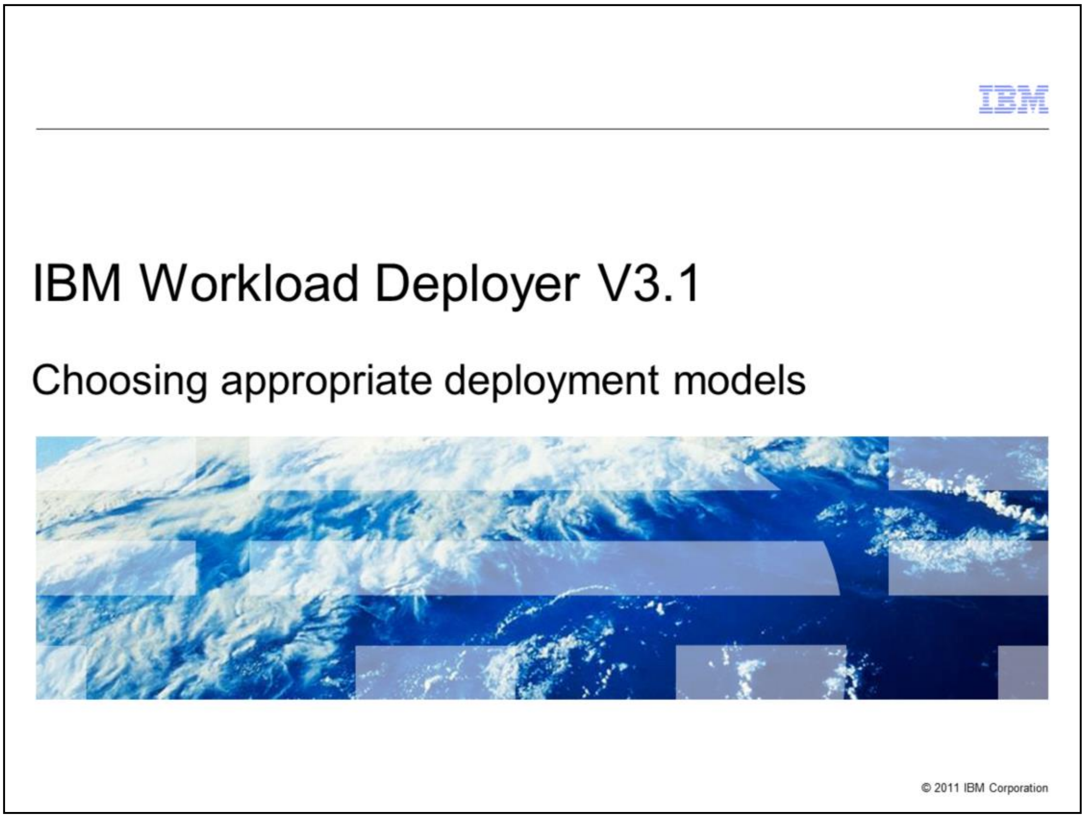

IBM Workload Deployer V3.1

Choosing appropriate deployment models

This presentation will cover deployment models in IBM Workload Deployer V3.1.

