

CC5212-1

PROCESAMIENTO MASIVO DE DATOS

OTOÑO 2020

Lecture 9

NoSQL: Overview

Aidan Hogan

aidhog@gmail.com

Hadoop/MapReduce/Pig/Spark: Processing Un/Structured Information



Information Retrieval:

Storing Unstructured Information

A word cloud of information retrieval terms. The words are arranged in a roughly circular pattern, with some larger than others. The colors of the words range from light blue to dark purple. The terms include: stop-words, information-overload, ranking, lemmatisation, compression, pagerank, heap's-law, keywords, tf-idf, zipf's-law, robots.txt, query, importance, relevance, site-map, DDoS, cosine, crawling, link-analysis, similarity, search, posting-lists, term-frequency, and elias-encoding.

stop-words information-overload
ranking lemmatisation
compression pagerank heap's-law
keywords tf-idf
zipf's-law robots.txt query
importance relevance
site-map DDoS cosine
crawling link-analysis similarity
search posting-lists
term-frequency elias-encoding

Storing Structured Information??



BIG DATA:

STORING STRUCTURED INFORMATION

Relational Databases





Relational Databases: One Size Fits All?

“One Size Fits All”: An Idea Whose Time Has Come and Gone

Michael Stonebraker
*Computer Science and Artificial
Intelligence Laboratory, M.I.T., and
StreamBase Systems, Inc.*
stonebraker@csail.mit.edu

Uğur Çetintemel
*Department of Computer Science
Brown University, and
StreamBase Systems, Inc.*
ugur@cs.brown.edu

Abstract

The last 25 years of commercial DBMS development can be summed up in a single phrase: “One size fits all”. This phrase refers to the fact that the traditional DBMS architecture (originally designed and optimized for business data processing) has been used to support many data-centric applications with widely varying characteristics and requirements.

In this paper, we argue that this concept is no longer applicable to the database market, and that the commercial world will fracture into a collection of independent database engines, some of which may be unified by a common front-end parser. We use examples from the stream-processing market and the data-warehouse market to bolster our claims. We also briefly discuss other markets for which the traditional architecture is a poor fit and argue for a critical rethinking of the current factoring of systems services into products.

of multiple code lines causes various practical problems, including:

- *a cost problem*, because maintenance costs increase at least linearly with the number of code lines;
- *a compatibility problem*, because all applications have to run against every code line;
- *a sales problem*, because salespeople get confused about which product to try to sell to a customer; and
- *a marketing problem*, because multiple code lines need to be positioned correctly in the marketplace.

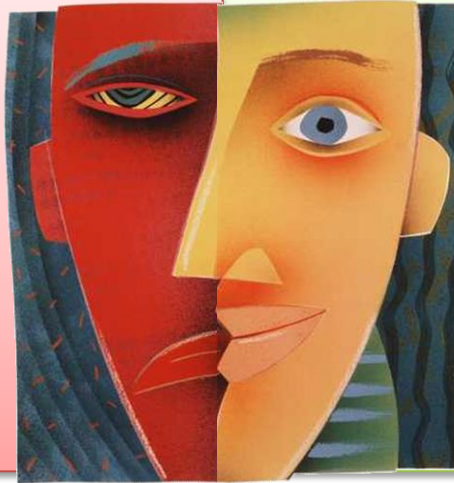
To avoid these problems, all the major DBMS vendors have followed the adage “put all wood behind one arrowhead”. In this paper we argue that this strategy has failed already, and will fail more dramatically off into the future.

The rest of the paper is structured as follows. In Section 2, we briefly indicate why the single code-line strategy has failed already by citing some of the key characteristics of the data warehouse market. In Section

SQL

Difficult to optimise

Difficult to distribute



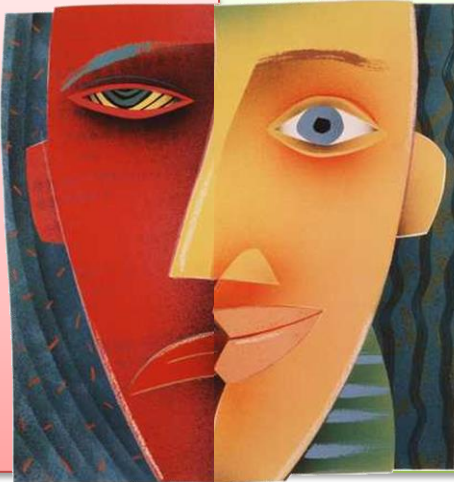
Declarative language

Expressive

ACID

Costly to implement

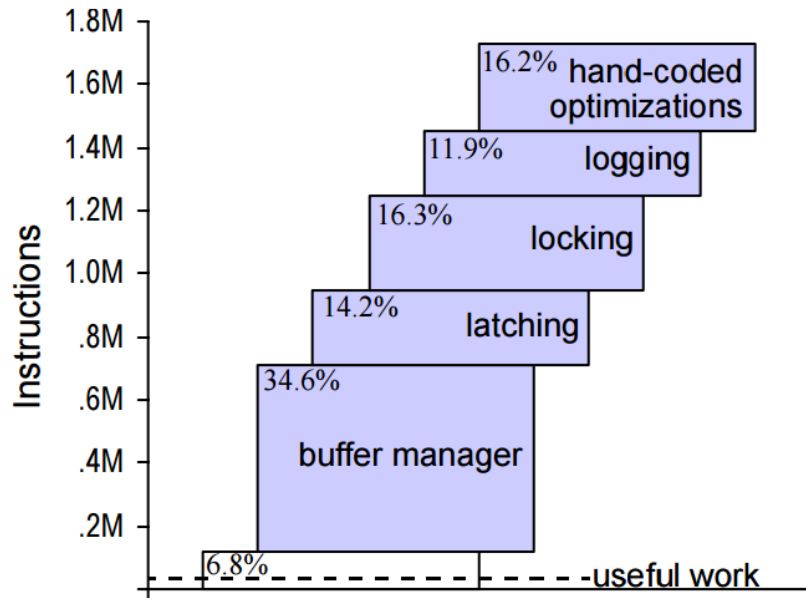
Difficult to distribute



Guarantees correct behaviour

Support transactions

Transactional overhead: the cost of ACID



- 640 transactions per second for system with full transactional support (ACID)
- 12,700 transactions per second for system without logs, transactions or lock scheduling

OLTP Through the Looking Glass, and What We Found There

Stavros Harizopoulos
HP Labs
Palo Alto, CA
stavros@hp.com

Daniel J. Abadi
Yale University
New Haven, CT
dna@cs.yale.edu

Samuel Madden Michael Stonebraker
Massachusetts Institute of Technology
Cambridge, MA
{madden, stonebraker}@csail.mit.edu

ABSTRACT

Online Transaction Processing (OLTP) databases include a suite of features — disk-resident B-trees and heap files, locking-based concurrency control, support for multi-threading — that were optimized for computer technology of the late 1970's. Advances in modern processors, memories, and networks mean that today's computers are vastly different from those of 30 years ago, such that many OLTP databases will now fit in main memory, and most OLTP transactions can be processed in milliseconds or less. Yet database architecture has changed little.

1. INTRODUCTION

Modern general purpose online transaction processing (OLTP) database systems include a standard suite of features: a collection of on-disk data structures for table storage, including heap files and B-trees, support for multiple concurrent queries via locking-based concurrency control, log-based recovery, and an efficient buffer manager. These features were developed to support transaction processing in the 1970's and 1980's, when an OLTP database was many times larger than the main memory, and when the computers that ran these databases cost hundreds of thousands to

ALTERNATIVES TO RELATIONAL DATABASES FOR BIG DATA?

NoSQL

Anybody know anything about NoSQL?




Not
Only **SQL**



Rank			DBMS	Database Model	Score		
Jun 2020	May 2020	Jun 2019			Jun 2020	May 2020	Jun 2019
1.	1.	1.	Oracle	Relational, Multi-model	1343.59	-1.85	+44.37
2.	2.	2.	MySQL	Relational, Multi-model	1277.89	-4.75	+54.26
3.	3.	3.	Microsoft SQL Server	Relational, Multi-model	1067.31	-10.99	-20.45
4.	4.	4.	PostgreSQL	Relational, Multi-model	522.99	+8.19	+46.36
5.	5.	5.	MongoDB	Document, Multi-model	437.08	-1.92	+33.17
6.	6.	6.	IBM Db2	Relational, Multi-model	161.81	-0.83	-10.39
7.	7.	7.	Elasticsearch	Search engine, Multi-model	149.69	+0.56	+0.86
8.	8.	8.	Redis	Key-value, Multi-model	145.64	+2.17	-0.48
9.	9.	11.	SQLite	Relational	124.82	+1.78	-0.07
10.	11.	10.	Cassandra	Wide column	119.01	-0.15	-6.17
11.	10.	9.	Microsoft Access	Relational	117.18	-2.72	-23.83
12.	12.	12.	MariaDB	Relational, Multi-model	89.79	-0.30	+4.59
13.	13.	13.	Splunk	Search engine	88.08	+0.33	+3.46
14.	14.	14.	Hive	Relational	78.65	-2.89	-0.40
15.	15.	15.	Teradata	Relational, Multi-model	73.28	-0.60	-3.36
16.	16.	20.	Amazon DynamoDB	Multi-model	64.87	+0.15	+9.61
17.	17.	21.	SAP Adaptive Server	Relational	53.09	-0.90	-2.03
18.	18.	16.	Solr	Search engine	51.26	-1.32	-9.22
19.	20.	19.	SAP HANA	Relational, Multi-model	50.82	+0.29	-5.56
20.	19.	18.	FileMaker	Relational	50.16	-0.80	-7.64
21.	22.	17.	HBase	Wide column	48.73	-0.99	-9.30
22.	21.	22.	Neo4j	Graph	48.27	-1.49	-1.28
23.	23.	24.	Microsoft Azure SQL Database	Relational, Multi-model	47.78	+5.03	+18.77
24.	24.	25.	Microsoft Azure Cosmos DB	Multi-model	30.80	+0.13	+2.56
25.	25.	23.	Couchbase	Document, Multi-model	29.14	+0.56	-4.22


NoSQL: features vs. scale/performance



NoSQL: common characteristics

- Often **distributed**
- Often **simpler** languages than SQL
- **Different flavours** (for different scenarios)

NoSQL: four main flavours



LIMITATIONS OF DISTRIBUTED COMPUTING: CAP THEOREM

What is CAP?

