Design of an enablement process for on demand applications

by K. Chang
A. Dasari
H. Madduri
A. Mendoza
J. Mims

In today's business and economic conditions, enterprise customers are demanding information technology (IT) solutions that are cheaper, less complex, and easier to install. At the same time, independent software vendors (ISVs) are seeing revenue from their core licensed offerings erode because of competition and market saturation. Many believe that the answer to these problems is to offer IT solutions by means of utility computing. Like an electric utility, software applications can be offered as on demand services, and customers pay only for what they use. Creating and implementing such utilities is by no means trivial. It requires some expert help and adherence to established standards and guidelines. In this paper we describe the design of such a process that we call the Application Enablement Program. The process helps ISVs transform their applications into on demand services. This process is structured, repeatable, and globally deployable.

Independent software vendors (ISVs) are seeing erosion of revenues because of competition and pressures from enterprise customers demanding applications that are cheaper, less complex, and easier to install. These ISVs must seek new ways to penetrate new markets and grow their customer base. Today's economic and business conditions are requiring that ISVs begin to offer their products as utilities. As utilities, they would be application-level services that are sold on a "pay-as-you-go" basis. To emphasize the on demand characteristics of these services, we call them "on demand services" or "ODSs" in this paper. Unfortunately, ISVs who traditionally sell their applications as licensed products are inundated by information about utility computing and the on demand "hype." They all ask the same question: How do I turn my application into an on demand service?

In this paper we describe the design of a process called the Application Enablement Program (AEP) for on demand services (ODSs), which are software utilities, and introduce ODS architectural concepts. On demand service is synonymous with on demand application. The AEP represents a streamlined, cost-efficient, and globally deployable process. Our design is based on our experience creating such a process for hosting licensed software applications in the IBM Universal Server Farm hosting facility, Version 5 (USF V5). The AEP for ODSs is expected to play the crucial role of transforming ISV applications into ODSs, identifying the tasks to enable the transformation, helping them make the changes to run in the on demand infrastructure and the ODS framework, and finally checking for ODS compliance with the utility platform standards and guidelines.

Background

Before we proceed with describing the AEP, we define some terms and concepts used in the process such as hosting models, utility computing, on demand

0018-8670/04/$5.00 © 2004 IBM

infrastructure, and the Universal Management Infrastructure to give the reader an idea of where ODSs fit.

Nonutility hosting. Traditional licensed software applications were typically purchased by enterprise customers and installed on their own premises. As software applications became more complex to manage, enterprise customers began moving these applications out of their premises and hired hosting companies such as IBM to host and manage these applications for them.

The majority of currently hosted applications fall into three (nonutility) hosting models.[1] Utility computing is discussed in the next major subsection.

Collocation. The application provider (the customer) collocates the machines running the application with a network bandwidth provider. The application provider is responsible for managing the application server machines remotely or on site. The provider integrates the offering into the existing infrastructure. The network bandwidth provider typically provides the application basic services such as electric power (as an uninterruptible power supply), network connectivity, and bandwidth (power, ping, and pipe).

Dedicated. An infrastructure service provider (InSP) supplies dedicated machines to run a fixed set of applications. The machines may be purchased outright or leased by the application provider. The InSP manages both the content and infrastructure. The infrastructure includes servers, storage, and network connectivity. The dedicated model is a simplified description of USF V5.

Both models typically serve one enterprise customer per single instance of the installation. They require application providers to procure, install, and deploy a new instance of the installation for each new company buying the service. The application usage is also predictable and linear. Both models have severe implementation and cost challenges when demand slows or grows rapidly in an unpredictable fashion.

Enterprise customers pay to maintain their respective resources. They can typically use as many resources as they need, limited only by what their infrastructure capacity permits.

Shared hosting or application model. The InSP provides the bandwidth, servers, storage, and software services for supporting a shared application. The application is shared among multiple enterprise customers with multiple users. The servers are typically owned or rented by the application providers.

This model leverages shared services across multiple customers and applications. The model has a lower incremental cost for each additional customer joining the service. Because the application is shared, the application provider has to ensure that the application scales well when additional customers are added.

Enterprise customers for these shared applications are typically billed on a monthly basis. The use of the services is not metered for usage, leaving the users to consume as much as they want for a fixed monthly fee. Service level agreements (SLAs) may exist, but limits on usage could be difficult to enforce.

Several factors are influencing the move away from these models. Server utilization, application usage, administration complexity, return on investment, total cost of ownership, and economic uncertainty are some reasons that are driving application vendors and InSPs to rethink the overall strategy for business hosting. This is where utility computing can help.

Utility computing and utility hosting. Utility computing in its simplest form is information technology (IT) presented as a utility. It is based on the model of conventional utilities such as telephone service and electricity. The enterprise taps into this utility for business and infrastructure services. The services are physically hosted in utility data centers that can be located either inside or outside the enterprise, or both. Like a traditional utility, the service is pay-as-you-go, charging only for use of the service. It is reliable and caters to the on demand nature of a utility.

Utility computing promises the following benefits:[2]

- Simplify IT by reducing complexity
- Turn IT from a fixed to a variable cost
- Reduce cost or operating expense

Current technological advances have turned the concept of IT delivered as a utility into reality. Some examples:

- Reduced cost of bandwidth enables the creation of new data services and high-speed network delivery of a variety of services to a broader range of customers.

- Distributed content and application architecture deployments shift delivery to the edge of the network.
- Server and storage virtualization enables new levels of shared infrastructures with the potential of reducing customer costs.

In utility-based hosting, application providers pay for the "IT infrastructure power" that they need to run their applications. The application providers pay only for what they use. The utility-hosting provider meters application usage metrics such as bandwidth, storage, and CPU use and bills the application owner accordingly. The utility-hosting provider owns all components of the infrastructure, including servers, storage, and network. This situation is in contrast to nonutility-hosting centers where application providers buy or lease IT infrastructure components such as servers and storage to provide themselves with IT infrastructure power. In a nonutility hosting model, the application provider is more like a consumer buying his or her own power generator instead of signing up with an electric utility provider.

