

# LinkedIn: Network Updates Uncovered

Ruslan Belkin

Sean Dawson

# Agenda



- Quick Tour
  - Requirements (User Experience / Infrastructure)
  - Service API
  - Internal Architecture
  - Applications (e.g., Twitter Integration, Email Delivery)
  - Measuring Performance
- 
- *Shameless self promotion*

## Network Updates

[See more updates »](#)



Sean Dawson preparing for my QCon presentation, Nov 20th  
9 days ago



Share

Today

### STATUS UPDATES (7)



**Elliot Shmukler** Leaving work. Weekend projects may involve OAuth & Named Pipes. 32 minutes ago [Reply privately](#) - [Add comment](#)



**Nick Dellamaggiore** is testing out some cool new features on LinkedIn 45 minutes ago [Reply privately](#) - [Add comment](#)



**Florina Xhabija** can't believe it's actually raining! 1 hour ago [Reply privately](#) - [Add comment](#)

[Show more...](#)

### RECOMMENDATIONS (2)



**Eishay (יש) Smith** was recommended by **Jen Granito**, "Eishay is a great engineer and teammate. Throughout our time together working on the news product, ..." 9 hours ago [Read more »](#)



**Jason Nolan** recommends **Elizabeth Moyer**, "Elizabeth has been an excellent colleague to work with at Ryerson. She bridged the faculty/admin div..." 11 hours ago [Read more »](#)

### People you may know



**Ben Konrath** (2nd)  
Software Developer

[invite](#) | [x](#)



**Cynde Bacher** (3rd)  
Technical College  
Recruiter/University Programs  
Specialist at Google

[invite](#) | [x](#)



**David Ascher** (2nd)  
CEO, Mozilla Messaging

[invite](#) | [x](#)

[See more »](#)

### Who's viewed my profile?

Your profile has been viewed by 4 people in the last 3 days. In the last 3 days, you have appeared in search results 22 times.

[See more »](#)

### Company Buzz

Topic: LinkedIn

[change »](#)



**@LinkedIn** hey, nice redesign!  
bradybd | 7 minutes ago | [Share](#)



Pratyush Academy New Posted a new entry


# The Stack

Environment	90% Java 5% Groovy 2% Scala 2% Ruby 1% C++
Containers	Tomcat, Jetty
Data Layer	Oracle, MySQL, Voldemort, Lucene, Memcache
Offline Processing	Hadoop
Queuing	ActiveMQ
Frameworks	Spring

# The Numbers



Updates Created	35M / week
Update Emails	14M / week
Service Calls	20M / day 230 / second



## Network Updates

[See more updates »](#)



Sean Dawson preparing for my QCon presentation, Nov 20th  
9 days ago

Today

### STATUS UPDATES (7)



**Elliot Shmukler** Leaving work. Weekend projects may involve OAuth & Named Pipes. 32 minutes ago [Reply privately](#) - [Add comment](#)



**Nick Dellamaggiore** is testing out some cool new features on LinkedIn 45 minutes ago [Reply privately](#) - [Add comment](#)



**Florina Xhabija** can't believe it's actually raining! 1 hour ago [Reply privately](#) - [Add comment](#)

[Show more...](#)

### RECOMMENDATIONS (2)



**Eishay (יש) Smith** was recommended by **Jen Granito**, "Eishay is a great engineer and teammate. Throughout our time together working on the news product, ..." 9 hours ago [Read more »](#)



**Jason Nolan** recommends **Elizabeth Moyer**, "Elizabeth has been an excellent colleague to work with at Ryerson. She bridged the faculty/admin div..." 11 hours ago [Read more »](#)

### People you may know



**Ben Konrath** (2nd)  
Software Developer

[invite](#) | [x](#)



**Cynde Bacher** (3rd)  
Technical College  
Recruiter/University Programs  
Specialist at Google

[invite](#) | [x](#)



**David Ascher** (2nd)  
CEO, Mozilla Messaging

[invite](#) | [x](#)

[See more »](#)

### Who's viewed my profile?

Your profile has been viewed by 4 people in the last 3 days. In the last 3 days, you have appeared in search results 22 times.

[See more »](#)

### Company Buzz

Topic: LinkedIn

[change »](#)



**@LinkedIn** hey, nice redesign!  
bradybd | 7 minutes ago | [Share](#)



Pratyush Academy New Posted a new entry

# Stream View



## Network Updates

[See more updates »](#)



**Sean Dawson** preparing for my QCon presentation, Nov 20th  
9 days ago



[Share](#)



**Adam Nash** has an updated AIM ID: [adamnashLINKEDIN](#), and other personal information 32 minutes ago [Update your personal information »](#)



**Elliot Shmukler** Leaving work. Weekend projects may involve OAuth & Named Pipes. 38 minutes ago [Reply privately](#) - [Add comment](#)



**Nick Dellamaggiore** is testing out some cool new features on LinkedIn 50 minutes ago [Reply privately](#) - [Add comment](#)



**Nick Dellamaggiore** now has a birthday listed: February 17 1 hour ago [Update your personal information »](#)



**Florina Xhabija** can't believe it's actually raining! 1 hour ago [Reply privately](#) - [Add comment](#)



**Joichi Ito** Arrived Singapore 2 hours ago [Reply privately](#) - [Add comment](#)


# Connection View



LinkedIn [Home](#) [Profile](#) [Contacts](#) [Groups](#) [Jobs](#) [Inbox](#) [More...](#) People

[All Updates](#) [Category View](#) **Connection View** [My Updates](#)

---


 **Jim Meyer** has updated their current title to Member, Board of Advisors at [Kinnexus](#) 18 hours ago  
[Send a message](#) - [Add comment](#)

**Jim Meyer** is reading about negotiauctions, the intersection of negotiation & auction.  
<http://hbswk.hbs.edu/item/6256.html> (via <http://xrl.us/MikeMcG>) 3 days ago [Reply privately](#) - 1 comment

**Jim Meyer** was recommended by [Alejandro Crosa](#), "Jim is a highly technical individual that can get really deep into the implementation details of sof..." 3 days ago [Read more](#) »

[See all Jim's updates](#) » [See Jim's 500+ connections](#) »

---


 **Andrey Petrov** was recommended by [Chia Hwu](#), "Andrey is quick and brilliant at writing code that works. In addition, he is mature, professional an..." 1 day ago [Read more](#) »

**Andrey Petrov** was recommended by [Jeff Mellen](#), "Andrey is a bright fellow and a clever engineer, and his experience and insight belies his age. He ..." 5 days ago [Read more](#) »

**Andrey Petrov** is now using LinkedIn for iPhone. [Learn more](#) » 28 days ago [Add comment](#)

[See all Andrey's updates](#) » [See Andrey's 99 connections](#) »

---

 **Brad Olcott** is now Director of User Experience at [Survey Monkey](#). 1 day ago [Send a message](#) - 6 comments

**Steve Ganz** 1 day ago  
Congrats!

[Show 5 previous comments](#) »

**Bryan Haggerty** 20 hours ago  
Congrats Brad. Catch you on the train.



# Profile



unit testing, build management/automation, version control, python, java, social networks, end-to-end product development

## Applications

[Add Application](#)

Tip: Add more to your professional profile by [adding applications](#) from our featured partners.