Three *guarantees* a distributed sys. could make

1. Consistency:

- All nodes have a consistent view of the system

2. Availability:

- Every read/write is acted upon


3. Partition-tolerance:

- The system works even if messages are lost

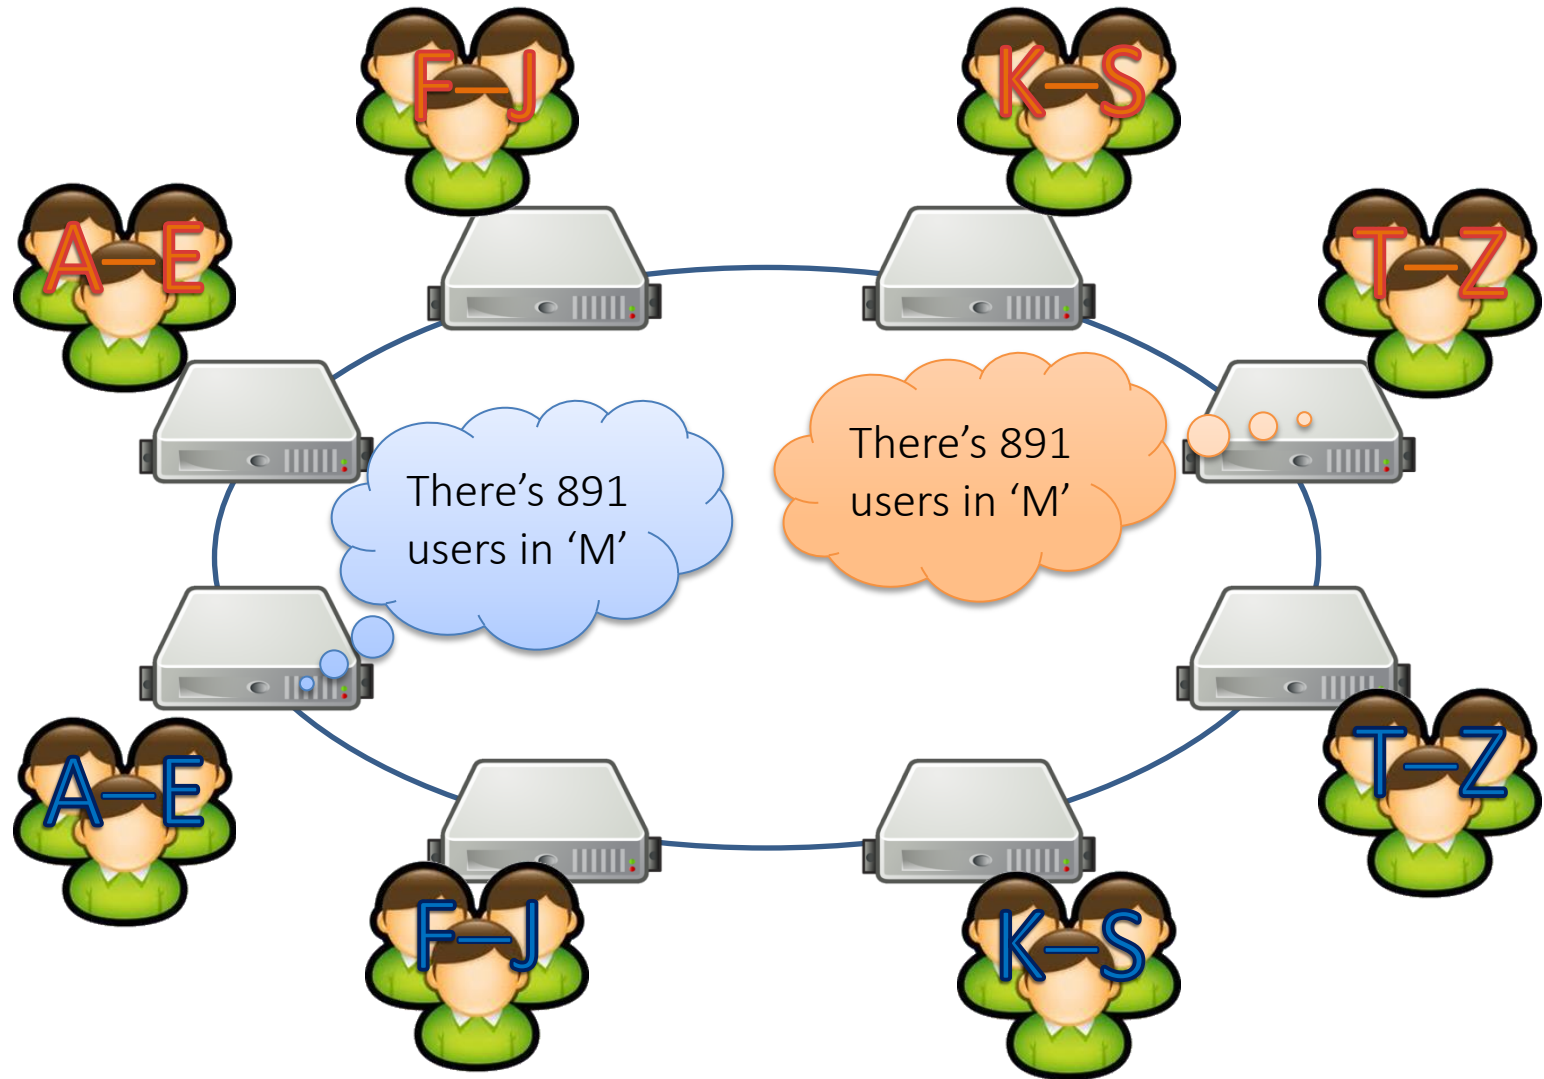
CA in CAP not the same as CA in ACID!!




A Distributed System (with Replication)



Consistency




Availability



How many users start with 'M'

Partition tolerance



The CAP Question

Can a distributed system guarantee

consistency (all nodes have the same up-to-date view),

availability (every read/write is acted upon) **and**


partition-tolerance (the system works if messages are lost)

at the same time?

What do you think?




The CAP Answer




The CAP Theorem

A distributed system cannot guarantee
consistency (all nodes have the same up-to-date view),
availability (every read/write is acted upon) **and**
partition-tolerance (the system works if messages are lost)
at the same time!

The CAP “Proof”




The CAP Triangle



CAP Systems

CA: Guarantees to give a correct response but only while network works fine (*Centralised / Traditional*)

CP: Guarantees responses are correct even if there are network failures, but response may fail (*Weak availability*)



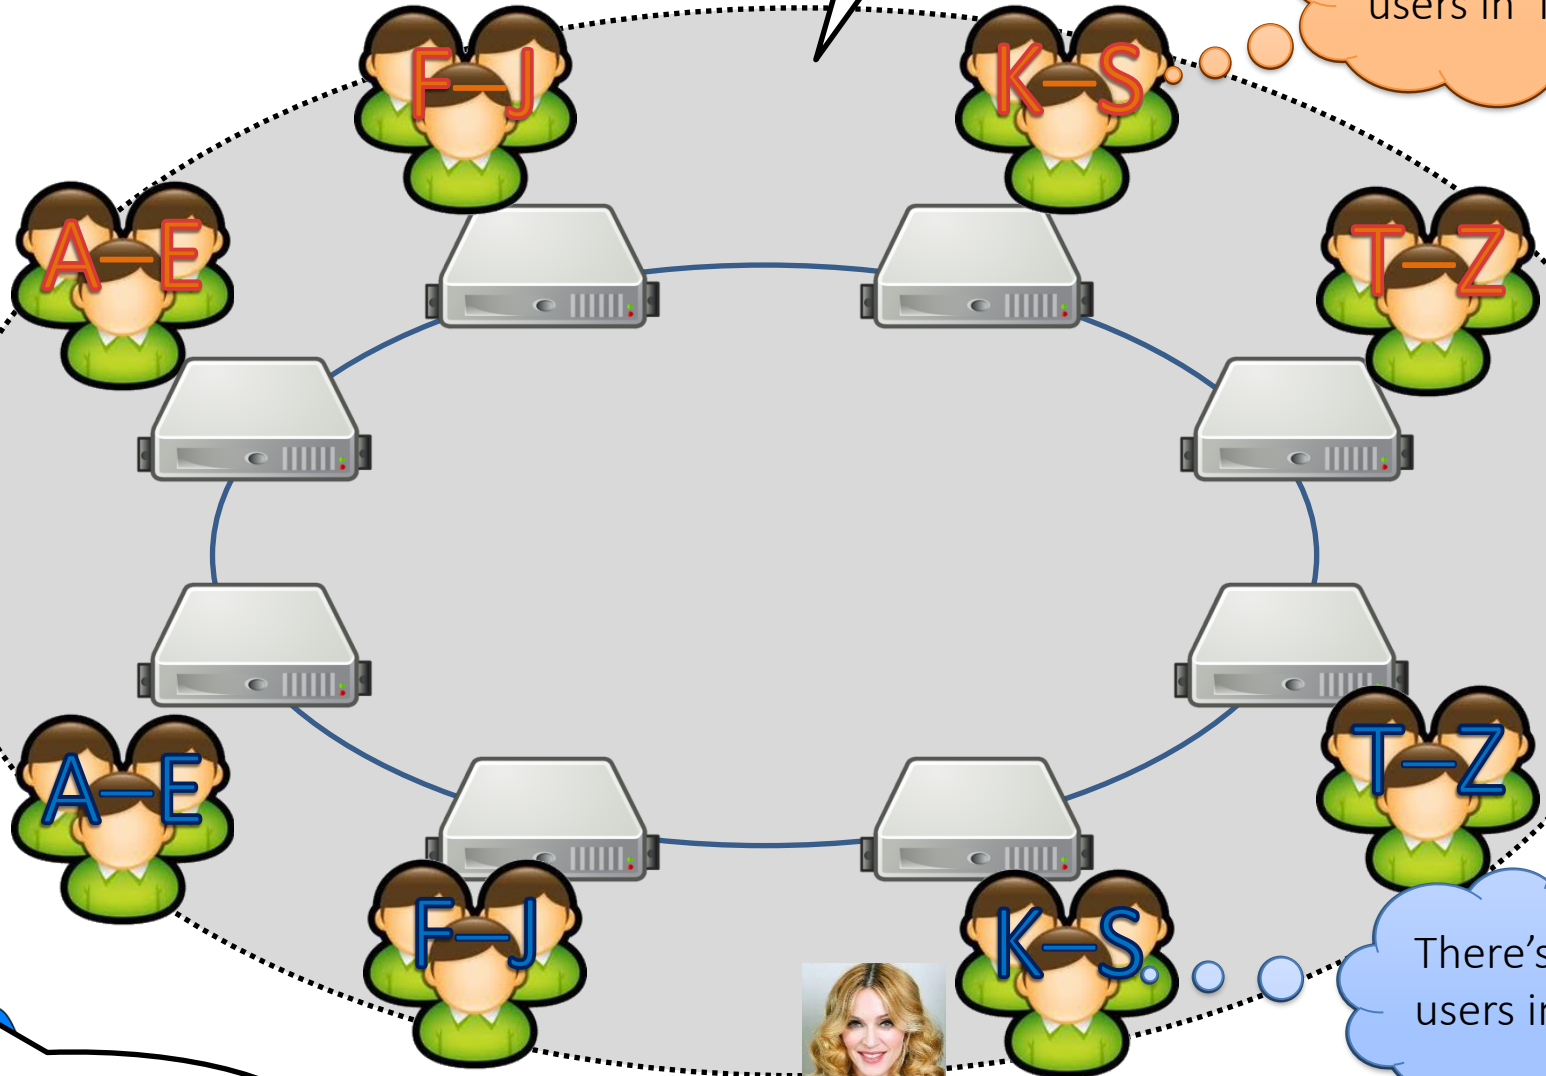
(No intersection)

AP: Always provides a “best-effort” response even in presence of network failures (*Eventual consistency*)

CA System

892

There's 892 users in 'M'




How many users start with 'M'

There's 892 users in 'M'

CP System

Error

There's 891 users in 'M'



How many users start with 'M'


There's 892 users in 'M'



AP System

891

There's 891 users in 'M'



How many users start with 'M'

There's 892 users in 'M'

BASE (AP)

- **B**asically **A**vailable
 - Almost always “up”
- **S**oft State
 - Replicated, cached data
- **E**ventual Consistency
 - Stale data tolerated, for a while

In what way does Twitter act as a BASE (AP) system?




High-fanout creates a “partition”

Rank ↕	Change (monthly)	Account name ↕	Owner ↕	Followers (millions) ↕	Activity ↕	Country ↕
1	—	@BarackObama	Barack Obama	120	Former U.S. president	United States
2	—	@justinbieber	Justin Bieber	112	Musician	Canada
3	—	@katyperry	Katy Perry	108	Musician	United States
4	—	@rihanna	Rihanna	97	Musician and businesswoman	Barbados
5	—	@taylorswift13	Taylor Swift	87	Musician	United States

Users may see retweets of celebrity tweets before the original tweet.

Later when the original tweet arrives the timeline will be reordered and made consistent.


