



CITO Research
Tell Us a Question.

AUGUST 2010

A Guide to BI for the Business User

Stop Guessing, Start Knowing

Sponsored by **QlikView**

Contents

Introduction	1
<hr/>	
Traditional Dashboards: The End of a Long BI Process	2
<hr/>	
Living Dashboards: The Beginning of an Exploration	3
<hr/>	
The Enablers of Living Dashboards	6
<hr/>	
The Effects of Living Dashboards	8
<hr/>	
QlikView: Where Living Dashboards Grow	9
<hr/>	



Introduction

The practice of business intelligence (BI) has suffered for many years from the pretense that the information needs of business users can be predicted in advance. Much of what BI delivers is predetermined. Data is extracted from many sources, consolidated into a data warehouse, and precomputed into answers. At the end, you get a dashboard or report that shows you what you wanted. An example could be sales by region. But what invariably happens is once the data is presented, it leads to more questions. If sales are down in the west, you need to know why. You need to see sales by rep and product, but typically, the dashboard or report can only show you what you originally asked for. If you want to ask additional questions or see new insights, you must go back to IT and ask them to add exactly what you want to the report or dashboard.

What is needed is a dashboard with the power to ask and answer *any* question on the fly. In other words, if you think it, you should be able to act on it. This would put business users in the driver's seat, with profound implications for any organization. Unlike traditional dashboards, these new dashboards are constantly being tuned and adjusted to meet up-to-the-second needs. Instead of information pushed to the business user in a predetermined way, the right kind of dashboard would allow users to pull information spontaneously based on a conversational process of analysis.

These *living dashboards* are the opposite of predetermined BI. They are powerful tools in their own right, able to access and manipulate information. Living dashboards can easily adapt to discovering and exploring new territory quickly at a low cost. Taken to their full potential, they can help find root causes, identify leading indicators, and then share them throughout the organization. This can increase accountability by exposing the impacts of decision-making and allowing employees to pull the development process along according to their own needs. When this happens, companies stop guessing and start knowing.

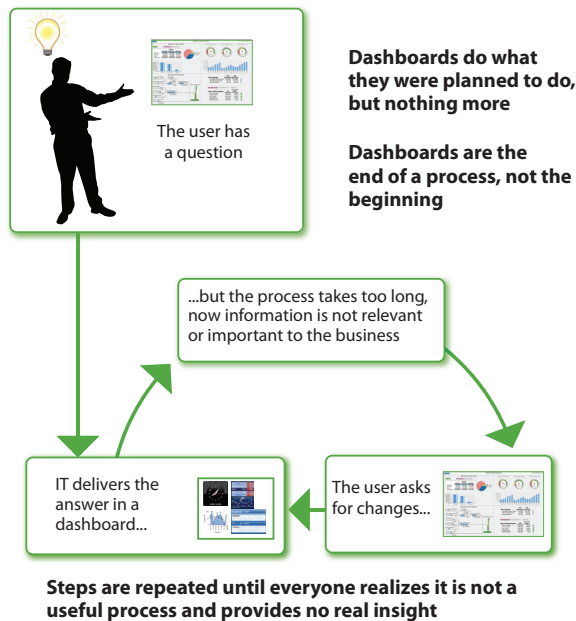
This paper explains how traditional dashboards cannot serve the needs of the business user. It also explores how living dashboards differ from their ancestors and how they empower business users to get better information so they can take their business success to the next level.





Traditional Dashboards: The End of a Long BI Process

Traditional dashboards resemble a snapshot behind glass—something to be glanced at occasionally and nothing more. A dashboard that doesn't change is a dashboard that doesn't get used. Such a dashboard can only tell you what it was designed to tell you and therefore quickly outlives its usefulness. It becomes obsolete even before it's ready, because conditions will have changed in the interim between the request and the finished product. The fatal flaw of these dashboards (besides being frozen in amber) is the intermediation by IT. Real-world metrics change too quickly for well-intentioned but resource-strapped IT departments to manage, leading to a disconnect between users' needs and their dashboards' capabilities. Efforts to then revise them inevitably run afoul of internal politics, prohibitive costs, and a failure to justify the ROI. When you are working with a spreadsheet, you don't have to ask somebody else every time you want to change something. But with most traditional dashboards that is exactly what you have to do.