## Reading List by Amazon [\[ Remove \]](#)

### ReadingList by amazon



**Born to Run: A Hidden Tribe, Superathletes, and the Greatest Race the World Has Never Seen**  
by Christopher McDougall

[See this book on Amazon »](#)

Sean has read this book

Recommended



**Into Thin Air: A Personal Account of the Mt. Everest Disaster**  
by Jon Krakauer

[See this book on Amazon »](#)

Sean has read this book

Recommended

## Sean's Activity

[edit](#)

### Friday

**Sean Dawson** has an updated profile ( Associations, Honors, Experience ) 2 days ago

**Sean Dawson** recommends **Born to Run: A Hidden Tribe, Superathletes, and the Greatest Race the World Has Never Seen** 2 days ago [Add comment](#)

**Sean Dawson** is reading **Once a Runner: A Novel** 2 days ago [Add comment](#)

**Sean Dawson** now has a birthday listed: September 30 2 days ago [Update your personal information »](#)

### Wednesday

**Sean Dawson** is now connected to **Paula Janicek** 4 days ago

### Wednesday, Oct 28

**Sean Dawson** preparing for my QCon presentation, Nov 20th 11 days ago - [Add comment](#)

**Sean Dawson** is now connected to **Travis Forden** and **Matthew Strimas-Mackey** 11 days ago

[See more »](#)

# Groups



Welcome, Sean Dawson · [Add Connections](#) · [Account & Settings](#) · [Help](#) · [Sign Out](#)

LinkedIn [Home](#) [Profile](#) [Contacts](#) [Groups](#) [Jobs](#) [Inbox \(1\)](#) [More...](#)

**My Groups** [Following](#) [Groups Directory](#) [Create a Group](#)

[FAQ](#)

## My Company Groups (1)

[Find a Group](#) | [Create a Group](#)

 **LinkedIn**  
Activity: [Discussions \(7\)](#) [News \(2\)](#)  
[Go to](#) ▾

## My User Groups (12)

[Edit the order of your groups »](#)

 **AIDS/LifeCycle**  
[Go to](#) ▾ [Actions](#) ▾

 **SF2G**  
[Go to](#) ▾ [Actions](#) ▾

 **JavaOne 2008**  
[Go to](#) ▾ [Actions](#) ▾

 **'09 Grads**  
Activity: [Discussions \(3\)](#) [News \(2\)](#) [Jobs \(1\)](#)

## People I'm Following

### LINKEDIN



**Ian McCarthy** and **Mario Sundar** commented on:  
**LinkedIn Is Getting a Redesign [Pics] - Mashable!**

"def! and thanks for the groups shout-out in that screenshot ;)"

3 hours ago 2 comments



**Sal Becerra** commented on:  
**Looking for real everyday LinkedIn members to run ideas by, use in beta tests or white lists?**

3 hours ago 1 comment

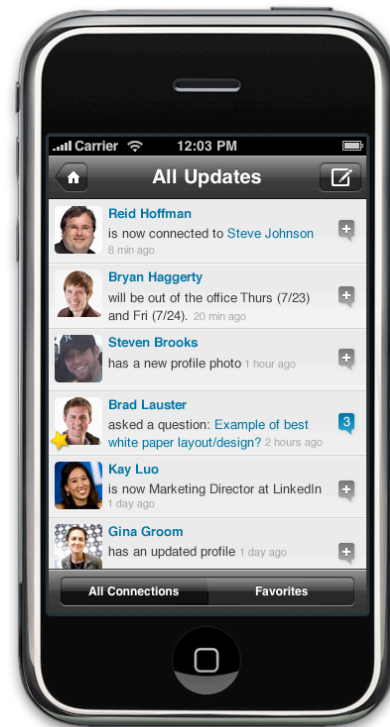
### LINKEDIN GROUPS PRODUCT FORUM



**Dylan (In Yong) Song** commented on:  
**New features this evening and a brief downtime**

3 hours ago 31 comments

# Mobile



Subject: LinkedIn Network Updates, 10/28/2009

From: LinkedIn Updates

## LinkedIn Network Updates

Oct 21 - Oct 28

### FEATURED UPDATES

[See more updates »](#)



#### PROFILE

[Esteban Kozak](#) has updated their current title to Principal Product Manager at LinkedIn  
[Candy Chastain Mielke](#) has updated their current title to Consultant at LinkedIn  
[Georg Puchta](#) has added new links: "[2D Barcode](#)"



#### CONNECTIONS

[Thanos Foufoulas](#) is now connected to [Geoff Wolfe](#)



#### STATUS

[Nancy Boever](#) Looking for a Purchasing Manager to join our Finance team! [Add a comment »](#)  
[Maisy Samuelson](#) Co. Profiles plug, "information provided on these profiles rivals more well-known business sites like Hoovers." <http://tinyurl.com/coProfile> [Add a comment »](#)  
[Konstantin Guericke](#) hatte viel Spaß, neue LinkedInsiders kennenzulernen. Morgen wieder in den Flieger zurück nach Kalifornien. [Add a comment »](#)



#### GROUPS

3 of your connections have joined the group [Friends of LinkedIn](#) including [Stefan Apitz](#), [Ian McNish](#),  
2 of your connections have joined the group [National Eagle Scout Association \(NESA\)](#),  
[BSA Eagle's Nest](#) including [Jakob Heuser](#) and [Scott Schlegel](#)  
[Basil Hashem](#) has joined [Ruby on Rails](#)



#### QUESTIONS & ANSWERS

[Chris Richman](#) [answered 2 questions](#) about Using LinkedIn  
[April Kelly](#) [answered 1 question](#) about Using LinkedIn

- Home
- Company Groups
- User Groups
- Profile
- Contacts
  - Connections
  - Imported Contacts
  - Network Statistics
- Inbox
  - Compose Message
  - Messages
  - InMail
  - Introductions
  - Invitations
  - Profiles
  - Q&A
  - Jobs
  - Recommendations
  - Groups

[Add Connections](#)

**Ruslan Belkin**  
Senior Director of Engineering at LinkedIn  
What are you working on?

**234 Connections**  
links you to 2,488,500+ professionals

**7,329 New People**  
in your network since October 21

Does your profile truly represent you? Update it today.

