



Data Infrastructure at LinkedIn

Lei Gao

<http://www.linkedin.com/in/gaolei>



Outline

- LinkedIn Products
- Data Ecosystem
- LinkedIn Data Infrastructure Solutions
- Next Play

LinkedIn By The Numbers

- 150M + users*
- ~ 4.2B People Searches in 2011**
- >2M companies with LinkedIn Company Pages**
- 16 languages
- 75% of Fortune 100 Companies use LinkedIn to hire***

* As of February 9th 2012

** As of December 31st 2011

*** As of September 30th 2011

Broad Range of Products & Services

Jobs Home | [Saved Jobs](#) | [Saved Searches](#) | [Advanced Search](#) | [Refer Connections](#)

[Sign up for the free webinar "Job Seeking on LinkedIn" for tips a](#)

Search for Jobs

Job Title, Keywords, or Company Name

Search

[Advanced](#)

Jobs you may be interested in **Beta** [Email Alerts](#) | [See More »](#)

- Software Architect**
Intuit - Menlo Park, CA
- Software Architect**
Nomis Solutions - San Bruno, CA
- Software Architect w/ Reverse Pr**
AT&T Interactive - San Francisco
- Network Software Architect**
AGCM Research - San Francisco
- Software Architect - MVC - Cloud**
CyberCoders - San Francisco, CA

Kara Swisher
[archives](#) | [ethics statement](#) | [bio](#)

Viral Video: Like P Lady, TouchPad's Get the HP Boot

AUGUST 19, 2011 AT 12:13 AM PT

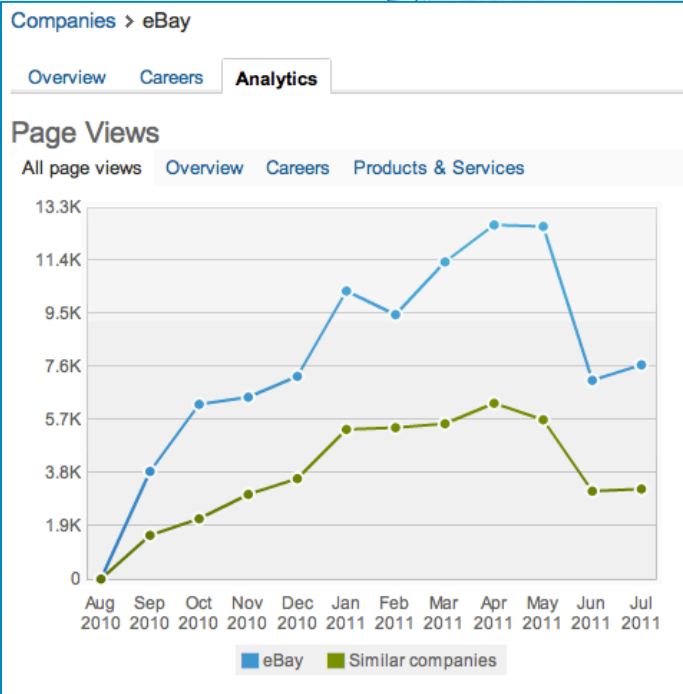
[Tweet](#) | [Like](#) | [Share](#) | [Share](#) | [Print](#)

INTERNET TODAY • AU

The most shared news by people in the Internet Industry. Take th

[FRONT PAGE](#) | [INTERNET](#) | [MARKETING &...](#)

Skype To Acquire Year-old G GroupMe



Inbox 92 | [Companies](#) | [News](#) | [More](#) | [People](#) | [Search...](#)

Invitations (92)

[Add Connections »](#)

- Wendy Pfeiffer**
Sr. Director, Enterprise Shared Services
- Alex Schwartz**
Site Reliability Engineer at LinkedIn
- Tony Ng**
Director of Architecture at eBay
Tony has indicated you are a Colleague at Yahoo!

Messages (0)

[Compose Message »](#)

- Christopher Brereton**
(Responded)
Hey Tom!
- Kevin Krawez**
Your call
- Aaron Tran**
VMware seeking references for Staff/Principal SW-Systems Management...

Neal Sample 1st

CTO - Open Commerce at eBay
San Francisco Bay Area | Internet

Current **CTO - X.commerce at eBay**

Past VP of Architecture, Technology Product Management, Developers Program at eBay
VP Social, Open and Participation Platforms at Yahoo!
Chief Architect (VPE/Fellow) at Yahoo!

[see all »](#)

Education Stanford University
University of Wyoming
Stanford University

Recommendations 33 people have recommended Neal

Connections 500+ connections

Public Profile <http://www.linkedin.com/in/nealsample>

[Share](#) | [PDF](#) | [Print](#) | [vCard](#) | [Flag](#)


User Profiles

LinkedIn Account Type: Pro

Home Profile Contacts Groups Jobs Inbox 92 Companies News More

People Search...

[Earn Your M.S. in Finance - 15 month program designed for busy working professionals. Click for info!](#) From: Saint Mary's College of California




Jeff Weiner
CEO at LinkedIn
Mountain View, California | Internet

1st

Suggest Connections

[Send Jeff a message](#)
[Suggest a profile update for Jeff](#)
[Save Jeff's Profile](#)

Jeff Weiner @dtunkelang adds great examples to recent talk re: the importance of understanding mutual fit when recruiting.



Dream. Fit. Passion. thenoisychannel.com · via Daniel Tunkelang
A few days ago, our CEO Jeff Weiner led a session at LinkedIn on how to "close" candidates — that is, how to persuade candidates to join your team once you have found and interviewed them. Since not everyone has the opportunity...

5 hours ago · Like (1) · Comment · Send a message · Share · See all activity

Current **CEO at LinkedIn**
Member, Board of Directors at DonorsChoose
Member, Board of Directors at Malaria No More

Past Executive in Residence at Accel Partners
Executive in Residence at Greylock
Executive Vice President Network Division at Yahoo!
[see all](#)

Education University of Pennsylvania - The Wharton School

Recommendations 8 people have recommended Jeff

Connections 500+ connections

Websites [Company Website](#)

Twitter [Follow](#) @jeffweiner

Public Profile <http://www.linkedin.com/in/jeffweiner08>


[Share](#) [PDF](#) [Print](#) [vCard](#) [Flag](#)

LinkedIn Premium

Get Hired Faster
with Job Seeker Premium

- Get noticed with a Job Seeker Badge
- Move to the top as a Featured Applicant
- Contact recruiters directly with InMail


[Learn More](#)



Jeff's Activity

Jeff Weiner likes this update:


Mike Gamson Another step in the right direction for the Khan Academy. . .



Khan Academy Integrates With Digital Textbooks mashable.com
The 12-minute lectures that Bill Gates has called "the start of a revolution" will now be linked with the material in some digital textbooks.

4 hours ago · Like (4) · Comment · Share

Jeff Weiner @dtunkelang adds great examples to recent talk re: the importance of understanding mutual fit when recruiting.



Dream. Fit. Passion. thenoisychannel.com · via Daniel Tunkelang
A few days ago, our CEO Jeff Weiner led a

Large dataset
Medium writes
Very high reads
Freshness <1s

Communications

The screenshot displays the LinkedIn inbox interface for a 'Pro' account. The top navigation bar includes links for Home, Profile, Contacts, Groups, Jobs, Inbox (91), Companies, News, and More. A search bar is located on the right. The left sidebar contains a 'Search Inbox' field, a 'Compose Message' button, and a list of folders: Inbox, Sent, Archived, and Trash. The main content area shows a list of messages under the 'Messages' tab, with an 'Invitations' tab (91) also visible. The messages list includes columns for a checkbox, a profile picture, the sender's name, the subject, the date, and action links (Reply, Forward, Archive, Delete). The messages are sorted by 'Newest'.

| Select | None | All Messages | Newest |
|--------------------------|--|--------------|--------|
| <input type="checkbox"/> | Christopher Brereton (Responded) | Jun 14 | |
| | Hey Tom! | | |
| | Reply • Forward • Archive • Delete | | |
| <input type="checkbox"/> | Kevin Krawez | Apr 19 | |
| | Your call | | |
| | Reply • Forward • Archive • Delete | | |
| <input type="checkbox"/> | Aaron Tran | Mar 30 | |
| | VMware seeking references for Staff/Principal SW-Systems Management developers | | |
| | Reply • Forward • Archive • Delete | | |
| <input type="checkbox"/> | Lynne Peters (Declined) | Mar 29 | |
| | NoSQL Architect Opportunity | | |
| | Reply • Forward • Archive • Delete | | |
| <input type="checkbox"/> | Courtney Gartin - We're Hiring!!! (Declined) | Mar 24 | |
| | Your experience is AMAZING! Great Opportunities & Your Expertise Requested - 7907 | | |
| | Reply • Forward • Archive • Delete | | |
| <input type="checkbox"/> | Catherine Porter | Mar 18 | |
| | Catherine Porter invites you to join LI for Education (LIFE) on LinkedIn | | |
| | Reply • Forward • Archive • Delete | | |
| <input type="checkbox"/> | Mark Gambirasi (Declined) | Mar 16 | |
| | Unique, Client-Facing Systems Engineer Position with Rising Star Cloud Computing Company | | |
| | Reply • Forward • Archive • Delete | | |

Large dataset
High writes
High reads
Freshness <1s

People You May Know







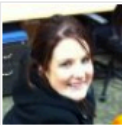

LinkedIn® Go to LinkedIn Recruiter »

Home Profile Contacts Groups Jobs Companies News More

People ▾ Search...