On demand infrastructure defined. An on demand infrastructure seeks to fulfill the IT functionality of today's enterprise business in a utility-like fashion. The on demand infrastructure delivers standardized processes, application function, and infrastructure over the network as a service.

The on demand infrastructure shares the following properties in common with utilities:

- Sharable—capable of serving many customers
- Standardized—requires (allows) little customization
- Flexible and scalable—use what you need and pay as you go

From a technical perspective, an on demand infrastructure has the following characteristics:

- Integrated—allows the integration of enterprise and legacy applications that transcends vertical industries through the use of open standards technology such as Web services.[3]
- Open—uses open specifications and standards to enable ease of integration of enterprise resources and applications.
- Virtual—uses server consolidation and capacity on demand technology to increase utilization of hardware resources.
- Autonomic—alleviates the need for skilled technical human resources to manage the complexity brought about by the rapid advance of technology. This characteristic also includes responding to customer needs for instant provisioning of resources.

The UMI platform. IBM Global Services (IGS) implements an on demand infrastructure by taking advantage of and integrating existing IBM hosting centers with the Universal Management Infrastructure (UMI). UMI is the infrastructure of tools and software that provides services to manage, meter, and provision resources at the hosting infrastructure level. The initial release of UMI concentrates on infrastructure provisioning and managing the provisioned resources. The current version of UMI consists of the following functions: auto-provisioning, monitoring, reporting, metering, billing, and SLA.

Shifting traditional IT infrastructure that is static or planned for peak demand to an on demand infrastructure clearly necessitates the use of an implementation such as UMI to manage the infrastructure. Although UMI addresses this shift, it does not offer anything that ISV application providers need so that their applications can be offered as utilities or services. This leads us to the definition of on demand services.

From here on, on demand infrastructure and UMI will be used interchangeably.

On demand service defined. An ODS is a utility. For the purpose of this paper, we define an ODS to mean a software application with on demand characteristics hosted in a utility environment. The business and technical characteristics of on demand systems are described in Reference 4.

An ODS is not the same as a software application offered as a service by traditional application service providers (ASPs). ASP is the name given to companies that delivered standard packaged application software to an organization on a rental basis during the dotcom boom.

There is one fundamental difference between an ODS and the traditional ASP software offering. In an ASP model the customer rents the *software* and pays a recurring subscription fee to use the application as much as he or she wants, whereas in the ODS model the customer pays for the *service* and does so based on usage (i.e., pay-as-you-go).[5]

Figure 1 The ODS and UMI architecture

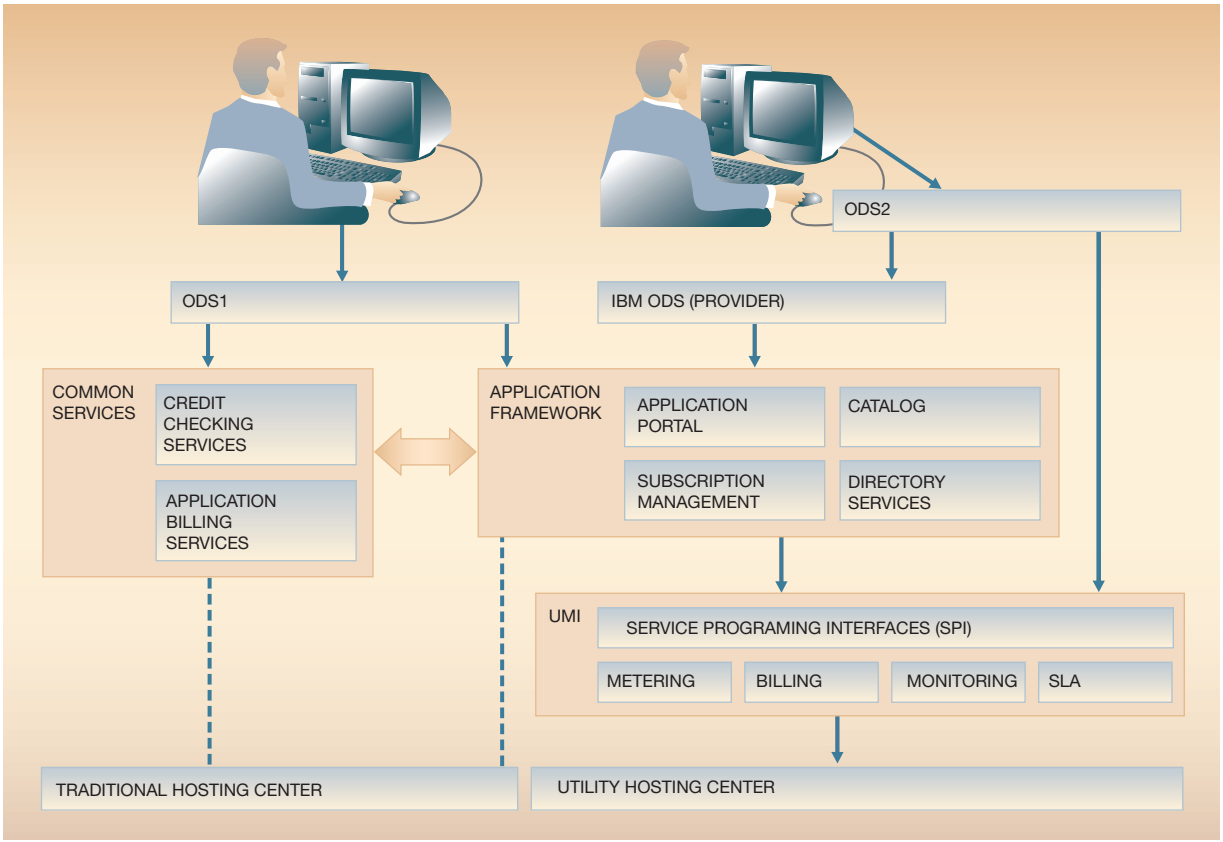

ODS architectural concepts. The AEP process, discussed in detail later, helps to transform the traditional ISV application into an ODS. Transforming an ISV application into an ODS means making the application adhere to the ODS architecture.

