


IBM Workload Deployer V3.0


New feature overview



© 2011 IBM Corporation

This presentation provides an overview of new product features in IBM Workload Deployer V3.0.

With IBM Workload Deployer V3.0, you get:



- New product branding to reflect broader-than-WebSphere scope
- Support for existing capabilities of WebSphere CloudBurst Appliance V2.0 and fix packs
 - Includes support for existing Hypervisor Edition images
- New capabilities for the V3.0 release, including support for workload-based virtual applications and more
- New physical appliance with more storage, memory, processing power, network capacity
 - Not compatible with earlier versions with WebSphere CloudBurst V2.0 hardware

IBM Workload Deployer V3.0 is the next evolution of the WebSphere CloudBurst Appliance. The re-branding of the product, to remove WebSphere from the name and focus on IBM, reflects the broader scope of IBM Workload Deployer as a deployment platform for multiple middleware environments that integrates across IBM's software brands. The IBM Workload Deployer product requires a new physical hardware appliance and is not compatible with earlier WebSphere CloudBurst Appliance hardware. The new hardware platform for IBM Workload Deployer offers expanded processing power, memory, network capacity, and storage. Additional storage is critical to enable you to store more images and patterns on your appliance that can be used to deploy tightly integrated, cross-product workloads. All of the existing capabilities of WebSphere CloudBurst Appliance V2.0 and its fix packs have been rolled into IBM Workload Deployer, including the ability to easily create custom middleware topology patterns using Hypervisor Edition images, and deploy those patterns to your private cloud. Your existing patterns, images, and custom content from WebSphere CloudBurst will continue to work with IBM Workload Deployer, including all of the product-specific Hypervisor Edition images provided by IBM. IBM Workload Deployer will come with migration tools that allow you to transition your existing WebSphere CloudBurst environments to the IBM Workload Deployer platform. In addition to WebSphere CloudBurst V2.0 functionality, IBM Workload Deployer V3.0 has been enhanced with many new product capabilities, including support for a new workload-based deployment model for virtual applications.

IBM Workload Deployer V3.0 – new capabilities

New Appliance Hardware

The appliance can hold more images and patterns with over 6x the storage capacity of WebSphere CloudBurst V2.0; also more network capacity and compute power

Autonomic Elasticity

Intelligent Management Pack-enabled virtual system patterns scale up / down according to observed demand

VM Mobility

User may request that a virtual machine is moved from its present location in order to optimize placement

Advanced VM Control

More control over IP address and name assignment for virtual system patterns; resource pools for sub-clouds; define hardware resources for VMs with add-ons

Workload Patterns

Enhanced patterns with cross-product integration and full life cycle management

IBM Workload Deployer V3.0 is built on an advanced hardware platform with over six times the storage capacity of the WebSphere CloudBurst Appliance V2.0. In addition, the new hardware platform has more compute power, memory, and networking capacity. IBM Workload Deployer supports two main deployment models – topologies and workloads. The next three items on this page – autonomic elasticity, virtual machine mobility, and advanced virtual machine control – are all related to topology deployments. Autonomic elasticity is a capability provided by the Intelligent Management Pack that allows your deployed virtual systems to automatically scale up and down by spinning up new virtual machines and server instances in response to observed demand. The new virtual machine mobility feature provides a way to manually load balance virtual machines and optimize placement post-deployment by explicitly moving a virtual machine from its current location to some other hypervisor. The advanced virtual machine control capabilities allow you to use environment profiles to control virtual machine naming, IP address assignment, and resource usage in the sub-cloud represented by the profile. You can also include specialized add-on components in your virtual machines to perform operating system-level customizations, such as creating user accounts or additional disks. Finally, workload patterns represent a significant new deployment model for IBM Workload Deployer. In addition to the existing topology model, workload patterns allow you to quickly define and deploy cross-product integrated solutions into your private cloud. There are two workload patterns available with IBM Workload Deployer V3.0 – the Pattern for web Applications and the Pattern for Database Applications.