People You May Know ^{beta} See people from different parts of your professional life

ORACLE Indian Institute of Technology, LinkedIn Stanford University Microsoft University of California, Berkeley YAHOO! Google

| | |
|---|--|
|  <p>Peter Carlin (2nd) at Microsoft Greater Seattle Area</p> <p>Connect 5 shared connections</p> |  <p>Dmitri Smirnov (2nd) at Microsoft San Francisco Bay Area</p> <p>Connect 7 shared connections</p> |
|  <p>Josh Walker (2nd) Principal Software Engineer at LinkedIn San Francisco Bay Area</p> <p>Connect 18 shared connections</p> |  <p>Subrahmanyam Sanka (Subbu) DW/BI at LinkedIn San Francisco Bay Area</p> <p>Connect 8 shared connections</p> |
|  <p>Walt Roloson (2nd) Business Operations Manager, Investor San Francisco Bay Area</p> <p>Connect 2 shared connections</p> |  <p>Vinoth Chandar (2nd) Engineer, Distributed Data Systems at LinkedIn San Francisco Bay Area</p> <p>Connect 17 shared connections</p> |
|  <p>Alison King (2nd) HR assistant at LinkedIn San Francisco Bay Area</p> <p>Connect 32 shared connections</p> |  <p>Ricky Ho (2nd) Principal Architect San Francisco Bay Area</p> <p>Connect 5 shared connections</p> |

Large dataset
Compute intensive
High reads
Freshness ~hrs


3X POINTS ON
AIRFARE

2X POINTS ON
GAS & GROCERIES

1X POINT ON
EVERYTHING ELSE

EARN 25,000
MEMBERSHIP REWARDS®
POINTS

\$0 INTRO ANNUAL FEE
FOR THE FIRST YEAR



[APPLY NOW](#)

Terms, Conditions, and Restrictions apply.

LinkedIn Today

LINKEDIN TODAY • MAR 30

The most shared news on LinkedIn [Take the tour »](#)

Powered by [twitter](#) [LinkedIn](#) [Google](#)

FRONT PAGE INTERNET IT WALL STREET... MONEY.CNN.CO... NPR.ORG BLOOMBERG.CO... MORE FOLLOWING

Blackberry maker posts \$125m loss



bbc.co.uk

[Share](#) [Save](#)

RIM's Playbook has suffered from poor sales Blackberry manufacturer Research in Motion (RIM) has reported a quarterly loss, due in part to falling revenues on the back of weak smartphone.


[Trending in Nonprofit Organization Management, Management Consulting, Information Technology and Services and Retail](#)




Five Career Lessons From Han Solo



The Cloud Will Cure Cancer



Blackberry maker posts \$125m loss




Facebook Delves Deeper Into Search



INTERNET INDUSTRY [SEE ALL »](#)


searchengineland.com

Are You In Danger Of An Online Traffic Crisis?



techcrunch.com

TechCrunch | Amazon's Appstore Generates More...



Will Pending Layoffs Put Final Nail In Yahoo Search?

searchengineland.com
77 shares



3 Reasons Facebook Brand Pages Are Good for Businesses

mashable.com
359 shares



Local Facebook pages outperform corporate pages in terms of reach, engagement percentage, study finds

incidefacebook.com

Search any topic

TOP SOURCES

WSJ Wall Street...

[Unfollow](#)

money.cnn.com »

[Unfollow](#)

npr.org »

[Unfollow](#)

bloomberg.com »

[Unfollow](#)

hbr.org »

[Unfollow](#)

nytimes.com »

[Follow](#)

[See all »](#)

SUGGESTED INDUSTRIES

Management...

Popular on LinkedIn

[Follow](#)

Marketing &...

Popular on LinkedIn

[Follow](#)

Online Media »


Popular on LinkedIn

[Follow](#)

[See all »](#)

News on the go

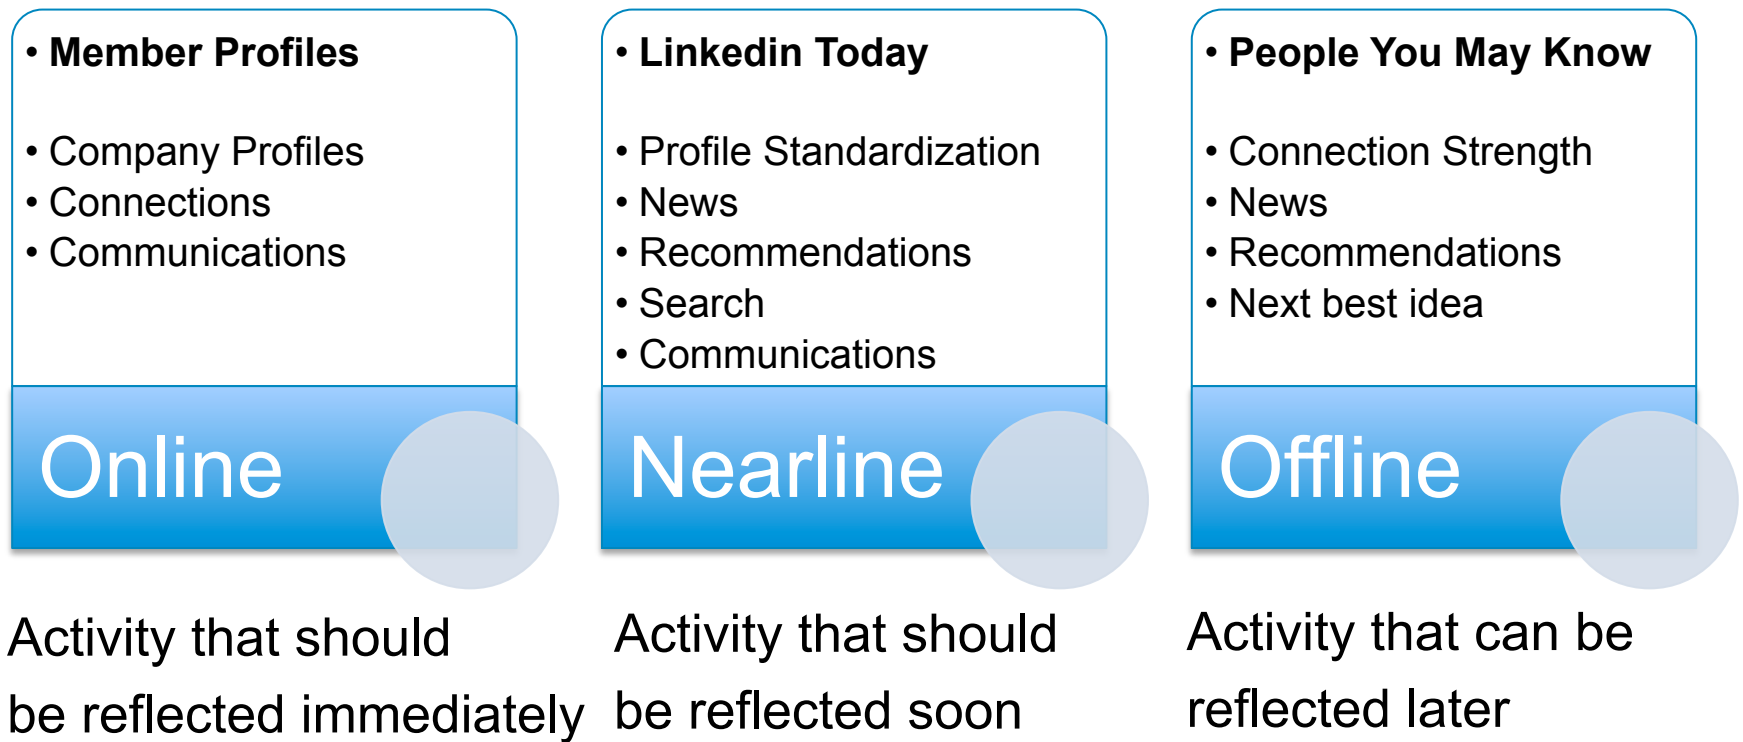
Get the new mobile app for iPhone and Android



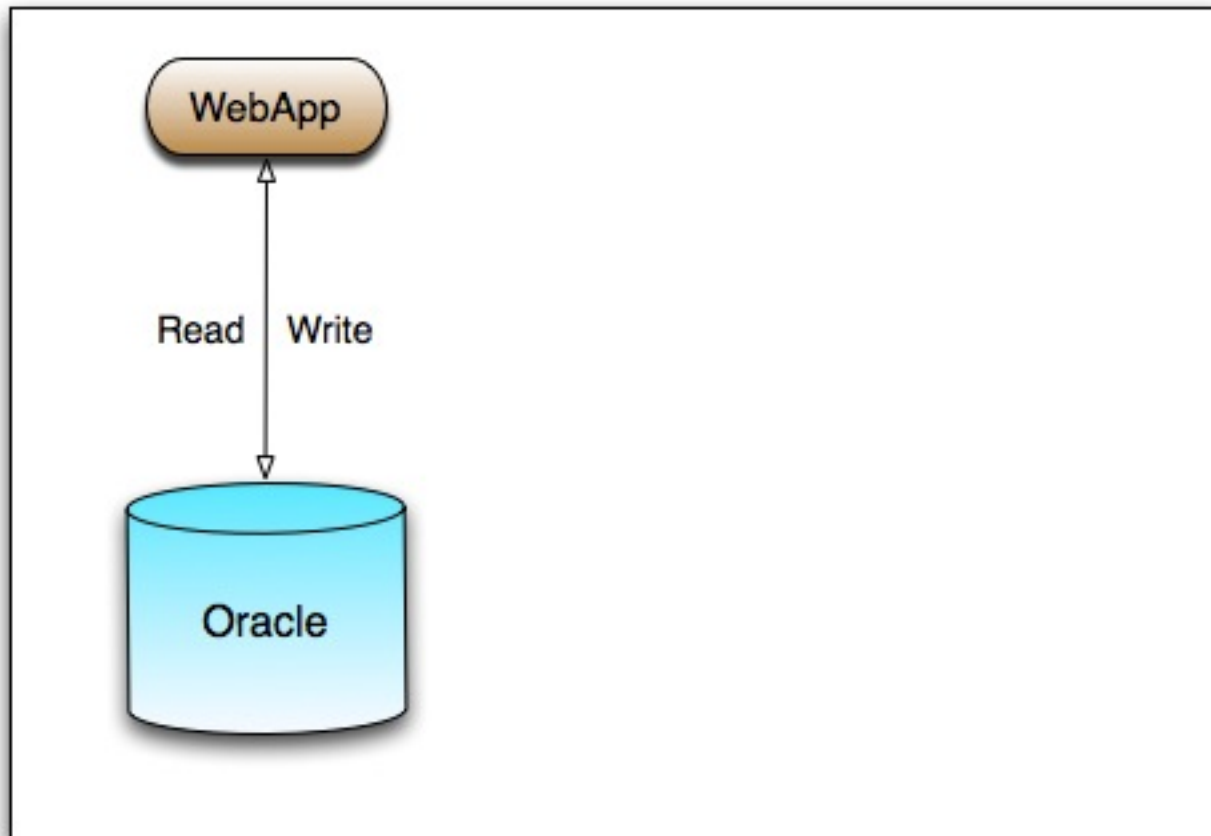
Outline

- LinkedIn Products
- **Data Ecosystem**
- LinkedIn Data Infrastructure Solutions
- Next Play

