WHITEPAPER

Everything You Know About Digital Measurement is Wrong...and What to Do About It

A Primer for C-Level and Marketing Execs on Digital Measurement

By Gary Angel, President and CTO, Semphonic

Overview

If you're a savvy marketing executive with your hands on digital, you may already know the key best practices in digital measurement:

- Focus on a small set of site-wide actionable KPI
- Measure your overall Site Satisfaction compared to your Industry
- Trend your NetPromoter scores to track your efficiency in creating brand advocates
- Use your Brands mentions in Social Media to track and measure yourself versus the competition

What you probably don't know is that every single one of these "best-practices" is wrong. Deeply, fundamentally, and completely wrong. What you probably don't know is that virtually every digital dashboard you look at is poorly constructed, deeply misleading and fundamentally flawed. Indeed, the common best practices in digital measurement are so ill conceived that it will take no more than the next 1500 words to convince you that you should flush not just the reports, but the people who make them, out of your organization.

www.infobright.com

www.semphonic.com

www.pentaho.com

The Challenge for Marketers

Let's start at the top. Your measurement department has almost certainly been preaching the conventional wisdom in digital measurement: don't overload on numbers. The key to successful dashboarding and reporting is finding a small set of site KPIs that are understandable and immediately actionable. And they've probably delivered exactly that - a small set of key metrics like Site Conversion Rate, Total Visits Trend, overall Site Satisfaction, etc. all laid out in big numbers with great fonts, pretty colors, big trend arrows and lots of Tufte-inspired whitespace.

Sadly, these reports deliver neither understanding nor actionability.

Suppose I walk into your office and tell you that your Site Conversion rate is up 5%. You'll probably be delighted. Now suppose I walk into your office and tell you that your Site Search Engine traffic is down 20%. That's bad, right? But would you realize that in all probability the two metrics are related and are, in fact, telling you exactly the same story? As you drive less early-stage traffic to your site via natural search, your Conversion Rate will go up. Understanding the system - the interrelationship between parts - is fundamentally different and more important than understanding the state of any single volumetric.

It turns out that a change in any single variable in a complex system can be explained in a variety of ways - some of which would be interpreted as positive and some as negative. It doesn't matter if the metric is traffic, conversion rate or even revenue.

Did I say revenue? Surely, it's impossible for a revenue increase to be bad?! Not only isn't it impossible, it's common. Company A creates a re-marketing program that sends a 10% discount offer to all cart abandons. Unfortunately, consumers quickly game the system and abandon the cart to get the coupon. Gross Revenue increases with the re-marketing, but Net Revenue declines due to reduced margins. Short-term revenue gains that sacrifice margin are frequent. Getting the wrong customers in the door (high-risk card applicants for example), paying too much for new customers, or simply reducing margins via discounting are all common marketing issues that are missed when you focus on a revenue KPI. Poor reporting systems make this type of system inter-dependency difficult or impossible to spot or understand.

What about Net Revenue? If Net Revenue is increasing, things have to be getting better! Not so. Many moons ago we worked with AOL in its ISP days. Now AOL was notoriously (infamously) difficult to cancel once you had signed up. The reason? Analysts at AOL had carefully measured the impact of difficult cancellation and could prove that making it hard to cancel improved Net Revenue. On the other hand, they didn't measure the resulting impact on brand and customer satisfaction. The brand literally eroded under a relentless program of short-term Net Revenue optimization.

There just is NO SINGLE METRIC that can be meaningfully interpreted when viewed in isolation.

When I walk into an executive's office and say something like "Traffic is up 5%", I expect him/her to ask "With whom?" Because if you don't know the audience behind that traffic increase, you don't know squat. Bad traffic is ubiquitous on the Internet and it's cheap. If that's what you measure, that's what you'll get. Once I've answered the "who" question, the next question I expect to be asked is "And why were they there?" Understanding what your customers are trying to accomplish on the Web is part and parcel of understanding whether you are successful or not. Your chances of making a product sale during a Customer Support visit? Zero. So why are those visits included in your Conversion Rate? If you're using a site-wide metric, all you're seeing is

noise.

We call this concept that every metric should be in the context of a “who” and a “what” Two-Tiered Segmentation - and if your dashboards aren’t built around this simple concept they probably aren’t very useful.

So here are the basic learnings:

- Individual KPIs aren’t actionable
- Site-wide metrics aren’t useful
- Dashboards that lack Audience and Visit Segmentation are useless

If there is so little utility in site-wide metrics reported from Web analytics systems like Omniture or Google Analytics, you might be tempted to steer your interest toward another staple of digital measurement - the online intercept survey. Survey research has a rich tradition with deep intellectual roots. So it stands to reason that the number it delivers are likely to be better than from Web analytics.