## Network Updates

Your Network Updates feed is being updated. Please check back soon.

### Who's viewed my profile?

Your profile has been viewed by 13 people in the last 11 days. In the last 9 days, you have appeared in search results 25 times.

[See more >](#)

### Answers : Software Development

- Q. Do somebody use any open source / commercial CMS for WAP portal management?
- Q. Question about Questions: At this moment in time, what are you top 5 questions that if you had answers to, could make a significant difference in your career or life?
- Q. We are planning to upgrade SAP 4.7 to ECC 6. There is a suspicion that SAP Script (Mostly for reports) may not be working. Does anyone know the compatibility or incompatibility?

[See more >](#)

### People : Java Software...

No results found. Please try again.

[Add an Application](#)

These modules are customizable.

# Expectations – User Experience



- Multiple presentation views
- Comments on updates
- Aggregation of noisy updates
- Partner Integration
- Easy to add new updates to the system
- Handles I18N and other dynamic contexts
- Long data retention

# Expectations - Infrastructure



- Large number of connections, followers and groups
- High request volume + Low Latency
- Random distribution lists
- Black/White lists, A/B testing, etc.
- Tenured storage of update history
- Tracking of click through rates, impressions
- Supports real-time, aggregated data/statistics
- Cost-effective to operate

# Historical Note



*(homepage circa 2007)*

## Network updates since Sep 18

**Send an update:** I have a new job

**Questions from your network:**

Asked by your connections:  
**Colleen Cole:** What are the best colocation facilities in the Bay Area?

Asked by friends of friends:  
What's a good advertising company which specializes in online... ?  
Where do most technology startups get their angel funding?  
Do minimum wage laws apply to all companies?  
[View more questions »](#)

---

**Profile updates:**

- Samuel Chung has new profile information. [View Profile.](#)  
[See all...](#)

---

**Chris Yee** is now a connection  
**Billy Bob Thorton** is now a connection

---

**Chris Yee** has just added 1 connection  
**Billy Bob Thorton** has just added 1 connection  
[See all...](#)

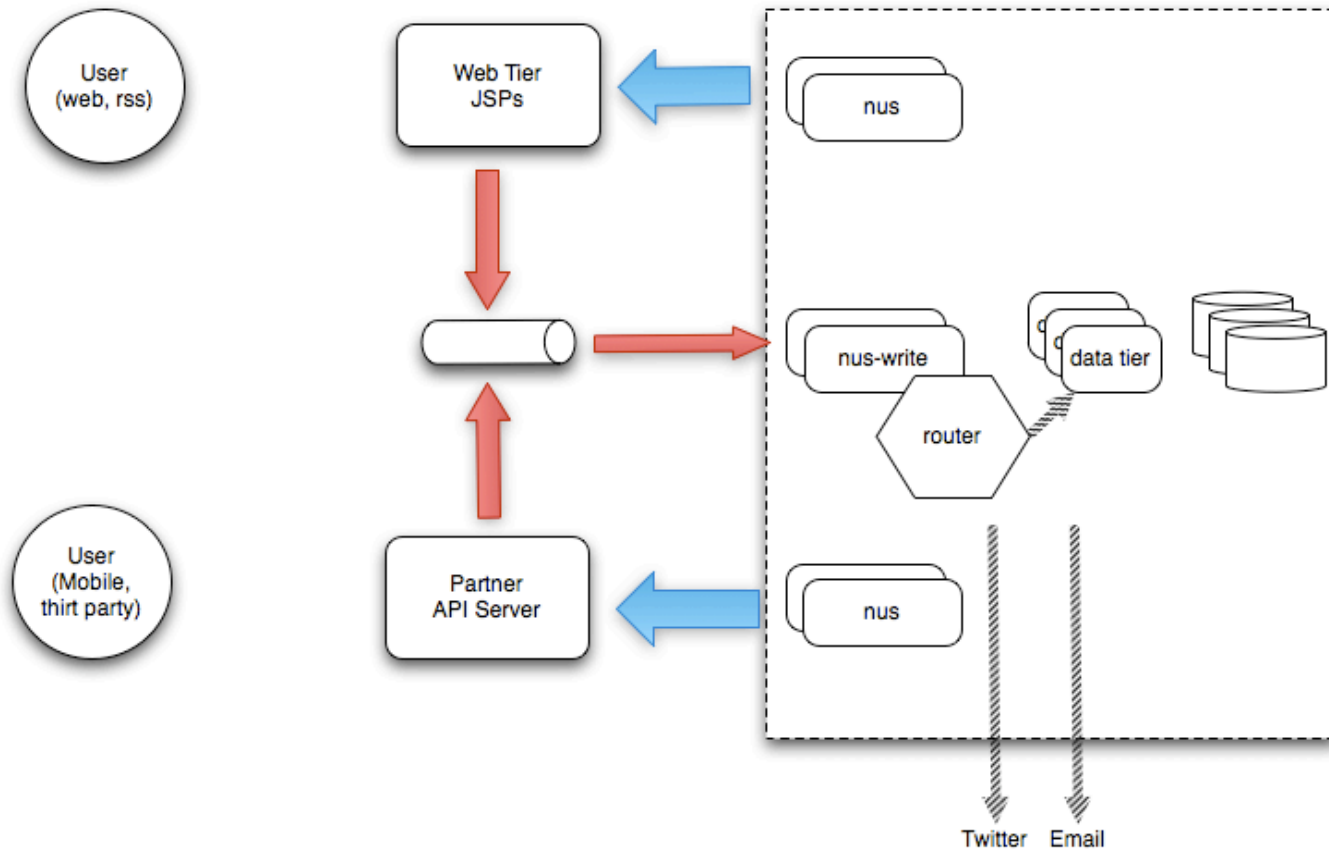
---

**6,017 new people** in your network since September 18

- Legacy “network update” feature was a mixed bag of detached services.
- Neither consistent nor scalable
- Tightly coupled to our Inbox
- Migration plan
  - Introduce API, unify all disparate service calls
  - Add event-driven activity tracking with DB backend
  - Build out the product
  - Optimize!



