

# Product Search

## collaborative product plan

**(Status: way outdated, except for Vision and Principles)**

### 1 Basics

#### Long-term vision: Shopping on Google

- Google as the web's best (easiest, fastest, most trustworthy, most comprehensive) place to start shopping for anything
  - not just where to buy, but also what to buy
- Goal: Provide a better experience for product queries on Google, growing the number of product queries on Google
- Acquire, clean, normalize, rank product data not just for use on product search, but also Google.com search results and ads
- Vast majority of users from google.com, but a vertical dedicated to organizing the world of products for sale must exist
- All products, all locales, all languages

#### Some guiding principles

- google.com integration as a foundation rather than an afterthought
- meet quality issues head-on as a primary focus: search quality, exposure quality, normalization quality, feed quality, attribute quality...
- start small, manageable and successful (most popular and easiest product verticals, online merchants only, in US English) but plan for much broader scope
- aggressively acquire and use vertical-specific product catalog data where available; best-effort algorithms for long tail
- start from search use cases and build out toward browse use cases eventually
- start from 'where to buy' use cases and build out toward 'what to buy' use cases eventually
- start from online inventory and build out toward local inventory (from big chains)
- start in US English and build out toward global eventually
- feed Checkout ecosystem because the more we know about merchants, the better job we can do
- separate infrastructure from user experience

### 2 Rough roadmap

#### Immediate focus

- Assemble the right cross-functional team to create a successful product
- Build EMG trust by quickly executing on set of Checkout-related enhancements and minor UI improvements to the existing Froogle site - see (3) below for mocks and details
- Begin determining how we want our feed specifications to change
- Begin executing on the short-term sales strategy articulated by Vince & team to increase coverage of good quality Checkout-enabled items
  - Checkout merchants in Base (lowest-hanging fruit)

- Base merchants in Checkout (some integration pains)
- Go after top merchants in neither Base nor Checkout by vertical
  - SKU coverage
  - Merchant quality
  - "Must-have" brands

## Medium-term focus

- **Great quality exposure of a Product Search experience for consumer electronics and media (books, CDs, DVDs, games) queries on google.com.** This could be via onebox, megaonebox, blended results, or whatever. Don't worry much about other verticals at this point.
  - Data acquisition
    - Acquire catalogs (CNET, Ingram, Muze, sources of UPC info, etc): human task - Content team.
    - Acquire buying guides for consumer electronics and media: human task? Super high precision needed.
    - Acquire reviews for consumer electronics and media: Seed with sites chosen by humans?
    - Acquire tax tables from merchants and third parties
    - Acquire shipping tables from merchant and third parties
    - Create new feed specs for consumer electronics and media products. Use catalogs to suggest fields.
    - Feed quality
      - Merchant credit card verification
      - Images
  - Infrastructure and search quality
    - Eval framework
    - Universal search infrastructure work
    - Product ranking work
    - Google.com triggering conditions (not just when to trigger, but which experience to trigger - navigational vs. category, etc)
    - Evals galore
  - Data understanding
    - Normalize offers into the ground-truth product buckets mentioned in the catalogs: precision more important than recall.
    - Categorize offers into the ground-truth category hierarchies mentioned in the catalogs
  - UI and features
    - Begin work on new XFE-based frontend
    - Determine reasonable ways to expose a Product Search experience on google.com for these product verticals. See some [toy mocks for product-navigational queries](#) and [toy mocks for category queries](#)
    - Define preliminary UI for "category" pages (Products > Electronics > Digital Cameras) and "product" pages (Products > Electronics > Digital Cameras > Canon Digital Rebel).
    - Checkout integration tasks
      - Tax/shipping/availability
      - Buy from a Google page
      - Fraud, popular items, etc
  - Improved snippets and result diversity on Google.com results pages for commerce queries

## Later (but still really important)

- Next several high-impact verticals.
- Revive structured data crawl for offers.
- Reintegrate local data.
- Work on best-effort normalization algorithms for long-tail.
- Work on best-effort categorization algorithms for long-tail.
- Integrate attribute extraction efforts such as Merlot.
- Invest more in webpage understanding to improve coverage of reviews, buying guides, etc.
- Build out better browse functionality on the vertical.
- Integration with other Google properties
- Internationalization

## 3 Details about short-term release

- Goals: Streamline interface, drop features with low usage and/or quality, debrand, enhance with Checkout, build team momentum and executive trust
- Methodology: Update existing CFFE
- Mocks (but see "Does not include" section below)
- Changes include:
  - Replacing the Froogle logo with a Google logo pointing to [google.com](http://google.com) and eliminating/replacing all other mentions of the Froogle name
  - Replacing Froogle green colors with Web Search blue colors throughout
  - Moving the block of refinements to the bottom of the page
  - Removing the Checkout restrict from the block of refinements and putting it at the top of the page with the alternative, more aggressive UI
  - Removing Local integration
  - Removing Grid View
  - Ensuring the Checkout restrict and Stores restrict at the bottom of the page play together correctly
  - (Maybe) adding promo badges beneath results from Checkout merchants for non-Checkout users
- Does not include:
  - Tax and shipping estimates
  - Ability to change shipping options
  - Buy buttons on results page
- Timeline
  - Rough demo in corp: Friday 3/16
  - Persistent demo in prod: Thursday 3/29
  - External launch: Wednesday 4/14 (debranding may slow us down due to dependencies on GWS)
- Marketing plan for rebranding (Tom to insert)