Except they aren’t. Online survey research is based on a sample of your Website visitors. Of course. But have you ever thought through the implications of this simple fact? Every time you launch a new marketing campaign, every time you improve (or worsen) your Search Engine Optimization, you effectively change the sample of visitors coming to your website and, therefore, the sample of visitors in your online survey.

There is no corollary to this in the traditional survey world where sampling and marketing were completely unrelated.
The upshot is that site-wide trends in satisfaction (and those are probably the only numbers you ever see on a Management Dashboard) are completely useless. Instead of tracking true changes to Site Satisfaction, you’re actually tracking changes to the site population caused by your marketing program. At the site-wide level, trends in online survey data are meaningless.

The situation is even worse if your measurement team has foisted off on you another staple of online survey satisfaction - the site satisfaction benchmark. In traditional (offline) primary research, these benchmarks are built using a single, large sample that drives a single survey instrument that includes questions about many brands. This works well. In the online world, however, the benchmark sample is built from the individual site samples of many customers. This means they are constantly changing and are completely at the mercy of a basket of individual marketing programs. And no, adding up a set of bad samples doesn’t produce a good one.

Here’s a simple example to show just how silly the use of online site satisfaction benchmarks really is.

Consider an example of four companies that are exactly alike in their business and are, therefore, 100% comparable. Each company has five Visit Types. Satisfaction for these Visit Types range from 59% to 71%. Each site is completely identical in structure and design except for the name of the company and each Visit Type has identical satisfaction scores on every site.

I hope you’ll agree that this benchmark set is an absolutely implausible best case in the real world. Any problem with this benchmark is certainly not in the chosen competitive set.

Suppose, however, that the distribution of Visit Types to these identical sites is as follows:

		% of Visits			
Visit Type	Visit Satisfaction	Site 1	Site 2	Site 3	Site 4
1	59%	20%	15%	25%	30%
2	63%	20%	18%	25%	25%
3	67%	20%	20%	20%	20%
4	69%	20%	22%	15%	15%
5	71%	20%	25%	15%	10%
	Site Total Satisfaction	66%	67%	65%	64%

Even with these perfectly identical sites, the average site satisfaction in our hypothetical example across all visitors ranges from 64% to 67%. In other words, even with four completely identical sites, with five completely identical sets of Visit Types, and the exact same satisfaction with every Visit Type on each site, a small difference in the distribution of Visit Types on the site can result in a significant variation in overall site success when measured across the entire visitor population.

Any decision-maker (this means you), looking at Site Total Satisfaction comparison, will believe that Site 4 is worse than Site 2 - even though the sites are completely idential in every respect including customer satisfaction by Visit Type and customer satisfaction across every single meaningful variable.

Surely this is a gross misrepresentation of reality. Of course, that gross misrepresentation of reality is on more than half the dashboards we see in Fortune 500 companies.

This effect is not limited to Visit Type. It is true for every single survey variable and it is true for any effort to “trend” the data. This is for identical sites, so imagine the implications for the real world where “comparable” sites are actually dramatically different in marketing drives, structure, function and audience.

If you’re using this type of benchmark on your Marketing Dashboards, you might as well - like ancient Greeks before battle - be shaking bones on the sand to predict the future. To evaluate your business based on such numbers is madness.

“Fortunately,” you may be thinking, “we use NetPromoter scores and these don’t suffer from any need of benchmarking.”

That’s true. Unfortunately, NetPromoter scores do suffer from all of the same problems with sampling. Trending your top-level NetPromoter score is no more valid than trending your top-level satisfaction. Devoid of careful controls on sub-populations, it, too, is completely devoid of meaning. Worse, unlike Site Satisfaction, NetPromoter tends to be a very poor intrinsic measure of Site experience. So if you’re using NetPromoter as a key top-line metric, you’re using a number based on a changing sample and that, even in a perfect world, probably has little relationship to site performance.

Lest you think that only your Web measurement teams are producing fancified garbage for your consumption, here’s a favorite example from the trendy new world of Social Media. If your marketing team is really on the ball, you probably have a shiny new

Social Media program along with the inevitable set of reports showing you how well it's doing. At the heart of those reports is something called Total Mention Counts - a simple measure of how often your brand is mentioned in the social realms. It might be compared to other brands (to give you "Brand Share") or trended over time (to give you Social Success).