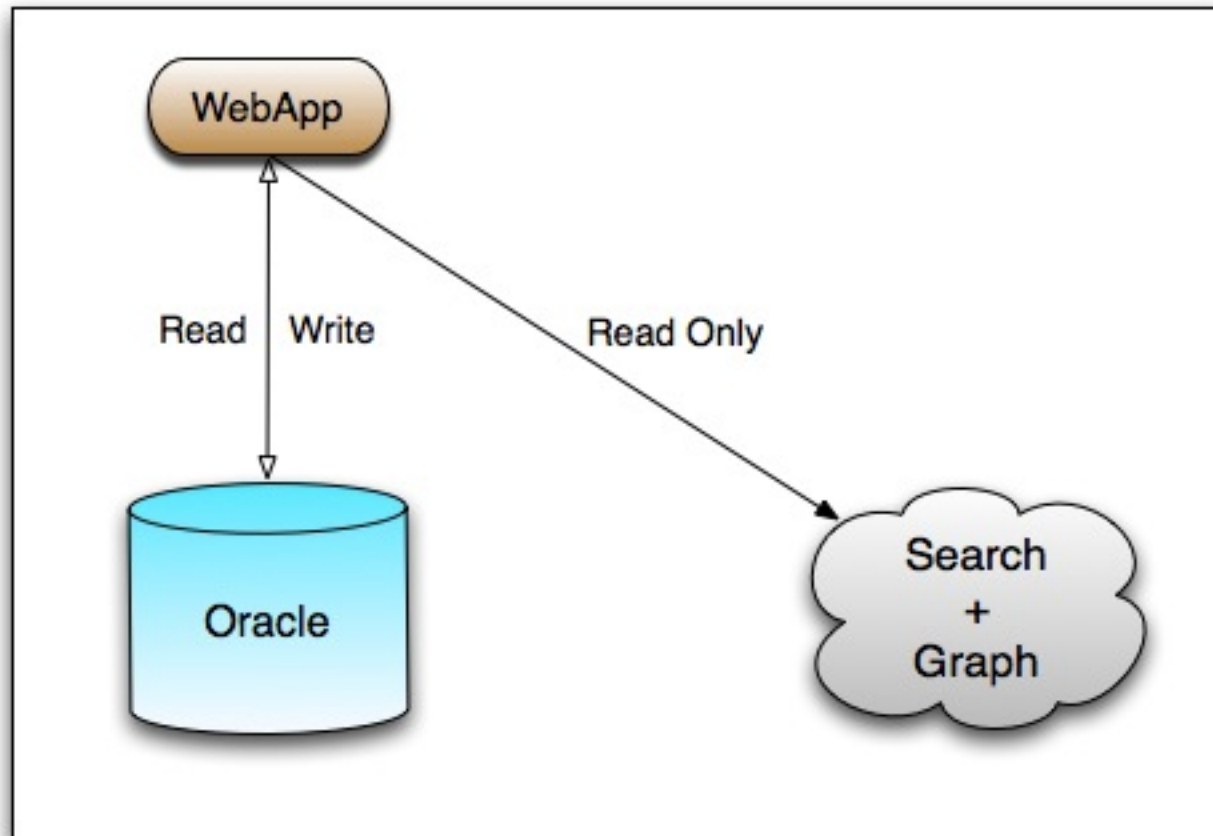
Three Paradigms : Simplifying the Data Continuum



