

Financial Services

Insurance

Health Care

Java/J2EE

Mobile Platforms

Microsoft .NET

HIGH-PERFORMANCE NEIGHBOURHOOD SOCIAL PLATFORM

Case Study



Thapovan
In pursuit of perfection

Investment Management

Investment Banking & Brokerage

Custody & Clearing Services

Corporate Services

The Client

The client is an innovative start up building a game changing social network platform for neighbourhoods, similar to Next Door (supported by Google VCs). Using this platform consumers, small businesses and local neighbourhood experts can interact and exchange trusted information that results in better quality of information flow. The platform leverages strength of other social networking portals like Facebook, Google Plus, LinkedIn, and Twitter by interfacing with them and identifying the existing friends network in the digital social world. An endorsement or caution for a business or a question answered in the forum can be traced back to a social friend making the endorsement/caution, or the answering the question posted more trust worthy

The Business Challenge

The client was looking for an application development partner to create a state-of-the-art solution for their unique social neighbourhood platform. The platform was required to handle several million users with concurrent use by several thousand customers from across the globe. The various applications that were part of the platform needed to provide outstanding user experience and response speed, as part of a true service oriented architecture (SOA). The platform needed to support web as well as the emerging technologies like mobile platforms. The platform needed to handle a variety of rich media including images, voice and video.

The specific functions that needed to be accomplished by the platform included:

- Ensuring a distributed architecture with server clusters.
- Onboarding over 25 million active businesses across US neighbourhood using FACTUAL database.
- Implementing clustered SOLR technology with multiple core to provide Google-like user experience for various searches.
- Having a user management module with the ability to onboard several millions of users Interfacing with other social networking platforms like Facebook, LinkedIn, Google+, Twitter, to enable smoother registration and login and providing better user experience for new users.
- Leveraging strengths of other social networking platforms to reach out to friends in the social world and make an activity go viral instantly.
- Integrating e-mail blasting and marketing functions.
- Providing powerful administration functions that include message casting, screening and reporting.

- Providing rule-based crawler and extractor engines to crawl and extract data from 200+ real estate feeder sites and update the client's real estate database
- Implementing FACET technology in search results
- Implementing robust auto complete search with Edge n-gram and e-gram filters
- Integrating geo location services for neighbourhood
- Creating a content delivery network
- Providing a mobile interface

The client was also interested in a partner who could:

- Leverage the strength of social networking.
- Offer reliable, long term partnership.
- Provide highly cost-effective, best of breed solutions.
- Leverage multiple cutting-edge technologies including emerging mobile technologies.
- Provide outstanding problem solving skills.
- Respond quickly to issues and operational challenges.

The Engagement

The overall solution for the Social Media Platform included a variety of modules and applications. Given here are some of the highlights:

SOA cloud implementation

The AWS cloud infrastructure was chosen to provide clustered, fault-tolerant, an on-demand, and scalable solution. Multiple instances of servers, each with different software support and hardware configuration were setup to support search, database, large-scale images, and content distribution

Social Network Registration and Inviting Friends

User-friendly and quick registration and login was provided with existing social logins like Facebook, Google+, LinkedIn and Twitter. Once logged in, the application retrieves the friends list from the social world the user is associated with and provides the user with the option to invite those friends to join the client application.

Onboarding business from Factual DB

An intelligent program was built to onboard over 25 million users based on state and city, input

the details into the client's database, and index the data in SOLR. The engine constantly keeps the data current by syncing with the Factual database.

neighbourhood feed

A user-friendly way, with auto complete feature was provided for selecting neighbourhood and to show what is currently trending. A combination multi-core search, caching and sharing engine was implemented for quicker response. An efficient interface with other social portals was provided to highlight the trending that the user's social friends (from FaceBook, Twitter, Google+, LinkedIn) are associated with. The trending would show the details like the questions posted, recommendation/caution made, and questions answered.

The trending engine is implemented through a combination of SOLR CORE, neighbourhood caching and sharding mechanism

The feed can be customized for the user to receive feed only from those categories, people or neighbourhood that the user is interested in.

neighbourhood Search Application

Leveraging multi core SOLR server cluster, neighbourhood search lists questions, answers, recommendations/cautions and business for the keyword searched or lists matching business profile for the category/keyword match.

For the logged in user the search application also brings in extra credibility by displaying the association of the users' social friends with the business recommended, questions they answered and/or questions they posted. That is, if a business is recommended or question posted or answered by a social friend of the user it is highlighted. The application has intelligence to connect the user associated with the activity of recommendation, question or answer with the logged in user by cross checking user's friends list across all the social portals with which the logged in user is associated.

Asking Questions and Seeking Recommendations

A highly intuitive and efficient application allows users to ask questions about the neighbourhood and seek recommendation by selecting the category and tagging the question/recommendation. The application interfaces with the social portals and shares the content in other portals like Facebook (FB Wall), LinkedIn, Twitter, Google+ thereby reaching a wider audience and making the activity go viral.

