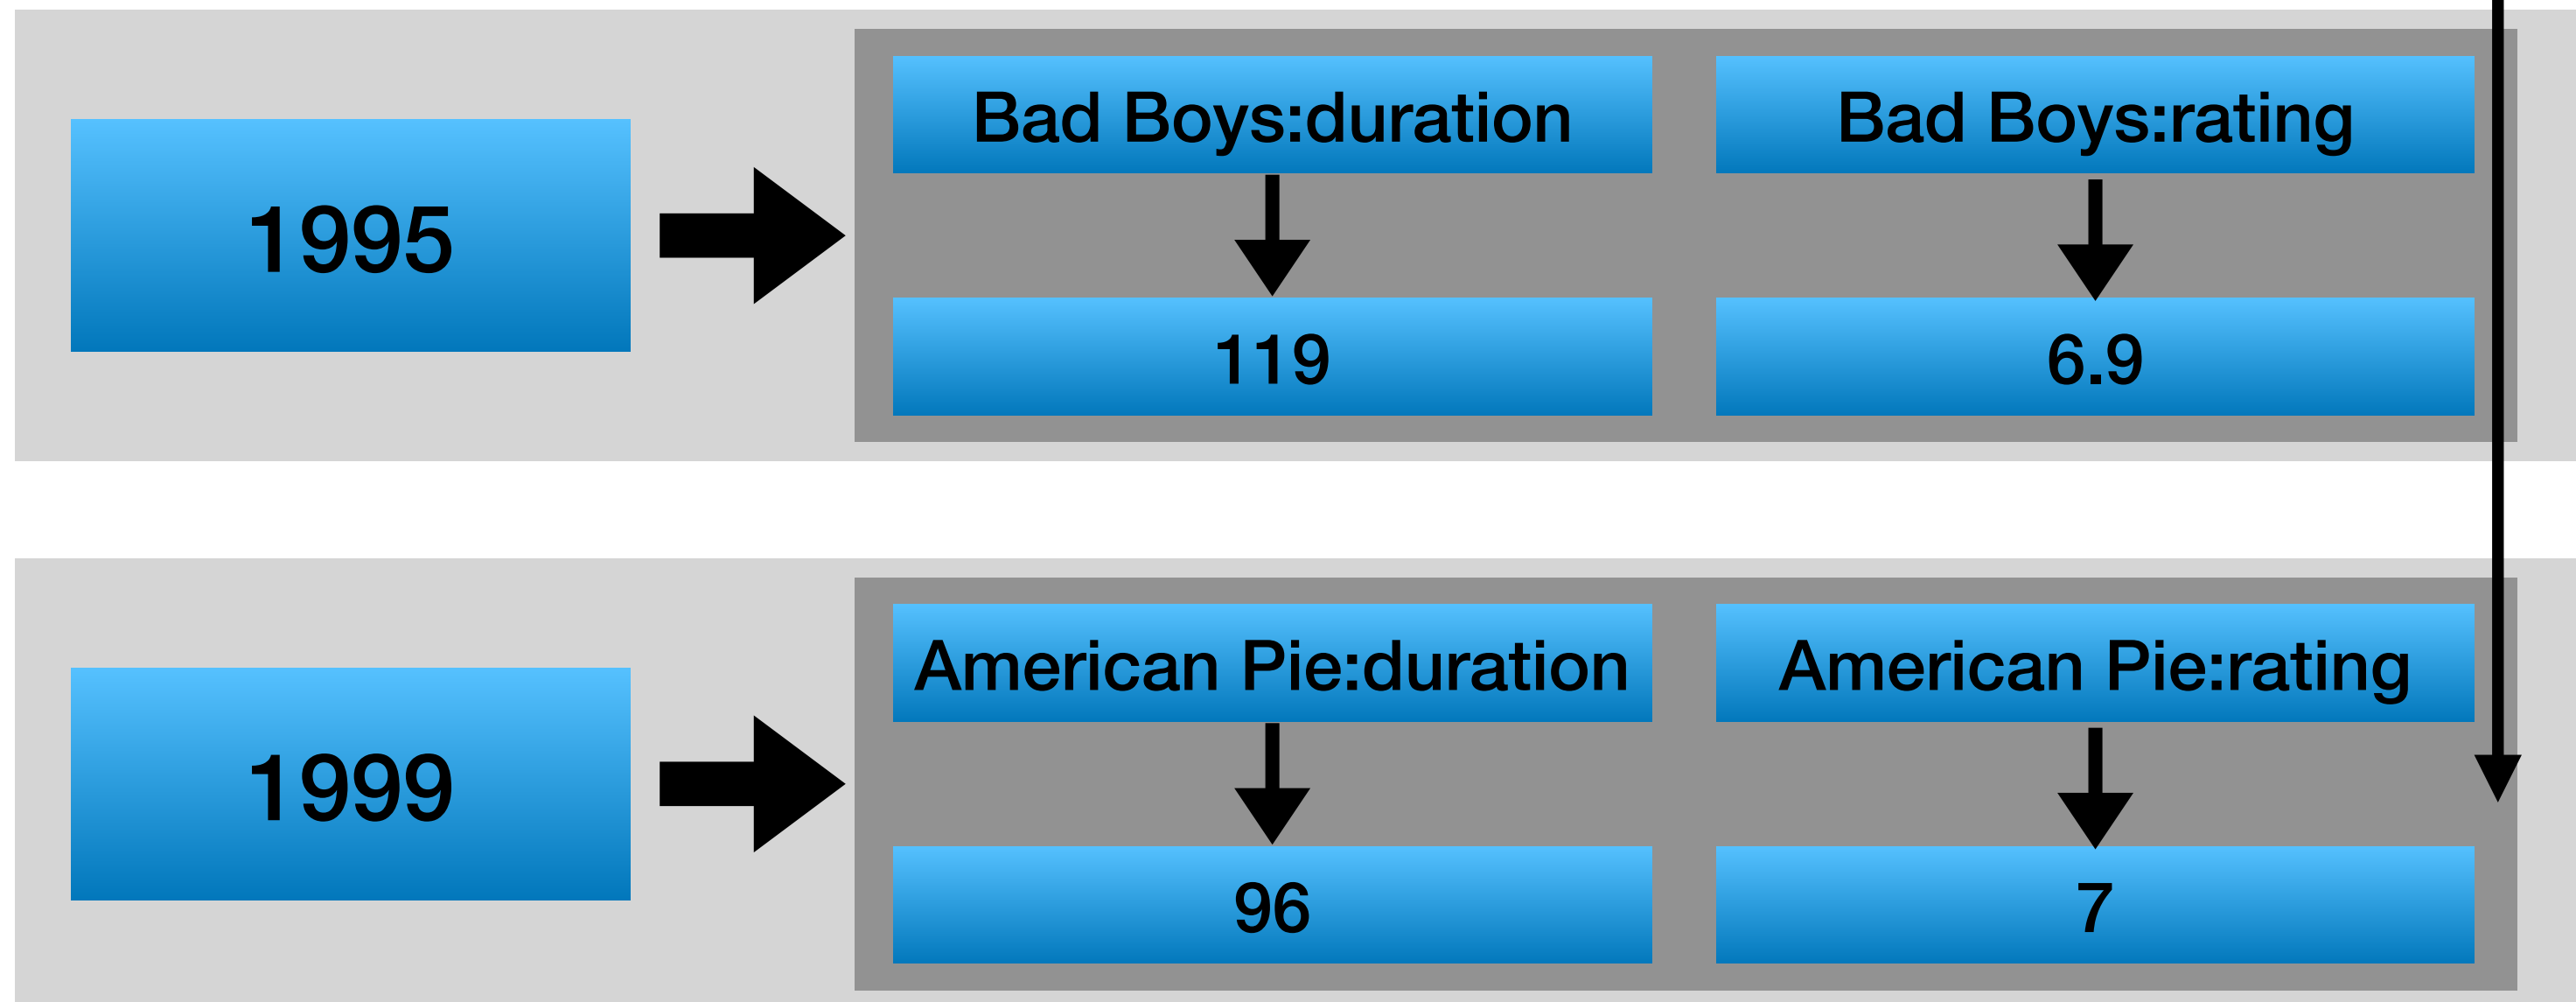
1995 →

Bad Boys:duration	Bad Boys:rating
↓	↓
119	6.9

1999 →

American Pie:duration	American Pie:rating	The Matrix:duration	The Matrix:rating
↓	↓	↓	↓
96	7	136	8.7

Clustering columns - format on disk

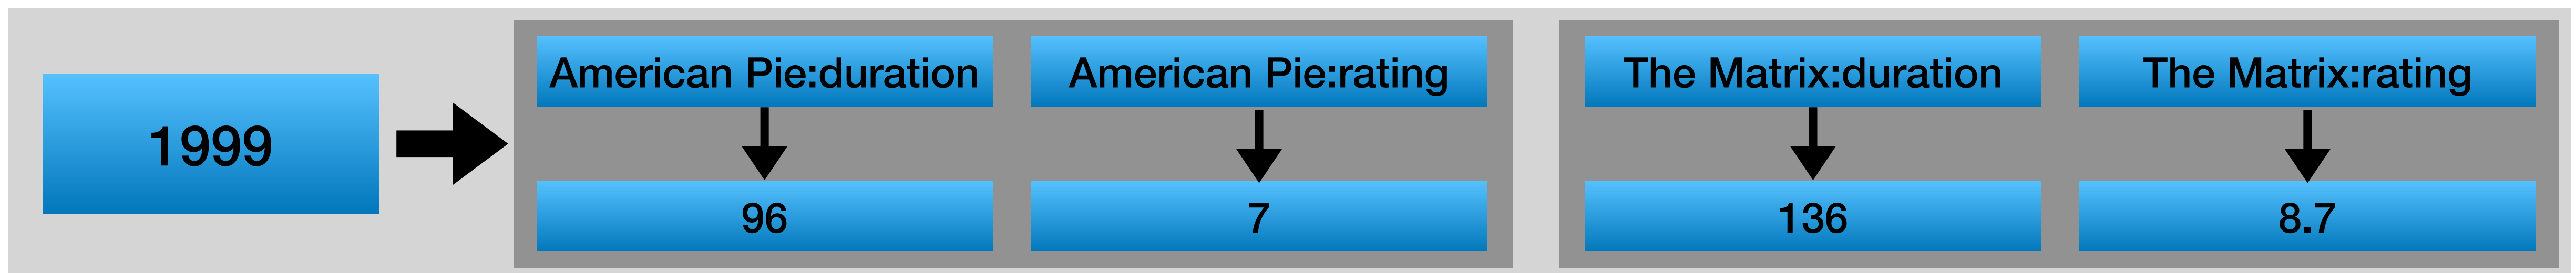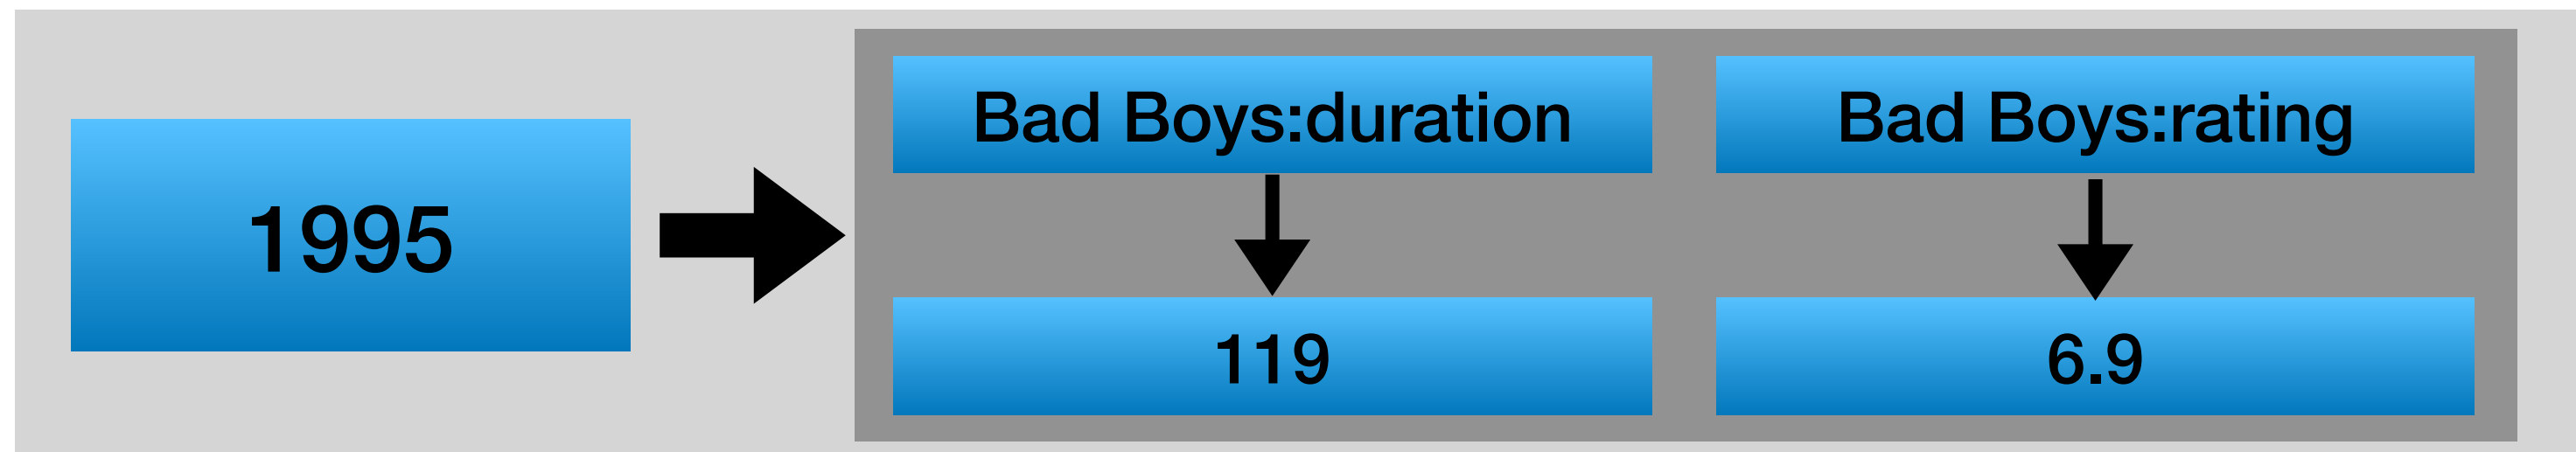| year	title	duration	rating
1995	Bad Boys	119	6.9
1999	American Pie	96	7
1999	The Matrix	136	8.7

```
CREATE TABLE "movies" (
    year            INT,
    title           TEXT,
    duration        INT,
    rating          DOUBLE,
    PRIMARY KEY ((year), title)
);
```

Can this statement generate wide partitions?

1995 →

Bad Boys:duration	Bad Boys:rating
↓	↓
119	6.9

1999 →

American Pie:duration	American Pie:rating	The Matrix:duration	The Matrix:rating
↓	↓	↓	↓
96	7	136	8.7

Clustering columns - format on disk

year	title	duration	rating
1995	Bad Boys	119	6.9
1999	American Pie	96	7
1999	The Matrix	136	8.7

```
CREATE TABLE "movies" (
    year            INT,
    title           TEXT,
    duration        INT,
    rating          DOUBLE,
    PRIMARY KEY ((year), title)
);
```

Can this statement generate wide partitions?

YES!
For example, what will happen if there are 10k movies in 1999?

oys:rating

6.9

1999

American Pie:duration	American Pie:rating		The Matrix:duration	The Matrix:rating
96	7		136	8.7

Clustering columns - format on disk

year	title	duration	rating
1995	Bad Boys	119	6.9
1999	American Pie	96	7
1999	The Matrix	136	8.7

```
CREATE TABLE "movies" (
    year            INT,
    title           TEXT,
    duration        INT,
    rating          DOUBLE,
    PRIMARY KEY ((year), title)
);
```

Can this statement generate wide partitions?

Clustering columns - format on disk

year	title	duration	rating
1995	Bad Boys	119	6.9
1999	American Pie	96	7
1999	Star Wars I	136	6.5
1999	The Matrix	136	8.7

```
CREATE TABLE "movies" (
    year              INT,
    title             TEXT,
    duration          INT,
    rating            DOUBLE,
    PRIMARY KEY ((year), title)
);
```
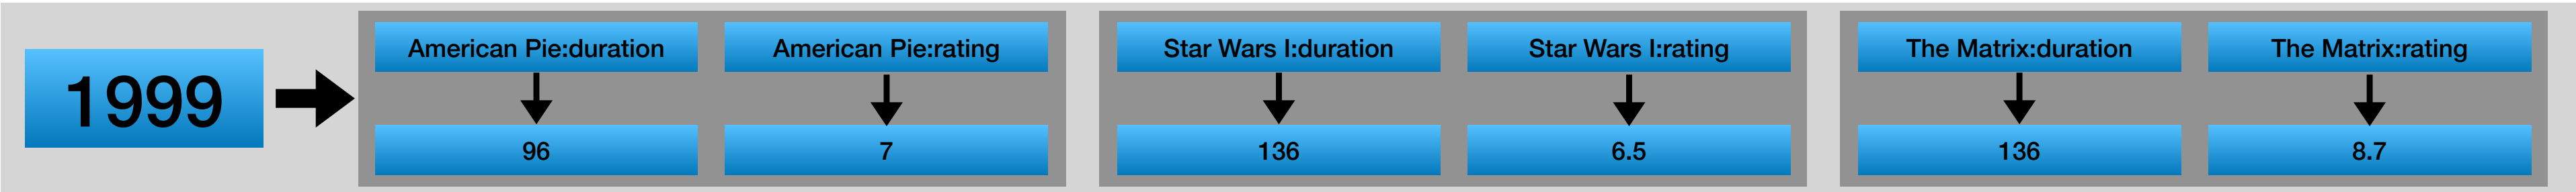
Can this statement generate wide partitions?

1995 →

Bad Boys:duration	Bad Boys:rating
↓	↓
119	6.9

1999 →

American Pie:duration	American Pie:rating	Star Wars I:duration	Star Wars I:rating	The Matrix:duration	The Matrix:rating
↓	↓	↓	↓	↓	↓
96	7	136	6.5	136	8.7

