Internationalisation of ATG

Paul Dyson, e2x limited

© e2x limited 2008. All Rights Reserved.

Summary

International commerce is now a common requirement for businesses of all sizes, particularly for online sales. Developing a single e-commerce system capable of delivering content and services to a number of different countries in a number of different languages has many business advantages including efficiencies in development, management, and maintenance; reduced time-to-market for new regions; low cost of entry for emerging markets; and so on.

Unfortunately, 'internationalisation' of an e-commerce systems is sometimes trivialised as being solved with some content translation and handling of different currencies and date formats. The reality is that serving customers from a variety of different cultures who live in vastly different economic and technical climates requires a lot more than just displaying the date with the day and the month the right way round. This whitepaper outlines some of the challenges of developing a 'truly international' e-commerce system and how ATG can be used to meet those challenges.

© e2x limited 2008. All Rights Reserved.

Contents

© e2x limited 2008. All Rights Reserved.

Introduction

Once the domain of large multinational corporations, international commerce is now a common requirement for much smaller businesses, particularly for online sales. The rise of a global economy and wide Internet availability make it imperative for all but the most local of sites to handle traffic from multiple countries, usually in a number of different languages. Business success often depends on serving the needs of customers in all these countries, in their native language.

ATG has been successfully used by a number of companies to provide fully internationalised e-commerce systems. e2x has been involved in delivering internationalised ATG solutions since 1999, ranging from systems that support a handful of countries and languages with identical product sets and functionality, to full 'international platforms' that support global operations and offer very different products and services to each country or region.

Internationalisation and localisation

Internationalisation of a system is the architecture, design, and implementation of mechanisms that allow the system to serve more than one 'locale.' A locale is an indication of the location of the user and the language he uses. For example:

- **United States/English**. The user is in the US and speaks (American) English.
- **California/Spanish**. The user is in California (and, hence in the US) and speaks (Mexican-derived) Spanish.
- **European Union/English**. The user is somewhere in the EU and speaks (‘International’) English.
- **Switzerland/German**. The user is in Switzerland and speaks (Swiss) German.
- **Switzerland/French**. The user is in Switzerland and speaks (Swiss) French.

In an internationalised system all of these users will be accessing the same codebase and, potentially, the same servers and database, but seeing very different content and possibly accessing different services.

Localisation is the process of taking an internationalised system and populating it with the correct content and configuration to deliver to each specific locale. Sometimes the localisation of a system will need to consider just the location (i.e. region or country), for example when determining business rules for taxation or shipping. Sometimes localisation will need to consider just language, such as the use of 'international' English to provide an English language variant for non-English speaking locations. Most commonly,

© e2x limited 2008. All Rights Reserved.

however, localisation will need to consider the combination of location and language: the locale.

How 'international' do you need to be?

Internationalisation is more than just populating resource bundles and introducing formatters for currency and dates. In formulating requirements for your internationalised system you need to evaluate the degree of internationalisation the system must support. The areas to be considered are:

- **Basic Content**. Does your basic content (the words, pictures, and digital artefacts) available on the site have to be different for different locales? E.g.: a multi-lingual website that provides one set of content and functionality but translated to several different languages.
- **Branding and Layout**. Does the system need to be presented completely differently for different locations? E.g. a company that wishes to be sensitive to the local customs of Asian countries and so uses different colours in the site design and supports vertically-written languages; or a company that is known by a totally different brand in one country vs. others.
- **Catalogue Content**. Do you need to sell a different set of products in different catalogue structures in different locations? E.g. a company that has localised catalogues for the different markets in Asian and European countries.
- **Marketing and Promotions**. Do you need to market and promote products differently in different locations? E.g. a company that runs different marketing campaigns tied to local events in the US vs. Europe.
- **User Profile**. Do you need to store different information about people depending on their location? E.g. a company that sells to the whole of the Americas needs to cater to the differences in addresses and personal details in North America vs. South America.
- **Fulfilment**. Do you need to calculate tax, shipping, etc. differently for different locations? Are different fulfilment providers used in different locales? E.g. a company that sells products in different European countries (each with their own tax regime) and uses local courier services to deliver goods in each country.
- **Other functionality and rules**. Do you need to vary the services you offer by locale? E.g. a customer support department that has different help-line numbers for different languages. Do you need to apply different validation or constraint rules per location? E.g. a company that provides access to download software but only for customers within the US.
- **Governance and compliance**. Do you need to adhere to different laws and compliance standards in different countries? Are these differences

© e2x limited 2008. All Rights Reserved.

something that can be handled by one set servers or do you need to physically separate out your deployments for different locations?

