


**Microsoft**



Microsoft® Research

# Faculty Summit 2012

Riviera Maya, Mexico | May 23-25 | In partnership with CONACYT



# Big Data and the Cloud Phenomenon

Pedro Celis  
Distinguished Engineer,  
Microsoft, United States

Date  
Thursday, May 24, 2012

# How much information is there?

- Almost everything is recorded digitally.
- Most bytes are never seen by humans.
- Data summarization, trend detection, anomaly detection are key technologies

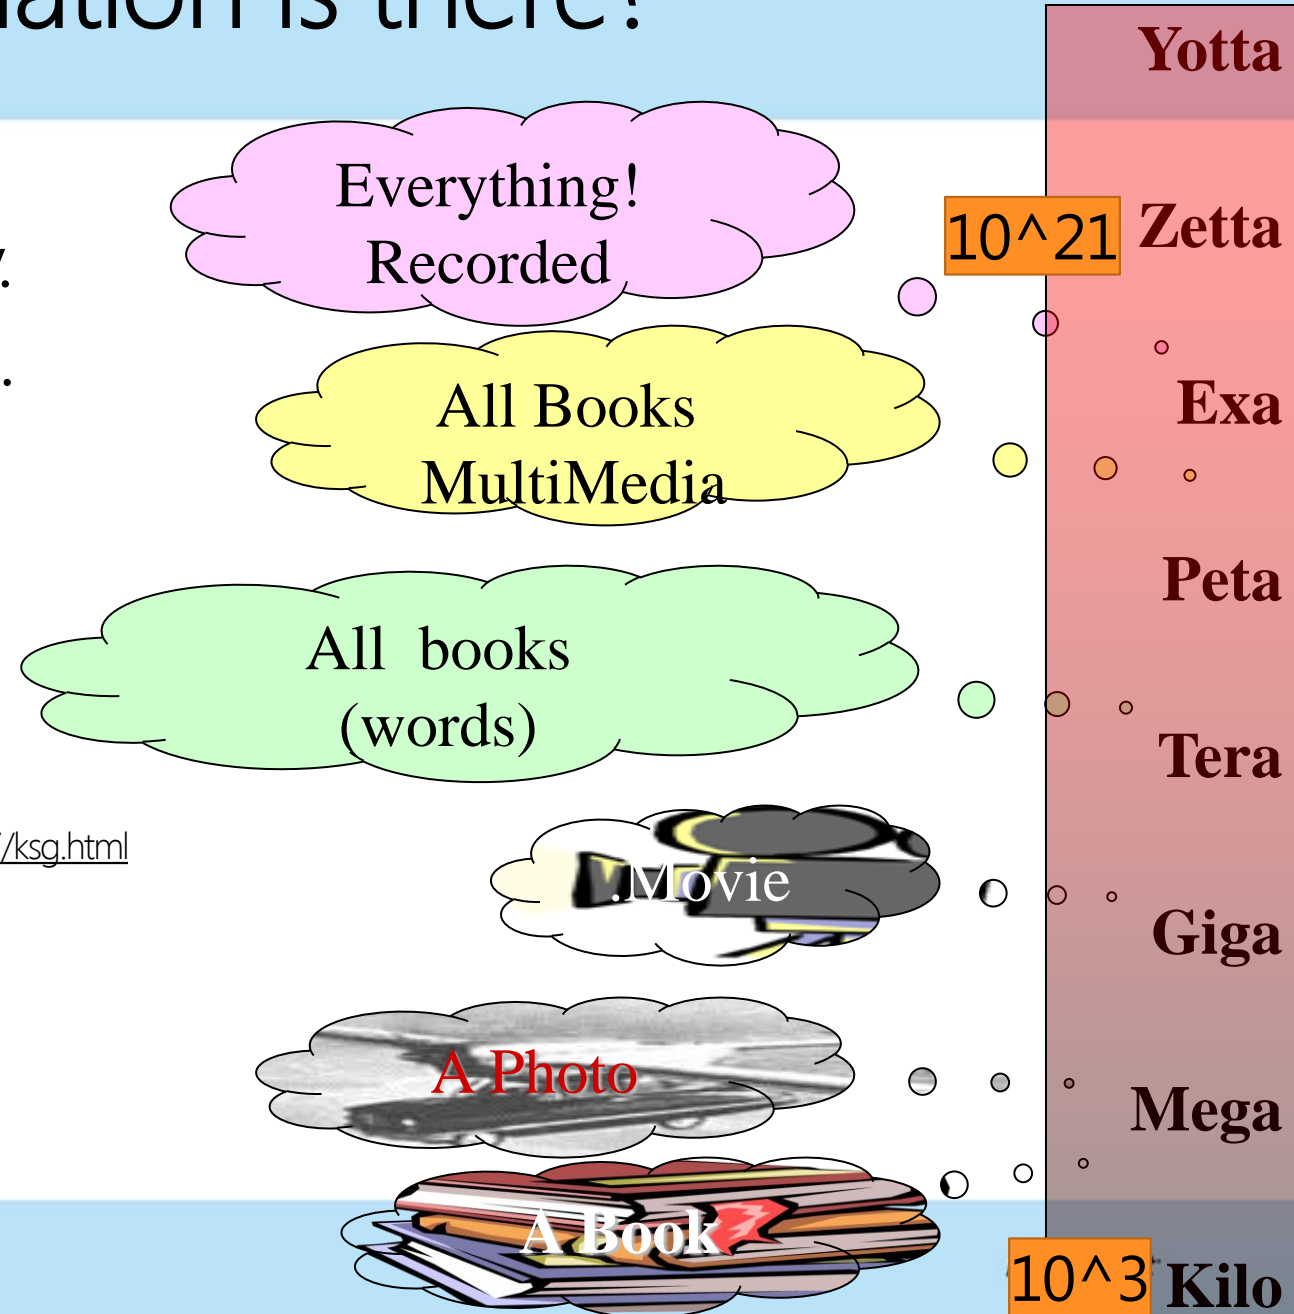
See Mike Lesk:

*How much information is there:* <http://www.lesk.com/mlesk/ksg97/ksg.html>

See Lyman & Varian:

*How much information*

<http://www.sims.berkeley.edu/research/projects/how-much-info/>





# Patterns

- AMBIENT DATA -> SHOEBOX PATTERN



# BIG DATA IS A DISRUPTIVE SHIFT

- DRIVES NEW OPPORTUNITIES AND NEW DESIGN PATTERNS
- AMBIENT DATA -> SHOEBOX PATTERN



# TECHNOLOGIES

“A SPREADSHEET WITH 300 BILLION ROWS?”

