

Responsive Content and The Bottomless Website

by Deane Barker for EPiServer

March 17, 2014

Abstract

One of the persistent challenges in marketing is addressing multiple audiences and contexts with the same content.

Media targeting has historically been used to increase the effectiveness of “push” marketing, but static marketing – like a corporate website – has suffered under its inability to target different audiences and contexts with the ideal content to engage visitors and prompt them to take action.

Unlike radio or TV, where an advertiser can choose when and in what context to push their marketing to the consumer, websites can be visited at any time, and therefore have to serve a variety of audiences with the same context.

Much like responsive design allows websites to visually adapt to their context (their browsing container), a new era of “responsive content” is allowing the content of websites to adapt to different audiences and contexts to more effectively target each visitor.

This paper is in three parts.

- **Part 1** is an examination of the historical problem and attempts to solve it.
- **Part 2** details the proposed solution – a comprehensive strategy for responsive content to allow a website to adapt to its audience.
- **Part 3** introduces EPiServer – a web content management system and marketing platform – and details its ability to manage and target responsive content.

Part 1: The Problem

Audiences and Contexts

The ultimate goal of web marketing is to offer content that engages your visitor and prompts them to take measurable, positive action (to “convert”).

However, catchall phrases like “the visitor” are disingenuous – there is no *one* visitor. On any given day, your website might be visited by thousands of people, each bringing different roles, needs, experiences, and browsing contexts that need to be addressed with the ideal content.

Thus, if we accept that an audience must have its needs met before they engage, and they must engage to convert, then we can restate our claim to this: the first goal of web marketing is to meet the needs of the audience in its current context.

Two definitions are necessary:

- **Audience:** an identifiable demographic group that shares similar needs. For example, people looking for a job from your company can be separated from people looking to hire your company. The needs of “Job Seekers” as opposed to those of “Potential Customers” are easily separated.
- **Context:** the combination of environmental factors that affect the current visitor. For example, a visitor sitting in their office down the street in on a Wednesday afternoon may have different content needs than the same visitor on their phone at 2 a.m. on a Saturday morning from halfway across the world.

The combinations of audience and context are potentially infinite. The same audience can differ based on the current context of individual members – a Job Seeker from your metropolitan area might be looking for a different position than a Job Seeker from three states away. The audience is the same, but their context is quite different.

Audience group and current context combine to define a specific visitor’s needs. Your website’s effectiveness is directly tied to your ability to react to and meet these needs. The closer you can meet the needs of our audience, the more likely they are to convert.

Servicing Infinite Needs with a Single Website

If we accept that there are never-ending combinations of audiences and contexts, we immediately identify a problem of delivering one website: we cannot possibly meet all the possible gradients of visitor needs with the same content.

What marketers tend to do is simply narrow down the domain of possible needs to the lowest common denominator – for example, both Job Seekers and Potential Customers need to know general information about your organization, and your organization needs to be presented in a positive light.

Clearly, this is a very crude attempt at engagement.

Alternately, you can encourage the visitor to self-identify by providing clear and easy-to-use navigation options. A link label “Career Opportunities,” for example, is an attempt to divert Job Seekers from other, less relevant content. At the other extreme, you might be even more overt in our attempts: some websites have drop-down elements labeled “I am interested in...” with options meant to identify that visitor and segregate them into an audience.

But, again, this is a very crude and imprecise method of identifying your visitor's needs.

In the last few years, there's been a strong tendency towards micro-sites and custom landing pages. Organizations might have dozens of sites and landing pages specifically targeting various audiences. (It's important to note that these audiences are not positively identified, other than assuming that only members of those audiences responded to the marketing media targeted to their demographic.)

The development and organizational costs of these sites can quickly spiral out of control. Conceiving, designing, developing, populating, and maintaining separate web properties for each audience is not a scalable strategy.

Note that this problem, while universal to marketing, is especially acute in website marketing. With television or radio, a science has developed around media targeting – the practice of pushing your content into a specific channel which is most likely to reach your target audience.

However, a website is less proactive – in many cases, it simply reacts to an anonymous visitor. While you can target advertising to drive visitors to your website, there is no guarantee that a visitor will respond to a campaign in the way you expect them to and by a method in which you can identify them.

For example, a search engine marketing campaign targeted at phrases used by Potential Customers is trackable, and we can use this to drive visitors to specific landing pages, targeted especially to their needs. However, nothing stops them from simply taking note of your company name (from that campaign or other source) and visiting the home page of your website at a later time, in a method which is now completely anonymous and totally circumvents any attempt at targeting. Since denying access to anonymous visitors clearly isn't an option, how do you address their needs when you have no way of knowing what audience they represent?