# Network Updates Service - Overview



# Service API – Data Model



**Network updates** since Sep 18

Send an update: I have a new job

**Questions from your network:** ask your network a question

Asked by your connections:  
**Colleen Cole:** What are the best colocation facilities in the Bay Area?

Asked by friends of friends:  
What's a good advertising company which specializes in online... ?  
Where do most technology startups get their angel funding?  
Do minimum wage laws apply to all companies?  
[View more questions »](#)

**Profile updates:** update your own profile

- Samuel Chung has new profile information. [View Profile.](#)  
See all...

**Chris Yee is now a connection**  
**Billy Bob Thorton is now a connection**

Chris Yee has just added 1 connection  
Billy Bob Thorton has just added 1 connection  
See all...

**6,017 new people** in your network since September 18

```
<updates>
  <NCON>
    <connection>
      <id>2</id>
      <firstName>Chris</firstName>
      <lastName>Yee</lastName>
    </connection>
  </NCON>
</updates>
```

# Service API – Post



```
NetworkUpdatesNotificationService service =
    getNetworkUpdatesNotificationService();

ProfileUpdateInfo profileUpdate = createProfileUpdate();

Set<NetworkUpdateDestination> destinations =
    Sets.newHashSet(
        NetworkUpdateDestinations.newMemberFeedDestination(1213)
    );

NetworkUpdateSource source =
    new NetworkUpdateMemberSource(1214);

Date updateDate = getClock().currentDate();

service.submitNetworkUpdate(source,
                            destinations,
                            updateDate,
                            profileUpdate);
```

# Service API – Retrieve



```
NetworkUpdatesService service = getNetworkUpdatesService();

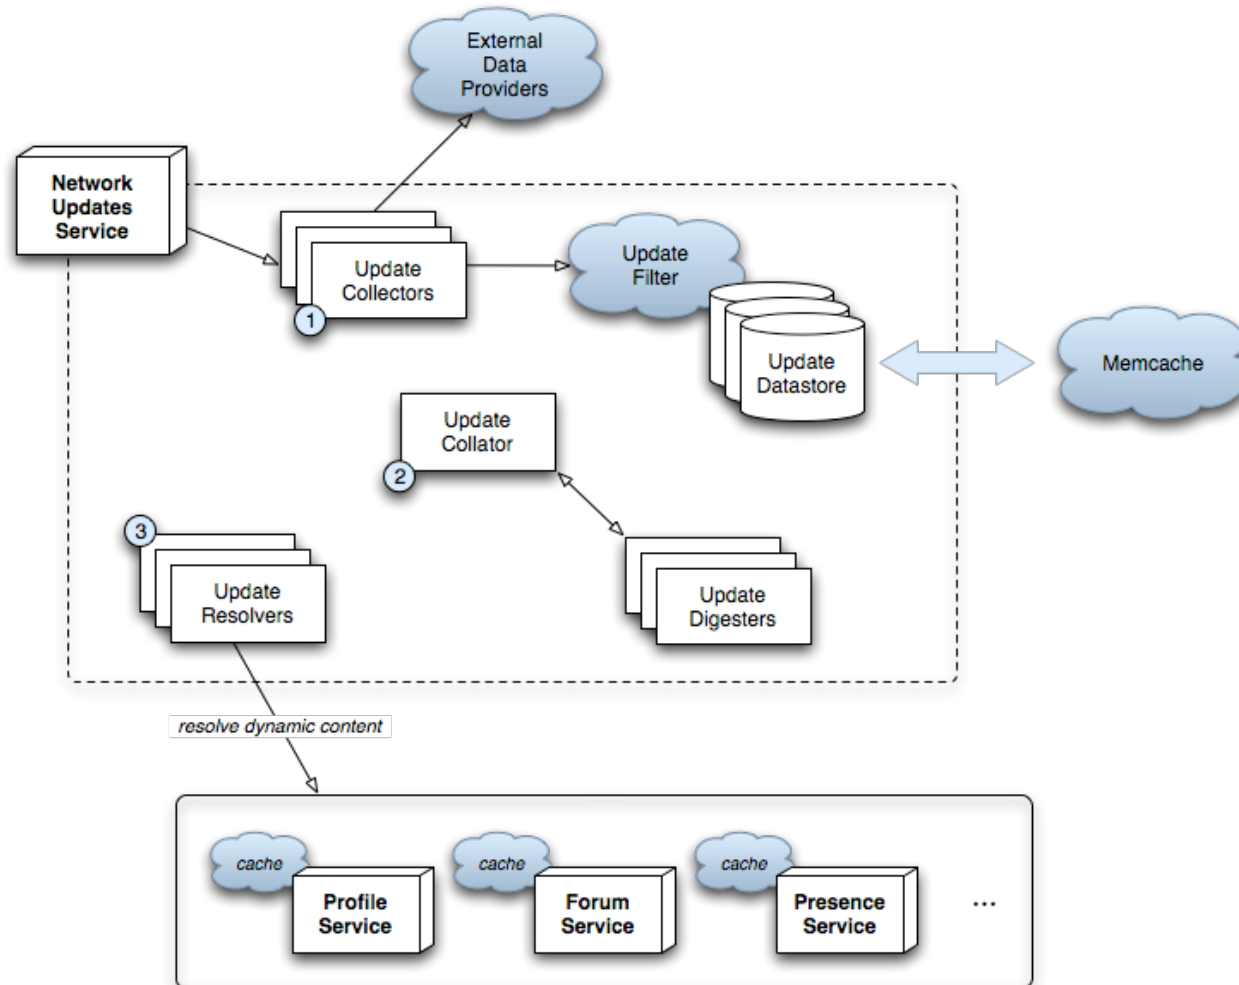
NetworkUpdateChannel channel =
    NetworkUpdateChannels.newMemberChannel(1213);

UpdateQueryCriteria query =
    createDefaultQuery().
    setRequestedTypes(NetworkUpdateType.PROFILE_UPDATE).
    setMaxNumberOfUpdates(5).
    setCutoffDate(ClockUtils.add(currentDate, -7));

NetworkUpdateContext context =
    NetworkUpdateContextImpl.createWebappContext();

NetworkUpdatesSummaryResult result =
    service.getNetworkUpdatesSummary(channel,
                                     query,
                                     context);
```

# System at a glance



# Data Collection – Challenges



- How do we efficiently support collection in a dense social network
- Requirement to retrieve the feed fast
- But – there a lot of events from a lot of members and sources
- And – there are multiplier effects

# Option 1: Push Architecture (Inbox)



- Each member has an inbox of notifications received from their connections/followees
- $N$  writes per update (where  $N$  may be very large)
- Very fast to read
- Difficult to scale, but useful for private or targeted notifications to individual users