ONE YEAR OF ADSENSE CLICKSTREAM DATA:

COOKED DOWN WITH MAP-REDUCE, 50X:20 T

VERTIPAQ SERVER SCALE-OUT, 33X: 600 GB

VERTIPAQ IN-MEMORY STORE, 40X: 15 GB

PROCESS IN EXCEL: 300 B ROWS

SINGLE-DIGIT  
HOURS

LESS, + QUICK  
RELOAD

0.3 S OPERATIONS

GARTNER: “WAIT, WHAT? WHAT DID WE JUST SEE?”



# THE COST OF 1 PB

COST OF DATA ENTRY FOR 1 PB: COST OF SERVERS FOR 1 PB:

MANUAL DATA ENTRY: \$1/kB  
1 PB = \$1 TRILLION

33 COMMODITY SERVERS @ \$3,000  
1 PB = \$100,000

$\$1 \text{ TRILLION} : \$100,000 = 10^7$



BING

FOR EACH SEARCH THEY SEND OUT DATA  
THEY ALSO COLLECT DATA ABOUT THE SEARCH AND THE USER  
THEY SAVE MORE DATA THAN THEY SEND BACK

BING IS AN INFORMATION MACHINE,  
CREATING DATA FOR THEIR CUSTOMERS –  
THE ADVERTISERS  
CREATING DATA FOR THEIR USERS –  
BETTER RELEVANCE AND RANKING ANSWERS  
PETABYTES PER DAY





