

7. Transit Web Site Technology Considerations

When the Work Plan for this J-09 Task Order was written, the project team had expected to obtain basic technical information about hardware and software requirements during the research. Since the research focused on three advanced Web site features (AIP, RTD, EMN), the project team inaccurately anticipated a greater depth of technical knowledge than it found.

However, the initial research showed, and the telephone interviews confirmed, that most of the information on transit Web sites focused on the *functions* rather than the technology. This was especially true for the three advanced Web features, which became the focus of the research. Typically, individuals responding to the telephone surveys had only a vague sense, if any at all, of what actually runs the Web site features. In retrospect, this finding was not surprising given that:

- Technologies for the Web, including information management and communications, are changing rapidly
- Most staff with policy-level responsibility for an agency's Web site typically do not have a background in information technology; and
- A wide variety of technological approaches can be used to achieve the same end on a Web site.

Nevertheless, the telephone interviews did demonstrate a definite pattern of transit agency concerns about Web site performance, system and data maintenance, requirements to add functionality, and the like. Using that as the starting point, the project team decided that an overview of fundamental Web hardware, software, and communications issues would be most valuable to the reader, in lieu of trying to ferret out specific technical issues of the participating transit agencies.

Therefore, the remainder of this Section will address general concepts and concerns related to the following technology and application issues:

- Application Design, Development, and Performance Issues;
- Functionality;
- Privacy Concerns;
- Security Concerns;
- Application Deployment; and
- Customer Relationship Management.

Although CRM received mixed reviews during the telephone interviews, the project team believes respondents showed enough interest, or apparent misconceptions about the complexity, cost, or value of CRM, that the topic should be covered in this Section. (CRM is also discussed in the best practices part of Section 9.) The project team found ample evidence of CRM functionality or design on transit Web sites to support the position that CRM may have an important future for the transit industry. While a range of telephone interview respondents focused on improving mapping capability, CRM could potentially be applied to offset some of the constraints of mapping described below.

7.1 Application Design, Development, and Performance Issues

Once a transit agency has decided to implement one or more advanced Web site features, a host of questions must be addressed concerning how to acquire or construct the required applications and how to deploy them. This section will discuss the following questions: (1) Should a transit agency buy or build its Web site features? and (2) how can the system architecture of a transit agency Web site or advanced customer information feature affect application development or performance?

Buy or Build?

The decision to buy applications or services presupposes that one or more appropriate candidate applications already exist. Transit agencies should consider not only which systems and components will meet all anticipated needs, but also applications that might be adapted or augmented to do so. Additionally, agencies should consider whether trading off a subset of expected needs that are available at an attractive price – versus buying a more feature rich but expensive product – is worth it. Buying an application or component can be particularly fruitful if the product’s functionality can be extended by integrating it with other information technology components.

For example, a commercial AIP product with excellent features, but an unsatisfactory customer interface, could be married with a custom-designed interface that organizes and presents the AIP functions in a better way. The key to successful integration is the availability of some type of efficient interface between the components. This could take the form of an API (i.e., application program interface),⁶⁴ XML document transfer,⁶⁵ or some other mechanism that provides access to the desired functionality from any additional component(s). Transit agencies that plan to deploy or upgrade customer information features on their Web sites should consider integration issues carefully.

The advantages of utilizing an existing system or product typically include cost (compared to building from scratch), time (it typically takes much longer to build a new system), and reduced cost and scheduling risk (since part of the work is already done, although this apparent advantage may be offset by the need to rely on vendors rather than staff under one’s direct control). The principal advantage of building a system from scratch, either with internal skills or contracted assistance, is that a transit agency can get what it specified rather than being limited to off-the-shelf products and the probable need to customize them. Mock-ups, prototyping, and pilot trials are often essential to close the gap between an initial Web site concept and a useful final design, since organizations often have difficulty envisioning how an abstract system will serve the needs of an undefined group of existing and potential customers.