An ODS architecture describes how the application interacts with and uses other applications (i.e., other ODSs and common services provided by IBM) as well as functions provided by the on demand infrastructure. The architecture includes the ODSs (i.e., transformed ISV applications), and the underlying hardware platform, which supports virtualized technologies and is managed by the UMI.

ISV applications are integrated into the architecture by using integration points provided by each external service that the application wishes to use. The integration points use open standards technology such as Web services.[3] This use allows an ISV application to easily move to other on demand infrastructure type implementations, thus avoiding lock-in.

Integration into the architecture is not the only requirement for transforming an ISV application to an ODS. Compliance of the application with certain attributes that are explained later in the AEP process section is also required.

Figure 1 illustrates the hierarchical nature of the UMI and ODS architecture, showing the relationships between components in the architecture. The arrows point to the component being used by the component from which the arrows originate. The arrows also represent the integration points between the components.

A hierarchy of ODSs. On demand services form a hierarchy of relationships and interactions. An ODS

may provide application-to-user functions directly to enterprise customers through a Web interface such as a browser or wireless Web device (ODS1 and ODS2 in Figure 1). Other ODSs may provide application-to-application services (be a "provider" ODS) through Web services (IBM ODS in Figure 1).

Common services and the application framework. A recent study of internally developed applications and outsourced applications favored the outsourced applications with respect to development time, maintenance, and cost.[6] Providing common services as well as an application framework can add value to the architecture by making available ready-made functionality that ISV applications can use.

Common services are ODSs provided by IBM that are available for use by other ODSs. These common services are just like any other ODS, except that they are owned and managed by IBM (Figure 1). Some common services can exist in the same layer as the ODSs. The goal for making common services available within the ODS architecture is to provide ready-made functionality that ISV applications can incorporate to become an ODS. An ODS developer could choose to use a common service rather than writing the functions into its application.

An example of a common service is application billing (as opposed to infrastructure billing). An application billing service essentially bills the enterprise customers of an ODS. The ODS owner may choose to subscribe to this service rather than building his or her own billing system. Another common service might be a credit check to qualify new customers. The added value of these common services is to provide the ODS owners with functionality that they may otherwise have to develop themselves.

An application framework will exist to provide non-business functions such as directory services, subscriber management, single sign-on, and other functionality still being defined as of this writing. Like common services, the application framework provides ready-made functionality that an ISV application can use instead of developing its own similar functionality.

Service Programming Interface. IBM will expose certain UMI functions through published interfaces known as Service Programming Interfaces (SPIs). An SPI sits directly on top of the UMI layer. It provides the interface for the ODS to communicate with UMI services such as monitoring, resource metering, auto-provisioning, and more. The SPIs are exposed as Web services that an ODS may invoke. The SPIs also isolate the ODS applications from changes in the underlying infrastructure.

Early on, we expect ODSs to use the SPI to take advantage of the UMI metering service. As ODS owners become more familiar with other UMI services we expect use of SPI to grow even more.

The hosting platform. The hosting infrastructures that support the UMI and ODS architectures may be located in various hosting centers to meet geographical needs. The servers, networks, and physical plants may be drawn from the existing traditional hosting centers (USF V5) and other resource pools to form the utility hosting centers.

The ODS development model. From an overall system perspective, both business and technical domains are affected during the transformation of an application from a traditional hosted offering to an ODS offering.

In the business domain, factors such as cost, return on investment, and time to market have to be evaluated by the ISV. A new business model that creates new market opportunities and generates revenue through usage and resource-based subscription has to be created and refined. In the technical domain, a gap analysis of the existing architecture with the ODS architecture will reveal the work effort and the skills required to transform existing applications. The work effort has to be sized and fed back to the business model.

The IBM AEP process mainly deals with the technical domain, but it provides guidance and input to the business model as well. Once a business model has been built and a decision has been made to convert the ISV application into an ODS (i.e., "utilitization"), the AEP-recommended development model should be followed. This model consists of a development process (a sequence of steps) that transforms the ISV application into an ODS, consulting methods to aid the transformation, and compliance procedures for quality assurance. All of these elements are incorporated into the AEP process described in detail in the following section.

The Application Enablement Program

The current hosting process only deals with traditional software applications that ISVs host on IBM tra-

Table 1 Comparison of AEP for USF V5 enablement and for ODS enablement

Traditional USV V5 Enablement	ODS Enablement
Focus is on enabling the ISV application for hosting dedicated instances for each customer. Enterprise customer orders dedicated hosting.	Focus is on ISV application to ODS transformation. Enterprise customer orders the ODS, not dedicated hosting.
Enablement task list is reused for each customer boarding.	Enablement is a one-time task to transform and certify the ODS.
Enterprise customer licenses ISV software.	ISV sells software as a service.
Enablement is performed during the boarding process (for each customer).	Similar tasks may be performed during the ODS boarding process, but the focus of this enablement is different. The focus is to transform and certify for compliance.
Enablement is focused on hosting the application; that is, monitoring, backup, remote application management, etc.	Enablement certifies that the application complies with utility architecture and standards. Note that monitoring, backup, etc., are still important and are done for the ODS as a whole.
Consulting is not a part of the classic AEP process.	Consulting is a significant part of the program—ISVs need assistance to convert applications to ODSs. It is offered as an option.
Assessment identifies effort hours to enable application for hosting.	Assessment identifies effort necessary for converting application to ODS.
Enablement report is created for use in subsequent application installations.	No enablement report is created—application is only boarded one time and then scales up and down based on demand. The steps required to board each additional customer onto the ODS are ODS-specific, and they typically do not affect the infrastructure needs.

ditional hosting centers, and does not deal with the fundamental questions that ISVs are asking, such as, how do I turn my application into an ODS? As ISVs begin to implement utility computing services through IBM utility hosting centers, a more structured and flexible business process should evolve from the current one.

IBM's goal of providing a utility computing infrastructure that hosts utility applications (ODSs) can only be fully realized by a well-defined, well-organized approach that does not compromise flexibility nor interfere with a successful boarding (provisioning service for a customer) or hosting process. The critical design goals of the process are to be repeatable and flexible, aiding in the transformation of traditional ISV applications to ODSs.