# BIG DATA IS A DISRUPTIVE SHIFT

- DRIVES NEW OPPORTUNITIES AND NEW DESIGN PATTERNS
- AMBIENT DATA -> SHOEBOX PATTERN

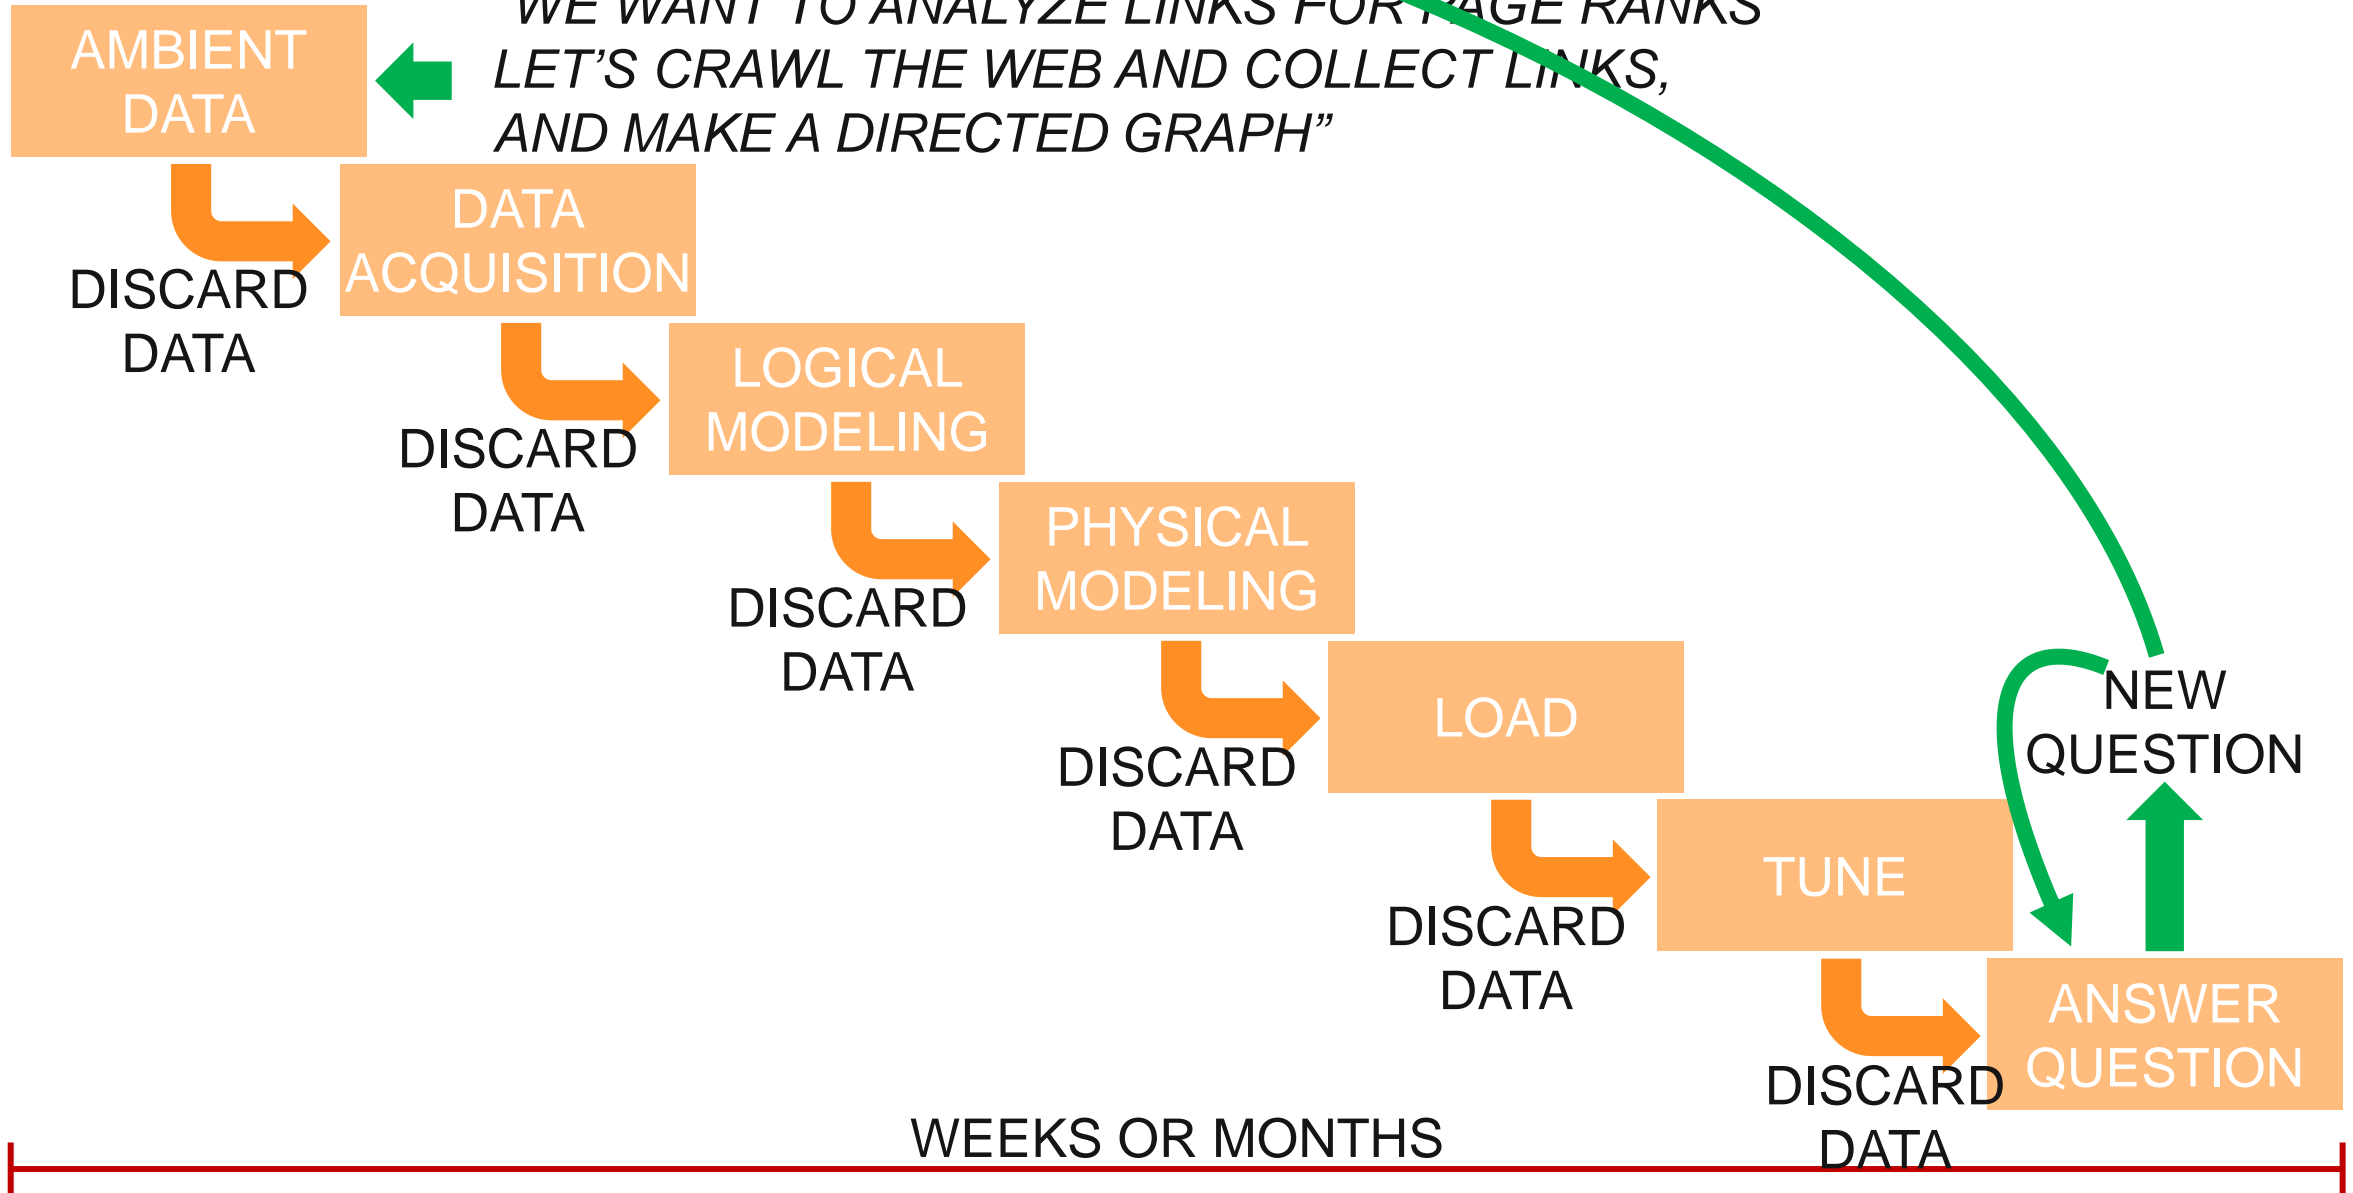


# EXAMPLE: THE DIGITAL SHOEBOX PATTERN

“RETAIN ALL  
POTENTIALLY VALUABLE  
AMBIENT DATA  
FOR SUBSEQUENT ANALYSIS”