Ultimately dashboards are the victims of a “push versus pull” dynamic between business intelligence and business users. In most organizations, BI pushes a set of reports (some as dashboards) to users such as sales by region. This information is no doubt useful, but always leads to more questions, such as why the North America region is behind on its forecast. Unfortunately business users lack the means to ask new questions or drill into the data as needed. Rather than being empowered by information, they are disempowered by uncertainty. Without the information they need, they don't know and end up all too often just guessing.

Figure 1. Traditional Dashboards Can Be a Dead End





Traditional Dashboards	Living Dashboards
IT-driven	User-driven
Creating the dashboard is end of the process	Creating the dashboard is the beginning of the process because users will be able to answer questions they haven't even thought of yet
Dashboards drive a centralized, top-down process	Dashboards drive a decentralized, bottom-up process
Dashboards are aimed at displaying a single predefined source of truth	Dashboards help people find what is true and important for their job
Dashboards show rolled up KPIs and scorecards	Dashboards show operational metrics and leading indicators
Dashboards evolve over weeks and months	Dashboards evolve over hours and days
Allow drill-down defined in advance	Allow drill-down based on path of inquiry

Living Dashboards: The Beginning of an Exploration

Business users will never stop guessing and start knowing until they can pull the flow of information in their direction, asking questions instead of waiting for already overburdened IT staff to do it for them. Once that happens, they can find their own answers instantly, not in weeks or months. For the first time, they have all the facts and a mandate to use them. The triumph of business users *pull* over BI *push* would overcome the limits of outdated tools and technology and amount to a BI revolution. The means for doing so is a living dashboard.

Unlike its predecessors, a living dashboard allows users to:

- **Ask and answer questions themselves**, no queries or IT middlemen required
- **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 configuration of the dashboard themselves** on the fly without the help of IT
- **Share their dashboard with others**. What's the use of creating new metrics and KPIs if you can't share them?





CITO Research
Tell Us a Question.

As its name implies, a living dashboard is a *living thing*—changing, evolving and maturing. This is important because how often does anyone get things right on the first try? The reason traditional dashboards proved so useless is because they only measured what could be measured, not what *should* be measured as dictated by a user's intuition. It's not the process that matters; it's what happens before and after it, or the inexplicable correlations between it and a seemingly random aspect of your business. It isn't until the second, third, fourth, or even fifth iteration of a dashboard that you're able to find those correlations, usually after much probing and massaging of data. For example, the creation of a living dashboard for purchasing at one company showed an unusual amount of buying from one supplier. It turned out there was fraud underway. What started as a purchasing dashboard became a fraud detection tool.

Living dashboards solve BI's main problem. Traditional BI can tell you only *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 KPIs; the latter is where insights are found. A living dashboard can change on a user's whim. When you start an analysis in a certain direction, a living dashboard allows you to follow leads like an investigative reporter.

It's also a **knowledge repository**. It captures discoveries as new metrics that can be propagated throughout the enterprise. Living dashboards turn the classical BI model on its head—instead of trickling down to users from the top, best practices begin to trickle up, or sideways, from department to department. It's a tool for packaging the lessons learned by your best employees and using them to bring everyone's performance up to the same level.

Living dashboards **change their users' behavior**. No longer limited to reading reports long after the fact, the users in charge of a process are able to see it for themselves in real-time, offering the chance to step in and make changes, test assumptions, eliminate waste, fix what's broken, and so forth. Increased visibility creates a feedback loop leading to continuous improvement—the ongoing iterations of the dashboard reflect iterative enhancements in the underlying process.

That's because living dashboards are built, refined and extended **by the people who use them**—not by IT. Users possess all of the tools necessary to fulfill a given task, monitor a process, or make a decision. Unlike traditional dashboards, which typically force users to switch applications, burrow through reports, and mine spreadsheets for a 360-degree view of the situation, a living dashboard has the best chance of incorporating all needed information in one place.