Announcing a Service

This feature is a smart way to onboard individual service providers and small businesses into the social portal. The application leverages the service provider's social contact and lets them announce their service to all friends and friends of friends in the social world and to the current user base of the neighbourhood platform there by making the announcement go viral.

Endorsements

The application encourages users to play an active role in the neighbourhood by creating an intuitive way to recommend or caution against a business and share it virally with social friends and friends of friends. The endorsement can be for a business existing in the system or for a business the user can introduce and recommend. When a business that is not part of the system is recommended, an email is automatically sent encouraging the business owner to be part of the neighbourhood platform.

Administration Application

An intranet application is designed for use by the client internal administrative team to manage and administer this neighbourhood social platform. The application includes, a Message Cast System to manage promotional campaigns, a Screening application to monitor and filter unwanted content, a Reporting System to support a wide range of preset and ad-hoc reports, and periodic mailers to reach out to the social friends to keep them in touch with the current trending

Real Estate Search

This sub-module is meant exclusively to search for a house/apartment on sale/rent. This module also provides a feature for the users to list their apartment for rent/sale. The crawler/extractor engine crawls and extracts data on a daily basis, through partnership, from real estate data providers like Trulia, Zillow, Backpage New York. The extracted information is then interfaced with Google Location API and Maonics to extract the neighbourhood details and lat/long information based on the address of the property listed.

The application also identifies the properties in Google Map. The application supports various forms of search supported by a robust ontology that helps in mapping keywords typed in by the user like balcony, sit-out, and terrace to the features associated with the property.

The robust crawler/extractor search platform was built with Hibernate, JAVA and multi core clustered SOLR setup.

CDN

In general, extensive media files are associated with profiles, businesses and real estate listings. As media files are large in size, we used CDN to serve the media files to the end user from the location nearest to the user thereby improving the response time. The images in CDN are kept in sync with real time data.

Ambassador Program and Payment Gateway Integration

Active user participation is closely tracked and encouraged by assigning points to each activity performed by the user. The points are accumulated and cash awards are given to the top 3 via payment gateway that deposits the cash awards in their respective financial institutions.

Caching/Sharding

To ensure a quick response and display multi keyword search results within a couple of seconds, the frequently searched queries/query results, frequently viewed profiles, and home neighbourhood are all cached. The least frequently used (LFU) algorithm is implemented. The combination of REDIS and the clustered MULTI CORE SOLR server is used to achieve the desired caching/sharding objectives

Maintaining Private Lists

This feature allows signed-in users to create exclusive private lists by leveraging their contacts in social networks. The application provides a facility for the users to interface with their current digital social network, select the desired friends from each of the social networking sites (Facebook, LinkedIn, Twitter, Google+) and create their own exclusive lists to share their activity in the client's social networking platform.

The private lists are especially useful when the users wants to caution about a business that they do not want to share in public.

Clustered Multicore SOLR Search

SOLR was chosen as the desired platform for its robust implementation of the search function. SOLR supports different types of keyword searches including search based on geo location using lat long, faceted search, and its ability to group the results all at an incredible speed.

Also, it helps in keeping the DB normalized by shifting the deformalized view to SOLR thereby

avoiding data redundancy in the DB. The use of SOLR minimizes the DB hit, thereby improving the performance manifold.

Maponics/Google Geo Location API

The geo location service strength of Maponics and Google's geo location API are leveraged to provide location based services like identifying the locations in the map, searching in a specific neighbourhoods, and distance calculation.

The Engagement

Thapovan was responsible for identifying and implementing the most appropriate technology for the various applications to meet the performance, cost and usability objectives set by the client.

Major Technology Components

- Language: Java, Php, Objective C, jQuery,
- Webserver: Apache, Tomcat
- Database: Mysql, Hibernate
- Search Engine Api: Solr
- Caching Technology: E-cache, Redis
- Architecture: Mvc – 2
- Frameworks: Spring, Zend

Key Technical Tools Used

E-CACHE, REDIS: Free, open source, high-performance, distributed memory object caching system used to reduce database load and improve performance.

CLUSTERED MULTI-CORE SOLR - High-performance, full-featured text search engine library written entirely in Java and PHP. Used for greatly enhancing text based searches resulting in outstanding response times.

-
- MAPONICS, GOOGLE API
- AJAX, JSON
- JSON for easy retrieval of the data in the PHP pages.
- Distributed Database: Better performance, fault tolerance and failover
- Session Tracker for auto logout of inactive sessions
- J-METER for performance testing

Results Achieved

The client achieved the following major benefits from the partnership:

- Building a complete state-of-the-art neighbourhood social platform.
- Rich functionality, high-performance, high availability, ease of administration.
- Highly cost effective and significantly better ROI.
- Excellent client customer adoption and user base growth.
- Joint development led to a better product brought to market faster.
- Efficient use of technology led to a wonderful user experience.