Databases in a Wireless World

David Yach

Senior VP Software

Research In Motion Limited

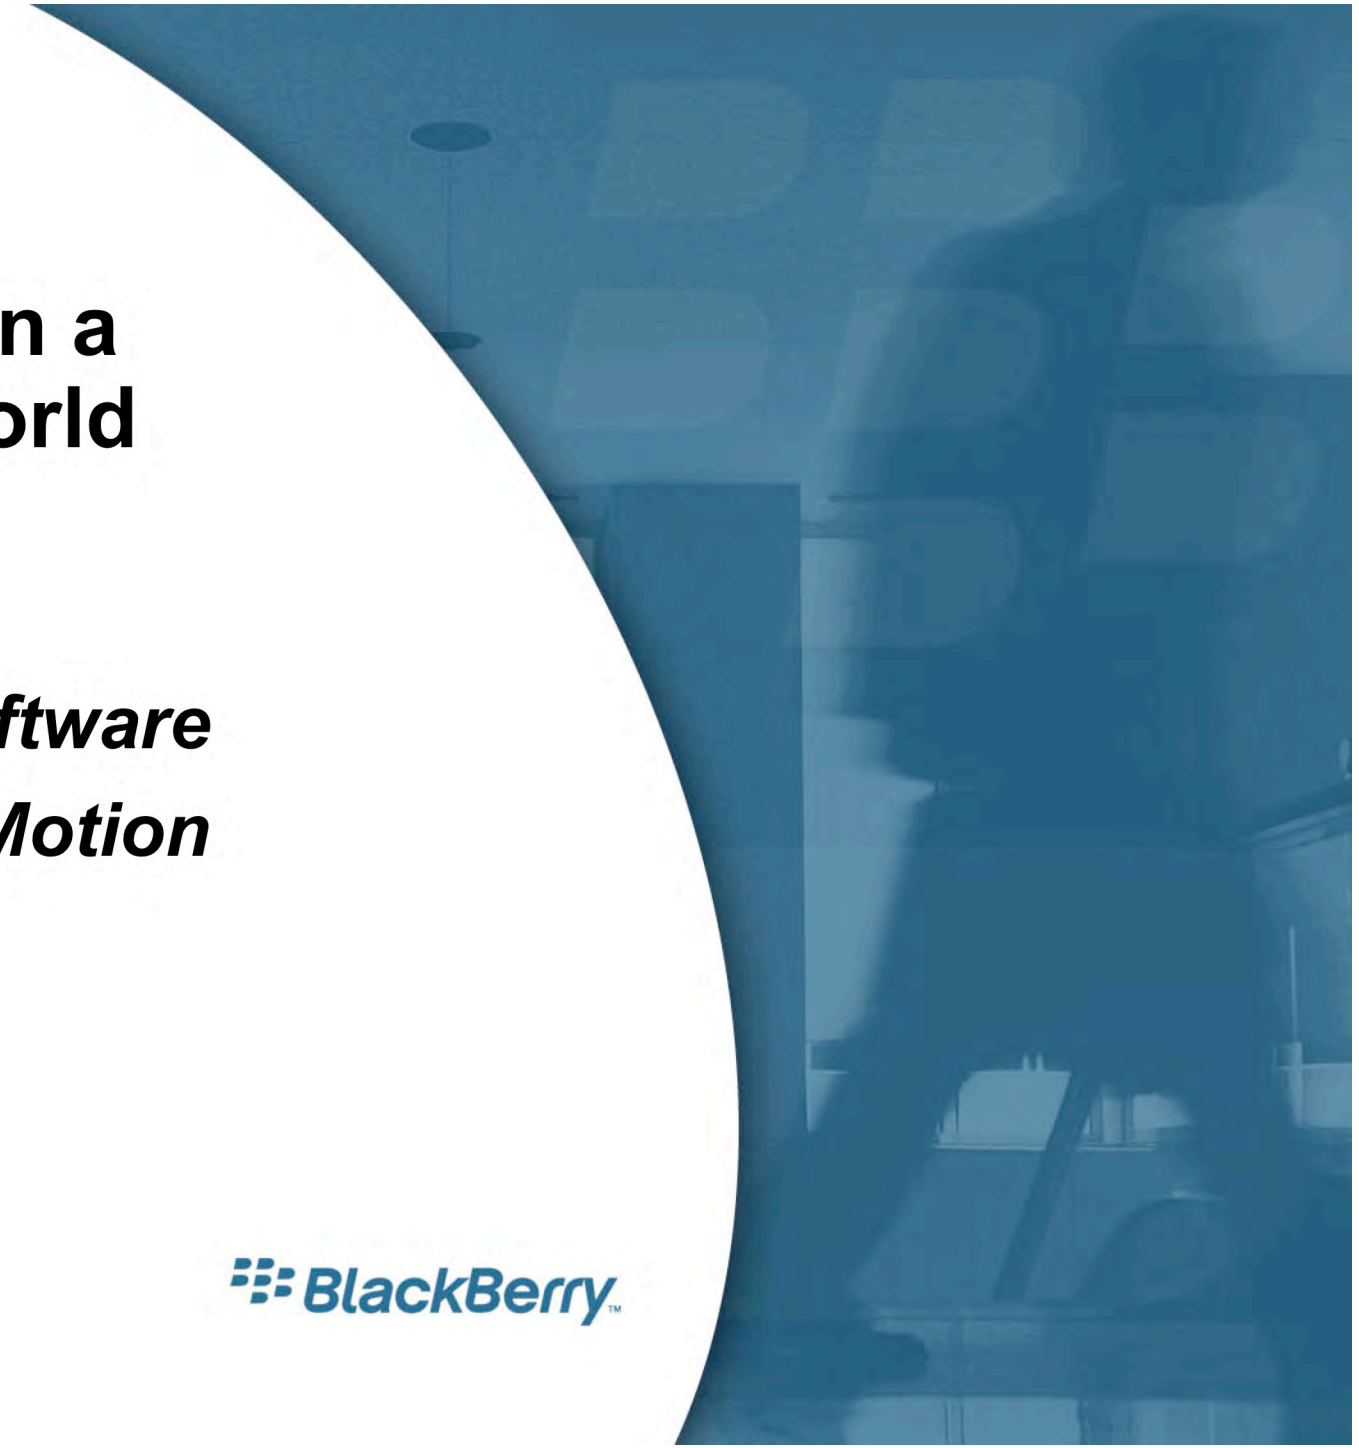

Agenda

- Problem Statement
- Wireless Background
- Alternatives and Analysis
- Going Forward

BlackBerry

Problem Statement

How should an application access a corporate database from a Wireless Device

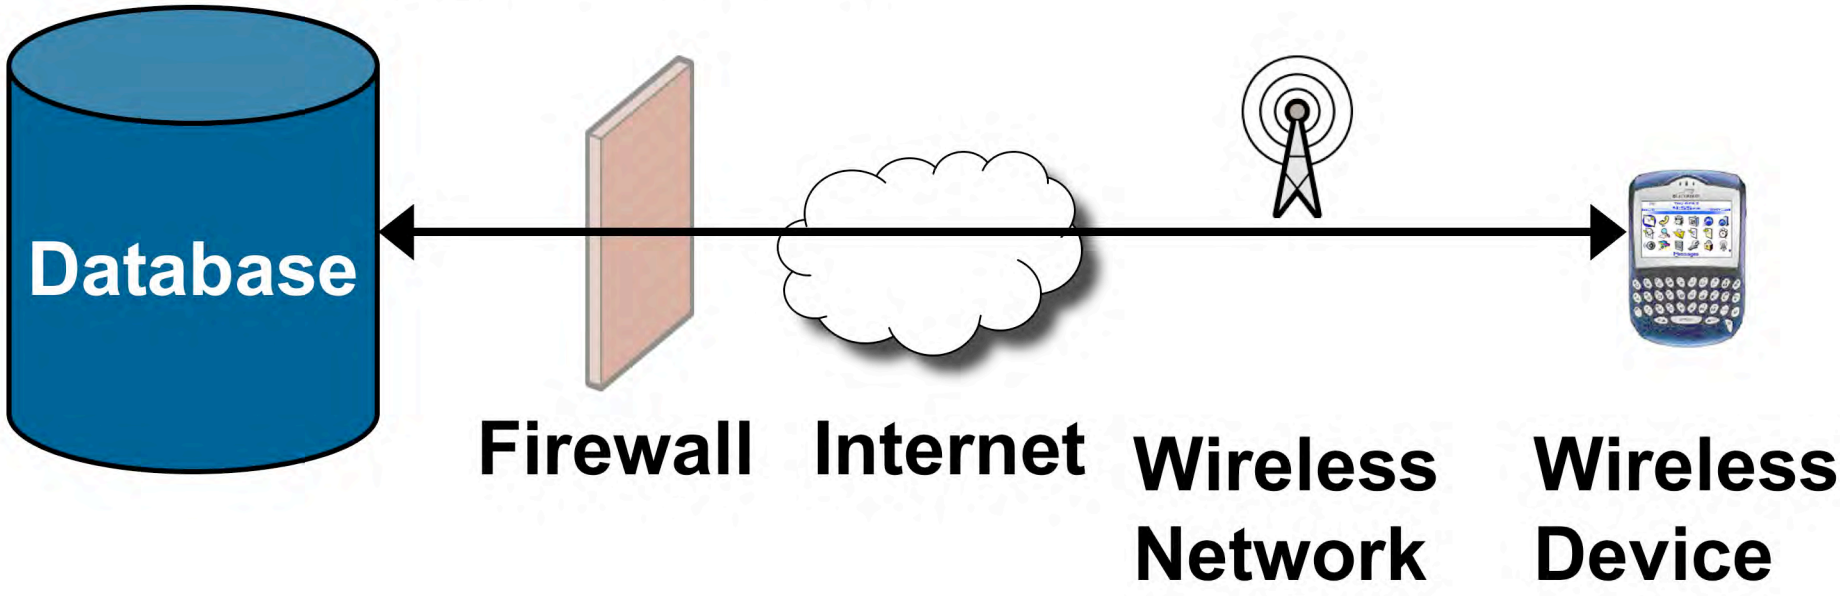

BlackBerry

Problem Statement

How should an application access a corporate database from a Wireless Device*s*

^bunch of

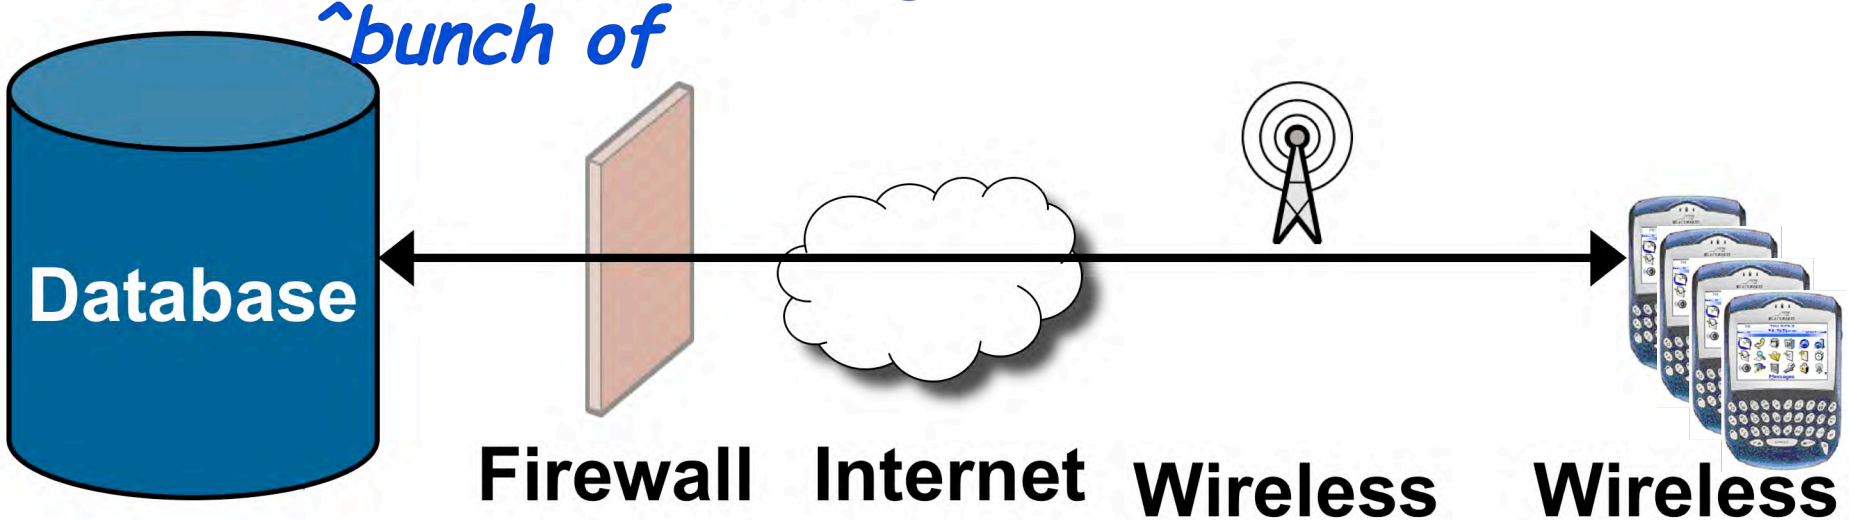

BlackBerry

Problem Statement

How should an application access a corporate database from a Wireless Device*s*

whole^bunch of

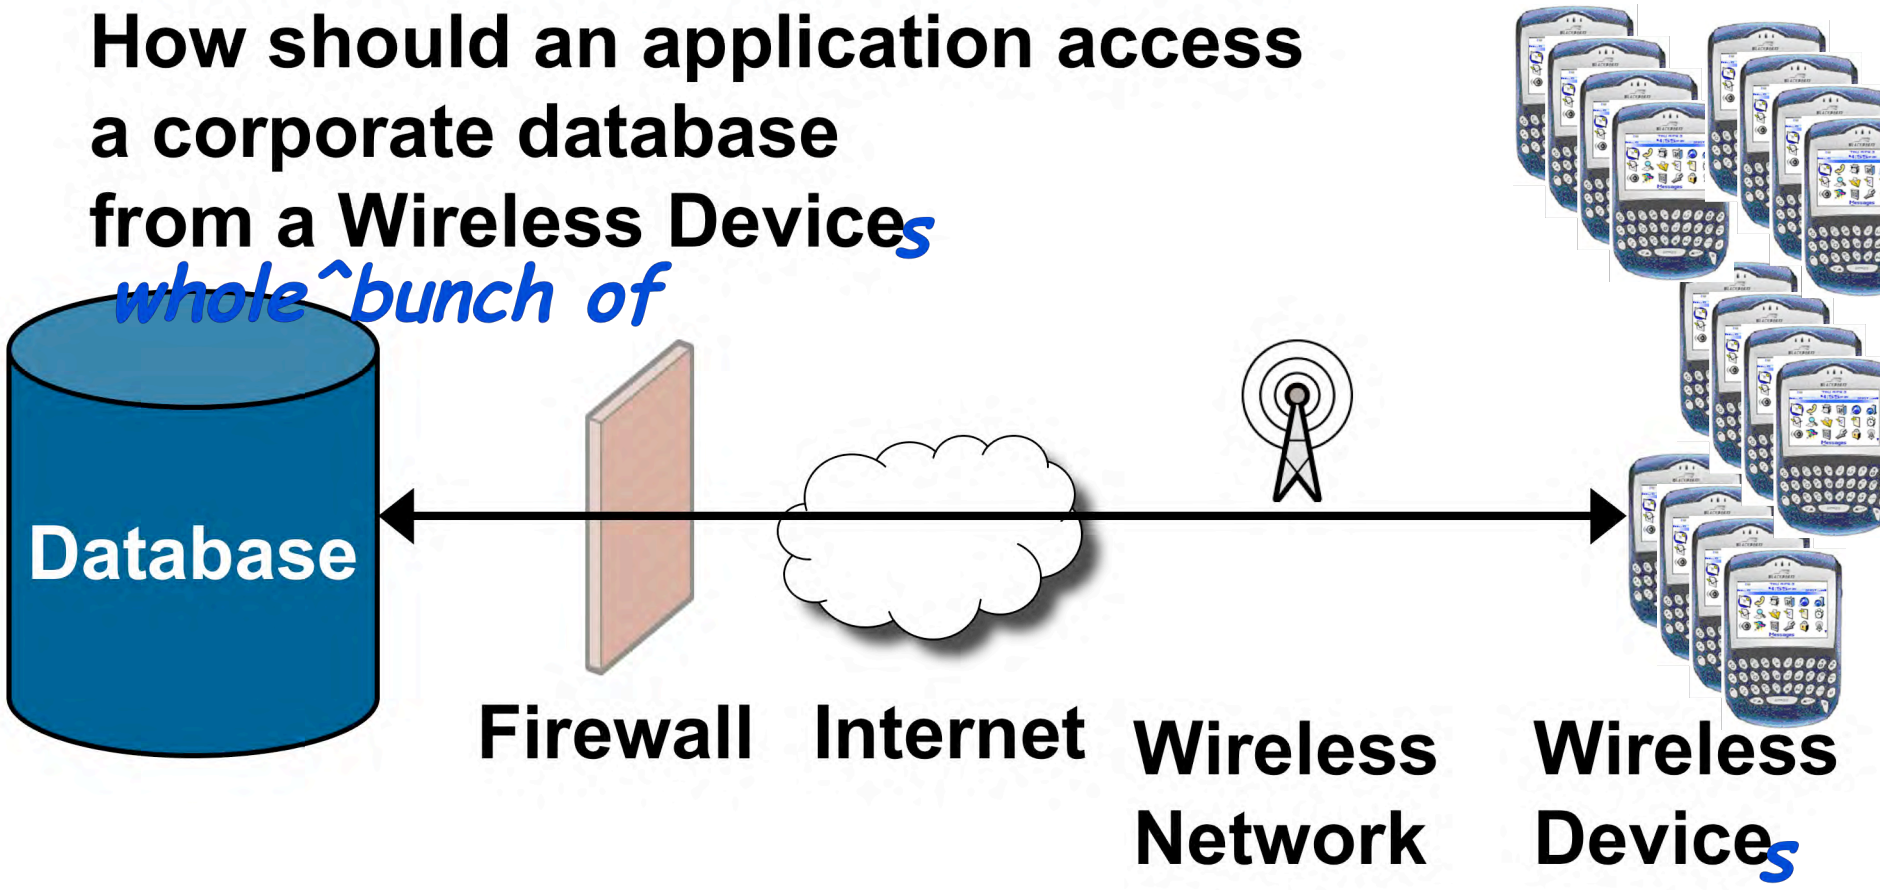

BlackBerry

Problem Statement

How should an application access a corporate database from a Wireless Device

^Enterprise

s

whole^bunch of

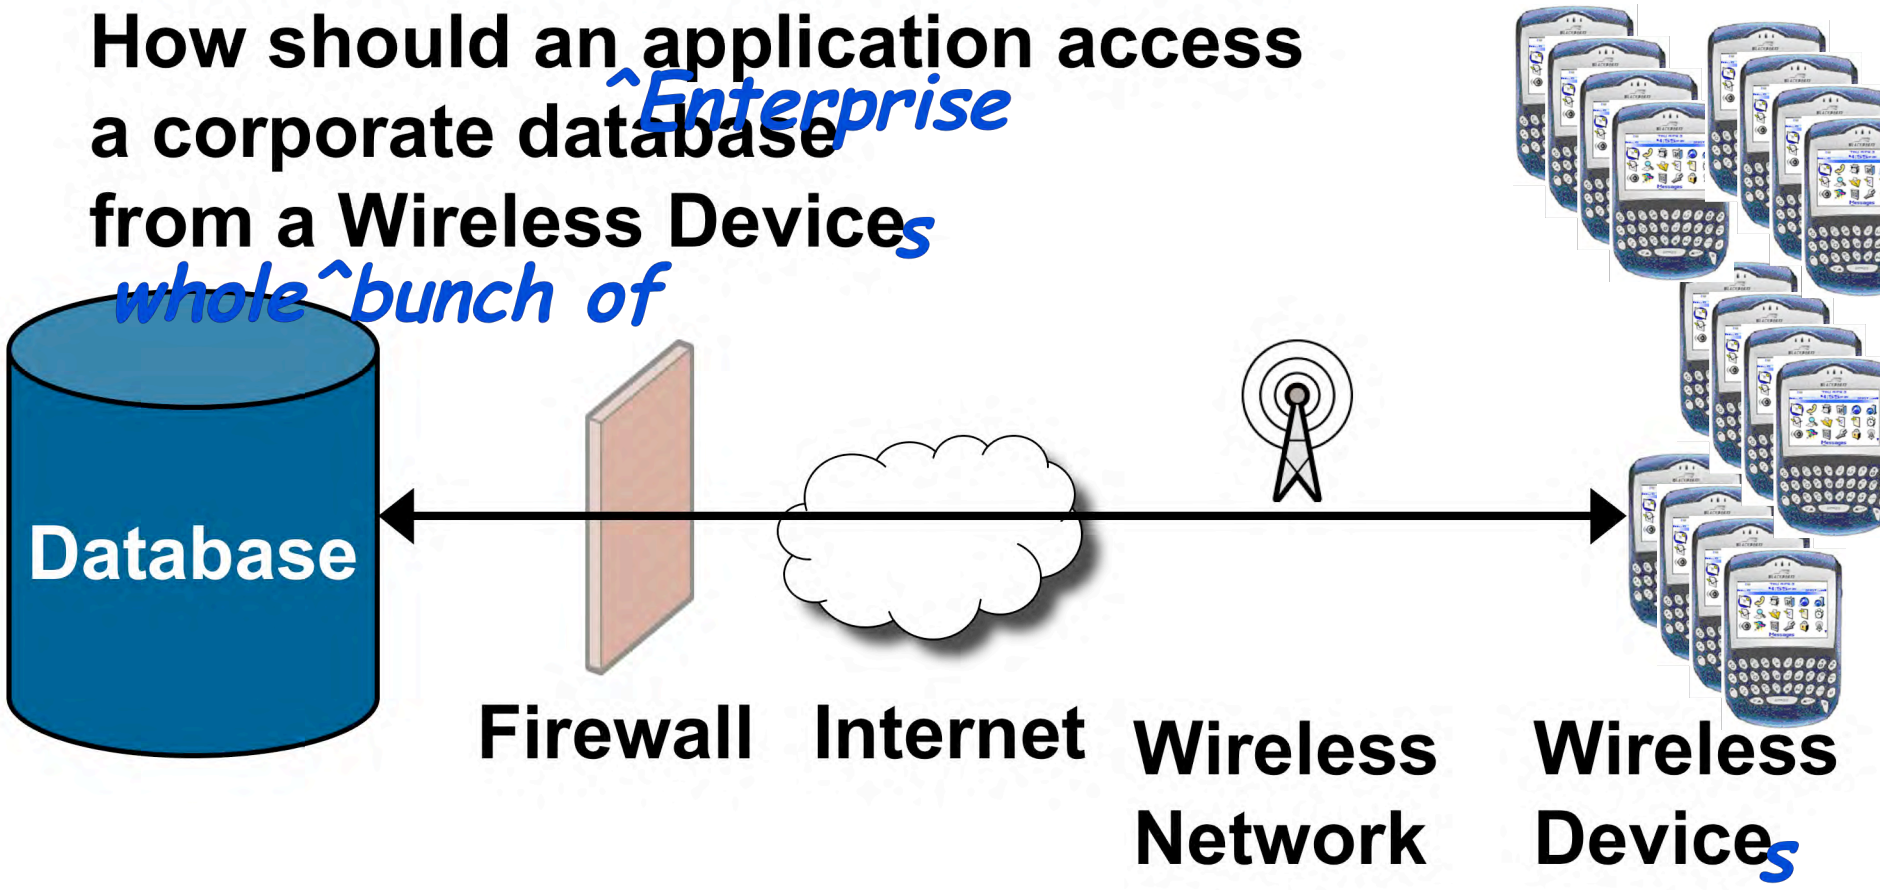

BlackBerry

Potential Solutions in a Wired World

- Browser/Thin Client
- Client/Server
- Distributed database
- Replication

What's Different About Wireless?