# CLASSICAL DATA WAREHOUSE PATTERN

*"WE WANT TO ANALYZE LINKS FOR PAGE RANKS  
LET'S CRAWL THE WEB AND COLLECT LINKS,  
AND MAKE A DIRECTED GRAPH"*



# CLASSICAL DATA WAREHOUSE PATTERN

AMBIENT DATA

DATA ACQUISITION

LOGICAL MODELING

PHYSICAL MODELING

LOAD


TUNE

ANSWER QUESTION

# A CULTURE OF SCARCITY

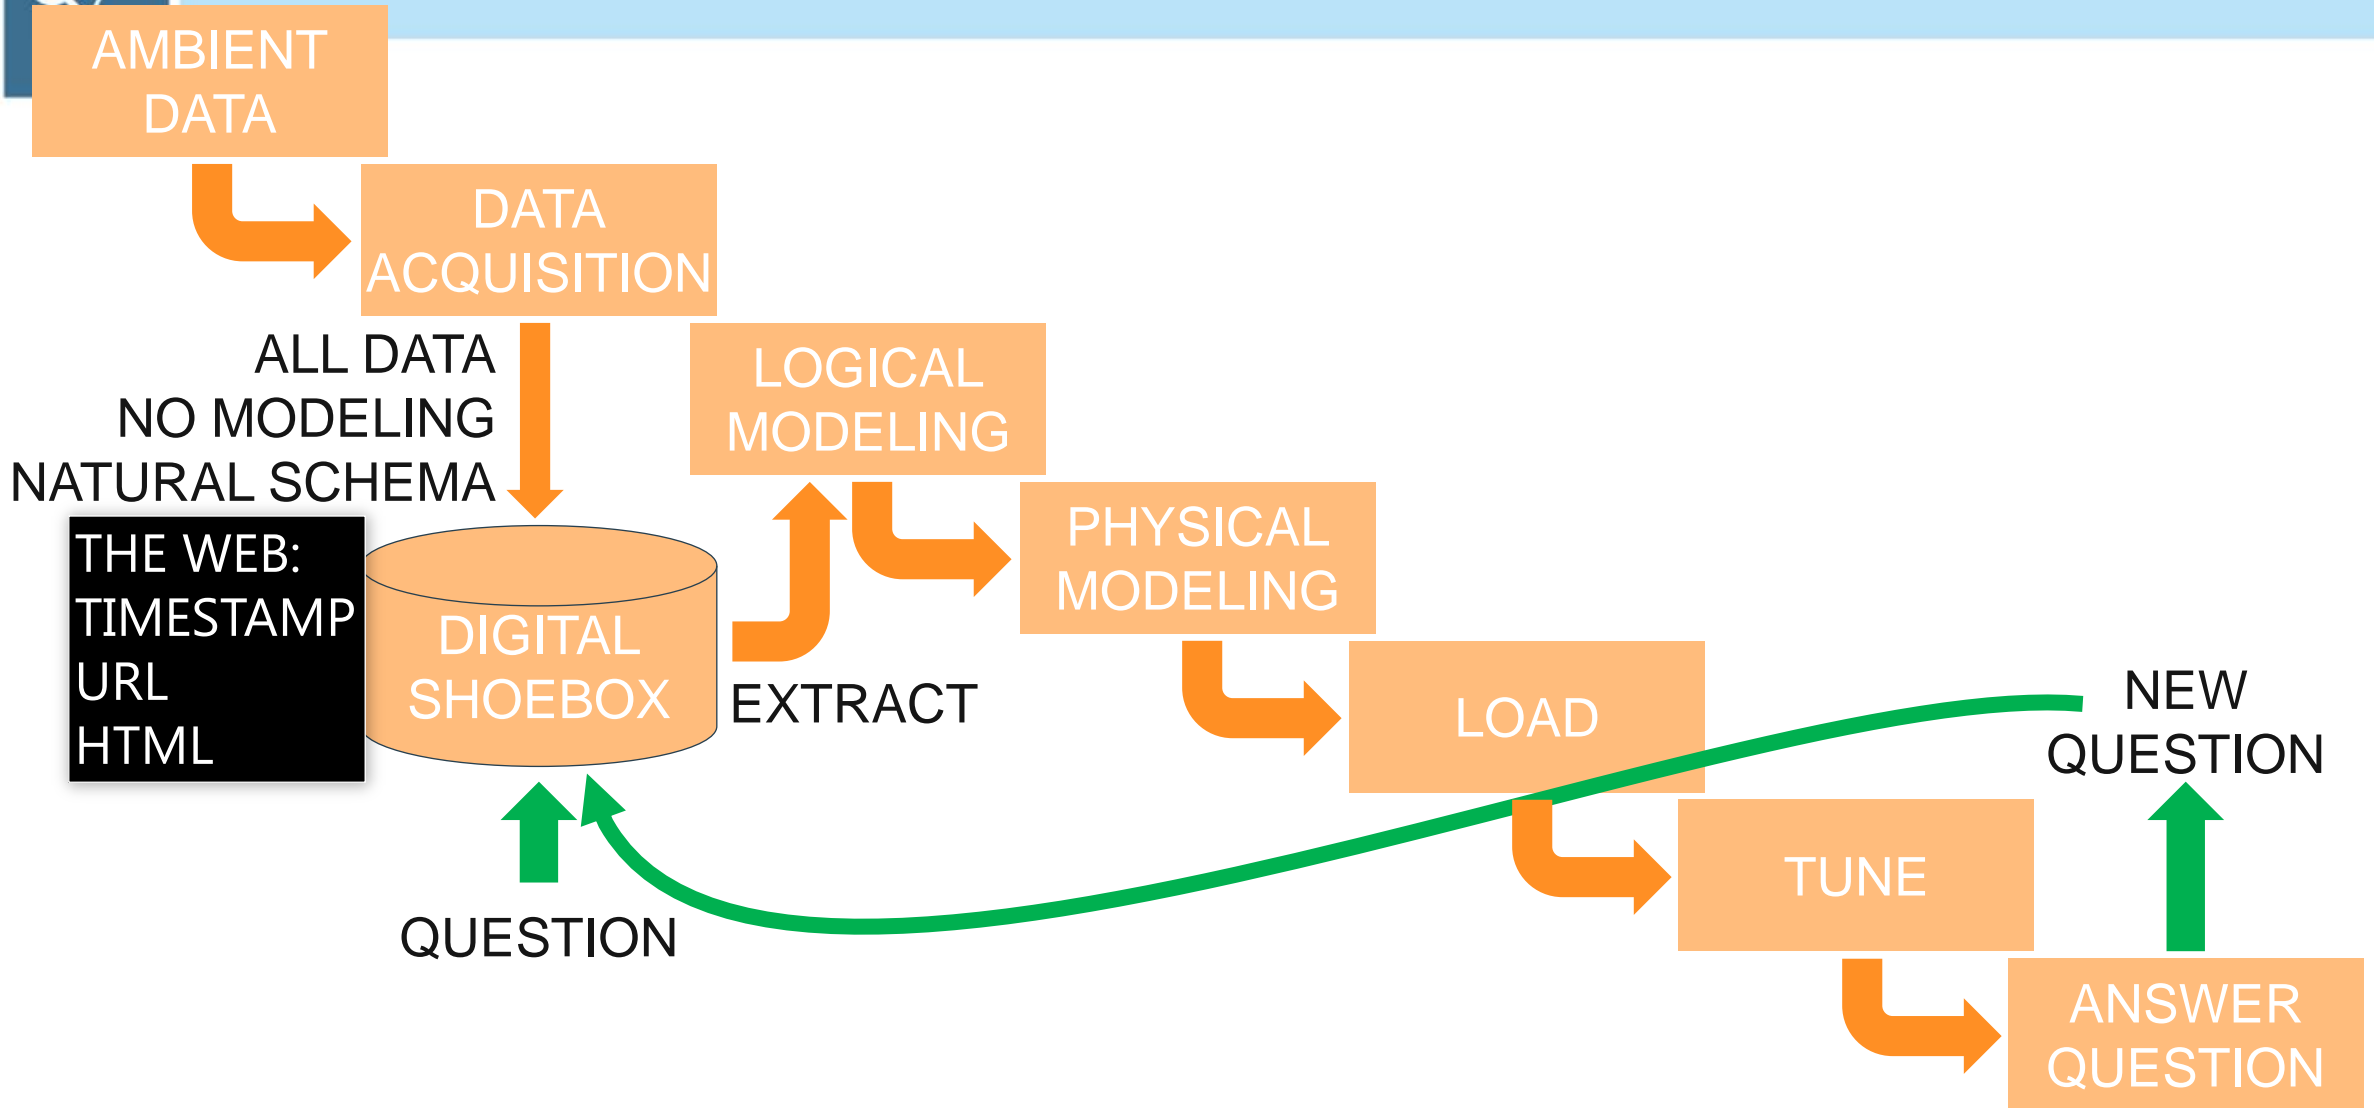
WEEKS OR MONTHS

DATA



# A CULTURE OF ABUNDANCE

# DIGITAL SHOEBOX PATTERN





# DIGITAL SHOEBOX PATTERN

WHEN IS THIS PATTERN APPROPRIATE?

$$\left( \frac{\textit{perceivedLatentValue}}{\textit{time}} \right) > \left( \frac{\textit{cost(dataAquisition)} + \textit{cost(dataStorage)}}{\textit{time}} \right)$$



# BIG DATA IS A DISRUPTIVE SHIFT

- DRIVES NEW OPPORTUNITIES AND NEW DESIGN PATTERNS
- AMBIENT DATA -> SHOEBOX PATTERN
- CONSTRUCTING DATA FROM DATA -> CORPUS PATTERN

## WHEN IS THIS APPROPRIATE?


1. PERF IS CRITICAL
2. STORAGE COST IS OK
3. STALE DATA IS OK



# CORPUS PATTERN

PROJECTION  
SCHEMA  
PARTITIONING  
INDEXES

AMBIENT  
DATA



CORPUS 1

SAME  
PARTITIONING

CORPUS 2'