Clustering columns - format on disk

year	title	duration	rating
1995	Bad Boys	119	6.9
1999	American Pie	96	7
1999	Star Wars I	136	6.5
1999	The Matrix	136	8.7
1999	…	…	…

```
CREATE TABLE "movies" (
    year            INT,
    title           TEXT,
    duration        INT,
    rating          DOUBLE,
    PRIMARY KEY ((year), title)
);
```
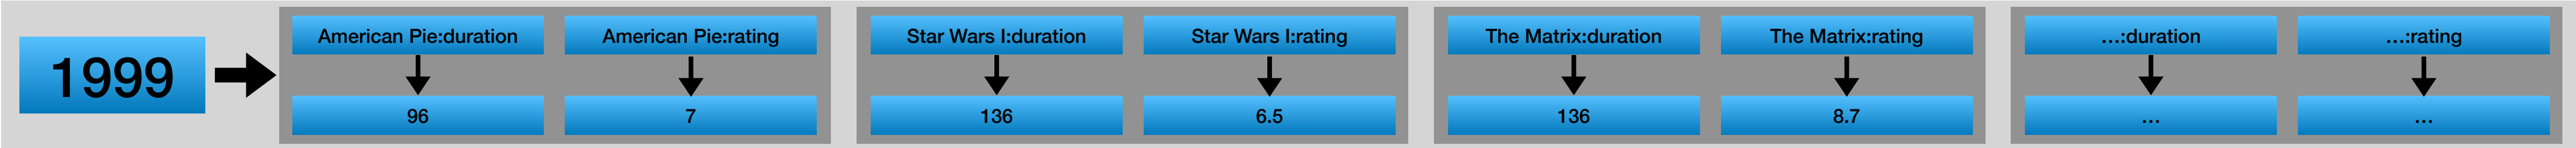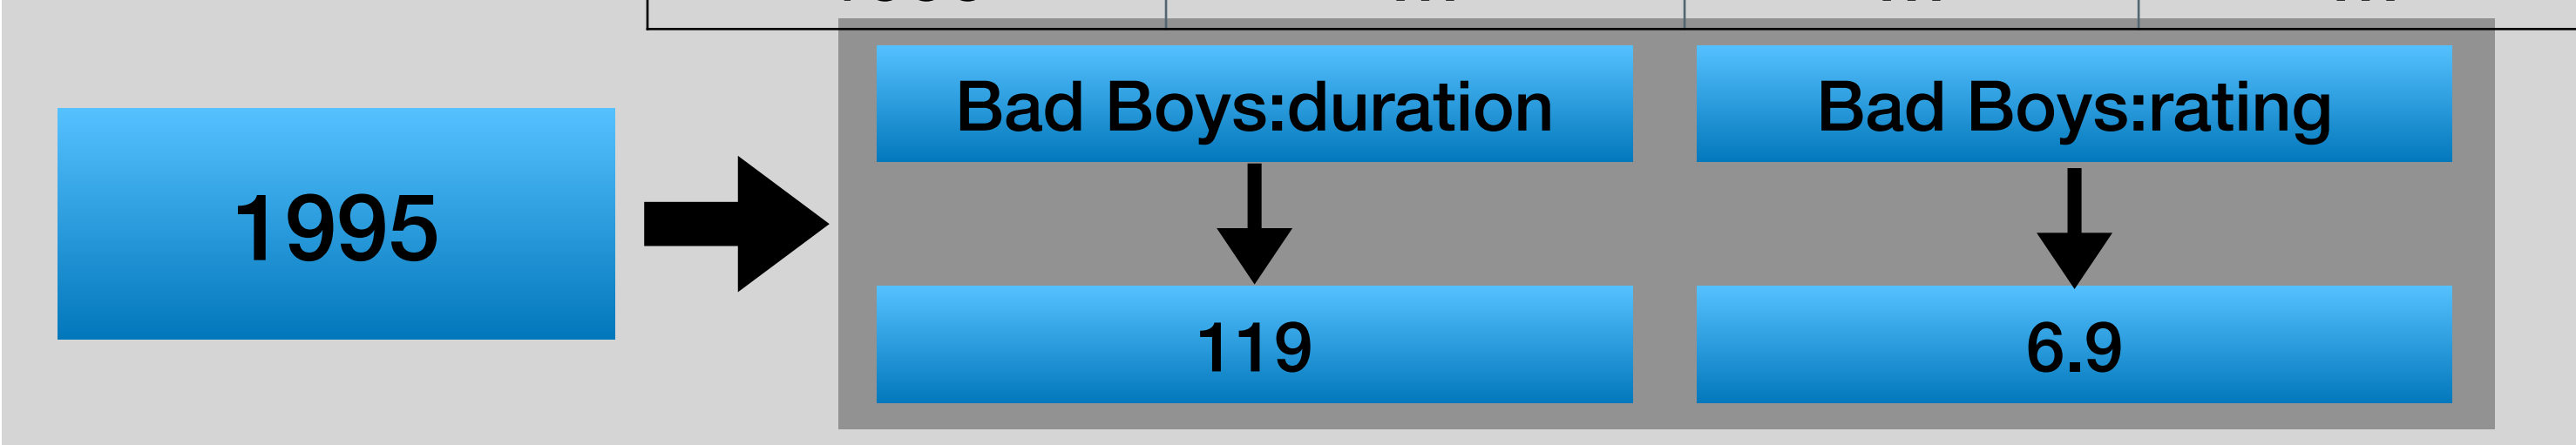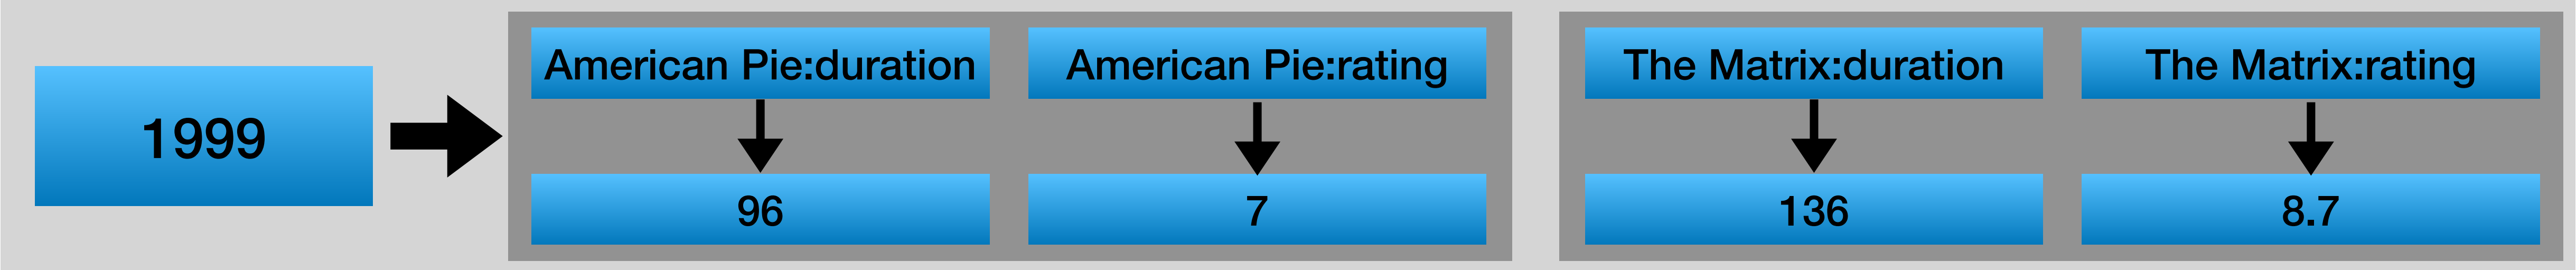
Can this statement generate wide partitions?

1995 → Bad Boys:duration → 119 | Bad Boys:rating → 6.9

1999 → American Pie:duration → 96 | American Pie:rating → 7 | Star Wars I:duration → 136 | Star Wars I:rating → 6.5 | The Matrix:duration → 136 | The Matrix:rating → 8.7 | …:duration → … | …:rating → …

Clustering columns - format on disk (2)

What is the difference?

year	title	duration	rating
1994	The lion king	119	8.5
2019	The lion king	118	6.9
1999	American Pie	96	7
1999	The Matrix	136	8.7

Clustering columns - form

1999	→	American Pie:duration	American Pie:rating		
		↓	↓		
		96	7	36	8.7

What is the difference?

The lion king	→	1994:duration	1994:rating	2019:duration	2019:rating
		↓	↓		
		119	8.5		

year	title	duration	rating
1994	The lion king	119	8.5
2019	The lion king	118	6.9
1999	American Pie	96	7
1999	The Matrix	136	8.7

```
CREATE TABLE "movies" (
    year            INT,
    title           TEXT,
    duration        INT,
    rating          DOUBLE,
    PRIMARY KEY ((title), year)
);
```

71

Clustering columns - form

1999 →

American Pie:duration	American Pie:rating
96	7

```
CREATE TABLE "movies" (
    year            INT,
    title           TEXT,
    duration        INT,
    rating          DOUBLE,
    PRIMARY KEY ((year), title)
);
```

36	8.7

When to use which version?

The lion king →

1994:duration	1994:rating
119	8.5

2019:duration	2019:rating

year	title	duration	rating
1994	The lion king	119	8.5
2019	The lion king	118	6.9
1999	American Pie	96	7
1999	The Matrix	136	8.7

```
CREATE TABLE "movies" (
    year            INT,
    title           TEXT,
    duration        INT,
    rating          DOUBLE,
    PRIMARY KEY ((title), year)
);
```

Clustering columns - form


```
CREATE TABLE "movies" (
    year              INT,
    title             TEXT,
    duration          INT,
    rating            DOUBLE,
    PRIMARY KEY ((year), title)
);
```
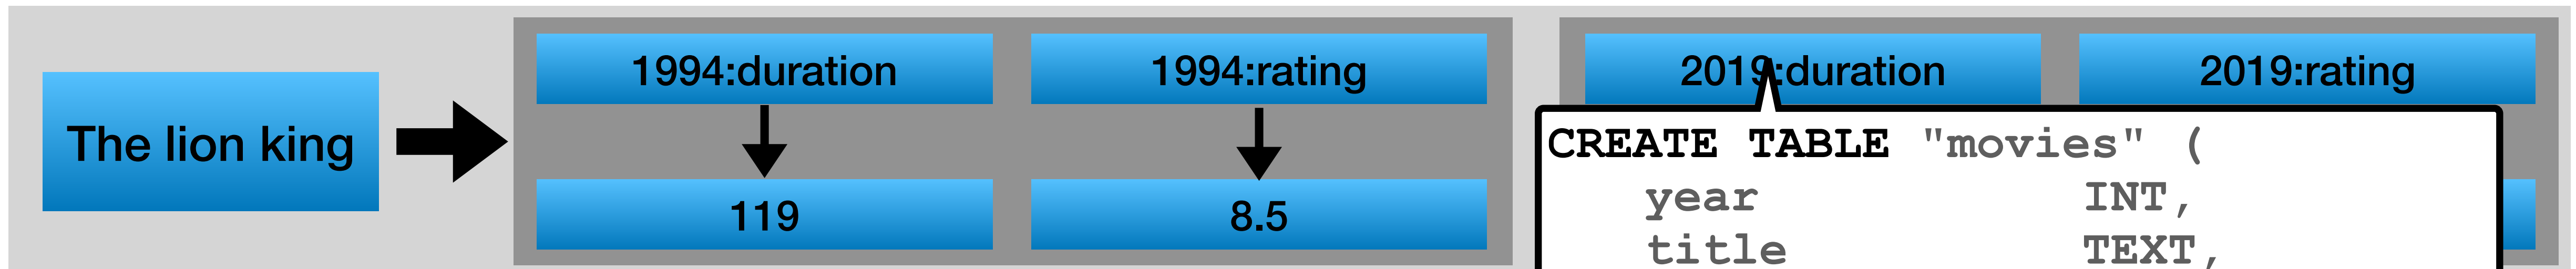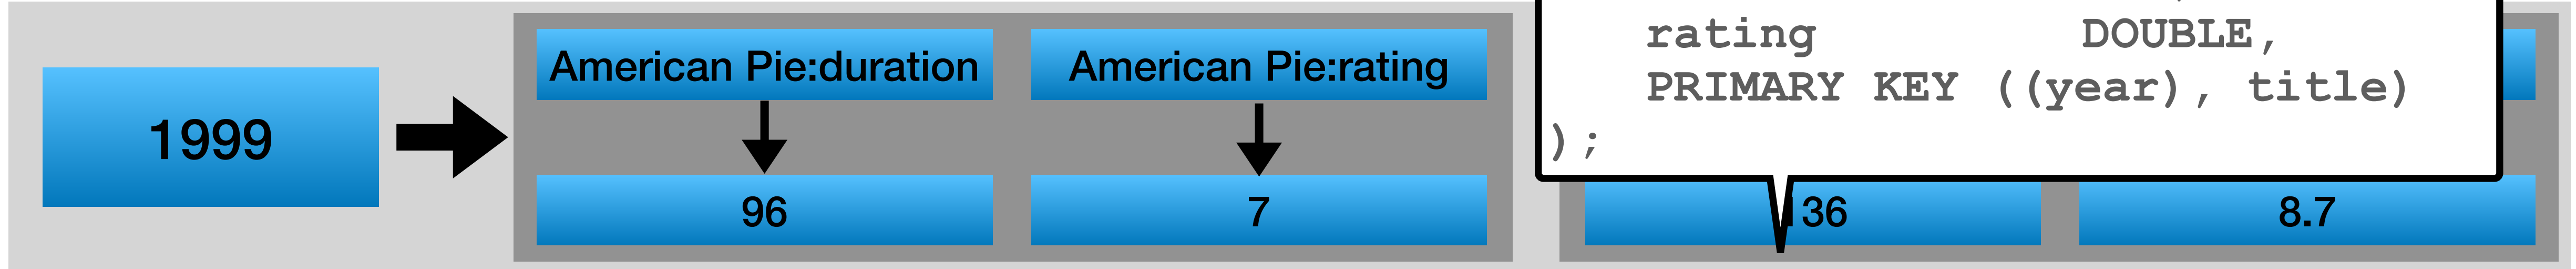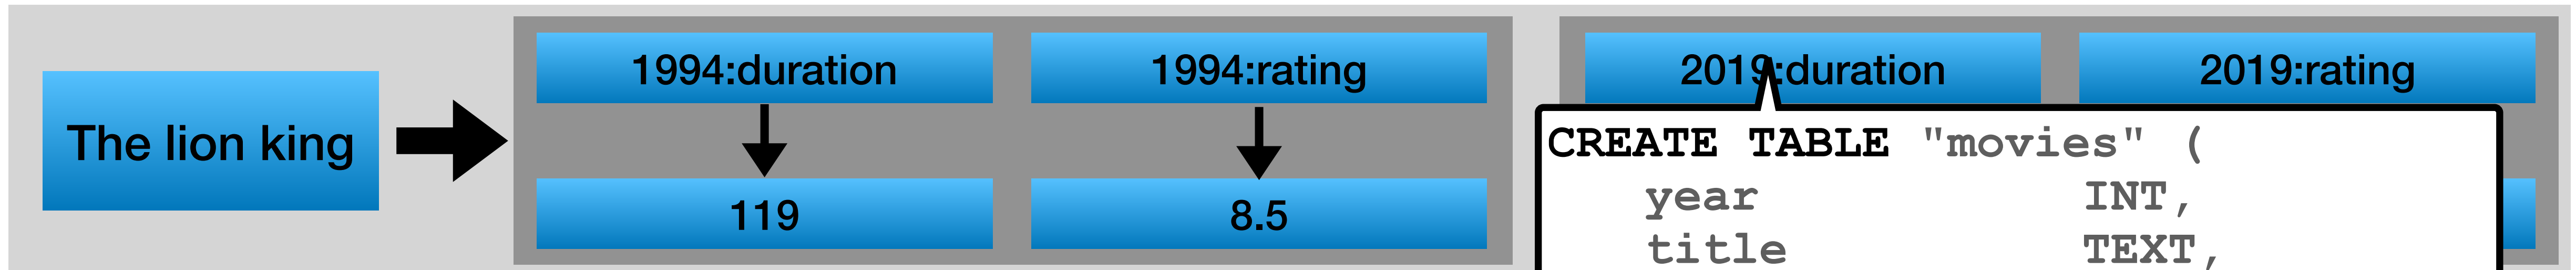
| | American Pie:duration | | American Pie:rating |
1999 → 96 | 7 | 36 | 8.7

When to use which version?

Points to consider
1. The way we will query the data
2. The size of the partition

The lion king →

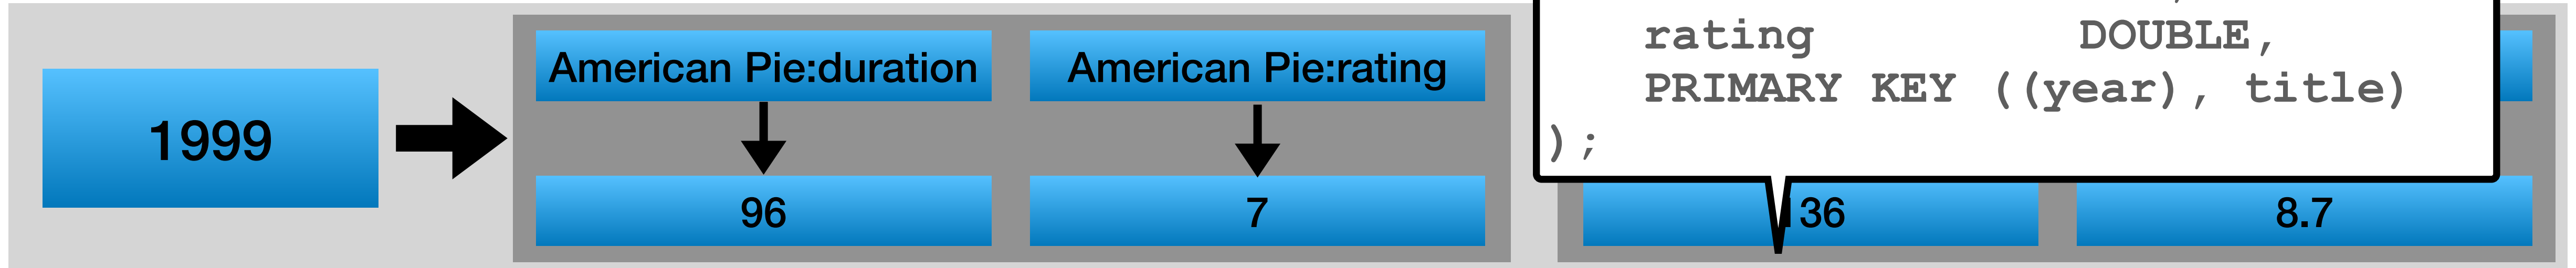| 1994:duration | 1994:rating | 2019:duration | 2019:rating |
| 119 | 8.5 |

```
CREATE TABLE "movies" (
    year              INT,
    title             TEXT,
    duration          INT,
    rating            DOUBLE,
    PRIMARY KEY ((title), year)
);
```

year	title	duration	rating
1994	The lion king	119	8.5
2019	The lion king	118	6.9
1999	American Pie	96	7
1999	The Matrix	136	8.7

73

Clustering columns - form

```
CREATE TABLE "movies" (
    year            INT,
    title           TEXT,
    duration        INT,
    rating          DOUBLE,
    PRIMARY KEY ((year), title)
);
```

| 1999 | → | American Pie:duration | American Pie:rating | | 36 | 8.7 |

| | | 96 | 7 | | | |

If we know the year and the title at query time, we can query both versions

...en to use which version?

Points to consider
1. The way we will query the data
2. The size of the partition

| The lion king | → | 1994:duration | 1994:rating | | 2019:duration | 2019:rating |

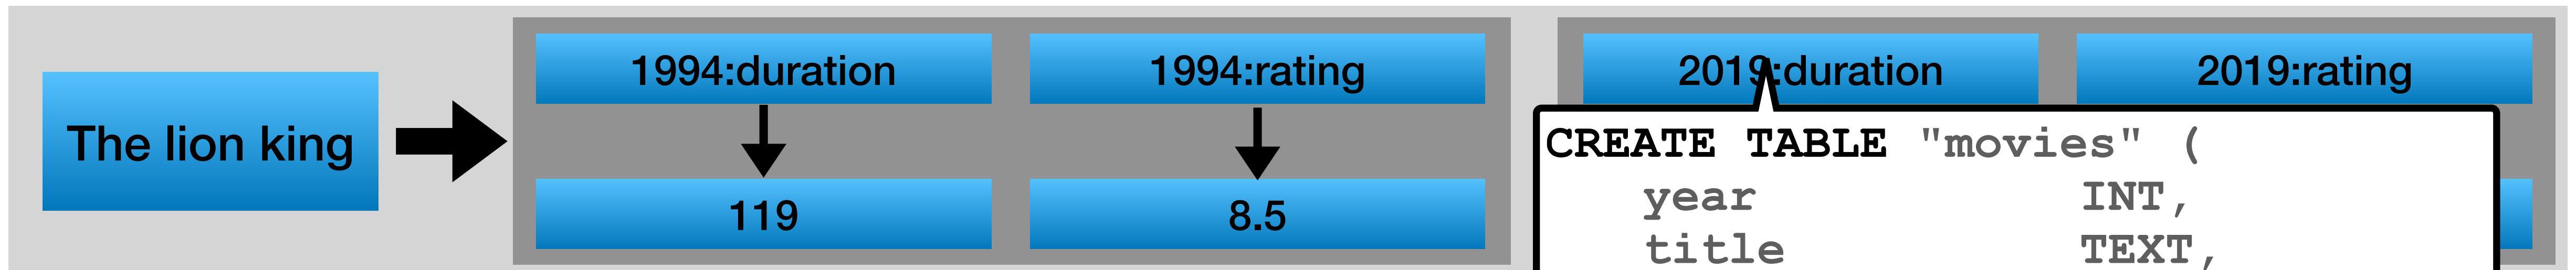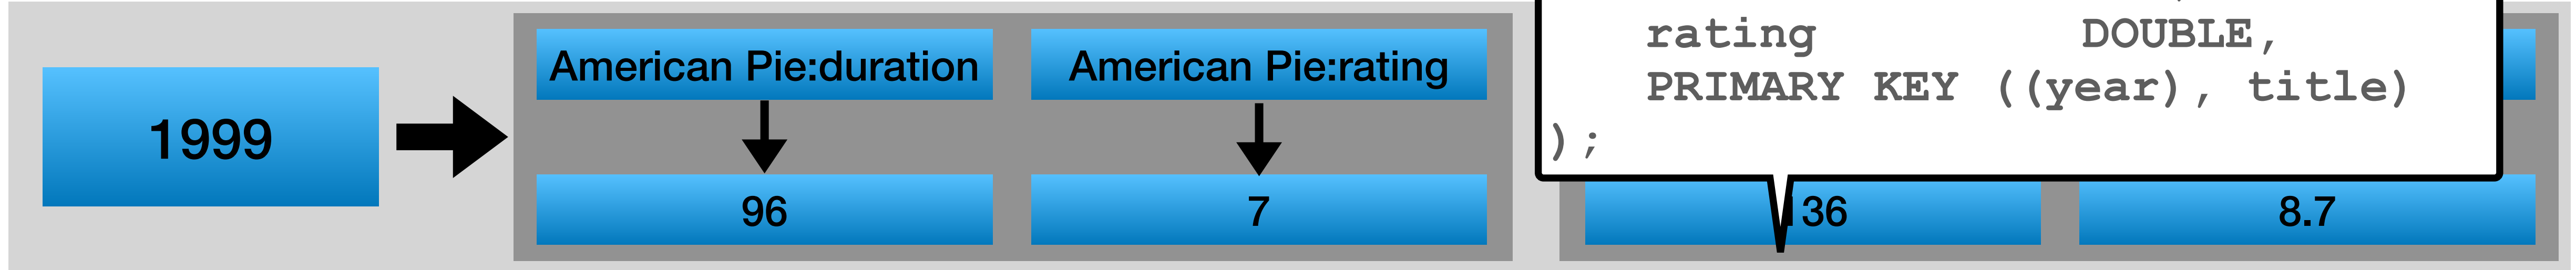| | | 119 | 8.5 | | | |

```
CREATE TABLE "movies" (
    year            INT,
    title           TEXT,
    duration        INT,
    rating          DOUBLE,
    PRIMARY KEY ((title), year)
);
```

year	title	duration	rating
1994	The lion king	119	8.5
2019	The lion king	118	6.9
1999	American Pie	96	7
1999	The Matrix	136	8.7

Clustering columns - form



```
CREATE TABLE "movies" (
    year            INT,
    title           TEXT,
    duration        INT,
    rating          DOUBLE,
    PRIMARY KEY ((year), title)
);
```

1999 → American Pie:duration → 96 | American Pie:rating → 7 | 36 | 8.7

If we want to query by year

en to use which version?