- Portability
  - Routing
  - In and out of coverage
  - Not reliable
- Constraints, constraints, constraints!
  - Battery Life
  - Bandwidth
  - Latency
- Asymmetric
  - Constraints
  - Managability and Control

BlackBerry

Wireless Basics

- Radio is essentially low frequency light
  - Travels at $c=3x10^8$ m/second
  - Information packed into waveforms
  - Never receive a clean signal
  - Limited spectrum range practical
- Usually only the last leg is wireless
  - Rest of network is wired
  - Dynamic routing to find the device
- Wireless bandwidth costs money

BlackBerry

Wireless Basics

- Radio is essentially low frequency light
  - Travels at c=3x10^8 m/second
  - Information packed into waveforms
  - Neve
  - Limite
- Usually
  - Rest
  - Dyna
- Wireles

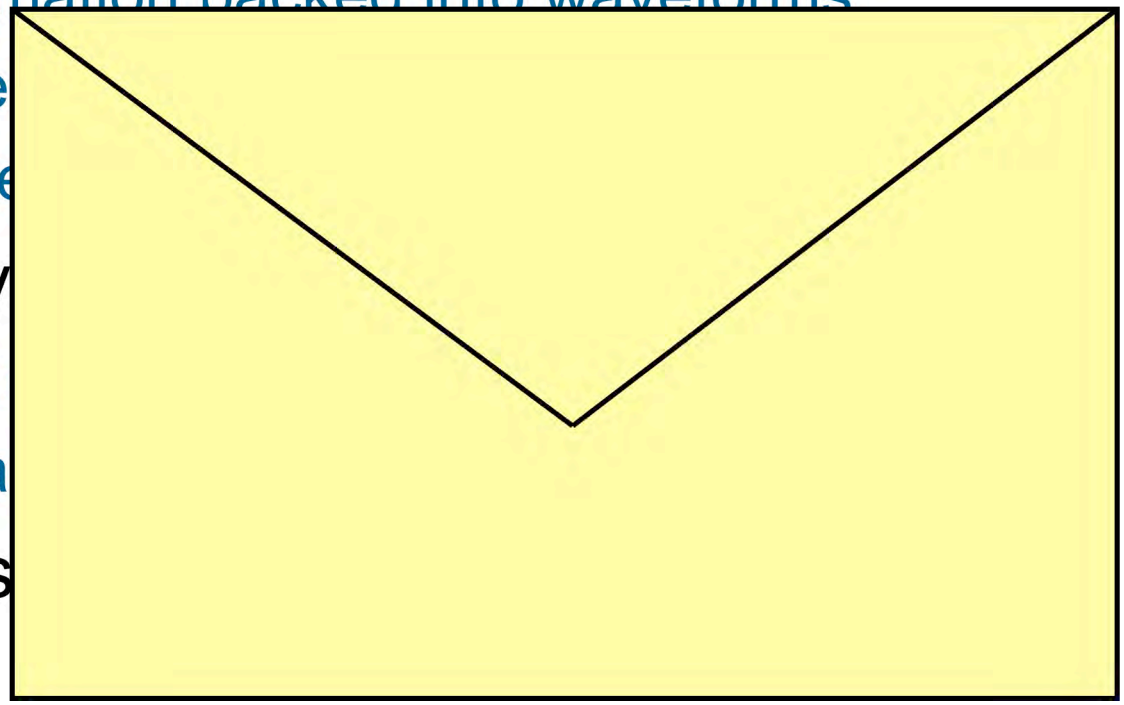

BlackBerry

Wireless Basics

- Radio is essentially low frequency light
  - Travels at c=3x10^8 m/second
  - Information packed into waveforms
  - Neve
  - Limite
- Usually
  - Rest
  - Dyna
- Wireles

Base facts

c = 300,000 km/second

Data spectrum
~800Mhz-5Ghz

Currently ~$2/MB

BlackBerry

Battery Life

- Transmit requires more power than receive
- Receive is on more than transmit
- Radio, screen/backlight, radio, audio, …
- Not enough power to listen all the time
  - Device frequently suspends

BlackBerry

Battery Life

- Transmit requires more power than receive
- Receive is on more than transmit
- Radio, [illegible]
- Not enc[illegible]
  - Devic[illegible]

Approximations

Transmit/bit = 10 x Receive/bit

Listen 5% of the time (50ms every second)

1 battery = 5000 packets

BlackBerry

Bandwidth

- The myth of wireless bandwidth
  - Quoted at peak, not at capacity
  - Measure # bits/second in available spectrum
  - Can't pack signals too close together
- Increasing capacity via smaller cells
  - When size is 1/n, cost is n^2
- Radio power roughly proportional to # bits
  - More data = more battery used

BlackBerry

Bandwidth

- The myth of wireless bandwidth
  - Quoted at peak, not at capacity
  - Measure # bits/second in available spectrum
  - Can't
- Increas
  - Wher
- Radio p s
  - More

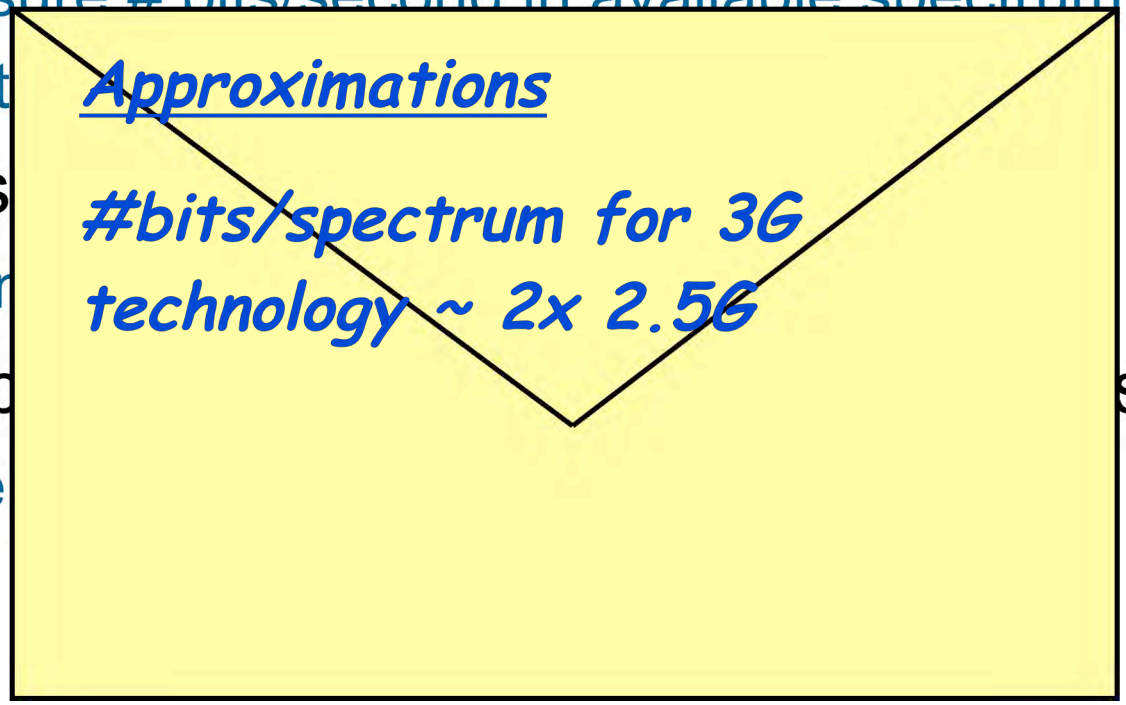

BlackBerry

Latency

- High bandwidth does not imply low latency
- Possible impacts of dynamic routing, battery life
- Speed of light can be a factor
- Devices sleep frequently
  - Heuristics can reduce this after first contact

BlackBerry

Latency

- High bandwidth does not imply low latency
- Possible impacts of dynamic routing, battery life
- Speed of light can be a factor
- Devices
  - Heuri

Approximations

Avg push latency in 2.5G or 3G ~0.5 seconds for ONE device

Other latency ~50ms

Half way around the world takes ~100ms at speed of light

BlackBerry

Not Good Things in Wireless

- Polling
- Multi-pass protocol
- Streaming

Not Good Things in Wireless

- Polling
- Multi-pass protocol
- Streami

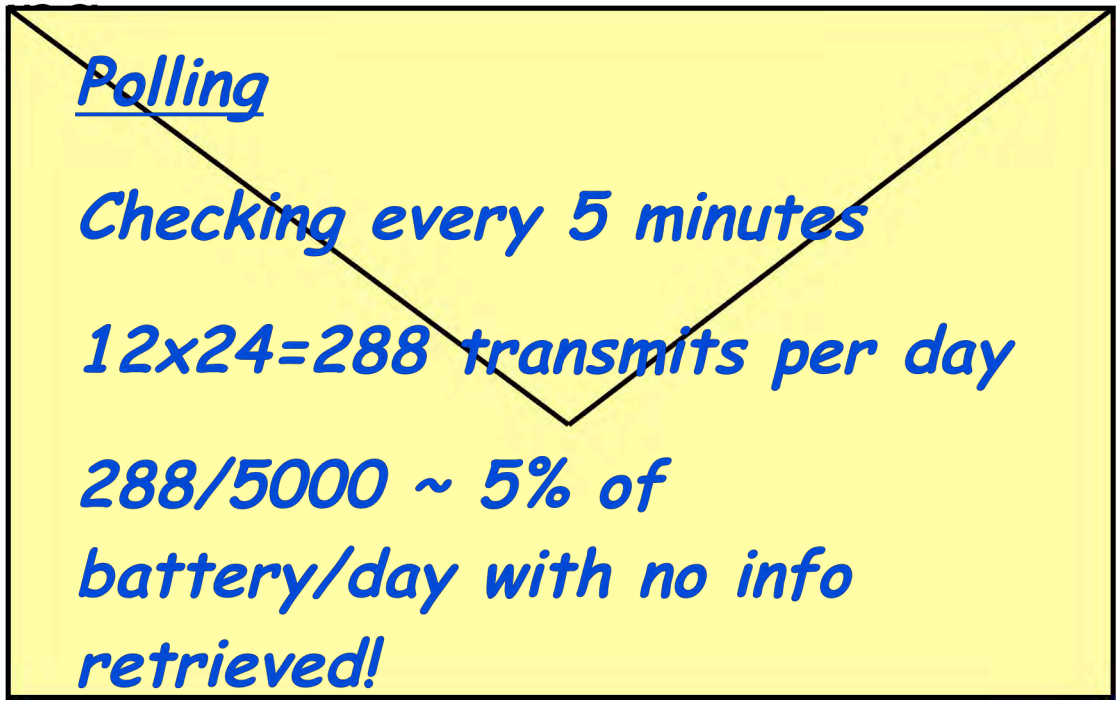

BlackBerry

Not Good Things in Wireless

- Polling
- Multi-pass protocol
- Streami

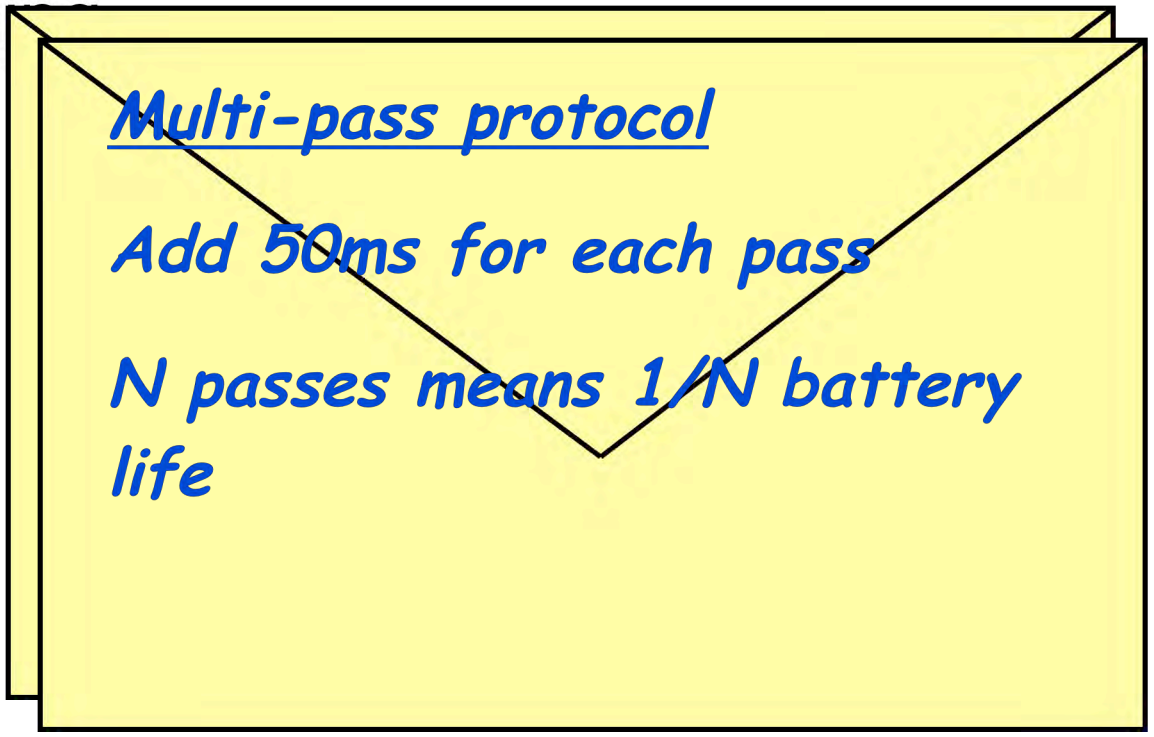

BlackBerry

Not Good Things in Wireless

- Polling
- Multi-pass protocol
- Streaming

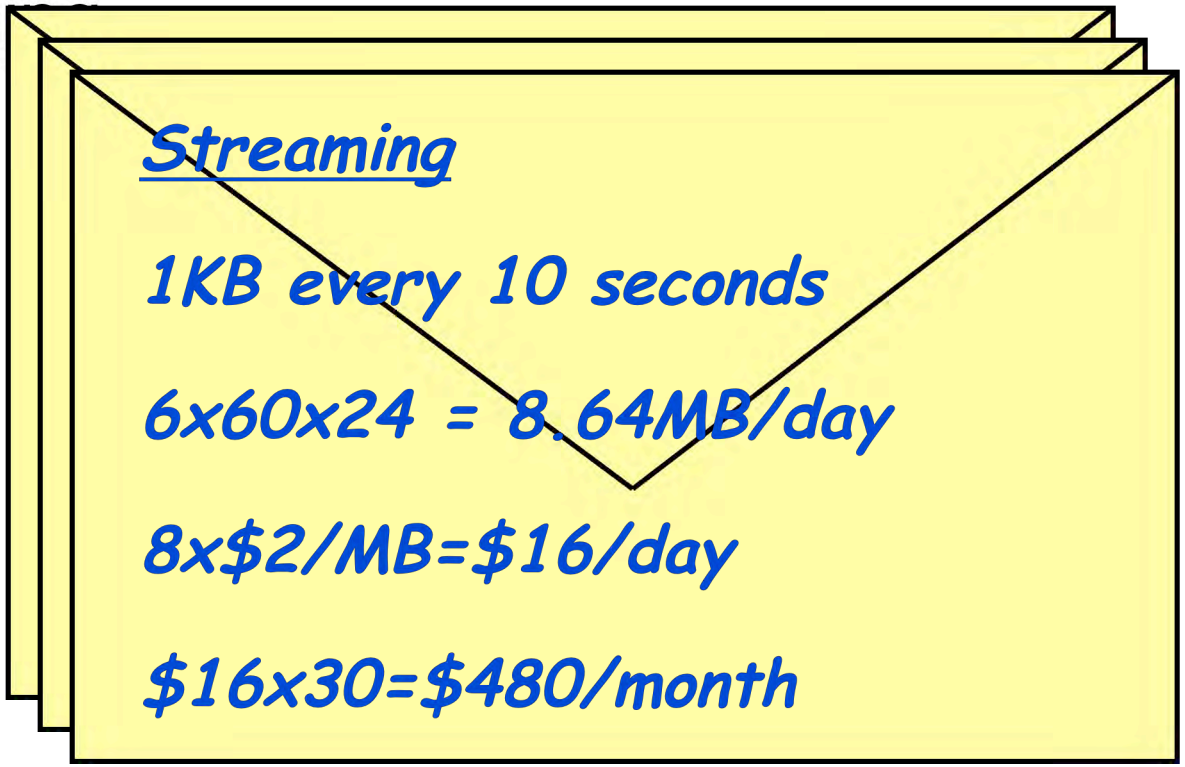

BlackBerry

Good Things in Wireless

- Push
- Transferring data only once
- Don’t count on coverage
  - Complicates Push
- Efficient data transfers
  - E.g. binary instead of XML

BlackBerry

Why Push Changes Everything

- Device
  - Device must be “Always On, Always Connected”
  - Application must be ready to receive data at any time
  - Must have proper background processing
- Infrastructure
  - Device may not be in coverage
  - Requires queuing
  - Reliability and retries are used frequently
    - Must be efficient!
    - Bad implementations can kill the network

BlackBerry

Potential Solutions in a Wired World

- Browser/Thin Client
- Client/Server
- Distributed database
- Replication

Potential Solutions in a Wired World

less?