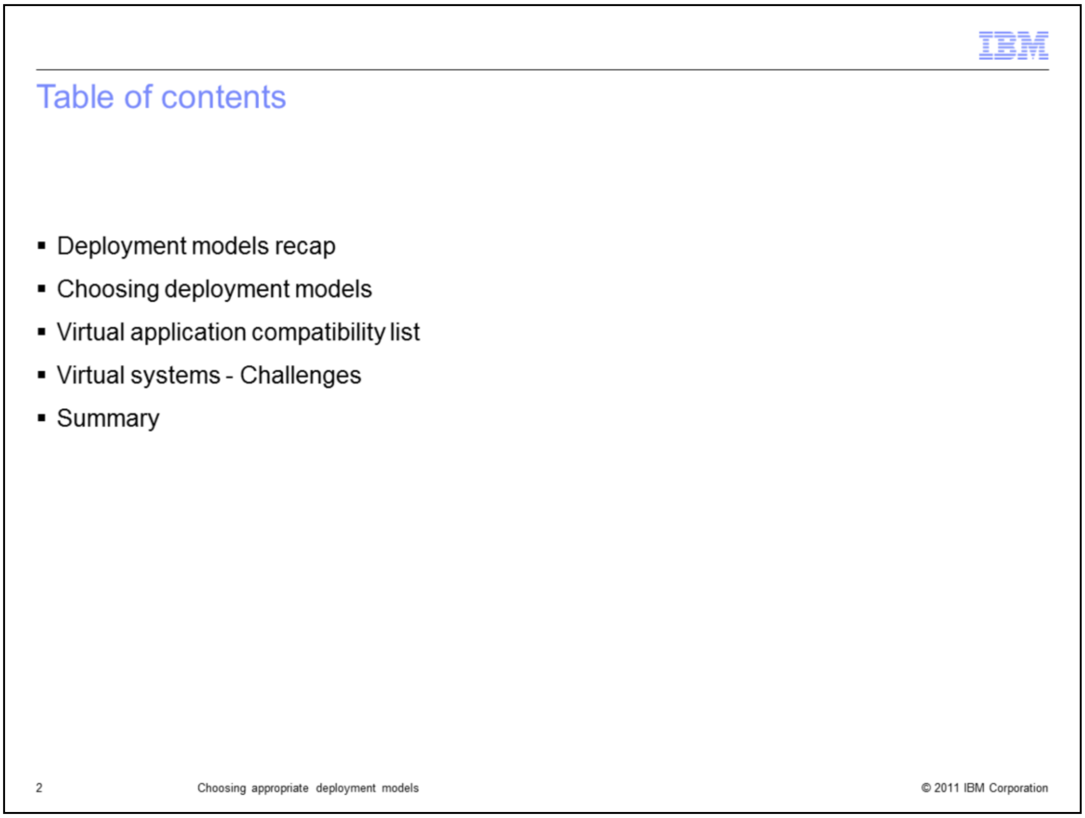

Table of contents

- Deployment models recap
- Choosing deployment models
- Virtual application compatibility list
- Virtual systems - Challenges
- Summary

This presentation will cover deployment models recap, choosing models, virtual application compatibility list, virtual systems challenges, and a summary.