⁶⁴ For a succinct definition of API, visit the Webopedia.com Web site at <http://www.pcwebopaedia.com/TERM/A/API.html>.

⁶⁵ For a complete definition of XML, visit the Whatis.com Web site at http://searchwebservices.techtarget.com/sDefinition/0,,sid26_gci213404,00.html.

System Architecture Issues

In direct contrast to the PC-centric distributed computing model that has dominated office applications, modern Web-based applications are host-based.⁶⁶ That is, most (if not all) computations are performed by one or more powerful computer servers, and the client (whether a personal computer at work or home, personal digital assistant (PDA, such as a Palm Pilot), cell phone, or other device) is primarily responsible for displaying output from the server and gathering input to pass to the server. This is known as a thin-client design, because so little takes place on the client.⁶⁷ For applications that must be accessed by the general public using a variety of devices, as is increasingly the case with transit customer information, such a design has important advantages such as it:

- Limits the amount of data that must be passed between the client and server. For wireless applications and dial-up Internet connections, in particular, minimizing bandwidth (see below) is a critical concern;
- Allows the application to run on devices with very little computing power, such as PDA's and possibly cell phones; and
- Can utilize a standard protocol for communications between the client and the server. If a browser exists for a particular device, that device can access the application.

With virtually all computing being performed by the servers, transit Web site design decisions must ensure that these servers can be very responsive to the demands of transit customers. The standard approach, which is used for many types of high volume applications including Web sites, is to disaggregate necessary functions so that multiple servers can share the workload. This can be accomplished by:

- Having each server perform all functions and designing the application so many servers can do this simultaneously;
- Breaking the work into distinctly different tasks and having each server specialize on a task; or
- Doing both, by using many specialized machines for each distinct task.

These various specialized tasks are typically divided into three types or "layers:"

- Presentation (i.e., the work that takes place on the user's personal computer desktop or PDA);
- Business logic (which services requests from the presentation layer; this layer includes all the algorithms required to perform calculations, flow logic, etc. to perform whatever functions are needed, such as calculating a transit itinerary); and
- Database management (reading and storing data to/from the business logic layer).

⁶⁶ For a brief overview of "distributed computing," visit the Whatis.com Web site at http://searchwebservices.techtarget.com/sDefinition/0_sid26_gci760724.00.html. For a brief definition of "host," visit the Webopedia.com Web site at <http://www.pcwebopaedia.com/TERM/h/host.html>.

⁶⁷ For a brief definition of "thin client," visit the Webopedia.com Web site at http://www.pcwebopaedia.com/TERM/t/thin_client.html.

In client/server or two-tier architecture, the presentation and business logic functions are performed on the user's PC, which makes requests directly to a database server across a network. This requires computing power on each desktop and considerable bandwidth (or data communication capacity), to the database server, since the PC needs all the data required to solve the problem.

Bandwidth is a proxy for data transmission speed, and the easiest way to visualize its impact may be to think in terms of a pipe carrying water: the bigger the pipe, the more water it can deliver. For a given size pipe, the less water you need to move, the quicker you will finish the job. (An example of bandwidth impacts that readers are probably familiar with is the slower Web page access over a typical dial-up Internet connection compared to that of a much higher speed connection provided over an agency's network.) Since each customer has only so much bandwidth available, limiting the amount required by an application makes that application faster for everyone.

In contrast to the client-server model, in a three-tier design, presentation is still handled by some sort of input/display device used by a customer (such as a PC, PDA, or cell phone), but it does not need to be a fast computer. Business logic, (such as the analysis to find the best trip itinerary preferences, is handled by a remote server via the Internet), and database management is handled by another server (typically located near the business logic server). The user, in this case a transit customer, makes a request at the presentation layer for information from the business logic server, which determines the necessary data and requests it from the database server. This database server responds and thus allows the business logic server such as an AIP program to answer the user via a Web browser.