How 'local' do you need to be?

Once you have an idea about how 'international' you need to be, you need to think about how you are going to produce the localised versions of your site. There are two basic organisational models for this:

1. The merchant maintains separate organisations for each locale, each working autonomously in preparing localised content and services.
2. There is a global organisation that produces localised content and services on behalf of each locale.

This is obviously a gross simplification and most organisations will operate some hybrid of the two. But the model you lean towards will affect your localisation process. For example, having separate local organisations tends to simplify the process of producing local content but complicate the internationalisation of marketing and promotions; the system and the processes need to support different campaigns being launched at different times with no central coordination of these activities. Equally, having a single global organisation complicates the process of creating and verifying local translations of content and producing 'locally sensitive' branding and layout, but makes managing lots of different campaigns much simpler.

The big decision: deployment architecture

The first and biggest decision that needs to be made is what deployment architecture you want for your system. This decision needs to be guided as much by the organisational considerations and internationalisation/localisation requirements as by technical considerations.

The choice to be made is whether you are going to build a single 'global platform' system capable of running all local variants from a single cluster of servers managed centrally, or whether you want an 'independent whitelabel' solution where each locale runs and manages its own cluster of servers.

© e2x limited 2008. All Rights Reserved.

The global platform solution

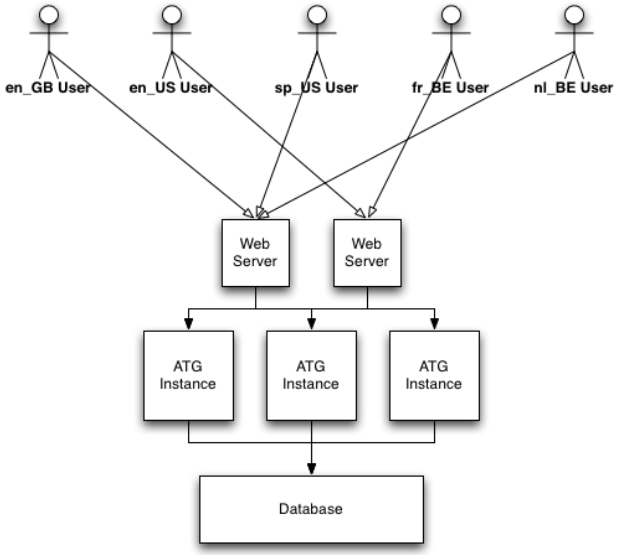

A global platform solution carries out all localisation dynamically. The internationalisation of the system is such that decisions about what content to serve, what look and feel the user sees, what products they browse, and so on are made on a per-request basis. A user viewing the Belgian help content in French may be using exactly the same server as a user purchasing a product from the UK catalogue in English. Effectively there is 'one system' serving all possible combinations of location and language.

The independent whitelabel solution

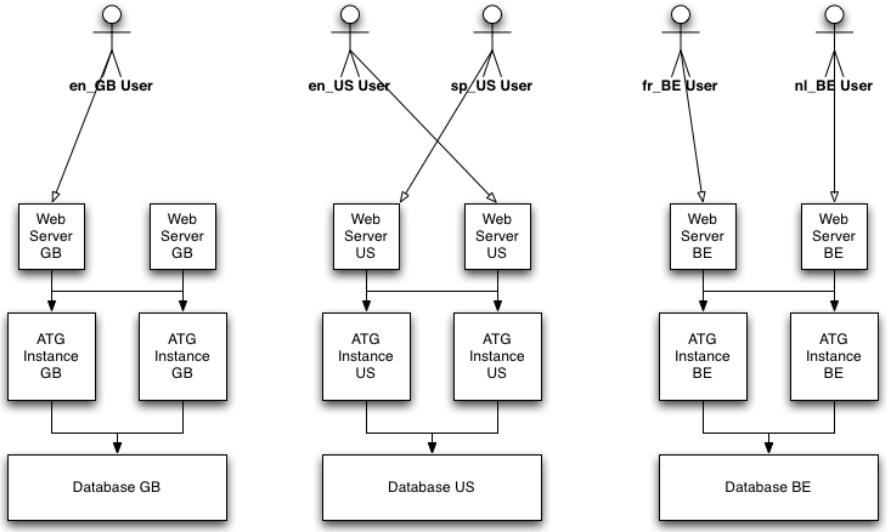

© e2x limited 2008. All Rights Reserved.