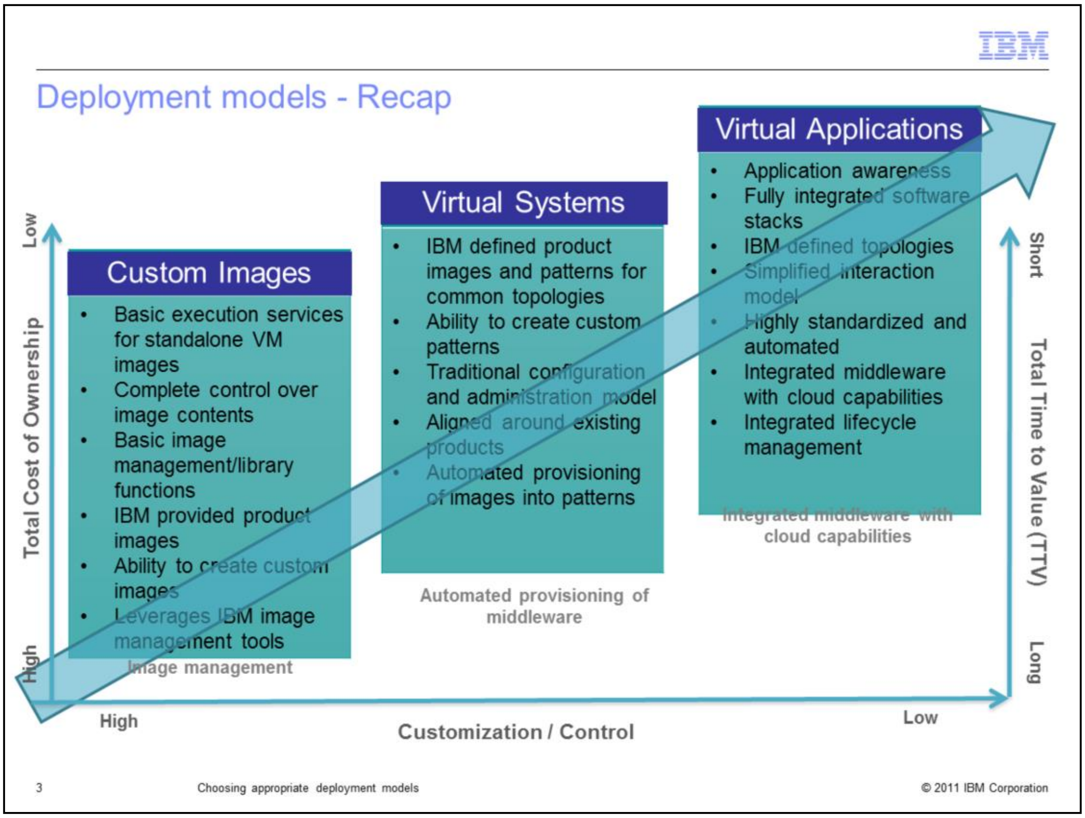

IBM

Deployment models - Recap

Custom Images

- Basic execution services for standalone VM images
- Complete control over image contents
- Basic image management/library functions
- IBM provided product images
- Ability to create custom images
- Leverages IBM image management tools

Image management

Virtual Systems

- IBM defined product images and patterns for common topologies
- Ability to create custom patterns
- Traditional configuration and administration model
- Aligned around existing products
- Automated provisioning of images into patterns

Automated provisioning of middleware

Virtual Applications

- Application awareness
- Fully integrated software stacks
- IBM defined topologies
- Simplified interaction model
- Highly standardized and automated
- Integrated middleware with cloud capabilities
- Integrated lifecycle management

Integrated middleware with cloud capabilities

3 Choosing appropriate deployment models © 2011 IBM Corporation

IBM Workload Deployer provides different deployment models, allowing you to choose which model is best suited for your needs. Consider image deployments. At the image level, you have a high degree of customization and control since you can define your own custom images. However, with that level of control comes both a long time to value – building your own images takes time; and a high cost of ownership – maintaining your own custom-built environment is typically expensive. By moving to the Virtual System model, you save time and ownership cost by relying on IBM’s pre-built Hypervisor Edition images for many product capabilities. You still have significant ability to customize your topology deployments – through image extension, by defining the specific topology and middleware levels for your application, and by using script packages to customize specific components. At the Virtual Application level, you have drastically reduced cost and time to value, because the solutions are pre-built and integrated for a specific use case. Rather than defining topologies, you provide your application artifacts, and the Virtual Application pattern will determine the appropriate underlying topology based on the service-level agreement that you provide. While using the Virtual Application model is the most cost-effective, its high level of standardization and cross-product integration results in fewer product configuration options being exposed for customization.

All deployment models – Client to logical to physical view

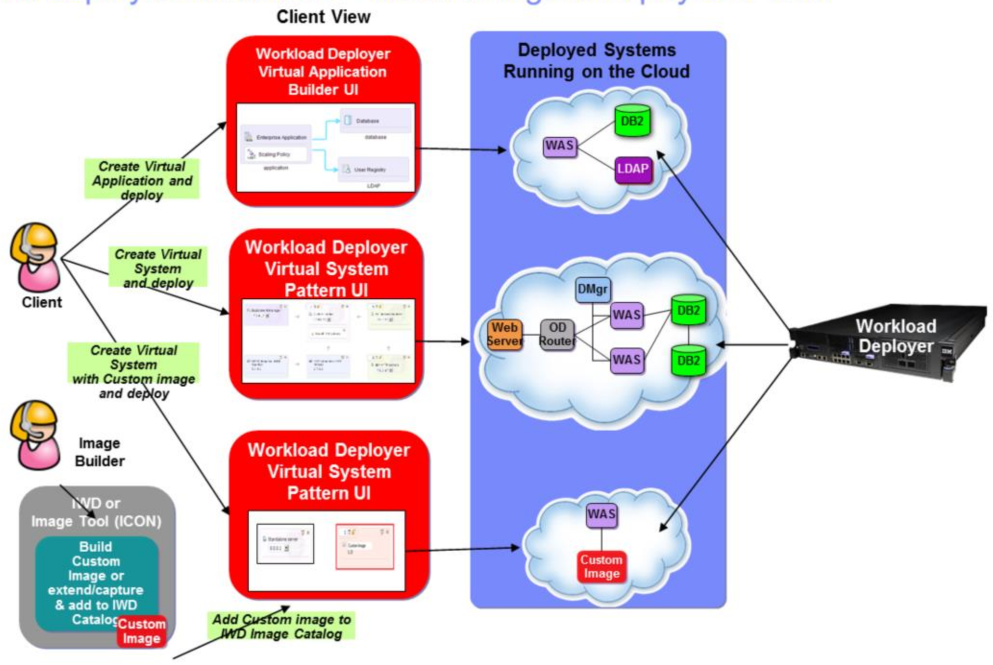
Client View
Workload Deployer Virtual Application Builder UI
Workload Deployer Virtual System Pattern UI
Workload Deployer Virtual System Pattern UI
Deployed Systems Running on the Cloud
WAS
DB2
LDAP
DMgr
Web Server
OD Router
WAS
WAS
DB2
DB2
WAS
Custom Image
Workload Deployer
Client
Create Virtual Application and deploy
Create Virtual System and deploy
Create Virtual System with Custom image and deploy
Image Builder
IWD or Image Tool (ICON)
Build Custom Image or extend/capture & add to IWD Catalog
Custom Image
Add Custom image to IWD Image Catalog