Oh is this a bad measure. Produced by tools like Radian6 or Buzzmetrics, it's been adopted by nearly every Social Media agency despite being almost completely devoid of meaning. It is flawed in so many ways that it's hard to enumerate them all but here are a few basic problems. Social Media chatter is nearly always deeply unclean. A great deal of commentary is driven by professional influencers with some stake in a given industry. When you include all these mentions in a count, you are measuring some impossible combination of PR messaging and actual consumer sentiment. Unless your team has aggressively pruned the counts of all professional posters and commenters, the number you are tracking has no correlation with actual consumer sentiment. Of course, it doesn't measure PR or Social Media effectiveness either. In that Total Mention Count you've generally added up mentions in the New York Times with tweets to your Customer Support reps. What's that addition supposed to represent? Nothing we can think of.

Two-Tiered Segmentation: Expanding on the Solution

For many years now, Segmentation has been the single most powerful data aggregation technique in marketing. You can segment along almost any dimension. The key to effective segmentation is finding the dimensions that deliver lots of value in a compact form. Traditional visitor segmentations typically start with the business relationship; at the highest level, we tend to think of "customers" vs. "prospects" vs. "non-qualified." One level beneath that we often find that organizations add a value dimension. Customers are often segmented into groups based on the strength, number, value and duration of their relationships. We think of "members," "longstanding customers," etc. Beneath this level, segmentations have typically focused on additional facts about the customer: their demographics, psychographics and interests. This type of Visitor Segmentation doesn't go away when it comes to online. You still care (a lot) about the relationship between you and customer.

As important as the Visitor Segmentation is, however, it doesn't solve our problems when it comes to digital data aggregation. The vast majority of digital data doesn't have anything to do with this type of Segmentation. Two completely different types of visitors might have identical Web behavioral records. Nor is the Web behavioral record easily read to create a visitor-level Segmentation when one doesn't exist externally. What's more, digital data is often completely anonymous - leaving you without an intelligible visitor-level Segmentation scheme.

What is needed is another type of segmentation that helps us aggregate digital data in a meaningful fashion. This second tier should allow the collapsing of a stream of web behavioral events (server calls or page views) into a single or small number of fields that capture the essential meaning of the tracked behaviors.

Semphonic has developed a second type of segmentation designed to do exactly that. This second tier of segmentation is Visit-Type Segmentation.

The theory is simple. If you walk in the door of an executive and say "Web traffic" was up 5%, we'd expect any sensible executive to ask two fundamental questions. Who were those visitors (traditional segmentation)? What were they trying to accomplish (Visit-Type Segmentation)? You can't meaningfully understand Web behavior or Web metrics unless you can answer those two questions. That means that you can't read the behavioral record to measure their success except by first understanding their interest. So it seems clear that this

type of Visit-Intent based segmentation is fundamental to getting meaning out of Web data. Here's a typical KPI matrix that illustrates how meaningful success metrics are tied to a Two-Tiered Segmentation:

The creation of a Visit-Type Segmentation is no simple task. Typically, visits must be "signatured" using a hierarchical set of rules that describe the behavioral patterns characteristic of that visit. These rules are almost never as simple as "contains a page." Instead, they frequently need to use the entire visitor behavioral stack to figure out what a visitor did most of and what a visitor did first or early in a session. In addition, because the rules are hierarchical, they must be executed in series or must contain exclusionary rules to eliminate alternate, higher precedence, assignments.

This makes for a complicated and programmatic ETL task - one that is entirely custom to every enterprise. On the positive side, it yields a fantastic aggregation set. At the Visit Level, we can collapse most of the meaning of a stream of Web activities down into a few simple fields: one field to capture the Visit Type and one or more fields to capture the metrics relative to that Visit Type's success and efficiency. In most cases, it's best to think of a Visit Type as having a primary success count, a primary success value, a secondary success count and value, and a primary and secondary measure of efficiency. This gives six metrics plus a Visit-Type code to represent the key meanings in a Web visit. For an eCommerce site, this might translate into something like this:

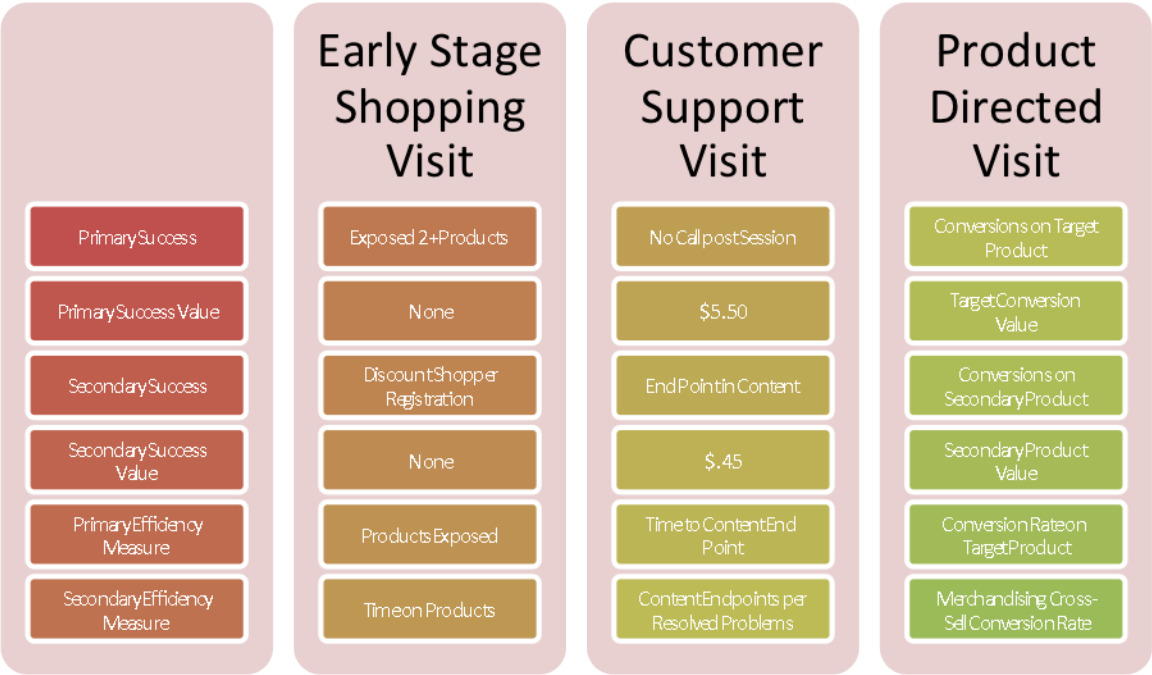

For most modern enterprise sites, we find that between 7-20 different Visit Types are necessary to fully and accurately represent site functionality. This doesn't mean you need 20 (Visit Types) x 6 (metrics) per visit. Because the rules are hierarchical, a visit needs only seven fields to be fully described.

Once you have Visit Segmentation in place, developing a Visitor Segmentation is much easier. The question of what to keep from the digital data at the customer level is now a whole lot easier to answer. What you want, at the Visitor-level, is a simple description of the entire customer's Web experience. What might have seemed almost impossible when first put in those terms now becomes much easier. To model the Visitor-level aggregations, we borrow a powerful concept from traditional Customer Analytics in the database marketing world. For many years, database marketers have modeled customers using RFM (Recency, Frequency and Monetary) metrics. It turns out that for many, many dimensions, if you can answer how often, how recently, and how successful someone was, you've answered many of the most important marketing questions about them. We modify the RFM formula because so many websites and Web Visit Types don't have a clear monetary return. So we tend to model the M as success count.

For each customer, the Visitor table is built to contain the RFM metrics for each Visit Type. With a minimum of three and a maximum of about seven fields to capture RFM, that means that you are looking at a fairly wide Customer Record - with somewhere between 21 and 120 columns to capture the website experience. That's pretty substantial, but once again, the yield is significant. With those fields, you've created a powerful and easy-to-use model of the entire customer experience on the Web.

Using this Two-Tiered Segmentation and RFM model, asking and answering questions like this from digital data become trivial: Which customers recently had an unsuccessful Customer Support Session?

- How many of those customers had a previous eCommerce-based Session?
- Which customers had a successful Early-Stage Shopping Experience but didn't return for a Potential Buying Session?
- Which types of Site Failures have the most negative impact on retention?
- Which Potential Buyers are the most successful - and are they also the most valuable?
- Which Customers are coming to the site more often? Which are coming less often?

The model not only makes it much easier for marketers to use the data to answer questions they have, it actually helps them frame meaningful questions. With digital data, this is vitally important. Digital marketing is still fairly new - and the methods to isolate and understand consumer segments just aren't well known. A powerful Visit-Type RFM model doesn't just expose the data in a performance-efficient manner, it helps marketers understand how the data fits together to form a meaningful picture of the customer's experience and the resulting marketing opportunity.