Points to consider
1. The way we will query the data
2. The size of the partition

The lion king → 1994:duration → 119 | 1994:rating → 8.5 | 2019:duration | 2019:rating
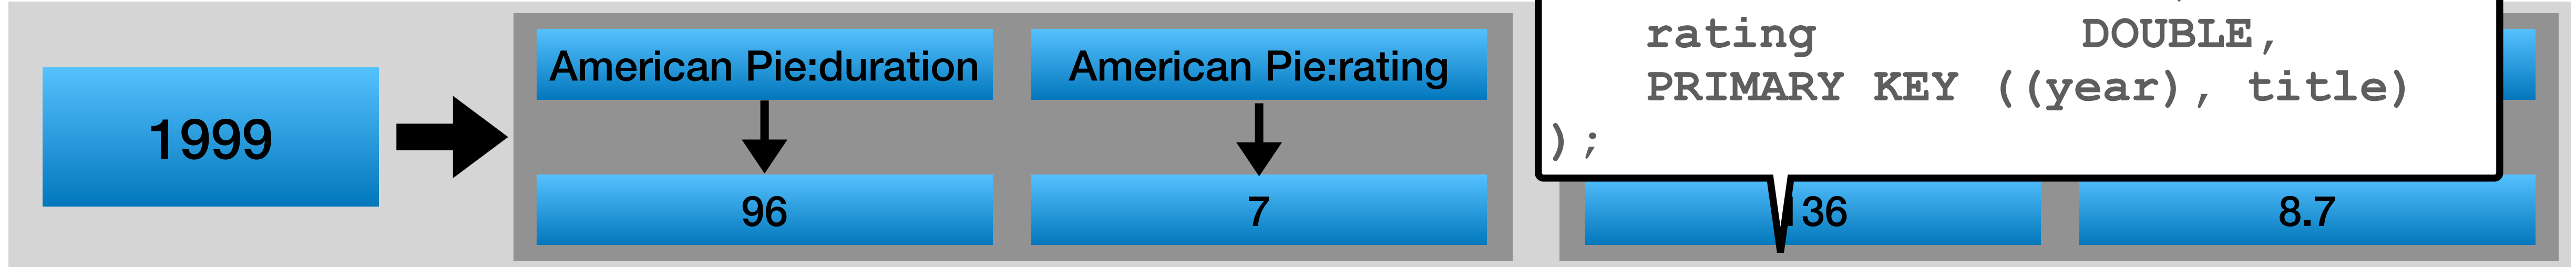
```
CREATE TABLE "movies" (
    year            INT,
    title           TEXT,
    duration        INT,
    rating          DOUBLE,
    PRIMARY KEY ((title), year)
);
```

year	title	duration	rating
1994	The lion king	119	8.5
2019	The lion king	118	6.9
1999	American Pie	96	7
1999	The Matrix	136	8.7

75

Clustering columns - form



```
CREATE TABLE "movies" (
    year            INT,
    title           TEXT,
    duration        INT,
    rating          DOUBLE,
    PRIMARY KEY ((year), title)
);
```

| 1999 | American Pie:duration | American Pie:rating | | 36 | 8.7 |

en to use which version?

If we want to query by title

Points to consider
1. The way we will query the data
2. The size of the partition

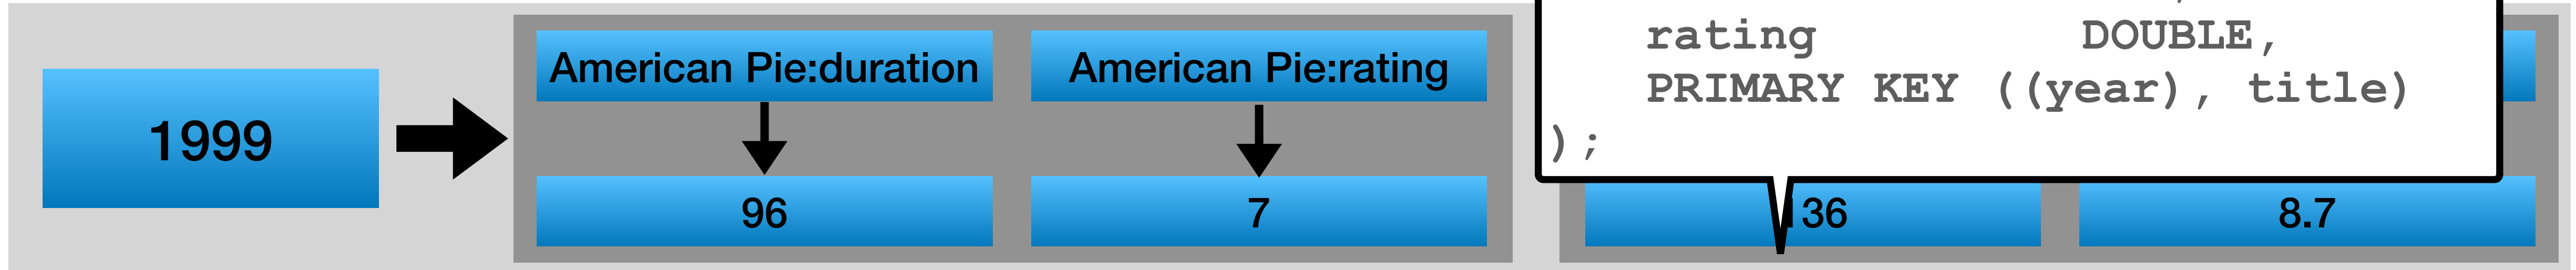| The lion king | 1994:duration | 1994:rating | | 2019:duration | 2019:rating |
| 119 | 8.5 |

```
CREATE TABLE "movies" (
    year            INT,
    title           TEXT,
    duration        INT,
    rating          DOUBLE,
    PRIMARY KEY ((title), year)
);
```

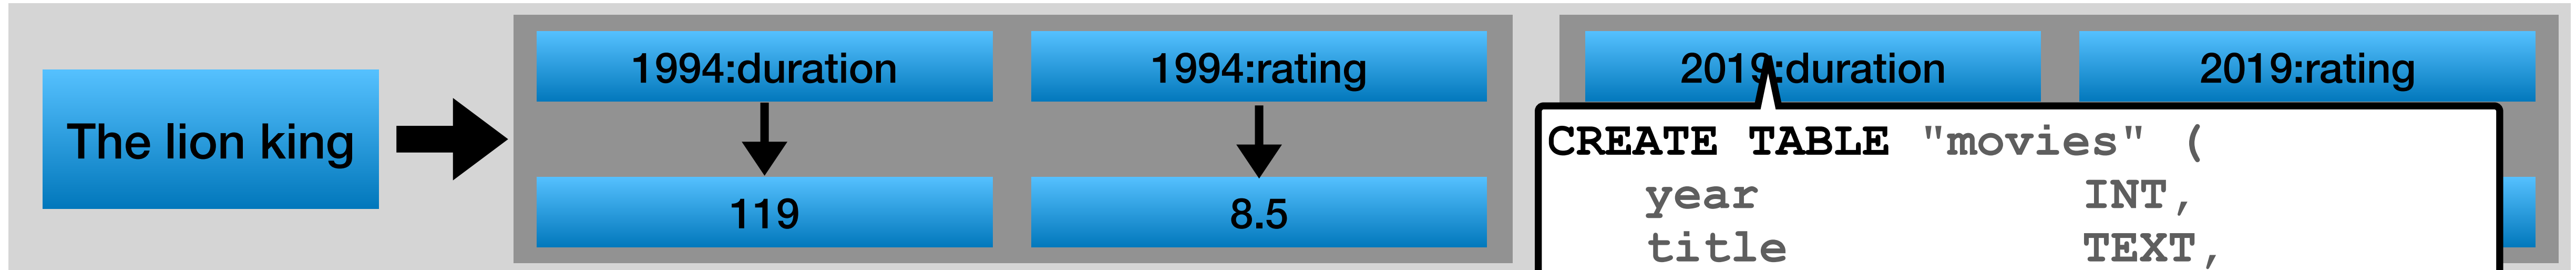| year	title	duration	rating
1994	The lion king	119	8.5
2019	The lion king	118	6.9
1999	American Pie	96	7
1999	The Matrix	136	8.7

76

Clustering columns - form

```
CREATE TABLE "movies" (
    year            INT,
    title           TEXT,
    duration        INT,
    rating          DOUBLE,
    PRIMARY KEY ((year), title)
);
```

1999 → | American Pie:duration | American Pie:rating | ... |
         | 96 | 7 | 36 | 8.7 |

Which version will probably create bigger partitions? ⟨

Points to consider
1. The way we will query the data
2. The size of the partition

The lion king → | 1994:duration | 1994:rating | 2019:duration | 2019:rating |
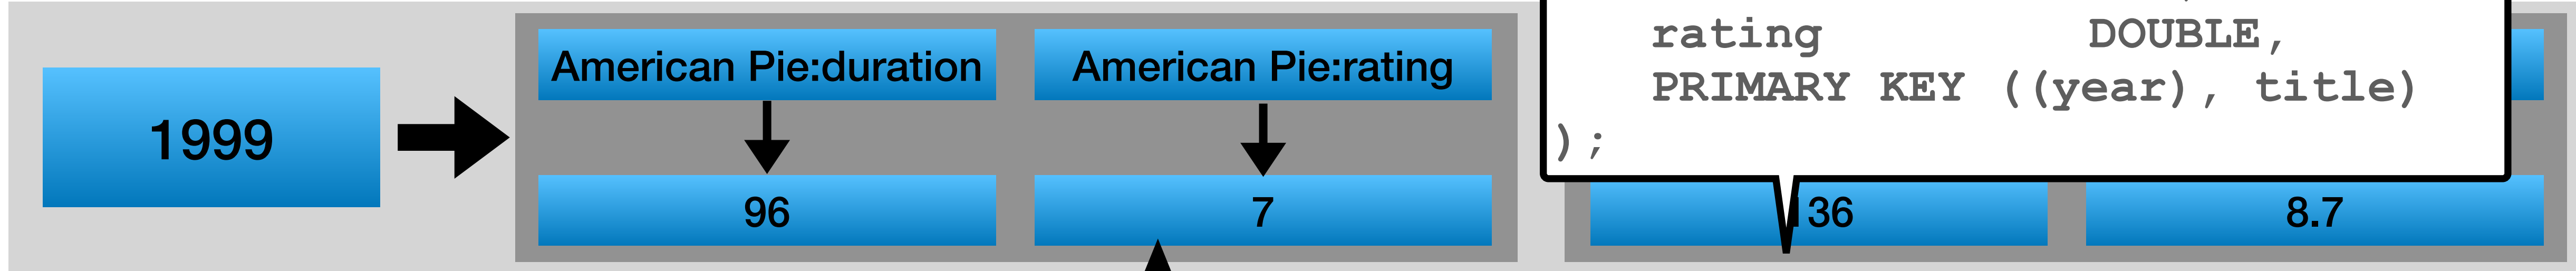 | 119 | 8.5 | | |

```
CREATE TABLE "movies" (
    year            INT,
    title           TEXT,
    duration        INT,
    rating          DOUBLE,
    PRIMARY KEY ((title), year)
);
```

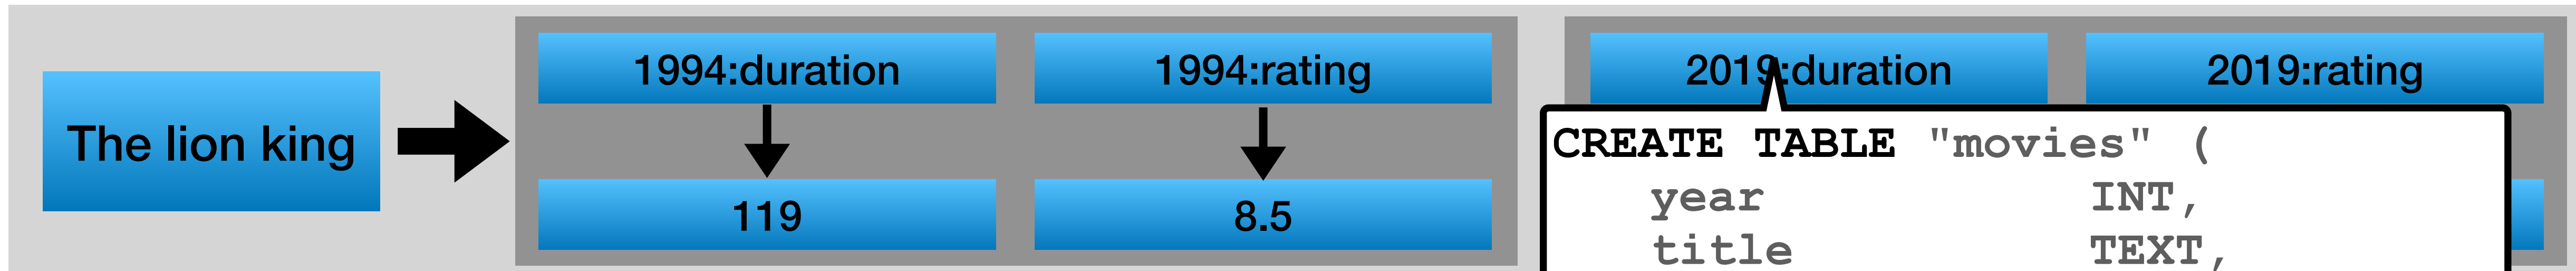| year	title	duration	rating
1994	The lion king	119	8.5
2019	The lion king	118	6.9
1999	American Pie	96	7
1999	The Matrix	136	8.7

77

Clustering columns - form


```
CREATE TABLE "movies" (
    year            INT,
    title           TEXT,
    duration        INT,
    rating          DOUBLE,
    PRIMARY KEY ((year), title)
);
```

1999 → American Pie:duration → 96 | American Pie:rating → 7 | 36 | 8.7

Which version will probably create bigger partitions?

Points to consider

1. The way we will query the data
2. The size of the partition

The lion king → 1994:duration → 119 | 1994:rating → 8.5 | 2019:duration | 2019:rating

```
CREATE TABLE "movies" (
    year            INT,
    title           TEXT,
    duration        INT,
    rating          DOUBLE,
    PRIMARY KEY ((title), year)
);
```

year	title	duration	rating
1994	The lion king	119	8.5
2019	The lion king	118	6.9
1999	American Pie	96	7
1999	The Matrix	136	8.7

78

Note on partition size

- The partition size is a crucial attribute for Cassandra performance and maintenance

- Ideally less than 100MB

  - Current versions supports much larger partitions (in GB)

- Can you think of a data model with a partition larger than 10GB?
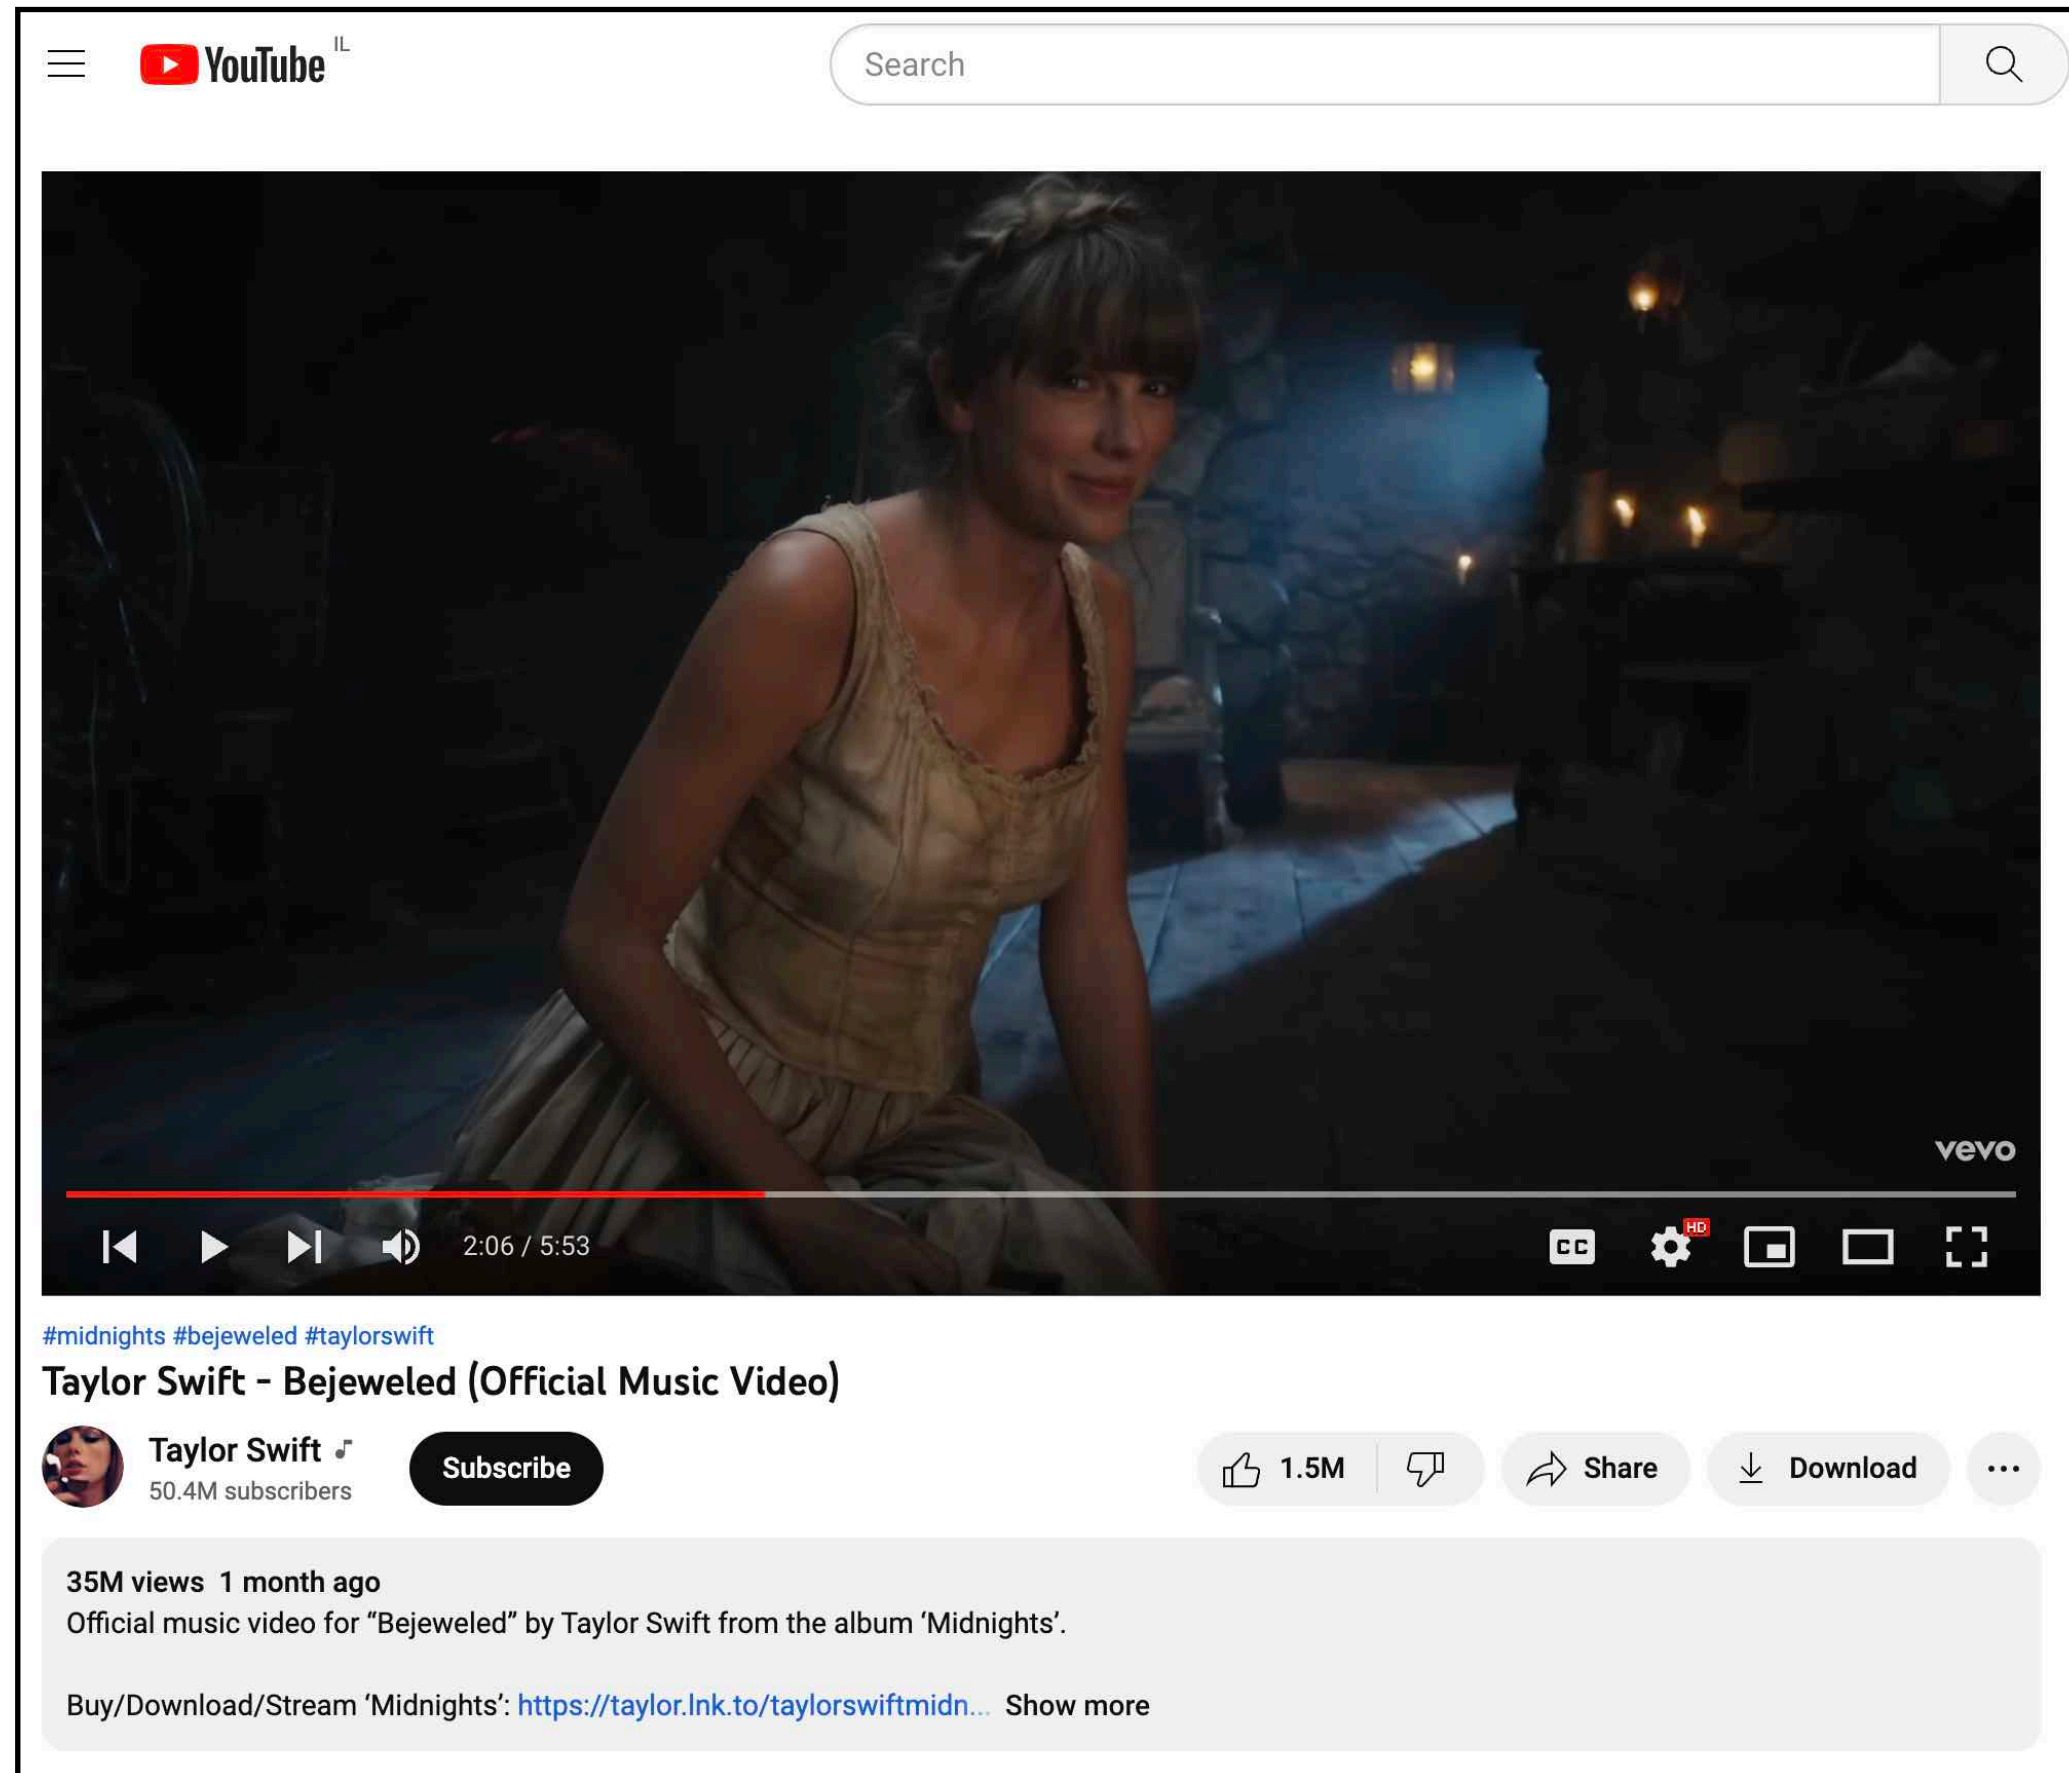
Note on partition size