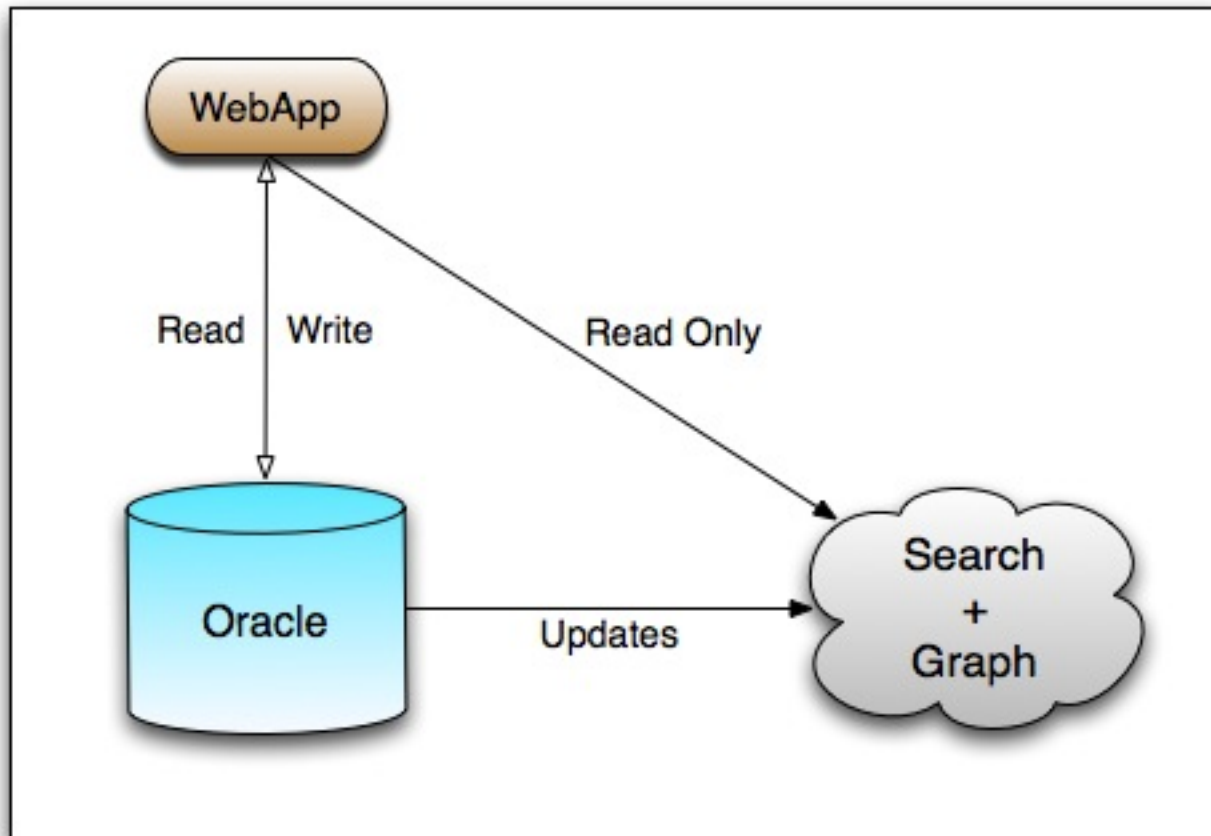
LinkedIn Product Architecture



LinkedIn Product Architecture



LinkedIn Product Architecture

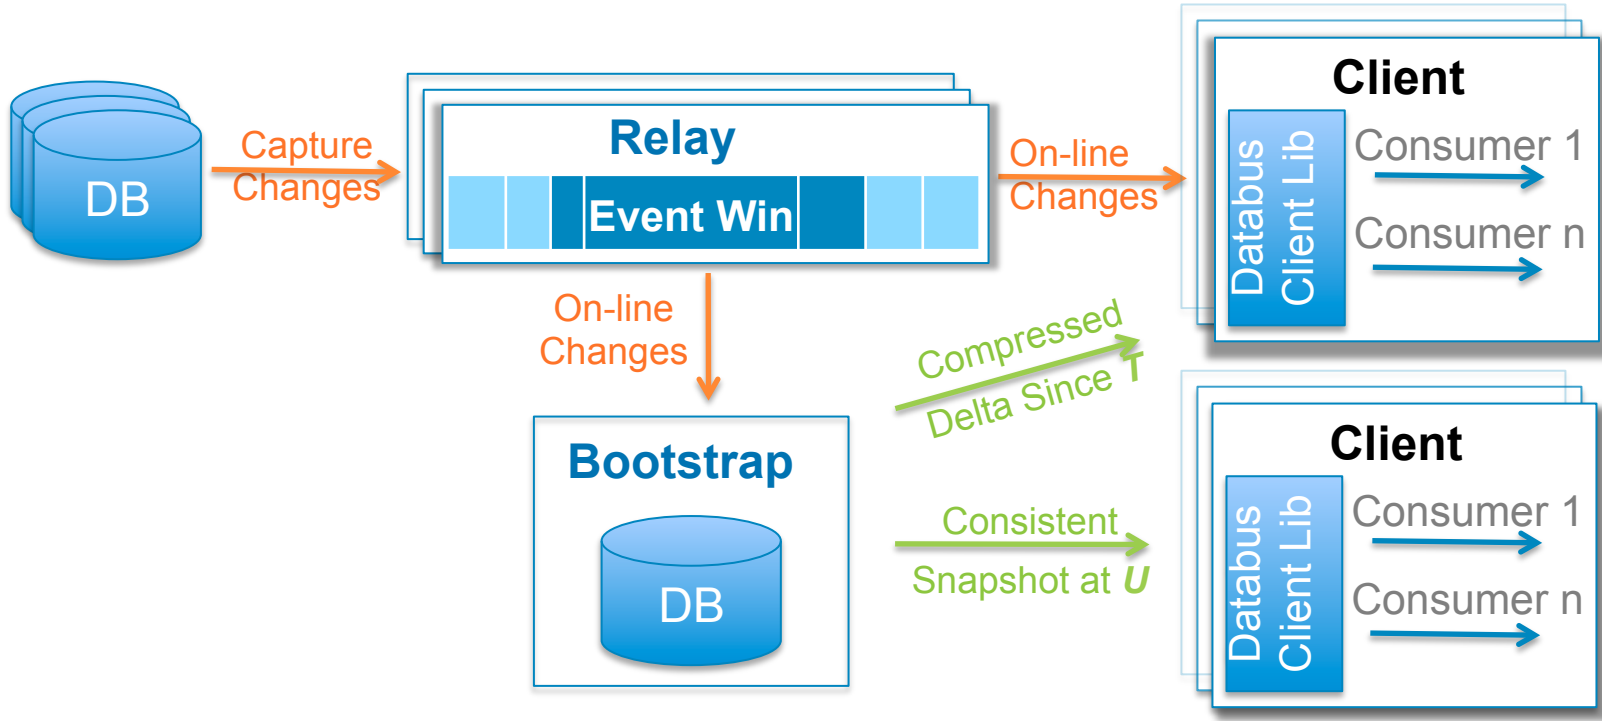


LinkedIn Data Infrastructure Solutions


Databus : Timeline-Consistent Change Data Capture



Databus at LinkedIn

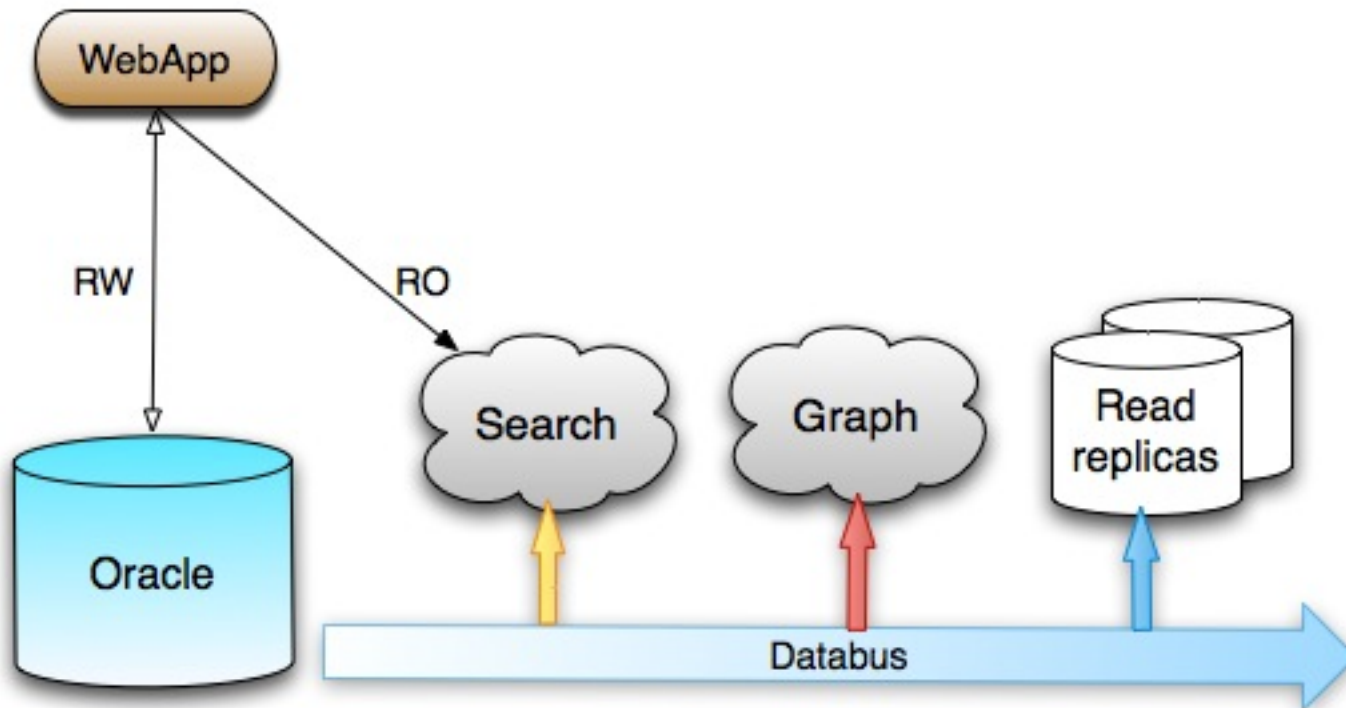


Databus at LinkedIn

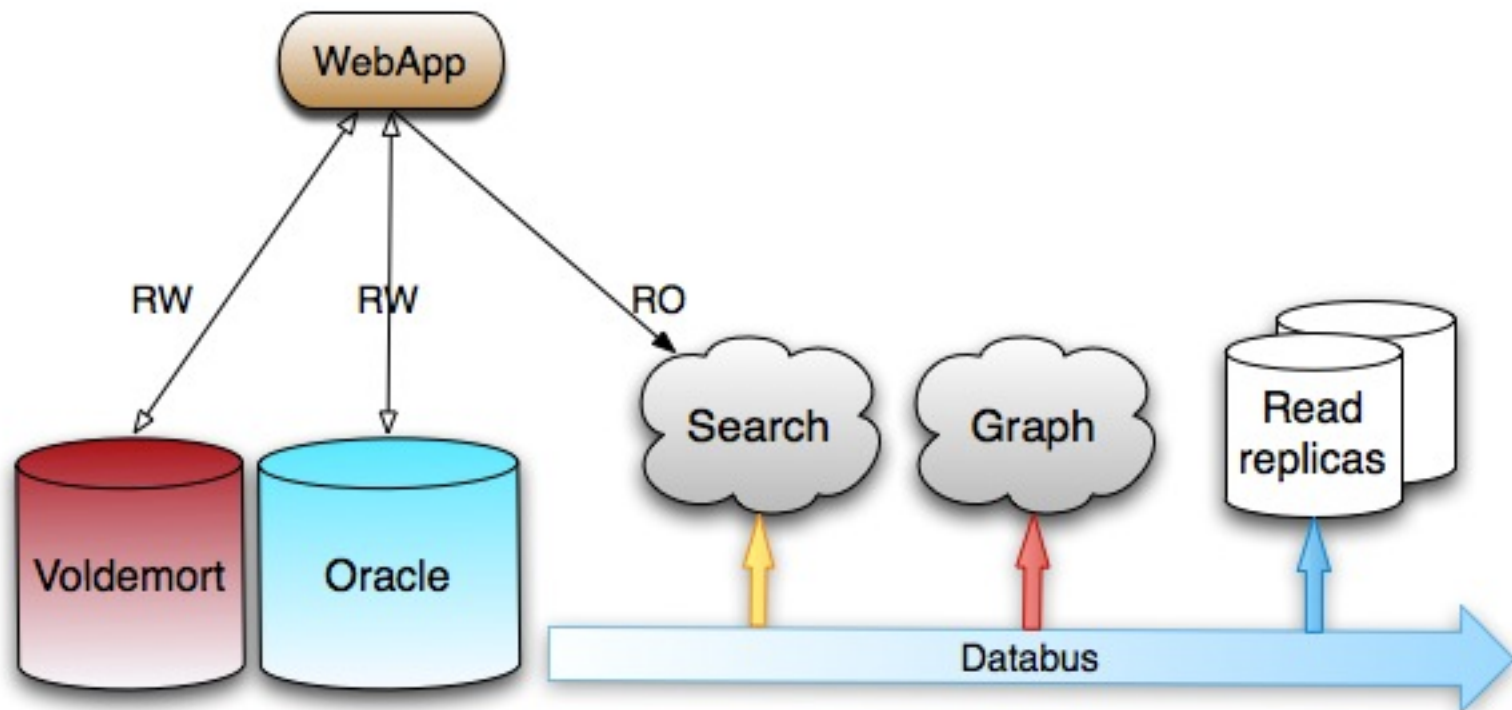


- Transport independent of data source: Oracle, MySQL, ...
- Transactional semantics
- In order, at least once delivery
- Tens of relays
- Hundreds of sources
- Low latency - milliseconds

LinkedIn Product Architecture



LinkedIn Product Architecture




LinkedIn Data Infrastructure Solutions

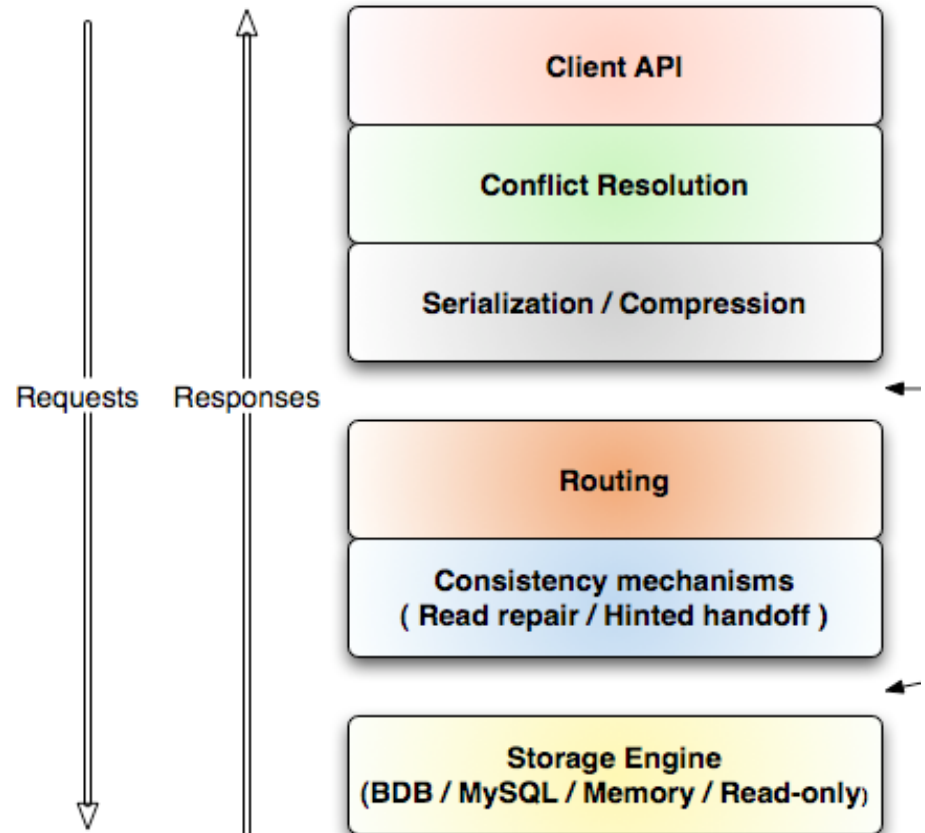
Voldemort: Highly-Available Distributed KV Store