## 4 Open questions

**Data acquisition and quality (Owners: Vince, Karen, Cynthia/Benita, Tim/Yukari)**

**Base Team Summary in bold; Base Team members individual input in various colors.**

- What % of our product listings in Base today currently come from Checkout merchants? (i.e. if we restrict to Checkout only, how much of the product index

- are we losing?) How does the raw number of unique offers compare to competitors such as Amazon, eBay, Yahoo Shopping and shopping.com?
- How shall we get rid of spam in Base? Can we simply do a merchant credit card verification?
    - **Base Response Summary: CC verification would definitely help, although it won't completely end spam/abuse.**
    - No. Credit card verification is not enough. With identity theft increasing, its very easy for people to get a hold of other people's credit card number and verification processes.
    - More than one verification- There needs to be more than one way to verify identities to eliminate spam.
    - Better aggregation of merchant reviews from review sites; give greater prominence to user reviews; make it easy to view and rate merchants based on buying experience
    - A lot of spam we receive comes from single item posts. We've developed signals that address this issue, but we really need to move to an account review system.
    - We make the system too easy. Some financial verification is good, but may not work for all merchants. In that case, and I firmly believe Base should be free, we should have them put down a small deposit (\$50-100) that will refunded after they meet some quality rank. Satisfied sales, sales, months active, etc.
    - I think credit card verification would be great. It won't be enough to completely *rid* Base of spam, but it'll likely help immensely to combat fraud, and your simple spammer isn't going to have access to a huge array of stolen identities. Perhaps cross-breed with Checkout and require people to have an active credit card or bank account registered in their Checkout account. This would create more Checkout sign-ups, too, and likely inspire people to sell their items using Checkout.
    - No tolerance spam policy with merchants. Clearly lay out guidelines on what counts as spam. If you are a first time offender, your account gets shut off for 2 months. If you do it again, you get shut off for 6 months. Do it again and you are banned for life.
  
  - What % of Checkout merchants currently give us data feeds via Base?
    - **Base Response Summary: Unknown, but moving forward, we would like to come up with a way to flag and measure.**
  
  - What % of Base merchants are Checkout-enabled?
    - **Base Response Summary: 13% (= 8454 thirty-day active CBG enabled merchants/65038 thirty-day active merchants).**
  
  - What changes should we make to our baseline Product feed spec? Should we really create another one, or can we simply accept a competitor's feed format that the merchant will already have ready?
    - **Base Response Summary: Support for alternative formats is great for those who want it. However, change is costly and confusing for providers; we have to be wary about making drastic changes to the core specification.**
    - If the idea is to gain large merchants, then perhaps it would make sense to come up with a system that is easier for them to submit the same feed format to us to put online
    - If we choose to do a different feed format, we should make it easier for merchants to do this through some tool that we provide.
    - The easier we make it for our customers, the better.
    - The better quality data we can get from our customers in the most efficient way, the better.

- While it is a competitor's feed format, this doesn't mean its bad. We need to look at the best way for them to give us content as well. And of course we want to get better data than our competitor, so if possible, we can use a similar format, but revise it a little so its google specific and gives us better quality and more structured information to return the best search results for users.
  - Change is hard and confusing for providers. Just changing from Froogle -> Base specs took a long time. Keep our specs the same unless there's strong case for change.
  - Improve Store Connector to better handle competitor feeds (eBay, Yahoo, Amazon and in that order of importance). Promote Store Connector more once it's more polished.
  - Definitely keep the same specs, but we should expand to include additional types of feeds. We constantly get people uploaded eBay / Yahoo store export files or shopping.com files. We should write the software to include these as well. If anything, we should be embracing any kind of upload and normalizing it on our end... make it as easy as possible for the providers and they will come to us.
  - More formats = better. While I'm sure that some of them would be a nightmare to program for, I bet a lot would be pretty simple to adapt into formats that we already accept. At the very least, a tool to convert other formats to our formats would be great. (Import a file, select "convert Bizrate to Base", go!)
  - I think making more changes will just confuse our providers even more. Our spec is already pretty complex. I think we should provide tools that converts other competitors' format into Google Base format supported format.
- How granular a product vertical deserves its own feed spec? Consumer electronics? Cell phones? How many feed specs do we need to do a good job on consumer electronics and media? Who creates feed specs? An ops team of "product vertical specialists"?
    - **Base Response Summary: More granularity is better for us, but we shouldn't overwhelm our merchants with choice. The more we can automate verticalization on our end, the better.**
    - The more granular the better. That way we can trigger more meaningful filters for specific product searches
    - We should look at most popular product searches on .com/Base/Froogle and start from there
    - Consumer electronics - TV, DVD player/VCR, Camera, Camcorder, audio players, phones, cell phones, video games, home theater systems, PDAs, computers, printers, car stereo, GPS devices
    - Ops team can create feed specs
    - We'll be able to categorize items better if we have our providers select "sub-verticals" under Products instead of having them type in whatever they want in "product\_type" attribute.
    - Up to this point, granularity is (supposed to be) achieved with attributes instead of verticals. For example, a merchant can supply "Cell phones" for the "product\_type" attribute.
    - The issue is that providers don't really supply very useful information for the "product\_type" category. However, with a extremeley granular system, where do items such as "iPod accessories" fit in? I think it could be useful to have specific feed specs but enforce our "required" attributes more strictly (i.e reject feeds with incorrect product\_types, etc).
    - If we're moving to a normalized system, categorizing products should be relatively easy depending on the catalog data that we have. We could have a pretty nice hierarchy set up. Things don't HAVE to only be in one

category, either. Headphones can be in "Consumer Electronics --> iPods --> Accessories" AND "Consumer Electronics --> Audio --> Headphones".