Driving Segmentation into Reporting

Is there a way to incorporate this kind of segmentation into reporting to solve the problems with KPIs and site-wide metrics? There is. By embedding a digital (Two-Tiered: Audience and Visit-Type) Segmentation into reporting and modeling systems instead of showing single un-related metrics, you can create truly meaningful dashboards.

Here's an example of a report designed to show how traffic to the website can be broken down into separate audience types. The flows provide a compact and powerful illustration of the critically important business information: what visitors the site attracts, how this is changing, and how successful the site is with each audience type.

In the next report, the health of each system (vs. forecast or historical trend) is shown in the up/down arrow in the Total Successes. Click there, and the report drills down one level into the Visit Segmentation for that audience type. The report explains how the overall success of the site with an audience can be further decomposed into the success of the site in meeting each type of visit that audience engages in. This type of model can even be weighted to reflect broader business goals.

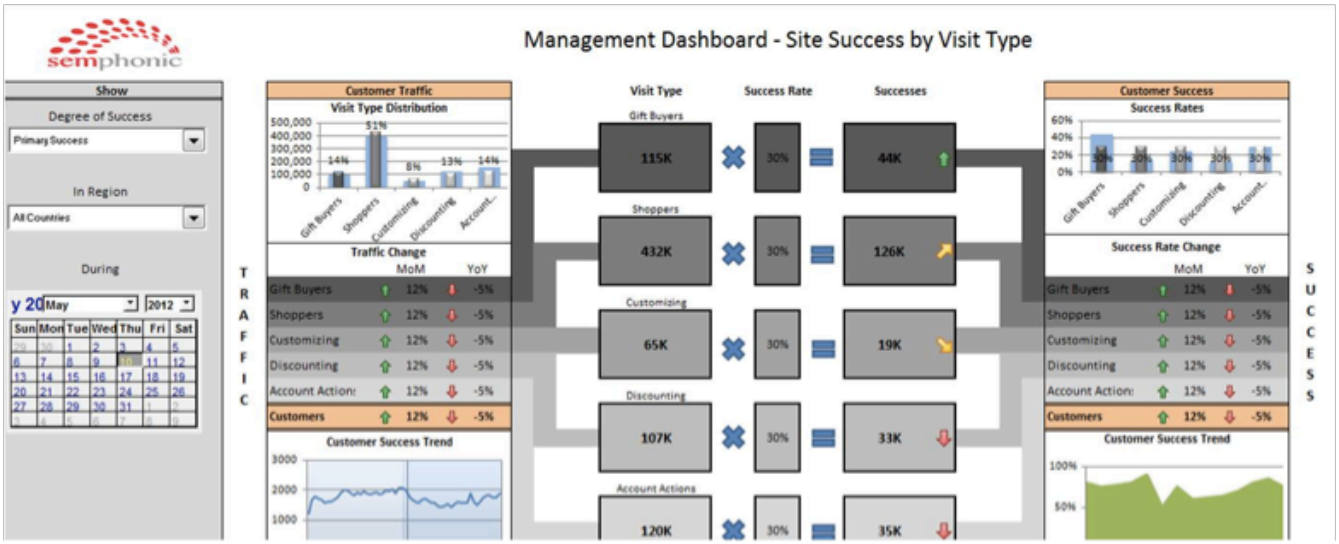

This is an elegant way to show how the different audiences and Visit Types fit together to form a true and accurate picture of site success. It's an elegant way to THINK ABOUT site success. If you've ever been asked "How is our site performing?" and realized at once how impossible a question it is, you should be able to appreciate this way of framing both the question and the answer. Because there isn't just one answer. A site may be doing very well with consumers but very poorly with industry professionals. It may be doing very well with high-net worth investors but not so well with average investors. Even if you could answer the uber-question - what meaning could it possibly convey? The site is doing great. Really? And what exactly does that mean? In creating report sets where Two-Tiered Segmentation is embedded in the model of the site, the answer to the question of "How the website is doing" carries both an explanation of why the answer has to be of a certain type and the conviction that the answer provided is a meaningful one.

These "systemic" dashboards are designed for senior audiences - people who need to know the high-level success of the site. But you can take the same approach (segmentation and systemic modeling) to more detailed views intended for line managers. Not only does this dramatically improve the quality of THOSE reports, it creates a powerful consistency of method and view across every level of the organization.