To help design such a process, we started with our experience of enabling ISV applications for traditional (USF V5) hosting. This AEP process (sometimes referred to as "AEP classic") is based on customer feedback as well as experiences through collaboration with different IBM business organizations and feedback from ISVs. Whereas the classic AEP process is aimed at enabling ISV applications for hosting (where each instance is typically dedicated to a single customer), the new AEP process (ODS enablement) is aimed at streamlining the transformation of a traditional ISV application into an ODS. Deploying a transformed ODS in an IBM utility hosting center is similar to the enablement steps in the classic AEP process. Because we based our design of the new AEP process on AEP classic, the flow of the process, which is divided into phases, will be similar. The phases, which are described shortly, will have their own similarities and differences. The similarities and differences between these two enablement processes are summarized in Table 1.

The AEP process for ODS provides the following:

- It assists ISVs in the small and medium business market place with the transition from dedicated hosting to a one-to-many, multitenant utility hosting environment.
- It will produce an efficient, structured, and repeatable procedure that will result in benefits to ISVs, IBM, and the partnership between them.
- It provides an avenue to explore details of an application through detailed assessments and interviews with the ISV before deploying it in an IBM hosting center. This will be explained in more detail in the next subsection on the AEP process.

- It will provide all the partners involved, that is, the ISV, IBM Strategic Outsourcing, IBM Application Management Services, and IBM Business Consulting Services, with a better understanding of customer needs and knowledge. Better understanding of customer needs and knowledge will in turn help reduce the risk of encountering major problems during the transformation and deployment of the service in a utility hosting environment.

We gained the following insights from the USF V5 enablement program that was developed during 2002:

- Service delivery organizations typically do not become involved with the details of the application operation. Early involvement of the delivery organization in defining the process, debugging it, and having it adopted is critical to the successful deployment of the program.
- Certain large applications cannot be enabled in a test laboratory, and hence, the very first customer boarding in a delivery center will be used to harvest the enablement process details.
- The enablement report will be of value for second and subsequent customers for the above case.

The design of the AEP process for utility enablement is described below. It incorporates the insights gained from the previous USF V5 program and is also designed to be flexible enough to accommodate the content changes in the UMI platform release, the set of available common services, and the variation in the degree of "utilitization" desired by the ISVs.

The AEP development team designed the process to be delivered in several phases: (1) assessment, which includes education in the form of presentations, (2) enablement, which includes consulting and compliance verification, and (3) deployment, which includes actual hosting and boarding on the on demand infrastructure. The AEP process begins with the assessment phase, which consists of analyzing the application architecture, technology, and implementation of the ISV, and leads to a quantifiable categorization of the degree to which they are compatible with what might be called an on demand architecture. Assessment also determines the various tasks required to become an ODS in the IBM utility computing centers (that implement UMI).

The contract signing in Figure 2 signifies that the ISV and IBM have entered into a partnership. As partners, the ISV and IBM will work toward transforming the ISV application into an ODS and host it on the IBM utility computing center.

After assessment and signing of the contract, the ISV application goes through the enablement phase, in which the application is transformed to meet the requirements of becoming an ODS. The ISV either engages IBM through consulting or does the transformation in-house. Once the necessary transformation is done and the ODS is ready to be deployed into the utility computing center, the enablement compliance process is conducted to ensure that the ODS meets all requirements of an ODS and that all the proper steps are taken to efficiently deploy and operate the service in the IBM utility hosting environment.

The AEP process in detail. In the following subsections we describe each step of the AEP process. As an example of its use, we apply the process to a hypothetical ISV called Ideas2Concepts, Inc., which has a time-tracking application that it sells as a licensed product. The stand-alone application enables consulting companies to track time spent by consultants on different projects. The interface is Web-based so that consultants can access the application through the Internet. The application also has a built-in billing system that automatically produces invoices and bills customers in a timely fashion. Over time, the application has become popular within the IT service industry. Ideas2Concepts has added interface points so that customers can integrate the solution into their own human services and accounting applications. In 2002, sales of the application started to decline because the market for such a product became saturated and competitive. Ideas2Concepts did some studies that convinced the ISV that turning its application into a utility would open new markets and increase revenue; therefore, the ISV decided to enter into a partnership with IBM through the AEP program.

Phase 1: The assessment. In AEP classic (for USF V5), the sole purpose of the assessment is to ensure that the application can be hosted at an IBM hosting center. In the AEP for ODS, additional application attributes are required for the application. These attributes distinguish ODSs from traditional licensed applications.

Every ISV application is unique. Each ISV uses different solutions, technical implementations, and business models to deliver a service or an application to address a market demand. The AEP focuses on

