



Dashboards as Easy to Use as Google

Sponsored by QlikView

In our personal lives, when we want to find information, our first stop is the simple and friendly home page of the Google search engine. We enter search terms in the box and then a list of documents appears. We look at the documents. If they are not what we want, we change the search terms. If another question arises, we search for related information. The feeling we get from this experience is satisfying, smooth, and effortless.

Technologist George Gilder once explained Google's success in the context of what he calls Gilder's Law: it wastes what is abundant (as determined by rapidly falling prices) and conserves what is scarce—time Google places a premium on people's time, aiming to return search results in as little as a twentieth of a second. Such performance not only demands peta-scale server farms, but also an utterly simple interface—starting with its famously uncluttered homepage—and the interactive elegance seen in applications like Google Maps. Users are readily able to answer any questions they ask.

In the business world, we are often frustrated because finding needed information, searching for answers, and performing analysis is so different from Google's simplicity. Instead of being presented with an intuitive, straightforward user experience, we are offered a confusing, cumbersome menu of precomputed answers to choose from. If we want to explore more on our own, we are left struggling with even more complicated tools or antiquated spreadsheets. The IT staff wants to do better, but its hands are tied by traditional business intelligence systems' hardwired limitations that require it to deliver predetermined, inflexible query output to users.

The business equivalent of the Google search interface is the dashboard, which broadly defined, is a collection of information and analysis tools that are focused on providing detailed business answers specific to an end user's role and needs. Compare Google's intuitive, accessible approach to the numbingly detailed reports and mind-bending queries of business intelligence (BI). To make the world of BI provide the same ease and



freedom in asking and answering questions that Google provides consumers, we must find a way to “Google-ize” dashboards. This means creating simple, intuitive yet powerful tools that are both customizable and scalable—a tall order for most BI systems. It will require abandoning the SQL query mindset and adopting new forms of interactivity, such as using in-memory associative search, which allows users to pull answers toward them rather than accept whatever IT pushes at them. The result is what’s known as *living dashboards*—applications that can help identify leading indicators and allow users to share them throughout the organization.

How can we adapt qualities of Google to build better dashboards? This paper explains the similarities and differences between Google searches and traditional BI and why Google-izing BI is the first step toward creating living dashboards.

The Difference between Searching the Web and Traditional Dashboards

Comparing Google with traditional dashboards requires more analysis than it might seem at first. Google does a great many things well, starting with searches, but the one thing it does not do well is actually answer questions. Unlike BI, which sifts through data and runs calculations, Google’s keyword searches return the entire Internet as their answer, albeit sorted by relevancy determined in large part by analyzing which pages are the targets of links. The reason this works in practice is because the Web is text and document based. But corporate data is typically numerical and highly dimensional—answers are hidden in the numbers and must be computed, not searched. Corporate data has no network of links to help tools figure out what’s important.

A simple example illustrates the dimensionality of business data. If a dashboard returned a list of unpaid invoices due a company, that wouldn’t answer the question of the total accounts receivable in the next 30, 60, and 90 days. The total of the receivables is one dimension, the date due is another, and both are linked in the invoice information.



Traditional BI is less a search engine than a computation engine. Most of BI infrastructure is aimed at computing answers to narrowly predefined questions from a set of laboriously assembled dimensional data. The answers are not pulled by users based on their emerging interests but pushed according to a predetermined set of needs (see Figure 1 for a comparison of Google and traditional BI). The architecture of traditional BI is left over from an era in which computing power was the scarce resource to be conserved. But everything has changed. Repeated cycles of Moore's Law¹ have made memory cheap and hardware exponentially more powerful. The number of 64-bit cores on an Intel processor is doubling every eighteen months while the cost per gigabyte of solid-state memory is plummeting. Now it's possible to load entire databases into memory on computers of any size for instant "in-memory" analysis, something that was once prohibitively expensive or even impossible. Traditional BI applications don't reflect this. Their architecture was designed to cope with resource bottlenecks; today, BI itself is the bottleneck.

¹ Formulated by Intel co-founder Gordon Moore, Moore's Law states that computing power, as measured by the number of transistors that can be efficiently placed on an integrated circuit, doubles every two years.