Voldemort: Architecture

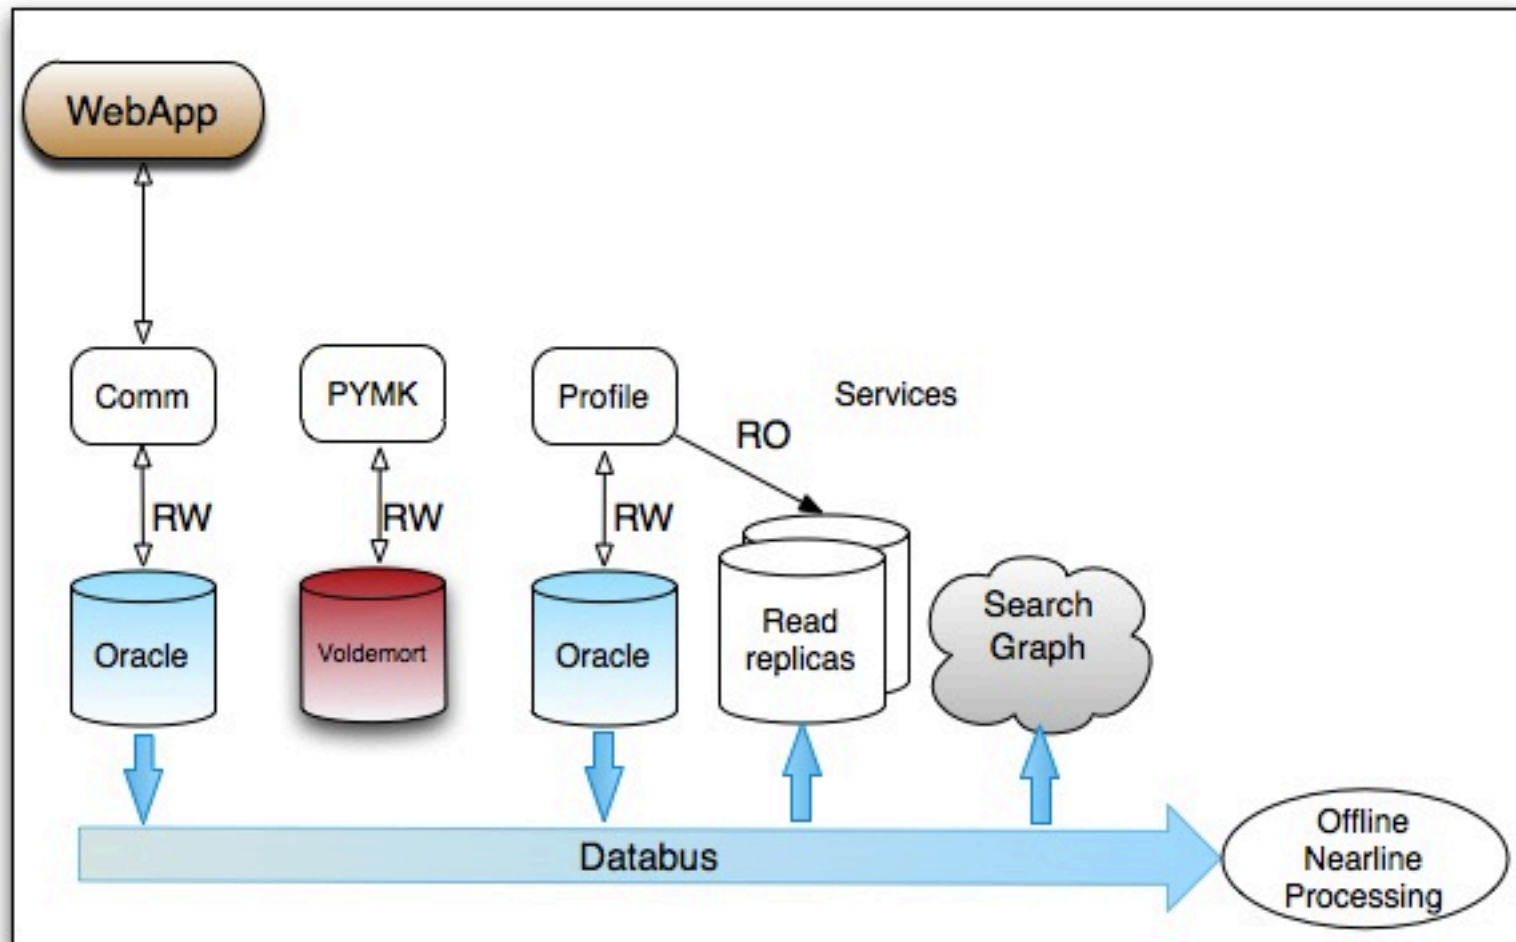


- Pluggable components
- Tunable consistency / availability
- Key/value model, server side “views”



- 10 clusters, 100+ nodes
- Largest cluster – 10K+ qps
- Avg latency: 3ms
- Hundreds of Stores
- Largest store – 2.8TB+

LinkedIn Product Architecture




LinkedIn Data Infrastructure Solutions

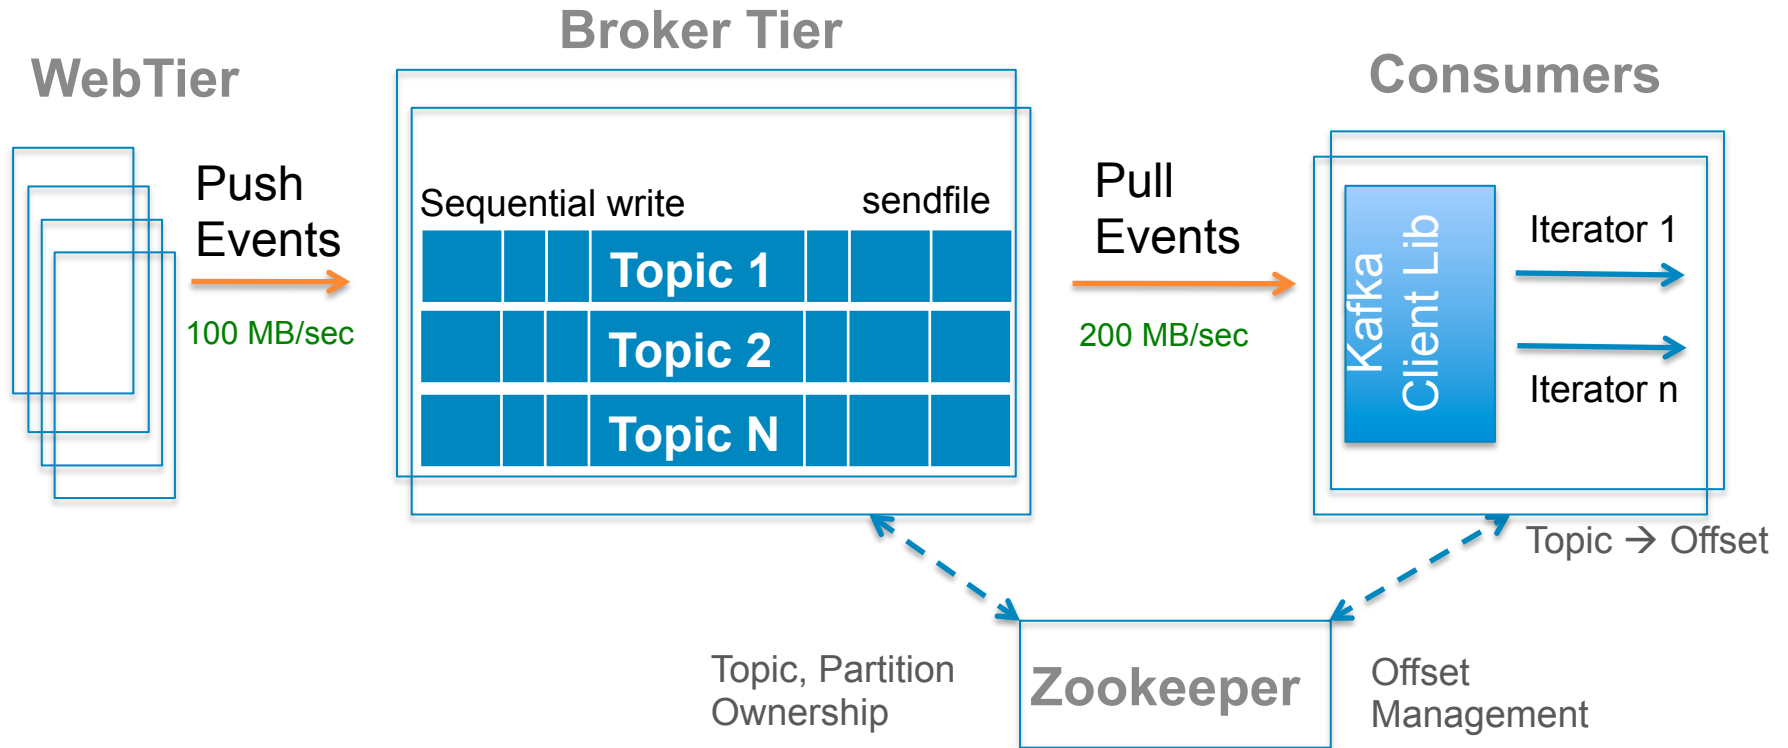
Kafka: High-Volume Low-Latency Messaging System