# Option 1: Push Architecture (Inbox)

Push model

1 mailbox per user = constant read time

each updated is pushed to as many users as specified - hence multiplier effect

used when the distribution list is unknown or unique



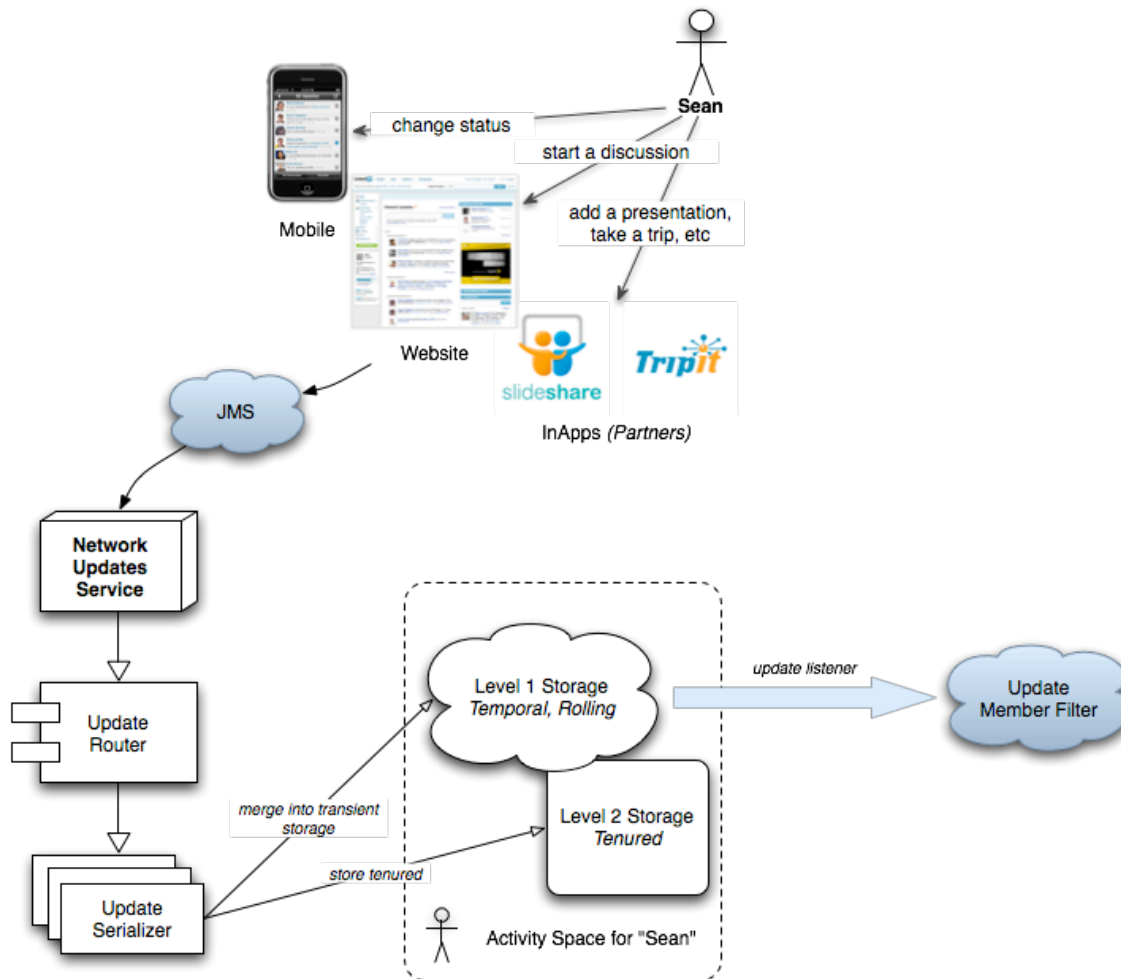


## Option 2: Pull Architecture

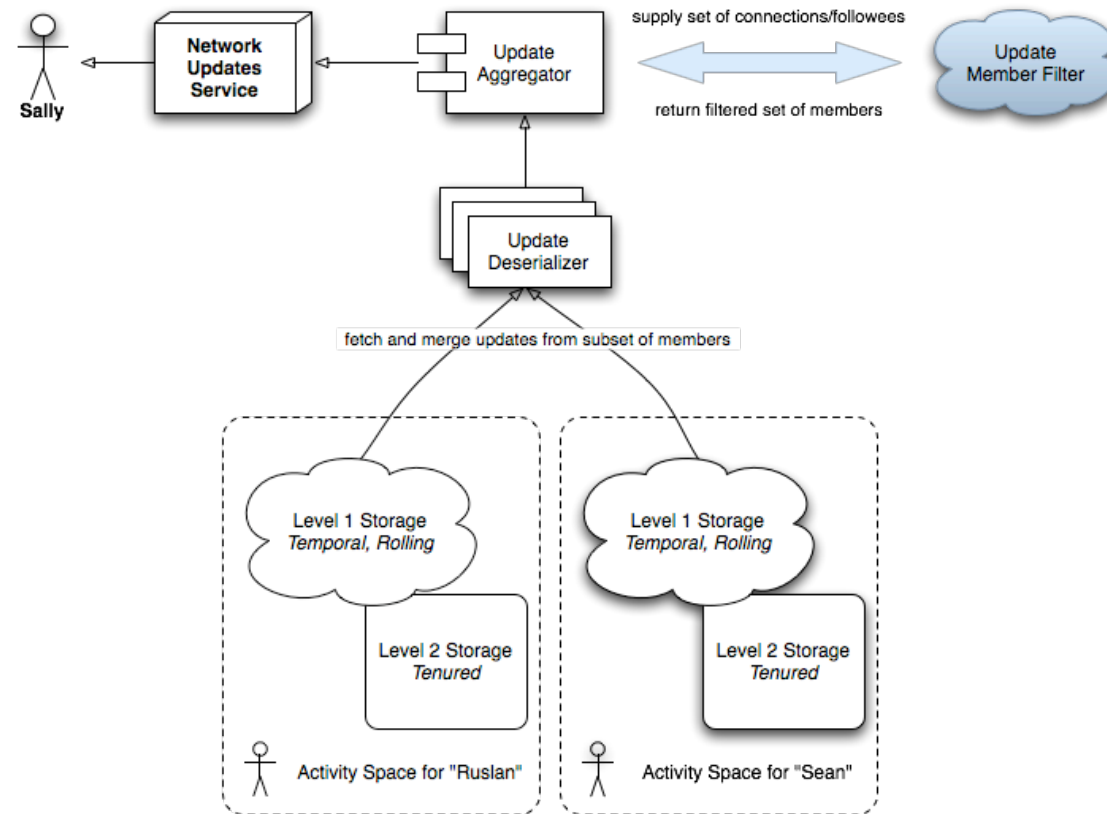


- Each member has an “Activity Space” that contains their actions on LinkedIn
- 1 write per update (no broadcast)
- Requires up to N reads to collect N streams
- Can we optimize to minimize the number of reads?
  - Not all N members have updates to satisfy the query
  - Not all updates can/need to be displayed on the screen
  - Some members are more important than others
  - Some updates are more important than others
  - Recent updates generally are more important than older ones