Figure 2 AEP process flow for ODS enablement

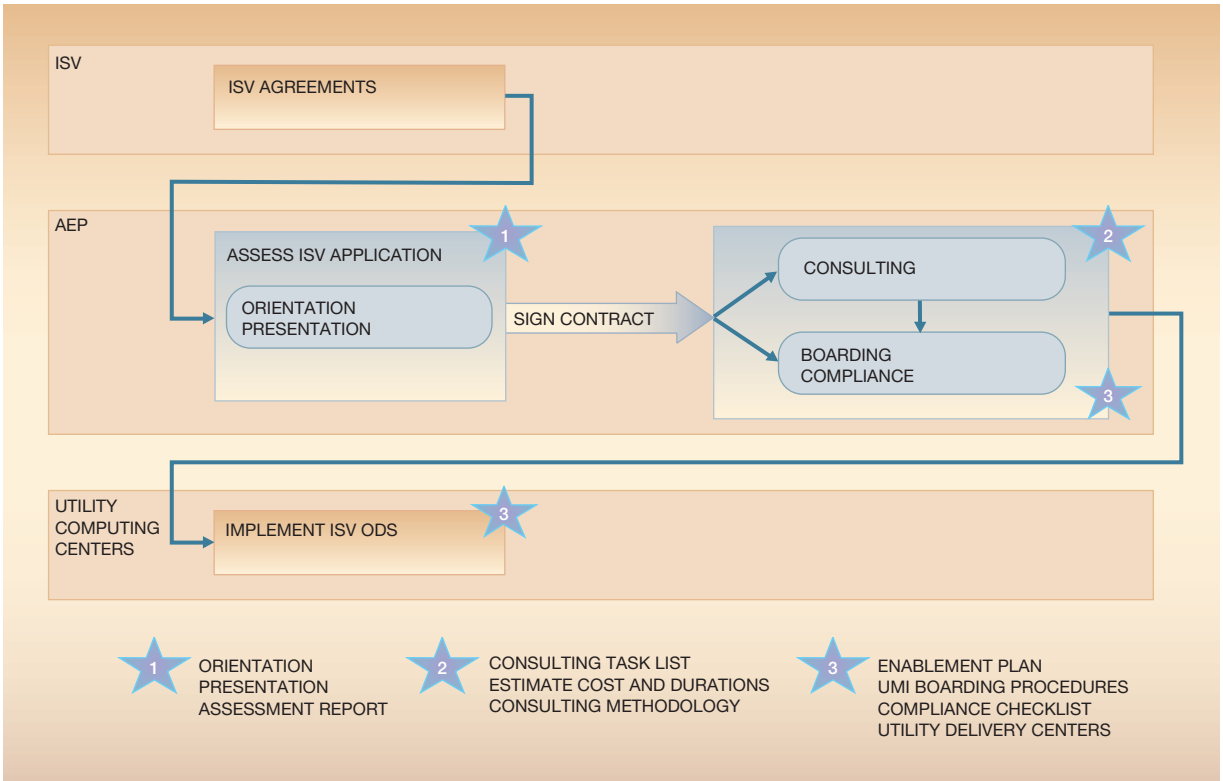

identifying and assessing various attributes relevant to the deployment and hosting of an ODS.

Assessment is necessary to determine the readiness of the application for a new way of doing business and offering a service within the technology provided by the IBM UMI. As mentioned earlier, the AEP services focus on the technical aspects of the ODS. Although the AEP does not focus on the business assessment of an ODS, such as market investigation or sector-specific assessment in terms of deployment based on Web services, those services are available and can be facilitated by IBM Business Consulting Services.

Through assessment, not only does AEP gather information on the technology and architecture supporting the application of the ISV, but it also provides an avenue for IBM to present the underlying technologies of the on demand infrastructure, the ODS architecture, and how the ISV can exploit those technologies. In the USF V5 assessment, the type of education presented was geared more toward the environment of an IBM hosting center.

The assessment produces an assessment report for the ISV. The report includes a task list that shows the ISV what it needs to change, modify, or add to the application to become an ODS.

The assessment does not examine how the application implements business functionality and is not overly concerned with what the application actually does. The assessment is used to discover whether the application can be hosted and how well it conforms to the requirements of the utility computing environment. With regard to the Ideas2Concepts application, the assessment concentrates more on knowing the architecture of the application and how the application uses middleware products and database requirements as opposed to knowing how the application helps solve a customer's time-tracking requirements.

The assessment happens over a series of four or five telephone interviews that last from one to two hours each. The first call orients the ISV with the AEP and process. The second call is the education presentation (discussed in more detail later in the next few subsections). The third and fourth calls are used to learn about the ISV application and architecture. A fifth call is made if necessary to go over the information again and clear up any remaining details.

It is important to note that the AEP process is designed and developed by the AEP development team and will be implemented by different delivery groups in IBM around the world. The current plan is for assessment to be conducted by consultants from the IBM Solution Partnership Centers (SPCs) worldwide.

The orientation and education presentation. The orientation presentation has two major topics. The first presents the hosting environment and describes how an application "lives" in the IBM hosting facility. The second topic presents technology used in the on demand environment, such as the UMI architecture and its exposed services, common services, and application framework, Web services, and other open-standards-based technologies used in the on demand infrastructure.

The orientation presentations are intended to provide the ISV with the basic knowledge and terminology of the IBM utility computing environment as well as to introduce the ISV to the functional services that its application can use when it decides to convert its application into a utility.

The template-driven interviews. The AEP development team has created a questionnaire template to guide a series of interviews with the ISVs. The template provides a consistent, repeatable process that can be deployed globally. The questionnaire is more than a simple checklist. It attempts to lead the ISV through a discussion of the technical details of its application architecture and implementation.

The interviews are designed to determine where the application of the ISV is today along a spectrum of "utility readiness" (i.e., readiness to be hosted and delivered as an ODS). The interviewer gains sufficient knowledge of the application to be able to identify what work efforts need to be completed for the application to be considered ready to exist in the utility computing environment and to exploit the services provided by the UMI and the application framework. The interview phase is also planned to be conducted worldwide by hosting consultants from the SPCs.

The interviews are conducted to draw out information relating to the attributes discussed next that an application will most likely need to become an ODS.

Recommended ODS attributes. Through several ISV interviews, market research, business models, and experience, we determined that the attributes described in this subsection are the top technical attributes that an ISV application must have in order to be a successful utility computing application.

One of these attributes is suitability to be hosted. An assessment is made as to whether the application has the rudimentary attributes to be hosted, such as by:

- Determining the manageability of mission-critical files or databases that are to be backed up and restored as needed
- Determining the capability of the application to issue the appropriate events or alerts so that the application can be monitored and managed effectively
- Determining Internet readiness of the application architecture and whether the application security implementation is sufficient for Internet delivery
- Determining whether the application can be administered remotely and securely in the hosting center (without physical access to the console or desktop)

In the case of the Ideas2Concepts time-tracking application, the application is assessed to be suitable for hosting. It is accessible through the Internet with adequate security and encryption protection. User and administrator access include levels of access rights, and both are accessible through the Internet.

Another of the attributes is suitability to be delivered as an ODS. An assessment is made as to whether the application has the characteristics or traits that are necessary to be a successful ODS as follows:

- Can the application scale horizontally and vertically to react to changes in demand? We define horizontal scaling as the ability to add application or middleware instances or servers when demand rises. Vertical scaling is defined as the ability to use features provided by the infrastructure, such as single sign-on, directory services, or digital cer-

tificate services, as needed to handle diverse business demand.