This is the problem that plagues website marketing: how do you identify the current visitor's needs by uncovering their audience group and context? Furthermore, once you have determined this, what can you do about it? In what ways can you dynamically manage the content of your website to address their needs, engage them with your content, and ultimately prompt them to action?

The Bottomless Website

The solution is easy to explain but hard to execute: you need a single website that can identify and adapt to *any* visitor. Your website must be able to analyze the inbound visitor, determine what their needs are, and re-arrange itself to best meet those needs.

To do this, you CMS must solve several problems:

1. It must be able to segregate (or “bucket”) the visitor in real-time into an audience based on clues that it detects and analyzes.
2. It needs to store and manage content in such a way that it’s not tied to the single-page aggregation model.
3. It needs to be able to mix-and-match content on-the-fly to deliver custom aggregations of content to your visitors based on that which you think will best meet their needs at that moment.

In this sense, you need a “bottomless website” from which you can produce an infinite number of visitor experiences from a single web destination, and deliver the most effective experience to the current visitor.

Part 2: Introducing Responsive Content

Sleuthing the Anonymous Visitor

Website personalization is not new – content management systems have been claiming this functionality for years.

However, it was always predicated on one thing: *knowing who the visitor was*, which was achieved by having the visitor log into an account (the ultimate in self-identification). In this sense, “personalization” was limited to targeting content to specific, known groups of users, and/or allowing users to customize their own experience.

However, most marketing scenarios, the visitor is unknown to you – their session is anonymous. How, then, do you achieve “anonymous personalization”?

To do this, you can draw clues from three general domains of information:

1. *Their Web Request*

When a visitor loads a page on your website, their browsing device has sent a request for content. In this request is quite a bit of information.

- **Their IP address**, from which you can geo-locate their rough location.
- **Their browsing device profile**. At the very least, this will identify their browser, and in many cases it will identify their viewing parameters (screen size, etc.)
- **Their language preferences**. Often, multiple languages are specified, in preferential order.

- **The website or referring URL** from which they linked to your site. Alternately, this is empty, which often indicates they are an “organic” visitor that proactively typed your URL in the address bar.
- **Any tracking or analytics code** in the URL which would tie this visit to a specific advertising or social media campaign.
- **Their connection speed** (while technically not sent with the request, it can often be derived).
- **The search keyword** they used to find your page in search engine results.

2. *Their Browsing Behavior*

After receiving the initial inbound request, you can track the visitor’s behavior during that session and answer such questions as:

- On what page did they initially enter your site?
- What content are they interested in? Specifically, how many pages have they visited in Content Area X, or how many pages have they visited that contain Tag Y?
- Have they visited Page X?
- Have they filled out Form X?
- How many total pages have they visited?
- What is their total time on your site?
- Did they use the search feature on your site? What did they search for?

3. *Global Information*

Beyond the visitor-centric information of their request and behavior, there is more site-centric, or global information available:

- What time of day is it?
- Is it a weekday or weekend?
- How is the site performing? How much load is it under?
- Do you have new content available since the last time they visited?

The above information can be detected, combined, and analyzed to provide remarkable insights on the current visitor’s audience and context.

Example: The Job Seeker

This visitor comes to your site with the specific intent to seek employment with your company.

- Was the visitor referred to your website from monster.com, careerbuilder.com, or another job board?
- Did they come from a search engine with the word “job” or “career” in their search?
- Have they visited your “careers” page on this visit, or a prior visit?

Example: The Prospective Customer

The visitor comes to your site with the specific intent to hire your organization for services or purchase the product you offer for sale.

- Has the visitor downloaded your whitepaper?
- Has our visitor visited X pages in the “products” section of your website?
- Has the visitor come to our website via a review of your product? (More important: which specific product was being reviewed?)
- Is this visit during business hours?
- Is this visit from your geographic area? (Important for geographically-limited services, such as property maintenance or medical services.)

Adapting Content Around the Visitor

At this point, you’ve been able to make educated assumptions about the visitor and their needs. Now what can you do with this information?

You adapt your content to their needs. You change your website in small or significant ways in real-time in order to present the visitor with content and an experience that matches their specific needs at that specific moment.