Here, for example, is a report of Marketing Channels shown as a system and optionally filtered by any of the audience and Visit-Type segmentations within the business:

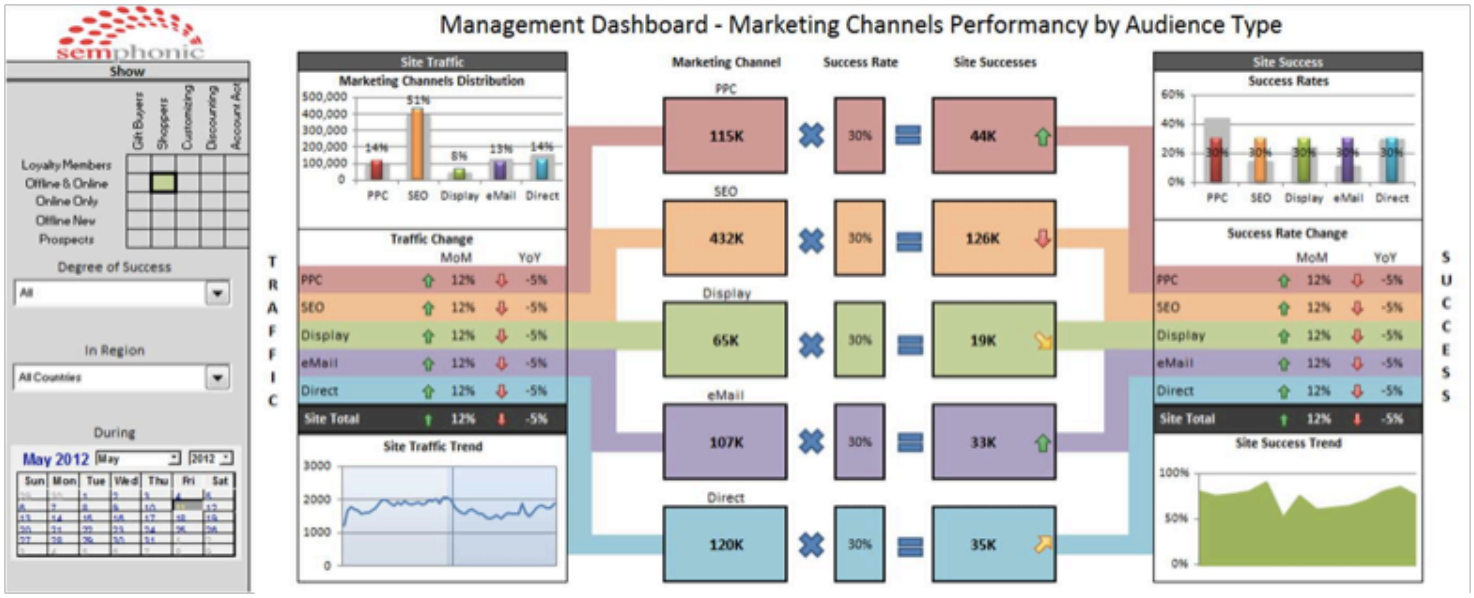

Website can be broken down into separate audience types. The flows provide a compact and powerful illustration of the critically important business information: what visitors the site attracts, how this is changing, and how successful the site is with each audience type.

With this type of report, you can measure the impact of each Marketing Channel in terms of both traffic generation and site success. But more importantly, you can understand how each Marketing Channel contributes to the audience make-up of your site and their corresponding Visit Types and success. This isn't just descriptive, it's prescriptive information that can be used to tune campaign creatives and testing programs.

The bottom line? You shouldn't have to choose between management dashboards that feature a few un-interpretable and confused metrics like Site-wide Satisfaction, Conversion Rate, and Visits or dashboards built from a daunting wall of complicated metrics. Creating site performance models based on Two-Tiered Segmentation makes for a style of dashboarding that is deeper, more interesting, more relevant, and more understandable than either of these all-too-common reporting choices.

The secret to great dashboards isn't Tufte-inspired white space, the right choice for line charts, or the best-looking font. It isn't finding the one magic KPI to rule them all. The secret to great dashboards is presenting information in a model that makes sense, that drives deeper business understanding, and that provides a real answer to the question: "How is my site doing?"

OK, What Next? Guide to Implementing the Analytics System

Using the methodology described above, Semphonic, Pentaho and Infobright have combined their respective technology and expertise to develop a reusable analytic infrastructure and dashboard that can be used as the basis for an organization's implementation. An interactive demonstration is available at www.webanalytics.infobright.com. Let's explore the expertise and technology used to build this use case.