- Browser/Thin Client
- Client/Server
- Distributed database
- Replication

BlackBerry

Browser/Thin Client

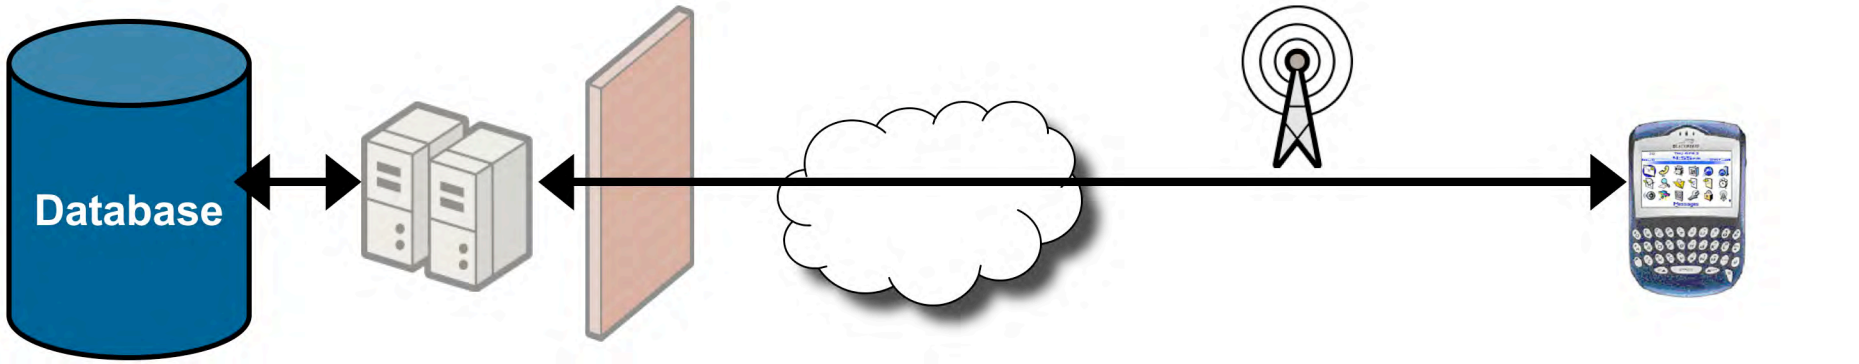

Application Server

Browser

BlackBerry™

Browser/Thin Client

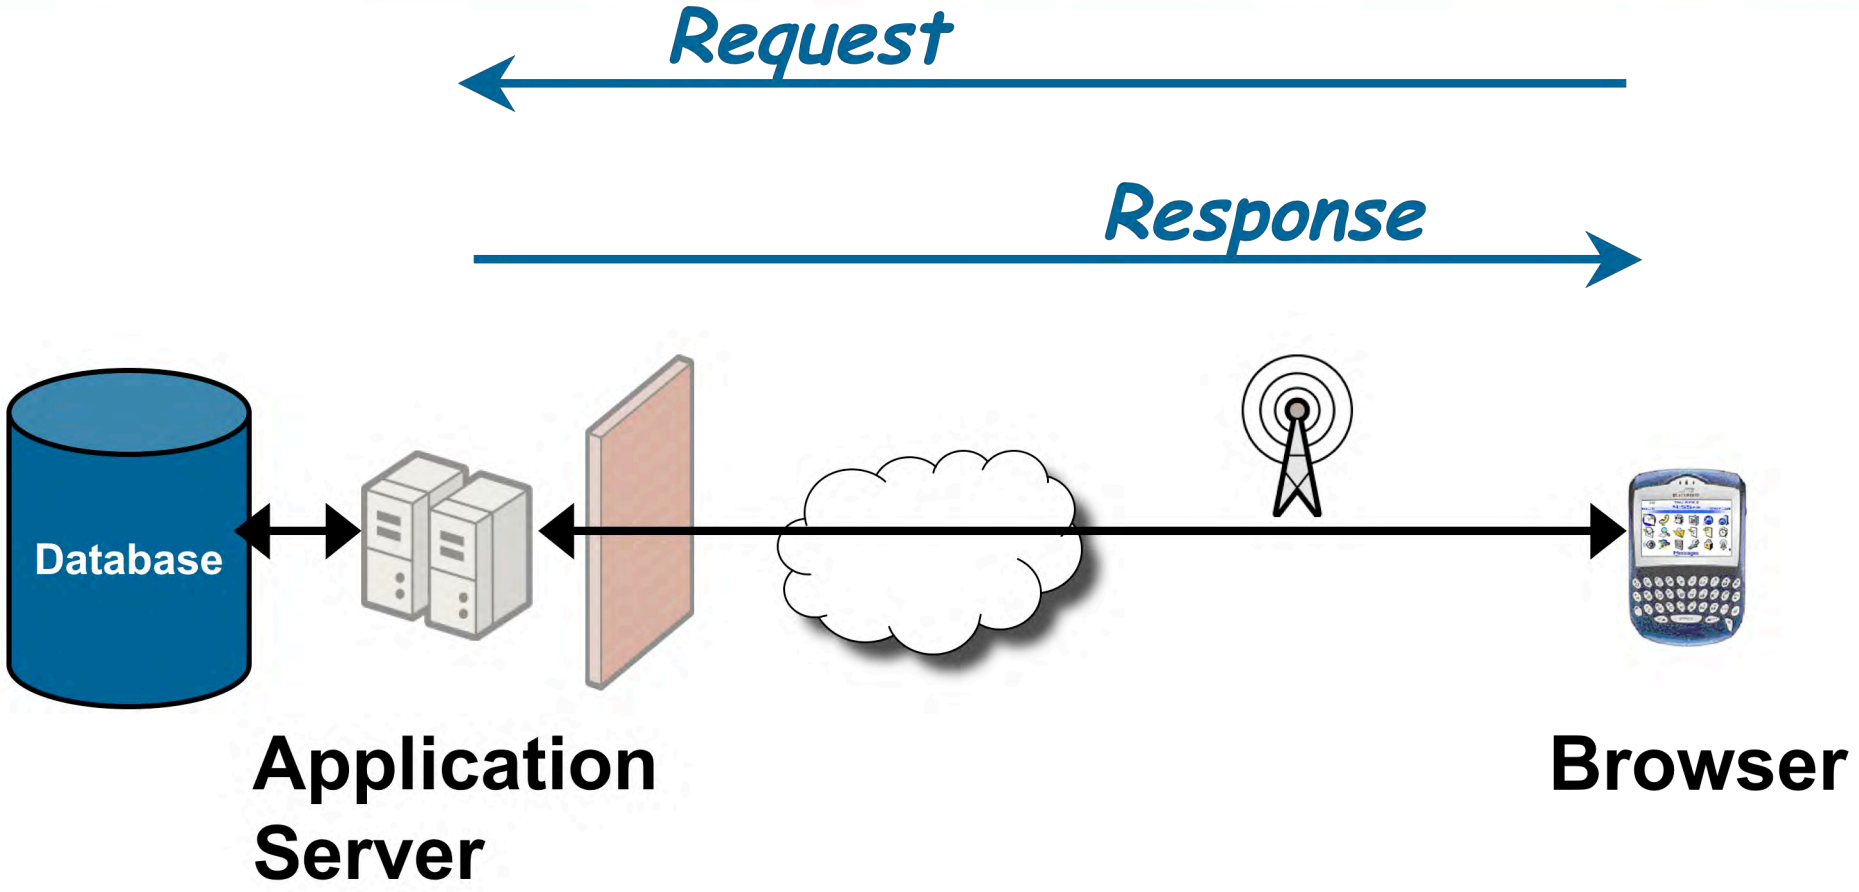

BlackBerry

Browser/Thin Client

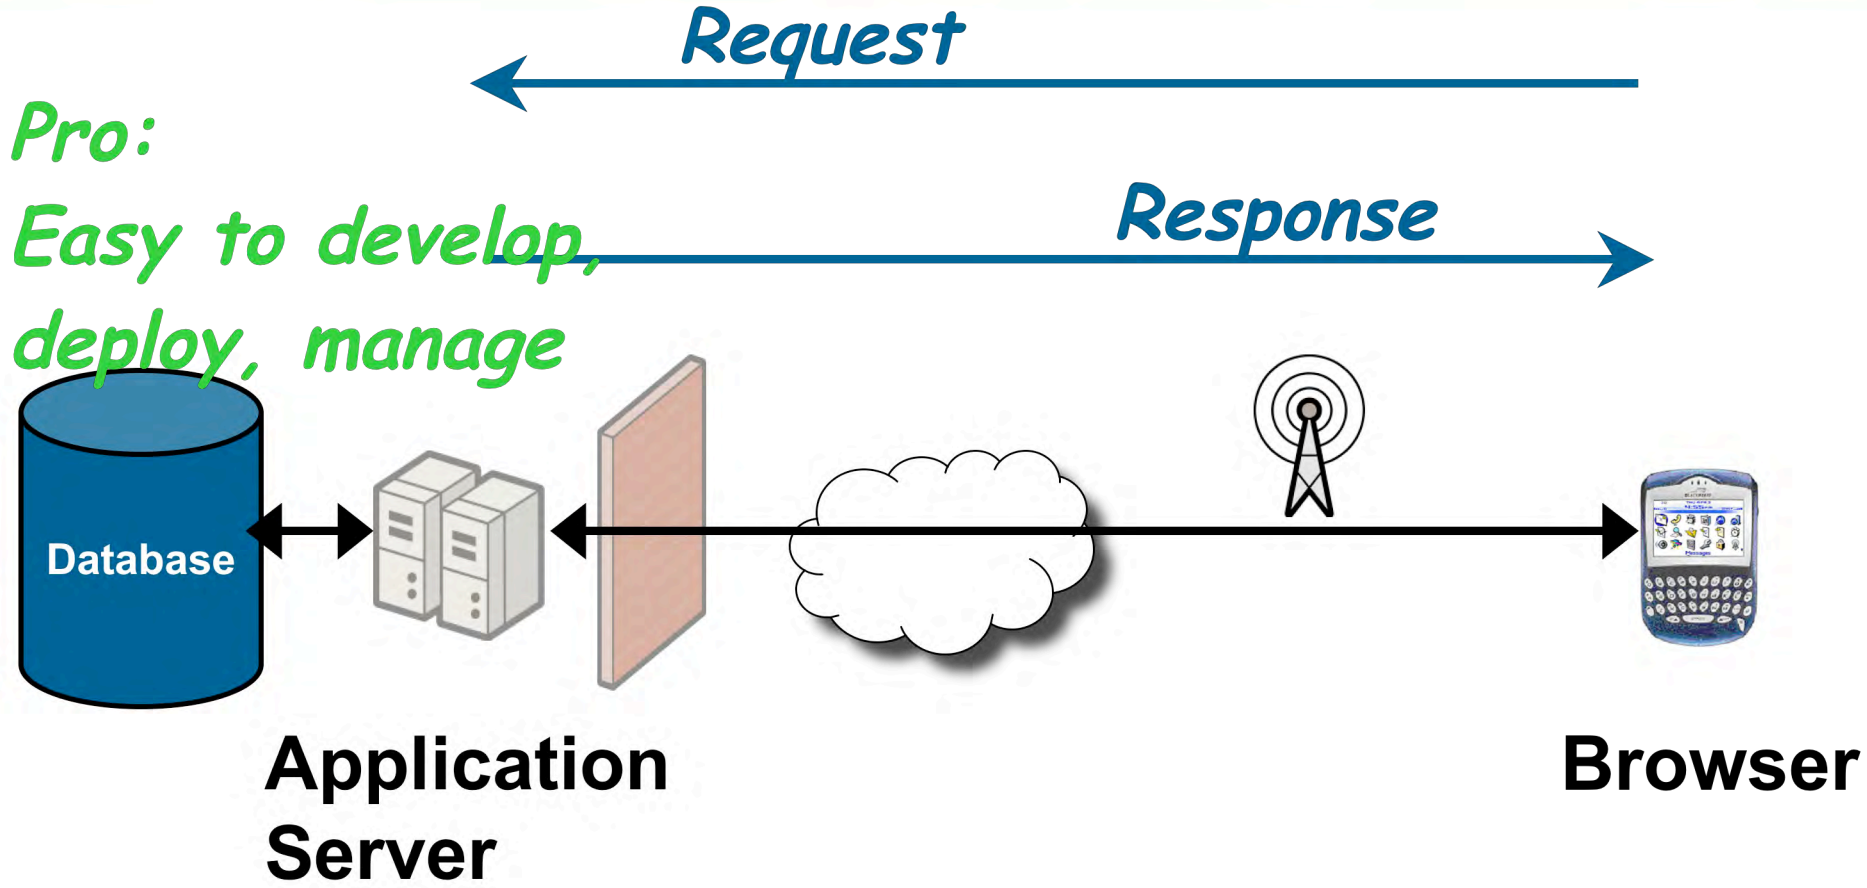

BlackBerry

Browser/Thin Client

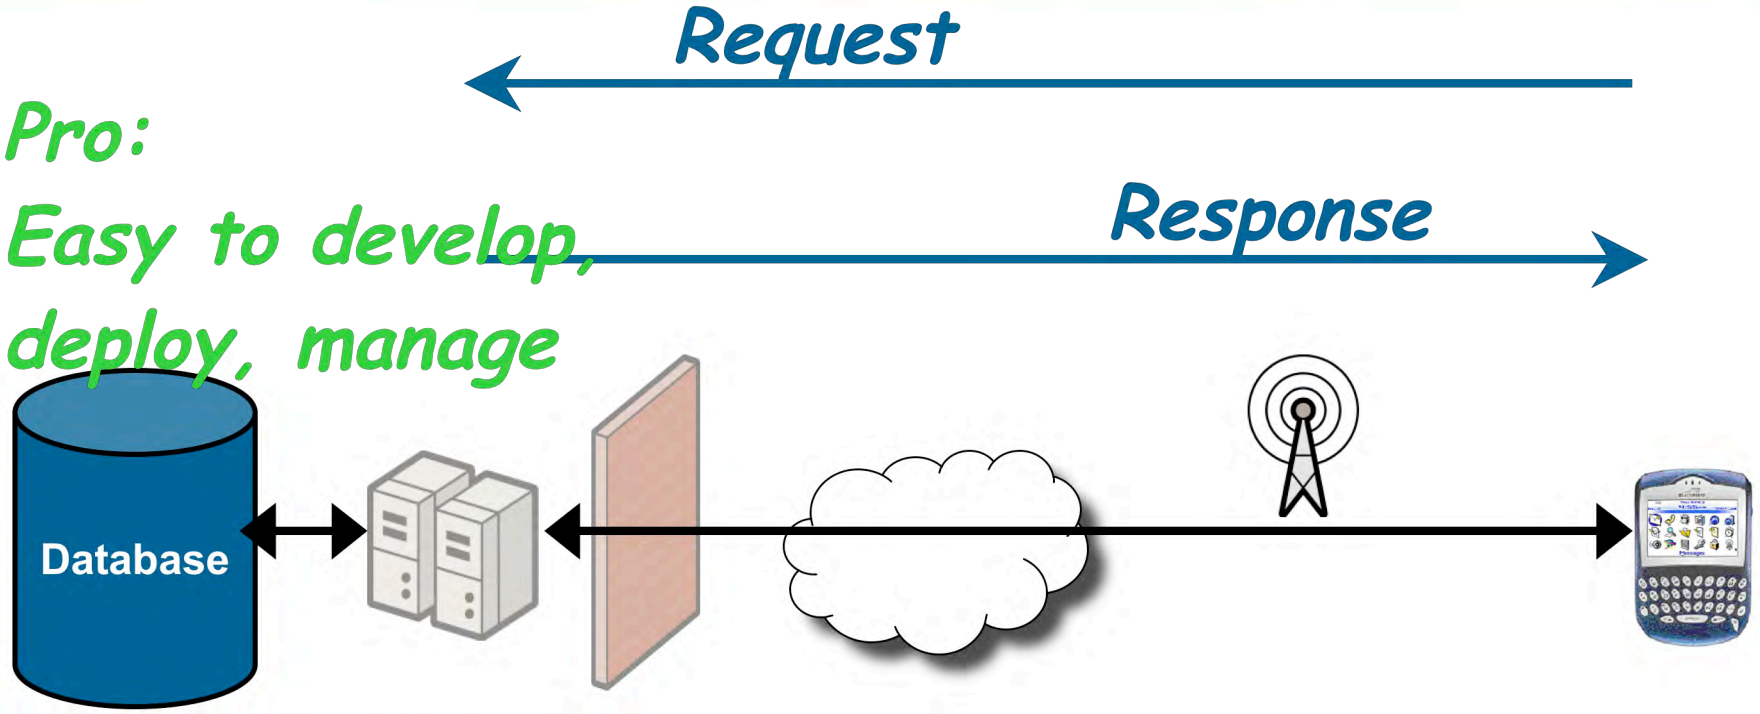

Browser/Thin Client

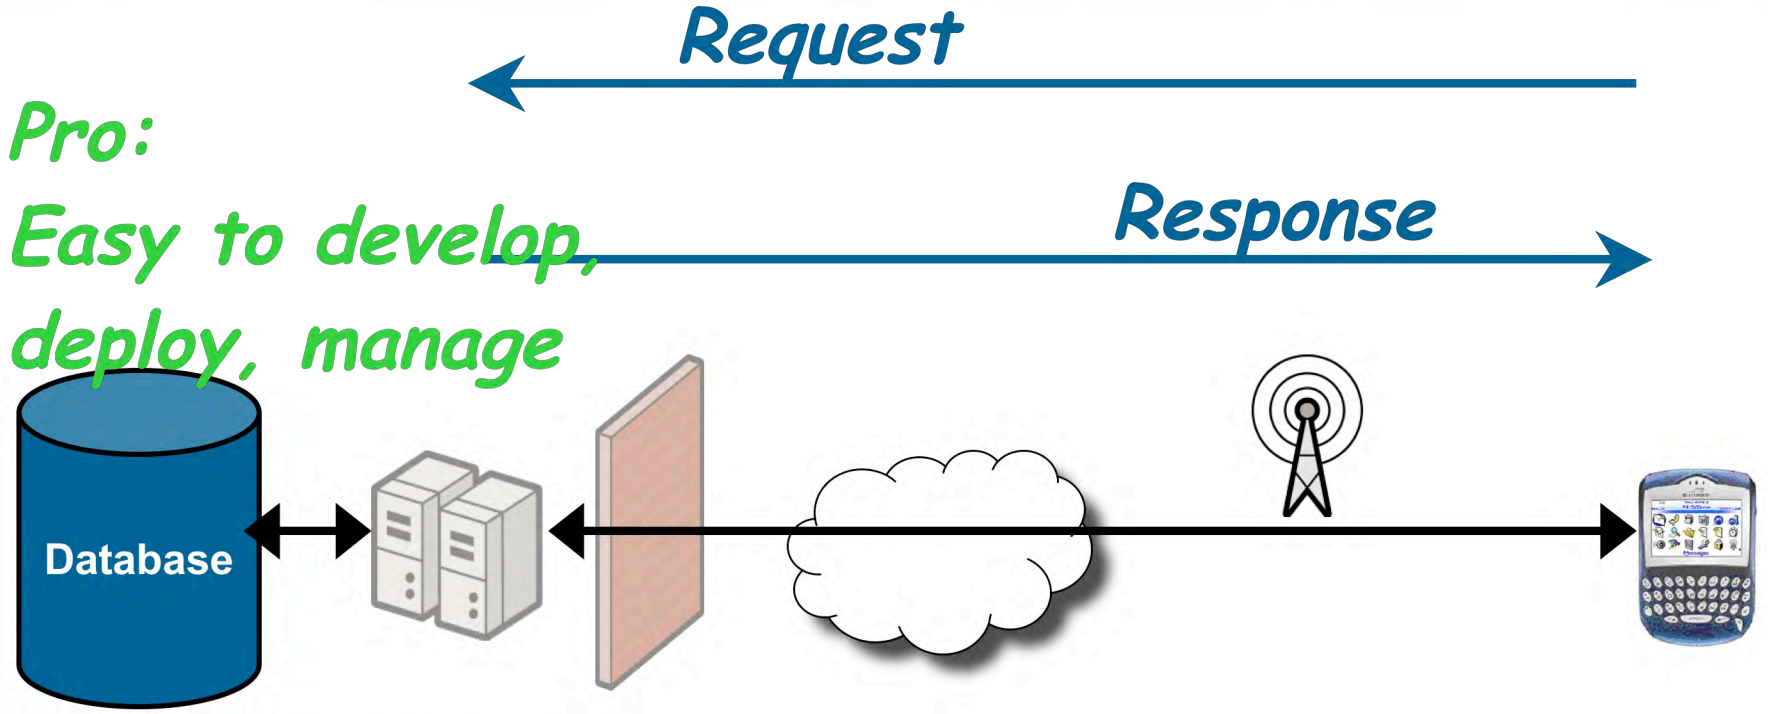

Browser/Thin Client

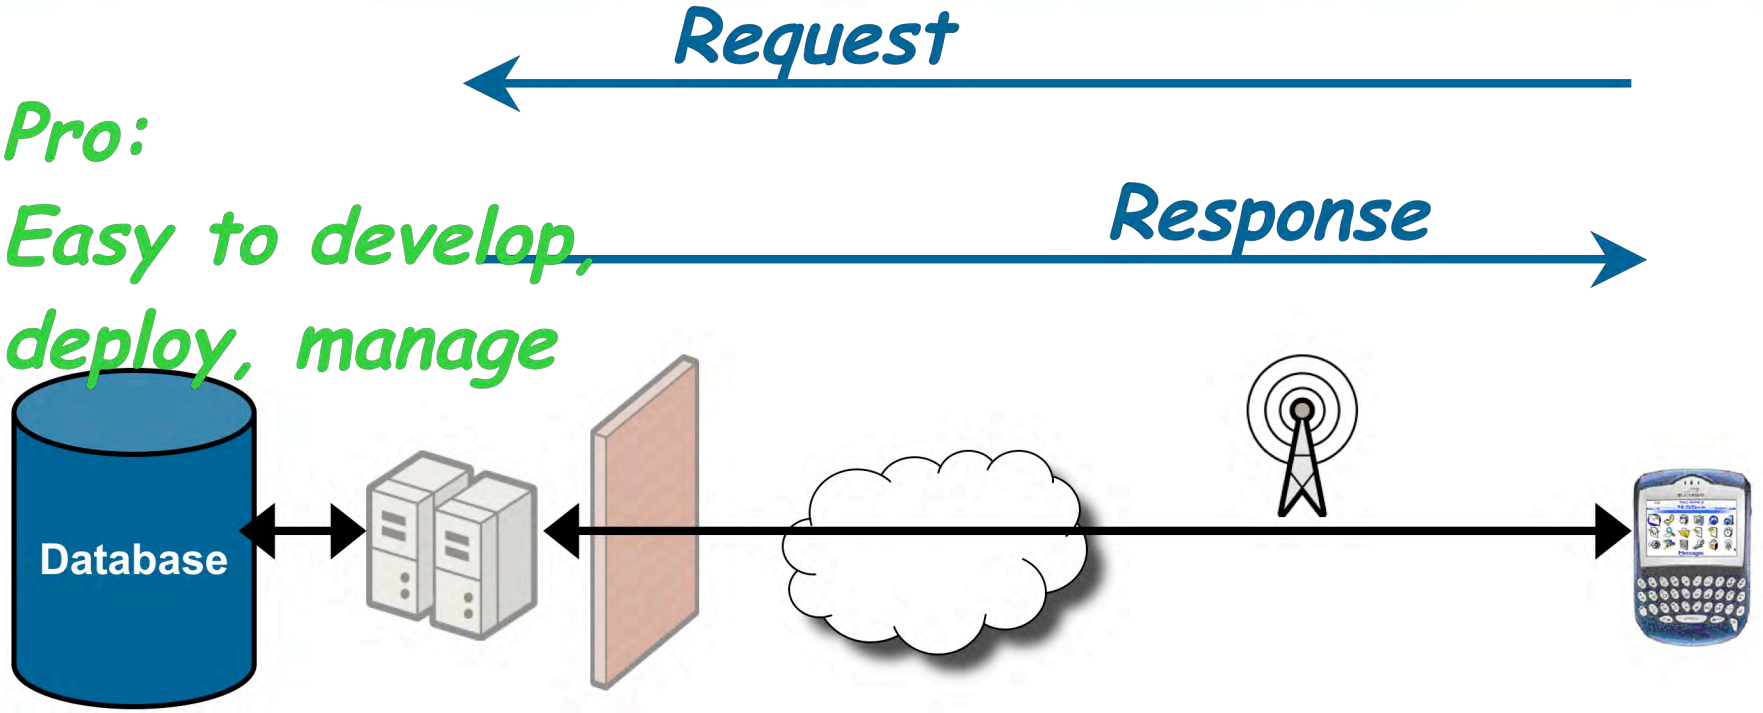

Browser/Thin Client

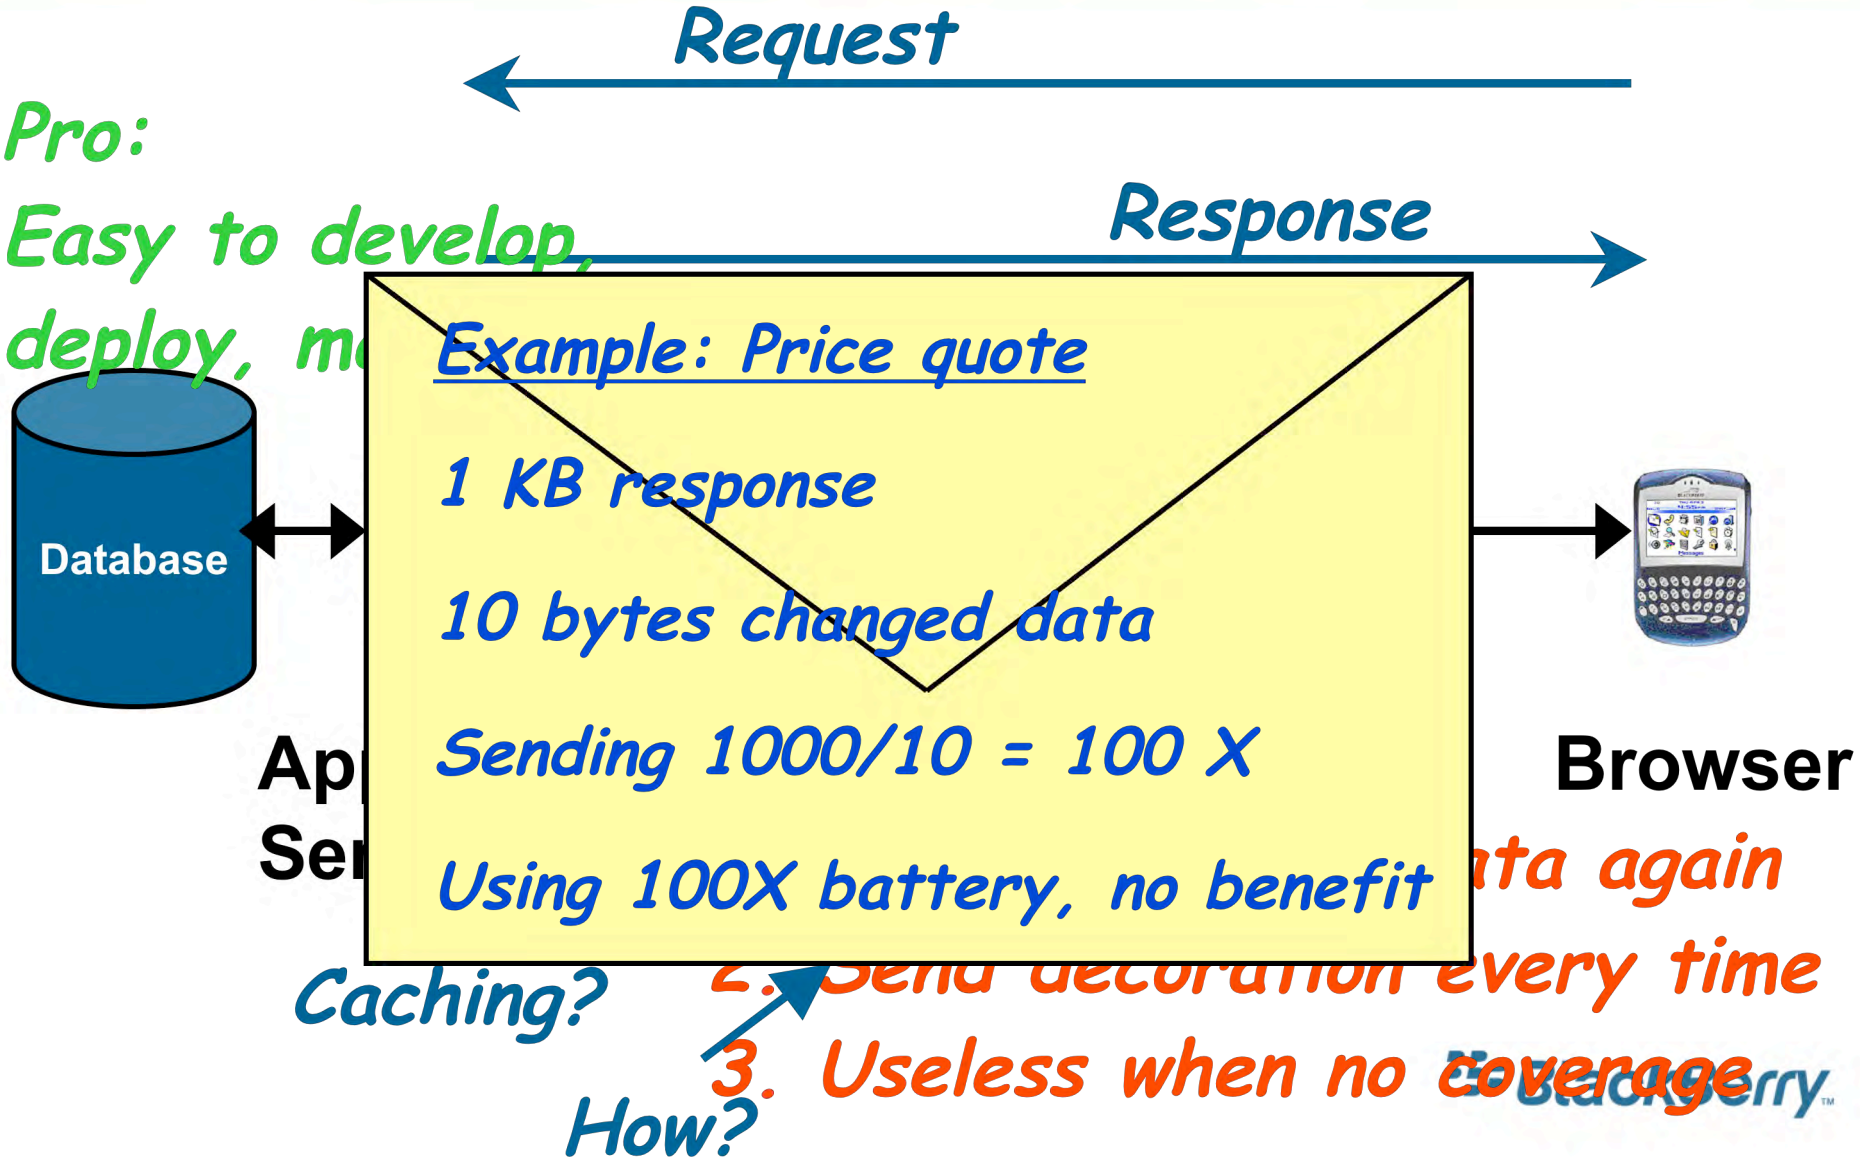

Client/Server

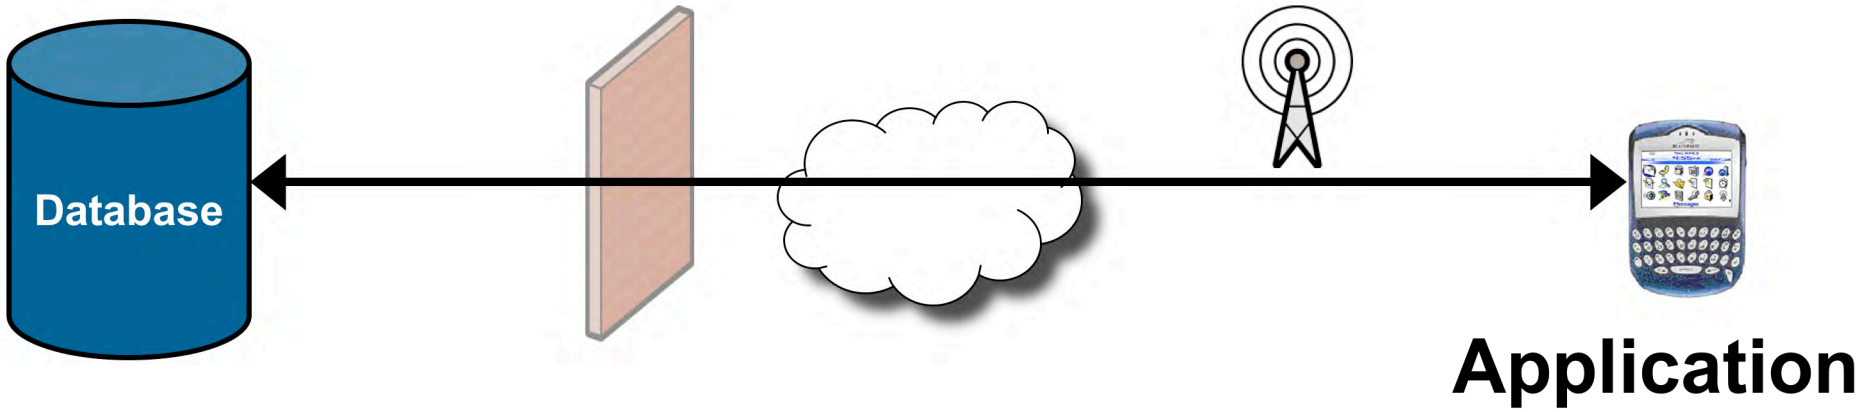

BlackBerry

Client/Server

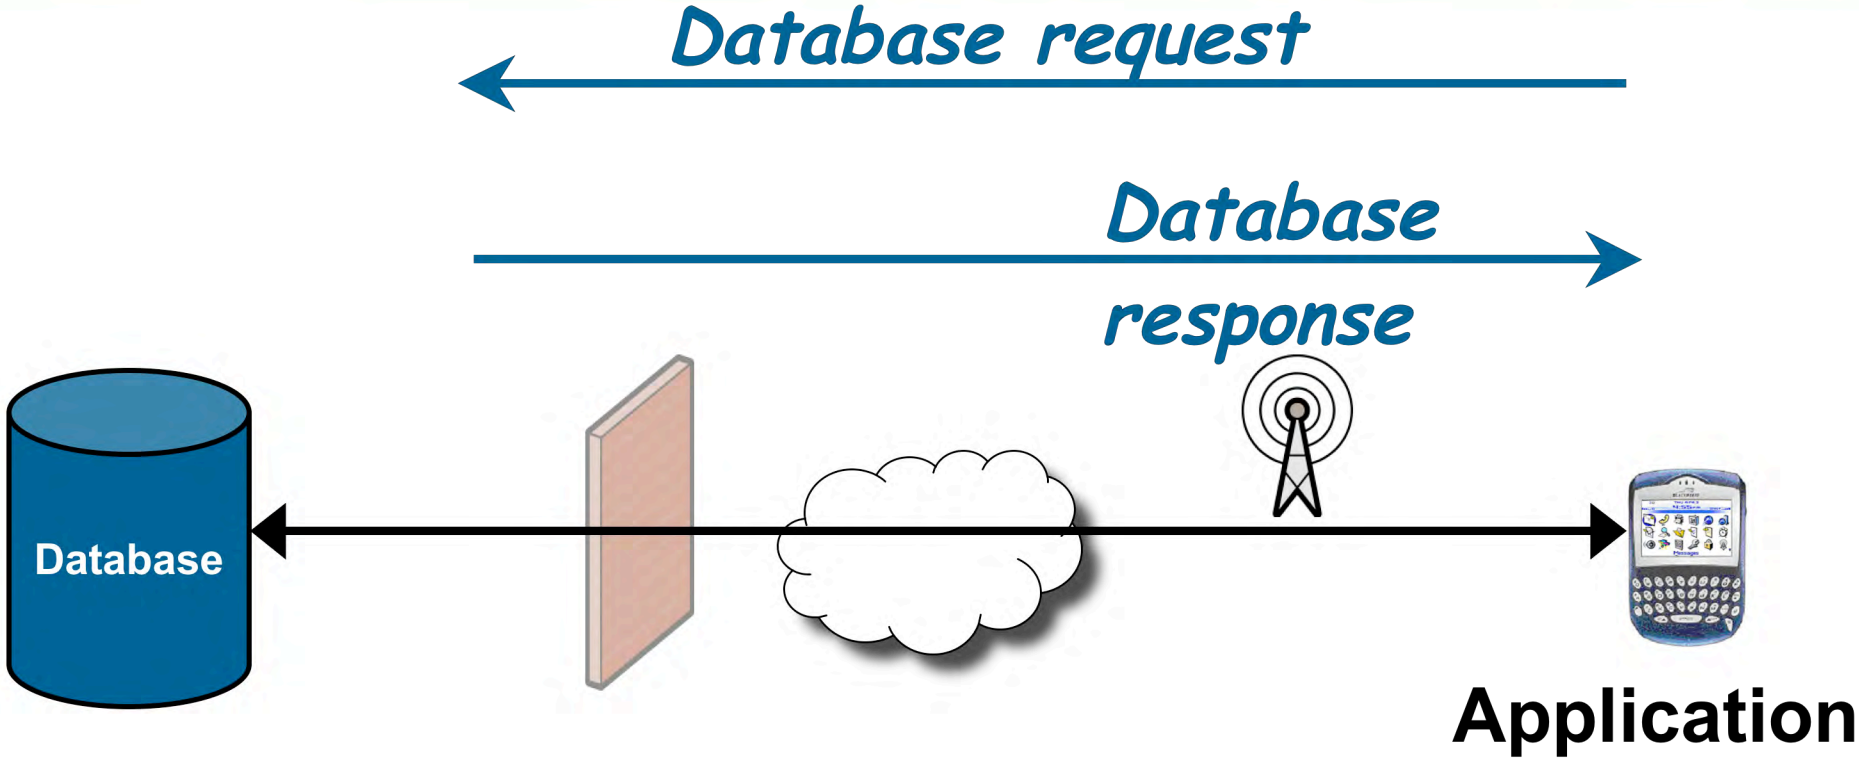

BlackBerry