- You replace default content elements with content created specifically for that audience.
- You rearrange navigation to highlight options they might be interested in.
- You direct them to custom initial landing pages that would appeal to them.
- You show or hide paragraphs of text in narrative sections.
- You alter graphical and photographic elements to those designed to appeal to them specifically.

In the process, your website ceases to become a single thing. There is a “default” view where – barring any detection or analysis of the current visitor – your website displays “default” content.

However, in the ideal situation, every page on your website becomes an on-the-fly aggregation of content specifically designed to appeal to *this* visitor in *this* context at *this* specific moment.

From a content perspective, your website becomes a bottomless well from which you assemble dynamic, targeted experiences which might be different for each visitor.

Managing Responsive Content

Clearly, identifying the specific audience and context of your visitor and adapting your content to their needs is highly effective. To achieve this, however, you must change how you manage content.

Specifically, content can no longer be managed at the page level. Managing content in aggregations this large gives you very little presentational agility to mix-and-match or “mash-up” content in real-time to address your current visitor.

Rather, content must be managed to the element level, defined as the “minimum reusable unit” suitable for dynamic placement throughout the site.

This includes:

- Promotional banners
- Self-contained paragraphs of copy
- Content teasers
- Embedded images

With a refined personalization framework, pages can be delivered to users via one of two models:

- **Pure Dynamic Pages:** the entire page is constructed in real-time. All that exists in the content management system is a placeholder (“Landing Page Container,” for instance) which is populated at request time.
- **Enhanced Static Pages:** the core content of the page (a news article, for instance) is managed, then enhanced by selected dynamic content elements.

Part 3: EPiServer as a Responsive Content Platform

About EPiServer

EPiServer CMS is a web content management and digital marketing platform based on Microsoft's .Net architecture. Under continuous development since the 1990s, EPiServer provides a highly usable, scalable platform for web content management and digital marketing.

Segregating Audiences and Contexts with Visitor Groups

The core of EPiServer's personalization architecture is the "Visitor Group." This is a conceptual "bucket" of visitors, defined by criteria. Fulfilling these criteria – both individually and in combination – designates a visitor as part of a given Visitor Group and allows content to be personalized for them.

Several criterion are available out-of-the-box. They include:

- **Geographic Location:** specified by continent, country, and region (state)
- **Geographic Coordinate:** specified as a distance radius around a specified point)
- **Time of Day:** specified as a day of the week, and an optional time range
- **Number of Visits:** specified as less or more than a supplied value in total, within a time range, or since a past point in time
- **User Profile:** specified as whether a value in the visitors ASP.Net profile matches, contains, starts or ends with a supplied value (note that this would require the visitor to login, which admittedly doesn't fit the "anonymous" model)
- **Visited Category:** specified as the minimum number of times the visitor has requested pages tagged with a specified category in EPiServer
- **Visited Page:** specified as a single page which the visitor has visited at some point in the current session (this could be a form confirmation page, which would signal form completion)



Name	Notes
Job Seekers	Referred from one of the job boards or se
Local to Chicago Metro	
On Phone	
On Tablet Device	
Prospective Customers	Downloaded the white paper, or visited mu
Referred by 2014-03-01 "FT" Review	Clicked the link in the third paragraph.
Returning Visitor	Visited at least once before.
Visiting During Business Hours	

Example Visitor Group definitions

- **Landing URL:** specified as whether the first site URL the visitor requested in the current session equals, contains, starts or ends with a supplied value
- **Referring URL:** specified as whether the domain name or URL from which the visitor entered the site equals, contains, starts or ends with a supplied value
- **Search Keyword:** specified as the inbound search keyword in the referring search engine URL

The screenshot shows a 'Criteria' configuration window with four rows of criteria:

- Row 1: Search Keyword criterion with the value 'career'.
- Row 2: Search Keyword criterion with the value 'job'.
- Row 3: Referrer criterion with 'Host name' selected, 'Equals' operator, and 'careerbuilder.com' value.
- Row 4: Referrer criterion with 'Host name' selected, 'Contains' operator, and 'monster.com' value.

An example of using criteria to define the boundaries of a Visitor Group.

A final meta-criterion provides a subtle but powerful way to further define Visitor Groups.

- **Visitor Group:** specified as the whether or not the visitor has already satisfied the criteria for another Visitor Group.

While seemingly superfluous, this is actually quite powerful as it allows the extension of an existing Visitor Group. A “base group” can be defined (call it Group A), then more specific groups (call them Groups B and C) can be defined by simply add one or two criterion in addition to the criterion of belonging to Group A. Thus, changing the criteria required for Group A automatically affects Groups B and C.