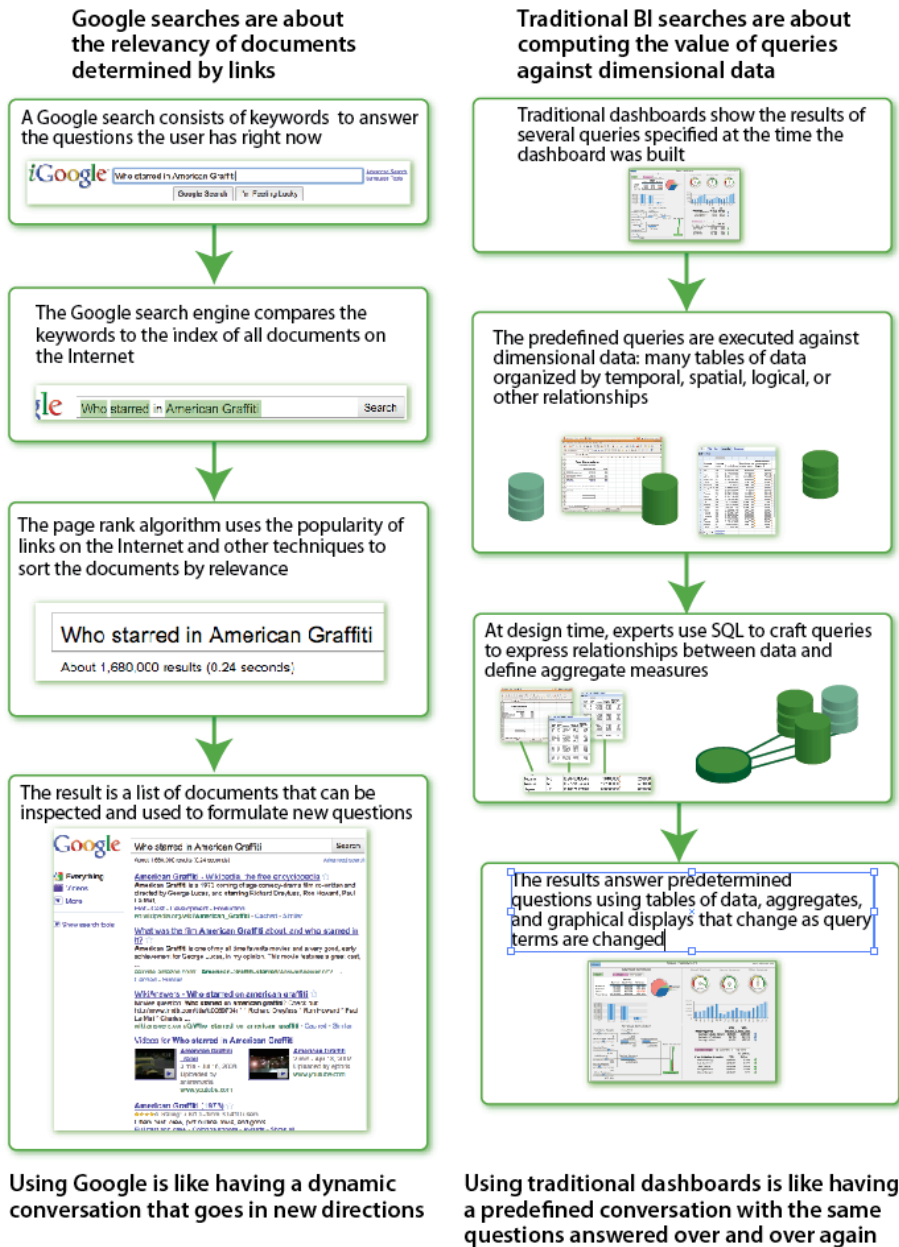


Figure 1: Google Searches Compared with Traditional Dashboards



What's needed is a fusion of the two approaches that offers BI's computational abilities through a dashboard as simple and as scalable as Google. A Google-ized dashboard would be able to explore sets of associated dimensional data in an interactive, responsive way that allows the conversation to go where the business user's ideas take it.

For example, in the unpaid invoice scenario described earlier, a Google-ized dashboard would allow the invoices to be sorted by criteria other than those planned. New aggregate measures and graphics could be added and incorporated into the dashboard. New data could be added by the business user, such as a database showing credit ratings, the last 30 days stock price, or delays in SEC filings.

To understand what a Google-ized dashboard is we must be specific. Here is a survey of which aspects of Google's interface must be kept and what new capabilities are necessary to make this vision a reality.

What Google-like Qualities Must Be Kept?

To create a Google-ized dashboard, the following qualities of Google must be kept:

- **User control of the analysis.** The back and forth flow of questions and answers should be determined by the user. Answers should be pulled out of data by the user, not pushed at the user by a predefined system.
- **No IT support required.** You just sit down and use Google without setup or training. A Google-ized dashboard should work the same way.
- **Simplicity and ease of use.** What could be simpler than a search box? Or more ubiquitous? Google's interface is querying stripped to its essence. Everything else is clutter and complications.
- **Speed of response.** Google returns results in as little as a twentieth of a second because users do not want to wait. Speed is everything. Nothing is more precious than time.