CAP in practical distributed systems



1. Fix **P**
2. Choose trade-off point between **C** and **A**

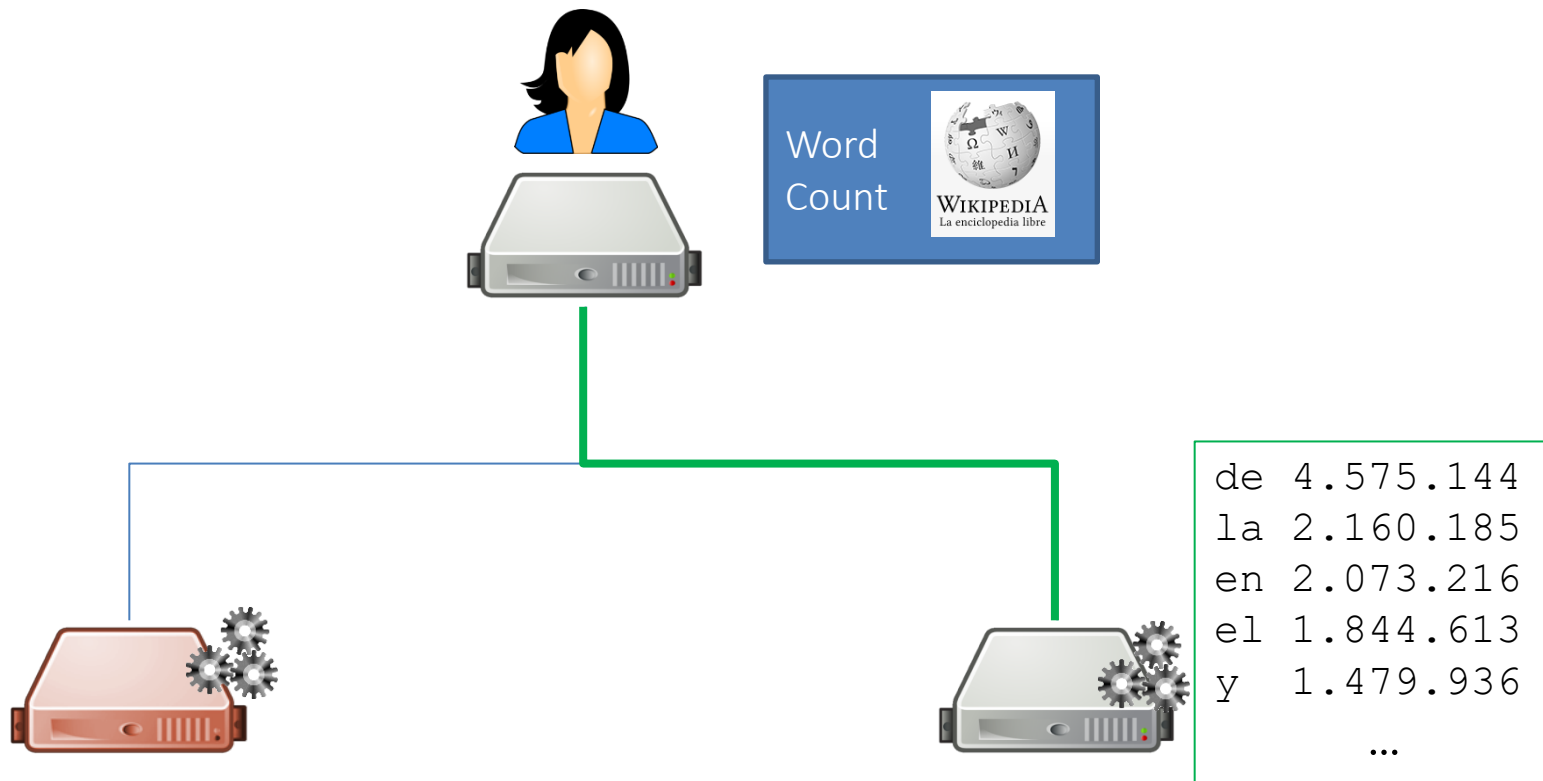
PARTITION TOLERANCE

Faults



Fail-Stop Fault

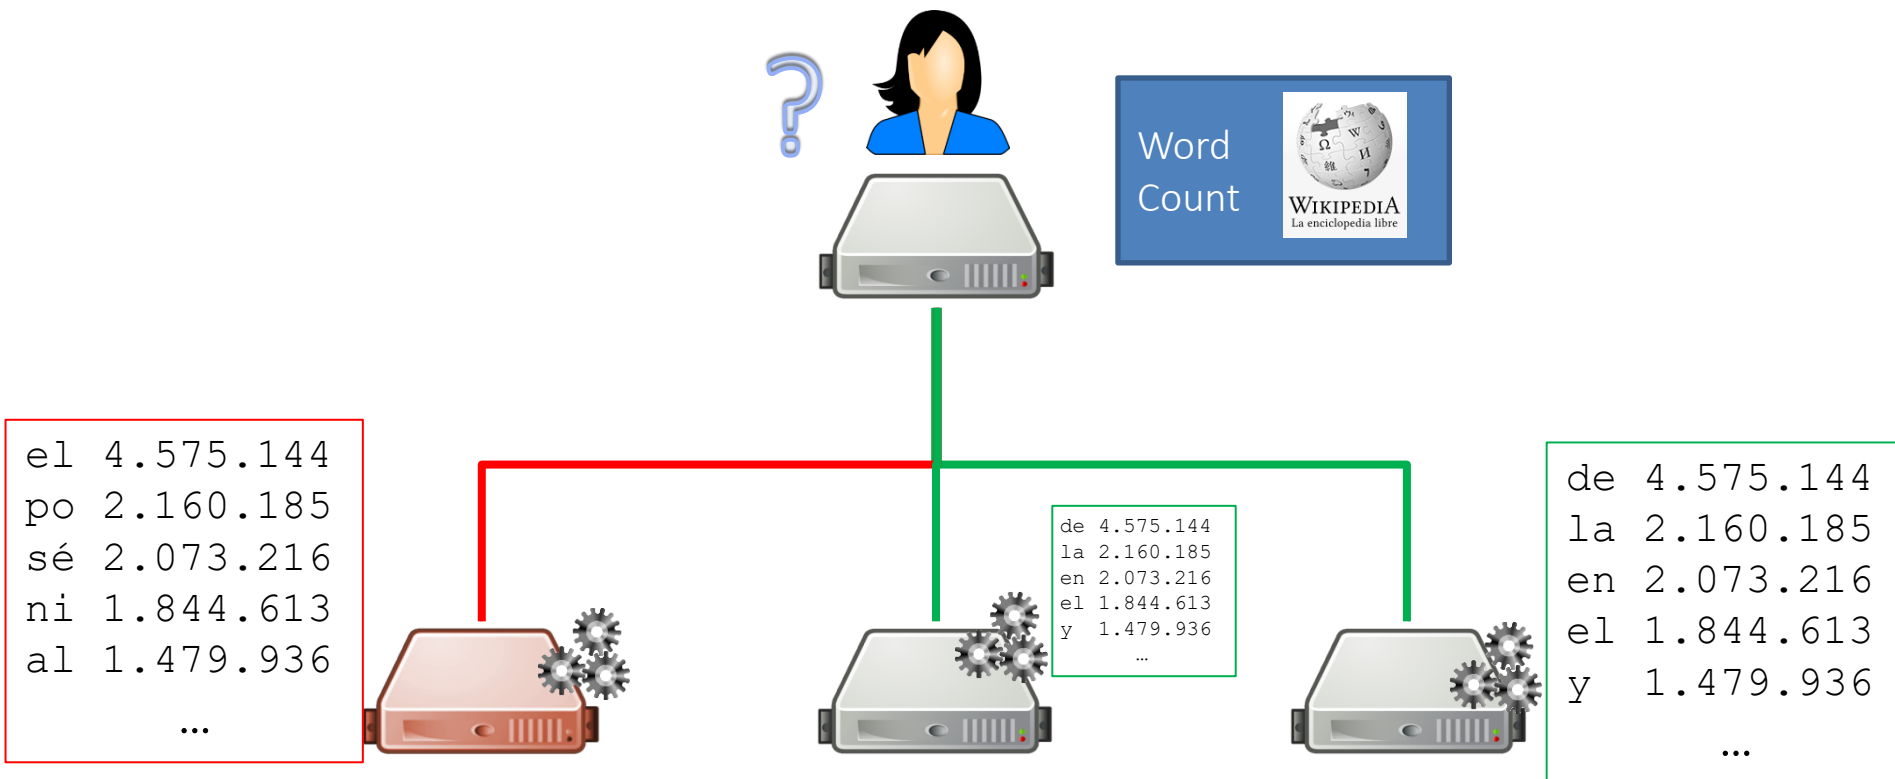
- A machine fails to respond or times-out
 - often hardware or load
 - need at least $f + 1$ replicated machines
 - f = number of fail-stop failures



Byzantine Fault


- A machine responds incorrectly/maliciously

How many working machines do we need in the general case to be robust against Byzantine faults?



Byzantine Fault

- A machine responds incorrectly/maliciously
 - Need *at least* $2f+1$ replicated machines
 - f = number of (possibly Byzantine) failures



DISTRIBUTED CONSENSUS

Distributed Consensus




Colour of the dress?