4 Choosing appropriate deployment models © 2011 IBM Corporation

The different deployment models have different client interfaces for building them. The Virtual Application Builder interface in IBM Workload Deployer allows you to focus on assembling your application artifacts without worrying about the underlying topology. The Virtual System builder interface lets you assemble various parts into the topology you want to use. Since the Virtual System builder is focused on the topology rather than the application, you will need to use script packages to configure your application and related artifacts. Custom images are built using either the ICON tool or by extending and capturing an existing image. These images can then be saved to the image catalog to be assembled into virtual systems.

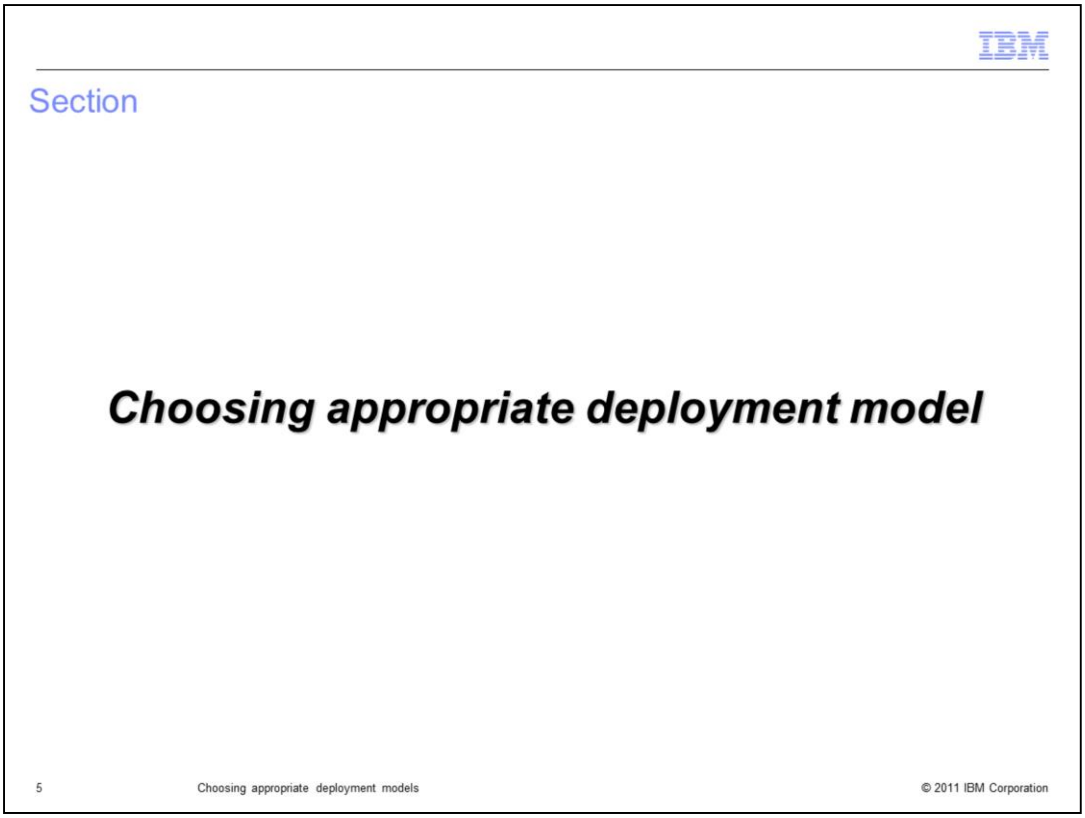

Section

Choosing appropriate deployment model

5 Choosing appropriate deployment models © 2011 IBM Corporation

This section will discuss how to choose between the available deployment models.