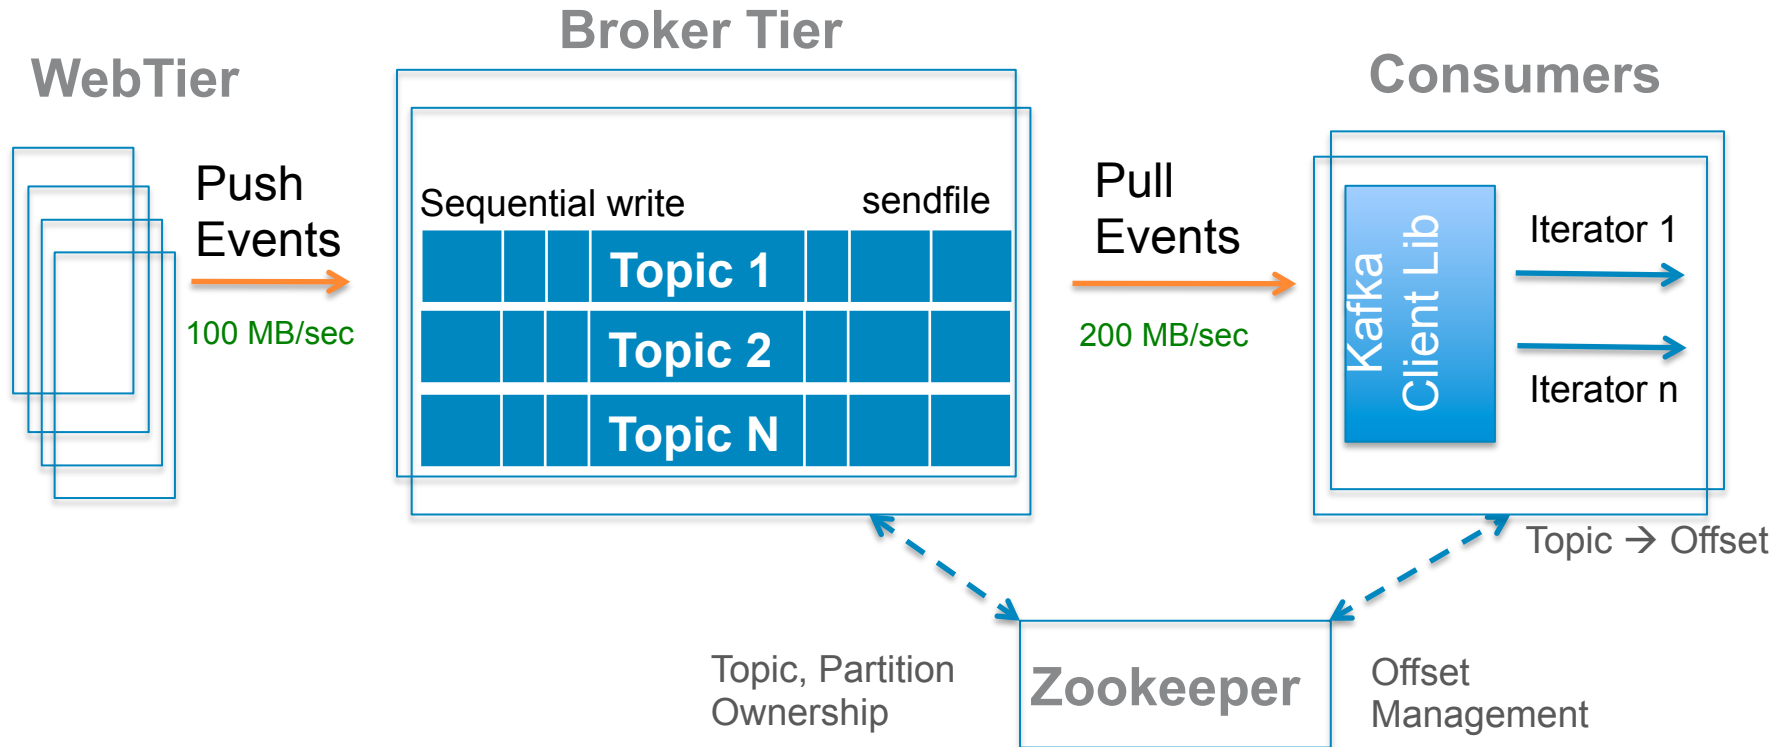
LinkedIn Product Architecture



Kafka: Architecture



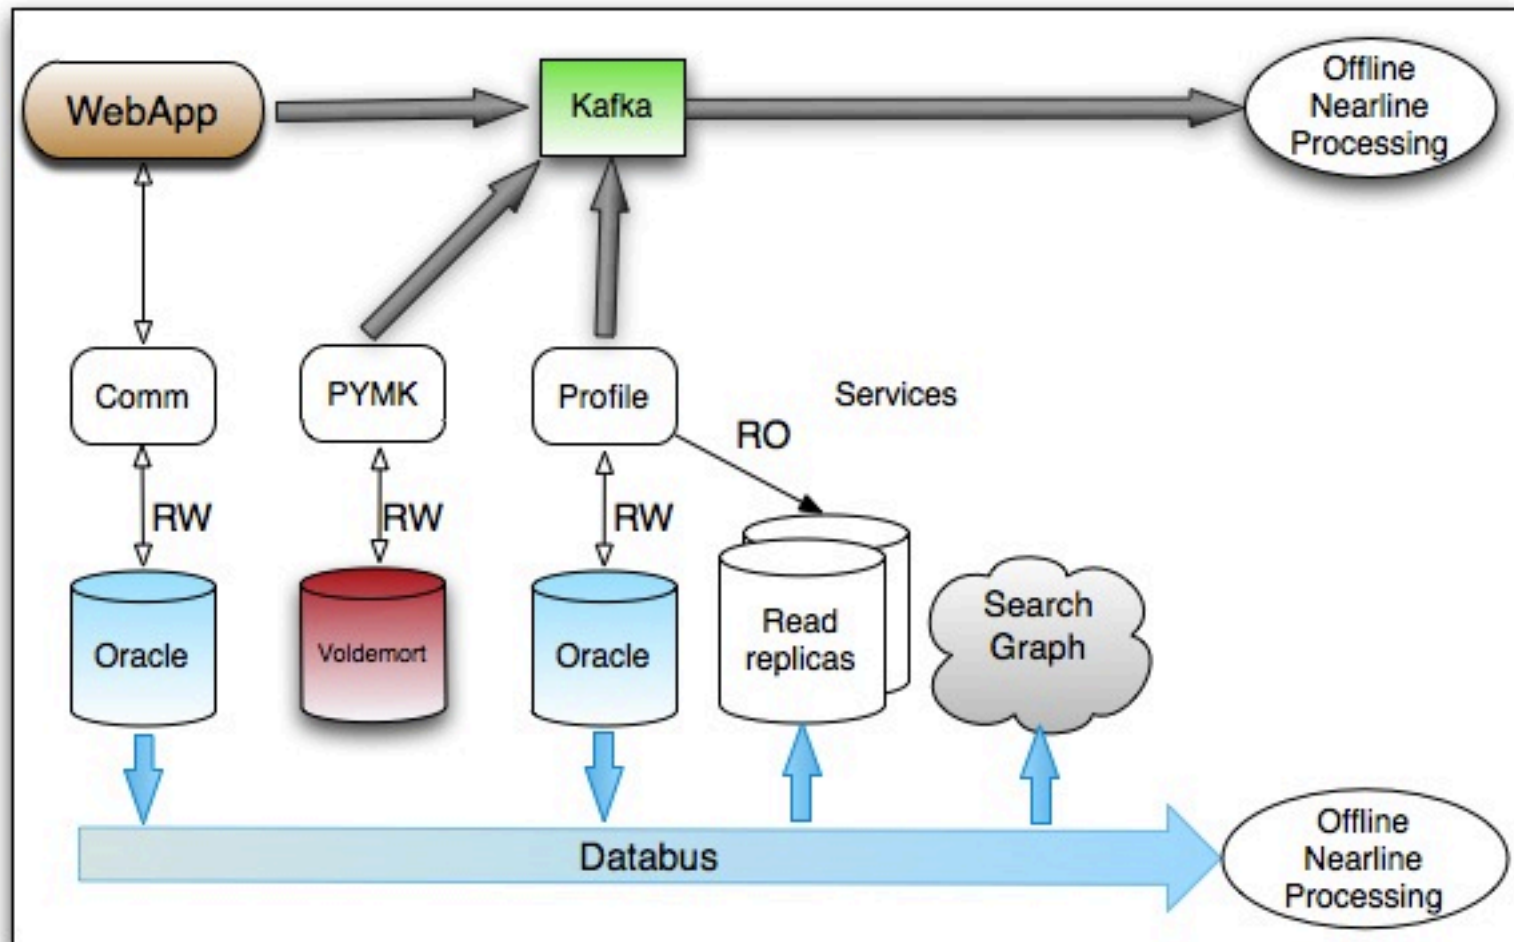
Kafka: Architecture



- At least once delivery
- Very high throughput
- Low latency
- Durability

- Billions of Events, TBs per day
- 50K+ per sec at peak
- Inter and Intra-cluster replication
- End-to-end latency: few seconds

LinkedIn Product Architecture



LinkedIn Data Infrastructure Solutions

Espresso: Indexed Timeline-Consistent Distributed Data Store



Application View

Mailbox Database

Message Metadata Table

| MemberId | MsgId | Value Blob |
|----------|-------|-----------------------------|
| bob | 1 | Invitation to join LinkedIn |
| bob | 2 | Job opportunity |
| bob | 3 | Request for referral |
| tom | 1 | Invitation to join LinkedIn |
| tom | 2 | Job opportunity |

Message Details Table

| MemberId | MsgId | Value Blob |
|----------|-------|---------------------------|
| bob | 1 | "Dear Bob,...." |
| bob | 2 | "Hello there,...." |
| bob | 3 | "Good morning, " |
| tom | 1 | "Hi Tom,..." |
| tom | 2 | "Interesting opportunity" |

Mailbox Aggregates Table

| MemberId | Value Blob |
|----------|----------------------|
| bob | unread:20, total:100 |
| tom | unread: 2, total: 25 |