Living dashboards are also **personally customized by motivated users**. The value of a living dashboard is defined by users, who bring their own data and understanding to bear on problems. In many companies that use living dashboards, there is no IT support. Power users help beginners who quickly learn to fend for themselves.

In this way, living dashboards are **conversational**. You ask questions, receive answers, ponder them, ask *more* questions, and engage in a dialogue with your data. First you ask to see the top ten suppliers. Then you notice a strange company. You look at what is being bought. It makes no sense. Boom! You've detected fraud. Traditional dashboards replay the same conversations over and over and are a one-sided dialogue at best.

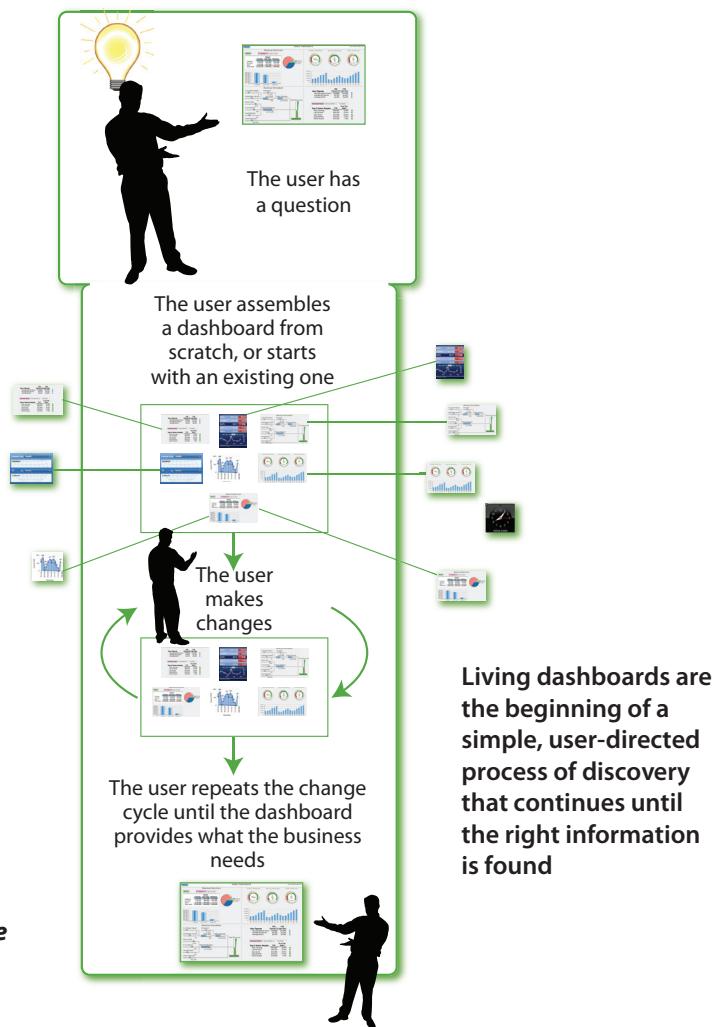


Figure 2.
Living Dashboards Are Always User-Directed





Is a Spreadsheet a Living Dashboard?

Superficially, a spreadsheet is a living dashboard. Spreadsheets possess many of the properties of living dashboards with respect to the user's ability to directly manipulate data and control the configuration. They've long been the measures of last resort for managers frustrated by canned BI reports and desperate for insights hidden in the data. But spreadsheets are weak at making and tracking associations between many different tables of data. Most of the time, they are a rigid two-dimensional table of columns and rows. Creating a connection between two data sets means combining them into a large table. Spreadsheets are hard to update, hard to use to consolidate large amounts of information, and rarely tell the whole story. When they do, it is only for a snapshot in time. They also limit users' ability to ask and answer their own questions. The spreadsheet interface has barely evolved since spreadsheets were introduced almost 30 years ago, and it continues to befuddle non-initiates. The hallmarks of a living dashboard include simplicity and intuitiveness, which are the root of its broad adoption. In that sense, no, a spreadsheet is not a living dashboard. Spreadsheets also become islands of information unto themselves, resulting in multiple versions of the truth. Adding more and more dimensions, that is, new tables of information, to a spreadsheet is difficult at best, compared to the ease with which new information is added to a living dashboard. Spreadsheets are not living dashboards, but they have played an enduring role as a bandage for traditional BI.