# Pull Architecture – Writing Updates



# Pull Architecture – Reading Updates



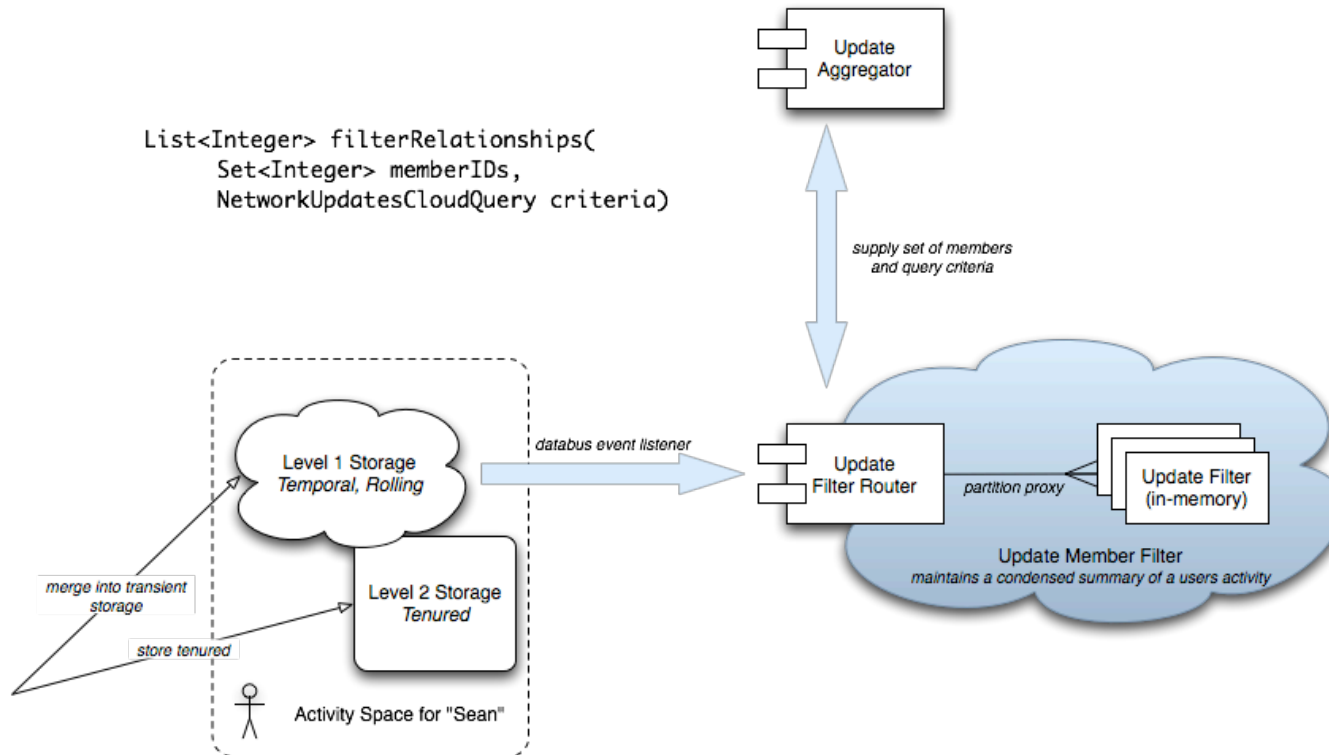
- L1: Temporal
  - Oracle
  - Combined CLOB / varchar storage
  - Optimistic locking
  - 1 read to update, 1 write (merge) to update
  - Size bound by # number of updates and retention policy
- L2: Tenured
  - Accessed less frequently
  - Simple key-value storage is sufficient (each update has a unique ID)
  - Oracle/Voldemort

# Member Filtering



- Need to avoid fetching N feeds (too expensive)
- Filter contains an in-memory summary of user activity
  - Needs to be concise but representative
  - Partitioned by member across a number of machines
- Filter only returns false-positives, never false-negatives
- Easy to measure heuristic; for the N members that I selected, how many of those members actually had *good* content
- Tradeoff between size of summary and filtering power

# Member Filtering



# Commenting



- Users can create discussions around updates
- Discussion lives in our forum service
- Denormalize a discussion summary onto the tenured update, resolve first/last comments on retrieval
- Full discussion can be retrieved dynamically