- Specialized knowledge of the vertical should be required for the feed spec development. Collaboration with key vendors in each vertical would also be helpful
- Do we need product specialists to help compile lists of good graduation gifts, Valentine's day gifts, etc?
  - **Base Response Summary: Humans can add value, especially as we're building out the feature. However, we shouldn't build features dependant on human support.**
  - I strongly believe some form of editorial intervention would provide a superior overall experience than focusing just on search. For example: We can compile weekly lists of "top Google Checkout deals" and other types of recommendation lists.
  - Or regularly "feature" a small-medium merchant to provide exposure and build a sense of community or at least put a human face on things.
  - One great feature of Amazon is that it provides a useful recommendation list based on items you've purchased before. Clearly, we won't have a history of users purchasing history, but we could provide some kind of "top selling" products list, which would definitely help improve the overall search experience for users (plus, those top sellers are more likely to be trusted providers).
  - I think putting one person on the task of compiling lists of "good" items forks away from the path that Google normal takes. I would rather see an automated view of popular items or building trends - a zeitgeist for products. Human involvement would be at a meta-level to choose which categories to use.
  - How much data do we have on our users? Do we keep track of the ads that we serve for them? How much search data do we keep on them? We might very easily be able to do an Amazon-like product recommendation based on all the information we already have.
  - We can look at what were the most popular product searches on Google for each season (easter, graduation, july 4th, christmas, valentines day, etc.) and we will be automatically able to generate a list of popular gifts (with some human intervention.)
- What should be our product hierarchy? Can we reuse a competitor's hierarchy or do we need our own?
  - **Base Response Summary: We should go with the most widely used hierarchy.**
  - I know there isn't a set standard product hierarchy that exists (especially with the various information we get from third-party providers), but we should go with the most widely used categorization. Amazon would probably be a safe bet, especially since they've put a lot of research behind their hierarchy.
  - There must be some industry standard for product hierarchy that we can use. If not, we should move towards creating a standard product hierarchy. Last resort Amazon.
  - Amazon's hierarchy would be the best to copy, in my opinion. It's easy to follow and very, very well-defined.
  - Use Amazon's as a basis and can add to it.
- How much human involvement do we need supervising normalization and categorization algorithms?
  - **Base Response Summary: This has to be automated, we used to categorize manually and it didn't work (even with a much smaller**

**index). Our users can provide a great "review" force to help us here.**

- I don't know a lot about algorithms, but from what I do know, we shouldn't need very much if the algorithm is designed correctly. Moving back to the old system of manually categorizing every product that comes in would **NOT** be a good thing.
  - While I agree that this should be done automatically, I remember Froogle's clustering and the nightmare that was. If the normalization is not done right, then it creates a huge headache for ops. Having most things done automatically but giving ops a tool (or information to give to CPs -- "put this UPC in your item") to assist with the outliers would be optimal.
  - We need to do this in the beginning of each vertical "launch" so that our classifier can learn.
- How shall we approach the problem of inventory tracking? Will merchants be able to supply the data we need?
    - **Base Response Summary: We have an instock field -- we should use it.**
    - The limited number of merchants that do have good inventory information will likely be hesitant to give us this data. Instock is a value that we are likely to get better data for. We could adapt this sort of attribute to find out if they have nothing in stock, limited availability, or a lot of stock. If it has limited availability, we can direct users to check the merchants site.
    - Exact inventories, like those needed for full checkout integration, will be next to impossible to get, especially for brick and mortar stores, which may only update their online inventories on a daily basis or so. A boolean value, such as the instock field, may be more useful in this case.
    - This is an extremely tough problem for most multichannel merchants. We can work on a channel partnership strategy with database companies to see if we can develop an easier way we can get the data from the source rather than have them extract and manually feed it to us. If we want to be extreme, we can outfit stores with our own POS system and can feed in all relevant data to Google
  - From whom can we get detailed tax tables and shipping tables? (Supposedly there are third parties large merchants contract to for this.)
    - **Base Response Summary: There are tables that exist on multiple sites. For shipping rates, we need to be careful because they are negotiated company rates. We probably should be on par with Amazon.**
    - Sales tax table mapped to zip code from Intuit - <http://www.eclipse.intuit.com/solutions/adv/tax.asp>
    - USPS Shipping tables: <http://pe.usps.gov/text/dmm300/ratesandfees.htm>
    - UPS are available starting at page 20: [http://www.ups.com/media/en/2007\\_retail\\_rsg.pdf](http://www.ups.com/media/en/2007_retail_rsg.pdf)
    - DHL and FedEx are probably available on their site, but I couldn't find them with a quick browsing.
    - We need to keep in mind, though, that although tax rates are standard, shipping rates are NOT. Just because USPS says it costs "blah" to ship an item, that doesn't mean that the store isn't charging "blah+\$20" or offering free shipping. Most shipping information should come directly from the companies themselves with the **option** of using *customizable* shipping calculators.
  - What should be our strategy for acquiring local data from merchants?