Announcements – WebSphere Application Server Hypervisor Edition

Red Hat Linux on System z

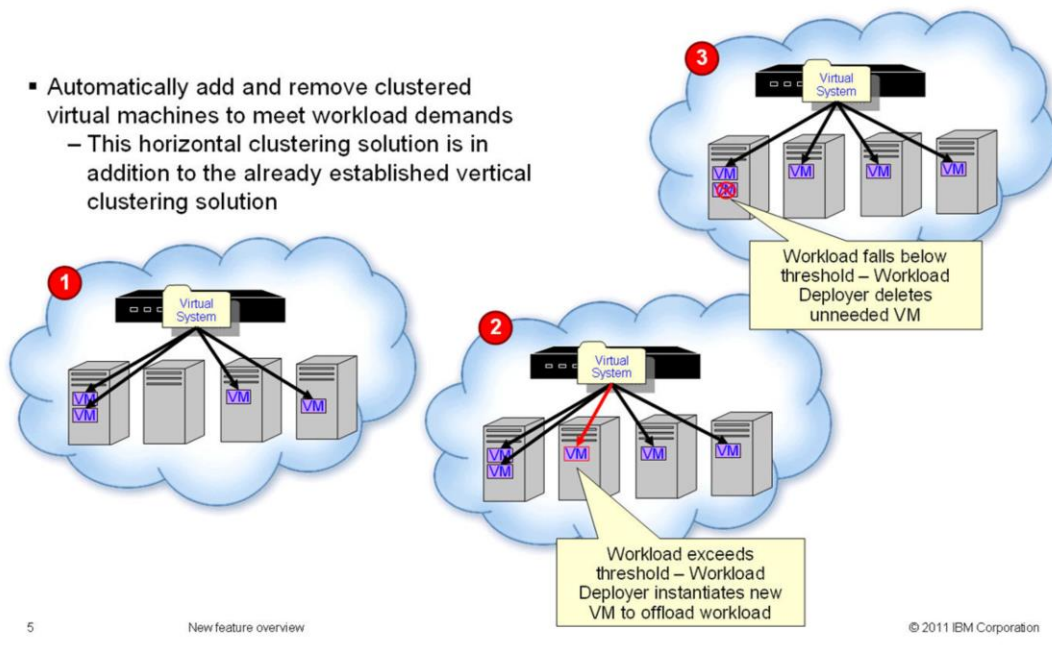
64-bit Linux support

Linux images are now available with an embedded 64-bit OS – VMware images still use a 32-bit application server

In addition to the base appliance functionality that is coming in IBM Workload Deployer V3.0, you will also have access to new WebSphere Application Server Hypervisor Edition images. For z/VM, there is a new Red Hat image that you can use to deploy virtual systems using Linux on System z. There are also 64-bit Linux images available. On the z/VM platform, the 64-bit images include both a 64-bit operating system and a 64-bit application server. On VMware ESX, the new 64-bit images contain a 64-bit operating system and a 32-bit application server.

Autonomic elasticity with Intelligent Management Pack


- Automatically add and remove clustered virtual machines to meet workload demands
 - This horizontal clustering solution is in addition to the already established vertical clustering solution



The Intelligent Management Pack is an optional add-on for WebSphere Application Server Hypervisor Edition that uses WebSphere Virtual Enterprise capabilities to support advanced topologies. New in V3.0, the Intelligent Management Pack supports autonomic elasticity for your virtual systems, allowing your WebSphere clusters to automatically add or remove clustered virtual machines to meet service-level policies. Say, for example, that your cluster begins with four virtual machines, as shown in step 1 on the left of this page. If the workload for your cluster exceeds the thresholds that you have defined, IBM Workload Deployer will automatically instantiate a new VM to offload that work and improve the responsiveness of your environment. You can see this in step 2 on this page, where a new VM is added to the hypervisor that was previously empty. Finally, step 3 shows an example of scaling down your cluster. If the workload falls below your defined thresholds, the appliance can delete VMs that are no longer needed to free up resource in your cloud.

Virtual machine mobility

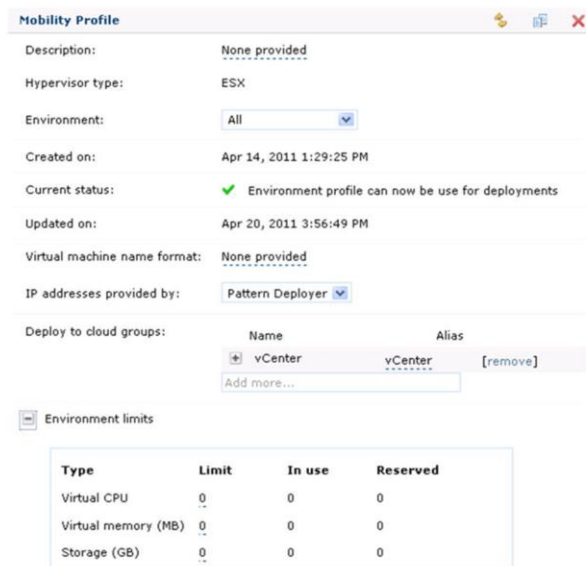
- Mobility provides the ability to initiate a live migration of a running virtual machine using IBM Workload Deployer
 - Supported on vCenter managed cloud groups; requires vMotion