Distributed Consensus

Strong consensus: All nodes need to agree




Distributed Consensus

Strong consensus: All nodes need to agree



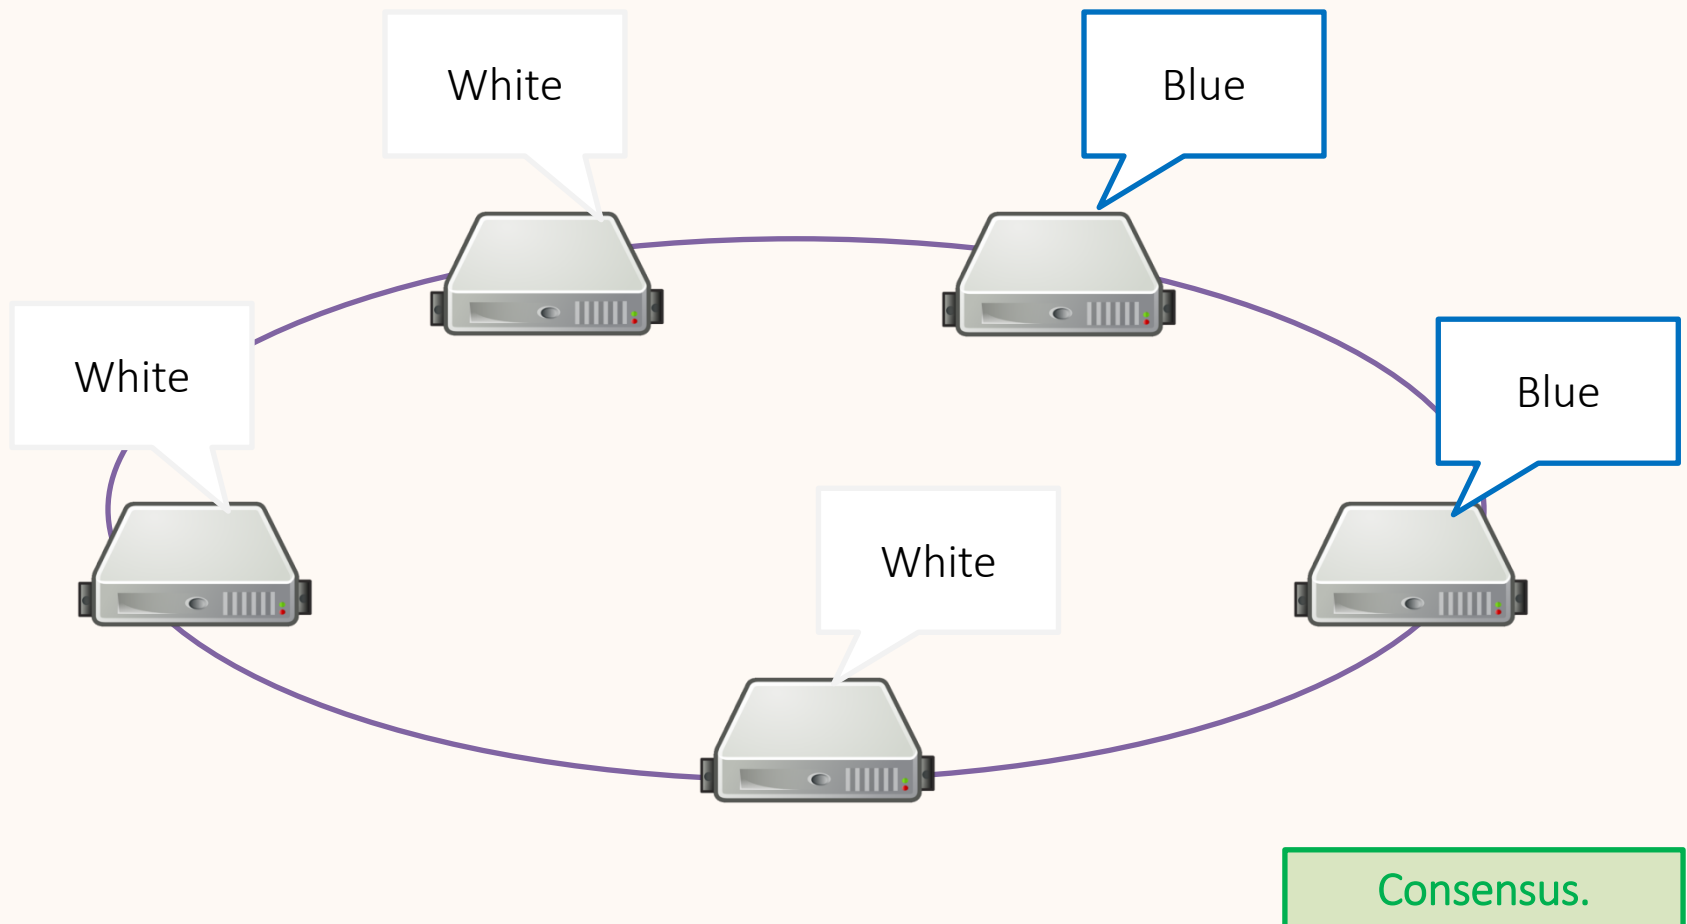
Distributed Consensus

Majority consensus: A majority of nodes need to agree



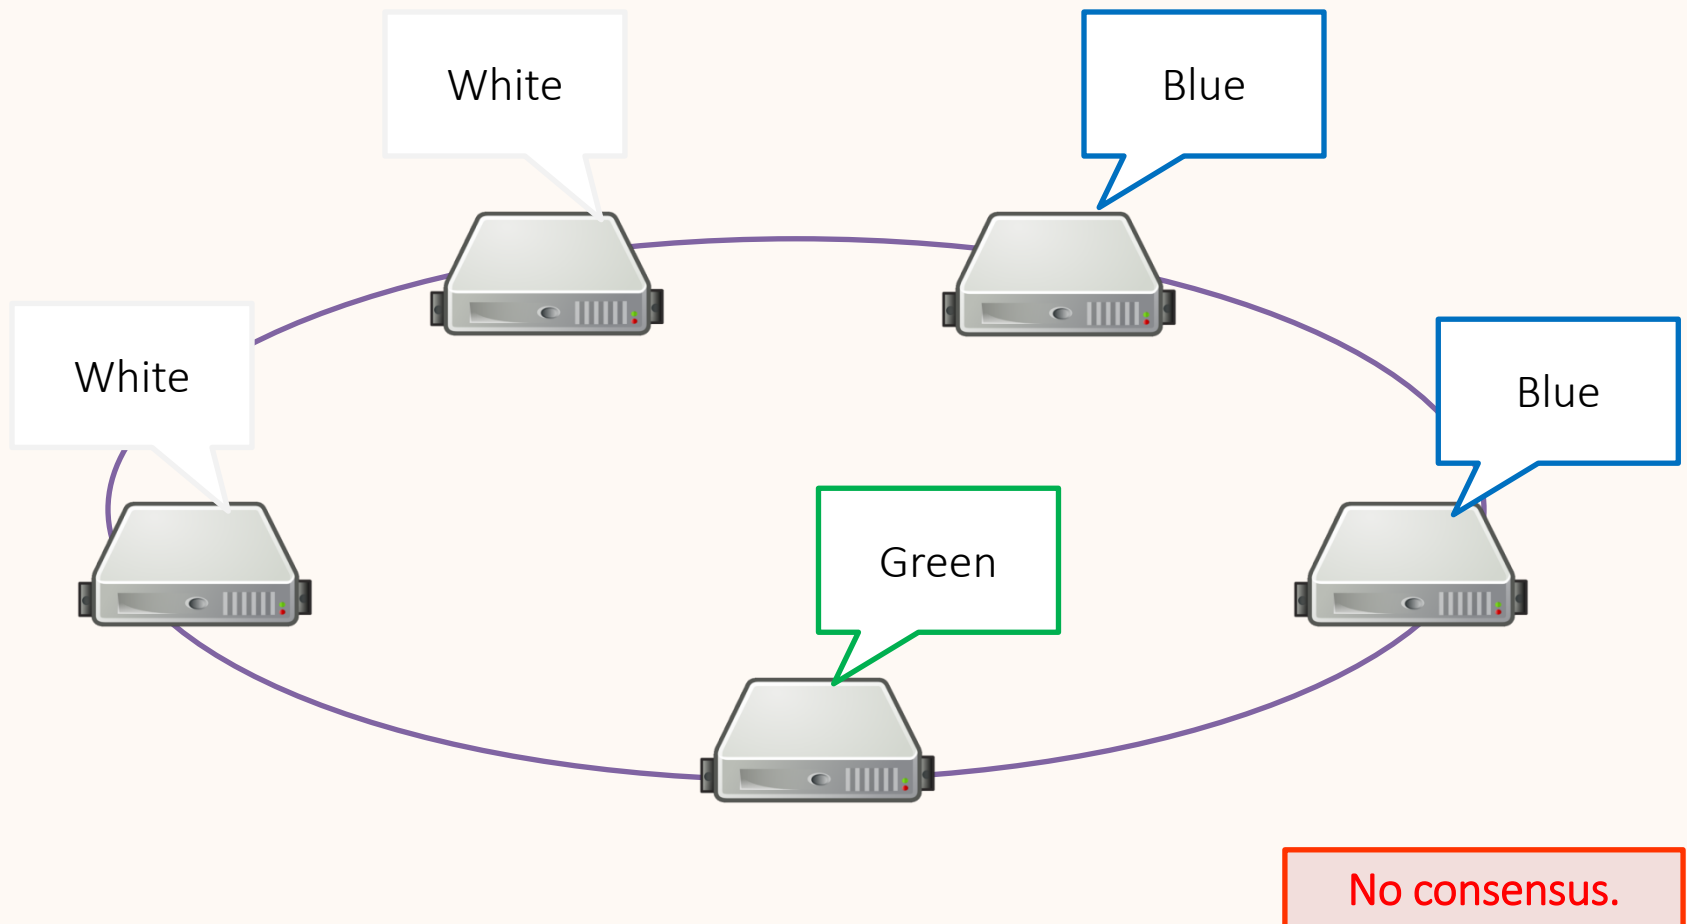
Distributed Consensus

Majority consensus: A majority of nodes need to agree



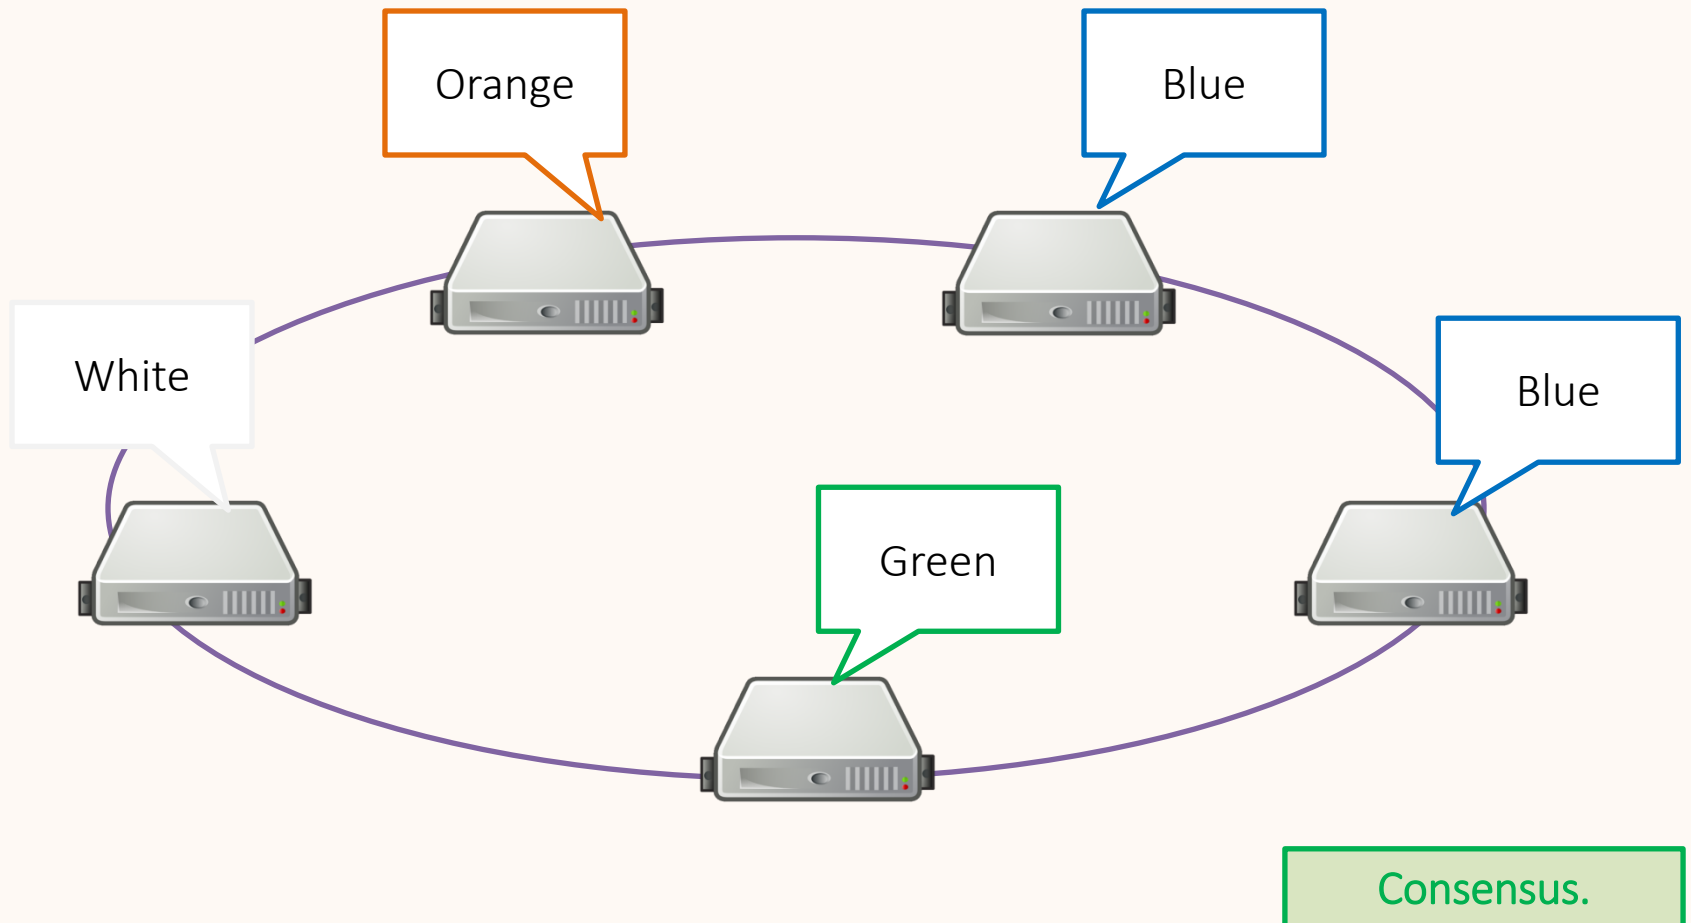
Distributed Consensus

Majority consensus: A majority of nodes need to agree



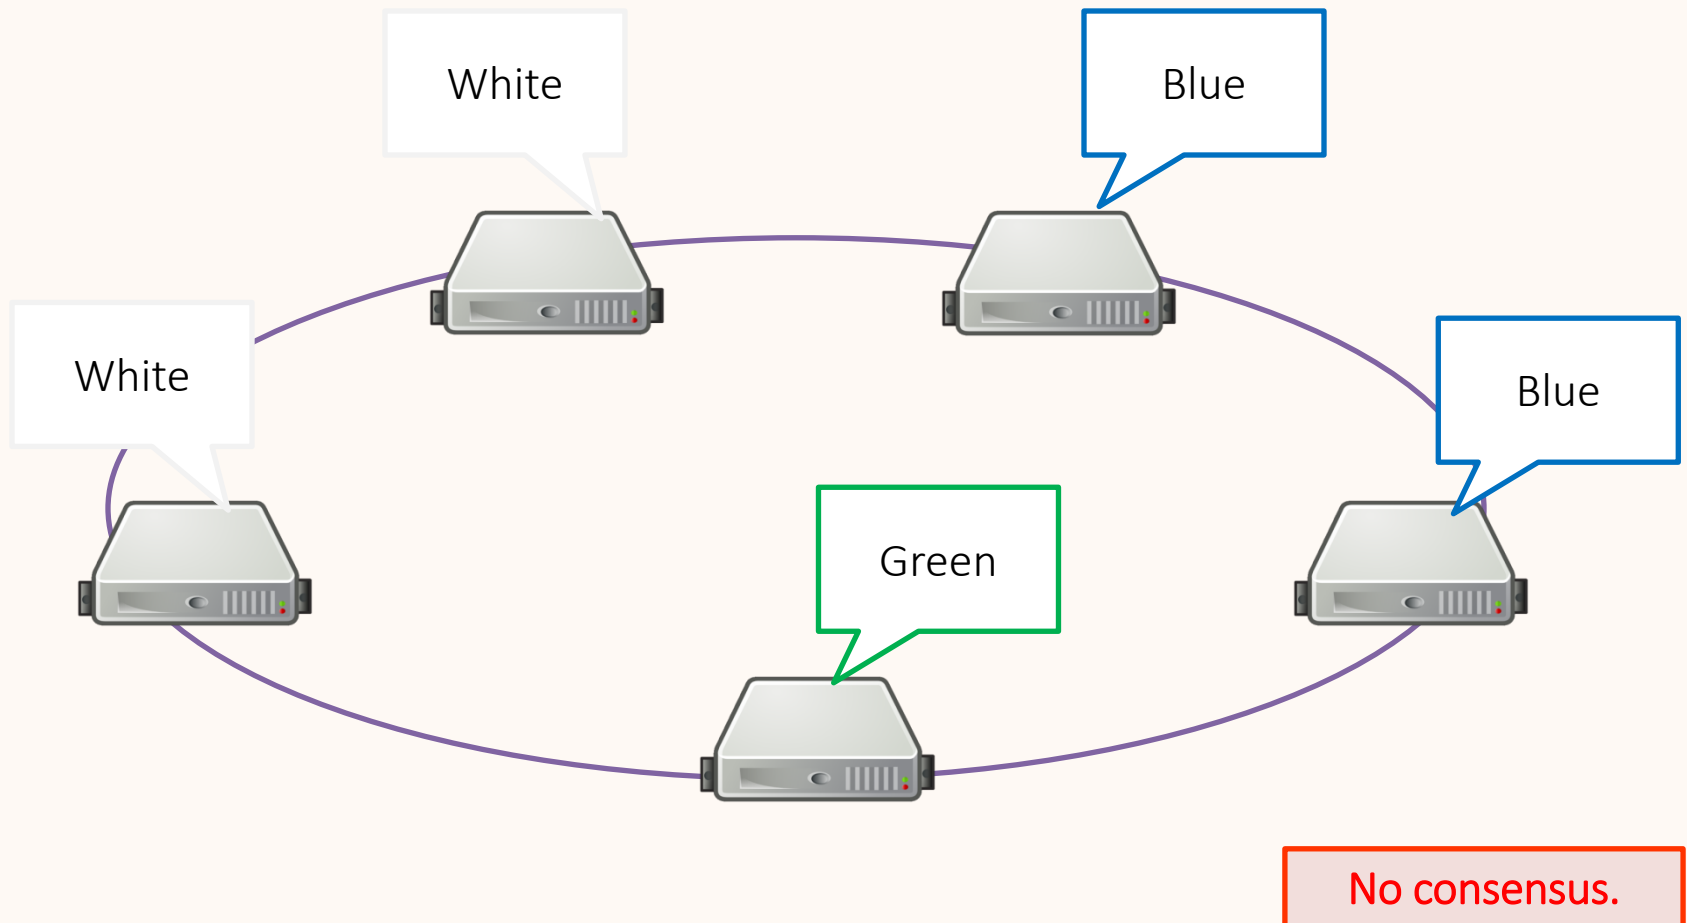
Distributed Consensus

Plurality consensus: A plurality of nodes need to agree



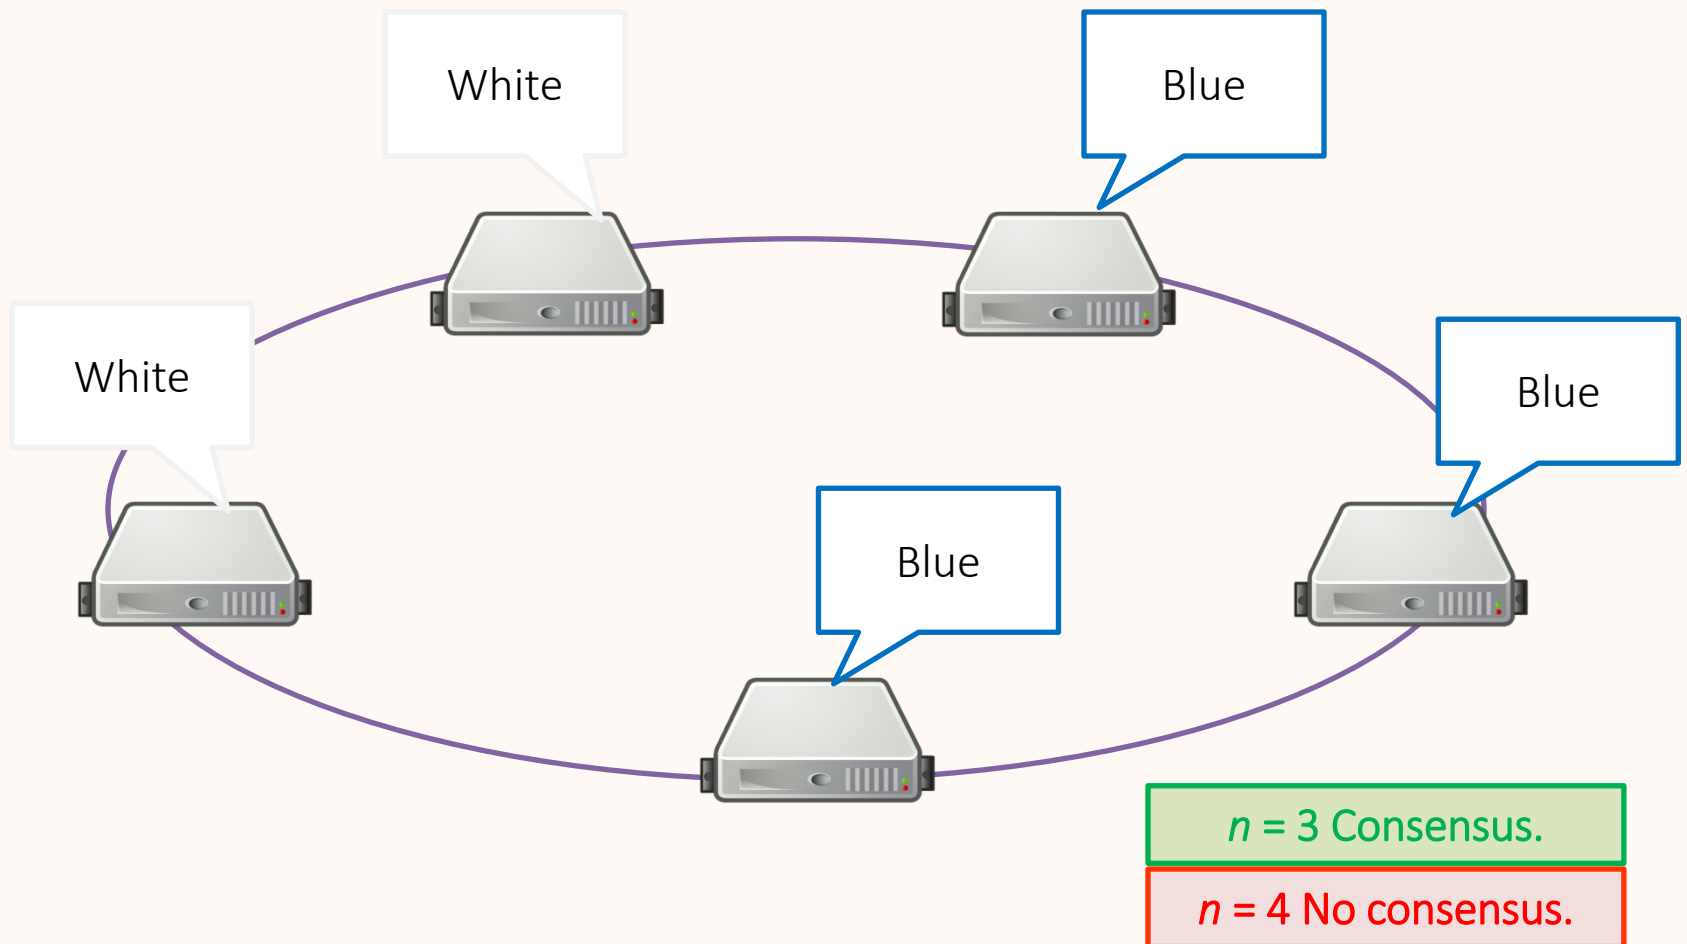
Distributed Consensus

Plurality consensus: A plurality of nodes need to agree



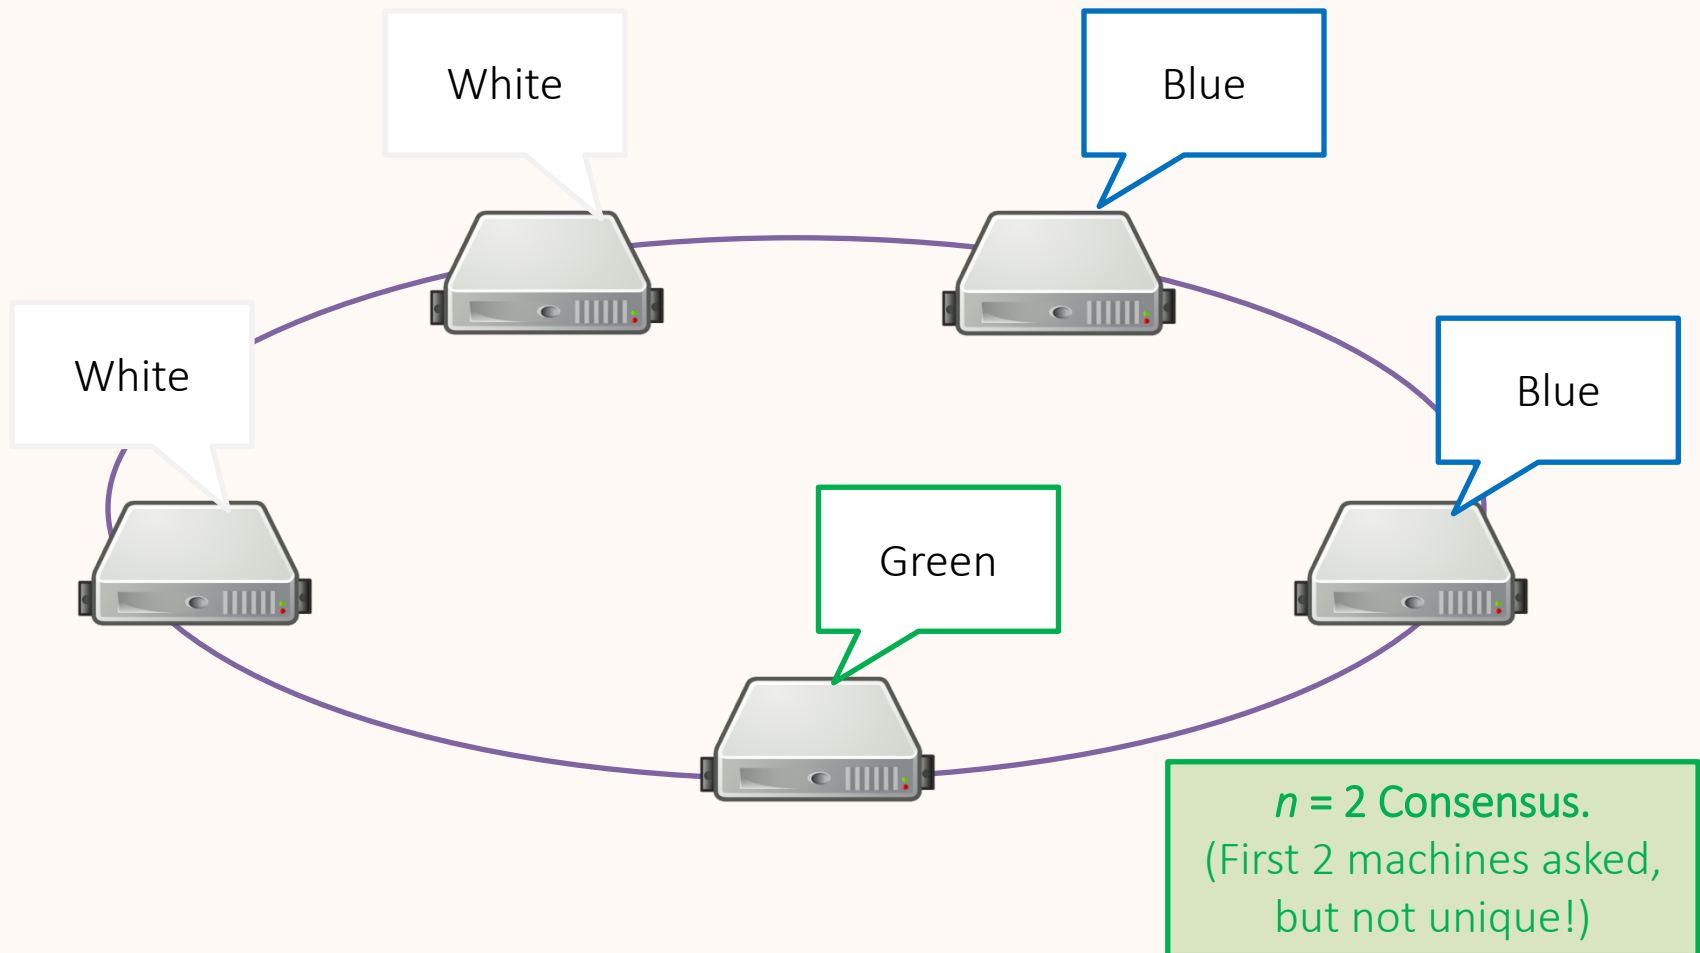
Distributed Consensus

Quorum consensus: n nodes need to agree



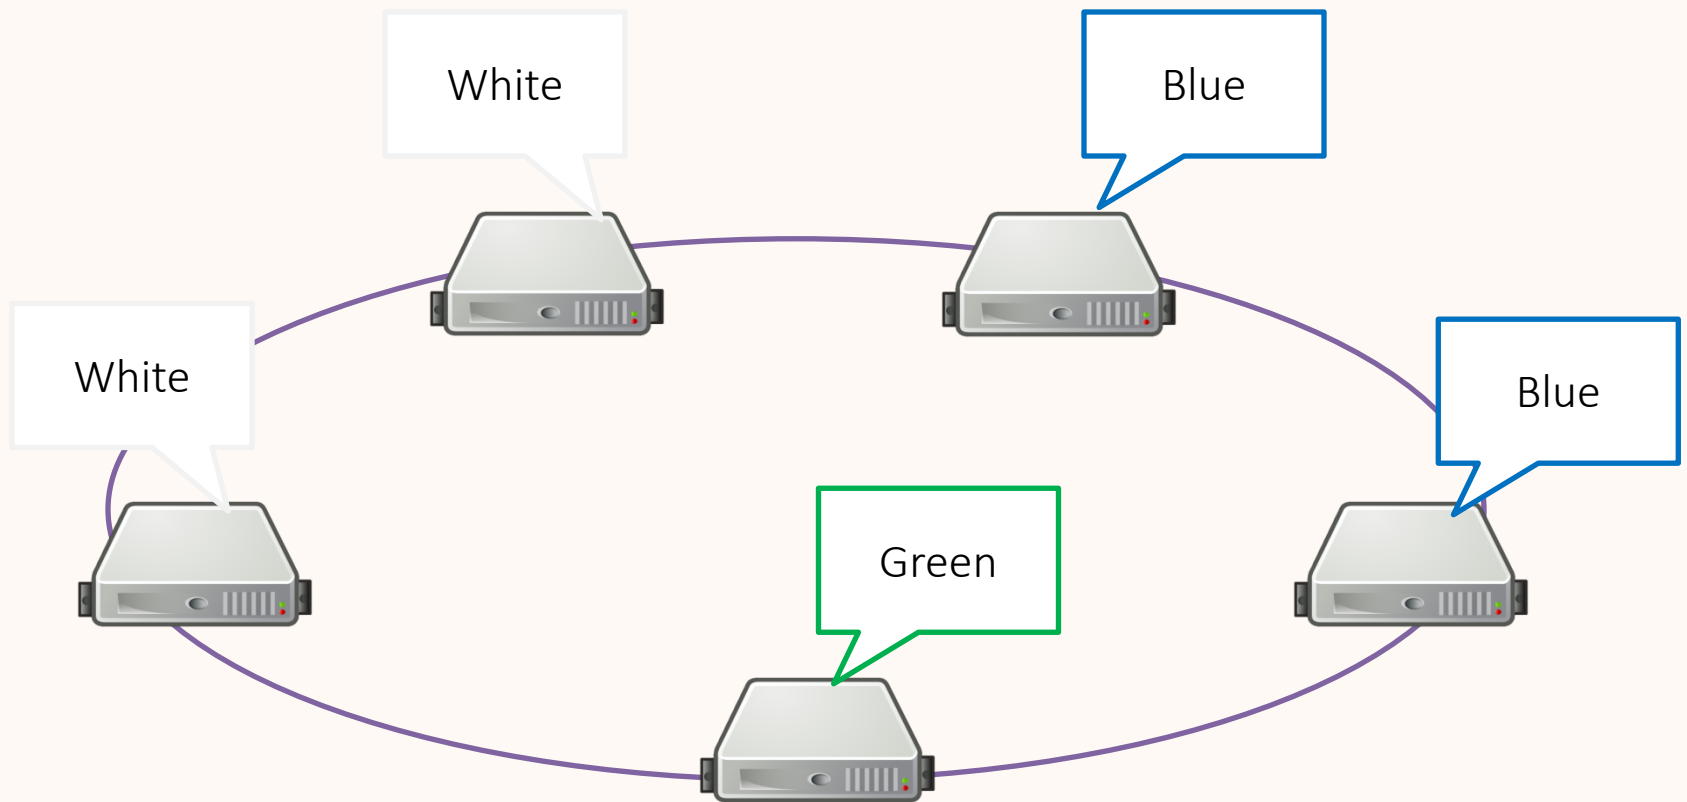
Distributed Consensus

Quorum consensus: n nodes need to agree



Distributed Consensus

Quorum consensus: n nodes need to agree



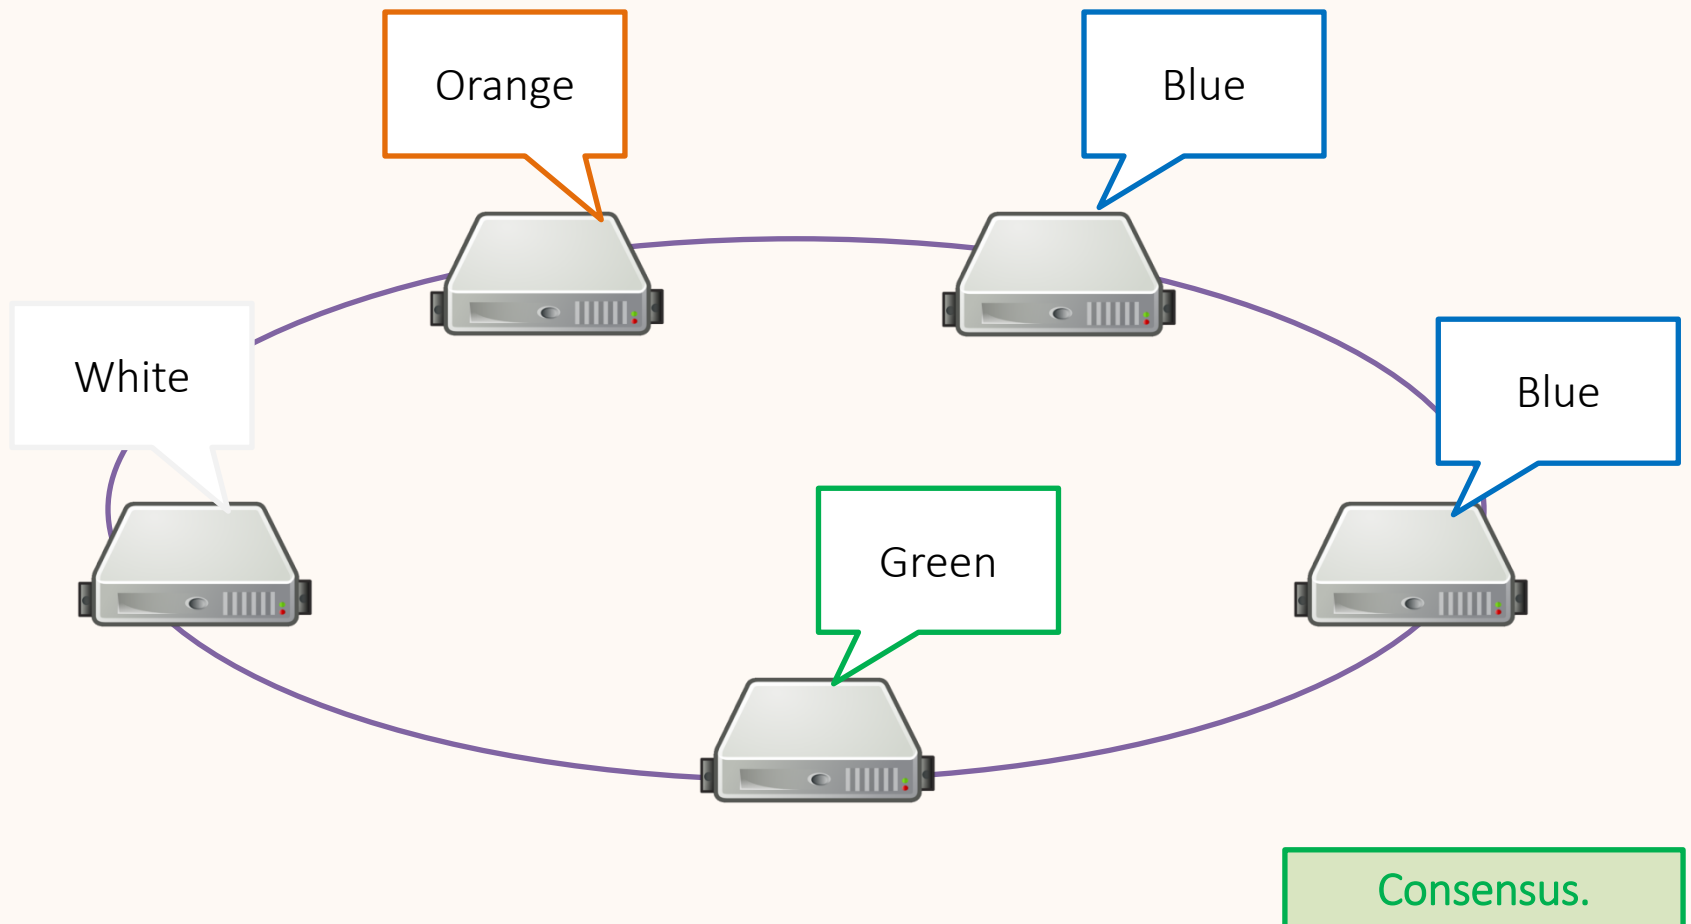
Value of n needed for unique consensus with N nodes?



$n > N/2$

Distributed Consensus

Consensus off: Take first answer



Distributed Consensus

CP vs. AP?



Strong consensus: All nodes need to agree

CP

Majority consensus: A majority of nodes need to agree

Plurality consensus: A plurality of nodes need to agree

Quorum consensus: “Fixed” n nodes need to agree

Consensus off: Take first answer

AP

Distributed Consensus

Scale?



Strong consensus: All nodes need to agree

More replication

Majority consensus: A majority of nodes need to agree

Plurality consensus: A plurality of nodes need to agree