CORPUS 2

PROJECTION  
SCHEMA  
PARTITIONING  
INDEXES

ETL

SERVICE 1

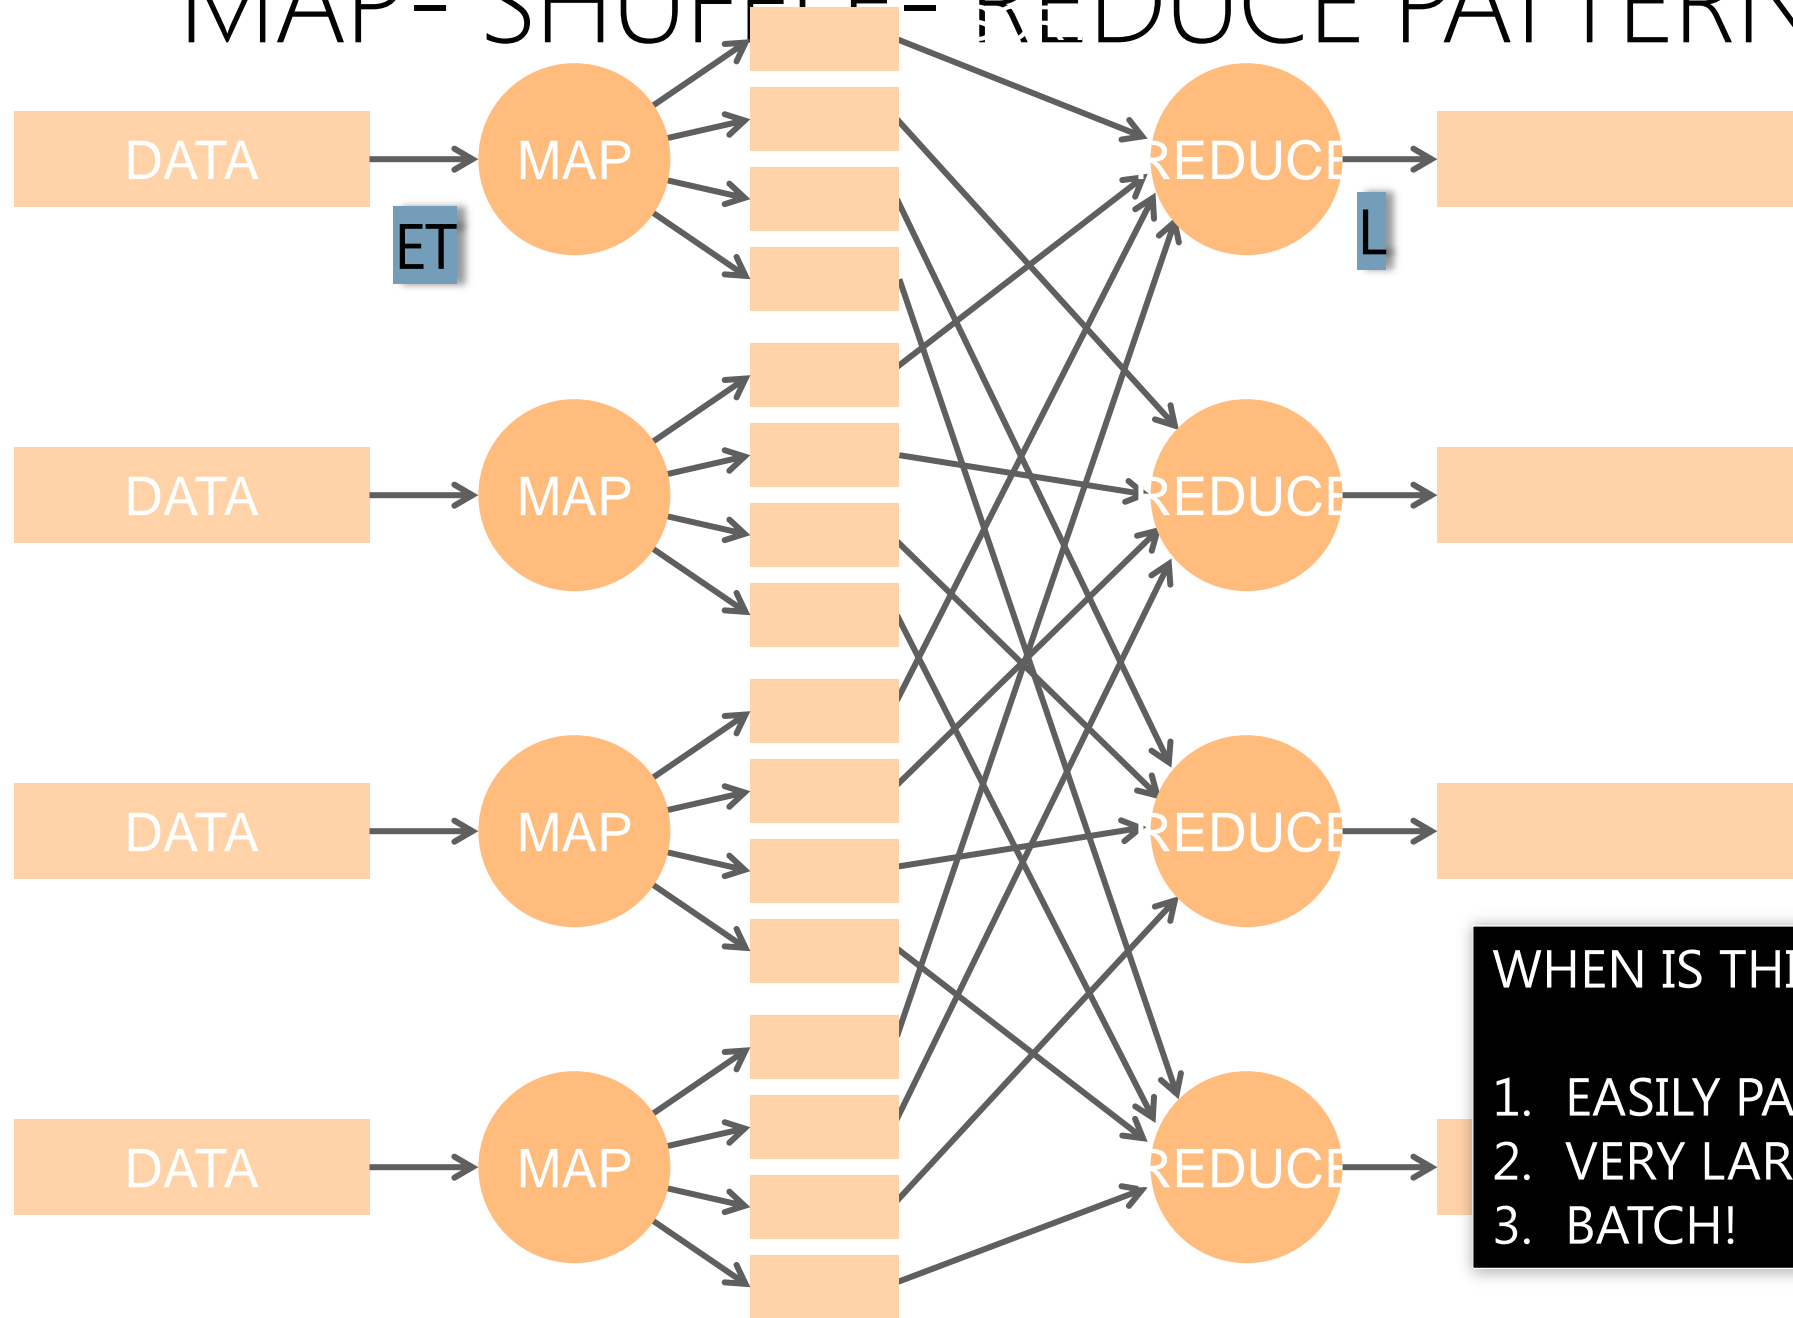
SERVICE 3

SERVICE 2

“WE’LL OPTIMIZE THE DATABASE FOR THIS DUAL USE”

- WHEN IS THIS APPROPRIATE?
- 1. PERF IS CRITICAL
  - 2. STORAGE COST IS OK
  - 3. STALE DATA IS OK

# MAP-SHUFFLE-REDUCE PATTERN



WHEN IS THIS APPROPRIATE?

1. EASILY PARALLELIZABLE
2. VERY LARGE SCALE
3. BATCH!




# BIG DATA IS A DISRUPTIVE SHIFT

- DRIVES NEW OPPORTUNITIES AND NEW DESIGN PATTERNS
- AMBIENT DATA
  - > SHOEBOX PATTERN
- CONSTRUCTING DATA FROM DATA
  - > CORPUS PATTERN
- MANY COPIES OF THE SAME DATA
  - > REPLICATION PATTERN