- From GPS: focus on big chain merchants (e.g. Best Buy, B&N, ...) rather than one-off small businesses
  - **Base Response Summary: Going after larger merchants is the first step. To go after the long-tail we'll need to have an easier way to provide this data.**
  - Small merchants need an easier way to submit their data to us. If we can help bridge that information gap for them, we can get their content easier.
  - Large merchants also need an easy way to give us data, but they will have a tech team responsible for doing this or a tech person.
  - We should push the Intuit initiative and large chains to get this data - a lot of good data could be extracted from new users of Quickbooks. We could also partner with a site like nearby.com or try to replicate what they have.
  - Add a link similar to [Add/Edit Your Business](http://maps.google.com/Add/Edit>Your/Business) on <http://maps.google.com/>. This link should say something like "List your products on Google."
  - This is an extremely tough problem for most multichannel merchants. We can work on a channel partnership strategy with database companies to see if we can develop an easier way we can get the data from the source rather than have them extract and manually feed it to us. If we want to be extreme, we can outfit stores with our own POS system and can feed in all relevant data to Google.
- To what degree do we need to involve humans in judging data quality?
    - **Base Response Summary: Humans can add a great deal of value to a heavily automated/optimized system (e.g. web search and SQE). However, the first step is working towards a system and is optimized automatically.**
    - Humans will be involved in data quality proportional to the ability of technology to meet the desired level of data quality. They will bridge the gap between what the automated systems can accomplish and the goal. It is my opinion that it is not possible to completely remove humans from the process, but that a system where a relatively small number of humans control powerful tools for data quality and perform higher level ("account" level to meta level) analysis is optimal.
    - Humans need to be involved in judging, but humans are also bias and inconsistent. The better we can regulate, streamline, and automate things based on effective signals, the better.
    - Humans judging data quality is limited as well because it takes us more time and can be inaccurate across different people
    - We should leave the manual judging to the very top cases where a human is needed.
    - Members of the Base Ops team has made a significant impact on data quality thus far. Human judgment is often necessary to determine spammy affiliate networks or sites, spammy IPs, etc. There's also no (good) landing page checking tool yet. Need to ensure correct prices, descriptions, working links, etc.
    - We can probably shift a lot of quality checks to an automated system. A lot of our current efforts are still pattern based which a computer should take care of.
    - With better merchant quality signals, we could wipe out a lot of poor quality items. Integration with Checkout data and AdSense data would be huge.
    - Many affiliate and dropship sites would be difficult to detect automatically. With any sort of creativity, people could bypass most automatic detections that we come up with. Considering that one of the most frequent complaints about Froogle and likely the top reason it's not being used is that it has low quality results, we should do whatever it



takes to keep quality as high as possible. I don't believe humans can be taken out of the equation. Both reactive and proactive measures should be taken.

- I agree with Marc that a lot of affiliate and dropship sites will be very hard to detect automatically. We need human to detect if they are the rightful content owner or they are just an affiliate uploading and earning commission. Also some policy violations such as borrowed content will be hard to detect automatically, and will need human enforcement.
- We need the human involvement in the beginning. We can teach the machines to do this eventually but we will have to make an investment of lots of human capital to overcome the existing sentiment that our data quality is poor.
- Besides merchant feeds, what other sources of data are available? Which catalogs would be helpful?
  - **Base Response Summary: Please refer to the table in the following trix**
  - <http://trix.corp.google.com/csc?id=o05015737939722549487.2476643432906340570.01370208361647569547.4956>
- Why are so few merchants/providers responding to the Base attributes edict?
  - **Base Response Summary: Merchants aren't clear on attributes (their value and/or how to use them) because we aren't clear on them internally. If we could demonstrate their power in our results, the industry would follow.**
  - Probably because of a (perceived) results vs. effort analysis. If a company thought that they would receive a "significant" number of sales (whatever that means to them) by following our advice, then I think they would do it. Companies seem to have no trouble adapting and bending over backwards to be a part of Google.com because they realize the importance. With product search to date, the pay off for merchants has been low, with some exceptions, especially in cottage merchants.??data??
  - Communication- We change our required attributes frequently, and then change the type of data values in these attributes frequently. Communication to our users needs to be solidified and defined. Us asking them for information and stressing its importance one month, and then us asking for different or revised information another month, is not effective in communicating standards and what needs to be done on the user end in order to achieve search results.
  - Attributes are a nuisance at times and hard to correctly format. They all have different formatting requirements, different character limits, and these are all very specific. While we see the point of this and why its necessary, perhaps we can explore a different and more effective way to get the same kind of information from users- without having to be so particular in how users manually submit their data to us.
  - Most product data is locked in an unstructured format. It's too much effort to pull this out consistently. We need to gather better product data from manufacturers or wholesalers. I'm not sure how much value merchants add to the quality of the data.
  - Upload process is very confusing already for many small merchants, and making drastic changes is just beyond their means. Most of them really only upload right now because it's A) Free and B) Google. If we really want to succeed, we need to make it as easy as possible for them to get products live, make updates, etc.
  - We need to enforce the changes and requirements we make-otherwise, there is no incentive for a merchant to change. We also need to more specifically spell out benefits of changing the feed

- What should we do about merchants who want to upload but use third party redirects for tracking? Policy question

### **Search and data understanding (Owners: Ashutosh, Jen, eng)**

- Is there anything we can do to improve our exposure via onebox in the short term?
- What do we need for a search quality evaluation framework?
- What kinds of signals from Checkout would be very useful in ranking?
- What are the major hurdles in google.com integration?
- What is a reasonable approach for product normalization in the long tail?
- What is a reasonable way to do product categorization for the long tail? (One of the Risk engineers has a preliminary approach with promising results)
- How to address the problem of attribute extraction
  - Some existing efforts at Google
- To what degree should humans be involved in feed understanding efforts?
- What feed understanding metrics should we track?
- What ranking metrics should we track?