Extracting, Transforming and Storing your Data

Semphonic developed the Two-tiered Digital Segmentation approach and it is fundamental to bringing customer-level analysis to digital data. As an industry leading consultancy, one of Semphonic's core competencies is identifying the correct data to use to meet business objectives. Through its years of experience with a broad range of Web tools, Semphonic has developed the methodology to extract, segment and transform this detailed Web data that is collected into a meaningful database layout. As one of the most difficult parts of an ETL (extract-transform-load) process is understanding how the data should be transformed to provide the intelligence needed by the business users, this is one of the key elements provided by this Semphonic-Infobright-Pentaho collaboration.

Infobright's high performance analytic database (Infobright © Enterprise Edition or IEE) is the intelligent repository for the data, with multiple options for loading the data into the database. Designed specifically for storing and analyzing large volumes of machine-generated data such as Web logs, IEE combines a column-orientation with intelligent technology that simplifies administration, eliminates the need to tune for performance, improves performance and data availability, and reduces total costs. In recent years, columnar orientation has become the de facto standard for high performance reporting and analytics. Additionally, Infobright's patent-pending compression algorithms result in average compression of 20:1 (raw data to size on disk), reducing storage requirements dramatically and letting companies store far more data history at much lower cost.

With high data compression, unmatched operational simplicity and low cost, Infobright is being used by hundreds of enterprises, SaaS and software companies in online businesses, financial services, telecommunications and other industries. Over 40% of Infobright customers are using the database for reporting on and analyzing Web data, including leading companies in online and mobile advertising, ecommerce and other online businesses.

How Infobright Works

IEE resolves complex analytic queries without the need for traditional indexes, data partitioning, projections, manual tuning or specific schemas. Instead, the Knowledge Grid architecture automatically creates and stores the information needed to quickly resolve these queries, as the data is loaded. The data is organized into two layers: the compressed data itself that is stored in segments called Data Packs, and information about the data which comprises the components of the Knowledge Grid. For each query, the Infobright Granular Engine uses the information in the Knowledge Grid to determine which Data Packs are relevant to the query before decompressing any data. By eliminating or reducing the need to decompress and read data to respond to a query, queries typically run significantly faster than using other technologies.

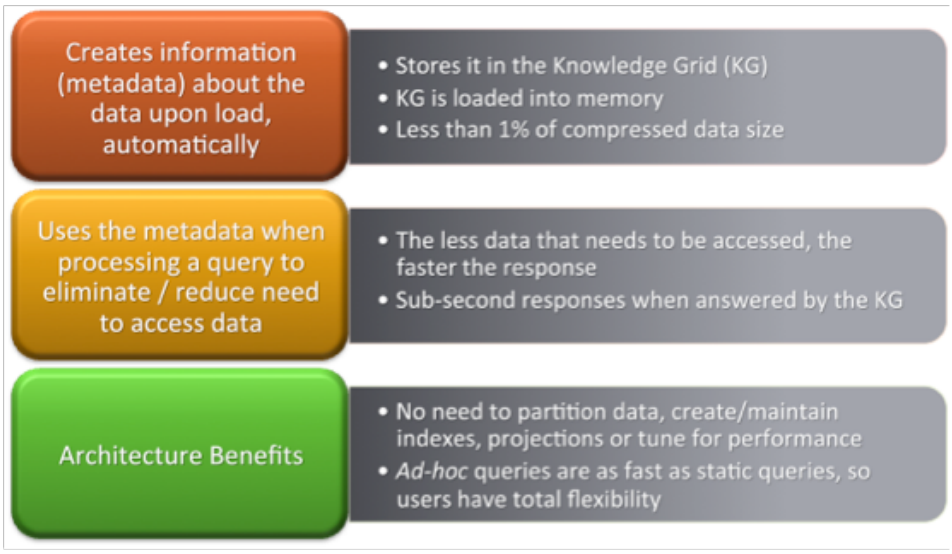

Business Analytics including Data Visualization and Dashboards

By tightly coupling data integration with business analytics, Pentaho brings together IT and business users to easily access, integrate, visualize, explore and mine all data that impacts business results. Pentaho's open source heritage drives continued innovation in a modern, unified, embeddable analytics platform that works with any data including big data and diverse data types. Pentaho Business Analytics provides a complete solution, is fast to deploy, easy to use, and extremely cost-effective — in short, delivering business analytics that work. The unified suite includes data integration, data discovery and exploration, and data mining.

Pentaho's Business Analytics platform has been used to build the sample Web analytics dashboard. The highly visual dashboard includes the capability to drill down based on different views of the data, and includes geo-mapping and heat grid presentation. The platform is featured below and includes: Data Integration, Data Discovery and Exploration, and Data Mining / Predictive Analytics.