- The partition size is a crucial attribute for Cassandra performance and maintenance

- Ideally less than 100MB

  - Current versions supports much larger partitions (in GB)

- Can you think of a data model with a partition larger than 10GB?
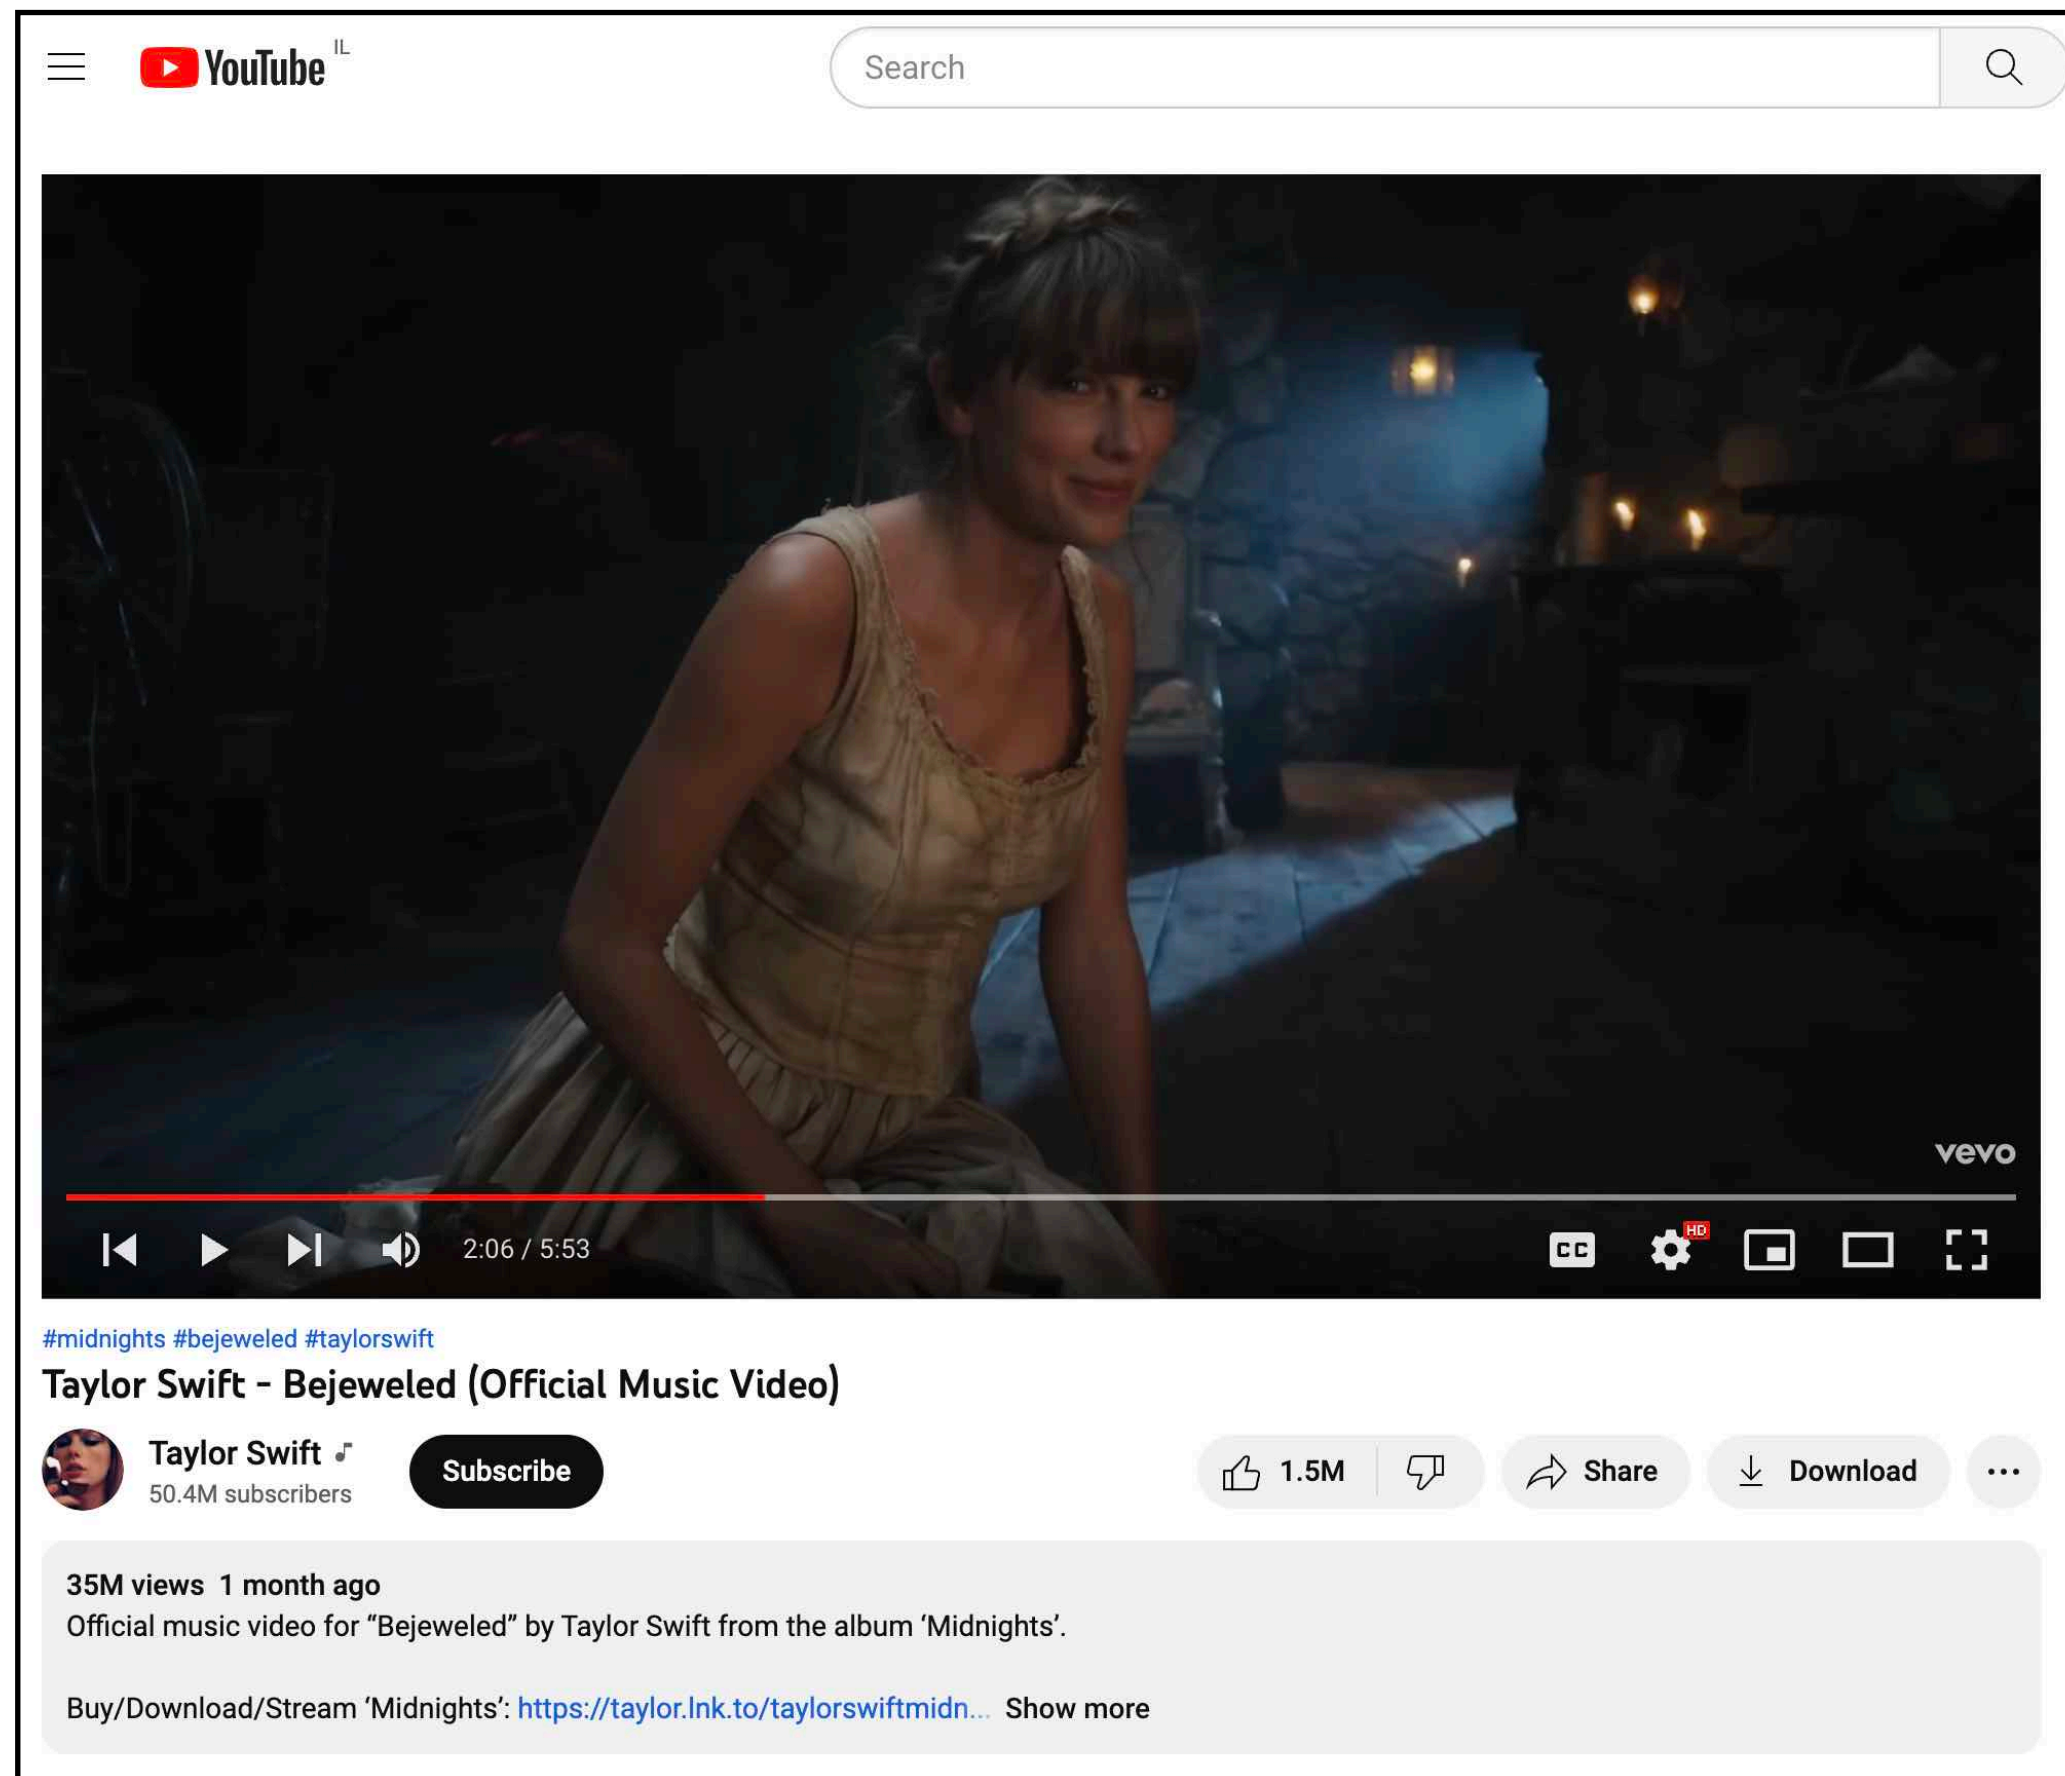
Note on partition size

Note on partition siz


```
CREATE TABLE "youtube_views" (
    video_id        BIGINT,
    view_id         TIMEUUID,
    user_id         BIGINT,
    percentage      INT,
    ...
    PRIMARY KEY ((video_id), view_id)
);
```

Assume YouTube saves for each video view some details (userID, device, time, percentage viewed…).

For simplicity, lets say 0.1k for each view

82

Note on partition siz

```
CREATE TABLE "youtube_views" (
    video_id        BIGINT,
    view_id         TIMEUUID,
    user_id         BIGINT,
    percentage      INT,
    ...
    PRIMARY KEY ((video_id), view_id)
);
```

Assume YouTube saves for each video view some details (userID, device, time, percentage viewed…).

For simplicity, lets say 0.1k for each view

#midnights #bejeweled #taylorswift

Taylor Swift – Bejeweled (Official Music Video)

Taylor Swift
50.4M subscribers

Subscribe

1.5M | Share | Download | ...

35M views 1 month ago
Official music video for "Bejeweled" by Taylor Swift from the album 'Midnights'.

Buy/Download/Stream 'Midnights': https://taylor.lnk.to/taylorswiftmidn... Show more

35m * 0.1k ~= 3.5GB

Note on partition size

35m * 0.1k ~= 3.5GB

Note on partition size

35m * 0.1k ~= 3.5GB

74m * 0.1k ~= 7.4GB

Note on

Taylor Swift – Bejeweled (Official Music Video)

#midnights #bejeweled #taylorswift

Taylor Swift – Shake It Off

Taylor Swift
50.4M subscribers

3.2B views 8 years ago
►Exclusive Merch: https://store.taylorswift.com

►Follow Taylor Swift Online Show more

Taylor Swift
50.4M subscribers

35M views 1 month ago
Official music video for "Bejeweled" by Taylor Swift from the album 'Midnights'.

Buy/Download/Stream 'Midnights': https://taylor.lnk.to/taylorswiftmidn... Show more
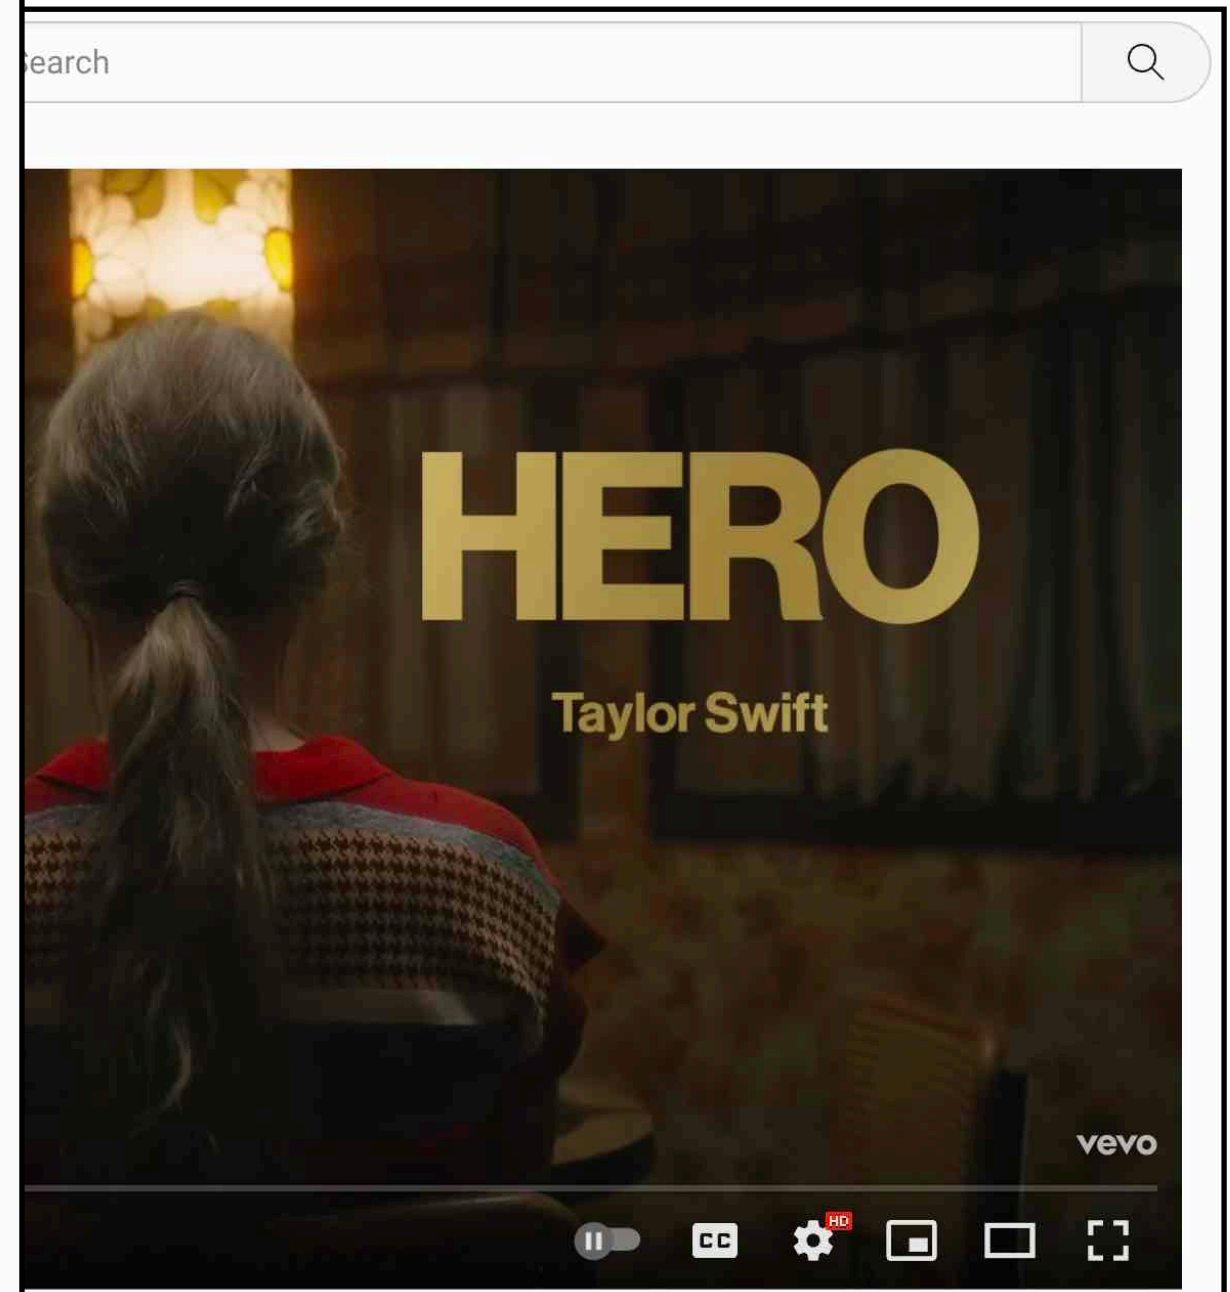
Taylor Swift
50.4M subscribers

74M views 1 month ago
Official music video for "Anti-Hero" by Taylor Swift from the album 'Midnights'.