### **Frontend (Owners: Dan, Jen, Jonathan, eng)**

- Should we develop our own, new frontend or rehabilitate CFFE or ASFE?
- How exactly does the ASFE change ranking? How is it different from CFFE?

### **Business, strategy, marketing (owners: Tom, Jeff, Vince, Benita, others)**

- What would be the impact on eBay and Amazon if we were to cannibalize their traffic from ads/search?
- What would be the reaction
- What is the B2B opportunity?
- How might our ads business need to change to accommodate Product Search?
- How would merchants/the industry react to Google selling from Google pages / disintermediation?
- How would merchants/the industry react to Google asking to place a consistent UI element (e.g. Checkout "add to cart" button) on the merchant site?
- What are the next countries in which we should expand?
- What metrics could we use to measure the success of Product Search driving Checkout vs competitors (both retailers and aggregators/agents) like Amazon or eBay? Possibilities (old):
  - Product Search merchant page clickthroughs vs. 'product page' views on a competing site (probably impossible to get good stats)
  - # conversions generated by Product Search merchant page clickthroughs vs. units sold on a competing site
  - GMV generated by Product Search merchant page clickthroughs vs. sales on a competing site
  - Product Search resultspages vs. resultspages on a competing site (probably impossible to get good stats)
  - Product Search searches vs. searches on a competing site
  - Product Search user sessions vs. user sessions on a competing site

### **Product (owners: Jeff, others)**

**Base Team Summary in bold; Base Team members individual input in various colors.**

- Shall we eventually accept jobs or services for sale? Cars? (ugh)
  - **Base Response Summary: We need to get retail products right before we try to incorporate related verticals.**

- No. Or at least, job/services/vehicles/housing/etc. differs enough from product search (and online purchasing) to be kept separate.
  - I think we should get online products right. Then move on to cars. I think it will be much more challenging to get good structured data for jobs and services.
  - Also, Services and Jobs are very location-dependant. You can't ship a house-cleaning service. These should probably be more tied in to a "Local / Maps" effort, down the road.
  - Compared to products, I think jobs and cars will be easier to implement. Useful attributes are limited for jobs and cars, so we can just come up with one feed spec for each. If we can implement products, we should be able to implement jobs and cars. Services may be a little bit more vague since prices tend to vary based on the scope for each service.
- We have a results page. What other page types will we need? Product reference pages? Category pages? "Comparison matrix" pages? Others?
    - **Base Response Summary: Start with a good results page. Review pages for the long term.**
    - We may need more pages in the future, such as review pages- and organizing this in a better way than our competitors obviously.
    - We should provide any necessary information in one place for our users, but also, we don't want to have our users going to multiple places to find information. It should be centralized and organized, and if that involves adding in more pages, then we should make sure its done effectively.
    - In Google fashion, its good to do the basic and necessities first in terms of things needed and then examine closely what the market wants and what its telling us. It will also be helpful to look at market trends from similar product searches to see what they are offering their users and what their users are saying about their product features.
    - Related searches- if someone is searching for a wii, the main results they get should be wii's, but there should be the option for the user to also look at wii accessories (games, remotes, etc.) The user shouldn't have to search AGAIN for wii-related things. So here, the way we organize related searches for the product is important.
    - User-submitted reviews / ratings would also be nice, and probably a highly-requested feature.
    - New pages might not be necessary beyond the search results page. Perhaps plus-boxes in results for ratings and product details? A "more info" link that fills the right side of the page with information about that product?
  - Do we need to maintain the ability to display non-Product structured data on google.com?
    - **Base Response Summary: Yes, if we can get this right for other verticles as well, it would be a huge win for the company.**
    - Yes, it's probably a good idea to maintain this ability. You never know when you might need it and how the information could benefit things down the line. If you can get people to send you structured data, that is a good thing.
    - Yes, especially since we spent so much effort acquiring well-structured content (especially housing, vehicles)
    - Definitely. Even though it's a small % of what we receive, any time we can improve the quality of data on our flagship product, that's a good thing.
    - Yes, yes, yes! The vision of █████ is amazing. If the execution can be pulled off right (housing gets close to it), it will be a very powerful feature that differentiates us from our competitors.