Virtual machine mobility provides the ability to initiate a live migration of a virtual machine from one hypervisor to another, using the IBM Workload Deployer interface. In IBM Workload Deployer V3.0, virtual machine mobility is supported for VMware only, and it requires that you are using vCenter-managed cloud groups with vMotion configured. Migrating virtual machines post-deployment allows you to alter the original placement decisions that the appliance made for your VMs or update hypervisors in response to future incoming load or hardware updates. For example, if you need to power down a hypervisor for maintenance, but want to maintain access to its deployed virtual machines, you can migrate those VMs off of the hypervisor before powering that hypervisor down for updates. Later, when you've completed your hypervisor updates, you can move the VMs back to their original host.

Advanced virtual machine control – environment profiles

- Use environment profiles to:
 - Explicitly assign IP addresses to pattern parts at deployment time
 - Customize names of virtual machines
 - Use replacement variables, prefixes, suffixes
 - Deploy a pattern to multiple cloud groups
 - Control resource usage for the sub-cloud
 - Virtual CPU, memory, storage, licenses



Mobility Profile

Description: None provided

Hypervisor type: ESX

Environment: All

Created on: Apr 14, 2011 1:29:25 PM

Current status: ✓ Environment profile can now be use for deployments

Updated on: Apr 20, 2011 3:56:49 PM

Virtual machine name format: None provided

IP addresses provided by: Pattern Deployer

Deploy to cloud groups:

Name	Alias
vCenter	vCenter
Add more...	

Environment limits

Type	Limit	In use	Reserved
Virtual CPU	0	0	0
Virtual memory (MB)	0	0	0
Storage (GB)	0	0	0

7

New feature overview


© 2011 IBM Corporation

Environment profiles provide advanced control of virtual machine configuration and placement in the cloud. During virtual system pattern deployment, you use an environment profile as the deployment target in the same way that you can use a cloud group. An environment profile contains a collection of cloud groups and IP addresses, along with other customization parameters that will apply to the virtual systems deployed using that profile. Since a profile can contain multiple cloud groups, when you deploy a pattern to an environment profile, you can explicitly place your virtual machines in different portions of your private cloud. For example, if you want to separate your web tier and application tier across different network segments, you can do so using an environment profile as the deployment target for your pattern.

Profiles allow you to explicitly pin IP addresses to pattern parts at deployment time, as opposed to having the appliance assign the IP addresses. You can also use environment profiles to customize the format of the virtual machine names, using hard-coded strings and a set of replacement variables. New in V3.0, environment profiles provide support for managing resource usage. You can set limits in a profile for virtual CPU, memory, storage, and product licenses. For example, you might have a portion of your private cloud that you want to assign to your test team, and of that sub-cloud, you only want to make 16GB of memory available for the test team to be able to consume with their virtual system deployments. To achieve this, set a virtual memory limit in your environment profile. Then, the appliance will restrict any deployments that would push the cumulative memory usage for the profile's hosted systems above the defined limit.

Advanced virtual machine control – add-ons

- An add-on is a built-in solution to customize and augment virtual machine configurations before deployment
 - Define network interfaces
 - Define the number and size of the disks
 - Define additional users



8

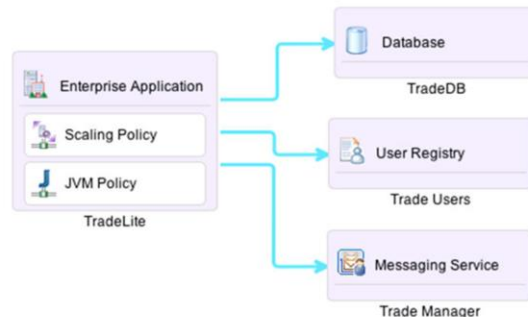
New feature overview

© 2011 IBM Corporation

Another new virtual machine customization capability available with IBM Workload Deployer V3.0 is a virtual machine add-on. An add-on is a special script designed to perform common operating system-level customization so that you no longer have to extend images or write your own scripts for some common scenarios. There are three types of add-ons available – disk, network interface, and user account. By defining a disk add-on in your pattern, for example, your deployed virtual machine will include an additional virtual disk, of the size and file system format you requested, mounted at a mount point that you specified. There are default add-on implementations provided with the appliance that include both hypervisor-level API calls and a higher level add-on implementation script package. That implementation package is customizable, so you can modify the default implementation to meet the specific needs of your configuration.