Quorum consensus: “Fixed” n nodes need to agree

Consensus off: Take first answer

Less replication

Distributed Consensus

Strong consensus: All nodes need to agree

Majority consensus: A majority of nodes need to agree

Choice is application dependent:


Plurality Many NoSQL stores allow you to choose
level of consensus/replication

Quorum consensus: “Fixed” n nodes need to agree

Consensus off: Take first answer

NoSQL: KEY-VALUE STORES

NoSQL: Key-Value Stores



Key–Value Store Model

It's just a Map / Associate Array / Dictionary 😊

- `put (key, value)`
- `get (key)`
- `delete (key)`

Key	Value
Afghanistan	Kabul
Albania	Tirana
Algeria	Algiers
Andorra la Vella	Andorra la Vella
Angola	Luanda
Antigua and Barbuda	St. John's
...	...

But You Can Do a Lot With a Map

Key	Value
country:Afghanistan	capital@city:Kabul,continent:Asia,pop:31108077#2011
country:Albania	capital@city:Tirana,continent:Europe,pop:3011405#2013
...	...
city:Kabul	country:Afghanistan,pop:3476000#2013
city:Tirana	country:Albania,pop:3011405#2013
...	...
user:10239	basedIn@city:Tirana,post:{103,10430,201}
...	...

... actually you can model any data in a map (but possibly with a lot of redundancy and inefficient lookups if unsorted).

THE CASE OF AMAZON

The Amazon Scenario

Products Listings: prices, details, stock

The screenshot shows an Amazon search results page for the query "presenter". The top navigation bar includes the Amazon logo, a search bar with "presenter" entered, and a "Shop by Department" dropdown menu. Below the navigation bar, the search results are displayed. On the left side, there is a sidebar with navigation links for "Office