Using AIP systems as an example, the transit customer uses a computer connected to the Internet to define travel preferences (presentation layer), this information is sent to an AIP program to formulate an itinerary (on the business logic server), and the AIP program requests additional information from the database such as the location of a landmark or particular transit route schedules (the database management layer). After the AIP program receives the schedule information from the database, it calculates the itinerary and then sends it back to the transit customer who views the results on a computer's Web browser.

This design approach can be extended to more than three tiers, for example, by having a business logic server that calls programs from another business logic server that then accesses the database. Also, because the technology supports having multiple business logic servers and multiple database servers, and can actively direct requests to balance the load across available machines, this approach is highly scalable. That is, the system can be expanded as needed to meet additional demand, without wasting the initial investment. Furthermore, when customers see performance slow (because of a bottleneck), it can often be resolved by adding another computer server to improve capacity for only the one bottlenecked function. This kind of design is essential for applications that may see very high demand, such as those from large and/or regional agencies, and with data intensive processes, such as providing interactive, GIS-based maps with a transit itinerary.

Servers can usually be assembled to guarantee a desired level of responsiveness. The more difficult problem is achieving that responsiveness over an Internet or wireless connection, where bandwidth is constrained. An application that processes requests in a

tenth of second, but takes 10 seconds to deliver those results to a customer is simply too slow. Indeed, the project team experienced significantly different response times on the many itinerary planning sites it reviewed and tested.

7.2 Functionality

Functionality can greatly both affect an application's usefulness and how much bandwidth it requires. Consequently, it is important to consider design features both in terms of what they offer the customer and what they take away, in the form of delays. In this context, some of the major considerations for transit Web site applications include the following items:

- Mapping;
- Data Push Versus Pull; and
- Use of Real-Time Data.

Mapping

Integrating maps into a transit Web site has both advantages and disadvantages, and was a frequent topic of discussion during the telephone interviews. In an AIP application, the map might be used to display a proposed trip or walking directions from a customer's last transit stop to the final destination. Maps can be very helpful to the transit customer; they provide an overview of the itinerary and indicate where transfer points are located, the proximity to landmarks, etc. A map displayed on a transit customer's Web browser can also be used as an input mechanism, similar in function to text fields for addresses, intersections, or landmarks found on AIP Web pages. With a map interface, for example, a customer could point to a particular street, intersection, or area without the need to know the cross street or building number to define the origin and destination of the desired transit itinerary.

Constructing the map is typically time and labor consuming. It requires the initial creation of an accurate map for the service area and then timely attention to adding new streets, subdivisions, landmarks, routes, etc. Depending upon the quality of the existing data and the frequency of changes to the local geography, this can be a significant labor burden.

A second consideration when incorporating a map is the simple fact that maps require much more bandwidth than text. As a consequence, Web pages containing maps download slower to the customer than those that avoid maps, especially for customers with dial-up Internet connections. One way to minimize this problem is to make map displays an option; that is, the customer would have to consciously request the map in order for it to display.

In any event, there are ways to speed up the loading process. For example, reducing the detail in the map and restricting its size reduces the amount of information required to draw the map and subsequently transmit it to the customer. If the map must be scalable (as in interactive or dynamic applications), it should be drawn as a vector, rather than a raster or bitmap image. A vector-based map graphic requires more processing power, but a bitmap image shows ragged edges when zoomed and typically looks unacceptable.

These issues apply whether the user uses a PC or a handheld device (although the applicability of maps to very small handheld displays is questionable).

Data Push Versus Pull

Data is said to be “pulled” when the user must request a new Web page or refresh an existing one in order to see new information. Thus if the Web site user does not make a such a request using the Web browser, the displayed data will be static and may be out of date. This would be the case, for instance, on a real-time transit information site where the map display of vehicle locations did not update itself automatically. Until the user requested refreshed data, the old data would continue to display. In contrast, a system that automatically updated its display (either on a regular schedule or when an event occurred) is an example of data that is “pushed”; the application continuously or periodically sends new data to the client. This is the case with a real-time, Web-based transit information service, such as KC Metro’s *BusView*, which automatically updates the information display at frequent intervals.