During the assessment it is found that the Ideas2Concepts time-tracking application uses the IBM WebSphere* Application Server to run the application logic and DB2* as its back-end database. For this case, horizontal scaling is dictated by the middleware used in the application. According to Ideas2Concepts, its application has been load-tested to scale well with 30 000 users on an eight-way configured IBM pSeries* machine running the AIX* operating system.

- Does it support one-to-many multitenancy? A multitenant application has the ability to share one application instance among several businesses or enterprise customers. The one-to-one, single tenant application is mostly sold as a stand-alone application. It is installed in a dedicated hosting center or on the customer premises and is used within the enterprise, even if the application is Web-enabled and Internet-accessible. A multitenant, one-to-many application needs a more robust level of security and requires more isolation but is able to be shared among multiple companies.

  In the case of the Ideas2Concepts time-tracking application, it is assessed to be a stand-alone application installed on the customer premises and made available through an Internet connection.

  Two important justifications to turn an application into a multitenant application are as follows:

1. Reduce management complexity—Multiple instances of the application require multiple machines. This requirement can add to the cost of maintenance and the hardware footprint. These costs are passed on to the enterprise customer, who then pays higher subscription fees. If the application were multitenant, a single instance of the application might serve multiple customers, resulting in lower costs for infrastructure maintenance, which could be passed along to customers.
2. Reduce cost of licensing middleware products—Depending on the licensing policies of middleware products, multiple instances of the application may result in multiple copies of middleware products and multiple license fees. Using one instance of middleware products not only saves on license fees but also minimizes the complexity of managing the middleware products.

  The assessment report will recommend that the Ideas2Concepts application be converted to a multitenant application.

- Can it use virtualized resources? Virtualization is a technology used in the on demand environment to respond to demand for server, storage, and network resources. The application of the ISV must not have any hardware- or platform-specific dependencies to be able to use virtual services. We consider the Linux** operating system on IBM zSeries*, Shared Processor Logical Partitioning on pSeries, and software from VMware, Inc. on xSeries* architecture as the type of virtual resource technologies most ISV applications will be using. Grid computing is not a consideration at this time; more rigor will be added to the assessment as the technology matures and becomes available in the utility hosting infrastructure.

  The Ideas2Concepts application has no hardware- or platform-specific dependency, which makes it a good candidate to use virtualized technologies currently offered in the utility computing centers. Use of grid technology will be mentioned as guidance for future planning only.

- Is the application packaging suited for auto-provisioning? To be able to auto-provision to meet rising and falling demand in the quickest possible way, the application packaging needs to be reconsidered for deployment in the on demand environment. The auto-provisioning criteria are based upon the UMI feature in which computing resources are built into server groups and distributed as a unit of provision. A server group can consist of server hardware, operating systems, network resources and topology, and applications. The application is assessed to determine what is needed in order for it to be deployed as part of a server group. As grid computing is incorporated into more applications, real-time auto-provisioning can handle more granular provisioning requests, such as allocating more CPU resources or disk drives to the running application.

- Can the application take advantage of a provided user subscription and management service? User subscription and management services are the functionality that will be provided by the application framework. User subscription functions include user authentication and access to resources based on user credentials. User management ser-

vices provide customer self-managed subscription services.

If Ideas2Concepts integrates with the application framework to take advantage of this functionality, the application will have the following capabilities:

1. The application framework will add user-subscription functions to the application. This includes single sign-on capability and the ability to restrict use of application resources according to access rights.

2. The application framework will add functionality that will allow Ideas2Concepts to delegate system administrator responsibilities to each respective enterprise customer. This would be very helpful if, for example, Ideas2Concepts had five enterprise customers with 2000 users each. Instead of Ideas2Concepts defining each of the 10 000 users to the application, the responsibility of subscribing the users is passed to the system administrators of the respective companies.

- Does the application have the ability to meter application usage? Application metering capability allows an application to monitor units of usage. Certain parts of the application can be monitored to produce usage and profiling information. Usage information can be fed to a billing system to bill subscribers for use of the application at a more granular level. Profiling information can give the application owner information that can be used to improve quality, performance, and features of a particular part of the application.

  The Ideas2Concepts application will have to implement metering after the application becomes multitenant. Currently, the application is sold as a stand-alone application, which does not require the use of metering technologies.

- Can the application incorporate and use Web services? Because the ODS architecture will use Web services extensively for its integration points, we assess the ISV application for amenability to Web services technology. Using Web services is a requirement only when the application has to communicate with the UMI, the application framework, or with other ODSs.

  The Ideas2Concepts application presents some of its time-tracking functionality as Web services so that human resources or accounting applications can readily integrate with it. The application is therefore assessed as Web services ready. No extra work is needed for this attribute.

Assessment benefits. The assessment process is an in-depth examination of the application architecture to identify where work has to be done for the application to "live" in the utility-computing center. It has benefits for both IBM and the ISV (i.e., Ideas2Concepts). It provides all involved parties with a clearer understanding of the status of the application in terms of on demand capability and what must be done to the application for it to be hosted in the utility computing environment. The assessment produces an assessment report for the ISV that consists of a task list showing the ISV what it needs to change, modify, or add to the application to become an ODS.

In the USF V5 AEP classic process, assessment was offered as a free service. It is possible that the assessment for ODS will also be offered as a free service. Because IBM may offer the assessment as a free service for ISVs, the assessment report identifies the "what" but not the "how." It will provide a task list of what needs to be done. It does not provide technical details on how to modify the application to meet the requirements to be an ODS.

Phase 2: The enablement: Consulting and compliance. After the assessment phase, Ideas2Concepts evaluates its business needs and decides whether to proceed. If it does proceed in the program, Ideas2Concepts makes a monetary commitment to the AEP process, and a contract is signed. At this point, the enablement phase is entered. The enablement phase consists of consulting, which is optional, and compliance, which is mandatory.

The rationale for offering optional consulting services is that some ISVs will choose to engage IBM for consulting assistance, whereas others may opt to do the transformation work themselves. Another reason may be that some applications already operate as utilities and use Web services for interoperability.