Products", "Computers & Accessories", "Cell Phones & Accessories", and "Software". Below these links, there is a "Refine by" section with filters for "International Shipping", "Eligible for Free Shipping", and "Brand". The main content area shows a list of products. The first product is "Point and zoom in presentations with Myo Armband" with a "Shop now" link. Below this, there is a "Related Searches" section with links for "logitech presenter", "mpow presenter", and "wireless presenter". The second product is "Logitech Wireless Presenter R400" by Logitech, priced at \$44.29 (was \$49.99) with Prime shipping, available by Wednesday, May 27. The third product is "Logitech Professional Presenter R800 with Green Laser Pointer" by Logitech, priced at \$57.49 (was \$79.99) with Prime shipping, available by Wednesday, May 27. The fourth product is "Kensington Wireless Presenter with Laser Pointer" by Kensington.

amazon
Try Prime

All ▾ presenter

Shop by Department ▾ Aidan's Amazon.com Today's Deals Gift Cards Sell Help

1-16 of 19,088 results for "presenter"

Show results for

Office Products >
Office Presentation Remotes
Office Presentation Pointers

Computers & Accessories >
Tablet Accessories
Computer Mice

Cell Phones & Accessories >
Cell Phone Accessories

Software >
Presentations


★ See All 29 Departments

Refine by


International Shipping
 Ship to Ireland


Eligible for Free Shipping
Free Shipping by Amazon


Brand
 Kensington
 Logitech
 Targus
 Satechi
 Infiniter
 August

 **Point and zoom in presentations with Myo Armband**
Shop now ▶

Related Searches: [logitech presenter](#), [mpow presenter](#), [wireless presenter](#).