Data needs to be pushed when the user needs to be alerted about an important event of which he may not be aware (e.g., the bus is late). One analogy to transit Web site features would be automatic e-mail notification system. The nature of the application determines how big a “price” one must pay in bandwidth to push the data. If the user simply gets a short message when important events occur, the demand on bandwidth is not high. If the application provides very dynamic, graphical information, the bandwidth requirements and associated costs could be high. In the extreme, this can resemble Internet video streaming, where the size, detail, and refresh rate of the moving image must be tailored to provide a pleasing image for any particular Internet connection.

Use of Real-Time Data

Transit schedule information can be categorized as either “planned” or “actual.” Planned (i.e., static) schedules are those published in public timetables and usually posted on a transit agency’s Web site. These are the typical basis for itinerary planning systems. Actual schedules reflect the status of current operations, such as those displayed on monitors in transportation terminals or used as the basis for personalized e-mail services. Airport arrival and departure monitors are one example. Actual schedule data is essentially real-time in nature.

Real-time data reporting relies on a data collection and processing infrastructure, such as Automatic Vehicle Location (AVL) systems, which can track vehicles and compare their actual and planned locations at selected points in time. Such information is the basis for computer-aided dispatch (CAD) systems used in the transit industry, real-time notification systems, and transit terminal displays. However, its integration into AIP systems has yet to be implemented, as far as the project team knows. For itineraries planned significantly in advance of departure, real-time deviations of transit vehicles from published schedules may not have much bearing on schedule adherence of much later trips. (In fact, if there were significant delays on a leg of a particular transit trip, the transit customer could find himself on a delayed earlier trip, which might then fit into the itinerary parameters.) The perspectives of transit industry representatives who participated in the telephone interviews were discussed in Section 7.

When headways are long, however, and the trip is planned shortly before departure, an AIP system that incorporated real-time schedule adherence data could alert the customer to delay their departure or that a planned transfer is not feasible. A simpler approach, already in use, might be to have the system access an incident database to determine if an itinerary calculated with planned schedules utilizes a transit route or roadway affected by a known incident, such as ongoing construction, and simply indicate a risk of delay.

The use of real-time data does not imply, by itself, significant use of communication bandwidth. Simple text updates every 30 seconds, for example, should not be a serious problem. By contrast, graphical updates that require retransmission of dynamically created maps, however, may be essentially infeasible (for practical purposes) until technology brings greater bandwidth to the typical transit customer.

7.3 Privacy Concerns

In the early days of the Internet, the public ignored or was unaware of issues about what types of information were being collected about them on-line. Either as a result of personal experience or through access to news articles decrying the situation, public awareness has now been raised, and people are both more reticent to knowingly share certain information and more cognizant that information might be shared without their knowledge.

Obtaining and maintaining trust about the privacy of personal information is critical to the customer relationship, and thus to the success of CRM services. Despite their status as public institutions, transit agencies that are planning CRM must have a clear policy and statement, for both internal reasons and to ensure customer confidence, about how information collected from their customers will be used and shared (if at all). One approach used on a few commercial Web sites is to provide customers the ability to choose for themselves how personal information is used and shared. Regardless of stated intentions, however, it is also necessary to ensure that such information is stored as securely as possible to protect it from internal misuse or from external “hackers.”

7.4 Security Concerns

Every application accessible via a public data network such as the Internet, regardless of its purpose, is a potential victim of an attack by hackers. The range of threats runs the gamut from denial of service attacks (in which the Web server is deliberately overwhelmed with machine-generated requests and therefore cannot service legitimate customer requests) to modifying or defacing the site. Theft of sensitive personal and/or financial information has occurred. Transit agencies cannot assume that they are immune from such malfeasance.