# Twitter Sync

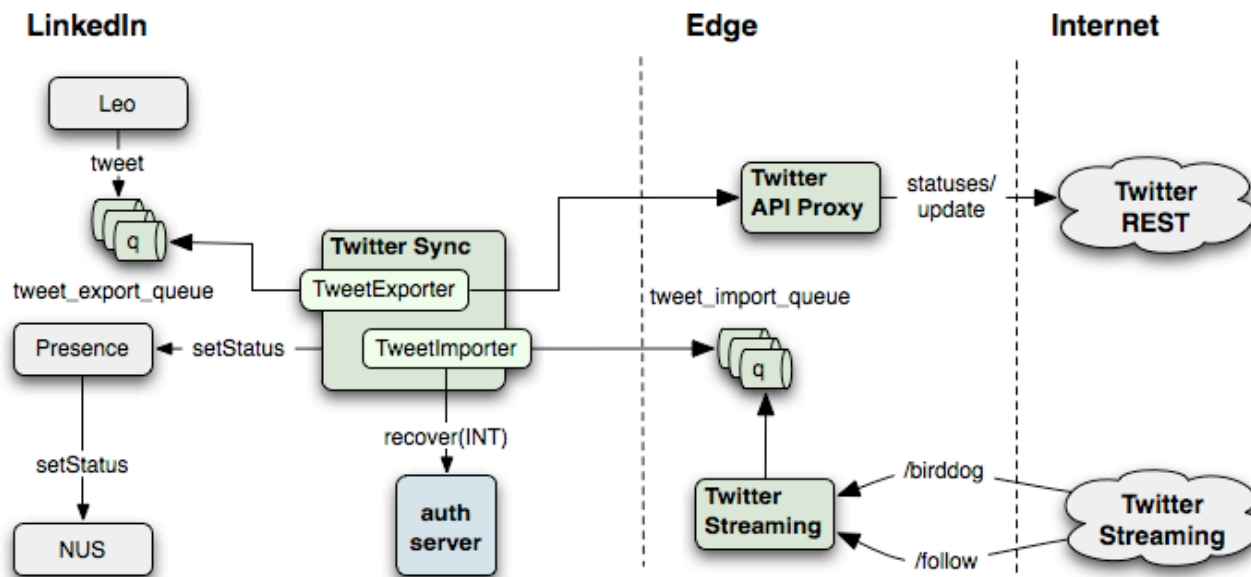


- Partnership with Twitter
- Bi-directional flow of status updates
- Export status updates, import tweets
- Users register their twitter account
- Authorize via OAuth



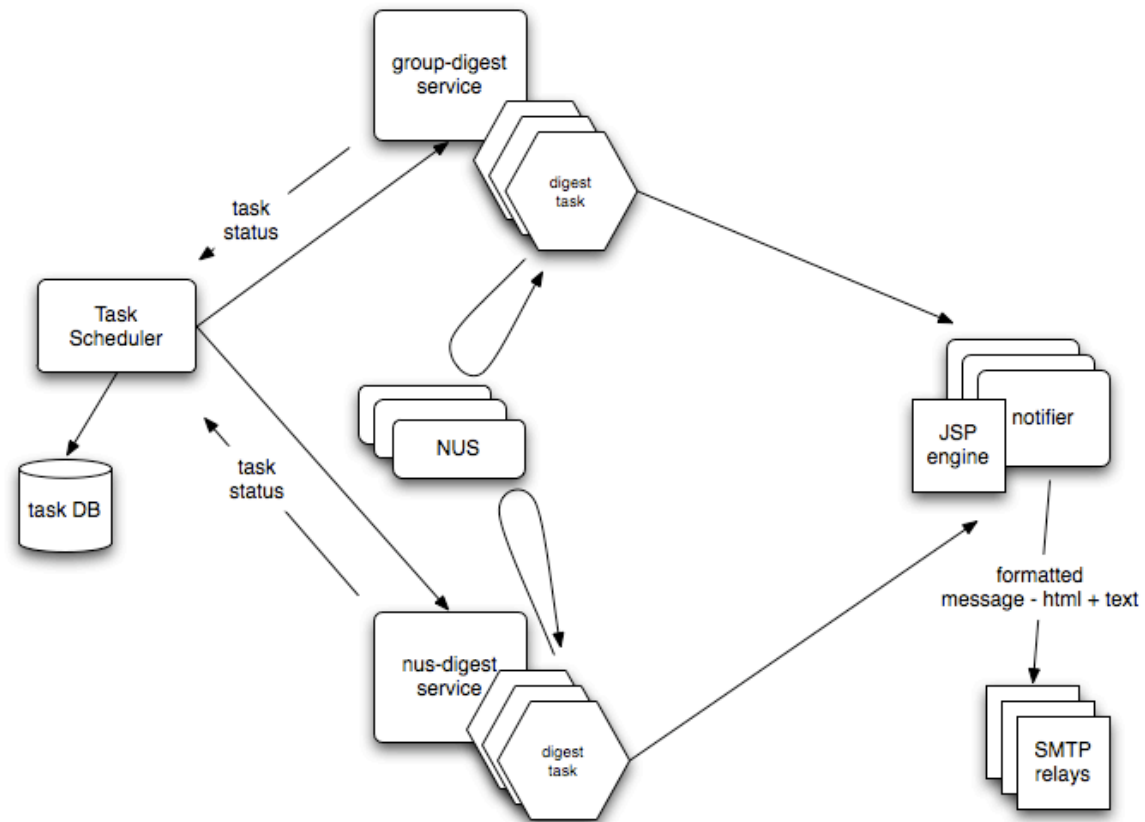


# Twitter Sync - Overview

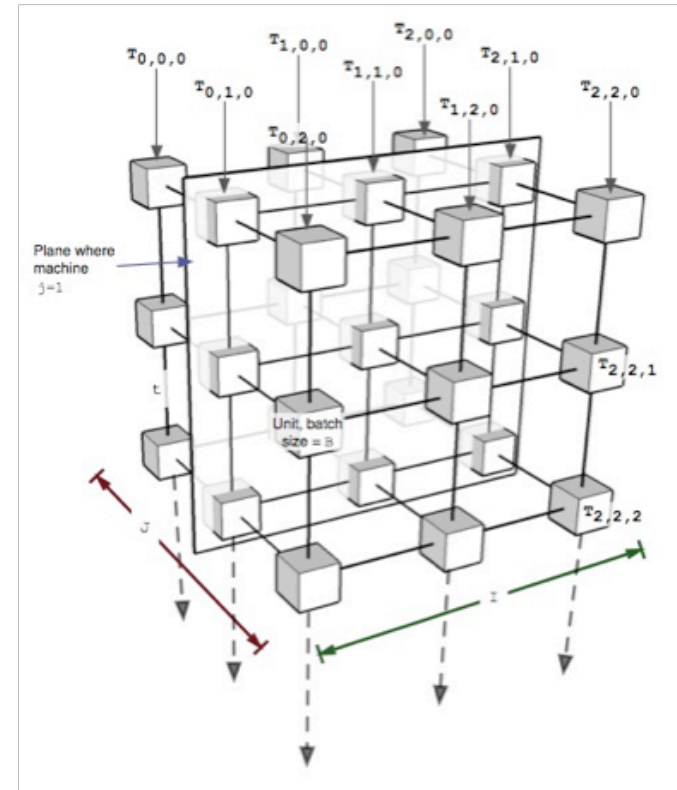
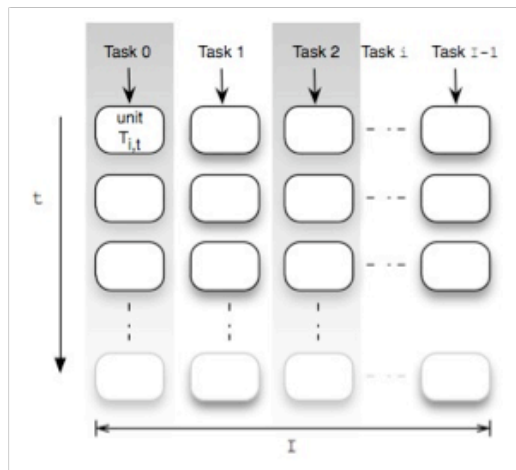