Criteria are completely extensible. EPiServer has provided a complete API whereby a developer can write custom code to define new criterion specific to the project. (See “Usage Examples” below for an example of custom criterion.)

Criteria can be combined to define the boundaries of a Visitor Group. Combinations of criteria can be designated as:

- **All:** all criteria must be satisfied for the user to be considered a member of the Visitor Group
- **Any:** any single criterion can be satisfied for the visitor to consider a member of the Visitor Group
- **Points:** each criterion can be assigned an arbitrary point value, and a threshold is specified. As criteria are satisfied, the visitor’s point total increases. At the moment the visitor’s point total crosses the threshold, they are considered a member of the Visitor Group. (Note that this might occur instantly, as they enter the site, or later in their session when their behavior on the site satisfies further criteria which increases their point total.)

Using Visitor Groups to Personalize Content

Once Visitor Groups have been defined, they can be used to alter the website in real-time. This might occur upon delivery of the first page (if their inbound information – for example, their geographic location – was enough to place them in a Visitor Group), or it might occur between page requests (if their behavior on the site causes them to fulfill enough criteria to place them in a Visitor Group).

There are several ways to use Visitor Groups to modify content:

- Editors working in EPiServer can drag individual content elements (called “Blocks” in EPiServer’s nomenclature) onto a page surface. These blocks are visible to all visitors by default, but can be specified as visible to only one or more specified Visitor Groups.
- The access rights of pages can be bound to Visitor Groups, meaning that membership in a Visitor Group affects the current visitor’s ability to see specific pages, which means that these pages might no longer be visible in the site’s navigation menus.
- Sections of narrative content inside EPiServer’s WYSIWYG editor can be visible only to one or more Visitor Groups. Content can therefore be shown or hidden down to the paragraph level.
- At the code level, the Visitor Groups API allows developers to make essentially unlimited changes to site behavior based on the visitor’s membership in a specified Visitor Group (or lack thereof).



Using Visitor Groups to specify content display rules.

Usage Examples

Here are some examples of using Visitor Groups, combinations of criterion, and personalized content to increase the effectiveness of your EPiServer-powered website.

Visitor Group: “Job Seeker”

- Visitor’s referring domain equals monster.com or careerbuilder.com
- Visitor’s inbound search keyword contains “jobs” or “career”

Default promotional content on the home page is replaced with a list of current job openings. Sidebar promotion content on each page is replaced with a career opportunities promotional banner.

Visitor Group: “Healthy Eater”

- Visitor has consumed three news articles assigned to the “Health and Nutrition” category
- Visitor has searched the site for “calorie,” “glycemic,” or a list of other nutrition-related words

A new group of pages on nutrition information becomes available and is added to the main navigation of the site.

Visitor Group: “Ice Cube”

- Visitor’s current temperature is less than 40-degrees F (this would be a relatively simple custom criterion which cross-references the visitor’s geo-located latitude/longitude with the World Weather Information Service API to determine the current temperature in that location)

A banner on the home page advertising the current special vacation package is replaced with an offer for Caribbean cruise packages. The color scheme of the site is globally altered to incorporate warmer color accents.

Conclusion

A single website is a crude way to market to a varied visitor base. With unlimited combinations of audiences and contexts, a single website only provides the ability to cater to the lowest, most common needs of all visitors.

By segregating visitors by audience demographic and current context, site elements can be displayed, hidden, and modified in such a way that a single website becomes theoretically “bottomless,” enabling the advanced marketer to retrieve unlimited combinations of content, each designed specifically to address the current visitor.

In the process, the website becomes agile and fluid – constantly adapting, changing, and shifting to ensure maximum effectiveness for each specific visitor.

Identifying new audiences or opportunities doesn’t involve new campaigns or microsites. Rather, it simply involves identifying and isolating new Visitor Groups and adjusting the site’s content in a matter of hours or days, rather than involving development teams and the associated time and expense.

The resulting competitive advantage is undeniable. Each visitor to your website receives a personalized experience based on their current needs which is specifically designed to prompt the desired action, whether it be making a purchase, completing a lead generation form, or visiting a promoted link.

About the Author

Deane Barker is a principal of Blend Interactive, a content management consultancy based in Sioux Falls, South Dakota.

Blend Interactive was EPiServer's first North American implementation partner, and has been specializing in EPiServer solutions for six years. Deane has been writing about content management for over a decade and speaks about related topics at conferences around the world.