No organization wants its Web services rendered useless by such attacks, or worse, to see privacy compromised or its machines subverted for improper use. Some of the measures that minimize the risk should be obvious to transit officials charged with developing Web applications. For example, the Web computer server should be protected by an effective firewall and the administrator should be vigilant about applying security patches as they

are provided by the manufacturer.⁶⁸ However, the struggle between information technology security professionals and hackers is a dynamic one, and it is not the purpose of this document to provide a treatise on Internet security. The bottom line is that security is a complicated issue that requires the attention of a knowledgeable IT professional who has a full understanding of how a transit agency Web application will be utilized and by whom.

7.5 Application Deployment

Given the availability of high-speed access to the Internet by both organizations and individuals, the question of where Web-based applications should reside is ultimately a matter of cost. In general, the decision can be divorced from that of which functions are to be provided and who develops the applications.

If a transit agency already hosts other Web-based applications (and thus has already addressed issues of security, support staff availability, and bandwidth), and the existing infrastructure can handle the additional load, the cost of adding one more application into such an environment would typically be less than that of housing the application off-site.

However, if the organization has not already addressed these issues and made the necessary investments in equipment and staff, consideration should be given to the alternatives of co-location and Web-hosting. In co-location, an organization purchases its own hardware platform(s) for the Web application (i.e., the computer or computers on which the application will run), but these servers reside off-site at the premises of a company whose responsibility is to securely house, administer, maintain/repair, and provide connectivity and security for the co-located hardware. The application itself, whether an AIP, RTD, or EMN system, is still the responsibility of the transit agency.

When a transit agency is not prepared to provide:

- Staff or arrangements for 24x7 system support;
- A secure data environment; and
- Sufficient Internet bandwidth to rapidly handle all customer requests;

co-location can be an attractive option, because the vendor takes on these responsibilities, but the transit agency's Web application still resides on a dedicated computer server or server(s) specifically configured and tuned for that particular customer information application. (Tuning refers to configuring the hardware and software to provide optimum performance of the functions it provides.)

Concerns about getting access to a co-hosted machine are normally unfounded, since local vendors are often available (to facilitate those occasions where physical access is required). Access to the application software can occur over the Internet and therefore should be no different for a co-located server than one maintained in-house by a transit agency, except that the connection to a co-hosted server would be limited to Internet bandwidth. This should not be a serious issue, since such access is usually required only to make changes.

⁶⁸ Review a definition of firewalls at <http://www.pcwebopaedia.com/TERM/f/firewall.html>.

When evaluating co-location providers (and even the strategy of co-location), transit agencies should carefully examine the specifics of the “Service Level Agreement” (SLA) supporting the applications, to ensure that it includes all of the services and the necessary responsiveness required.

Web hosting provides the same types of advantages as co-location, except that the vendor owns the hardware. This can be a useful alternative if the application’s hardware requirements are ordinary, and particularly if the traffic does not require dedicated computers; in this case, the hardware capital and maintenance costs are spread over several companies, and there can be significant cost savings. As with co-location, the contents of the SLA should be critically evaluated. In addition, as 2002 proceeds, it is clear that the financial health of some smaller Web hosting companies is becoming critical. Transit agencies, therefore, factor in the probability of uninterrupted service into the decision to select a specific Web hosting company.

7.6 Customer Relationship Management

The concept of CRM is well established in the commercial, *eCommerce* marketplace. Companies such as Siebel Systems, e.piphany, Oracle, PeopleSoft, and others provide enterprise-level systems designed to help companies manage their relationships with customers. These systems provide a wide variety of information about a customer’s product needs, prior and planned purchases or activities, and the potential for purchasing additional products.

The CRM concept is that, by better understanding customers and addressing their needs, an organization can eventually increase sales to those customers and thus increase profits. CRM has burgeoned in popularity as the Web has grown more important in the general economy. Even though the reader may associate, at first glance, CRM solely with commercial activities, many of the essential concepts can be applied to transit, as so aptly demonstrated by UTA and NJT (in Section 6 above).