Hierarchical data model

Rich functionality on resources

- ✓ Conditional updates
- ✓ Partial updates
- ✓ Atomic counters

Rich functionality within resource groups

- ✓ Transactions
- ✓ Secondary index
- ✓ Text search

Partitioning

Mailbox Database

| MemberId | MsgId | Value Blob |
|----------|-------|-----------------------------|
| bob | 1 | Invitation to join LinkedIn |
| bob | 2 | Job opportunity |
| bob | 3 | Request for referral |
| tom | 1 | Invitation to join LinkedIn |
| tom | 2 | Job opportunity |

| MemberId | MsgId | Value Blob |
|----------|-------|---------------------------|
| bob | 1 | "Dear Bob,...." |
| bob | 2 | "Hello there,...." |
| bob | 3 | "Good morning, " |
| tom | 1 | "Hi Tom,...." |
| tom | 2 | "Interesting opportunity" |

| MemberId | Value Blob |
|----------|----------------------|
| bob | unread:20, total:100 |
| tom | unread: 2, total: 25 |

Mailbox Database - Partition 1

| MemberId | MsgId | Value Blob |
|----------|-------|-----------------------------|
| bob | 1 | Invitation to join LinkedIn |
| bob | 2 | Job opportunity |
| bob | 3 | Request for referral |

| MemberId | Value Blob |
|----------|----------------------|
| bob | unread:20, total:100 |

| MemberId | MsgId | Value Blob |
|----------|-------|--------------------|
| bob | 1 | "Dear Bob,...." |
| bob | 2 | "Hello there,...." |
| bob | 3 | "Good morning, " |

Mailbox Database - Partition 2

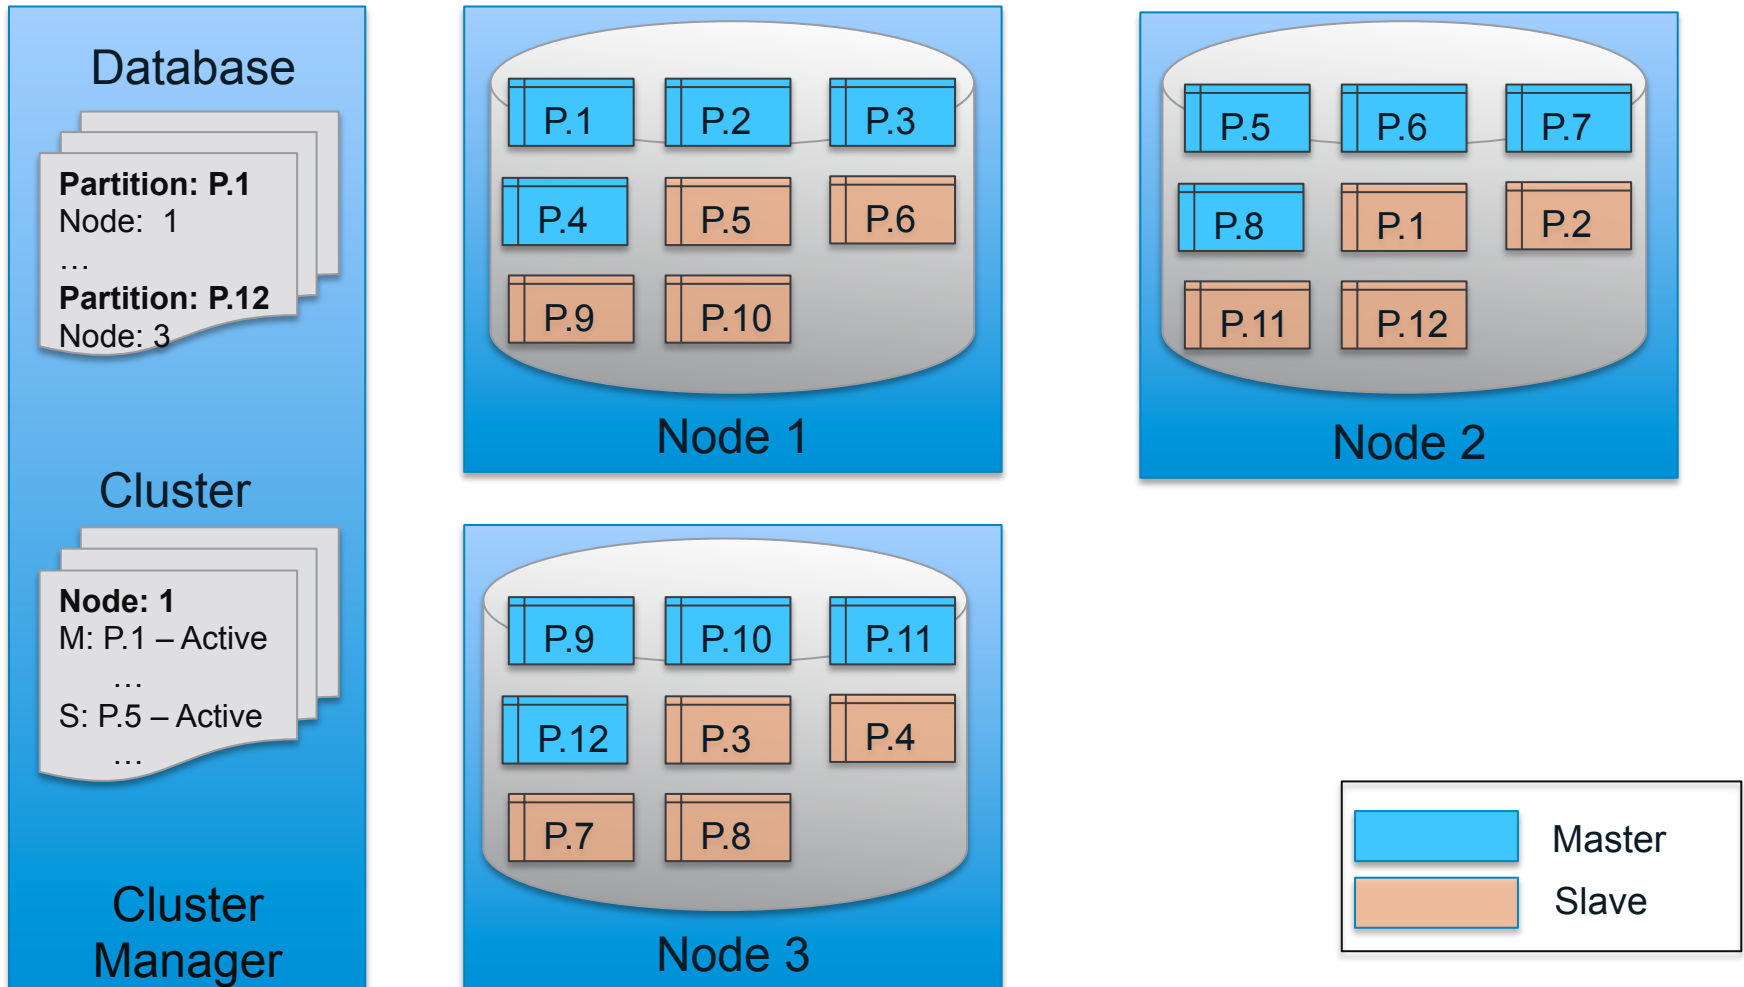
| MemberId | MsgId | Value Blob |
|----------|-------|-----------------------------|
| tom | 1 | Invitation to join LinkedIn |
| tom | 2 | Job opportunity |

| MemberId | Value Blob |
|----------|----------------------|
| tom | unread: 2, total: 25 |

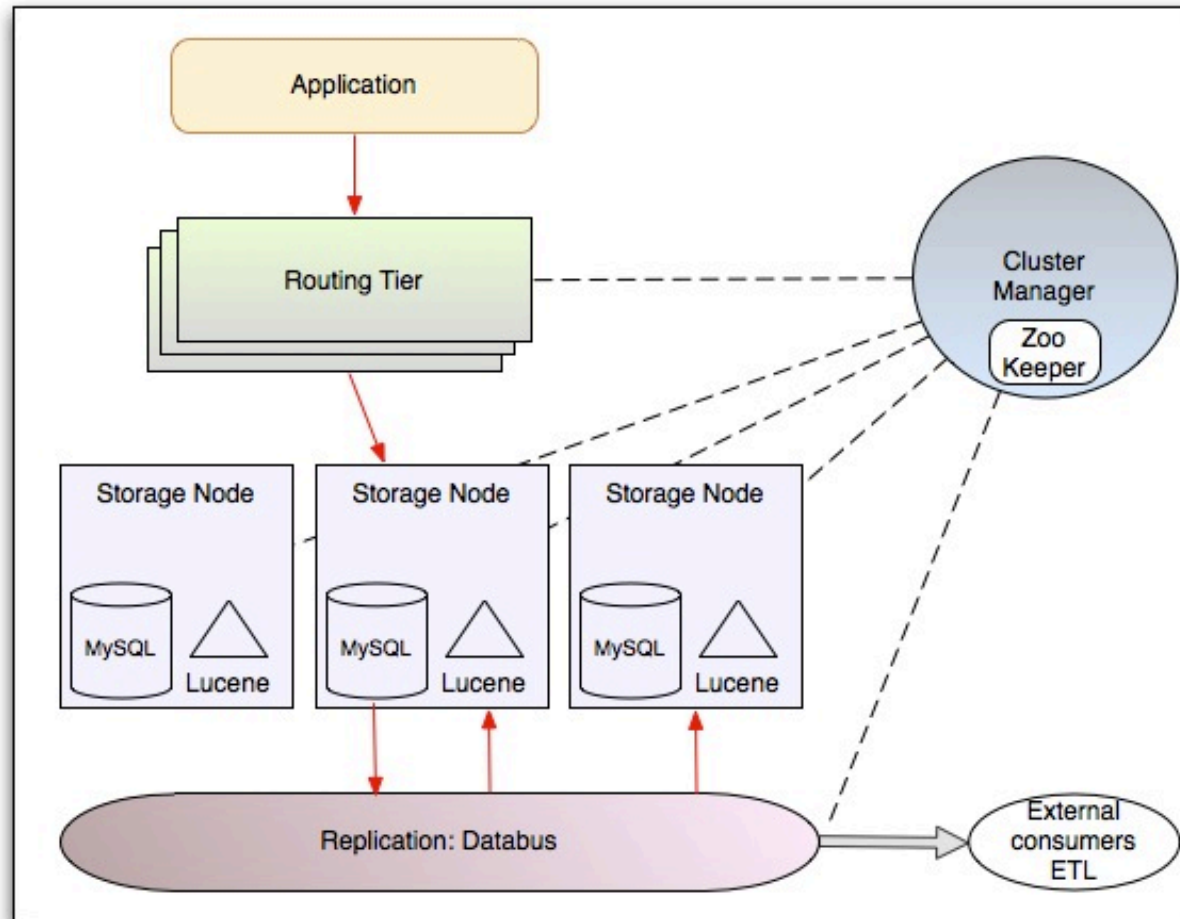
| MemberId | MsgId | Value Blob |
|----------|-------|---------------------------|
| tom | 1 | "Hi Tom,...." |
| tom | 2 | "Interesting opportunity" |