Virtual applications

- IBM Workload Deployer Pattern for web Applications is a platform-as-a-service (PaaS) solution in which your application takes center stage
 - You start with your application and define its specific requirements, such as services (or middleware) and policies (or QoS)
 - Appliance takes care of configuring the environment your application needs to run
- There are three parts to a virtual application pattern:
 - Your application
 - EAR or WAR
 - Dependent services
 - For example, DB2, LDAP
 - Policies



Virtual application patterns are new in IBM Workload Deployer V3.0. These patterns are defined using application components and related services, rather than topology components. For example, this virtual application pattern contains one enterprise application that connects to three services – a database, a user registry, and a messaging service. The blue arrows are links from the application to its required services. Virtual applications can connect to existing cloud resources, or deploy new resources to meet the application's needs. When defining the pattern, you choose whether to create a new database, for example, or connect to an existing database in your enterprise. The components in a virtual application have a fixed set of exposed configuration parameters that you can customize. For example, the enterprise application component requires that you provide the EAR file that represents your application. You can also define qualities-of-service for your application using service policies. In this example, the enterprise application components includes both a scaling policy and a JVM policy. These policies allow you to define how your application will scale in response to load, and to set JVM options, such as heap size and generic JVM parameters.

Virtual application features

Automated Scaling

Managed environments scale up and down based on observed utilization of compute resources

Failover

Failed virtual machines are replaced with new VMs which are configured with the old VM's identity

Load Balancing

Requests coming into workload pattern environments are load balanced

Security

ACLs for application sharing and management access, LDAP integration for application security

Monitoring

All components of workload pattern environments are monitored by IBM Workload Deployer

Virtual applications offer a very simple approach to building and deploying application serving environments using a policy-based scheme to control non-functional requirements and qualities of service. The scaling policy provides automated scaling capabilities, allowing virtual application environments to scale up and down in response to resource utilization thresholds and parameters that you define in the policy for your application. Virtual applications also support automatic failover, replacing a failed virtual machine with a new virtual machine that takes over the original identity of the failed system. Requests coming in to workload pattern environments can also be load balanced, using the built-in shared proxy service for the Pattern for web Applications. Virtual applications use the built-in security model of IBM Workload Deployer to manage access to application resources on the appliance. The Pattern for Web Applications also supports LDAP for application security – either by deploying a new LDAP instance as a part of your application pattern, or by connecting to an existing LDAP server. The IBM Workload Deployer interface provides tools to monitor all of the components of your workload pattern environment, including operating system and middleware monitoring, updating the application configuration, and troubleshooting.

User view of virtual applications

- As a user, you:
 - Work with application assets for a solution (for example, WAR, EAR, DDL)
 - Have a limited, fixed set of software configuration exposed for the solution
 - Select which policies to apply to your application (for example, scaling, high availability)
 - Specify connections to remote services
- As a user, you do not control or know:
 - How many virtual machines are actually created to support your application
 - How the software is installed onto the virtual machine
- As a user, after deployment, you:
 - Do not have direct access to the virtual machines
 - Do not have direct access to the individual product consoles (for example, no application server administrative console)

Virtual applications are designed to be very simple to build, deploy, and manage. As a virtual application user, you work with the assets that define your application – for example, the EAR file and database schema for a web application – and build your application using a graphical editor that exposes a limited, fixed set of configuration parameters. You can also add policies to your application to define qualities of service – for example scaling and high availability – and connect your application components to external resources, such as a database or messaging service.

Your view of the application definition and the deployed application environment is limited by design, to make it very simple to deploy tightly integrated solutions. You will not control exactly how many virtual machines are created to support your application, or how the supporting software components are installed on each virtual machine. Post-deployment, IBM Workload Deployer does not provide direct access to the virtual machine operating system or middleware. Instead, you use the IBM Workload Deployer interface to access a deployment inlet for your application that allows you to monitor the environment, update the configuration, and access logs.