Buy/Download/Stream 'Midnights': https://taylor.lnk.to/taylorswiftmidn... Show more

35m * 0.1k ~= 3.5GB

74m * 0.1k ~= 7.4GB

Note on

3.2b * 0.1k ~= 320GB (!)

Taylor Swift – Shake It Off

Taylor Swift ♪
50.4M subscribers

Subscribe

👍 12M 👎 ↱ Share ⬇ Download ⋯

3.2B views 8 years ago
▶Exclusive Merch: https://store.taylorswift.com

▶Follow Taylor Swift Online Show more

#midnights #bejeweled #taylorswift

Taylor Swift – Bejeweled (Official Music Video)

Taylor Swift ♪
50.4M subscribers

Subscribe

👍 1.5M 👎 ↱ Share ⬇ Download ⋯

35M views 1 month ago
Official music video for "Bejeweled" by Taylor Swift from the album 'Midnights'.

Buy/Download/Stream 'Midnights': https://taylor.lnk.to/taylorswiftmidn... Show more

35m * 0.1k ~= 3.5GB

Taylor Swift ♪
50.4M subscribers

Subscribe

👍 2.6M 👎 ↱ Share ⬇ Download ⋯

74M views 1 month ago
Official music video for "Anti-Hero" by Taylor Swift from the album 'Midnights'.

Buy/Download/Stream 'Midnights': https://taylor.lnk.to/taylorswiftmidn... Show more

74m * 0.1k ~= 7.4GB

Note on

3.2b * 0.1k ~= 320GB (!)

REMINDER
Each partition should ideally be less than 100MB
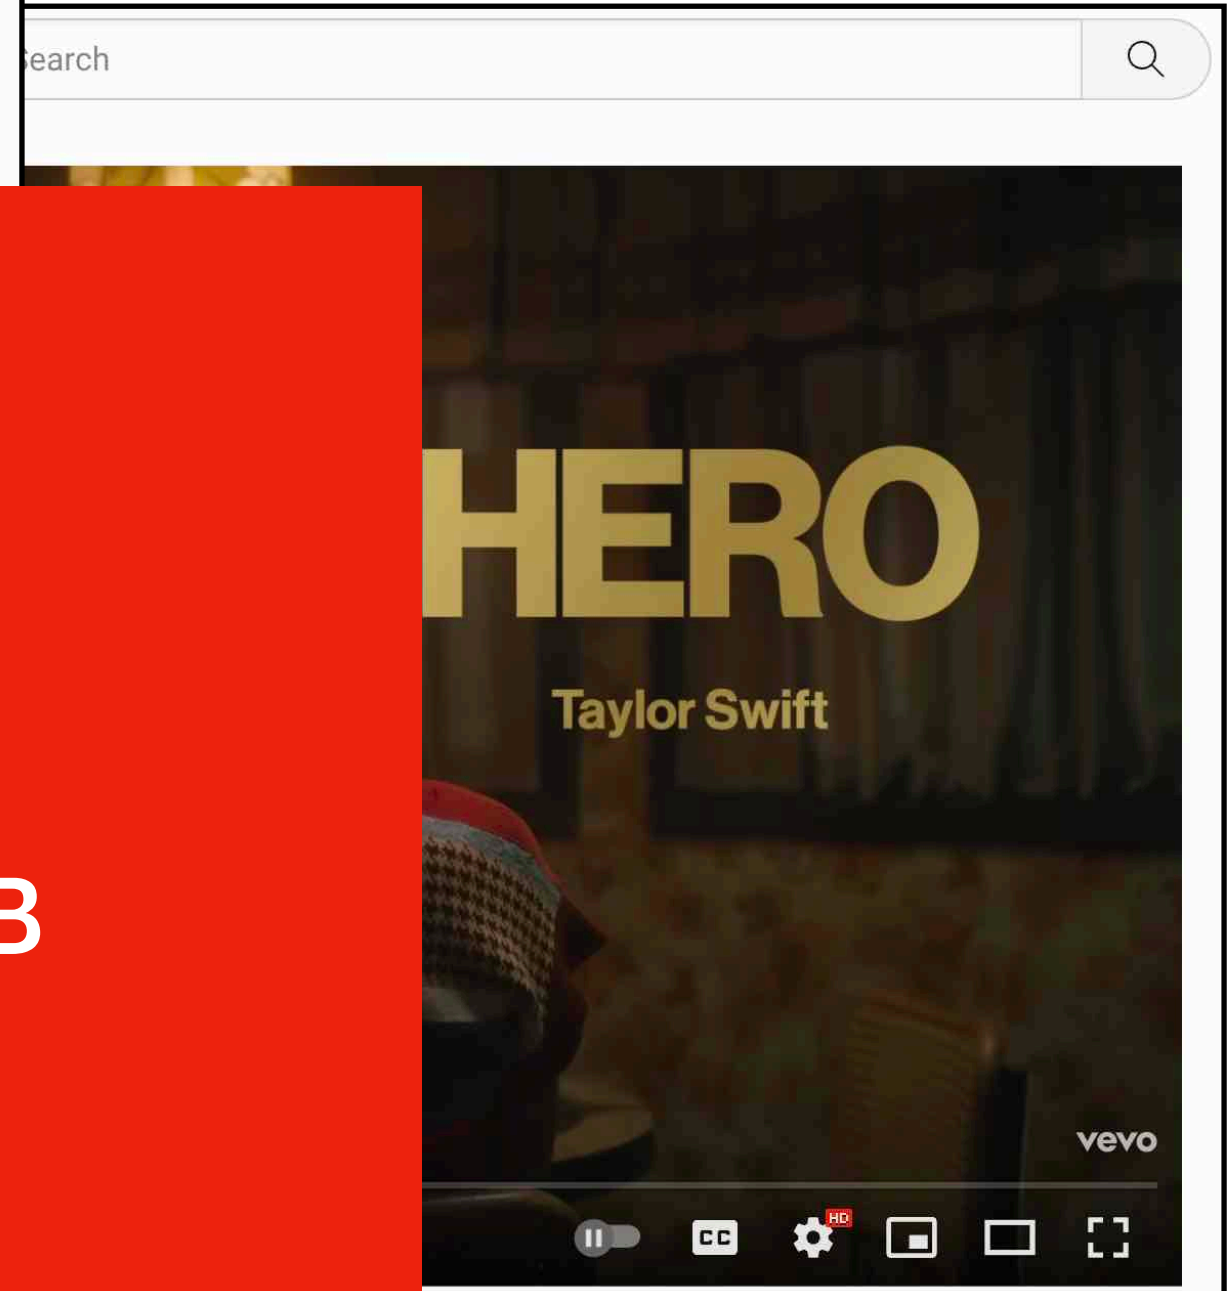
Taylor Swift - Bejeweled (Official M

Taylor Swift ♪ [Subscribe]
50.4M subscribers

35M views 1 month ago
Official music video for "Bejeweled" by Taylor Swift from the album 'Midnights'.

Buy/Download/Stream 'Midnights': https://taylor.lnk.to/taylorswiftmidn... Show more

35m * 0.1k ~= 3.5GB

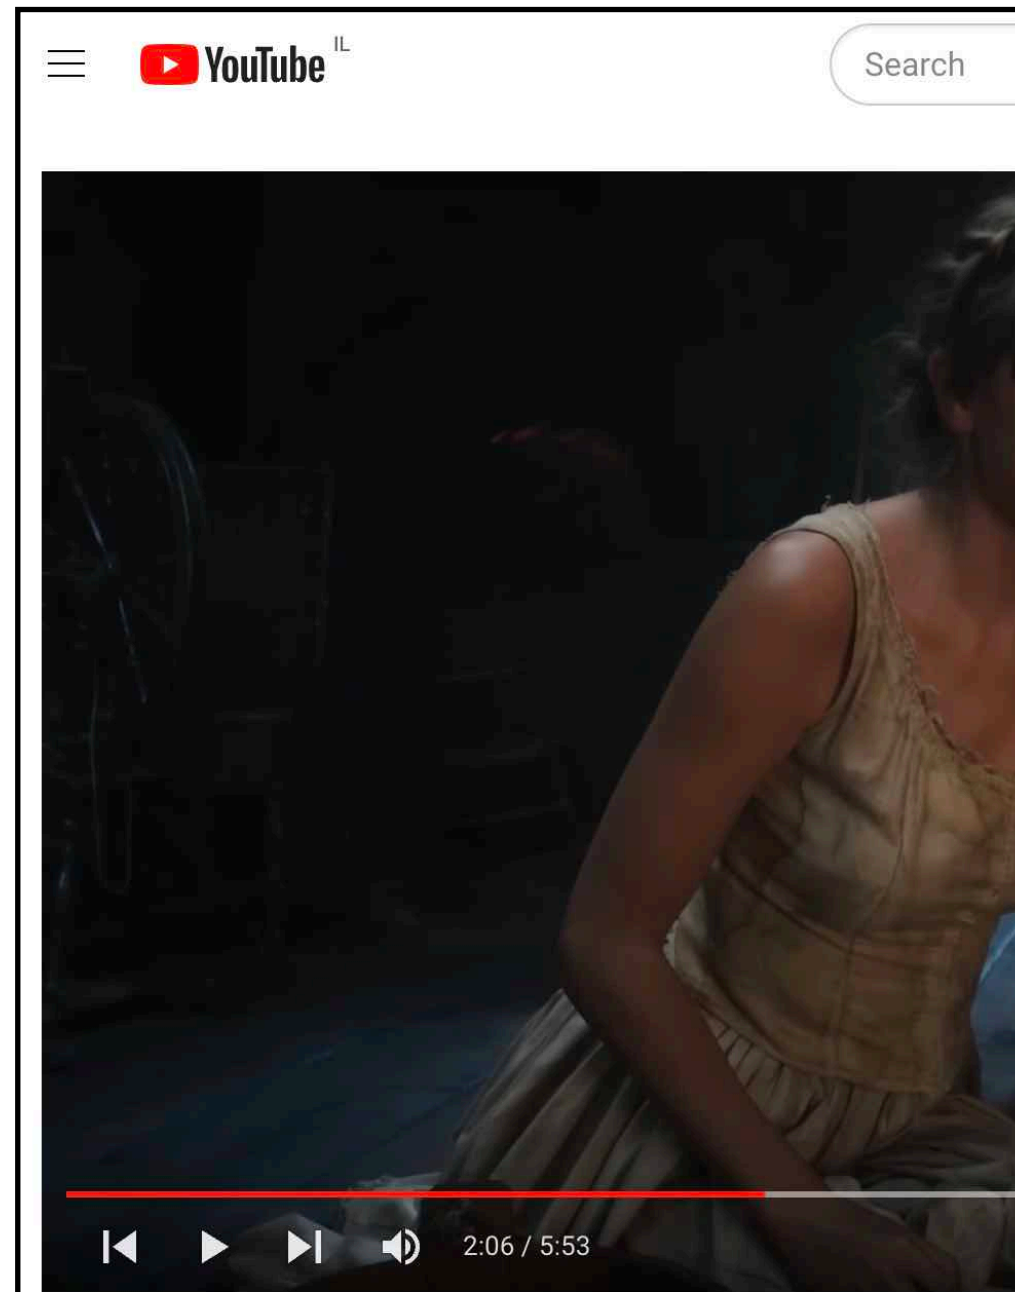**74M views 1 month ago**
Official music video for "Anti-Hero" by Taylor Swift from the album 'Midnights'.

Buy/Download/Stream 'Midnights': https://taylor.lnk.to/taylorswiftmidn... Show more

74m * 0.1k ~= 7.4GB

Top videos

The following table lists the top 30 most-viewed videos on YouTube, with each total rounded to the nearest 10 million views, uploader, and publication date. Note that some videos may not be available worldwide due to regional restrictions in certain countries.[6]

	Video name	Uploader	Views (billions)	Date	
1	Baby Shark Dance[7]	Pinkfong Baby Shark - Kids' Songs & Stories Youtube ↗	15.36	June 17, 2016	[A]
2	Despacito[10]	Luis Fonsi	8.60	January 12, 2017	[B]
3	Johny Johny Yes Papa[18]	LooLoo Kids - Nursery Rhymes and Children's Songs ↗	6.99	October 8, 2016	
4	Bath Song[19]	Cocomelon - Nursery Rhymes	6.93	May 2, 2018	
5	Wheels on the Bus[20]	Cocomelon - Nursery Rhymes	6.82	May 24, 2018	
6	See You Again[21]	Wiz Khalifa	6.49	April 6, 2015	[C]
7	Shape of You[26]	Ed Sheeran	6.38	January 30, 2017	[D]
8	Phonics Song with Two Words[29]	ChuChu TV Nursery Rhymes & Kids Songs	6.14	March 6, 2014	
9	Uptown Funk[30]	Mark Ronson	5.41	November 19, 2014	
10	Gangnam Style[31]	Psy	5.37	July 15, 2012	[E]
11	Learning Colors – Colorful Eggs on a Farm[36]	Miroshka TV	5.20	February 27, 2018	
12	Axel F[37]	Crazy Frog	4.90	June 16, 2009	
13	Dame Tu Cosita[38]	Ultra Records	4.88	April 5, 2018	
14	*Masha and the Bear* – Recipe for Disaster[39]	Get Movies	4.61	January 31, 2012	
15	Baa Baa Black Sheep[40]	Cocomelon - Nursery Rhymes	4.30	June 25, 2018	
16	Lakdi Ki Kathi[41]	Jingle Toons	4.27	June 14, 2018	
17	Shree Hanuman Chalisa[42]	T-Series Bhakti Sagar	4.19	May 10, 2011	
18	Sugar[43]	Maroon 5	4.15	January 14, 2015	
19	Waka Waka (This Time for Africa)[44]	Shakira	4.10	June 4, 2010	
20	Counting Stars[45]	OneRepublic	4.10	May 31, 2013	
21	Roar[46]	Katy Perry	4.08	September 5, 2013	
22	Humpty the train on a fruits ride[47]	Kiddiestv Hindi - Nursery Rhymes & Kids Songs	4.02	January 26, 2018	
23	Sorry[48]	Justin Bieber	3.89	October 22, 2015	
24	Dark Horse[49]	Katy Perry	3.86	February 20, 2014	
25	Perfect[50]	Ed Sheeran	3.85	November 9, 2017	
26	Thinking Out Loud[51]	Ed Sheeran	3.83	October 7, 2014	
27	Let Her Go[52]	Passenger	3.75	July 25, 2012	
28	Faded[53]	Alan Walker	3.73	December 3, 2015	
29	Girls Like You[54]	Maroon 5	3.72	May 31, 2018	
30	Lean On[55]	Major Lazer Official	3.70	March 22, 2015	

As of December 3, 2024

Top videos

The following table li... ...est 10 million views, uploader, and publication date ... ...in certain countries.[6]

YouTube IL

Search

Baby Shark Dance | #babyshark Most Viewed Video | Animal Songs | PINKFONG Songs for Children

Pinkfong Baby Shark - Kids' Songs & St... ♪
63.2M subscribers

Subscribe

38M

Share

Download

11B views 6 years ago
⭐ [NEW CHANNEL] Subscribe 🐣 Bebefinn - Nursery Rhymes & Kids Songs ⭐
https://www.youtube.com/channel/UC-cY...
Show more

YouTube Premium IL

Search

Baby Shark Dance | #babyshark Most Viewed Video | Animal Songs | PINKFONG Songs for Children

Pinkfong Baby Shark - Kids' Songs ... ♪
72.7M subscribers

Subscribe

42M

Share

Download

14,078,923,150 views 7 years ago
⭐ Pinkfong Sing-Along Movie 3: Catch the Gingerbread Man - In Cinemark September 9 (USA Only) ⭐

Get your ticket HERE: ...more

26	Thinking Out Loud[51]	Ed Sheeran	3.83	October 7, 2014
27	Let Her Go[52]	Passenger	3.75	July 25, 2012
28	Faded[53]	Alan Walker	3.73	December 3, 2015
29	Girls Like You[54]	Maroon 5	3.72	May 31, 2018
30	Lean On[55]	Major Lazer Official	3.70	March 22, 2015

As of December 3, 2024

Clustering columns - ordering

- Default order is ascending (descending optional)

```
CREATE TABLE "movies" (
    year                INT,
    title               TEXT,
    duration            INT,
    rating              DOUBLE,
    PRIMARY KEY ((year), title)
);
```
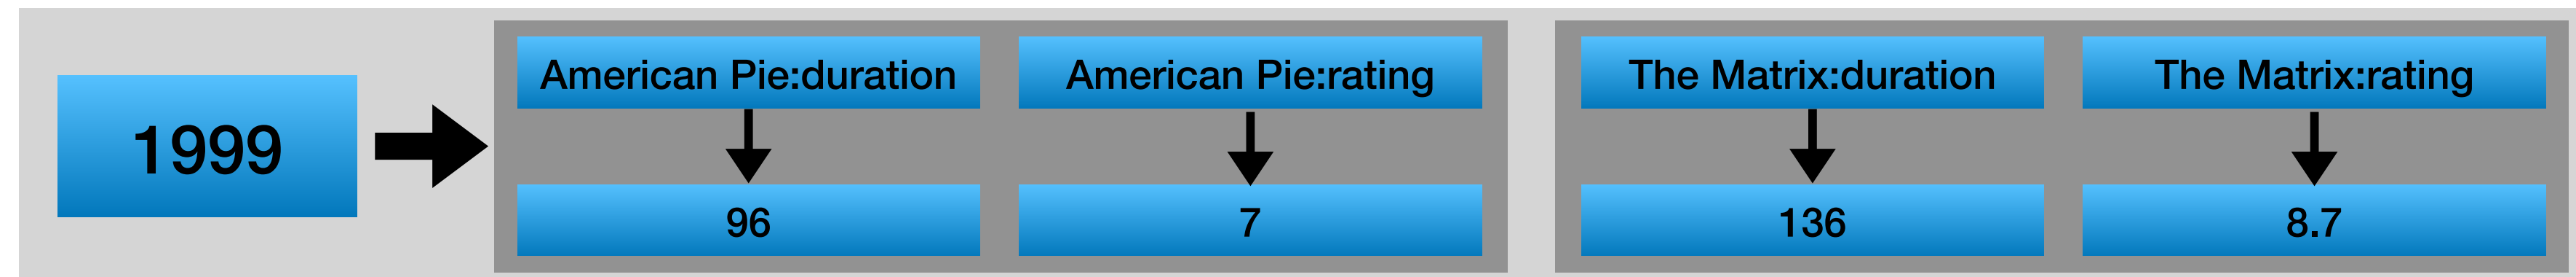


```
CREATE TABLE "movies" (
    year                INT,
    title               TEXT,
    duration            INT,
    rating              DOUBLE,
    PRIMARY KEY ((year), title)
) WITH CLUSTERING ORDER BY (title DESC);
```
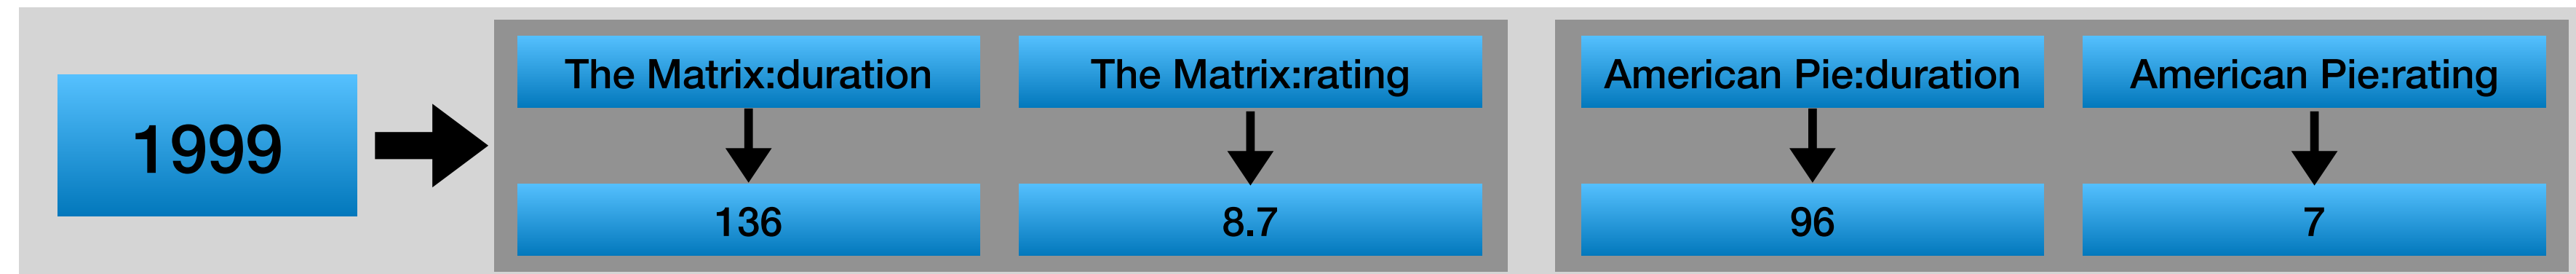
Clustering columns - queries