An independent whitelabel solution is localised at configuration and data-population time. The internationalisation of the system is such that it is capable of serving any locale's content and catalogue but, once populated, is only capable of serving one location at time, possibly with multiple language translations. A user viewing the Belgium help content in French will be served by the Belgian server and a user purchasing a product from the UK catalogue will be served by the UK server.

Which is the best solution?

Both have advantages and disadvantages. The global platform solution is the easiest to manage and maintain and is the most flexible. In a platform solution all the code and content are in one place and only one cluster of servers needs to be managed. Introducing a new locale is simply a matter of deploying all the relevant content and configuration (on one platform system we built, deploying a totally new country took less than an hour) and there are many options for 'tuning' the system from configuring it to favour certain high-value locales to creating new clusters dedicated to serving particular regions in an independent whitelabel-like manner (e.g. a European platform, a US platform, and an Asian platform). Platform systems also enable a user to register once and then access all locations without the need for custom 'single sign-on' integration.

The cost of this flexibility and ease of maintenance is implementation complexity. Although ATG lends itself well to building platform systems, they are still inherently harder to build than whitelabel systems.

Whitelabel systems simply have to be architected and implemented to avoid restrictions on the types of content, catalogue, and other aspects that they support. They may also be easier to integrate to local systems or customise for local requirements. They will probably need to be multi-lingual but, in all other aspects, they are like any other non-internationalised e-commerce system.

In addition to the implementation and organisational considerations, there are potentially legal issues that will guide whether a platform or whitelabel solution is the most appropriate. Some countries have strict laws governing where 'sensitive' data (most likely user profile and order payment information in an e-commerce system) can be stored.

Internationalisation of ATG

Basic design approach

Whether you're building a platform or whitelabel system, the question, “How might this be different in different locales?” needs to be asked about every requirement to be implemented and every piece of infrastructure to be built. Even something as simple as a JSP needs to be written in a way that ensures

© e2x limited 2008. All Rights Reserved.

it will display correctly in all intended languages and with potentially different data sets populating it.

Although the internationalisation question needs to be asked constantly, it is entirely possible to design the system such that, in most cases, the developers can safely ignore the possibility of local variation as long as some basic implementation rules are followed. The simplest example of this is the use of content identifiers: as long as developers know they must not hard code text on a JSP but instead use a content identifier and a content look-up component, they don't need to know how different pieces of content are translated or in which language the user is currently viewing the system.

ATG offers many 'hooks' into its product suite for adding components to promote this 'transparent' approach to internationalisation, including the servlet pipeline, the GSA repository mechanism, Nucleus, etc. The following sections give some high-level guidance about how to use these hooks to internationalise your system.

Locale identification

In order to provide the user with the correct localised content and services, you need to know what locale she is in (or acts as though she is in). Don't be distracted by systems that claim to identify a user's location by the IP address from which the HTTP request originates. These are far from 100% accurate and they don't account for the case where the user is travelling and wishes to access her 'home' locale. A far simpler and more effective approach is simply to let the user make the decision.

The first part of that decision process has to be the domain name used to access the site. If a user navigates to myorg.nl then he will expect to see the Netherlands site in the Dutch language. myorg.com is a slightly more difficult prospect, as it is not clear whether the user is specifically looking for the US site or is using .com as the default 'global' domain. In a whitelabel system this is easy: each domain's DNS entry points to the server cluster for the location. The system should present the default language for that location. In a platform system a component needs to translate the requested DNS to a supported locale, usually placed into the servlet pipeline.

Users will expect to be able to change their locale while using the system. In a platform system all valid locales can be presented to the user to make a choice. Every time a new locale is deployed it will automatically be added to the list. In a whitelabel system there either needs to be a separate 'global' server that allows the user to choose his locale, or else every local cluster needs to know about all other valid locales. The latter approach obviously makes the process of deploying a new locale much more complex.

In both platform and whitelabel systems changing the language but not the location should not affect the user's session information. She should just see the system in a new language. But what if she changes the location as well? In a platform system there is the possibility to retain all session information

© e2x limited 2008. All Rights Reserved.

(because the user is still on the same server and working in the same session) but this might not be wholly appropriate. For example the user may have placed a product in her shopping cart whilst in a German locale and then switched to an Austrian locale that does not sell the same product. In whitelabel systems migrating session information is much more difficult.