External resources for Pattern for web Applications

External component	Supported products
Database	DB2 for z/OS V9.0, V9.1, or V10 DB2 for Linux, UNIX, and Windows (versions supported with WebSphere Application Server V7.0) Oracle Database 11g Standard or Enterprise Release 1 or 2
Messaging	IBM WebSphere MQ V6.0.2 or V7.0
LDAP servers	IBM Tivoli Directory Server V6.0 and future releases Microsoft Active Directory 2003 and future releases
CICS	IBM CICS Transaction Gateway for z/OS V7.0, V7.1, V7.2, or V8.0 IBM CICS Transaction Gateway for Multiplatforms V7.0, V7.1, V7.2, or V8.0

The IBM Workload Deployer Pattern for web Applications enables you to quickly and easily connect your applications environments to existing resources in your enterprise. This page lists some of the external resources that are supported with the Pattern for Web Applications, including databases, IBM WebSphere Message Queue, LDAP servers, and IBM CICS Transaction Gateway. Notice that the Pattern for Web Applications integrates with both IBM and non-IBM software – for example, you can connect to Oracle Database and use Microsoft Active Directory as your LDAP server. The list of supported external resources is subject to change, so always consult the product documentation for the latest information on supported components and versions.

Scaling policy example


- **Example:** Application pattern, no scaling policy




- **Example:** Possible instantiated solution



- **Example:** Application pattern, with scaling policy



- **Example:** Possible instantiated solution



13

New feature overview

© 2011 IBM Corporation

Building virtual application patterns is very simple. First, assume that you are creating a very simple application environment, with an enterprise application and a database, as shown on the left of this page. You can do this in the graphical Virtual Application Builder by dragging the application components into your pattern, then providing your EAR file and a database schema. When you deploy that pattern, it will likely be deployed as two virtual machines – an application and a database, as shown in the lower left of this page.


Consider a second example, as shown on the right of this page, in which you also include a scaling policy in your application definition. With the scaling policy, you can define constraints for load balancing and scaling in your environment. The process of adding a scaling policy to the application is very simple – just choose the policy in the editor and select from a handful of scaling options to choose the ones that meet your needs. By taking this one extra step of including a policy in your application, the resulting deployed topology will likely look much different. For example, you might end up with an environment that is fronted and load balanced by a shared proxy service, and multiple application server instances to share your application load. The power of the virtual application model is that you are able to very quickly and easily deploy advanced topologies to support your application by making simple policy-level updates to your application pattern.

Scaling with failover and load balancing


- 1) When a virtual machine becomes unavailable, IBM Workload Deployer replaces it automatically



- 2) The new virtual machine is created with the failed system's identity; requests continue to flow normally



- 3) If the application is not meeting its service-level agreement, IBM Workload Deployer will automatically scale up



14

New feature overview

© 2011 IBM Corporation

In addition to being very simple to build and deploy, the entire life cycle of a virtual application is managed by IBM Workload Deployer. The sequence shown on this slide represents the same application that you saw on the previous page – an enterprise application defined with a scaling policy, that connects to a database – and shows how the environment automatically fails over, balances load, and scales to meet demand.

When a virtual machine becomes unavailable, IBM Workload Deployer replaces it automatically – step 1 represents a deployed virtual machine becoming unavailable in your cloud. In step 2, a new virtual machine is created with the failed system's identity. Throughout this process, application requests continue to come in through the proxy service and flow normally to the application servers that are still available and handling requests. In step 3, if an application is not meeting its service-level agreement as defined in the scaling policy, IBM Workload Deployer will automatically scale up to meet demand.

Monitoring virtual applications

- From the deployment inlet, monitor virtual hardware, OS, middleware
- Perform administrative tasks on the application environment
- Access log files for the entire application environment

IBM Workload Deployer Deployment ID :d-4d94c1bb-3552-4dc4-bc7c-b355b2e868fe

Virtual Machine Monitoring Middleware Monitoring **Operation** Logging

Role Operation

Beta_09_DB2-db2.DB2	DB2
Enterprise_Application-was.WAS	WAS

Set WebSphere Application Server trace level

Description: The operation sets the trace level of the WebSphere Application Server. For example: *=INFO:com.ibm.websphere.*=FINEST. You can use trace logs to assist you in monitoring system performance and diagnosing problems.

Trace String:

Submit

Generate javacore

Description: The operation will generate two java core dumps of the WebSphere Application Server. The dumps will be generated one minute apart to aid in troubleshooting.

Submit

Update configuration

Description: The operation will update the parameters under this role dynamically.