- Clustering columns are ordered —> **fast queries**

```
SELECT * FROM movies
WHERE    year  = 1999 AND
         title = "The Matrix"
```

```
CREATE TABLE "movies" (
    year             INT,
    title            TEXT,
    duration         INT,
    rating           DOUBLE,
    PRIMARY KEY ((year), title)
);
```


1999	American Pie:duration	American Pie:rating	The Matrix:duration	The Matrix:rating
	96	7	136	8.7

1000+ nodes
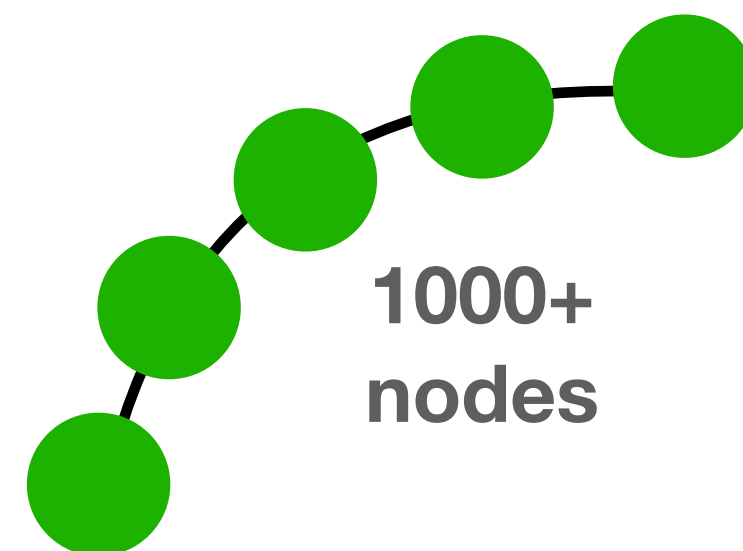
Clustering columns - range queries

- Lexicographic order for text

```
SELECT * FROM movies
WHERE    year  = 1999 AND
         title >= "The Matrix"
```

```
CREATE TABLE "movies" (
    year                INT,
    title               TEXT,
    duration            INT,
    rating              DOUBLE,
    PRIMARY KEY ((year), title)
);
```

1999	→	American Pie:duration	American Pie:rating		The Matrix:duration	The Matrix:rating
		96	7		136	8.7

1000+ nodes
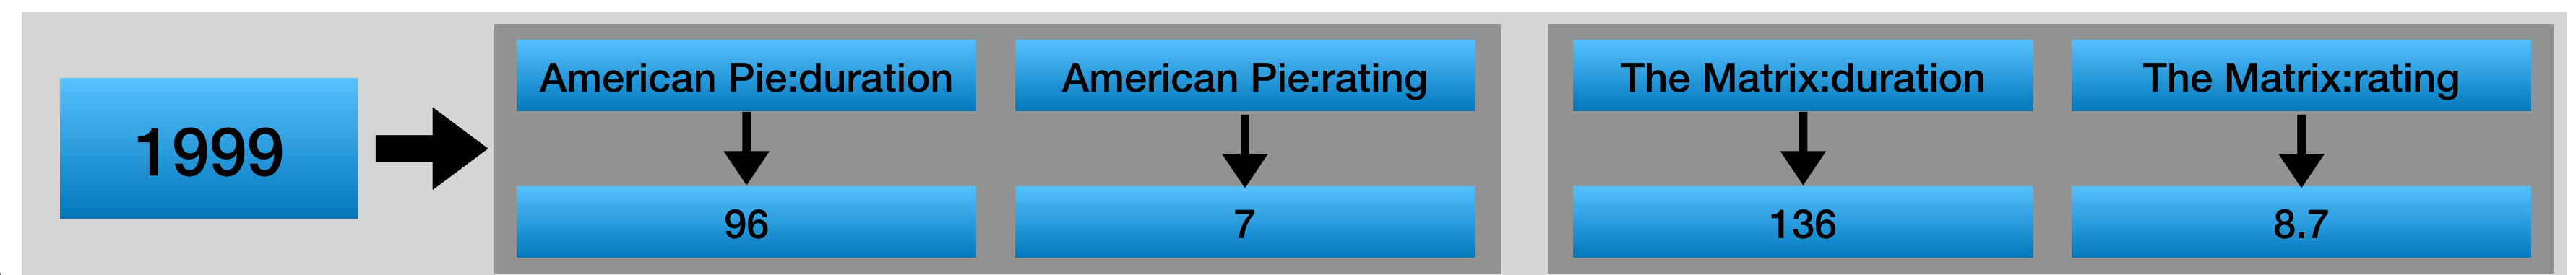
Clustering columns - range queries

- Lexicographic order for text

Can you think of an example for this query?

```
SELECT * FROM movies
WHERE    year  = 1999 AND
         title >= "The Matrix"
```

```
CREATE TABLE "movies" (
    year              INT,
    title             TEXT,
    duration          INT,
    rating            DOUBLE,
    PRIMARY KEY ((year), title)
);
```

| 1999 | → |

American Pie:duration	American Pie:rating		The Matrix:duration	The Matrix:rating
96	7		136	8.7

1000+ nodes
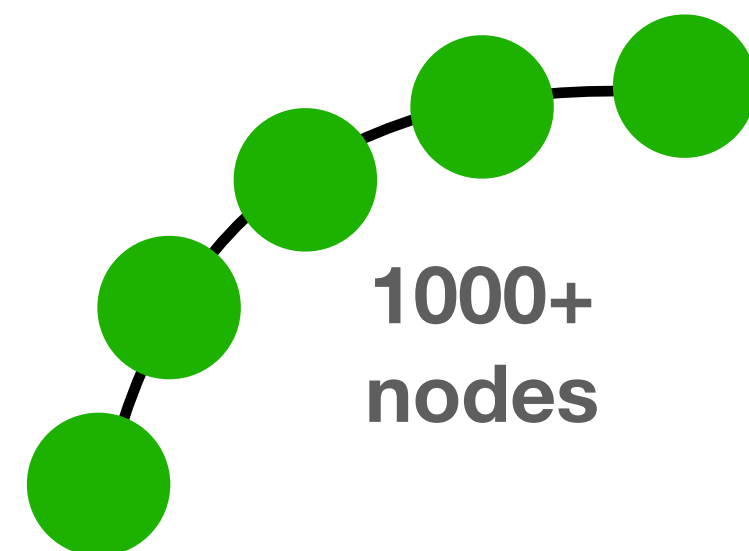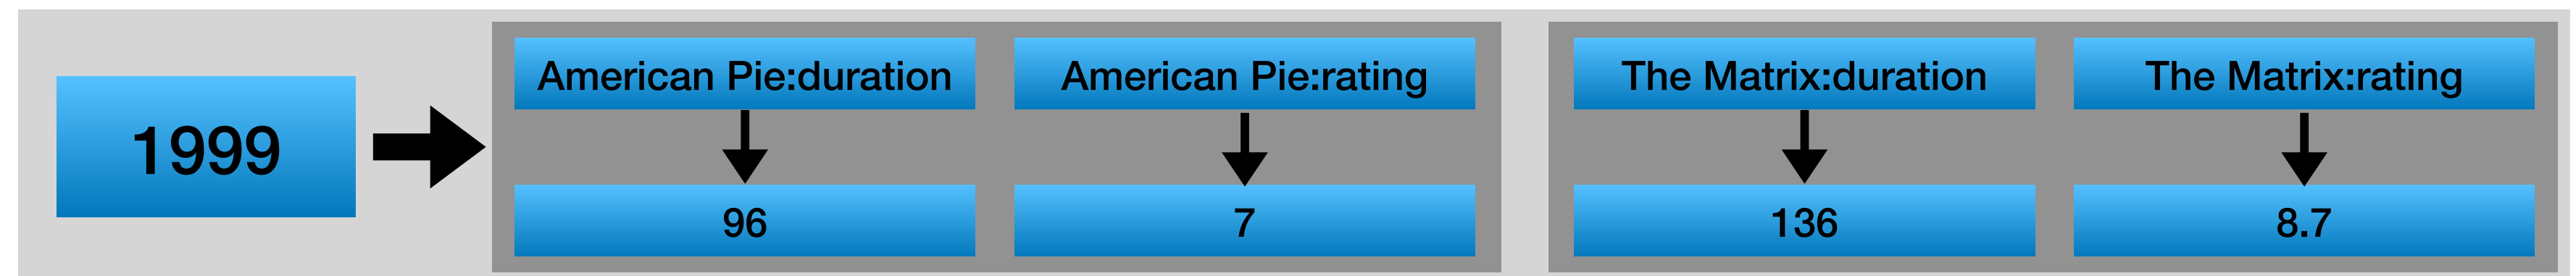
Clustering columns - range queries

- Lexicographic order for text

Can you think of an example for this query?
Showing movies by year **with paging**

```
SELECT * FROM movies
WHERE    year  = 1999 AND
         title >= "The Matrix"
```

```
CREATE TABLE "movies" (
    year              INT,
    title             TEXT,
    duration          INT,
    rating            DOUBLE,
    PRIMARY KEY ((year), title)
);
```

1999	→	American Pie:duration	American Pie:rating		The Matrix:duration	The Matrix:rating
		96	7		136	8.7

1000+
nodes

Clustering columns - more than one

```
CREATE TABLE "logs" (
    year            INT,
    month           INT,
    day             INT,
    hour            INT,
    log             TEXT,
    PRIMARY KEY ((year, month), day, hour)
) WITH CLUSTERING ORDER BY (day DESC, hour DESC);
```

Clustering columns - more than one

```
CREATE TABLE "logs" (
    year              INT,
    month             INT,
    day               INT,
    hour              INT,
    log               TEXT,
    PRIMARY KEY ((year, month), day, hour)
) WITH CLUSTERING ORDER BY (day DESC, hour DESC);
```

Clustering columns - more than one

```
CREATE TABLE "logs" (
    year               INT,
    month              INT,
    day                INT,
    hour               INT,
    log                TEXT,
    PRIMARY KEY ((year, month), day, hour)
) WITH CLUSTERING ORDER BY (day DESC, hour DESC);
```

When querying - WHERE needs to follow the order of the clustering columns

30:11:log	30:07:log	26:18:log	24:22:log
"log…"	"log…"	"log…"	"log…"

2021:11

2021/11/30, 11am   2021/11/30, 7am   2021/11/26, 6pm   2021/11/24, 10pm

Clustering columns - more than one

```
CREATE TABLE "logs" (
    year              INT,
    month             INT,
    day               INT,
    hour              INT,
    log               TEXT,
    PRIMARY KEY ((year, month), day, hour)
) WITH CLUSTERING ORDER BY (day DESC, hour DESC)
```

When querying - WHERE needs to follow the order of the clustering columns

Is this valid?

```
SELECT * FROM logs
WHERE    year  = 2021 AND
         month = 11 AND
         hour  = 18
```


2021:11

| 30:11:log | 30:07:log | 26:18:log | 24:22:log |
| "log…" | "log…" | "log…" | "log…" |

2021/11/30, 11am   2021/11/30, 7am   2021/11/26, 6pm   2021/11/24, 10pm

Clustering columns - more than one

```
CREATE TABLE "logs" (
    year             INT,
    month            INT,
    day              INT,
    hour             INT,
    log              TEXT,
    PRIMARY KEY ((year, month), day, hour)
) WITH CLUSTERING ORDER BY (day DESC, hour DESC)
```

When querying - WHERE needs to follow the order of the clustering columns

Is this valid?
NO - day is required

```
SELECT * FROM logs
WHERE    year  = 2021 AND
         month = 11 AND
         hour  = 18
```

2021:11	→	30:11:log	30:07:log	26:18:log	24:22:log
		↓	↓	↓	↓
		"log…"	"log…"	"log…"	"log…"

2021/11/30, 11am   2021/11/30, 7am   2021/11/26, 6pm   2021/11/24, 10pm

100

Clustering columns - more than one

```
CREATE TABLE "logs" (
    year              INT,
    month             INT,
    day               INT,
    hour              INT,
    log               TEXT,
    PRIMARY KEY ((year, month), day, hour)
) WITH CLUSTERING ORDER BY (day DESC, hour DESC)
```

When querying - WHERE needs to follow the order of the clustering columns

Is this valid?

```
SELECT * FROM logs
WHERE    year  = 2021 AND
         month = 11 AND
         day   = 26
```

2021:11	30:11:log	30:07:log	26:18:log	24:22:log
	"log..."	"log..."	"log..."	"log..."

2021/11/30, 11am   2021/11/30, 7am   2021/11/26, 6pm   2021/11/24, 10pm

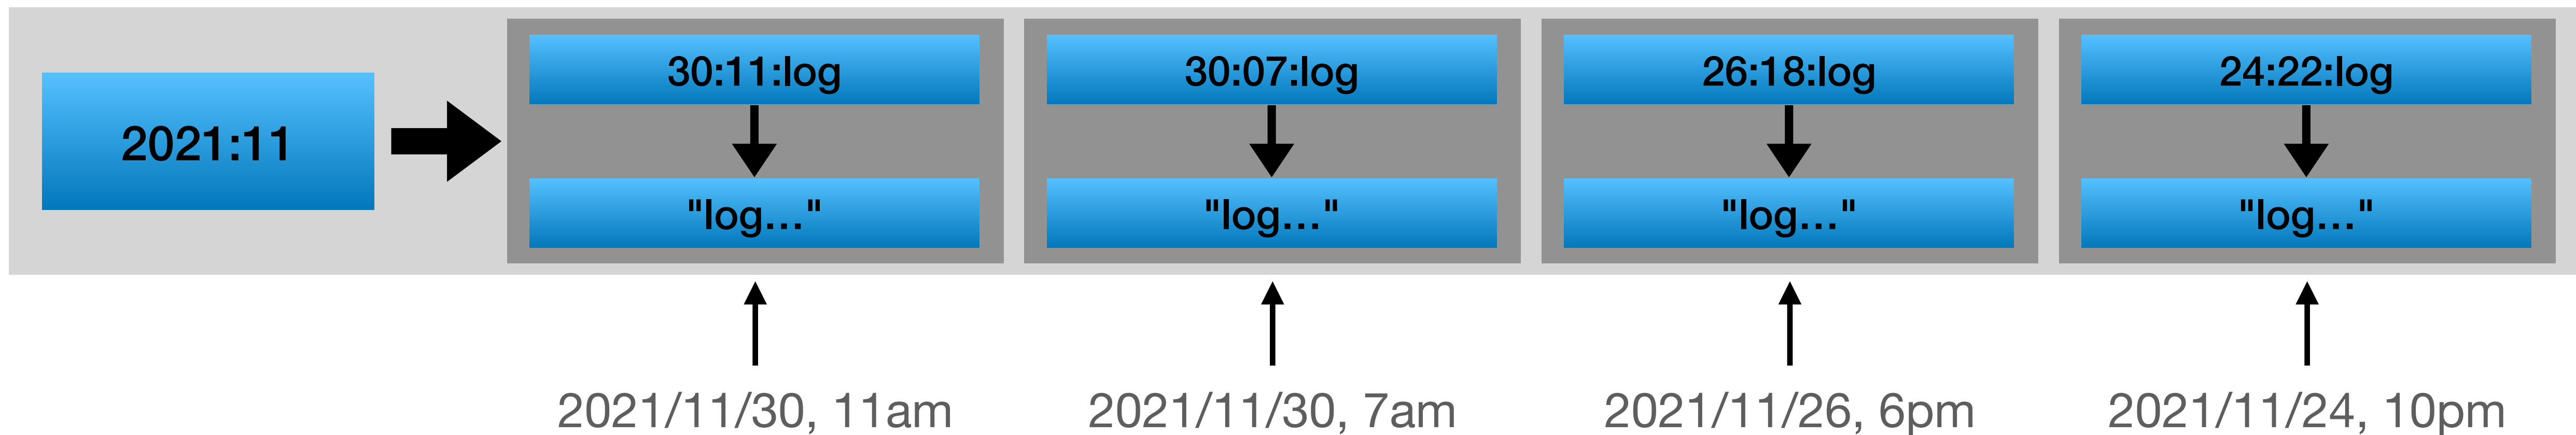# Clustering columns - more than one

```
CREATE TABLE "logs" (
    year              INT,
    month             INT,
    day               INT,
    hour              INT,
    log               TEXT,
    PRIMARY KEY ((year, month), day, hour)
) WITH CLUSTERING ORDER BY (day DESC, hour DESC)
```
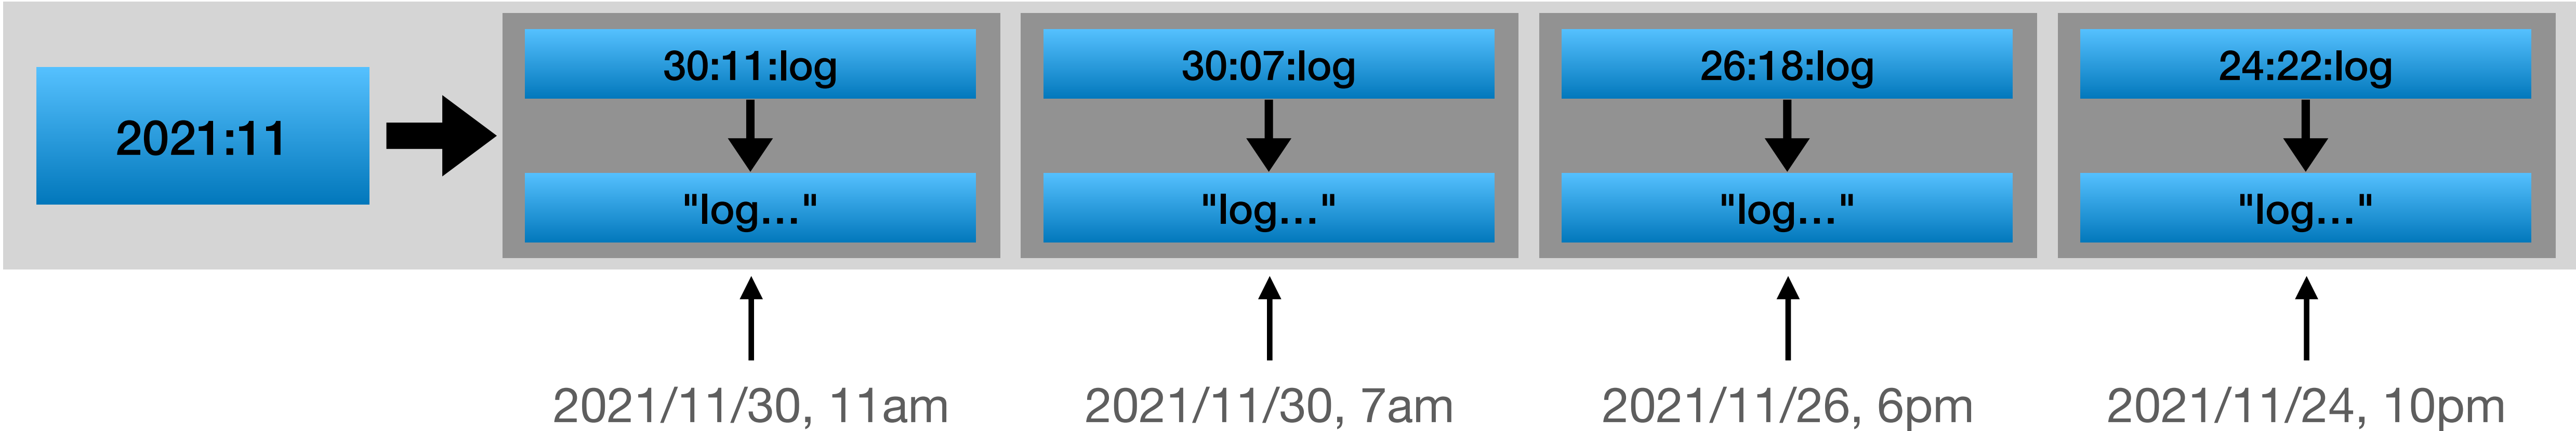
When querying - WHERE needs to follow the order of the clustering columns