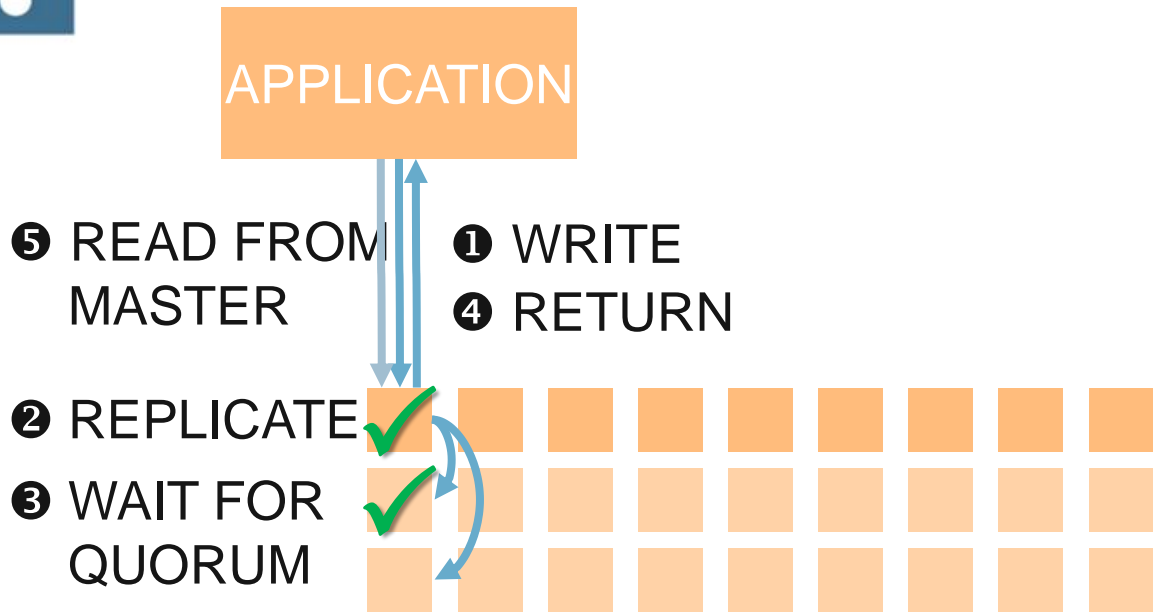
# SQL AZURE REPLICATION PATTERN

REPLICAS  
FOR  
AVAILABILITY



1. NO TRANSACTIONS ACROSS PARTITIONS
2. AUTOMATIC RECOVERY ON NODE FAILUR

# SQL AZURE REPLICATION PATTERN



- 1. NO TRANSACTIONS ACROSS PARTITIONS
- 2. AUTOMATIC RECOVERY ON NODE FAILURE
- 3. ACID TRANSACTIONS WITHIN A PARTITION
- 4. CONSISTENCY
- 5. READ PERFORMANCE LIMITED BY MASTER