WAR/EAR File: Browse Delete

15 New feature overview © 2011 IBM Corporation

IBM Workload Deployer does not provide direct access to the operating system or deployed middleware for your virtual application environments. Instead, you monitor and manage those environments directly from the IBM Workload Deployer interface using the deployment inlet for your application. From the deployment inlet, you can monitor the behavior and performance of the virtual machine and installed middleware, using displays that provide a variety of performance metrics. IBM Workload Deployer includes a panel for performing administrative tasks on the environment, such as updating the application configuration and updating the trace specifications for the deployed middleware. The deployment inlet also offers a centralized location to view and monitor all of the logs for your deployed workload.

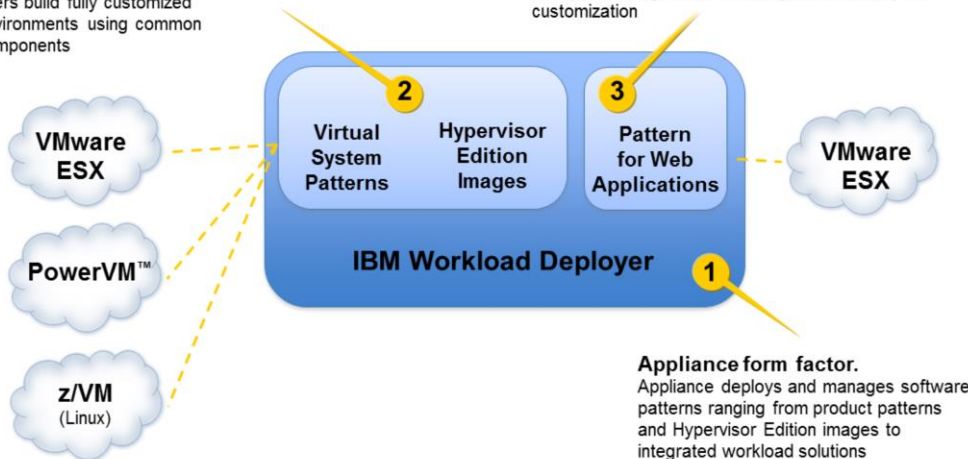
IBM Workload Deployer architecture

Virtual system patterns.

Set of optimized product configurations (Hypervisor Edition products + base patterns) using IBM best practices; users build fully customized environments using common components

Virtual application patterns.

Integrated solutions, such as the Pattern for Web Applications, delivered with integrated monitoring, life-cycle management, elasticity; limited ability for customization



16

New feature overview

© 2011 IBM Corporation

IBM Workload Deployer V3.0 is the next evolution of the WebSphere CloudBurst Appliance. The product integrates all of the existing capabilities of WebSphere CloudBurst Appliances and expands on the platform by offering new, cutting edge integrated workload deployments.

IBM Workload Deployer is delivered in an appliance form factor with enhanced processing power, networking, and storage capabilities. The appliance acts as a central deployment and management endpoint for your software patterns, ranging from existing product patterns and Hypervisor Edition images to integrated workload solutions. Virtual system patterns are built using Hypervisor Edition images and are standardized using IBM best practices. You can use these images and patterns to build and deploy fully customized environments based on common, standardized components. Virtual system patterns can be deployed to multiple hypervisor platforms, depending on the images you are using. IBM Workload Deployer introduces a new deployment model for integrated workload solutions that can be deployed to VMware ESX hypervisors. There are two workload patterns available with IBM Workload Deployer V3.0 – the Pattern for web Applications and the Pattern for Database Applications. These workloads are delivered as fully integrated solutions, with built-in monitoring, elasticity, and life-cycle management. They offer limited customization capabilities, but are extremely easy to build and deploy, offering exceptional value.

Trademarks, disclaimer, and copyright information

IBM, the IBM logo, ibm.com, CICS, CloudBurst, DB2, System z, Tivoli, WebSphere, z/OS, and z/VM are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of other IBM trademarks is available on the web at "[Copyright and trademark information](http://www.ibm.com/legal/copytrade.shtml)" at <http://www.ibm.com/legal/copytrade.shtml>

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, and the Windows logo are registered trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. IN ADDITION, THIS INFORMATION IS BASED ON IBM'S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE. IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION. NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, NOR SHALL HAVE THE EFFECT OF, CREATING ANY WARRANTIES OR REPRESENTATIONS FROM IBM (OR ITS SUPPLIERS OR LICENSORS), OR ALTERING THE TERMS AND CONDITIONS OF ANY AGREEMENT OR LICENSE GOVERNING THE USE OF IBM PRODUCTS OR SOFTWARE.

© Copyright International Business Machines Corporation 2011. All rights reserved.