Client/Server

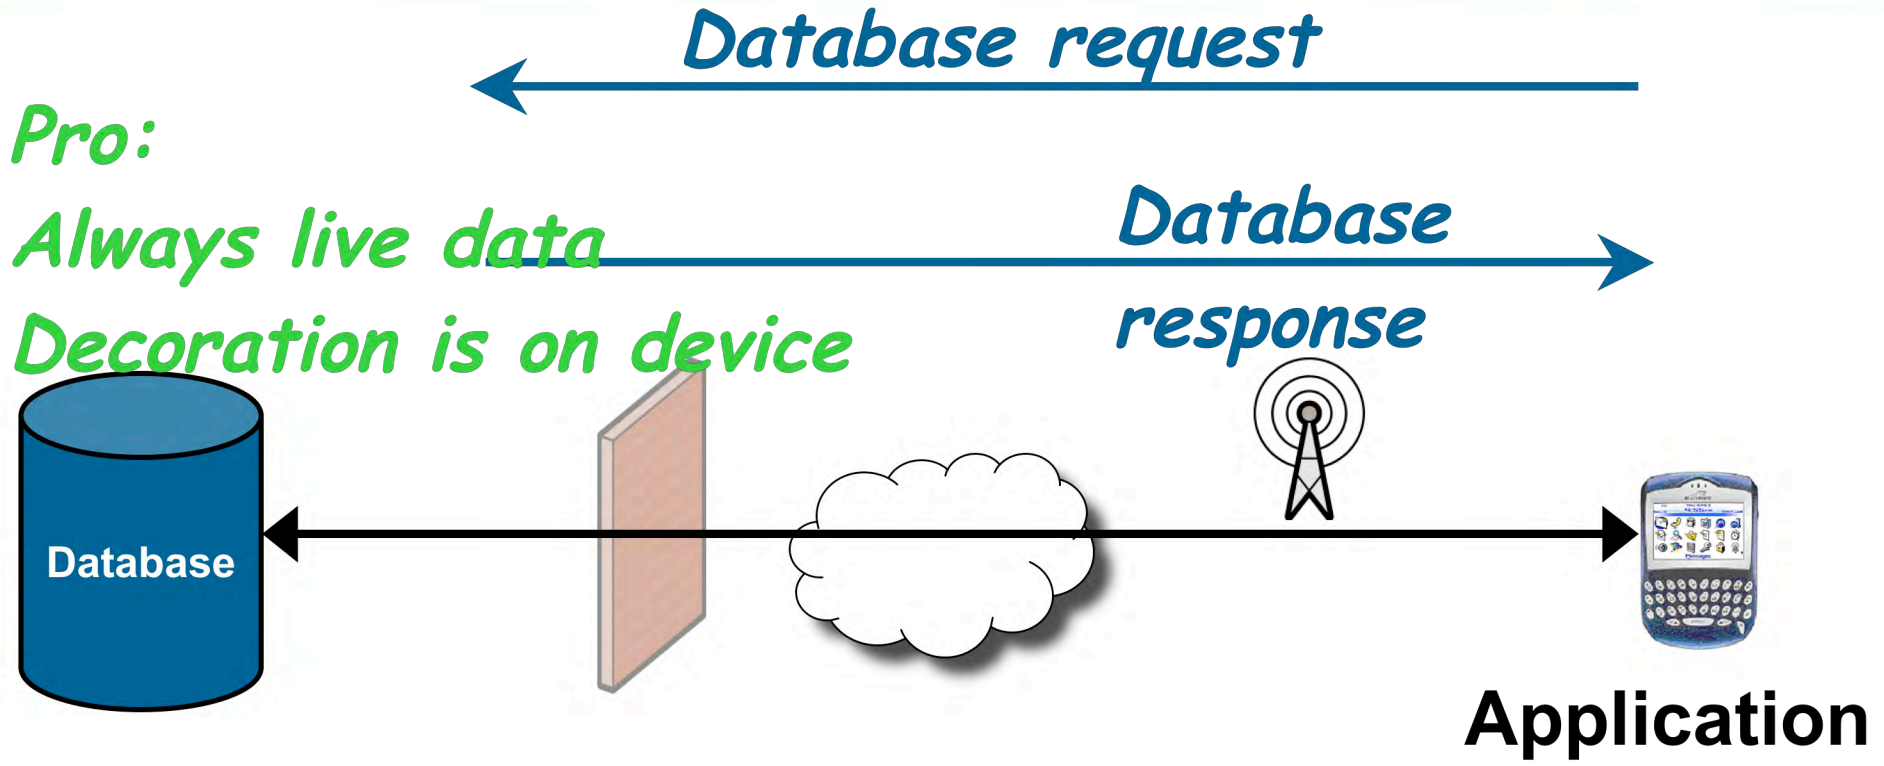

BlackBerry

Client/Server

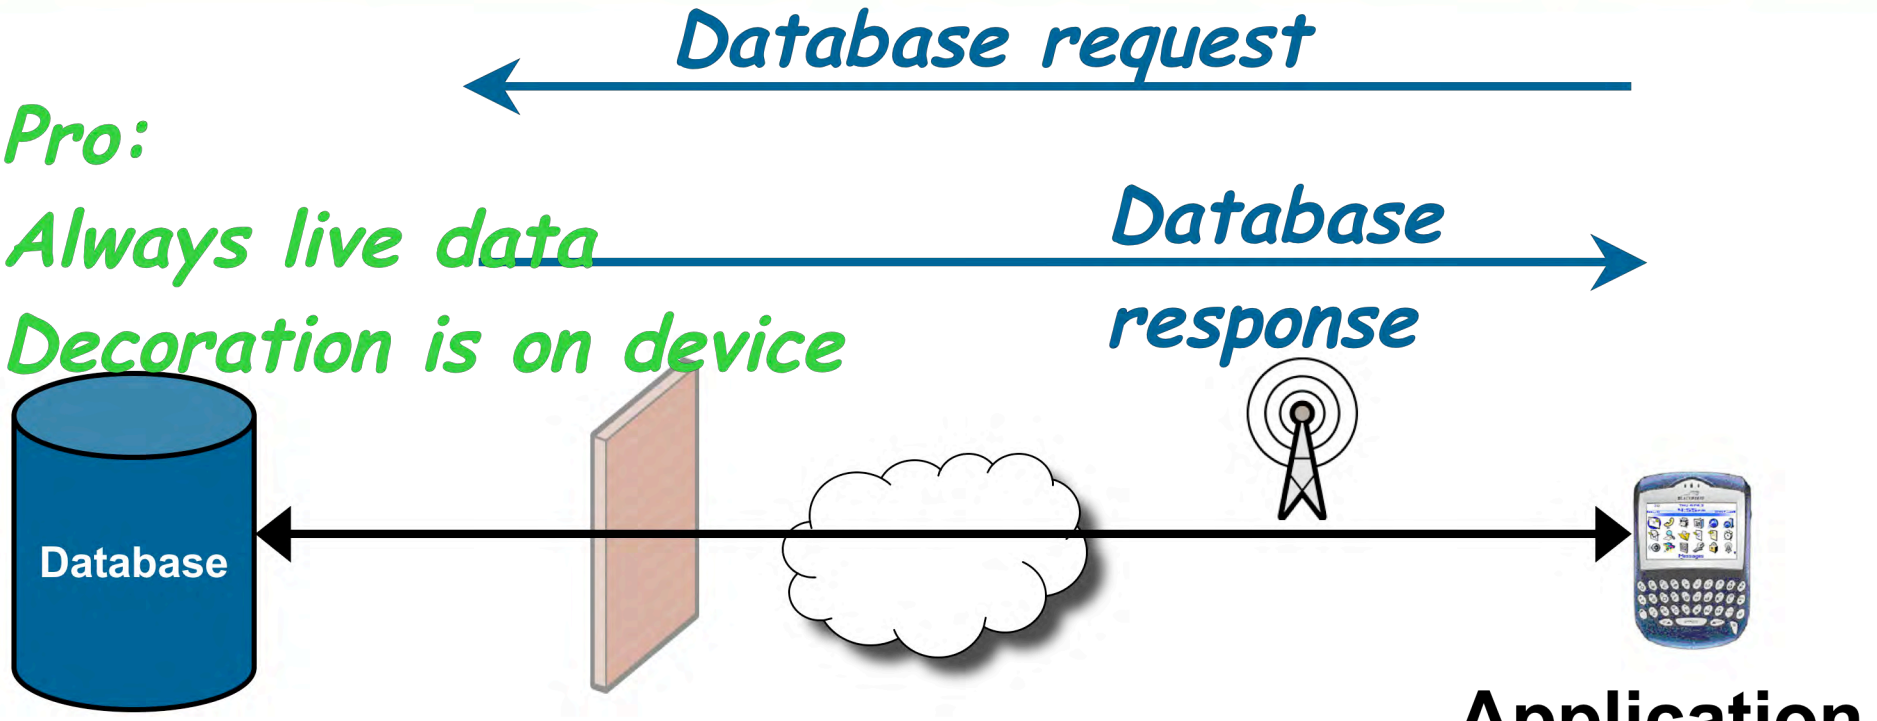

Cons:

1. May send same data again
2. May be many passes
3. Useless when no coverage

Client/Server

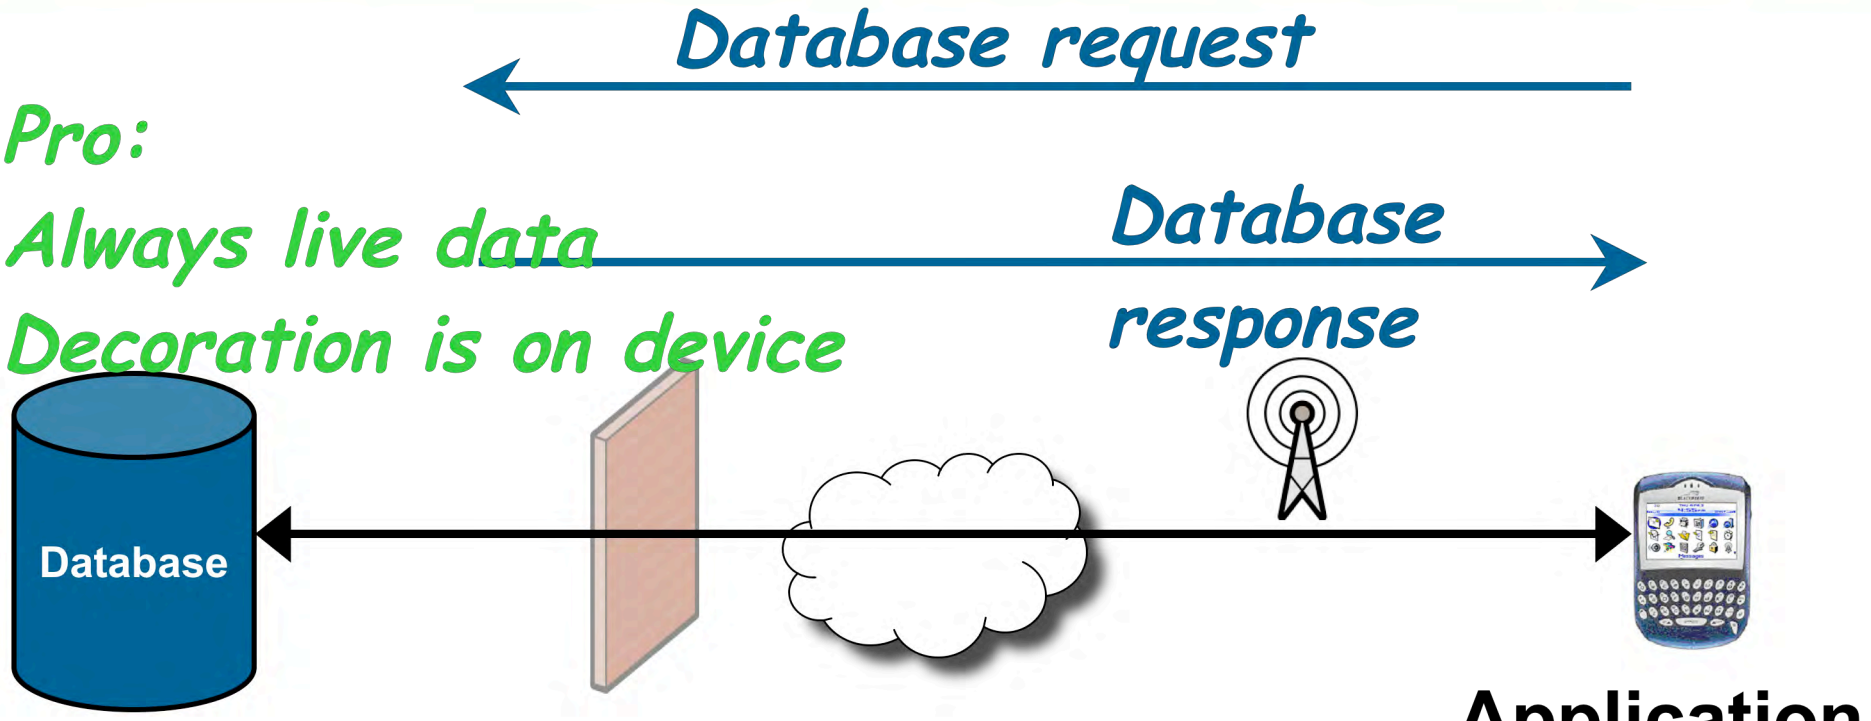

Cons:

1. May send same data again
2. May be many passes
3. Useless when no coverage

How to know?

Client/Server

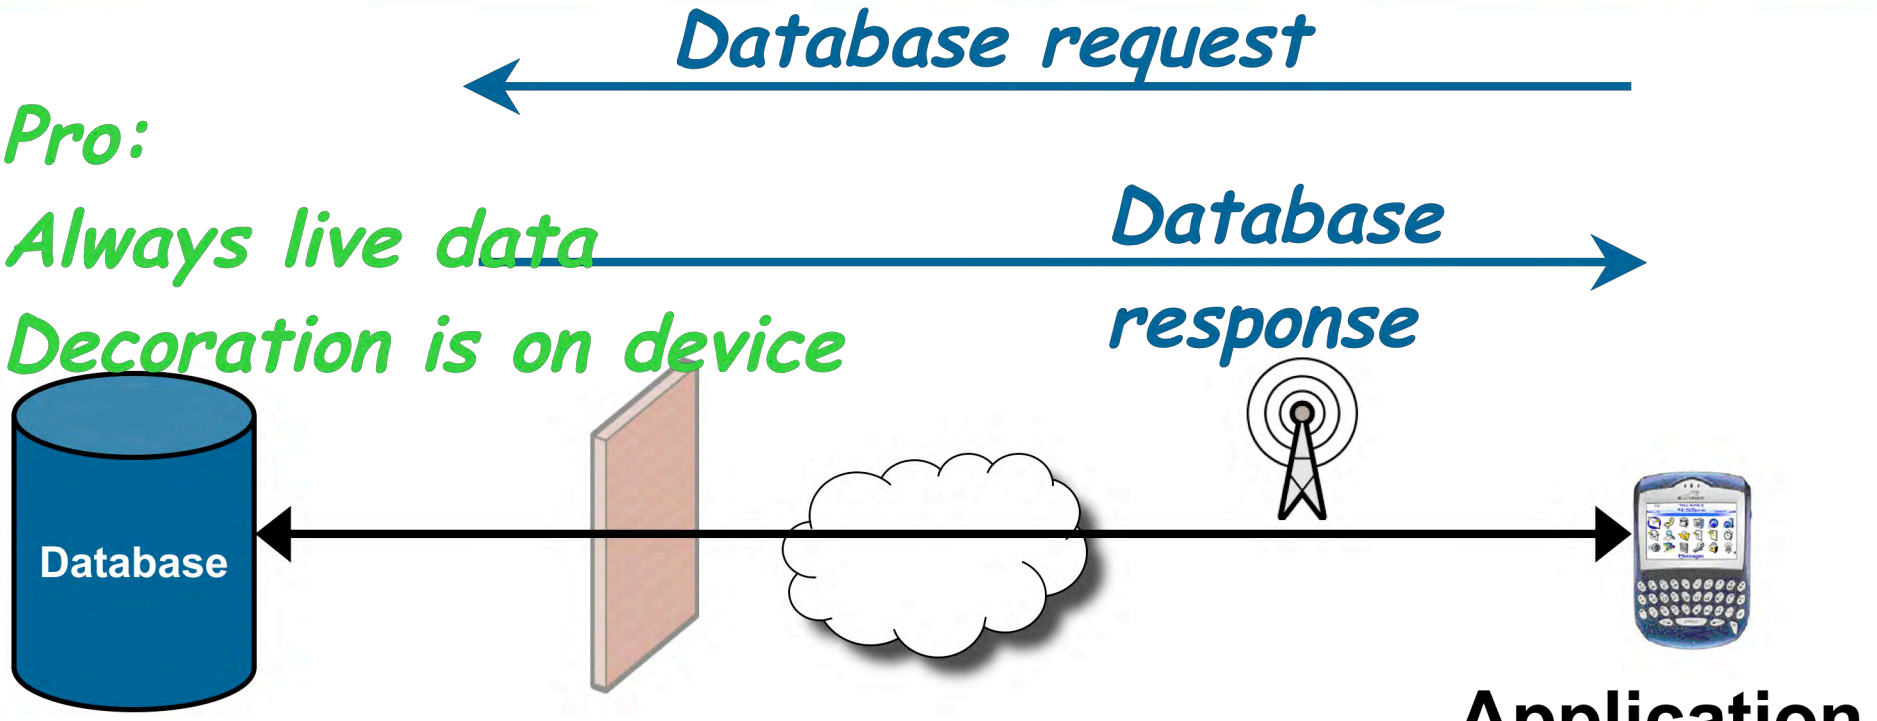

Cons:

1. May send same data again
2. May be many passes
3. Useless when no coverage

How to know?

Batch operations?

Client/Server

Database request

Pro:

Always live data

Database

Decoratio

Example: Price list

Fetching 1000 rows

1000 requests

1000/5000 = 20% of battery

And usually no changes!