- For what product verticals should we even bother attempting to normalize product offers? (No sense in normalizing flower bouquets...)
  - **Base Response Summary: We should only normalize products with unique identifiers (SKUs, ISBNs, Barcodes) etc.**
  - The obvious ones: electronics, books / magazines, media, home and garden, toys, sporting goods. Some that don't make as much sense: jewelry, antiques / furniture.
  - Should normalization be tied to vertical or rather just to the products themselves? We should get as much normalization information as we can get and use it wherever we can fit it. If there ARE a few standard flower bouquets, then by all means normalize them (if I just want to buy a dozen red roses, why shouldn't I be able to see that in a normalized view?).
  - We should only normalize products with unique identifiers (SKUs, ISBNs, Barcodes) etc.
- Ultimately, what features should distinguish Google Product Search from its competitors?
  - **Base Response Summary: If our goal is to be as good as Amazon, look at Amazon and set that as the benchmark in search result, speed, quality, personalization, tax/shipping, etc.**
  - google.com intergration
  - speed - pages load faster
  - quantity - more products, more sellers, more data, etc. "a marketplace that sells everything to everyone everywhere".
  - quality - better data, less spam
  - UI and organization- we have more data, but when we display it to our users it needs to appear simple and streamlined. Users shouldn't be overwhelmed when they conduct product searches. It needs to be easy for them.
  - fraud- little or \*no fraud. Related to quality.
  - Personalized search- this can play a role for users looking for certain products. We should customize it for our users with product search as well.
  - Internationalization- we should give our users the ability (eventually) to search for products in certain languages and have results returned to them in certain languages. Perhaps having a translator service available so users could seamlessly translate pages in english to their language and vice versa if needed. Ultimately, we should make internationalization a way to condense currency/language barriers to the market so its really one market for everyone. This would potentially involve establishing relationships between world banks, etc, and would mess in with policies, but if we are able to do this for our users, we'd be the first.
  - Market research: Study our customers and the industry. We are lucky in the fact that we won't be the first to do product search. Amazon came before us. We should look at the strengths and weaknesses of already existing product searches and use it to make ours stronger and better. We should use our predecessors as an example and take advantage of the fact that we can learn from what they are doing and we can learn from their users.
  - Product Search is very, very well done on our competitors' sites. We should put a huge concentration on getting the most information we can, since that is one of our biggest advantages. A fast, Googley interface and instant purchase with Checkout will also help. First, we should concentrate on getting up to the same level as our competitors, though - tax/shipping, ratings, etc.
- What are the biggest issues in internationalizing the product?

- **Base Response Summary: Lack of people support in the new markets - this include CSRs, and people who can shape the product for the market (research local laws, create appropriate feed specifications, etc.)**
  - inertia: no one ever internationalized Froogle beyond DE and UK. deadlines came and went each year like the seasons. The question is if we can pull our resources together in a more timely fashion and stick to our deadlines.
  - support: quality (how easy to use, how reliable, etc.) of product will determine amount of support needed. will have to coordinate support activity over multiple offices.
    - We need to make sure we streamline processes and trainings across offices for the product. This needs to be done ahead of time and not progressively as the launch in another country occurs.
  - multi-language helpcenter
  - international laws. how do foreign countries classify Checkout? avoid any applicable PayPal pitfalls.
  - US outcome: How things go in the US is very important. If the product is unstable (bugs all the time) and not defined in its vision and direction, this will be a problem when we want to launch it in other countries.
  - Internationalized launches should not necessarily be Checkout-dependant. When are we going to have Checkout in China? Who knows, but I'll bet lots of Chinese want to buy stuff, and the faster we can get an efficient and quality product search engine into that market, the better. Checkout integration, which I'm sure has many more legal hoops to jump through, can come later.
- What Checkout features are needed to boost Product Search?
    - consistent cart appearance &| checkout UI from merchant pages
    - add-to-cart &| checkout from google.com pages
    - tax and shipping estimates
    - inventory
    - trusted seller stuff
    - ability to change shipping options
    - Ability to purchase multiple items with checkout
    - Internationalization
    - User flexibility: People are fickle. Shopping online is not as easy sometimes because your sale is generally final, returning things is a pain b/c it involves shipping, you may have clicked on or selected the wrong size, etc. The ability to go back and edit an order after you have processed it. A lag time (maybe 15 minutes or something else) where you have a grace period to cancel an order and not worry that your credit card will still be charged, or change your shipping address, or change the color of the shirt your ordered. And a grace period for the seller to also make changes, if needed, to the order. Google can help facilitate this for users and put them in more direct contact with merchants. For example, buy.com and amazon allow you to cancel orders up until the time they are shipped. This might not be as feasible with Checkout, but it should be our goal.
    - Communication: Facilitate communication between the buyer and seller. Enable chat perhaps to discuss orders if necessary- this would involve the opt in on the buyer and seller sides.
    - Contact: Providing all our users with a valid contact from whom they are buying. How many times do you have to search for a contact address or phone number to call. If Google can put this up front for users, this would save some time for customers and help us gain more trust from users.

- Merchant / Product ratings and reviews.