Espresso Partition Layout: Master, Slave

3 Storage Engine nodes, 2 way replication

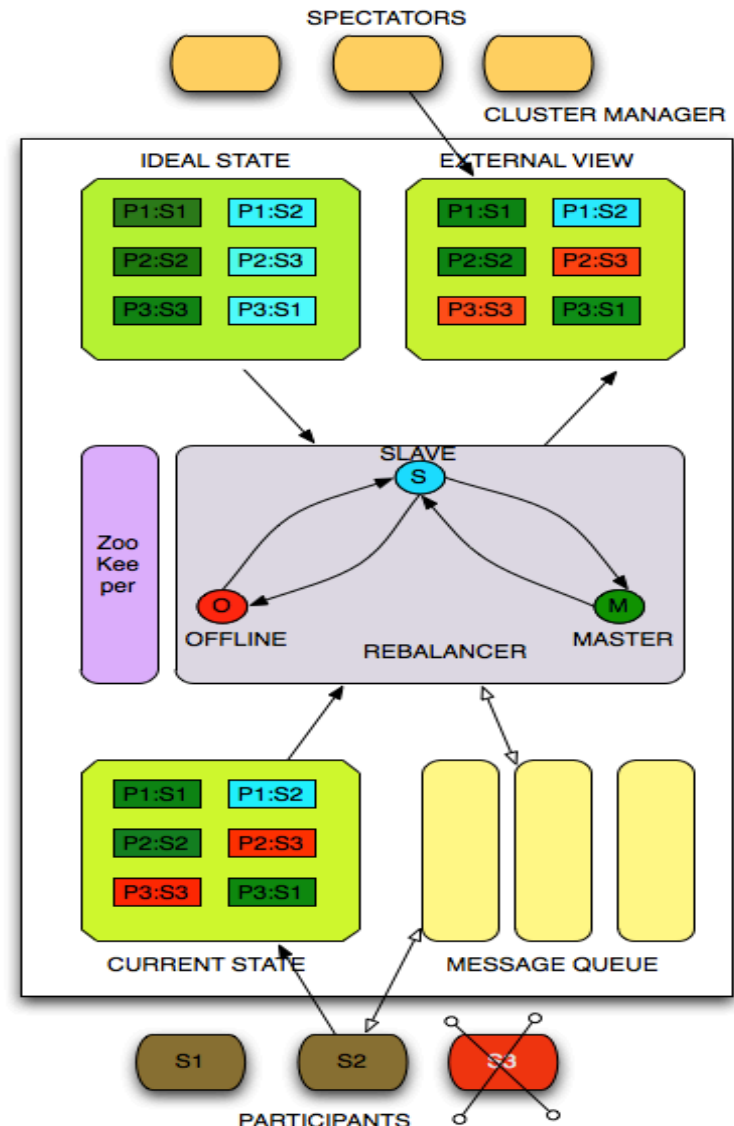


Espresso: System Components



Generic Cluster Manager: Helix

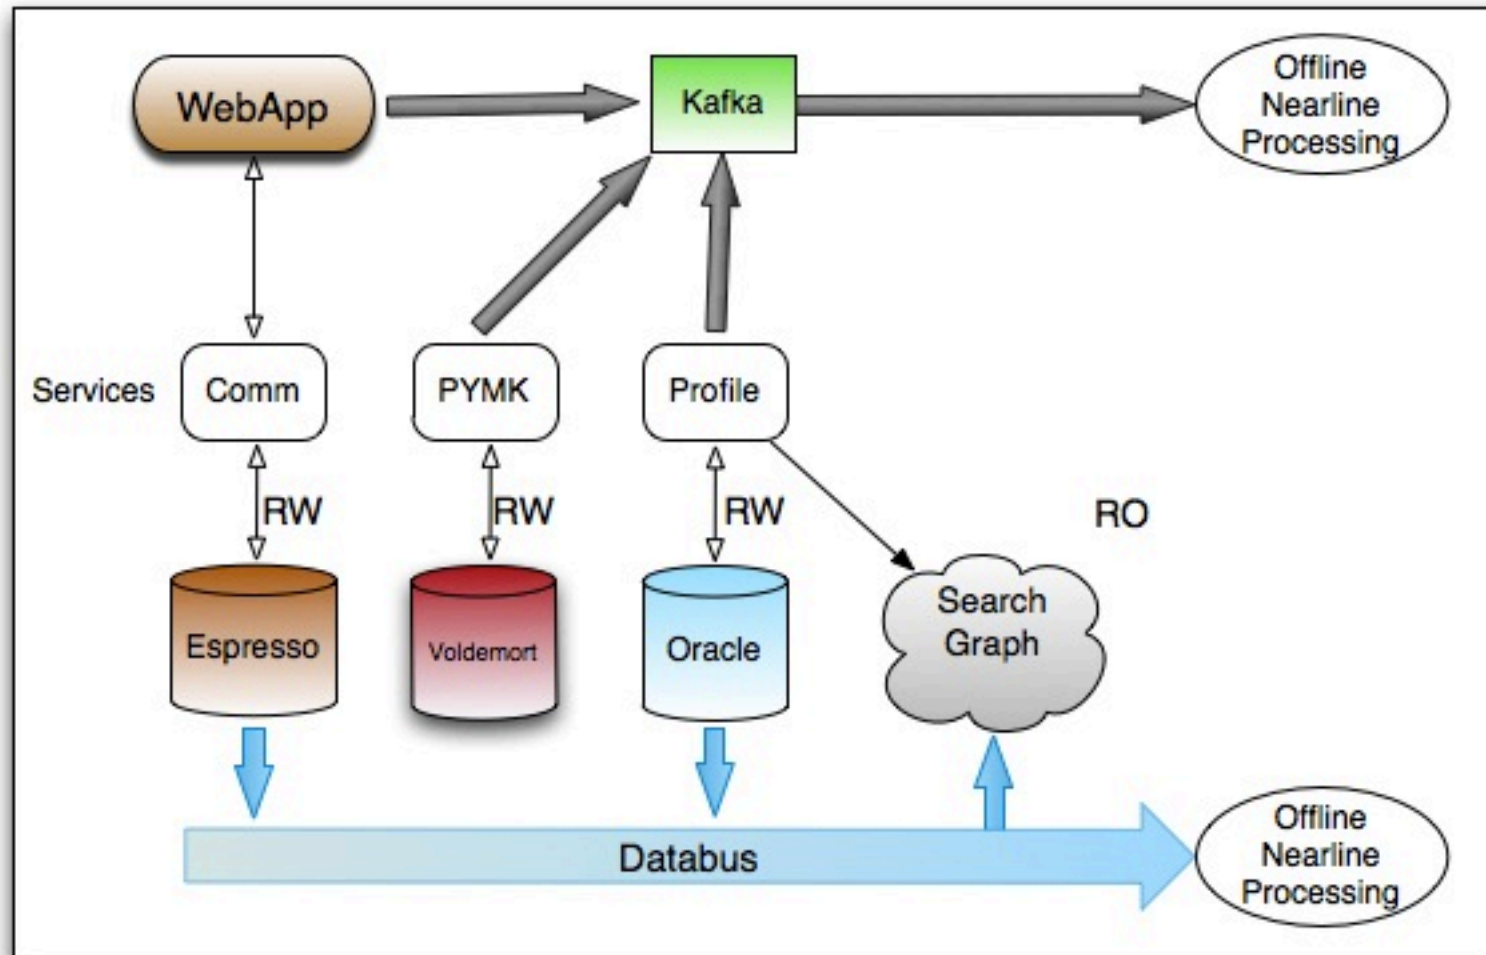
- Generic Distributed State Model
- Centralized Config Management
- Automatic Load Balancing
- Fault tolerance
- Health monitoring
- Cluster expansion and rebalancing
- Espresso, Databus and Search
- Open Source Apr 2012
- <https://github.com/linkedin/helix>



Espresso@Linkedin

- Launched first application Oct 2011
- Open source 2012
- Future
 - Multi-Datacenter support
 - Global secondary indexes
 - Time-partitioned data

LinkedIn Product Architecture



Acknowledgments

Siddharth Anand, Aditya Auradkar, Chavdar Botev, Vinoth Chandar, Shirshanka Das, Dave DeMaagd, Alex Feinberg, John Fung, Phanindra Ganti, Mihir Gandhi, Lei Gao, Bhaskar Ghosh, Kishore Gopalakrishna, Brendan Harris, Rajappa Iyer, Swaroop Jagadish, Joel Koshy, Kevin Krawez, Jay Kreps, Shi Lu, Sunil Nagaraj, Neha Narkhede, Sasha Pachev, Igor Perisic, Lin Qiao, Tom Quiggle, Jun Rao, Bob Schulman, Abraham Sebastian, Oliver Seeliger, Adam Silberstein, Boris Shkolnik, Chinmay Soman, Subbu Subramaniam, Roshan Sumbaly, Kapil Surlaker, Sajid Topiwala, Cuong Tran, Balaji Varadarajan, Jemiah Westerman, Zach White, Victor Ye, David Zhang, and Jason Zhang

Questions?