Is this valid?
YES

```
SELECT * FROM logs
WHERE   year  = 2021 AND
        month = 11 AND
        day   = 26
```

2021:11	→	30:11:log	30:07:log	26:18:log	24:22:log
		↓	↓	↓	↓
		"log…"	"log…"	"log…"	"log…"

2021/11/30, 11am   2021/11/30, 7am   2021/11/26, 6pm   2021/11/24, 10pm
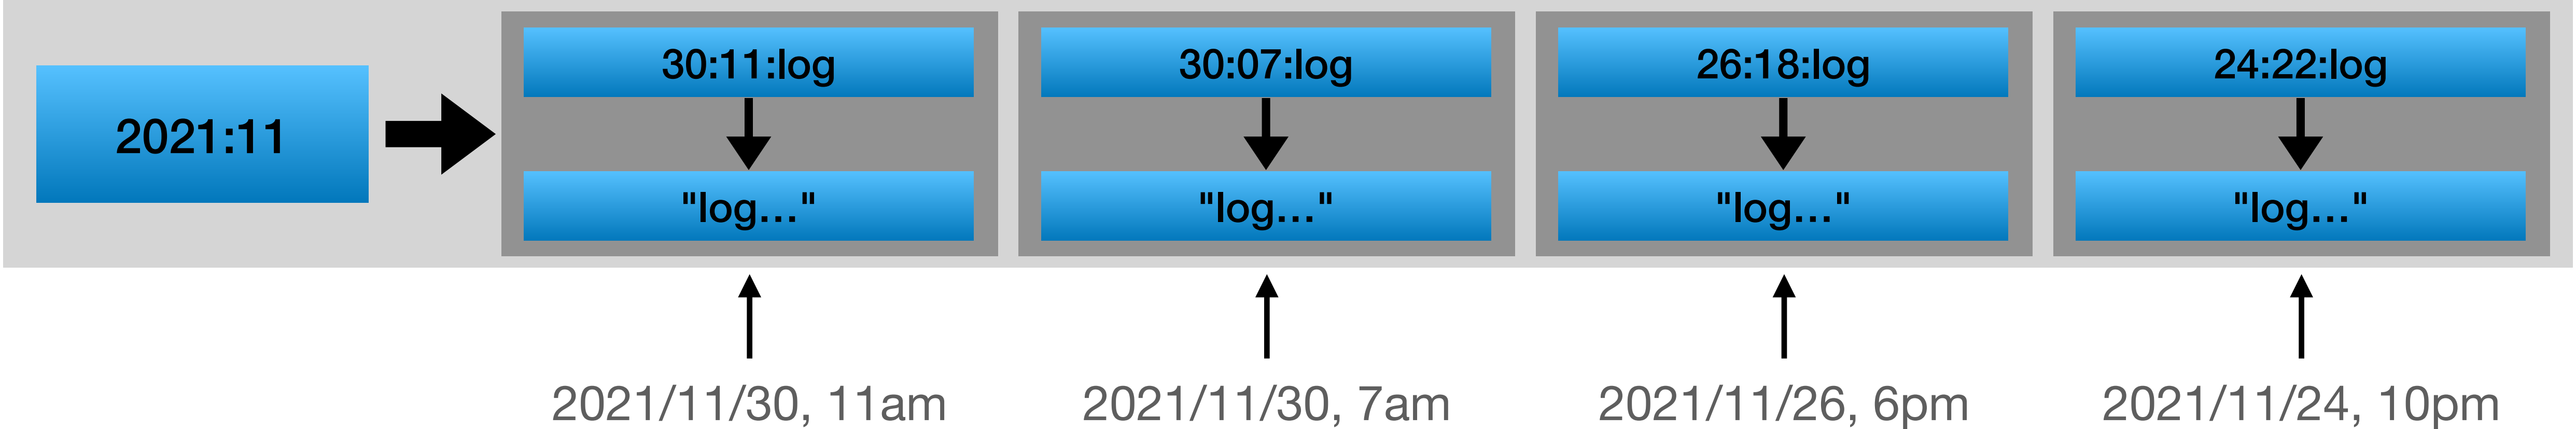
Agenda

- History

- Architecture

- Data model (Original)

- Data model (CQL)

- **Examples**

Example (1)

user_id	video_id	video_ts	duration	device
user_12	video_55	1638198718	344	iPhone
user_12	video_84	1638258828	877	Apple TV

- Saving user viewing history

```
CREATE TABLE "user_viewing_history" (
    user_id      TEXT,
    video_id     BIGINT,
    view_ts      TIMESTAMP,
    duration     INT,
    device       TEXT,
    PRIMARY KEY ((user_id), video_id)
);
```

user_12	→	video_55:view_ts	video_55:duration	video_55:device		video_84:view_ts	video_84:duration	video_84:device
		↓	↓	↓		↓	↓	↓
		1638198718	344	iPhone		1638258828	877	Apple TV
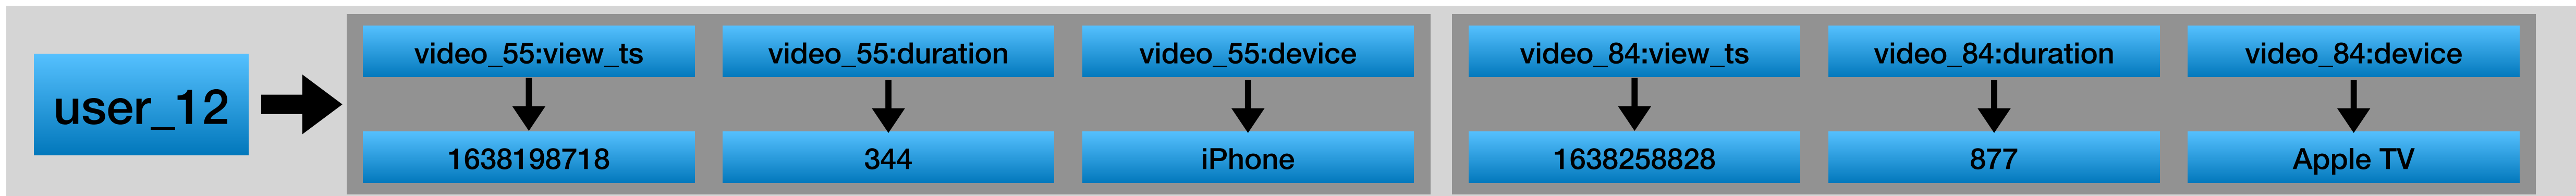
Example (1)

user_id	video_id	video_ts	duration	device
user_12	video_55	1638198718	344	iPhone
user_12	video_84	1638258828	877	Apple TV

- Saving user viewing history

> How would you support saving more than one view per video?

```
CREATE TABLE "user_viewing_history" (
    user_id      TEXT,
    video_id     BIGINT,
    view_ts      TIMESTAMP,
    duration     INT,
    device       TEXT,
    PRIMARY KEY ((user_id), video_id)
);
```

user_12	video_55:view_ts	video_55:duration	video_55:device	video_84:view_ts	video_84:duration	video_84:device
	1638198718	344	iPhone	1638258828	877	Apple TV

Example (1)

user_id	video_id	video_ts	duration	device
user_12	video_55	1638198718	344	iPhone
user_12	video_84	1638258828	877	Apple TV

- Saving user viewing history

```
CREATE TABLE "user_viewing_history" (
    user_id      TEXT,
    video_id     BIGINT,
    view_ts      TIMESTAMP,
    duration     INT,
    device       TEXT,
    PRIMARY KEY ((user_id), video_id, view_ts)
);
```

user_12	→	video_55:1638198718:duration	video_55:1638198718:device		video_84:1638258828:duration	video_84:1638258828:device
		↓	↓		↓	↓
		334	iPhone		887	Apple TV

Example (1)

user_id	video_id	video_ts	duration	device
user_12	video_55	1638198718	344	iPhone
user_12	video_84	1638258828	877	Apple TV

- Saving user viewing history

What is the order the results are returned?

```
CREATE TABLE "user_viewing_history" (
    user_id      TEXT,
    video_id     BIGINT,
    view_ts      TIMESTAMP,
    duration     INT,
    device       TEXT,
    PRIMARY KEY ((user_id), video_id, view_ts)
);
```
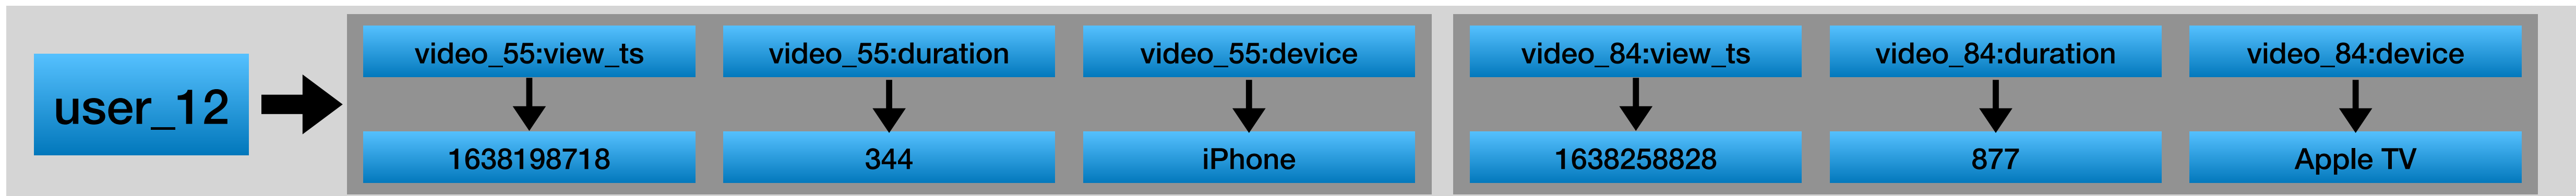
user_12	video_55:1638198718:duration	video_55:1638198718:device	video_84:1638258828:duration	video_84:1638258828:device
	334	iPhone	887	Apple TV

Example (1)

user_id	video_id	video_ts	duration	device
user_12	video_55	1638198718	344	iPhone
user_12	video_84	1638258828	877	Apple TV

• Saving user viewing history

How would you support getting the results by the view_ts and not by video_id?

```
CREATE TABLE "user_viewing_history" (
    user_id      TEXT,
    video_id     BIGINT,
    view_ts      TIMESTAMP,
    duration     INT,
    device       TEXT,
    PRIMARY KEY ((user_id), video_id, view_ts)
);
```
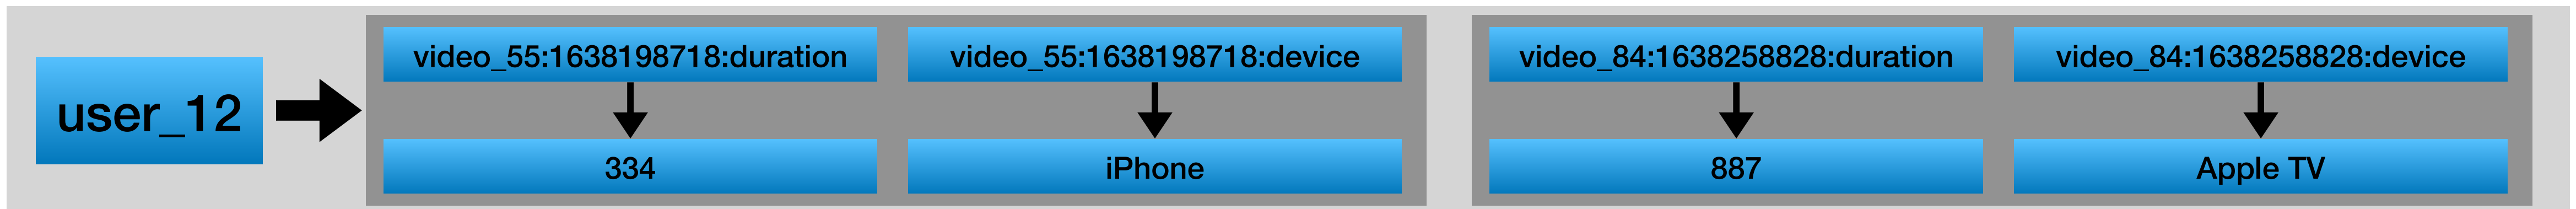
user_12	video_55:1638198718:duration	video_55:1638198718:device	video_84:1638258828:duration	video_84:1638258828:device
	334	iPhone	887	Apple TV

Example (1)

user_id	video_ts	video_id	duration	device
user_12	1638258828	video_84	877	Apple TV
user_12	1638198718	video_55	344	iPhone

- Saving user viewing history

```
CREATE TABLE "user_viewing_history" (
    user_id     TEXT,
    video_id    BIGINT,
    view_ts     TIMESTAMP,
    duration    INT,
    device      TEXT,
    PRIMARY KEY ((user_id), view_ts, video_id)
) WITH CLUSTERING ORDER BY (view_ts DESC);
```
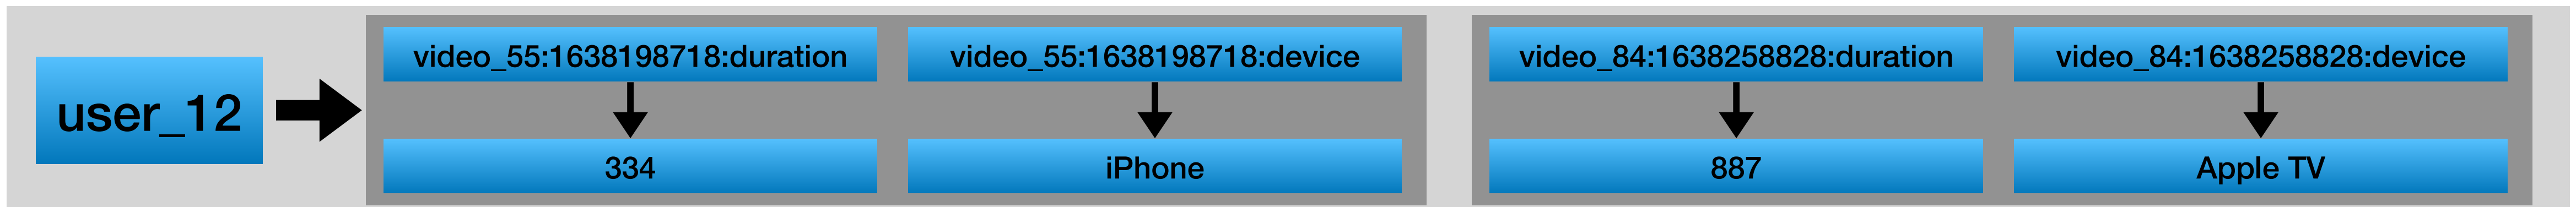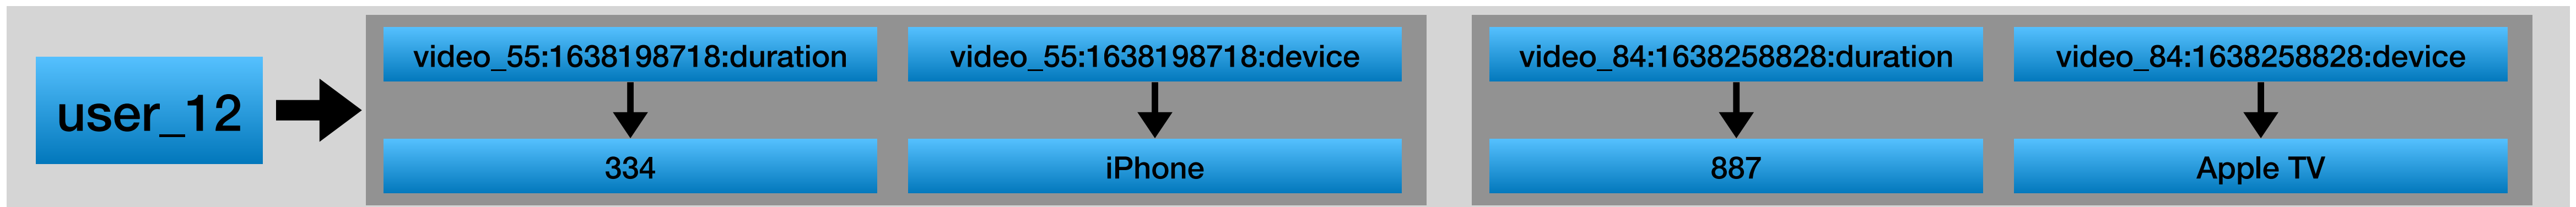
user_12	→	1638258828:video_84:duration	1638258828:video_84:device	1638198718:video_55:duration	1638198718:video_55:device
		↓	↓	↓	↓
		887	Apple TV	334	iPhone

Example (2)

video_id	video_ts	user_id	duration	device
video_84	1638258828	user_12	877	Apple TV
video_84	1638245477	user_44	644	Apple TV

- Saving user viewing history for a video

```
CREATE TABLE "video_viewing_history" (
    user_id      TEXT,
    video_id     BIGINT,
    view_ts      TIMESTAMP,
    duration     INT,
    device       TEXT,
    PRIMARY KEY ((video_id), view_ts, user_id)
) WITH CLUSTERING ORDER BY (view_ts DESC);
```
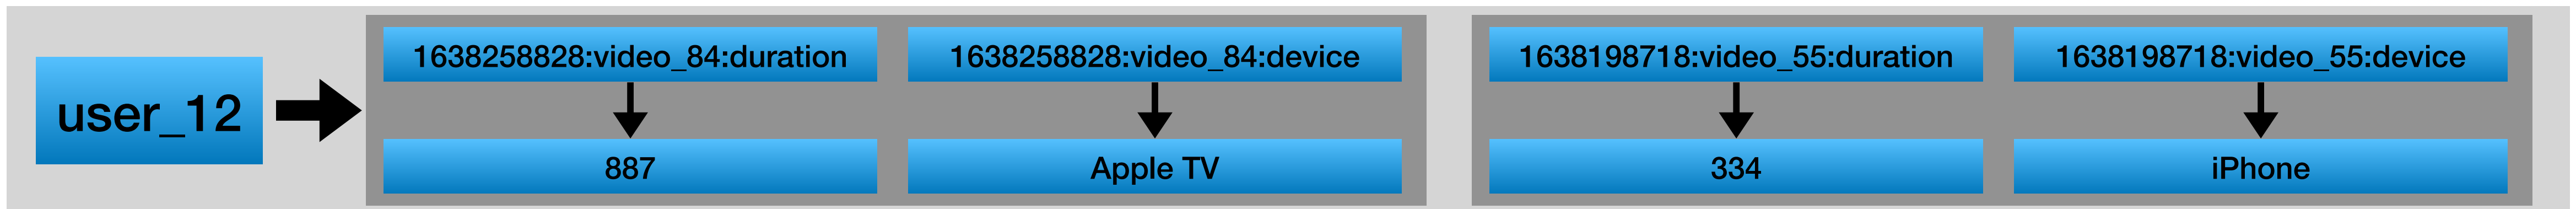
video_84 →	1638258828:user_12:duration	1638258828:user_12:device		1638245477:user_44:duration	1638245477:user_44:device
	↓	↓		↓	↓
	887	Apple TV		664	Apple TV

Example (2)

video_id	video_ts	user_id	duration	device
video_84	1638258828	user_12	877	Apple TV
video_84	1638245477	user_44	644	Apple TV