### **Organizational**


- What is the relationship between the Product Search organization (in eng, ops, sales, marketing, ...) and the Base organization?
- Staffing! See the [Staffing trix sheet](#)


### **Other desiderata**

- % English google.com queries that are commerce-related [appears to be 5-8%, based on looking at random Google.com queries]
- distribution of commerce-related English query types on google.com [looks like roughly equal category and navigational queries, plus some brand queries and store queries]
- % google.com queries that are commerce-related in {de,fr,es,nl,it} [need help from language specialists]
- top 100 froogle standalone queries [mostly consumer electronics]
- daily # of clicks sent to merchant sites from checkout-enabled items on froogle [may not be logging this currently, jeff to estimate from froogle coverage]
- daily # of clicks sent to merchant sites from all items on froogle [about 1M daily]
- proportion of google.com users with a (US?) location in their cookie (so we can calculate shipping) [looks like 1.1%]
- proportion of google.com users with a Checkout account
- histogram of # pages seen for users with merchant clickthroughs
- histogram of # pages seen for users without merchant clickthroughs
- clickthroughs on froogle onebox topline vs. smaller links below
- usage of refinement links [see below]

## **5 Appendices (Data, old stuff, and stuff that needs to be integrated/updated)**

### **Data for usage of Froogle features**





## **Proposed milestones [needs update - Jeff]**

- Initial team staffed
- Clear product definition and short- and medium-term mockups created and agreed upon by key parties
- (separately) Launch "Checkout products" onebox for existing Checkout users
- First set of changes to existing Froogle site: clean up UI, Checkoutify
- Launch de/rebranding changes and do press cycle when FFE and utilities are repaired to previous quality
- Demonstrate inclusion of Checkout quality signals into ranking and normalization for product-navigational queries and category queries in the top 3 product verticals, e.g. consumer electronics, books, music.
- Demonstrate good quality ranking and normalization for product-navigational queries\* and category queries in the top 3 product verticals, for USA, for online merchants (and large chain merchants with local outlets?)
- Create simple product reference pages for all products; good quality for top 3 product verticals (all web resources pertaining to the product; reviews if possible (but capacity to collect reviews depends critically on normalization quality)
- 1% google.com traffic to Product Search index
- [Universal search integration] Demonstrate good quality exposure for product-navigational and category queries in the top 3 product verticals (100% google.com traffic to Product Search index)
- Demonstrate inclusion of Checkout quality signals into ranking and normalization for next N product verticals
- Demonstrate good quality ranking, normalization and product reference pages for next N product verticals...
- Demonstrate good quality exposure for next N product verticals...
- [Repeat for all product verticals we identify]

- [Repeat for all international markets we understand]
  - Additional vertical features...
- \* Note: Choice of whether to focus initially on product-navigational queries or category queries to be made on basis of google.com and froogle.com query analysis and technical feasibility. Category queries may actually be easier in product verticals with catalog information.

## **Louis/Louis staffing estimates [to be incorporated? - Jeff]**

- Base team is about 25 engineers. It takes probably about that number for the basic mechanics, acquisitions, searching, multiple frontends, and production work, and it takes up to two times as many for the machine learning/classification/cluster work.
- Current team is in MV. Could identify parts that can be done remotely.
- **Data Acquisition (3)** (Feed system): must scale, be maintainable. Backwards compatible data format, merchant push, much later maybe get data from crawl. Must support at least sub 1s inventory updates. Currently long chain of processes, need shorten and add directly to the repository, rather than many steps.
- **Feed frontend (3)** Merchant must be able to manage their feeds, manage attributes, etc. PFE does this, but is heavy, based on ads code, etc.
- **Repository (3)** Bigtable used for user interaction, a replica that is used for annotation and cleansing and to build Mustang repositories. Most work here is probably production work, keeping it up, etc.
- **Product Search System (6)** attribute based search, browsing, multiple catalogs. Using mustang is mostly sufficient, just not very fast. Need to massively improve quality to just match everybody else's. 20 people 1 year (Joel)
  - histograms + accurate counting, fast. Mustang can't do this
  - ranking/selection rules
  - behavioral ranking/selection
  - grouping of results
  - diversity
- **Shopping frontend (6)** On top of product search system, needs to support browsing of categories, parametric search, product and merchant reviews, recommendations Generally support of user interaction. Current Froogle code does a lot of postprocessing of Mustang results. Base has a lot of strangeness, too, and is very general, but looks more modern. Need to rewrite the whole thing. Frontend should not do very much, and may be different for different categories.
- **Annotations, Analysis & Cleansing (50?)** (Udi thinks this requires 50 people, with 5 we fail, not worth doing if only with 5. Udi also thinks this is the main part; everything else is limited by data quality.) These initiatives can be done in parallel, and in different locations. This teams controls mostly how fast we can roll out from selected large merchants only to smaller merchants, and lower quality data.
  - category classification. Which node in taxonomy.
  - catalog mapping. Which canonical item something belongs to.
  - attribute extraction
  - spam
  - catalog/attribute fodder
- **Tools (3)** Need internal toolsets for all of these for manual editing, fodder approval, training set editing, quality eval, etc. Some of them external, too



- **Universal Search support (1)** system that figures out how commercial a query is, whether we should put a link to product search, or even list a few products. Can e.g. do this manually for x queries that cover y% of search traffic temporarily.
- **Reviews (3)** Need product and merchant review system. Collecting reviews, processing them, giving merchants/vendors opportunity to change them (a la checkout), displaying reviews, etc.
- **Recommendation engine:** People who bought this also bought this, your friends bought this, etc.

## Competitive data

From Comscore:  
Worldwide monthly pageviews

	November 2006	Jan 2007
eBay sites	11B	9.5B
Amazon site	1.9B	1.6B
Froogle	30M	30M

That is, all eBay sites together have 317M average daily pageviews and the Amazon sites have 53M average daily pageviews. Google has 2.3B daily searches, so this is 13.8% and 2.3% of Google searches respectively.