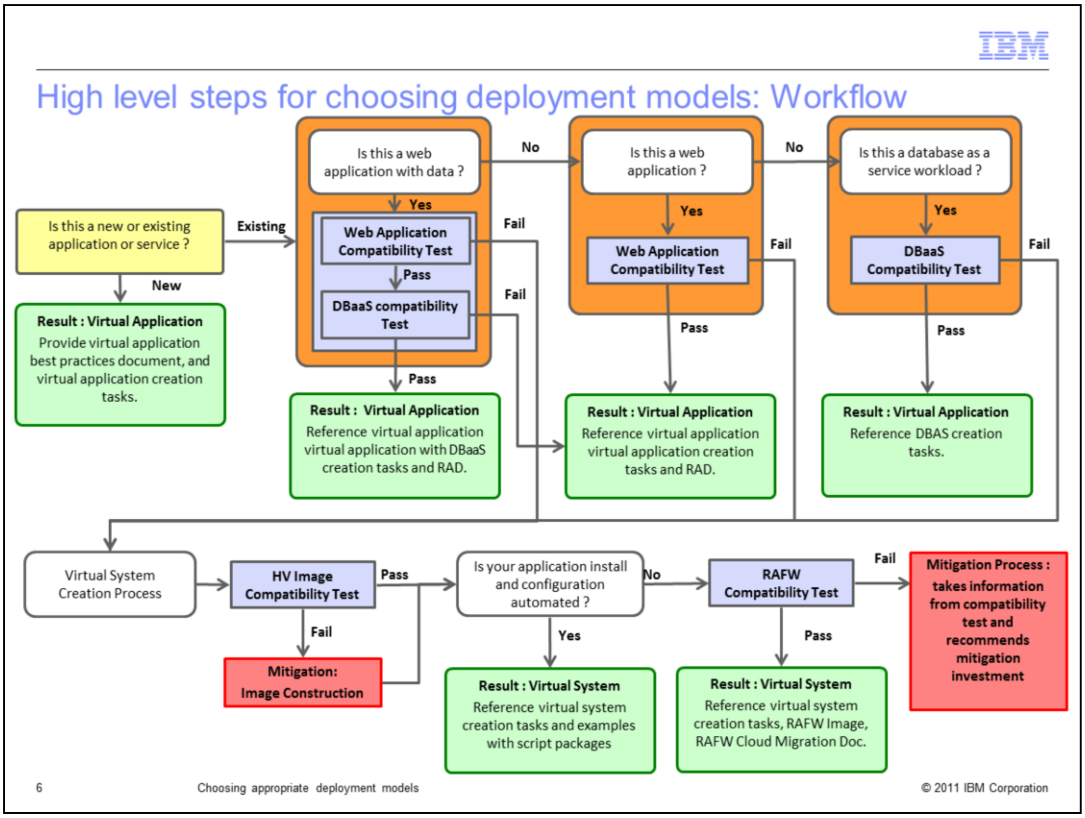

High level steps for choosing deployment models: Workflow

Is this a new or existing application or service ?

Existing

New

Result : Virtual Application
Provide virtual application best practices document, and virtual application creation tasks.

Is this a web application with data ?

Yes

Web Application Compatibility Test

Pass

DBaaS compatibility Test

Pass

Fail

Fail

Result : Virtual Application
Reference virtual application virtual application with DBaaS creation tasks and RAD.

No

Is this a web application ?

Yes

Web Application Compatibility Test

Pass

Fail

Result : Virtual Application
Reference virtual application virtual application creation tasks and RAD.

No

Is this a database as a service workload ?

Yes

DBaaS Compatibility Test

Pass

Fail

Result : Virtual Application
Reference DBAS creation tasks.

Virtual System Creation Process

HV Image Compatibility Test

Pass

Fail

Mitigation:
Image Construction

Is your application install and configuration automated ?

Yes

Result : Virtual System
Reference virtual system creation tasks and examples with script packages

No

RAFW Compatibility Test

Pass

Result : Virtual System
Reference virtual system creation tasks, RAFW Image, RAFW Cloud Migration Doc.