Once the user has selected or re-selected their locale through the URL or user choice it should be stored against his Profile. For registered users it means that, having selected the locale that is meaningful to them, they will be returned to that locale unless they use a different URL or change their selected locale. Having chosen French as the language in which he wishes to view the site, our Belgian user will expect to see the site in French the next time he visits.

Internationalised content

'Content' refers to any text, images, documents, or other assets that the user sees or downloads. For both platform and whitelabel systems the anticipation has to be that every locale will have a completely different set of content.

In ATG systems content tends to be split according to type:

- Display elements - welcome messages, button values, copyright messages, etc. - tend to be placed in resource bundles and media repositories
- Catalogue content (product and SKU descriptions, category names, etc.) lives in the ProductCatalog repository
- Promotional and marketing content delivered by scenarios or targeters tends to be in ATG repositories
- Almost every ATG implementation has one or more custom repositories, some of which may well contain user-visible content

For most e-commerce systems the bulk of the content is going to be in the ProductCatalog repository. The definitions of all repository items in the catalogue and other repositories need to be internationalised so that property values are keyed off the user's locale. Hooking into the GSA repository mechanism can make locale-based look-up transparent to developers. For example, using a custom GSA property descriptor means that developers can still refer to the displayName property of the product item type. The repository will automatically return the correct displayName for the user’s current locale.

Standard resource bundles can be used and even administered with ATG Content Administration (CA) but we recommend placing resource bundles in an SQL repository as this makes administration tasks, such as querying whether all resource bundles have been localised for a particular locale, much simpler. Deployment is also easier, as there is no need to ensure the class loader reloads the resource bundle file; the standard CA deployment is sufficient.

© e2x limited 2008. All Rights Reserved.

The internationalisation of content is the same for both platform and whitelabel systems, as even a whitelabel system should be capable of delivering more than one language for a single location.

Locale-derived templates

Rendering internationalised content in a JSP can be sufficient if all of your different content fits well within a single template and all the locales you support are happy with the same UI look and feel. However, you may well want to support both Western and Asian locales, which will require different UI conventions. Many Asian pictographic languages are traditionally written vertically, although most Asian users are comfortable with reading these horizontally from right to left. Arabic is written right-to-left with 'anglicised words' and all numbers written left-to-right. More significant will be use of colours and symbology. Asian cultures put significant emphasis on the luckiness (or otherwise) of certain colours and images.

Selecting templates and colour schemes dynamically is really only necessary for platform systems. In such systems, the CSS and JSP fragments to be displayed need to be derived from the locale. For whitelabel systems a new set of JSPs could be produced for each location the system is deployed to if necessary.

Location-specific catalogues

If you need to internationalise your product catalogue on a platform system such that every location has a different set of products (and, potentially, a different category hierarchy), ATG provides the perfect mechanism out of the box: custom catalogues. A single 'global' product catalogue can be defined that contains all products available to all locations and then a custom catalogue is defined for each individual location with its own category hierarchy and product set. The locale selected by the user is used to identify the correct custom catalogue for her to view.

Partitioning scenarios and promotions

Scenarios are a fundamental ATG mechanism and a powerful part of any e-commerce system. There are three ways in which scenarios tend to be internationalised:

1. **'Global' scenarios delivering local content.** In this case the scenario is the same for all locations but the content delivered by the scenario is selected based on the user's locale.
2. **'Global' scenarios with local variations.** In this case there is one scenario with different branches for different locations.
3. **Local scenarios.** In this case there is a scenario for a particular location and there may be no equivalent in any other location.

© e2x limited 2008. All Rights Reserved.

The ATG scenario mechanism handles all of these out of the box without any need for enhancement. The main consideration is that scenarios must be written to be locale aware in both their flow and the content they deliver.

ATG promotions and vouchers can also work out-of-the-box in situations where the same promotion or set of vouchers is going to be applied globally. If you want to handle location-specific promotions or vouchers, the out-of-the-box mechanism will need to be extended.

The 'super profile'

The user profile is another key out-of-the-box ATG component. Internationalising the profile is really a case of ensuring that the user item itself and all associated items (such as addresses) are capable of holding the variant information that different locales will need. For example, in Latin American countries 'Barrio' is an important part of an address that has no equivalent in most other societies.