- **The ability to quickly see and inspect results.** Google may do a poor job of returning answers, but it does a fantastic job of ranking, sorting, and presenting results. It gives you everything you need to decide the relevance of a result and nothing more.
- **The ability to quickly refine questions based on viewing results.** Didn't find what you were looking for? Type something else into the search box and try again. There's no need to construct an elaborate query.
- **No penalty for volume.** There's no such thing as too broad a Google search. Type the letter "A" into the search field, and instead of returning an error, Google instantly lists the first of 15,490,000,000 results (starting with the entry for "A" in Wikipedia). This is a far cry from BI, where too broad a query can bring the system to its knees. A Google-ized dashboard doesn't penalize you for thinking big.

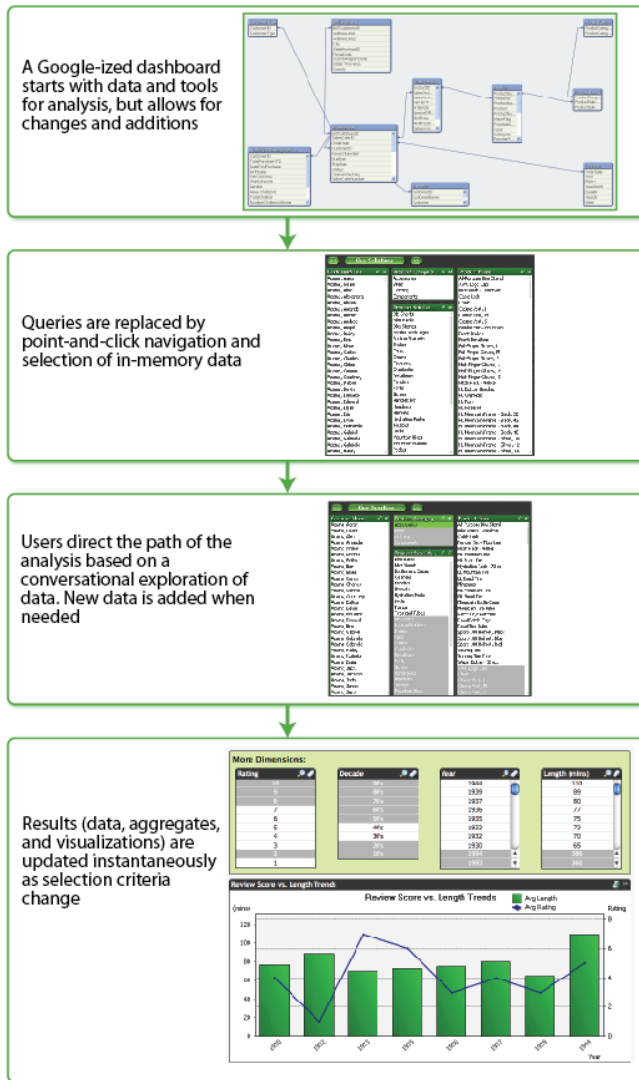
What's Different?

The obvious, for starters: searching for a string of text is a wholly different proposition than computing metrics from relational databases. To understand database records through visual inspection in the same way that web pages are understood is next to impossible. You need to aggregate data and graph it to see the story told by tens of thousands of records. In addition, instead of just looking at one document at a time, dimensional data involves connecting many tables. An invoice for a company may be connected to the database showing payments using a company identifier as a key. A Google-ized dashboard would make connecting tables together much easier, automatically providing a list of associations between data sets loaded into its own memory. Users can define the dimensions of data, that is, collections of related tables, based on the needs of their analysis.

A Google-ized dashboard of the type described in Figure 2 would solve BI's enduring dilemma: BI can tell you *known unknowns*—common, dependable metrics such as sales



and customers by region—but it can't tell you *unknown unknowns*, the questions you haven't thought to ask. The former is the province of traditional dashboards and KPIs; the latter is where insights are found. A Google-ized dashboard would be able to change search criteria on a user's whim and track down answers pronto. If the starting point for a sales dashboard showed products sold by region, a Google-ized dashboard would quickly allow examination of aggregated data for groups of regions or inspection of the individual invoices for each product by region. One reason Google's interface is so simple is because its results are simple. But a Google-ized dashboard would offer users more options, such as the ability to easily aggregate sums, averages, statistical functions, and equations with a few clicks so that custom metrics can be created, saved and reapplied. Being able to choose from many different ways of displaying data enables the dashboard to be simple, to show only what is needed for a particular situation. In other words, a dashboard would have dozens of different charts, tables, graphs, and other dynamic visualizations available.



The ease of use of Google is preserved in an environment for understanding dimensional data

Figure 2: A Google-ized Dashboard



The Results

As Google and so many others have found, the real value lies in empowering users to control the means of analysis. No IT department is as scalable as all of the colleagues it serves. What they need isn't Google, *per se*, but what we are calling a *living dashboard*, a dashboard that preserves the direct simplicity of Google's interface, but has the mechanisms needed to handle dimensional data.