 **Logitech Wireless Presenter R400**
by Logitech
\$44.29 ~~\$49.99~~ ✓Prime
Get it by **Wednesday, May 27**

 **Logitech Professional Presenter R800 with Green Laser Pointer**
by Logitech
\$57.49 ~~\$79.99~~ ✓Prime
Get it by **Wednesday, May 27**

 **Kensington Wireless Presenter with Laser Pointer**
by Kensington





The Amazon Scenario

Customer info: shopping cart, account, etc.

 **Shopping Cart** Already a customer?
[Sign in](#)

[See more items like those in your Cart](#)


subtotal = \$88.77
Make any changes below? [Update](#)

Shopping Cart Items--To Buy Now		Price:	Qty:
<small>Item added on May 22, 2009</small>	The Principles of Beautiful Web Design - Jason Beard; Paperback Condition: New In Stock	\$26.37 You Save: \$13.58 (34%)	<input type="text" value="1"/>
Save for later	Delete	 Eligible for FREE Super Saver Shipping	
	Add gift-wrap/note  (Learn more)		
<small>Item added on May 22, 2009</small>	Don't Make Me Think: A Common Sense Approach to Web Usability, 2nd Edition - Steve Krug; Paperback Condition: New In Stock	\$26.40 You Save: \$13.60 (34%)	<input type="text" value="1"/>
Save for later	Delete	 Eligible for FREE Super Saver Shipping	
	Add gift-wrap/note  (Learn more)		

The Amazon Scenario

Recommendations, etc.:

Customers Who Bought This Item Also Bought



Oliver Twist (Dover Thrift Editions)

> Charles Dickens

★★★★☆ (213)

Paperback

\$3.50



David Copperfield (Dover Thrift Editions)

> Charles Dickens

★★★★☆ (196)

Paperback

\$5.00



JANE EYRE

> Charlotte Bronte


★★★★☆ (1,045)

Paperback

\$2.99

The Amazon Scenario

- Amazon customers:




The Amazon Scenario



The Amazon Scenario

Databases struggling ...



But many Amazon services don't need:

- SQL (a simple map often enough)

or even:

- transactions, strong consistency, etc.

Key–Value Store: Amazon Dynamo(DB)

Dynamo: Amazon's Highly Available Key-value Store


Giuseppe DeCandia, Deniz Hastorun, Madan Jampani, Gunavardhan Kakulapati, Avinash Lakshman, Alex Pilchin, Swaminathan Sivasubramanian, Peter Vosshall and Werner Vogels

Amazon.com

ABSTRACT

Reliability at massive scale is one of the biggest challenges we face at Amazon.com, one of the largest e-commerce operations in the world; even the slightest outage has significant financial consequences and impacts customer trust. The Amazon.com platform, which provides services for many web sites worldwide, is implemented on top of an infrastructure of tens of thousands of servers and network components located in many datacenters

One of the lessons our organization has learned from operating Amazon's platform is that the reliability and scalability of a system is dependent on how its application state is managed. Amazon uses a highly decentralized, loosely coupled, service oriented architecture consisting of hundreds of services. In this environment there is a particular need for storage technologies that are always available. For example, customers should be able to view and add items to their shopping cart even if disks are failing, network routes are flapping, or data centers are being



Goals:

- Scalability (able to grow)
- High availability (reliable)
- Performance (fast)

Don't need full SQL, don't need full ACID


Key–Value Store: Distribution

How might we distribute a key–value store over multiple machines?



Key	Value
country:Afghanistan	capital@city:Kabul,continent:Asia,pop:31108077#2011
country:Albania	capital@city:Tirana,continent:Europe,pop:3011405#2013
...	...
city:Kabul	country:Afghanistan,pop:3476000#2013
city:Tirana	country:Albania,pop:3011405#2013
...	...
user:10239	basedIn@city:Tirana,post:{103,10430,201}
...	...

$$\text{mod}(\text{hash}(key), m)$$



Key–Value Store: Distribution

What happens if a machine leaves or joins afterwards?




How can we avoid rehashing everything?



Key	Value
country:Afghanistan	capital@city:Kabul,continent:Asia,pop:31108077#2011
country:Albania	capital@city:Tirana,continent:Europe,pop:3011405#2013
...	...
city:Kabul	country:Afghanistan,pop:3476000#2013
city:Tirana	country:Albania,pop:3011405#2013
...	...
user:10239	basedIn@city:Tirana,post:{103,10430,201}
...	...

$$\text{mod}(\text{hash}(\textit{key}), m)$$




Consistent Hashing

Avoid re-hashing everything

- Hash using a ring
- Each machine picks n pseudo-random points on the ring
- Machine responsible for arc after its point
- Objects mapped to ring
- If a machine leaves, its range moves to previous machine
- If a machine joins, it picks new points


How many keys (on average) would need to be moved if a machine joins or leaves?



Amazon Dynamo: Hashing




- Consistent Hashing (128-bit MD5)




Amazon Dynamo: Replication

- A set replication factor (e.g., 3)
- Commonly primary / secondary replicas
 - Primary replica elected from secondary replicas in the case of failure of primary



Amazon Dynamo: Replication

- Replication factor of n ?
 - Easy: pick n next buckets (different machines!)



Amazon Dynamo: Model

- Named table with primary key and a value
- Primary key is hashed / unordered

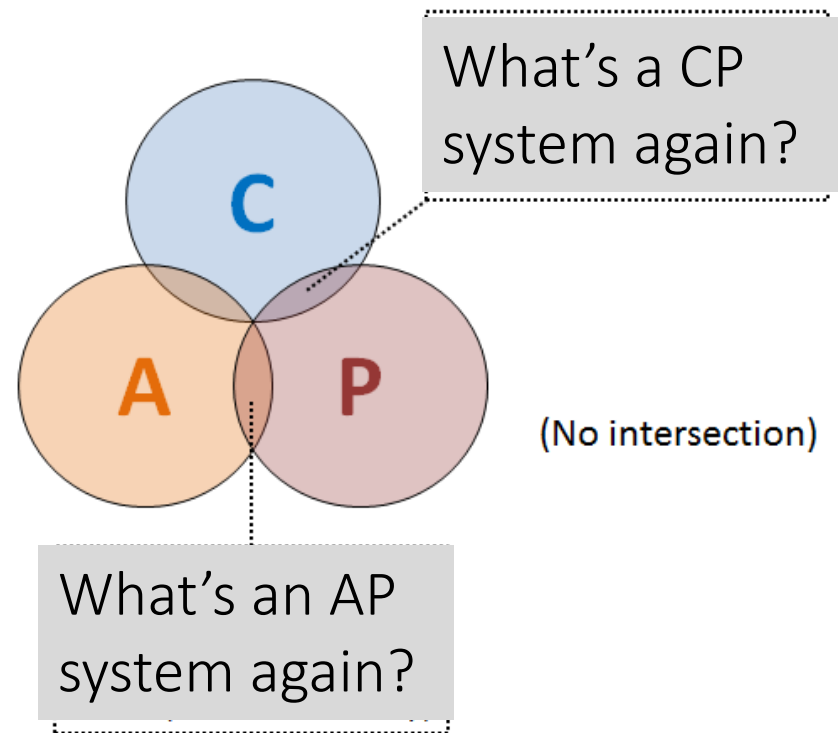
Countries	
Primary Key	Value
Afghanistan	capital:Kabul,continent:Asia,pop:31108077#2011
Albania	capital:Tirana,continent:Europe,pop:3011405#2013
...	...

Cities	
Primary Key	Value
Kabul	country:Afghanistan,pop:3476000#2013
Tirana	country:Albania,pop:3011405#2013
...	...

Amazon Dynamo: CAP

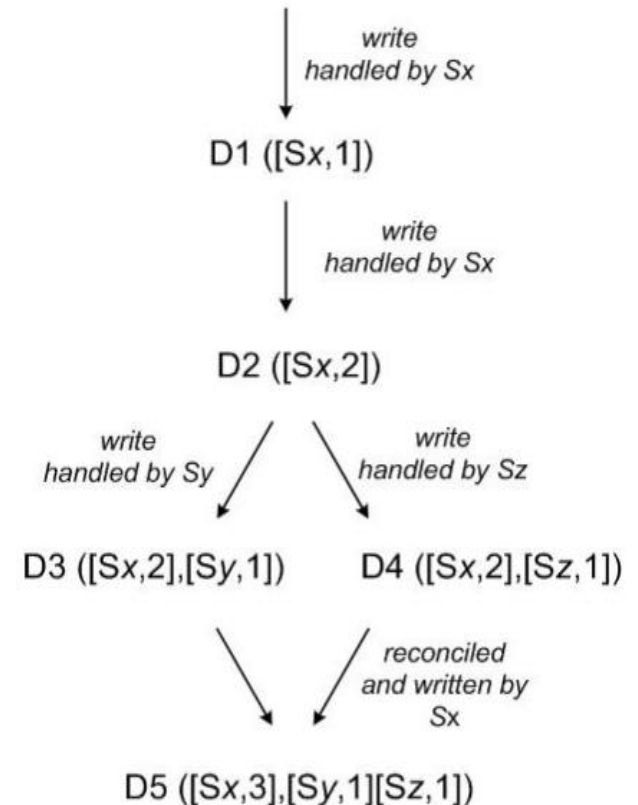
Two options for each table:

- **AP**: Eventual consistency,
High availability
- **CP**: Strong consistency,
Lower availability



Amazon Dynamo: Consistency


- **Vector Clock:**
 - A list of pairs indicating a node and operation count
 - Used to track branches of revisions




Amazon Dynamo: Consistency

- Two versions of one shopping cart:

Shopping Cart

	Price	Quantity
 WD My Passport Ultra 2TB Portable External USB 3.0 Hard Drive with Auto Backup - Red by Western Digital In Stock Eligible for FREE Shipping <input type="checkbox"/> This is a gift Learn more Delete Save for later	\$90.99 You save: \$49.00 (35%)	1
 Logitech Wireless Presenter R400 by Logitech In Stock Eligible for FREE Shipping <input type="checkbox"/> This is a gift Learn more Delete Save for later	\$44.29 You save: \$5.70 (11%)	1
Subtotal (2 items): \$135.28		
Total savings: \$54.70		

Shopping Cart

	Price	Quantity
 AKG Perception P120 Professional Studio Microphone, Silver by AKG Pro Audio Only 2 left in stock. Shipped from: Sam Ash Gift options not available. Learn more Delete Save for later	\$99.00 You save: \$30.00 (23%)	1
 Logitech Wireless Presenter R400 by Logitech In Stock Eligible for FREE Shipping <input type="checkbox"/> This is a gift Learn more Delete Save for later	\$44.29 You save: \$5.70 (11%)	1
Subtotal (2 items): \$143.29		
Total savings: \$35.70		


How best to merge multiple conflicting versions of a value
(known as reconciliation)?




Application knows best

(... and must support multiple versions being returned)

Amazon Dynamo: Consistency



Key	Value
country:Afghanistan	capital@city:Kabul,continent:Asia,pop:31108077#2011
country:Albania	capital@city:Tirana,continent:Europe,pop:3011405#2013
...	...
city:Kabul	country:Afghanistan,pop:3476000#2013
city:Tirana	country:Albania,pop:3011405#2013
...	...
user:10239	basedIn@city:Tirana,post:{103,10430,201,408}
...	...