Another important example is honorifics: the 'title' by which we refer to people. In western societies a simple Mr/Mrs/Miss/Ms choice will often do and personalised messages usually refer to the user by his first name ('Welcome Paul') so the choice of honorific is optional. In India there is a plethora of honorifics and one person may have several. In Japan the use of honorifics is mandatory and highly important in terms of showing an appropriate level of respect.

Rules about the usage of profile properties and the differences in what is mandatory or not between different locales can be handled through UI templating and internationalised services (see below), but the profile itself needs to be a 'super profile' capable of storing user information appropriate to any locale the system is likely to be localised to.

In platform systems, one of the key advantages of a super-profile is that the user can access any location without the need for custom 'single sign-on' integration. He can update his profile with location-specific information which will only be shown in the relevant locales. In whitelabel systems the main purpose for developing a 'super profile' is to reduce the amount of location specific development needed for locations that need to extend the 'core' profile.

ATG engines and custom services

ATG provides a number of 'engines' – for pricing and fulfilment for example – that can be used as-is or customised to a particular set of business rules. In a platform system there potentially needs to be one version of each engine for each of the locations the system supports. Whether this is true or not depends on your business; many businesses have a global set of pricing rules (with local variations in unit price and tax calculation) but different rules about fulfilmentfulfillment, due to having different fulfilment partners in each location.

© e2x limited 2008. All Rights Reserved.

In order to maximise the use of the out-of-the-box code and minimise the impact of internationalisation on the developer, we recommend configuring as many different individual instances of each engine as you need and using ATG ServiceMaps to map between the location or locale and the correct engine for that locale. This approach also works for custom-developed services that implement business rules, validation, and so on specific to your business.

'Opt out' of locale-specific functionality

The most advanced internationalisation of a platform system is one that allows specific locations to opt out of or turn off certain functionality. For example, a service such as a dealer locator may only be possible in countries for which there is an established dealer network and all the relevant data is available for use in the system. In the US and Europe there may well be sufficient data for the dealer locator service to be integrated with a mapping service allowing users to easily find their closest local dealer. In Africa, however, this level of service may simply not be possible and so the dealer locator service itself should not be available but, perhaps, just a static page of content about how to find a dealer.

Turning off or opting out of a locale-specific service requires more than just not implementing that service. All links to the service need to be removed and the user needs to be actively prevented from accessing it. Implementing a 'site map repository' that contains a content-managed description of pages and features on offer allows total control over the services offered to users by location.

Localisation with ATG

Localising an internationalised system requires two things: localisation of the configuration by the development or operations teams and population of all the localised content by editors, marketers, sales, etc.

Configuration

The Nucleus component mechanism makes configuring an internationalised system extremely simple. For a whitelabel system the locale-specific services can be configured using Nucleus's standard mechanisms for overriding the default 'global' configuration: replace all the internationalised settings with localised variants in a location-specific configuration layer.

For a platform system Nucleus is the primary mechanism for configuring the mapping between locations and localised services.

Content Management

ATG Content Administration and ATG Merchandising provide excellent out-of-the-box support for localising content. Merchandising is used to create and manage new custom catalogues and CA can handle both file-based and

© e2x limited 2008. All Rights Reserved.

database-stored content. The combination of CA and Merchandising allows you to manage all locale-specific variants of content as well as to write internationalised scenarios and promotions and to set locale-specific opt-outs of functionality.

Conclusion

Internationalisation of a system is all-too-commonly trivialised as being solved with some content translation and handling of different currencies and date formats. If you need a system that supports many different countries and regions across the globe, serving customers from incredibly diverse cultures, for an organisation operating in wide variety of financial and technological climates, you have a difficult problem to solve.

ATG provides good out-of-the-box support for building and deploying truly internationalised systems as well as, perhaps more importantly, excellent support for extending its features to cater for internationalisation and localisation. Today there are a number of ATG customers that run multi-national or global systems that address all the kinds of issues described above.

However, before embarking on the internationalisation and localisation of an ATG system there is a key decision any ATG customer must make: do you want to organise your management of the system and its content globally or locally? This will have a fundamental effect on the high-level architecture of the system and the way in which you approach the implementation of your internationalisation/localisation mechanism.

About the author

Paul Dyson has been delivering internationalised ATG systems since 1999 and e-commerce systems in general since 1996. Paul is the author of a book on architecting e-commerce systems - Architecting Enterprise Solutions: Patterns for High-Capability Internet-based Systems (Wiley & Sons, 2004) – and is a partner at e2x limited.

paul@e2x.co.uk

www.e2x.co.uk

© e2x limited 2008. All Rights Reserved.