## Traffic baseline

### Top conclusions from analysis of current Froogle traffic

- **Traffic from onebox (82%) dominates standalone traffic (18%)**
- **0.34 merchant clicks per Froogle resultspage (surprisingly high - probably 'info' clicks more than purchase clicks)**
- **Median user session length once on a results page is 1 (mean 1.73)**
- **Some commerce-related queries much more popular than the top onebox-generating query do not generate oneboxes**
- **Spot checks suggest onebox triggering is not very reliable (about ~60-75% of the time for a given English query)**
- **We show 86% of Froogle oneboxes to English traffic but from spot checks, only about ~60% of top onebox generating queries came in English, so possible int'l exposure opportunity**

Considering Thursday, Feb 22 2007:

4,589,098 total PVs

3,926,624 English PVs (86%)

4,185,310 total resultspages

112,644 resultspages with checkout restrict applied (2.9%)

187,453 PVs to /

151,569 PVs to frghp (from web tabs with an empty textbox)

**1,459,061 PVs to froogle\_url (merchant page clickthrough) [uncertain]**

45,089 PVs to froogle\_cluster (product clickthrough in grid mode)

11,491,219 onebox impressions

1,877,903 onebox clickthroughs (16%)



9,780,358 onebox impressions with English displaylang (85%)  
1,710,835 onebox impressions with German displaylang (15%)


Top query generating a onebox impression was [nintendo ds] which was searched 13,247 times, 6345 with English displaylang (47%), 1688 in French (13%), 1435 in Spanish (11%), 1397 in German (11%), 562 in Dutch (4%). Similar distributions for the rest of the top 10 queries generating a onebox impression (the queries came in English, German, Spanish, French, and Dutch, in that order).

Of 6345 onebox impressions for English traffic we generated a onebox only 3658 times (58%).

Alas, [nintendo wii], which was searched 25,154 times, and [wii], searched 57,185 times, did not trigger oneboxes at all.

Top 10 queries containing wii vs. top 10 queries containing wii and generating a onebox:





Of the top 120 queries generating onebox impressions on 2/22, 41 (33%) were obviously generated by scripts (or offsite links being followed) than organic search traffic. **The top 120 English onebox-generating queries only accounted for 70,710 onebox impressions (0.7% of the total 9,780,358 onebox impressions with English displaylang[!]), i.e. current Froogle onebox exposure has an absurdly long tail.**

If we use a hit to "/", frghp, or a onebox clickthrough as a proxy for a user session and assume standalone users and onebox users have similar session lengths (allowing for the initial hit to "/"): [REDACTED] mean session length is 1.73 pages starting from a results page. (Note there will be some sessions much longer than 3 but no sessions shorter than 1, so given the mean of 1.73 the median must be 1. Makes sense that this distribution is highly skewed.)

## Froogle issues


Product normalization - hurts popular queries like [wii]  
Ranking - hurts popular queries like [wii]

Refinements (wireless routers don't have MB, and Volts is hard to argue for when the values are things like 3-8, 70-175, etc.)  
Lots of Froogle merchants using a fake picture that says "No image"  
No notion of how much shipping is  
No notion of when something might arrive  
No notion of what checkout logo means  
Local results have very low coverage, aren't normalized with online results  
Sparse reviews  
No help figuring out what to buy

[wii] is a disastrous query due to product normalization / accessory issues-- as of Feb 24 2007, the first actual Wii console is on page 4 and the next is on page 7, then the next on page 9. That is, sorting by relevance (default) gives you 3 legitimate results in the first 90, a 97% irrelevance rate. Sort-by-price-descending turns out to be the only way to find the actual wiis. See 214 local results for wii gives a bunch of useless crap, but now in Richmond. Ads are relevant: one ebay landing page with tons of spam and other user generated crap for \$0.01 one from lowestdeal.com which has the wii in stock. On Amazon, [wii] gives great results: the actual wii is first, the wii remote controller is second, the wii nunchuck controller is third, and zelda: twilight princess is fourth.

Base Team Comment: Maybe we can penalize peripherals and accessories by giving lower rankings. For example, if they specify the product\_type of "game accessories" etc. we can give a multiplier less than 1 and penalize their rankings.

## Data



## Resources

[Current Base feed specification](#)  
[Whale/Froogle Consumer Front End PRD 2005](#)

### **Base Team Additional Comments-Ranking Merchants Braindump:**

overall rank = (pagerank) + (% items with rich attributes/information) + (ratio of "good/bad" Checkout transactions) + (add-to-cart conversion rate) + (unique content i.e. titles/descriptions/landing page URLs) + (upload/refresh frequency) + (merchant rating scores) - (# of bad offer complaints) - (eBay items)\*

\*subtract from overall rank if items are mostly eBay because a common complaint is eBay items dominate search results. A lot of merchants/users searching on Base want to AVOID eBay.

total number of clicks, history on Base/Froogle (doesn't have to be linear - step function)

Learn from other products like news - marking sites as golden and super golden  
We can maintain a list of trusted aggregators

Demotion of provider sites that have bad content

Number of attributes

% of errors in a file (how well formatted their file is)

diversity of payment options - wider variety = trusted providers

uniqueness of items could be factored in - eg: braided cable knit poncho

Digg it - you can rate it after you buy it, integrate it with GAIA