While the thrust of providing travel information to the transit customers is slightly different from that of a sales force trying to maximize profits, there are significant parallels. Like the commercial salesman, the transit information provider wants to provide additional value to current and potential customers to increase market share, encourage additional product use (ridership), improve product image, etc. Thus, while major commercial CRM products are not designed for the niche of transit agency Web sites, the concept of CRM clearly has applicability to this market.

Indeed, not only do the UTA and NJT Web sites show how the CRM umbrella can be used in the transit industry, many of the other transit Web site features reviewed by the project team have elements of CRM. A good example is the subscription page for KC Metro's *Transit Alerts!* e-mail notification. Figure 7.1 shows (part of) the notification preferences a customer can specify. This preference information is stored in a database, so that when an incident occurs, the agency can send each customer specific information according to the preferences specified. This is a very basic but effective component of CRM.

Figure 7.1: King County Metro Rider Alerts

King County Alert!
Bus alerts

METRO *online home*

Get bus service disruptions via e-mail or text-page.*

Enter your e-mail or text-page address(es) below.:**

Please review all instructions, disclaimers and the [FAQ](#) before subscribing.

E-mail/Pager No. 1
(required):

E-mail/Pager No. 2
(optional):

** Due to limited message space, text-page recipients will receive a summary 'headline,' and, possibly, the beginning of the regular e-mail message. For complete details, it will be necessary to check e-mail, log onto Metro OnLine, or call 206-553-3000/TTY 206-684-1739*

*** Please enter addresses carefully. Most unsuccessful e-mail transmittals are due to addresses entered incorrectly by subscribers.*

Choose a Transit Alert that meets your needs:

If you do not receive a confirmation to your subscription, you will also not receive Transit Alert! messages. If you do not receive a confirmation within 24 hours, please e-mail TransitAlert@metrokc.gov

General Alerts with no route-specific information.
If you want information about changes in multiple bus routes but prefer to limit the number of e-mails you receive, this is your best option. You'll be alerted to service disruptions, and advised to check [Transit Alert!](#) for the most current information posted on the Internet.

Lite Alerts regarding specific bus routes.
If you want information about changes in a specific bus route and have reliable web access, this is for you. The messages you receive will contain few details and a link to [Transit Alert!](#) for the latest details posted on the Internet. A word of warning: The more routes you select, the more mail you'll receive.

Detailed Alerts regarding specific bus routes.
If you prefer to receive detailed, route-specific e-mail messages, choose this. Each message you receive will

CRM in the transit industry context involves providing accurate and appropriately current travel information to transit customers based on knowledge accumulated about that customer via a Web site. If the customer is encouraged to open an “account” or “subscribe” to CRM on an AIP Web site (and has accepted a privacy policy that allows their data to be used on their behalf), it is technologically simple to record that person’s itinerary planning requests. As suggested by the UTA and NJT telephone interviews, the cost of setting up a CRM data system is not prohibitively high. Theoretically then, a transit agency could tie that information to other data about customer travel habits, needs, or requirements that might be collected elsewhere in the Web site, such as interest in a transit pass program, driving directions to a Park & Ride lot, information about accessibility, etc. All of this information supports the development of a customer profile.

Certainly, taken taken together, this information is valuable for a host of planning activities. For example, from an AIP system, one could determine both desired origin and destination pairs and well as those for which an acceptable transit itinerary could not be calculated – a measure of unsatisfied demand. This is similar to how UTA reports using some of its CRM data.

The challenge in adapting CRM to transit, however, is to utilize this profile information at the individual customer level. At its most basic level of CRM, a transit agency could allow the its customers to customize their view of the agency’s Web site, which would provide easier and personalized access to a customer’s preferred features. A transit CRM system might also provide alert messages about transit status, traffic incidents, or promotional programs that might affect a transit customer, such as known sources of delays and alternative routes, fare pass or other programs that might save them money, etc.