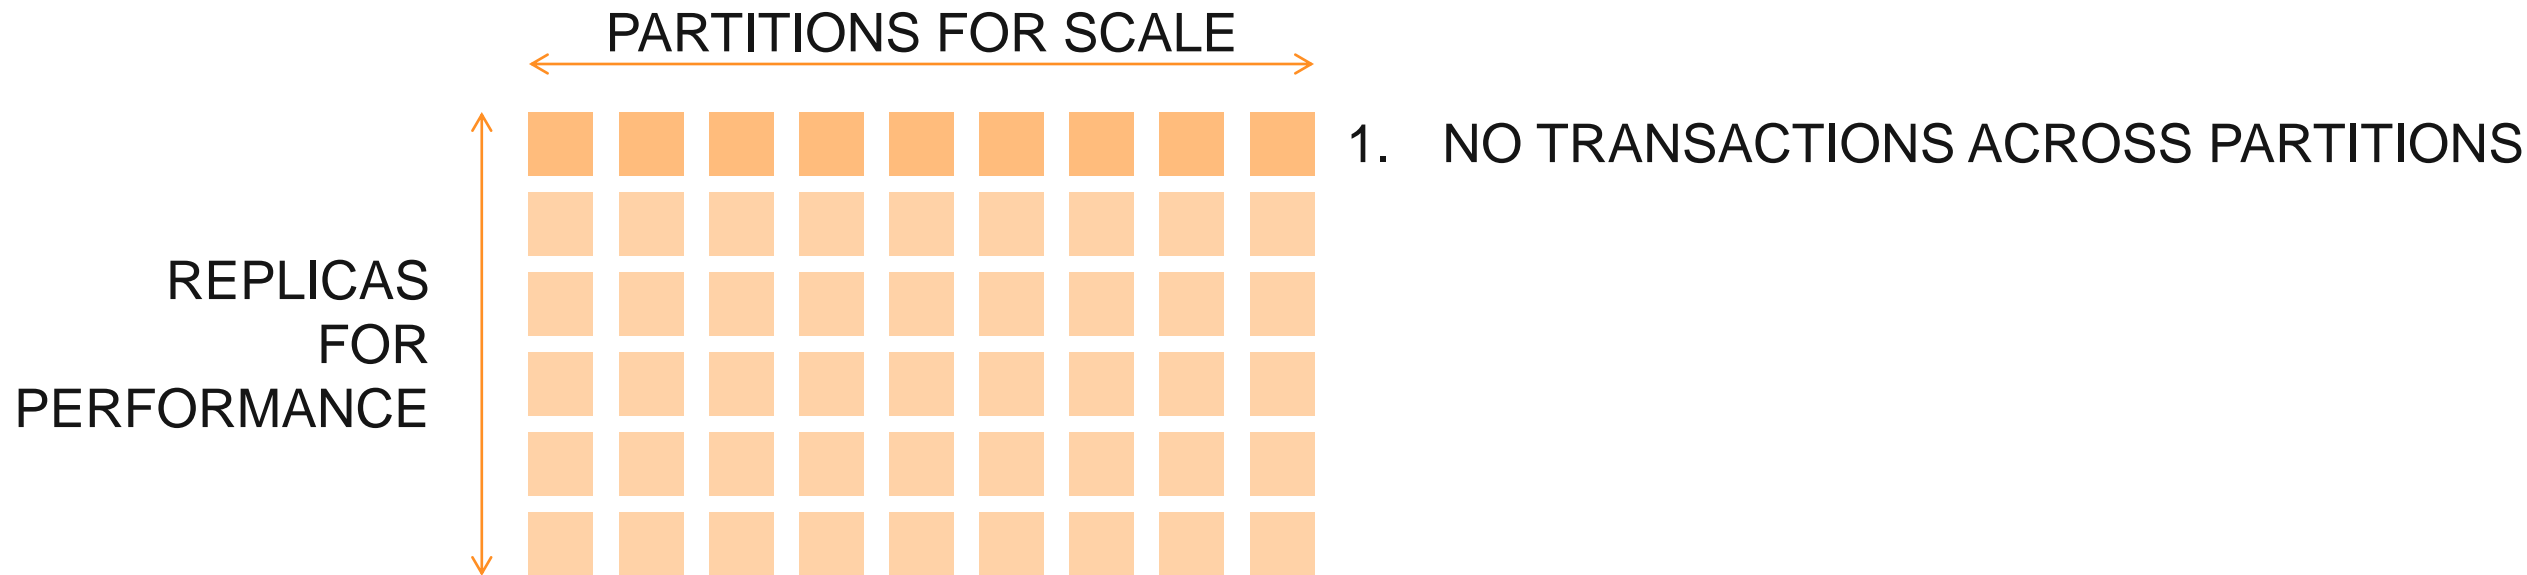
**WHEN IS THIS APPROPRIATE?**

- 1. APP NEEDS CONSISTENCY
- 2. READ AND WRITE LOADS SIMILAR

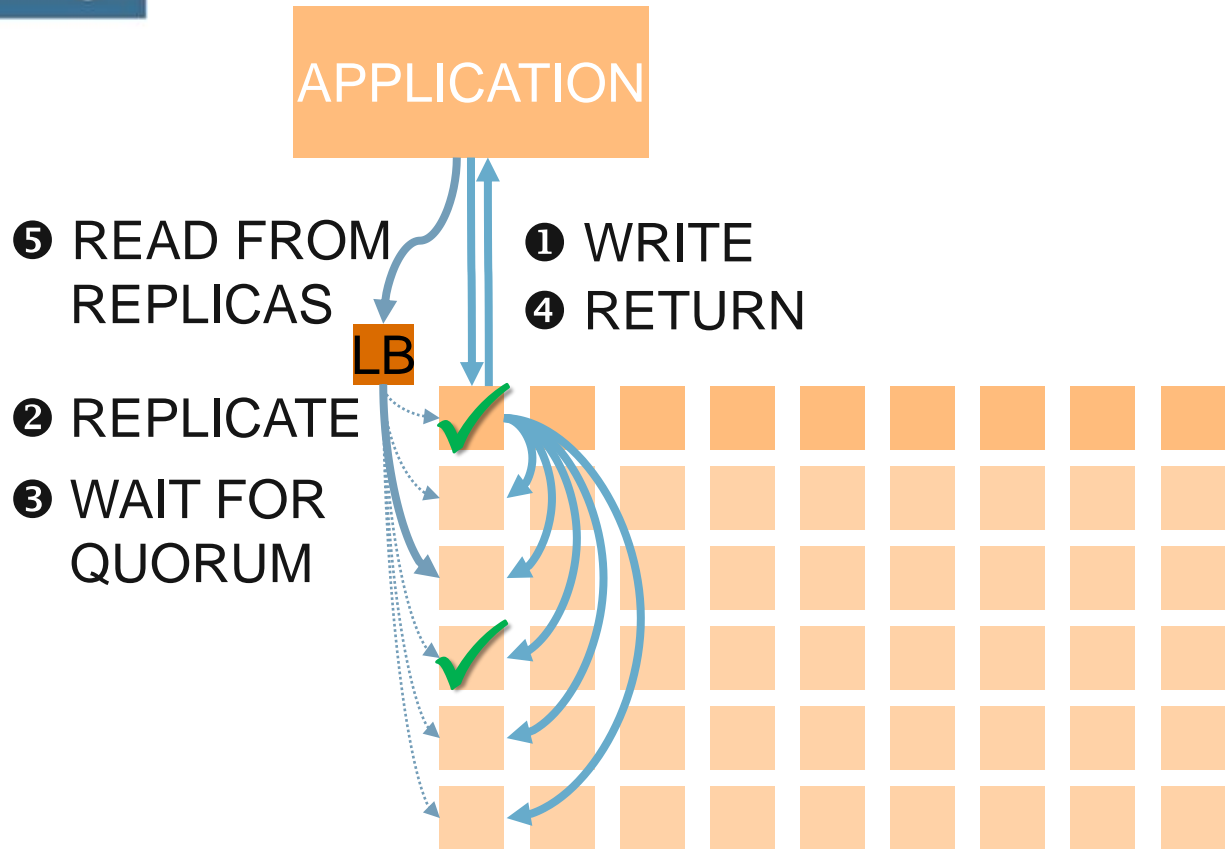
Consistency.  
The applications sees all the effect of its transactions as they occurred in chronological order



# BING REPLICATION PATTERN



# BING REPLICATION PATTERN



SQL Always ON (HADRON) with Readable Secondaries


1. NO TRANSACTIONS ACROSS PARTITIONS
2. ATOMIC TRANSACTIONS WITHIN A PARTITION
3. NO READ CONSISTENCY ACROSS PARTITIONS
4. EVENTUAL CONSISTENCY

WHEN IS THIS APPROPRIATE?

1. READ FOCUSED LOAD
2. PERF IS CRITICAL
3. WEAK CONSISTENCY OK



# BING REPLICATION PATTERN



ALLOWS TUNING READ PERF  
BY ALLOCATING RESOURCES  
TO HOT PARTITIONS





# CHOOSE DESIGN PATTERN

- READ/WRITE LOADS
- SCALE REQUIREMENTS
- PERF REQUIREMENTS
- TOLERANCE FOR WEAK CONSISTENCY
- TOLERANCE FOR STALENESS
- BATCH OR LIVE
- AD HOC OR PRODUCTION



“WE HAVE BEEN SUCCESSFUL AS A  
DATABASE COMPANY FOR SO LONG,  
PEOPLE DON’T RECOGNIZE THIS IS A NEW ERA.”



WE HAVE THE TECHNOLOGY  
WE HAVE THE DATA

DO WE HAVE THE IMAGINATION?

***Microsoft***<sup>®</sup>

© 2012 Microsoft Corporation. All rights reserved. Microsoft, Windows, and other product names are or may be registered trademarks and/or trademarks in the U.S. and/or other countries.

The information herein is for informational purposes only and represents the current view of Microsoft Corporation as of the date of this presentation. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information provided after the date of this presentation. MICROSOFT MAKES NO WARRANTIES, EXPRESS, IMPLIED OR STATUTORY, AS TO THE INFORMATION IN THIS PRESENTATION.