Unlike traditional BI and its rigid, predefined analytics a living dashboard applies the best practices of Google to exploring dimensional data. Living dashboards enable business users to:

- **Ask and answer questions themselves**, without complicated SQL queries or IT middlemen
- **Add new data sources** at will, in any combination required for the task at hand
- **Get answers quickly**, almost instantly. No more waiting for reports to run
- **Change the dashboard themselves** on the fly without the help of IT
- **Share the dashboard with others**. Because what's the use of creating new metrics if you can't share them?

Living dashboards find results in time to matter, and this newfound visibility increases accountability. The ability to see the outcomes of decisions increases the understanding of what works and what does not. Ineffective behavior and successful practices can be more easily identified. As processes are better understood, leading indicators of problems or opportunities come to light. With the right metrics in place, a manager can hold his team accountable for what matters in an open and transparent manner, improving everyone's performance. If good metrics change awareness, then awareness invariably changes behavior. And this awareness can emerge anywhere in the company—an individual, a team, a division—and quickly propagate via their dashboards.

Over time, a virtuous circle is created. Because these dashboards are prototyped, improved, and shared by individual users, they no longer require the attention and



resources of IT. As the speed and ease of configuration increase, the cost of implementation falls and ROI rises accordingly. The speed of adoption becomes paramount. As Google has found, the faster it scales, the greater and more rapid the return on investment. The cheaper and easier it becomes, the more willing people are to experiment with it—to discover more leading indicators and more previously hidden insights.

The virtuous circle picks up speed. Resources once pushed to the business users by a central committee are pulled by the users themselves, who have the best understanding of how to deploy them and a real appreciation of their benefits. In the end, the entire notion of business intelligence is flipped upside down: it's fundamentally about discovery, not delivery; it's bottom-up, not trickle-down; and it's putting the tools for innovation in everyone's hands, not just the experts. It's simplicity itself.

QlikView: Helping You Google-ize Your Dashboards

QlikView is a new kind of business intelligence software that lets you stop guessing and start knowing how to make faster, smarter decisions. Figure 3 shows how QlikView creates a living, Google-ized dashboard:



List boxes allow data to be selected with a few clicks

What is selected and not selected is shown instantly

Graphics are updated instantly as selected data changes

Users can change the list boxes and graphics

QlikView returns results immediately because all data is stored in memory

Each table is connected to the other tables related to it when the data is loaded

Users can add tables on their own without the help of IT

Because all data is in memory and related, it can be explored in a hands-on way, by inspection and experimentation, not by queries

The screenshot shows a QlikView dashboard with a 'More Dimensions' panel on the left containing list boxes for 'Director', 'Genre', and 'Actor'. The main area displays a table of dimensions (Cing, Decade, Year, Length (mins)) and a bar chart titled 'Review Score vs. Length Trends' showing a score of approximately 100 for the year 2003.

Figure 3: A QlikView Google-ized Dashboard

Sophie Jasson-Holt 4/8/10 14:51

Comment: Ika-will you be able to change this graphic and add SQL to “not by SQL queries”



QlikView's goal is simple: to provide the tools for living dashboards, helping customers explore and visualize data that enables them to personally find answers and drive innovation. QlikView takes an alternative view of BI that can be summarized in three steps:

- **Consolidate:** Identify related data sets, map the associations between them, and load it all into memory
- **Search:** Explore the data using point-and-click controlled “list boxes,” which display selected and deselected data as well as aggregates. Information is displayed and updated instantly as selection criteria change
- **Visualize:** Maps, charts, and assorted graphics can be created and instantly updated

QlikView aims to increase your chances of making genuine discoveries and eliminates much of the grind:

- **No SQL queries, no middleman:** You click to select data and click again to deselect it
- **Matching and non-matching data are displayed:** You can see what is highlighted, and what is not highlighted. You can see which data was excluded and see what happens when you mix it in
- **No waiting:** The answers are right in front of you

The idea behind QlikView is that a simple process for asking questions and encouraging individual exploration leads to better answers, insights, and innovations. It does this by replacing database queries and cubes with its associative in-memory architecture. Rather than precalculating answers, the software loads data sets into memory and maps the associations between them. The user's job is to understand the data, not to grapple with the technology. You can ask as many questions as you like until you find the important insights. Because it's all done in-memory, the answers are returned instantly and updated continuously. Your BI is no longer as good as your IT department's last cube—it's as good as the questions you ask.

CITO Research

CITO Research is a source of news, analysis, research, and knowledge for CIOs, CTOs,



and other IT and business professionals. CITO Research engages in a dialogue with its audience to capture technology trends that are harvested, analyzed and communicated in a sophisticated way to help practitioners solve difficult business problems.