Database

Application

ta again

2. May be many passes

3. Useless when no coverage

How to know?

Batch operations?

Client/Server

Database request

Pro:

Always live data

Database

Decoratio

Batched Price list

1000 rows X 100 bytes

10X/day, 1000 sales reps

1MB/day X 1000 X 20 days/month X $2/MB

$40,000/month!

Database

pplication

ta again

es

How to know

Batch operations?

3. Useless when no coverage

BlackBerry

Distributed Database

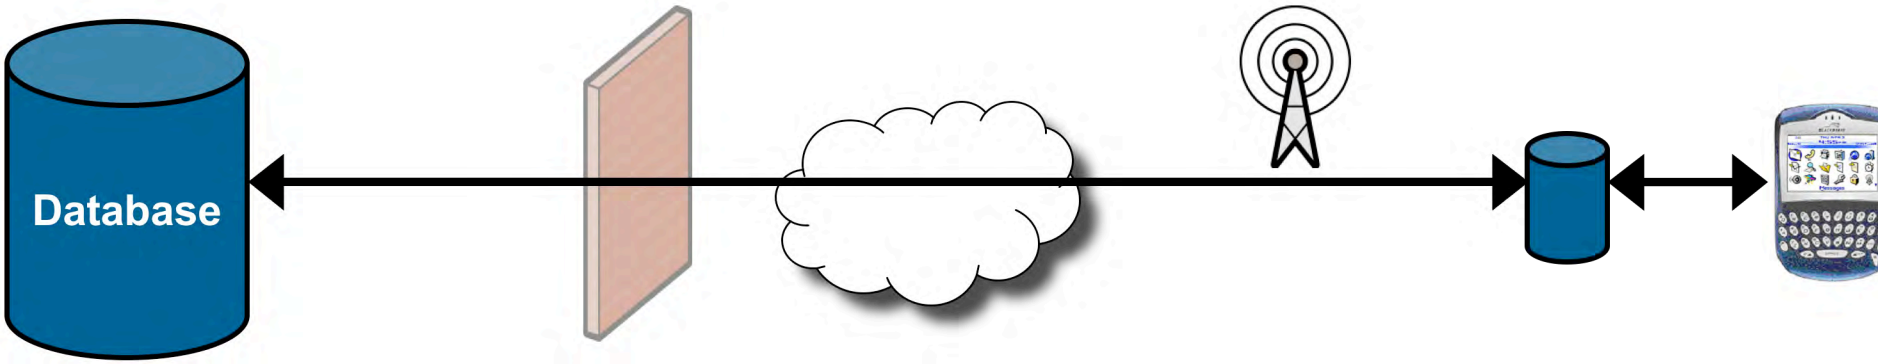

Application

BlackBerry

Distributed Database

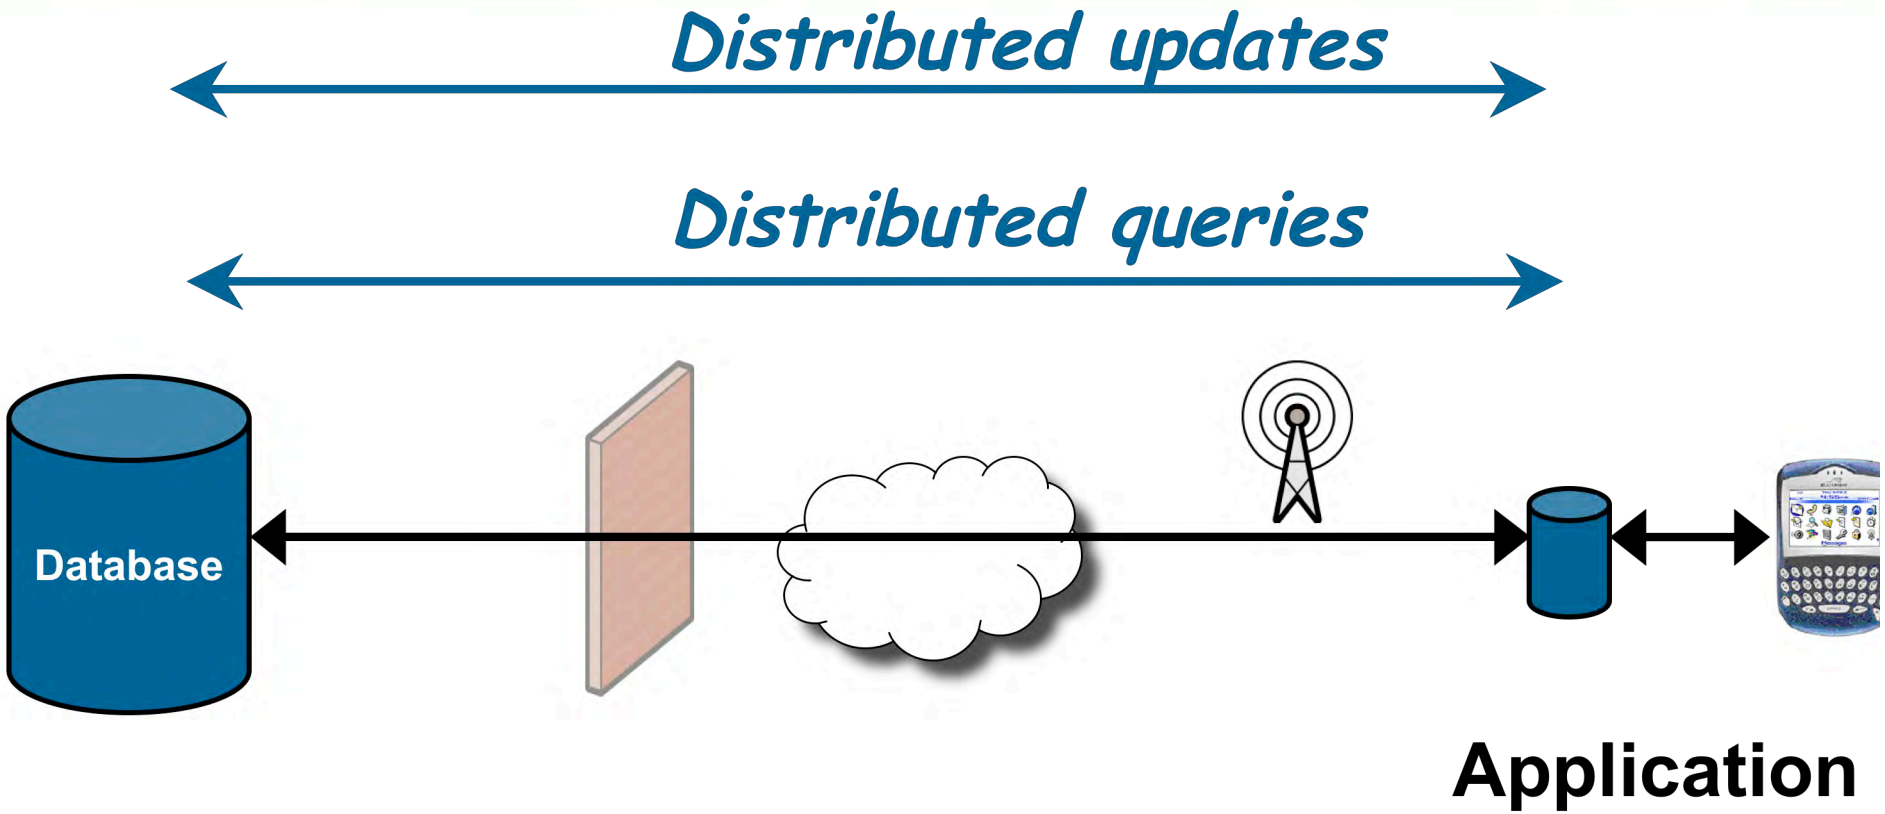

BlackBerry

Distributed Database

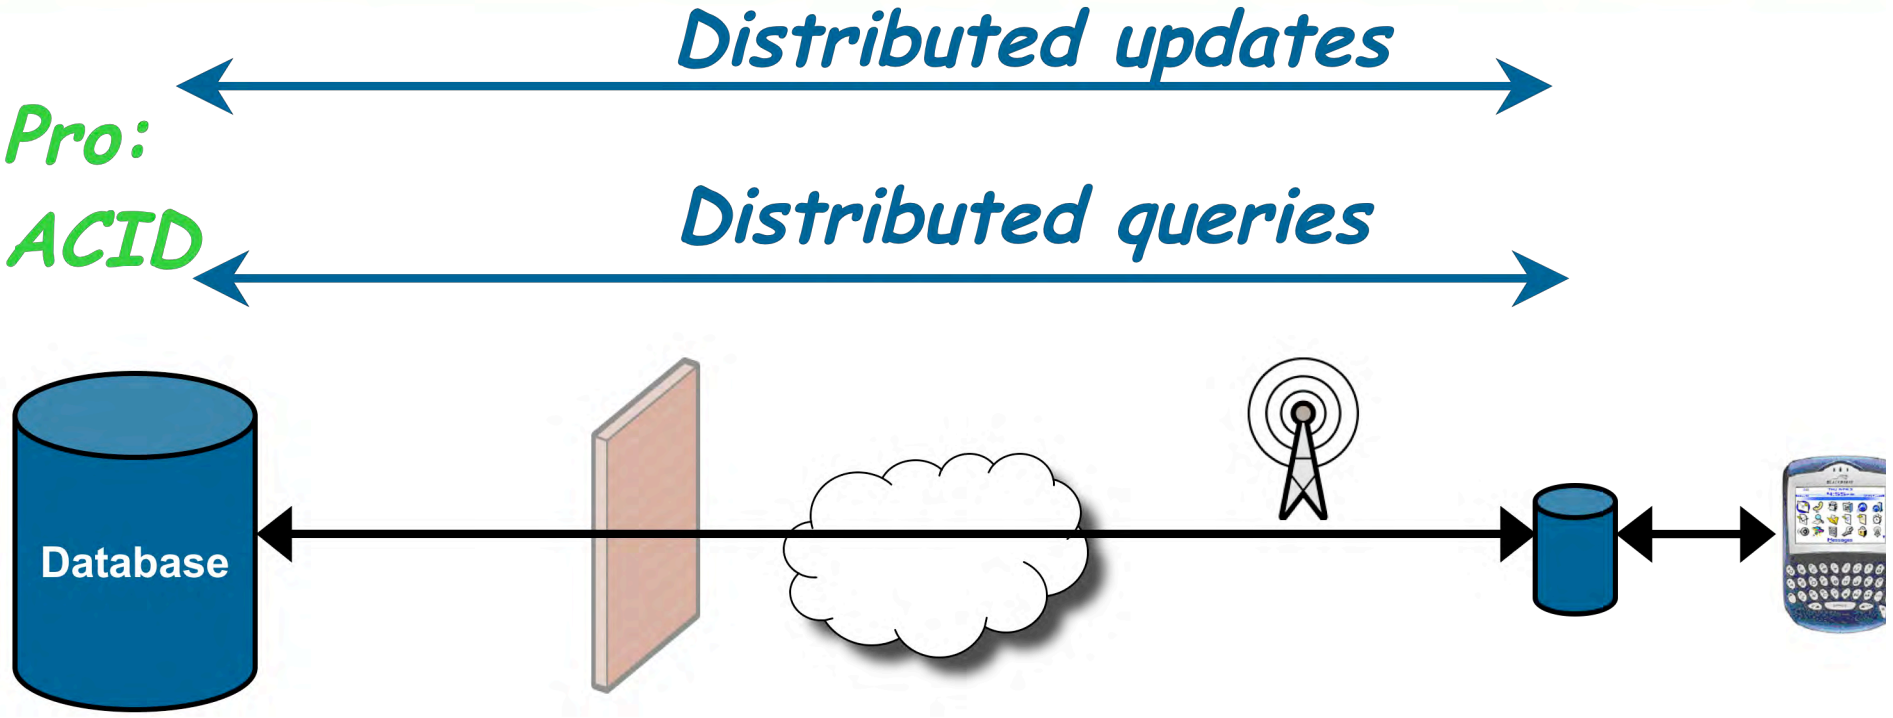

BlackBerry

Distributed Database

Distributed updates

Pro:

ACID

Distributed queries

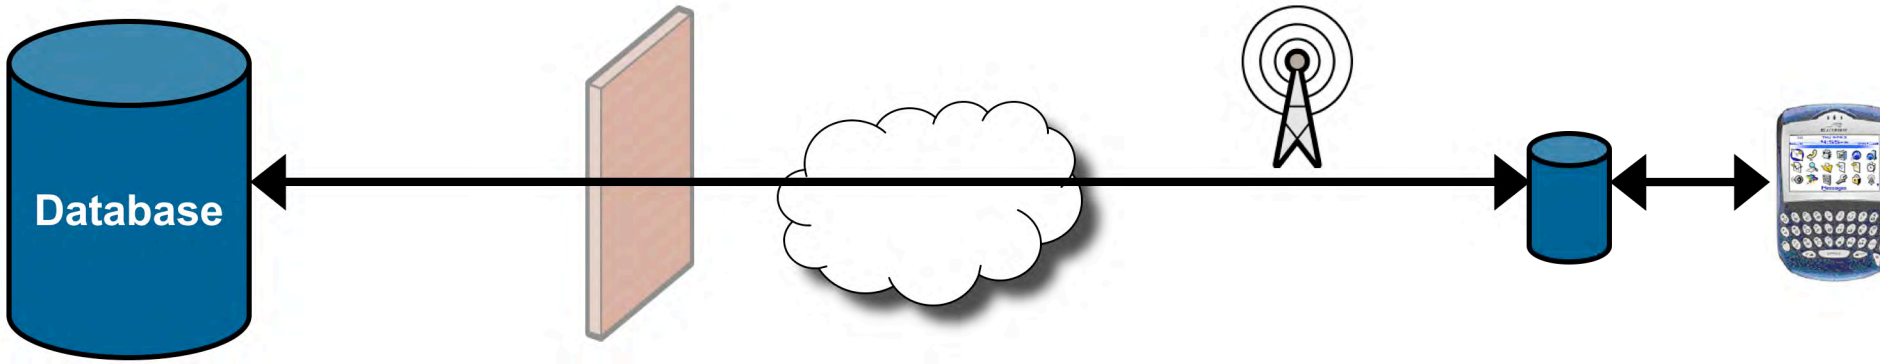

Application

Cons:

1. Data on unreliable media
2. Too many nodes
3. Problems when no coverage

Distributed Database

Distributed updates

Pro:

ACID

Distributed queries

Example: Distributed Query

1000 nodes

5 distributed queries/day

5000/day = 1 battery!

Database

Application

media

2. Too many nodes

3. Problems when no coverage

Distributed Database

Distributed updates

Pro:

ACID

Distributed queries

Example: 2PC

1000 nodes

1 node will have maximum latency = 1s

All distributed commits = 1s

Database

pplication

media

3. Problems when no coverage

Replication

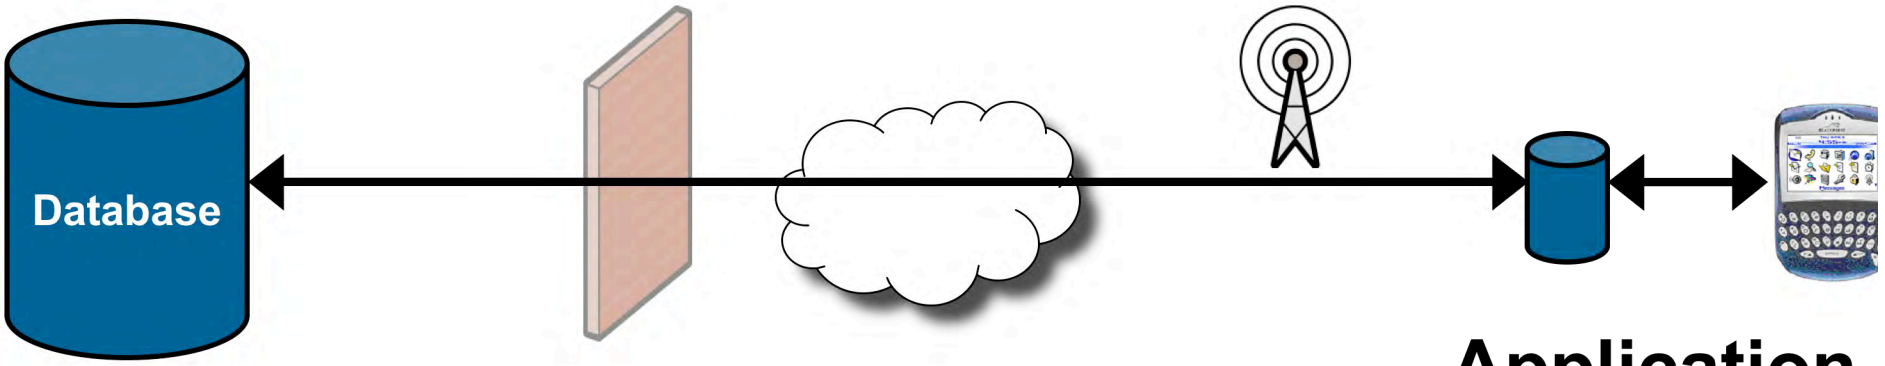

BlackBerry

Replication

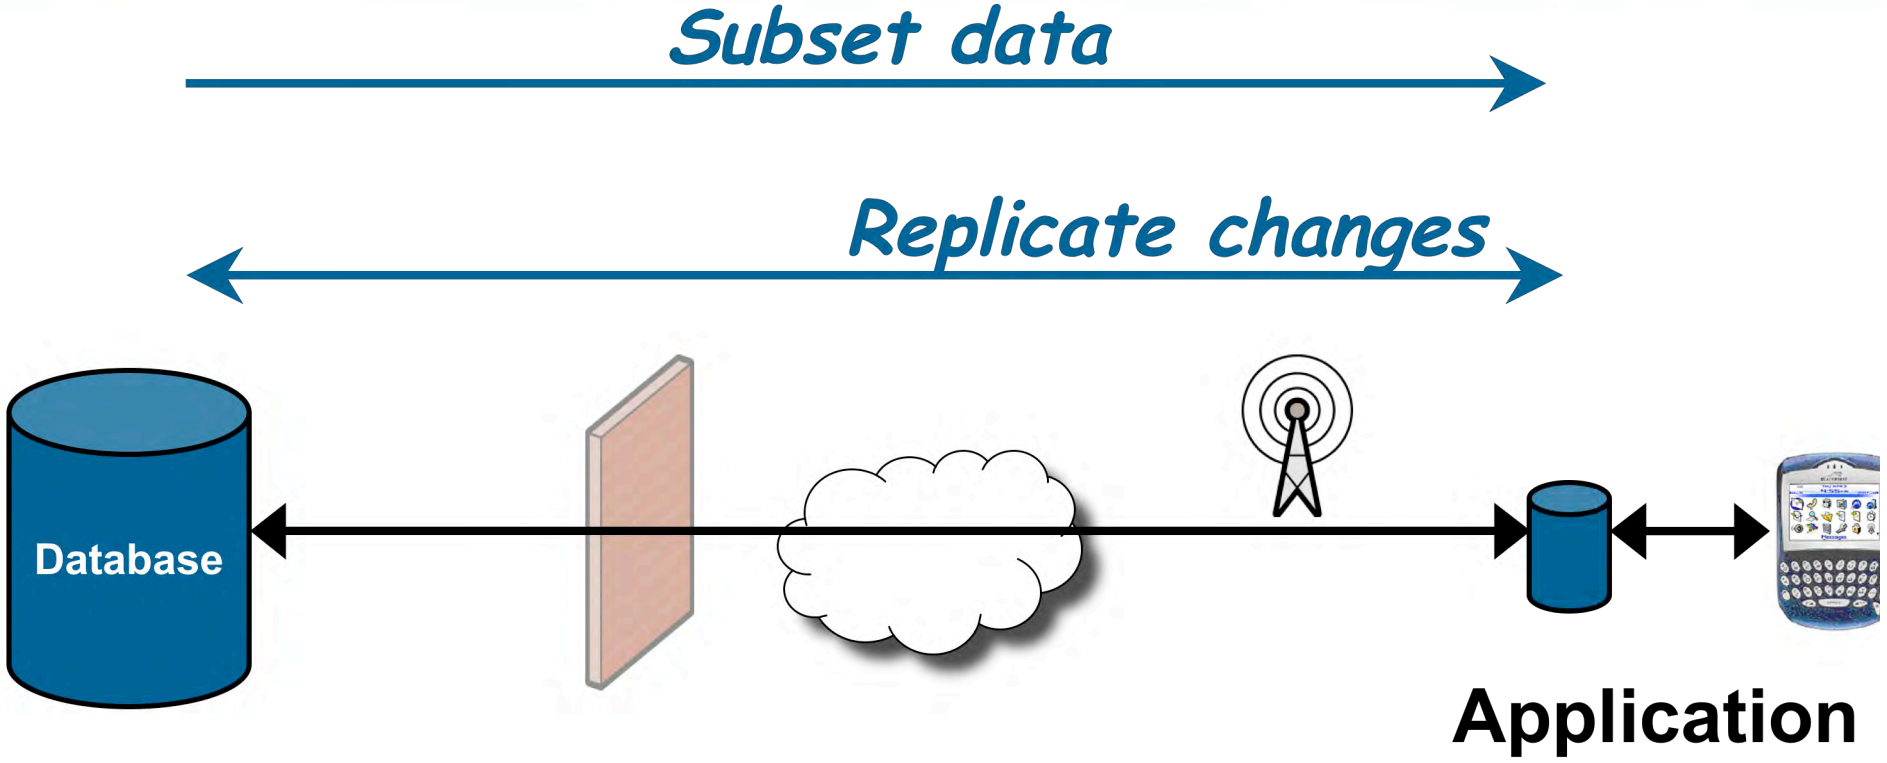

BlackBerry

Replication

Pro:

Can be push based

Minimizes data

Works out of coverage

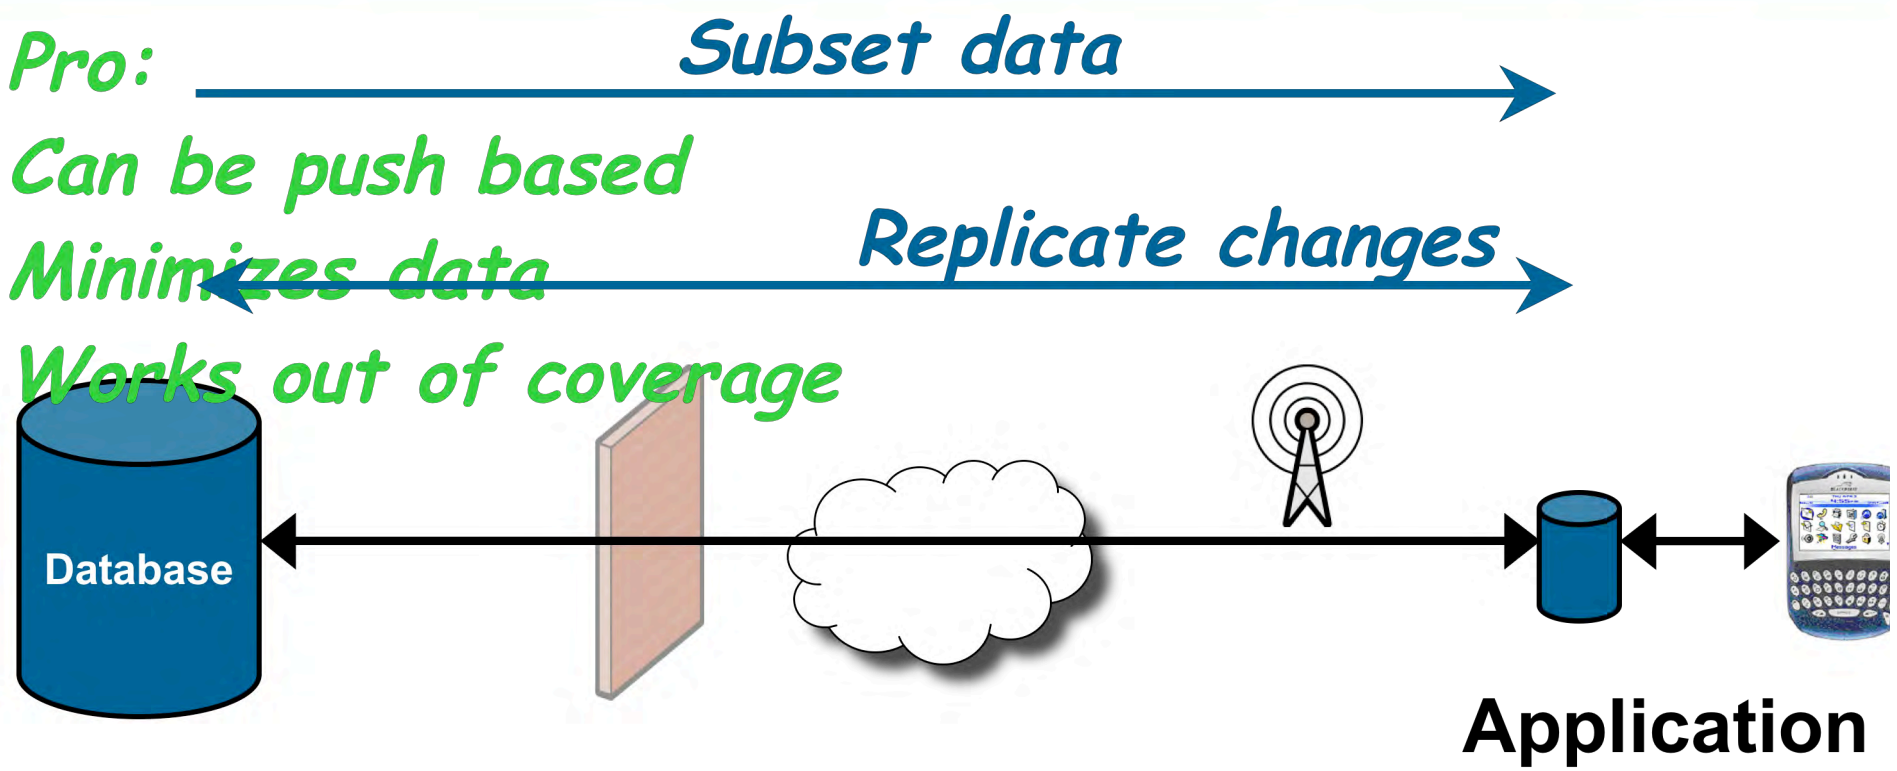

BlackBerry

Replication

Pro:

Can be push based

Minimizes data

Works out of coverage

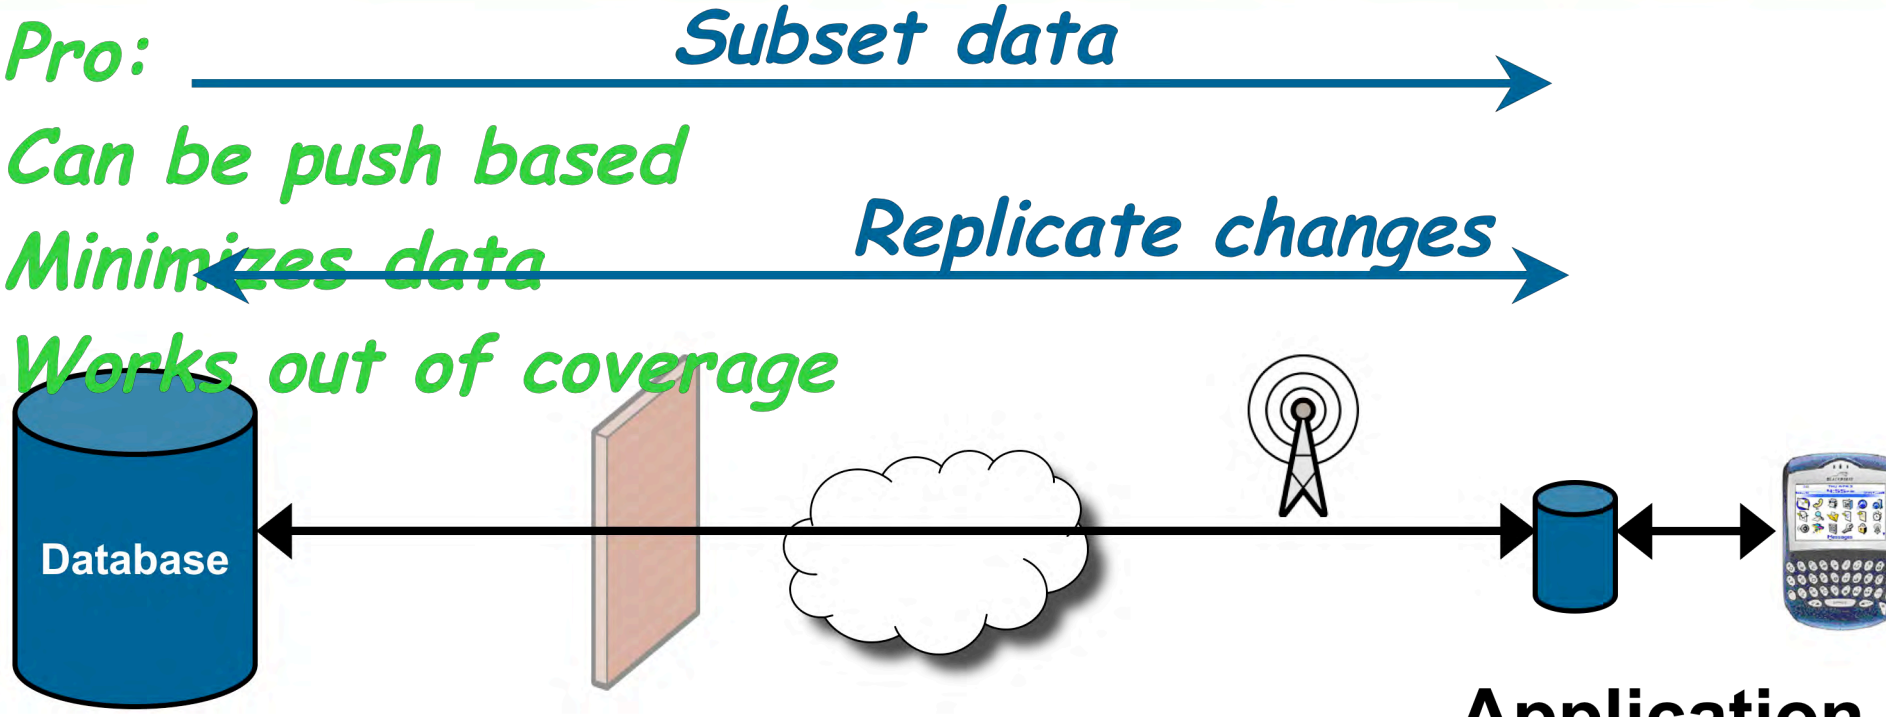

Cons:

1. Not ACID
2. Conflict Resolution
3. Device has limited capacity
4. Fine grained updates

Replication

Pro:

Can be push based

Minimizes data

Works out of coverage

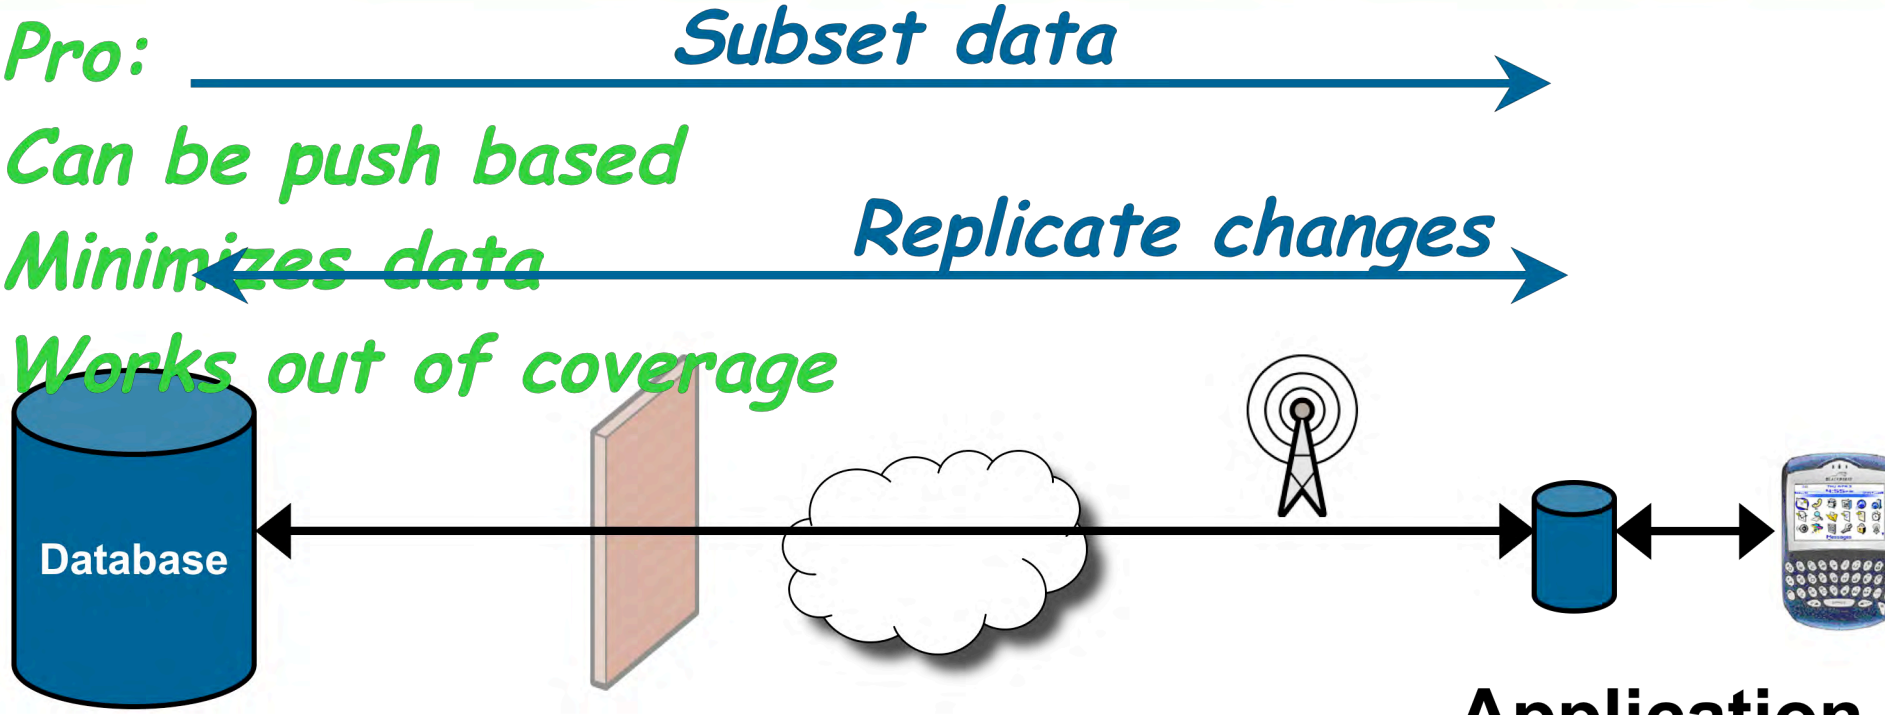

Cons:

1. Not ACID
2. Conflict Resolution
3. Device has limited capacity
4. Fine grained updates

Application specific?

Replication

Pro:

Can be push based

Minimizes data

Works out of coverage

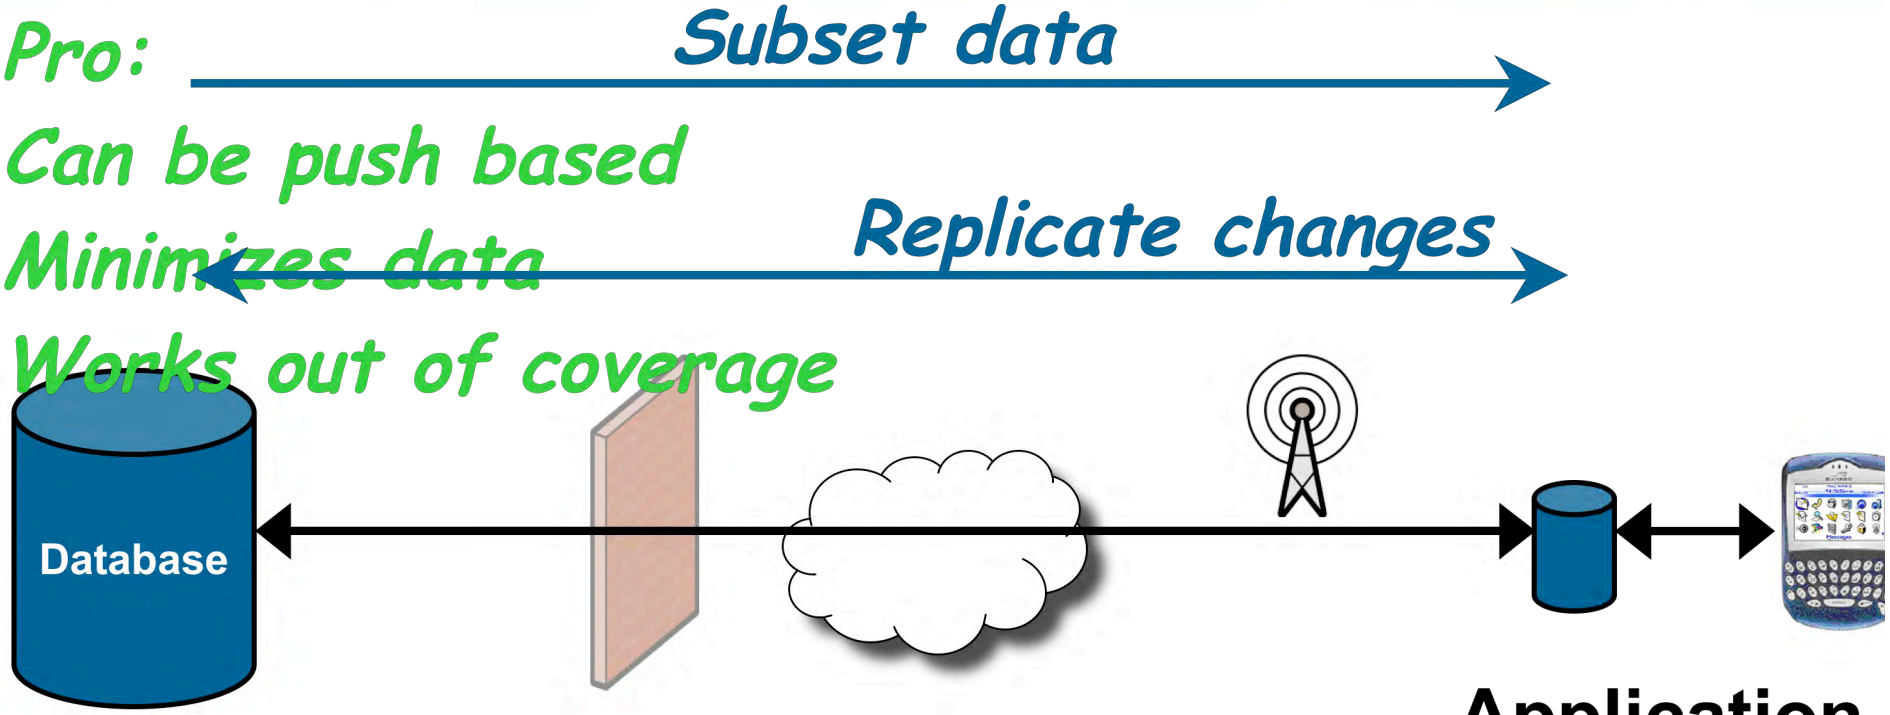

Cons:

1. Not ACID
2. Conflict Resolution
3. Device has limited capacity

Application specific?

Replication

Pro:

Can be push based

Minimizes data

Works out of coverage

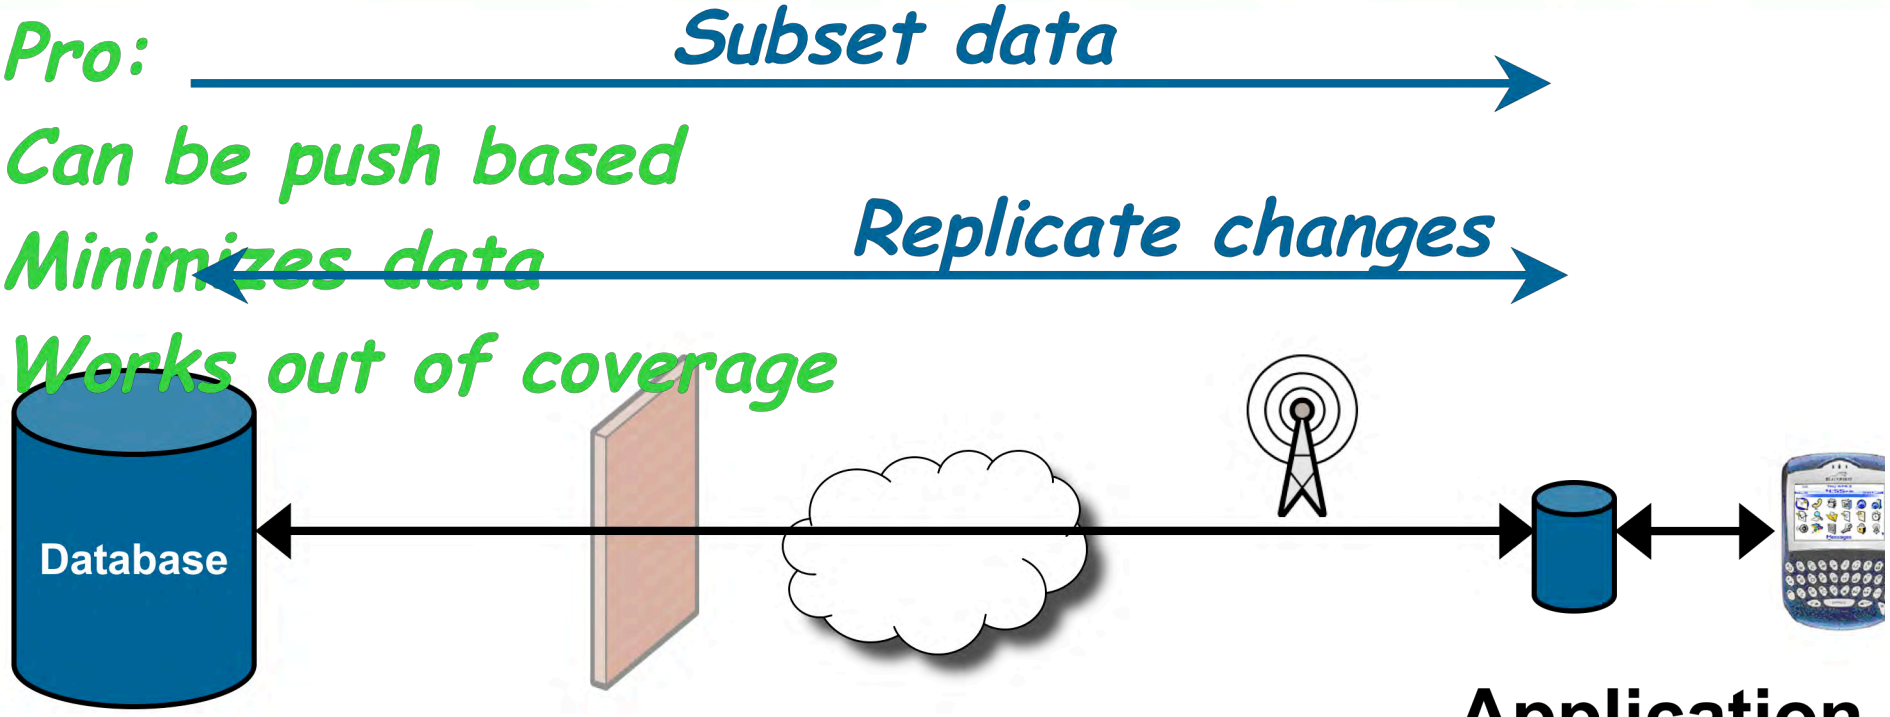

Cons:

1. Not ACID
2. Conflict Resolution
3. Device has limited capacity

Application specific?

Replication

Pro:

Can be push based

Minimizes data

Works out of coverage

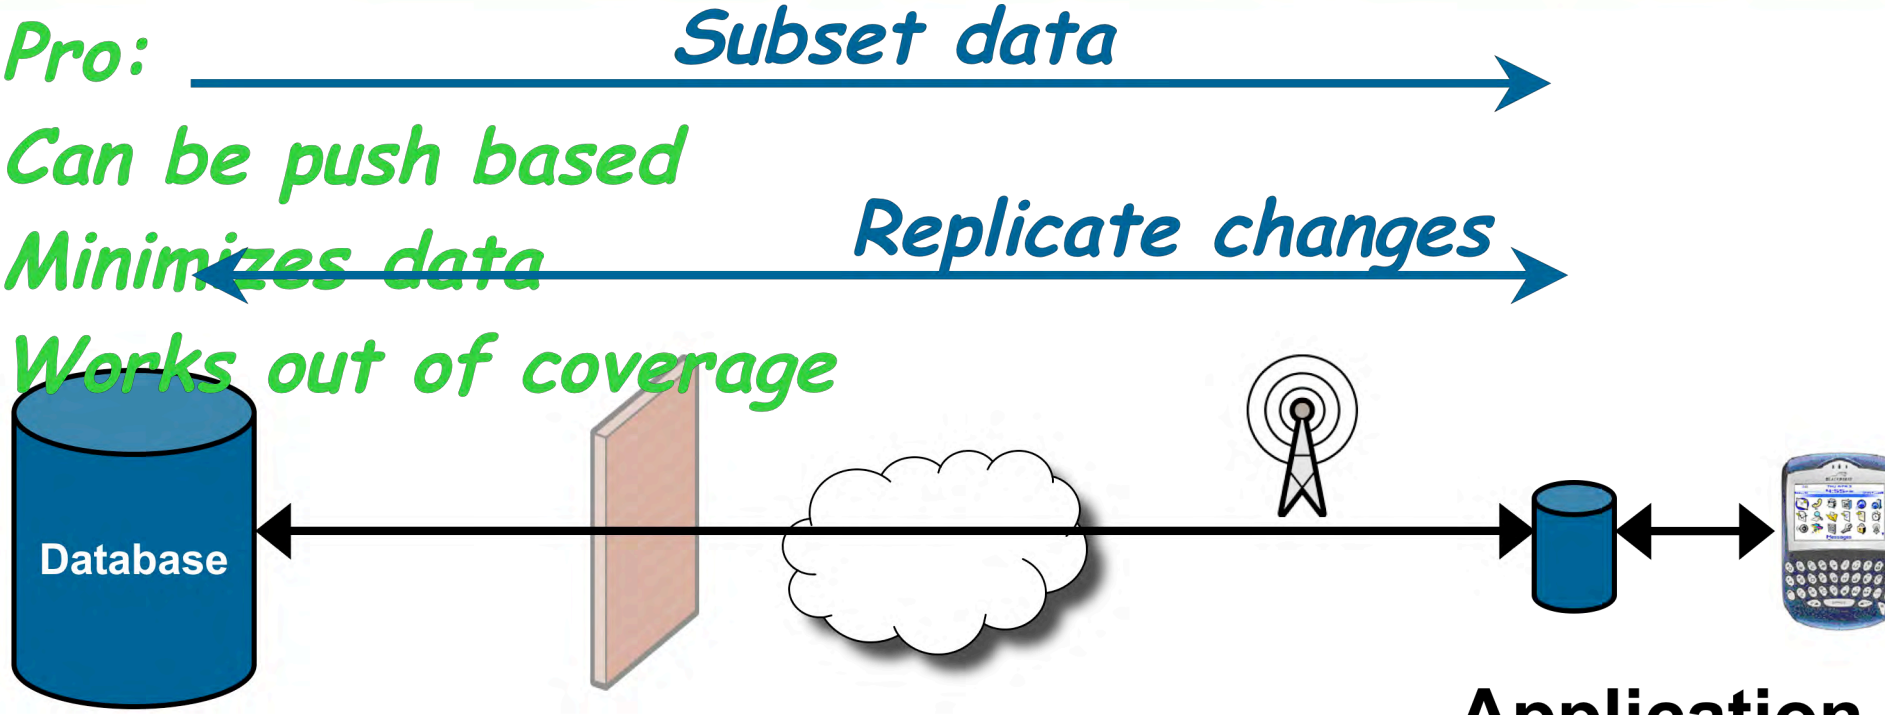

Cons:

1. Not ACID
2. Conflict Resolution
3. Device has limited capacity
4. Fine grained updates

Application specific?

Big verbs?