- Multiple concurrent email generating tasks
- Each task has non-overlapping ID range generators to avoid overlap and allow parallelization
- Controlled by task scheduler
  - Sets delivery time
  - Controls task execution status, suspend/resume, etc
- Caches common content so it is not re-requested
- Tasks deliver content to *Notifier*, which packages the content into an email via JSP engine
- Email is then delivered to SMTP relays

# Email Delivery



# Email Delivery



# What else?



Brute force methods for scaling:

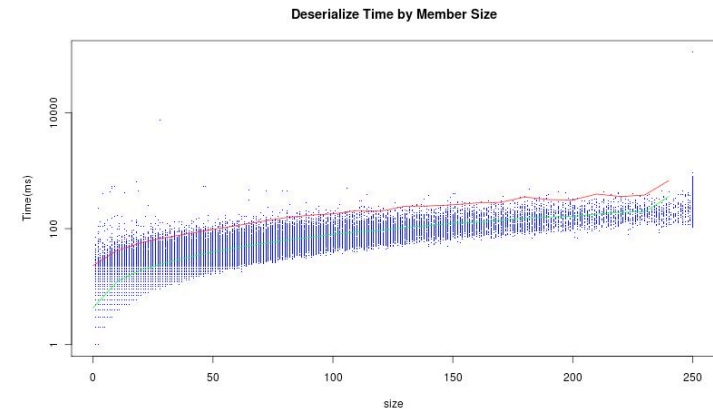

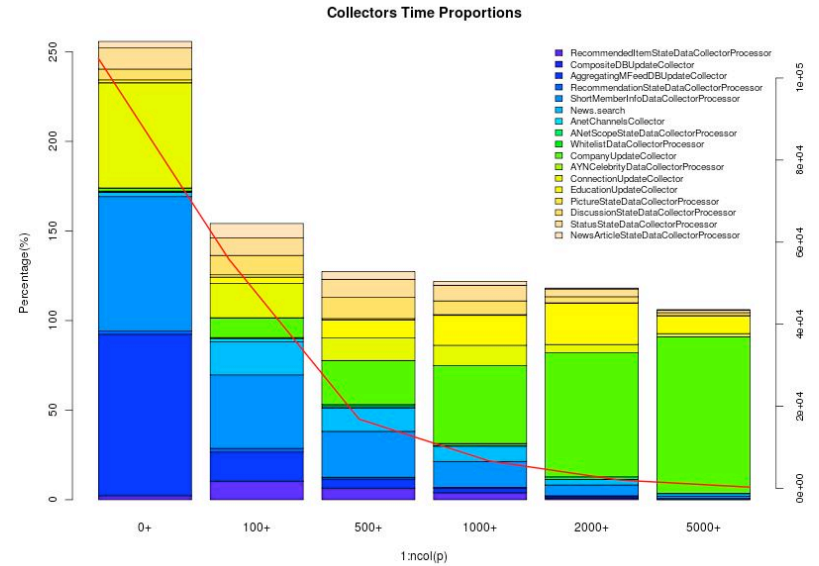
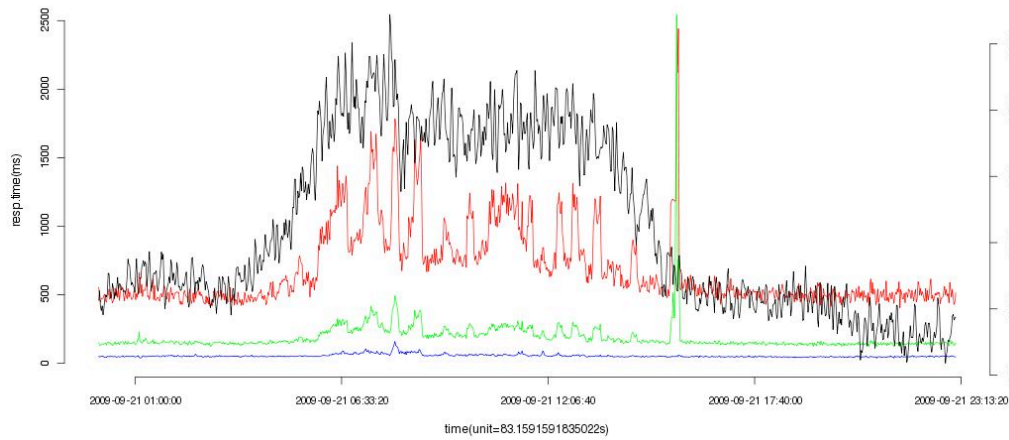
- Shard databases
- Memcache everything
- Parallelize everything
- User-initiated write operations are asynchronous when possible

# Know your numbers



- Bottlenecks are often not where you think they are
- Profile often
- Measure actual performance regularly
- Monitor your systems
- Pay attention to response time vs transaction rate
- Expect failures

# Measuring Performance



# Another way of measuring performance



It's your f\*\*\*ing site!

**David Henke** 2<sup>nd</sup>  
VP Engineering/Operations at LinkedIn  
San Francisco Bay Area | Internet



<b>Current</b>	<ul style="list-style-type: none"><li>VP Engineering/Operations at LinkedIn </li></ul>
<b>Past</b>	<ul style="list-style-type: none"><li>Senior Vice President, Engineering/Operations at Yahoo!</li><li>VP Engineering/Operations at AltaVista Company</li><li>Director of Engineering at Silicon Graphics</li></ul> <p><a href="#">see all...</a></p>
<b>Education</b>	<ul style="list-style-type: none"><li>University of California, Santa Barbara</li></ul>
<b>Recommendations</b>	1 person has recommended David
<b>Connections</b>	93 connections
<b>Public Profile</b>	<a href="http://www.linkedin.com/in/drhenke">http://www.linkedin.com/in/drhenke</a>




# LinkedIn is a great place to work





► [home](#) [culture](#) [benefits](#) [technology](#) [team](#) [jobs](#)

## Are you ready to build the next big thing?

LinkedIn seeks intelligent problem solvers who are inspired and motivated to change the world for good.

You  ?  
Apply Now!



 **Elizabeth Reaves**  
Assoc. Product Manager

Meet **Riccardo**

## Questions?

Ruslan Belkin (<http://www.linkedin.com/in/rbelkin>)

Sean Dawson (<http://www.linkedin.com/in/seandawson>)