Fail

Mitigation Process : takes information from compatibility test and recommends mitigation investment

6 Choosing appropriate deployment models © 2011 IBM Corporation

This flow chart details some high level steps for helping you choose a deployment model for existing applications web applications that are compatible with the Database as a service model are generally well-suited to the Virtual Application Deployment model. Applications that do not fit the Virtual Application model can be deployed on Virtual Systems, provided that they are compatible with the the Hypervisor Edition images, and that they can be configured automatically using script packages that already exist, or are generated by Rational Automation Framework for WebSphere. Applications that do not fit either model can still be run as custom virtual images, though you might want to consider migrating them to a configuration that is compatible with a different deployment model to fully use the benefits of IBM Workload Deployer. If you are developing a new application, it is recommended that you try to develop an application that fits the Virtual Application deployment model.

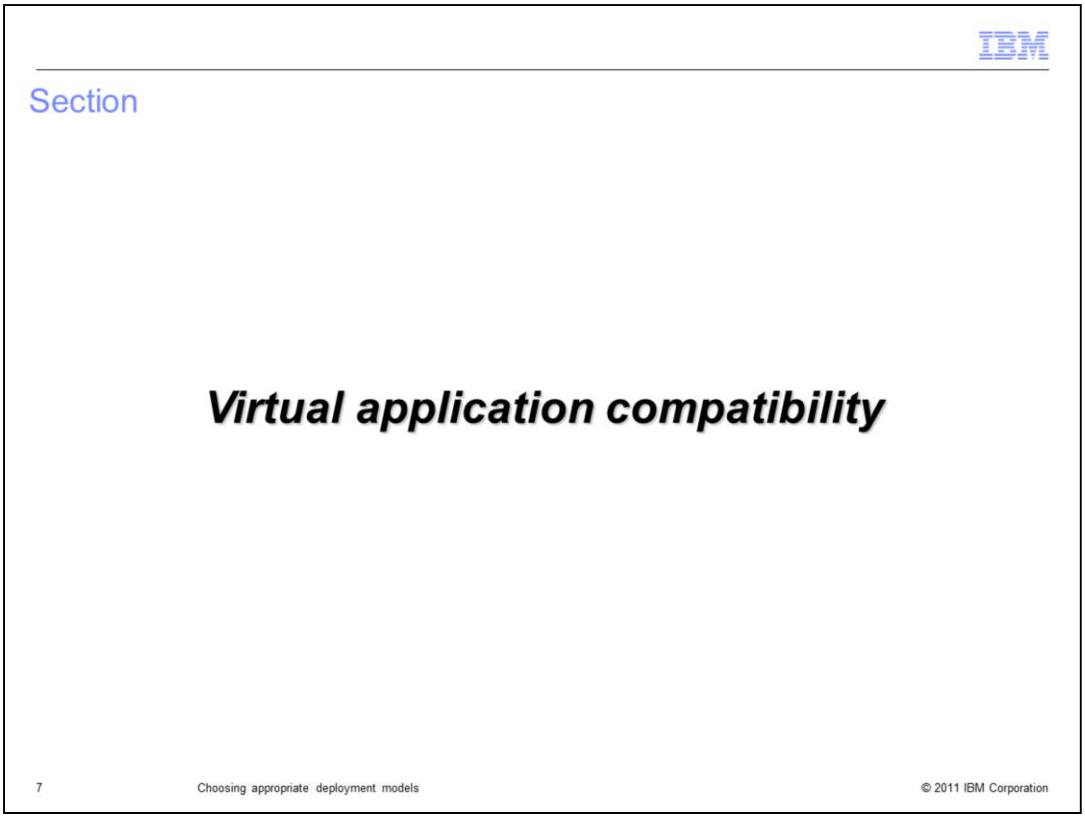

Section

Virtual application compatibility

7 Choosing appropriate deployment models © 2011 IBM Corporation

This section will discuss virtual application compatibility.