Replication

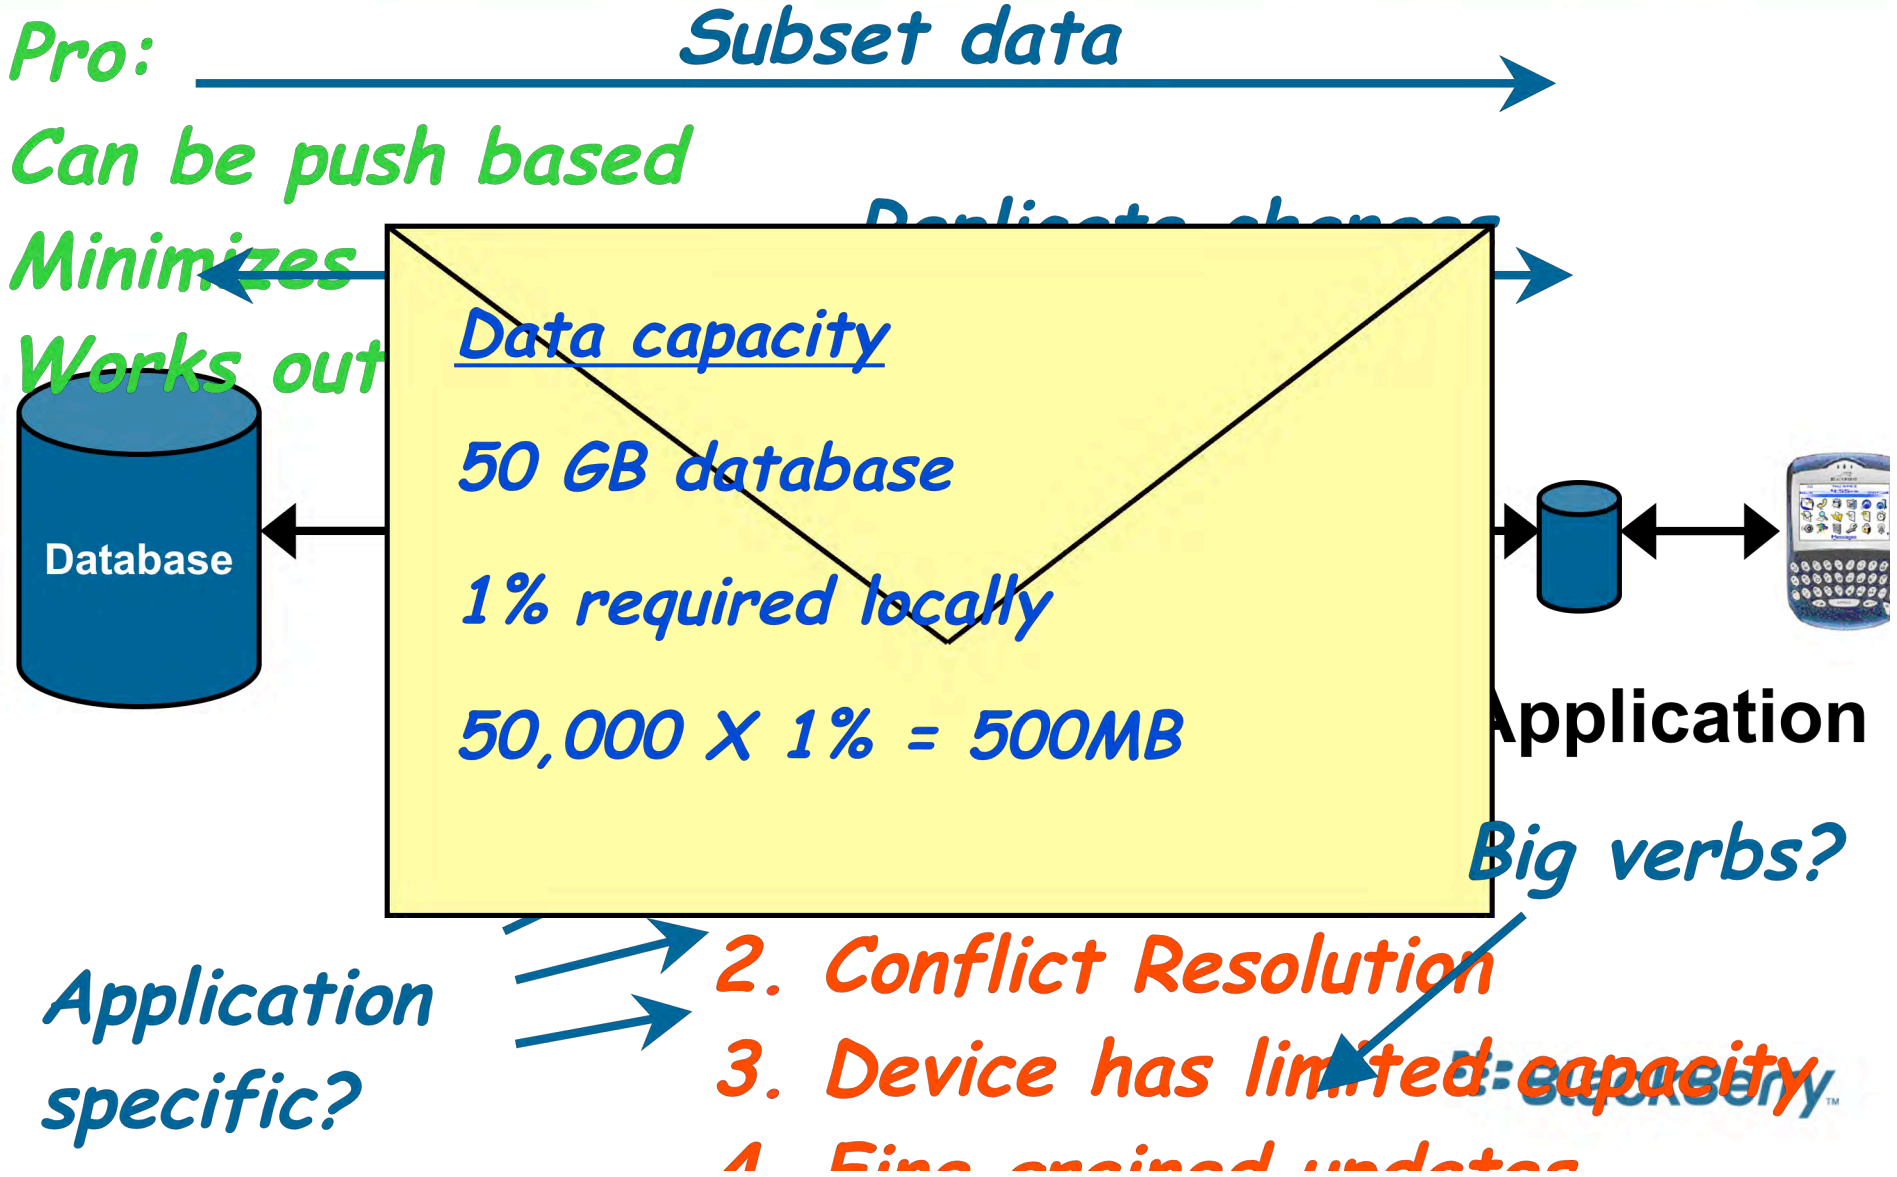

Proposed Solution Characteristics

- Push based, optimize data transfers
- Local utility when out of coverage
  - Read
  - Changes
- Data might not be on device
- Preserve integrity of Enterprise data
- Usable
- Maintainable
- Scalable
- Affordable

BlackBerry

Proposed Solution

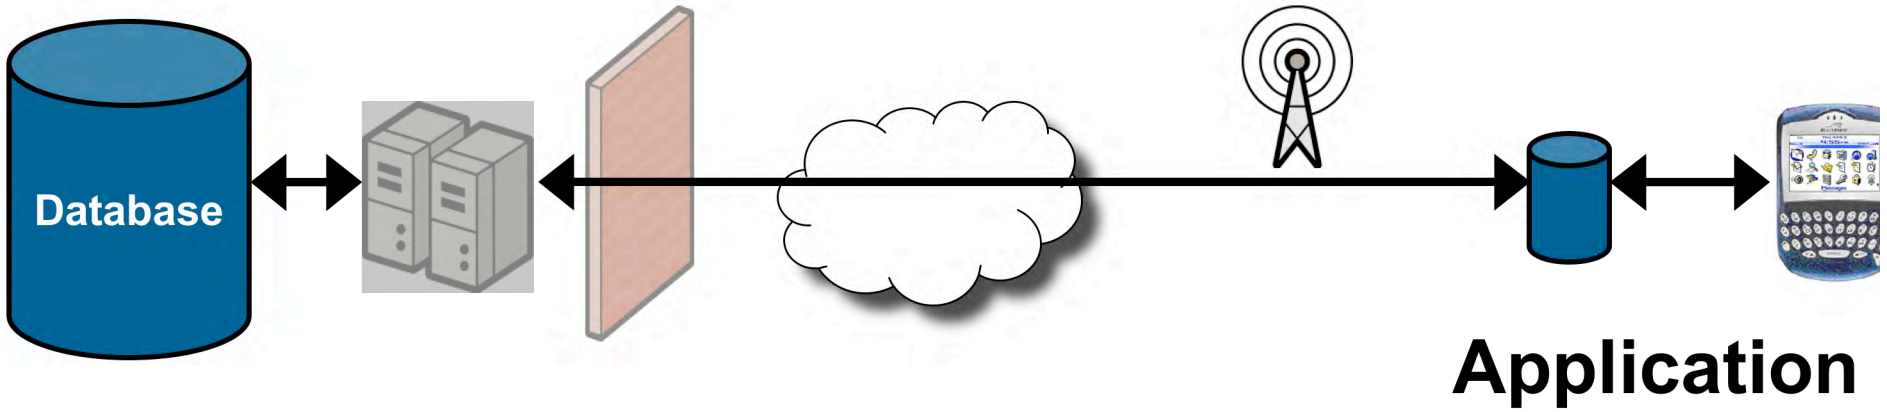

BlackBerry

Proposed Solution

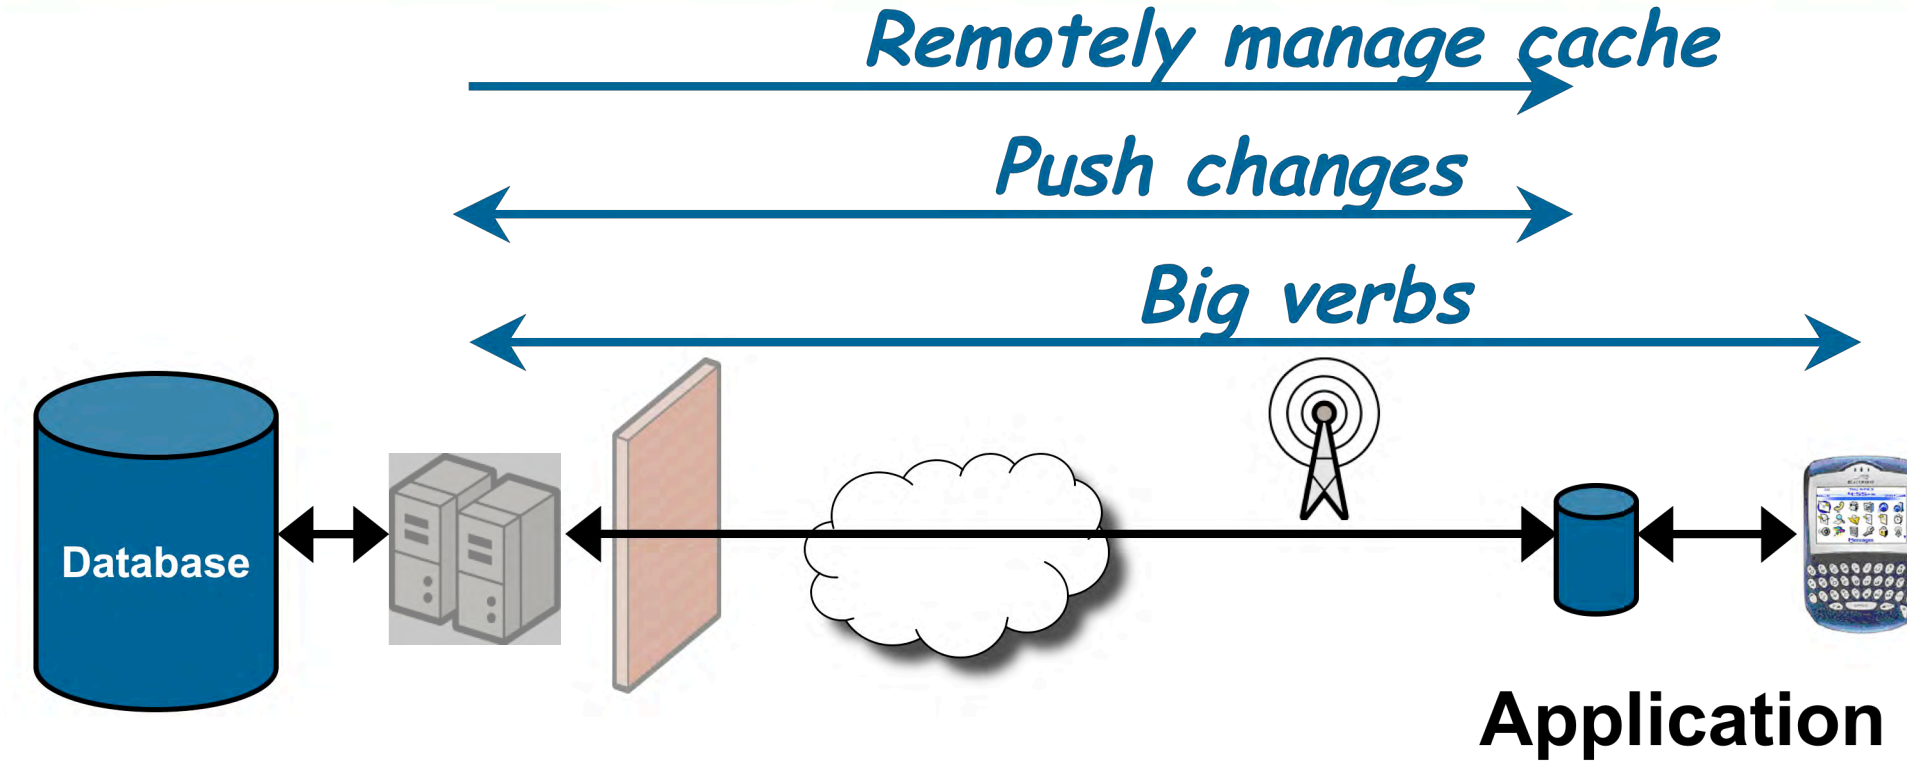

BlackBerry

Proposed Solution

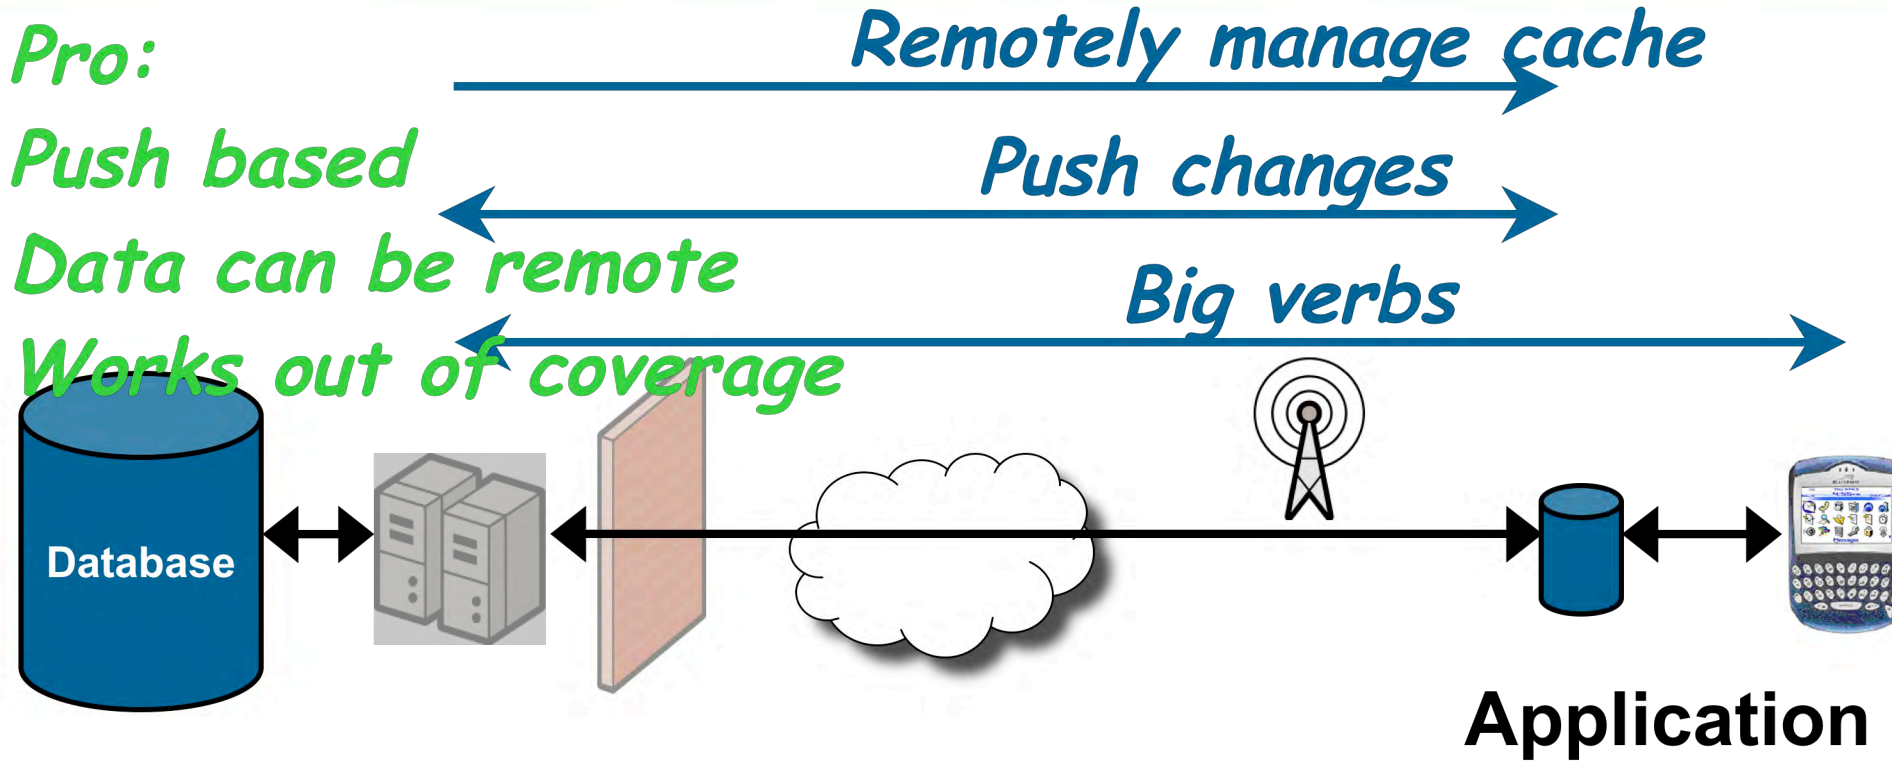

BlackBerry

Proposed Solution

Pro:

Push based

Data can be remote

Works out of coverage

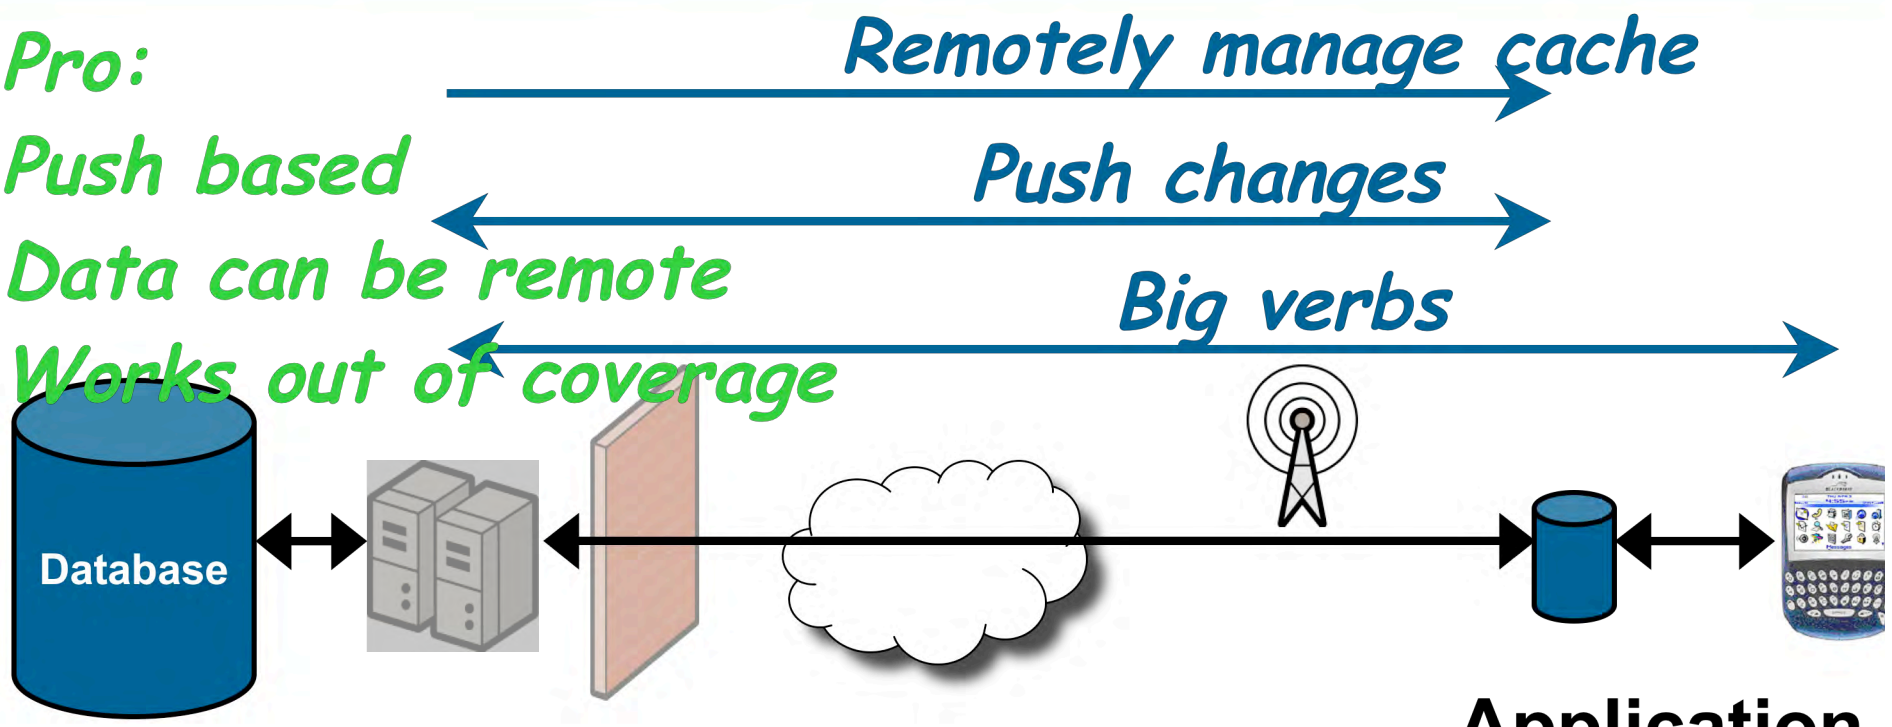

Cons:

1. Server has lots of state
2. Optimization heuristics
3. When is missing data okay?

BlackBerry

Challenges in Wireless Database World

- Server based remote cache management
- Verb aggregation
- Transaction semantics with old data
- Distributed query optimization on a wireless device

BlackBerry

Summary

BlackBerry

Summary

- Wireless is constrained in many dimensions

BlackBerry

Summary

- Wireless is constrained in many dimensions
- Push is everything, and changes everything

BlackBerry

Summary

- Wireless is constrained in many dimensions
- Push is everything, and changes everything
- Difficult to find a solution that satisfies all constraints

Summary

- Wireless is constrained in many dimensions
- Push is everything, and changes everything
- Difficult to find a solution that satisfies all constraints
- Do the math!

BlackBerry

Databases in a Wireless World

David Yach

Senior VP Software

Research In Motion Limited