The Enablers of Living Dashboards

What does it take to build a living dashboard? What capabilities are needed and what technical work is necessary to make this dream (now 30 years old) a reality? The good news is that the technology already exists. Traditional dashboards were designed in a computing universe of scarcity—slow processors, limited storage capacity, and ruinously expensive memory, all focused on delivering a predetermined set of answers. 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

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.





Intel processor is doubling every 18 months, while the cost per gigabyte of solid-state memory is plummeting. Now it's possible to load entire databases into memory for instant (known as "in-memory") analysis, something that was once prohibitively expensive. Traditional BI applications don't reflect this. Their architecture was designed to cope with resource bottlenecks; today, BI itself is the bottleneck.

With the most daunting technical obstacles removed, what does a living dashboard look like? At the very least, it possesses the ability to:

- **Ask and answer questions directly without help from others.** This demands an utterly intuitive interface, as simple as pointing and clicking. Its conventions are more reminiscent of Google's or Apple's applications than spreadsheets or traditional dashboards, as both examples have proven remarkably successful at creating user-friendly apps with rich features and broad adoption
- **Use the fruits of the BI process: data stored in data warehouses.** New plumbing is necessary to extract data from the warehouses, load it onto the living dashboards' servers, and refresh it in real-time
- **Add data sets to the collection of data being examined and make associations.** One of BI's failings is that it's hard-wired to produce reports from the same data sets over and over. Why? Because adding new data sets requires painstaking SQL joins and elaborate queries and sub queries. (Entire books have been written on how to build better queries.) Adding a new table to a living dashboard has to be easy. The system has to provide help
- **Make associations automatically** and offer users the chance to define these associations for themselves and consolidate all of the data in memory. This associative in-memory approach is faster and easier than manually joining data sets scattered across numerous servers
- **Define aggregates.** Aggregate measures such as counts, sums, averages, medians, other statistical functions, and the results of equations can be combined on the fly in any order to create custom metrics that can be saved and reapplied
- **Create visualizations.** A living dashboard has dozens of different charts, tables, graphs, and other dynamic visualizations in its repertoire, rather than a single, static view. When you create a chart and then change the data that is selected, the chart should change in real time. You can see how aggregates change as you massage the data
- **Share the dashboards you've created.** They can be virtualized, cloned, and reassembled to share an entire dashboard, a clever chart visualization, or even a newly invented metric created after much experimentation





The Effects of Living Dashboards

If a company were to suddenly find itself equipped with living dashboards, it would gain the insight to create beneficial change immediately. Living dashboards jump-start an organic process of continuous improvement. Experimentation leads to a deeper understanding of business processes. You are able to see true leading indicators, those tested by the experience of those on the front lines of your business. False metrics fall into disuse. You can see the results on sites like KPILibrary.com, where more than 4,100 KPIs and benchmarks have been added by contributors. Creating KPIs is not just for analysts or senior executives anymore.

Over time, a virtuous circle results—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. The faster people use them and the more people use them, the greater and more rapid the return on investment. And the cheaper and easier they become, the more willing people are to experiment with living dashboards—to discover more leading indicators and more previously hidden insights.

The virtuous circle picks up speed. Resources once pushed to the users by a central committee are pulled by 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 about putting the tools for innovation in people's hands, not in experts' ones.

The best metrics can't predict the future and dashboards can't run your company for you, but the visibility they provide will offer you a chance to stop guessing and start knowing the impacts of your decisions. What you do with that knowledge is up to you.





QlikView: Where Living Dashboards Grow

QlikView is a new kind of business intelligence software that lets you stop guessing and start knowing how to make faster, smarter decisions. 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 "list boxes," which display the unique values in each field and can be highlighted and aggregated. Both information included and excluded from the selection criteria is displayed and updated instantly
- **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

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.