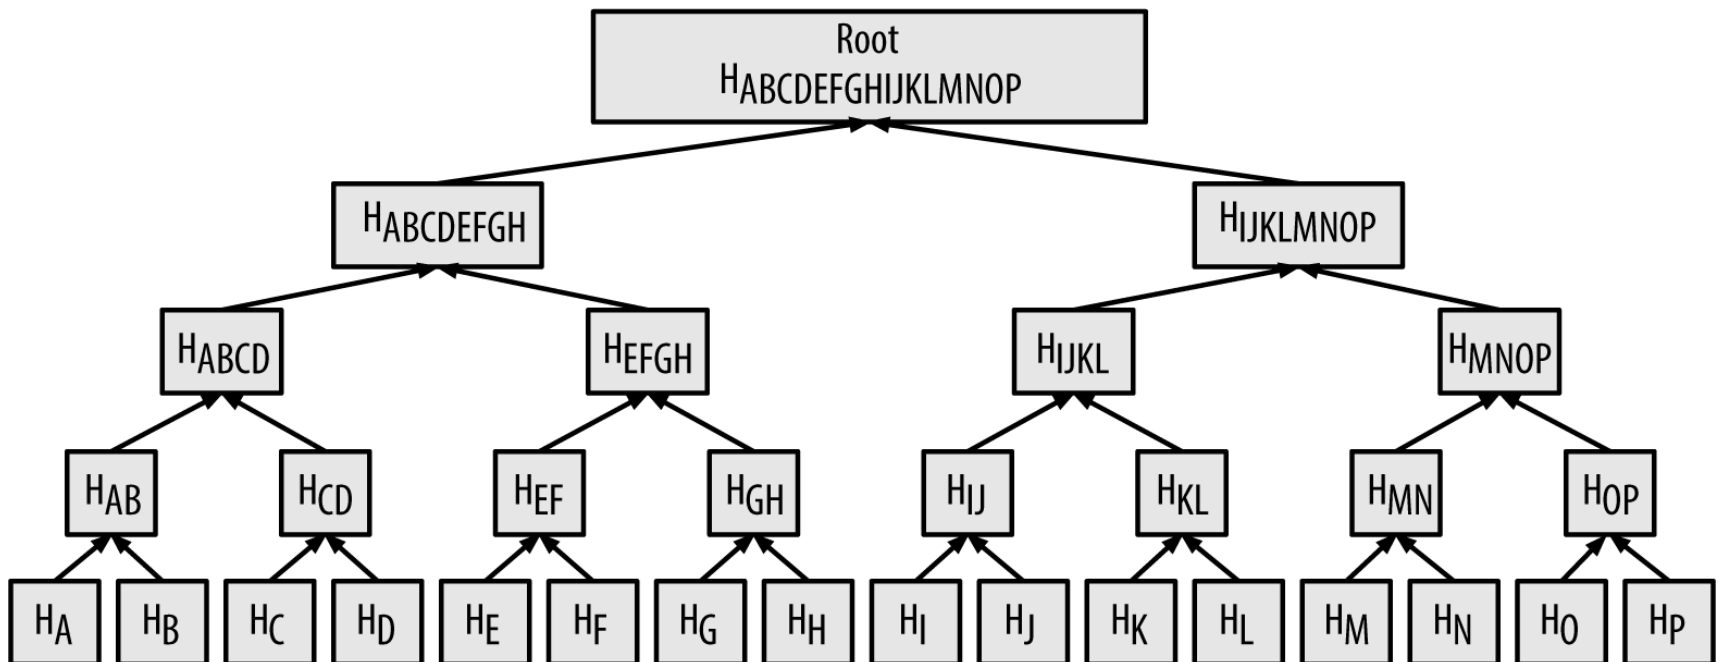
Key	Value
country:Afghanistan	capital@city:Kabul,continent:Asia,pop:31108077#2011
country:Albania	capital@city:Tirana,continent:Europe,pop:3011405#2013
...	...
city:Kabul	country:Afghanistan,pop:3476000#2013
city:Tirana	country:Albania,pop:3011405#2013
...	...
user:10239	basedIn@city:Tirana,post:{103,10430,201}
...	...

How can we efficiently verify that two copies of a block of data are the same (and find where the differences are)?



Amazon Dynamo: Merkle Trees

- **Merkle tree:** A hash tree
 - Leaf node compute hashes from data
 - Non-leaf nodes have hashes of their children
 - Find differences between two trees level-by-level



Read More ...



Dynamo: Amazon's Highly Available Key-value Store

Giuseppe DeCandia, Deniz Hastorun, Madan Jampani, Gunavardhan Kakulapati,
Avinash Lakshman, Alex Pilchin, Swaminathan Sivasubramanian, Peter Vosshall
and Werner Vogels

Amazon.com

ABSTRACT

Reliability at massive scale is one of the biggest challenges we face at Amazon.com, one of the largest e-commerce operations in the world; even the slightest outage has significant financial consequences and impacts customer trust. The Amazon.com platform, which provides services for many web sites worldwide, is implemented on top of an infrastructure of tens of thousands of servers and network components located in many datacenters


One of the lessons our organization has learned from operating Amazon's platform is that the reliability and scalability of a system is dependent on how its application state is managed. Amazon uses a highly decentralized, loosely coupled, service oriented architecture consisting of hundreds of services. In this environment there is a particular need for storage technologies that are always available. For example, customers should be able to view and add items to their shopping cart even if disks are failing, network routes are flapping, or data centers are being

OTHER KEY-VALUE STORES

Other Key–Value Stores



Other Key–Value Stores




redis



StackExchange 



Other Key-Value Stores



cassandra




Instagram
Fast beautiful photo sharing

accenture
High performance. Delivered.


Answers.com




Disney



Other Key-Value Stores



cassandra



Instagram
Fast beautiful

accenture
High performance. Delivered.

Answers.com

Evolved into a
tabular store ...




DISNEY



TABULAR / COLUMN FAMILY

NoSQL: Column Family Stores



Key–Value = a Distributed Map

Countries	
Primary Key	Value
Afghanistan	capital:Kabul,continent:Asia,pop:31108077#2011
Albania	capital:Tirana,continent:Europe,pop:3011405#2013
...	...

Tabular = Multi-dimensional Maps

Countries				
Primary Key	capital	continent	pop-value	pop-year
Afghanistan	Kabul	Asia	31108077	2011
Albania	Tirana	Europe	3011405	2013
...

Bigtable: The Original Whitepaper

MapReduce
authors

Bigtable: A Distributed Storage System for Structured Data

Fay Chang, Jeffrey Dean, Sanjay Ghemawat, Wilson C. Hsieh, Deborah A. Wallach,
Mike Burrows, Tushar Chandra, Andrew Fikes, Robert E. Gruber

{fay,jeff,sanjay,wilsonh,kerr,m3b,tushar,fikes,gruber}@google.com

Google, Inc.

Abstract

Bigtable is a distributed storage system for managing structured data that is designed to scale to a very large size: petabytes of data across thousands of commodity servers. Many projects at Google store data in Bigtable, including web indexing, Google Earth, and Google Finance. These applications place very different demands on Bigtable, both in terms of data size (from URLs to web pages to satellite imagery) and latency requirements (from backend bulk processing to real-time data service). Despite these varied demands, Bigtable has successfully provided a flexible, high-performance solution for all of these Google products. In this paper we describe the simple data model provided by Bigtable, which gives clients dynamic control over data layout and format, and we describe the design and implementation of Bigtable.

achieved scalability and high performance, but Bigtable provides a different interface than such systems. Bigtable does not support a full relational data model; instead, it provides clients with a simple data model that supports dynamic control over data layout and format, and allows clients to reason about the locality properties of the data represented in the underlying storage. Data is indexed using row and column names that can be arbitrary strings. Bigtable also treats data as uninterpreted strings, although clients often serialise various forms of structured and semi-structured data into these strings. Clients can control the locality of their data through careful choices in their schemas. Finally, Bigtable schema parameters let clients dynamically control whether to serve data out of memory or from disk.

Section 2 describes the data model in more detail, and Section 3 provides an overview of the client API. Sec-

Bigtable used for ...



Google Analytics



+Amit Web Images Videos Maps News Shopping Gmail More

Amit Singhal 2 Share

Google chikoo

Search 50 personal results and 419,000 other results (0.61 seconds)

Manilkara zapota - Wikipedia, the free encyclopedia
en.wikipedia.org/wiki/Manilkara_zapota
Sapodilla is known as **chikoo** ("चिकू," or chiku, "चोफ़,") in India and Pakistan and sapota in some parts of India (Tamil Nadu, Kerala, Karnataka, Andhra ...
Description - Other names - See also - References
You've visited this page 3 times. Last visit 12/4/11

Images for chikoo - Report images

You	Shilpa Singhal	You	Shilpa Singhal

Chikoo - a simple file organizer for the Mac
codingturtle.com/chikoo/
Chikoo, a simple file organizer for the Mac. Download (30-day trial). Version 0.9.1 ...



orkut by Google

Bigtable: Sorted Keys

Primary Key	capital		pop-value		pop-year	
Asia:Afghanistan	t ₁	Kabul	t ₁	31143292	t ₁	2009
			t ₂	31120978		
			t ₄	31108077	t ₄	2011
Asia:Azerbaijan
...
Europe:Albania	t ₁	Tirana	t ₁	2912380	t ₁	2010
			t ₃	3011405	t ₃	2013
Europe:Andorra
...

S
O
R
T
E
D

Benefits of sorted vs. hashed keys?



Range queries and ...

Bigtable: Tablets

	Primary Key	capital		pop-value		pop-year	
ASIA	Asia:Afghanistan	t ₁	Kabul	t ₁	31143292	t ₁	2009
				t ₂	31120978		
				t ₄	31108077	t ₄	2011
	Asia:Azerbaijan

EUROPE	Europe:Albania	t ₁	Tirana	t ₁	2912380	t ₁	2010
				t ₃	3011405	t ₃	2013
	Europe:Andorra

Benefits of sorted vs. hashed keys?



Range queries and ...

... locality of processing

A real-world example of locality/sorting



Primary Key	language		title		links	
com.imdb	t ₁	en	t ₁	IMDb Home	t ₁	...
			t ₂	IMDB - Movies		
			t ₄	IMDb	t ₄	...
com.imdb/title/tt2724064/	t ₁	en	t ₂	Sharknado	t ₂	...
com.imdb/title/tt3062074/	t ₁	en	t ₂	Sharknado II	t ₂	...
...
org.wikipedia	t ₁	multi	t ₁	Wikipedia	t ₁	...
			t ₃	Wikipedia Home	t ₃	...
org.wikipedia.ace	t ₁	ace	t ₁	Wikipèdia bahsa Acèh
...




Bigtable: Distribution

Asia:Afghanistan	t ₁	Kabul	t ₁	31143292	t ₁	2009
			t ₂	31120978		
			t ₄	31108077	t ₄	2011
Asia:Azerbaijan
...

Europe:Albania	t ₁	Tirana	t ₁	2912380	t ₁	2010
			t ₃	3011405	t ₃	2013
Europe:Andorra
...

Split by tablet




Horizontal range partitioning

Bigtable: Column Families

Primary Key	pol:capital		demo:pop-value		demo:pop-year	
Asia:Afghanistan	t ₁	Kabul	t ₁	31143292	t ₁	2009
			t ₂	31120978		
			t ₄	31108077	t ₄	2011
Asia:Azerbaijan
...
Europe:Albania	t ₁	Tirana	t ₁	2912380	t ₁	2010
			t ₃	3011405	t ₃	2013
Europe:Andorra
...

- Group logically similar columns together
 - Accessed efficiently together
 - Access-control and storage: column family level
 - If of same type, can be compressed

Read More ...



Bigtable: A Distributed Storage System for Structured Data

Fay Chang, Jeffrey Dean, Sanjay Ghemawat, Wilson C. Hsieh, Deborah A. Wallach
Mike Burrows, Tushar Chandra, Andrew Fikes, Robert E. Gruber

{fay,jeff,sanjay,wilsonh,kerr,m3b,tushar,fikes,gruber}@google.com

Google, Inc.

Abstract

Bigtable is a distributed storage system for managing structured data that is designed to scale to a very large size: petabytes of data across thousands of commodity servers. Many projects at Google store data in Bigtable, including web indexing, Google Earth, and Google Finance. These applications place very different demands on Bigtable, both in terms of data size (from URLs to web pages to satellite imagery) and latency requirements (from backend bulk processing to real-time data serving). Despite these varied demands, Bigtable has successfully provided a flexible, high-performance solution for all of these Google products. In this paper we describe the simple data model provided by Bigtable, which gives clients dynamic control over data layout and format, and we describe the design and implementation of Bigtable.


achieved scalability and high performance, but Bigtable provides a different interface than such systems. Bigtable does not support a full relational data model; instead, it provides clients with a simple data model that supports dynamic control over data layout and format, and allows clients to reason about the locality properties of the data represented in the underlying storage. Data is indexed using row and column names that can be arbitrary strings. Bigtable also treats data as uninterpreted strings, although clients often serialize various forms of structured and semi-structured data into these strings. Clients can control the locality of their data through careful choices in their schemas. Finally, Bigtable schema parameters let clients dynamically control whether to serve data out of memory or from disk.

Section 2 describes the data model in more detail, and

Tabular Store: Apache HBase



Tabular Store: Cassandra



cassandra



Questions?