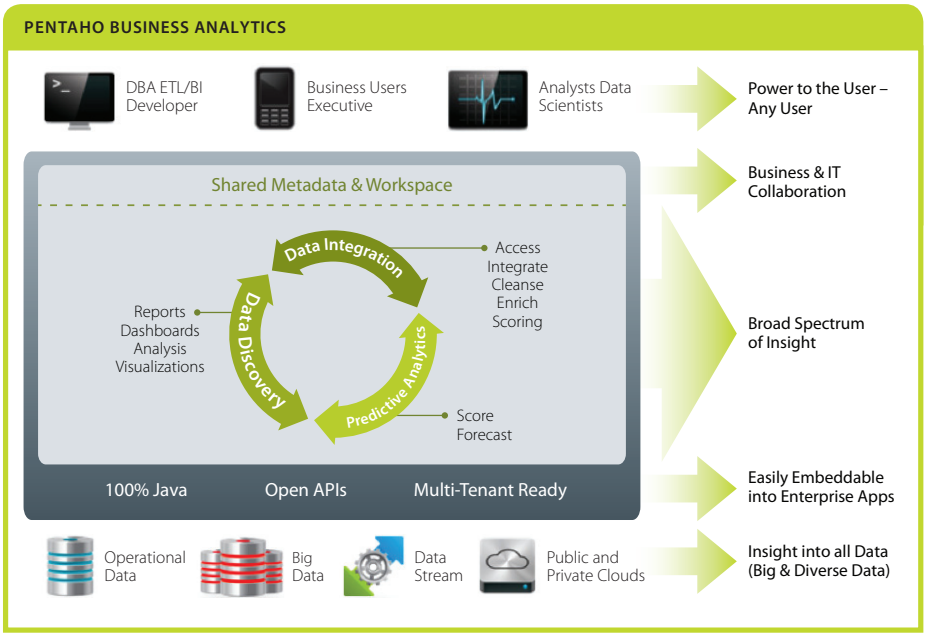

Pentaho Data Integration

With Pentaho Data Integration organizations can extract data from complex and heterogeneous sources and diverse data types to produce consistent, high quality ready-to-analyze data for powering business analytics. With a rich graphical user interface and a parallel processing engine, Pentaho Data Integration offers high performance ETL (extract, transform and load) that covers all data integration needs, including big data.

Data Discovery and Exploration

Pentaho Business Analytics provides a highly interactive and easy to use Web-based interface for business users to access and visualize data, create and interact with reports and dashboards, and analyze data across multiple dimensions, without depending on IT or developers. For IT, Pentaho Business Analytics is built on a modern lightweight high-performance platform and can be flexibly deployed on-premise, in the Cloud, or seamlessly embedded into other software applications.

Predictive Analytics

The powerful, state-of-the-art machine learning algorithms and data processing tools in Pentaho Business Analytics enable users to uncover meaningful patterns and correlations that may otherwise be hidden with standard analysis and reporting. These sophisticated, advanced analytics help plan for future outcomes based on a better understanding of prior business performance.

About Semphonic

Semphonic is the world's largest independent Web analytics consultancy, with headquarters in the San Francisco Bay Area and offices in Boston, New York, Washington, DC, Portland, OR and Berlin, Germany. Founded in 1997, the company has helped leading corporations, government agencies and non-profits achieve measurable improvement in the performance of their web channel. Clients include American Express, Charles Schwab, Genentech, the National Cancer Institute, Sears and Turner Broadcasting. Semphonic is also the driving force behind the premier web analytics conference, X Change. For more information about Semphonic, please visit: www.semphonic.com

About Infobright

Infobright's high-performance database is the preferred choice for applications and data marts that analyze large volumes of "machine-generated data" such as Web data, network logs, telecom records, stock tick data and sensor data. Easy to implement and with unmatched data compression, operational simplicity and low cost, Infobright is being used by enterprises, SaaS and software companies in online businesses, telecommunications, financial services and other industries to provide rapid access to critical business data. For more information, please visit www.infobright.com or join our open source community at www.infobright.org.

About Pentaho Corporation

Pentaho is delivering the future of business analytics. Pentaho's open source heritage drives our continued innovation in a modern, integrated, embeddable platform built for the future of analytics, including diverse and big data requirements. Powerful business analytics are made easy with Pentaho's cost-effective suite for data access, visualization, integration, analysis and mining. For a free evaluation, download Pentaho Business Analytics at pentaho.com/get-started.