Enablement compliance is required because IBM must ensure the quality of the shared utility computing environment. Compliance determines the degree to which the application conforms to UMI standards and ODS architecture requirements.

Consulting (optional). The consulting phase begins when Ideas2Concepts decides to enlist IBM's exper-

tise on hosting and utility computing. Service Delivery Center technical architects or hosting consultants will work with the ISV on the basis of the activity list contained in the assessment report. The consulting architect will determine the scope of work and will document what has to be done. Service Delivery Center consultants may provide tools and guidance to Ideas2Concepts, and the IGS Application Management Services (AMS) organization may be involved to provide programming services for development and testing of the Ideas2Concepts application.

The AEP development team is creating a globally consistent consulting methodology that will facilitate the delivery of ODS consulting services. This methodology will include current best practices and processes learned from IBM AMS field experiences, the UMI and SPI development organizations, and other IBM software development teams. The methodology will have a step-by-step process to transform the Ideas2 Concepts application into an ODS, with task lists, work products, skill descriptions, and level of effort estimates and costs. The methodology will borrow from and be aligned with the appropriate IGS methods.

As part of the transition from development to delivery, the AEP development team will conduct a series of skills transfer sessions with the Service Delivery Center teams as the program is deployed worldwide. The skills transfer sessions will educate the teams in the details particular to the utility computing environment and to delivering applications as ODSs. Sessions will be provided for the teams, including technical architects and transition managers, and for the AMS group.

Consulting benefits. IBM AEP consulting will improve the degree of confidence that the application of an ISV will easily pass compliance and be hosted on the IBM utility platform in a timely manner. Those ISVs that do not have in-house expertise in utility computing would benefit from contracting with IBM for AEP consulting.

Compliance process. In the classic AEP for USF V5, a compliance process was not required because applications were deployed primarily on customer-owned or rented servers. Another reason for not requiring compliance is that applications were mostly standalone, not requiring any interaction with IBM-owned infrastructure components.

The compliance process for AEP for ODS is a required step before boarding in the IBM utility-computing infrastructure, independent of whether the ISV used AEP consulting services. Compliance ensures that all work done to the application is correct according to ODS architectural requirements. The rationale behind the requirement is that IBM is responsible for resources deployed in the utility-computing infrastructure and must have a high degree of confidence that deployed applications will interact correctly and also that high Quality of Service can be provided to the entire ISV community.

As the UMI technology matures and the application framework takes shape, the AEP development team will create a series of checklists that Service Delivery Center architects will use to conduct the ODS compliance process. The compliance architect will use these checklists to determine answers to questions such as:

- Can the application be delivered as a utility (e.g., multitenancy, scalability, etc.)?
- Does the application conform to the on demand architecture (e.g., does it use open standards technologies such as Web services)?
- Does the application have all the minimum requirements to be boarded (backup and restore plan, monitoring plan, packaging plan)?

Compliance details. There are two areas of compliance:

1. Compliance for boarding on UMI—meets hosting requirements, auto-provision application package requirement, and others
2. Compliance as ODS—multitenant, proper use of SPI (if ODS uses the SPI), proper architecture to provide Web services to other ODSs, and proper use of the application framework through known integration points.

At the time of writing, the ODS compliance process is a paper exercise that depends on receiving satisfactory answers from ISVs to the checklist questions. In the future, the compliance process could require the ISV to run the application against one or more UMI "simulators" and application framework components on a nonproduction hosting platform in the utility center to ensure interface compliance.

The Services Delivery Center architect conducting the compliance process will go over the checklist with Ideas2Concepts to ascertain whether each requirement is met. If any requirements are not met, the reason for noncompliance must be determined in or-

der to assess the degree of impact on other applications or UMI services. Depending on the severity of the noncompliance, appropriate remedies will be recommended. Successfully completing the compliance checklist will result in a conclusion that the ODS can be safely boarded in the utility hosting center.

Compliance benefits. If the ISV application is in compliance, it improves efficiency of deployment because all required steps are validated as complete, reduces the risk of encountering major problems as a result of having knowledge about IBM ODS transformation, and increases compatibility of deployment in the hosting environment.

Benefits of the AEP process

Although ODS owners may choose to convert their current applications to an ODS themselves, going through the AEP process offers significant benefits to the owners.

One benefit is having a "one-stop shop" for experts on utility computing subject matter. The AEP group has extensive knowledge of the on demand infrastructure and application framework architectures. These technologies are at the center of the IBM utility computing infrastructure. Combining this knowledge with their experience with ISV applications produces the business and technical "know how" to assist in the conversions of traditional hosted applications to ODSs on the IBM utility computing infrastructure.

A second benefit is having optional consulting services for converting applications to on demand services. IGS Service Delivery Center technical consultants can recommend cost-effective, flexible, extensible, and standards-based solutions that are backed by IBM resources. IBM will also recommend best practices on how to turn ISV-hosted applications into on demand services. Utilizing IBM consulting services should also improve the potential for certification.

A third benefit is having the compliance process make sure ODSs adhere to IBM standards and guidelines used in UMI and the application framework. Through the compliance process, the ODS owner is assured that the application adheres to standards used in the on demand infrastructure environment. Certification offers a greater chance of successful boarding and reduces the possibility of encountering problems. It offers a comprehensive checklist of necessary tasks that an ODS owner must do to successfully operate on the IBM utility computing infrastructure.

Conclusion and summary

Several business reasons to generate new revenue compel ISVs to adopt new technologies or change their business model.[7] Such is the case in the example of the Ideas2Concepts Corporation. Changing its business model from a licensed software model will allow it to reach new customers, expand its existing partnership or build new ones, and expose existing offerings to new delivery channels.

We introduced ODS architectural concepts to understand where an ISV application fits in as well as why the ISV application has to be transformed into an ODS. We described the design of an AEP process as it evolved from an existing one (traditional hosting on USF V5), which answers the fundamental question an ISV will ask: How do I turn my application into an ODS? The process has three stages: education and assessment, enablement, and compliance and deployment.