- Saving user viewing history for a video

For games of thrones this partition could be too big. What would you do?

```
CREATE TABLE "video_viewing_history" (
    user_id      TEXT,
    video_id     BIGINT,
    view_ts      TIMESTAMP,
    duration     INT,
    device       TEXT,
    PRIMARY KEY ((video_id), view_ts, user_id)
) WITH CLUSTERING ORDER BY (view_ts DESC);
```
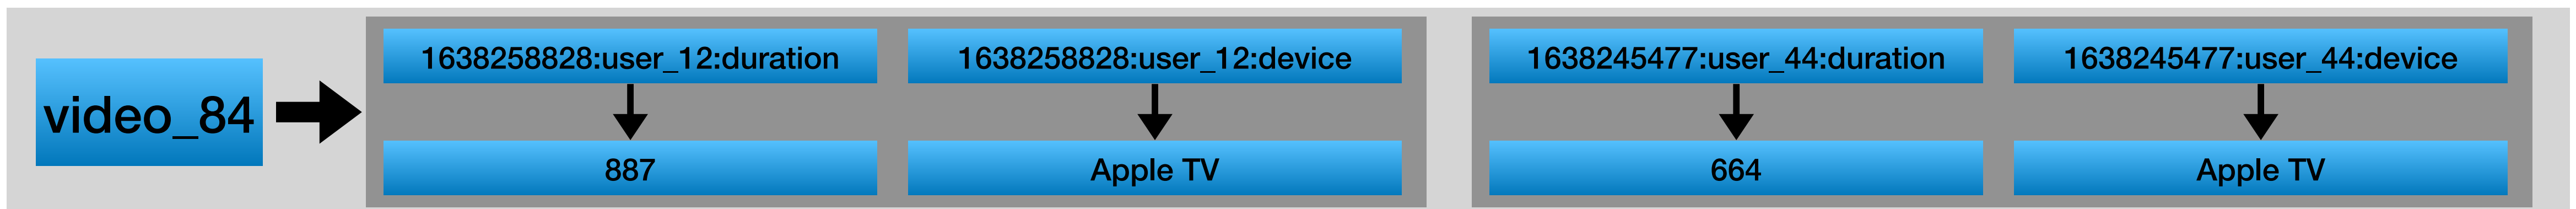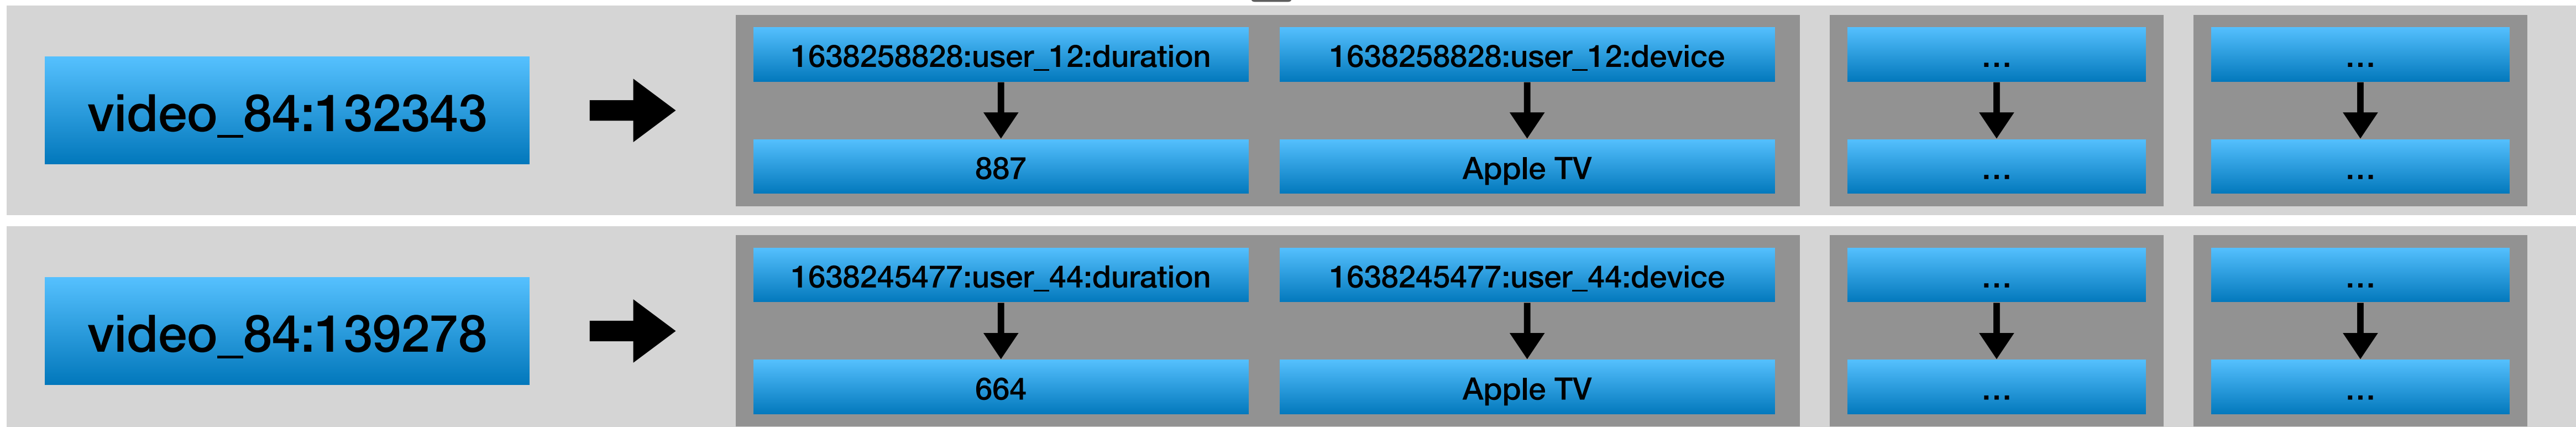
video_84 →

1638258828:user_12:duration	1638258828:user_12:device	1638245477:user_44:duration	1638245477:user_44:device
↓	↓	↓	↓
887	Apple TV	664	Apple TV

Example (2)

video_id	view_day	video_ts	user_id	duration	device
video_84	132343	1638258828	user_12	877	Apple TV
video_84	139278	1638245477	user_44	644	Apple TV

- Saving user viewing history for a video

```
CREATE TABLE "video_viewing_history" (
    user_id      TEXT,
    video_id     BIGINT,
    view_day     INT,
    view_ts      TIMESTAMP,
    duration     INT,
    device       TEXT,
    PRIMARY KEY ((video_id, view_day), view_ts, user_id)
) WITH CLUSTERING ORDER BY (view_ts DESC);
```

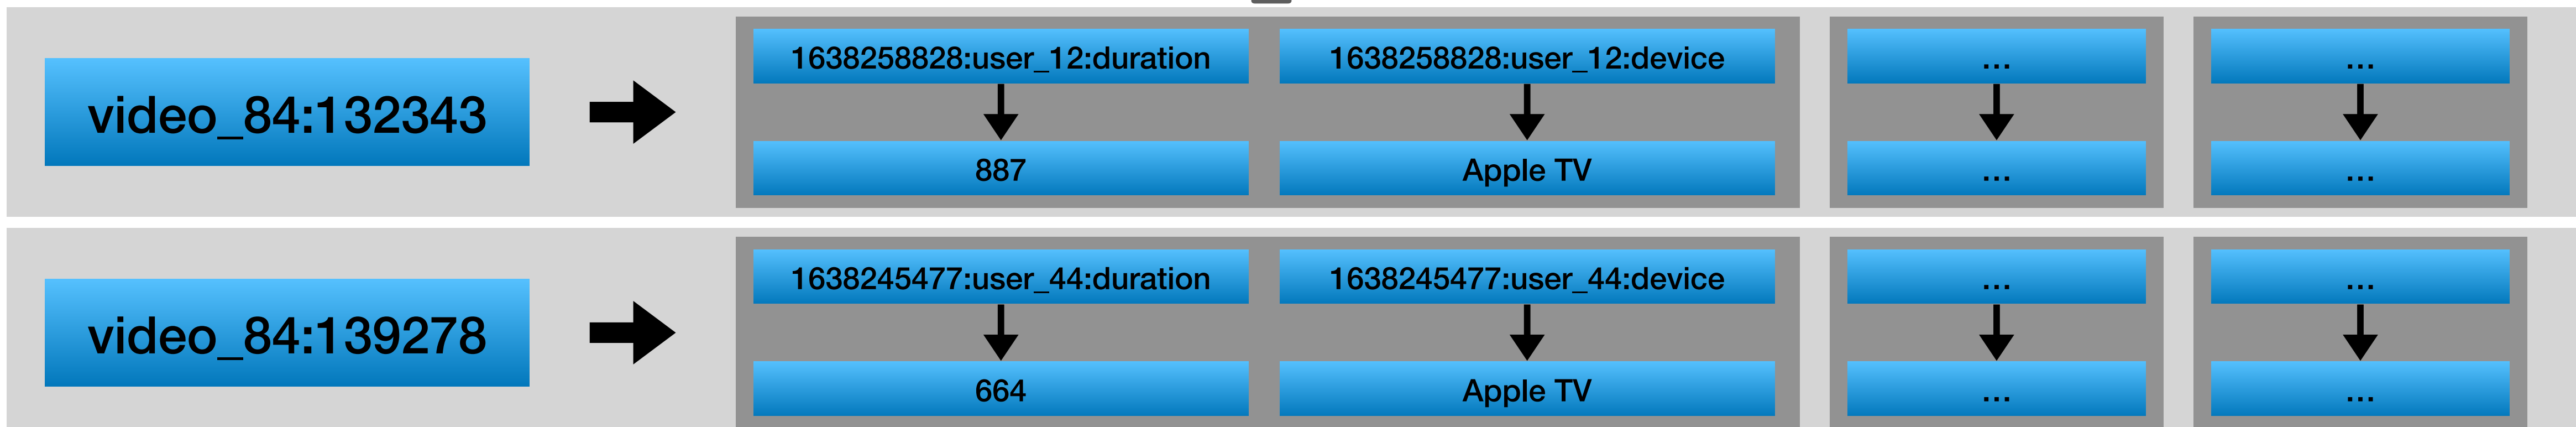view_day = same as timestamp, just by days and not by milliseconds

Example (2)

video_id	view_day	video_ts	user_id	duration	device
video_84	132343	1638258828	user_12	877	Apple TV
video_84	139278	1638245477	user_44	644	Apple TV

- Saving user viewing history for a video

```
CREATE TABLE "video_viewing_history" (
    user_id      TEXT,
    video_id     BIGINT,
    view_day     INT,
    view_ts      TIMESTAMP,
    duration     INT,
    device       TEXT,
    PRIMARY KEY ((video_id, view_day), view_ts, user_id)
) WITH CLUSTERING ORDER BY (view_ts DESC);
```

But what if 10m users watch a new episode on the day it is first released?

A lot more on "data modeling" later...

video_84:132343	➡	1638258828:user_12:duration	1638258828:user_12:device	...	...
		↓	↓	↓	↓
		887	Apple TV	...	...

video_84:139278	➡	1638245477:user_44:duration	1638245477:user_44:device	...	...
		↓	↓	↓	↓
		664	Apple TV	...	...

Example (3)

- Raw map of maps (old data model)

row_key	k	v
123	title	Bad Boys
123	year	1995
123	duration	119
123	rating	6.9
125	title	American Pie
125	year	1999
125	duration	96
125	rating	7

```
CREATE TABLE "test_map_of_maps" (
  row_key       TEXT,
  k             TEXT,
  v             BLOB,
  PRIMARY KEY (row_key, k)
);
```
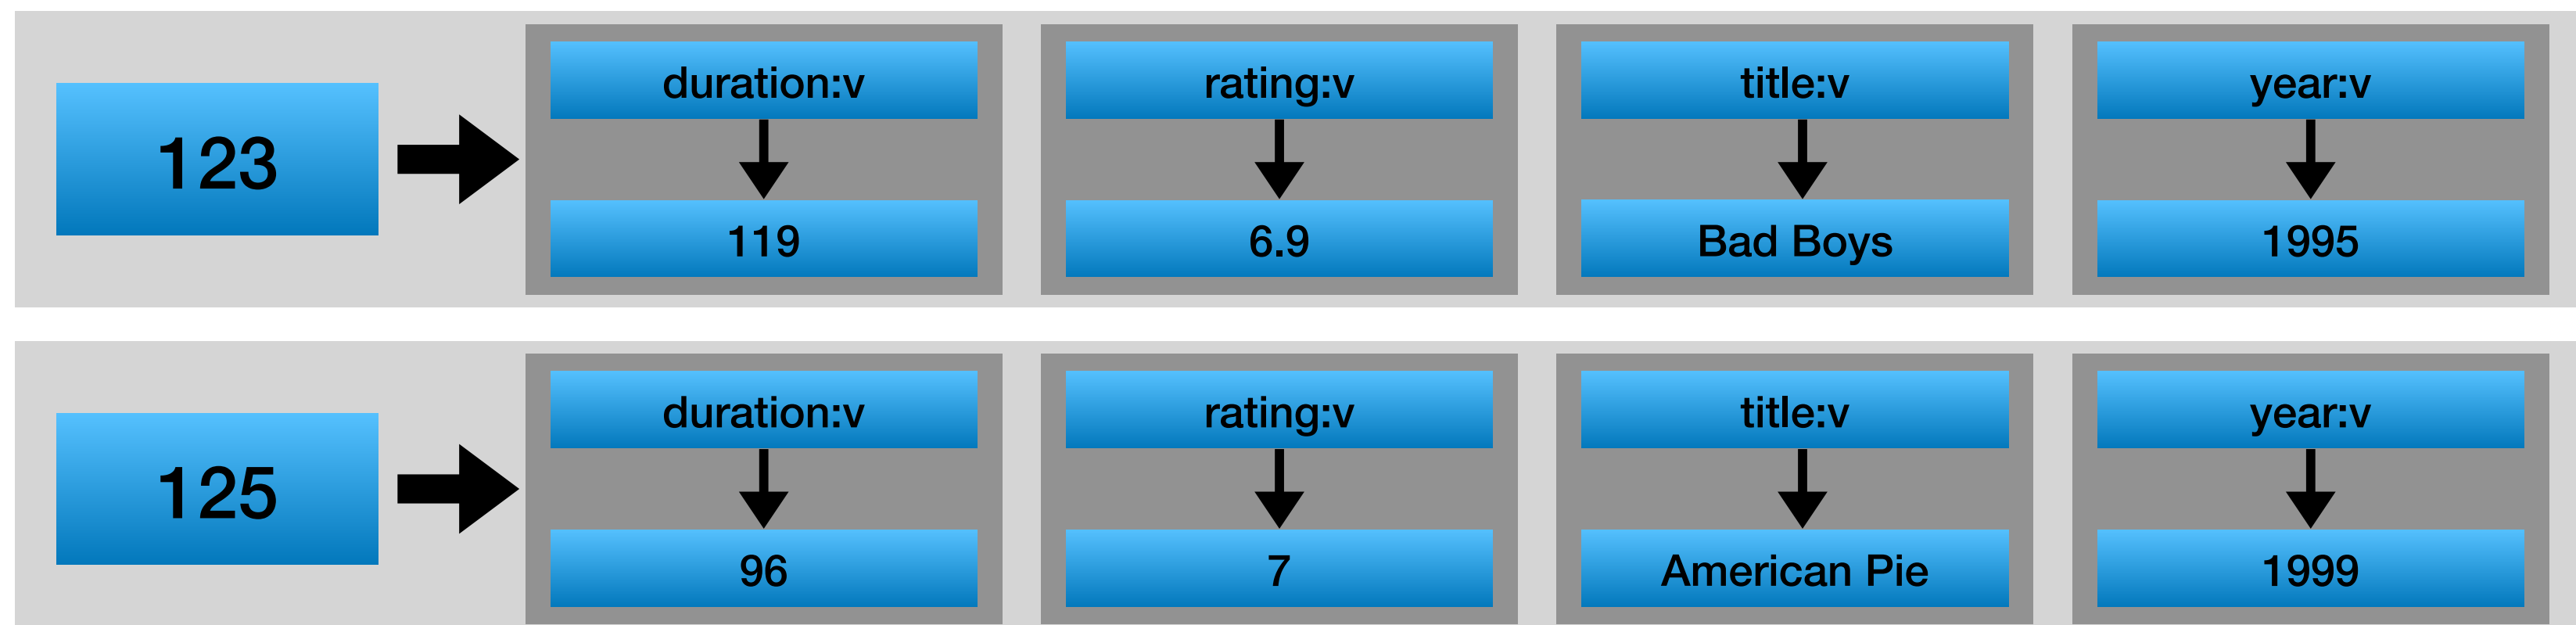
Example (3)

- Raw map of maps (old data model)

row_key	k	v
123	title	Bad Boys
123	year	1995
123	duration	119
123	rating	6.9
125	title	American Pie
125	year	1999
125	duration	96
125	rating	7

```
CREATE TABLE "test_map_of_maps" (
    row_key        TEXT,
    k              TEXT,
    v              BLOB,
    PRIMARY KEY (row_key, k)
);
```

Can you simulate google's Bigtable model (that is, with versions)?
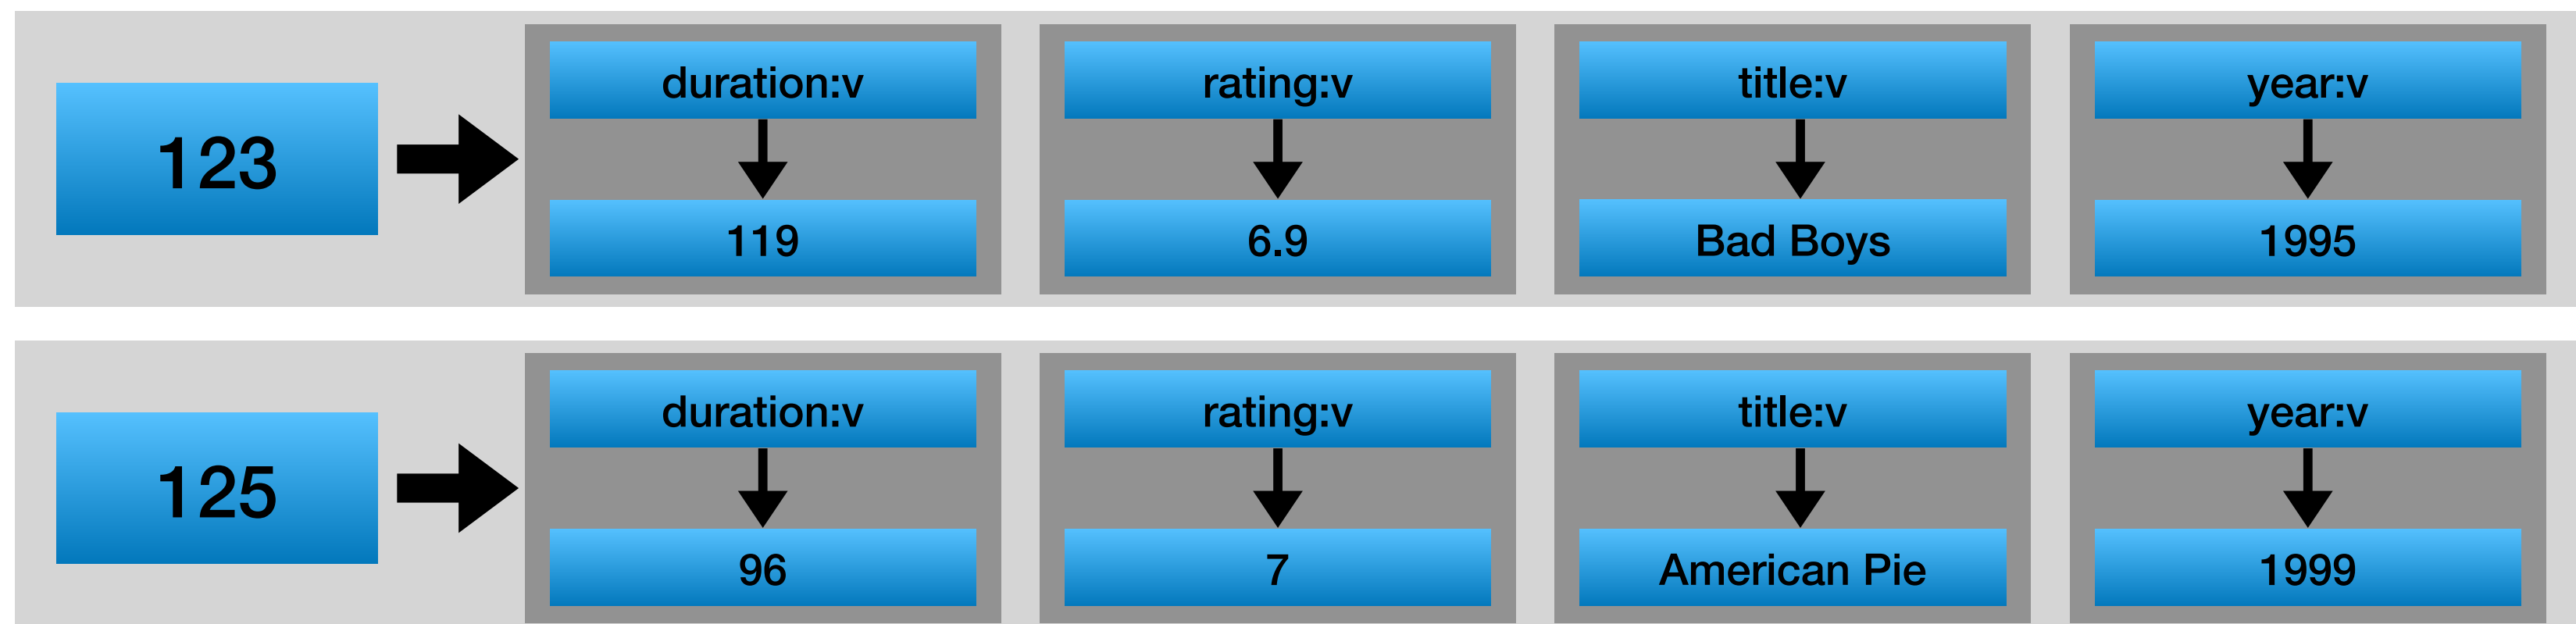
Example (3)

- Bigtable data model
  without sorting on row keys

```
CREATE TABLE "test_map_of_maps" (
  row_key        TEXT,
  k              TEXT,
  t              INT
  v              BLOB,
  PRIMARY KEY (row_key, k, t)
);
```

row_key	k	t	v
123	title	0	Bad Boys
123	year	0	1995
123	duration	0	119
123	rating	0	6.9
123	rating	1	7.2
123	rating	2	7.1
125	title	0	American
125	year	0	1999
125	duration	0	96
125	rating	0	7

t == version