Although the AEP process is unique and a prerequisite before an ISV application is hosted in IBM utility centers implementing UMI, it does not lock the ISV into any IBM technology. Technologies used to enable the ISV to become an ODS are based on open standards that allow the ISV application to adapt easily to any other vendor implementation of an on demand infrastructure.

The AEP process is a streamlined, repeatable, and globally deployable process. It offers several benefits, including risk reduction, guided implementation, faster time to market, and quality assurance to the ODS owner. The AEP also offers the benefits of cost reduction, reduced risk, faster customer enablement, better utilization of resources, and better customer satisfaction to the utility platform owner.[8]

As for the future of the AEP process, we see it evolving to accommodate new advancements in technology for utility computing.

*Trademark or registered trademark of International Business Machines Corporation.

**Trademark or registered trademark of Linus Torvalds.

Cited references and note

1. S. Nickolas, *IBM Framework for e-business: Application Hosting Services*, developerWorks, IBM Corporation, http://www.ibm.com/developer.

2. G. Thickins, *Utility Computing: The Next New IT Model* (April 2003), http://www.utilitycomputing.com/news/258.asp.
3. *IBM and Strategic Potential of Web Services, Assessing the Customer Experience*, IDC, Framingham, MA (May 2002).
4. A. Gutierrez, *e-business on demand: A Developer's Roadmap*, developerWorks, IBM Corporation (February 2003), http://www.ibm.com/developer.
5. L. McCabe, *Out of the Box: Top Nine Net-Native Software-as-Services Design Differentiators,* Market Strategy Report, Summit Strategies, Inc. (September 2002).
6. A. Mizoras, *In-house versus Outsourced: A System Dynamics Comparison*, #27743, IDC, Framingham, MA (August 2002).
7. D. Gisolfi, *Web Services Architect, Part 2: Models for Dynamic E-business* (April 2001), http://www-106.ibm.com/developerworks/webservices/library/ws-arc2.html.
8. The descriptions of UMI and the application framework and its implementation are all to be viewed as intended directions and not as commitments to be delivered in IBM products. The actual functions and features delivered in UMI releases may vary considerably from what has been described or referenced in this paper.

Accepted for publication September 25, 2003.

Kay Chang *IBM Global Services, 11501 Burnet Road, Austin, Texas 78758 (kchang@us.ibm.com).* Ms. Chang is a senior software engineer working on the Application Enablement Program. Among previous assignments, she worked on delivering the mySAP™ hosting solution, mail, and messaging hosting solutions using wireless devices. She was a business development manager for telecommunications server management, using IBM and Nokia wireless products. She was also an AIX system architect working on expanding AIX user management to the LDAP-based, federated user management in the area of identification and security. Ms. Chang helped to develop various Internet-based methods for security, such as ipsec (IP Security Protocol), PKI (Public Key Infrastructure), and Kerberos, as well as the next generation of Internet Protocol (IP V6) for the AIX system. She has a Master's degree in computer engineering from Wright State University.

Anthony Dasari *IBM Global Services, 11501 Burnet Road, Austin, Texas 78758 (adasari@us.ibm.com).* Mr. Dasari is an IT architect in the Application Enablement Program for small and medium business ISVs. He has been active in customer consulting engagements for the past 12 years in various industry sectors, such as health care management, oil and refinery, telecommunications, publishing, automobile, criminal investigation, and the state and federal government. Prior to joining IBM, he was the technical director of a profitable multimillion dollar consulting and integration company. He performed various architectural roles such as data architect, information architect, solutions architect, infrastructure architect, and business process architect. The experience that he gained from multiple industry sectors and various architectural roles has given him unique insight into the various perspectives of an e-business architecture and the e-business on demand solution areas in various industry sectors. Mr. Dasari has an M.S. degree in computer science from Oklahoma City University and a B.S. degree in technology with specialization in electronics and communication from Nagarjuna University.

Hari Madduri *IBM Software Group, 11501 Burnet Road, Austin, Texas 78758 (madduri@us.ibm.com).* Dr. Madduri has an M. Sc. degree in physics from Osmania University (India) and an M. Phil. degree in computer applications from the University of Hyderabad (India). He also has Master's and Ph.D. degrees in computer science from the University of Wisconsin, Madison. Prior to being with IBM, he was a researcher at Honeywell Inc., working on fault-tolerance and distributed systems. Since joining IBM in 1990, he has worked on distributed object technologies, data mining, and hosted service development. His more recent projects have been electronic bill presentment and payment service, data transformation service, UMI, UMI metering, the application framework for UMI, and the Application Enablement Program for IBM server farms and UMI, most of which are examples of the "software as a service" concept. A Senior Technical Staff Member, Dr. Madduri is currently an architect in Tivoli. He is working on Tivoli orchestration and provisioning technologies with a focus on metering and service provisioning.

Alfredo (Freddie) Mendoza *IBM Global Services, 11501 Burnet Road, Austin, Texas 78758 (mendoza1@us.ibm.com).* Mr. Mendoza is an advisory software engineer in the Application Management Services Group within IBM Global Services. He currently works in the application enablement program that involves creation of processes and procedures that will transform ISV applications to on demand services. Prior to this position, he worked in AIX kernel development as the component lead for libpthreads, and before that he worked in the IBM Solutions Developer Organization, Tivoli, and AIX service and support as an independent consultant. Before coming to work for IBM in the United States, he worked for IBM in Wangaratta, Australia as a product engineer with the task of learning and transferring manufacturing technology from the United States for such products as the former IBM ROLM telephone switches and IBM Micro Channel® 370. Mr. Mendoza has a B.Sc. ECE from De La Salle University in the Philippines.

John Mims *IBM Software Group, 11301 Burnet Road, Austin, Texas 78758 (mimsj@us.ibm.com).* Mr. Mims has been with IBM since 1989 and is currently a senior software engineer with Tivoli in the IBM Software Group. He holds a Bachelor's degree from the University of Texas at Austin and is an IBM Certified IT Specialist. He has spent most of his IBM career in mainframe software development and technical support and has presented technical seminars and workshops around the world. Prior to an assignment with IBM Global Services, he was the author of two IBM redbooks and a number of technical white papers. He is currently assigned to Tivoli Education Development in Austin.