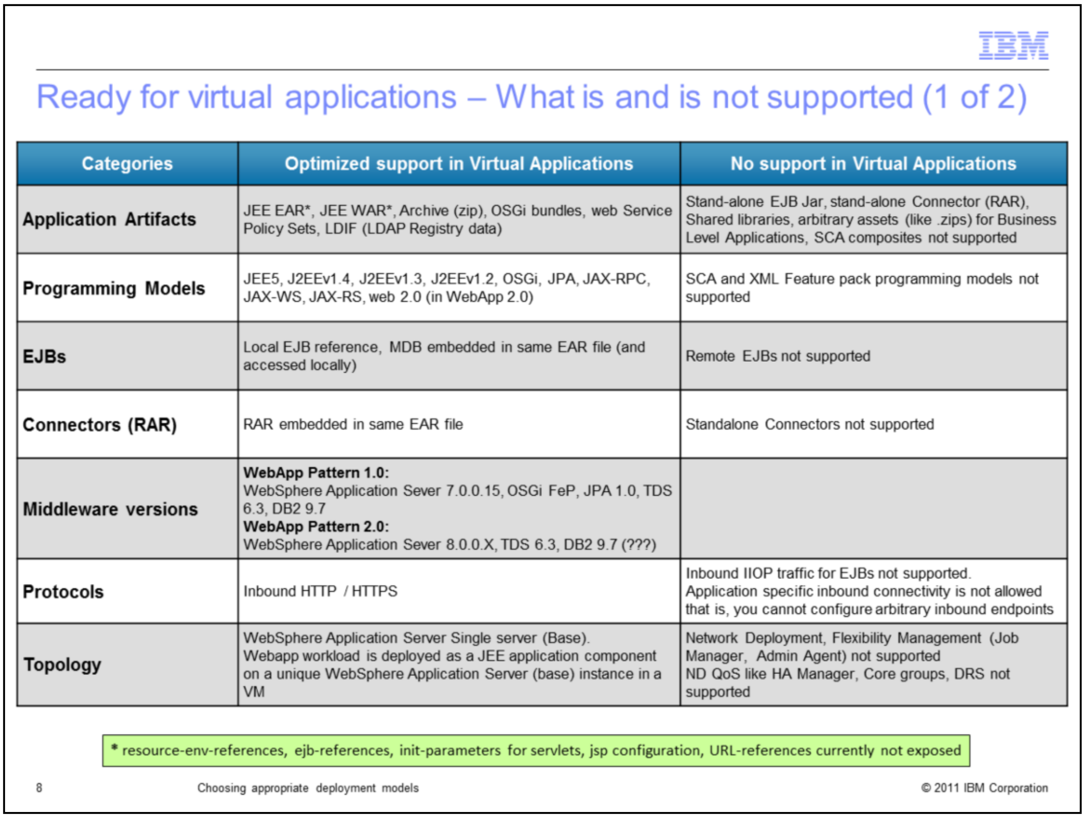

Ready for virtual applications – What is and is not supported (1 of 2)

Categories	Optimized support in Virtual Applications	No support in Virtual Applications
Application Artifacts	JEE EAR*, JEE WAR*, Archive (zip), OSGi bundles, web Service Policy Sets, LDIF (LDAP Registry data)	Stand-alone EJB Jar, stand-alone Connector (RAR), Shared libraries, arbitrary assets (like .zips) for Business Level Applications, SCA composites not supported
Programming Models	JEE5, J2EEv1.4, J2EEv1.3, J2EEv1.2, OSGi, JPA, JAX-RPC, JAX-WS, JAX-RS, web 2.0 (in WebApp 2.0)	SCA and XML Feature pack programming models not supported
EJBs	Local EJB reference, MDB embedded in same EAR file (and accessed locally)	Remote EJBs not supported
Connectors (RAR)	RAR embedded in same EAR file	Standalone Connectors not supported
Middleware versions	**WebApp Pattern 1.0:** WebSphere Application Sever 7.0.0.15, OSGi FeP, JPA 1.0, TDS 6.3, DB2 9.7 **WebApp Pattern 2.0:** WebSphere Application Sever 8.0.0.X, TDS 6.3, DB2 9.7 (???)	
Protocols	Inbound HTTP / HTTPS	Inbound IIOP traffic for EJBs not supported. Application specific inbound connectivity is not allowed that is, you cannot configure arbitrary inbound endpoints
Topology	WebSphere Application Server Single server (Base). Webapp workload is deployed as a JEE application component on a unique WebSphere Application Server (base) instance in a VM	Network Deployment, Flexibility Management (Job Manager, Admin Agent) not supported ND QoS like HA Manager, Core groups, DRS not supported

* resource-env-references, ejb-references, init-parameters for servlets, jsp configuration, URL-references currently not exposed

8 Choosing appropriate deployment models © 2011 IBM Corporation

The JEE application components deployed with Workload Deployer can only accept http/https as inbound traffic into the workload. This means it is not possible to have inbound iiop/EJB traffic and hence one cannot have EJBs hosted on Workload Deployer deployed workloads that need to be accessed from outside. It is also not possible to have a JEE application EAR to call EJBs in another JEE application where both are deployed in an Workload Deployer workload.

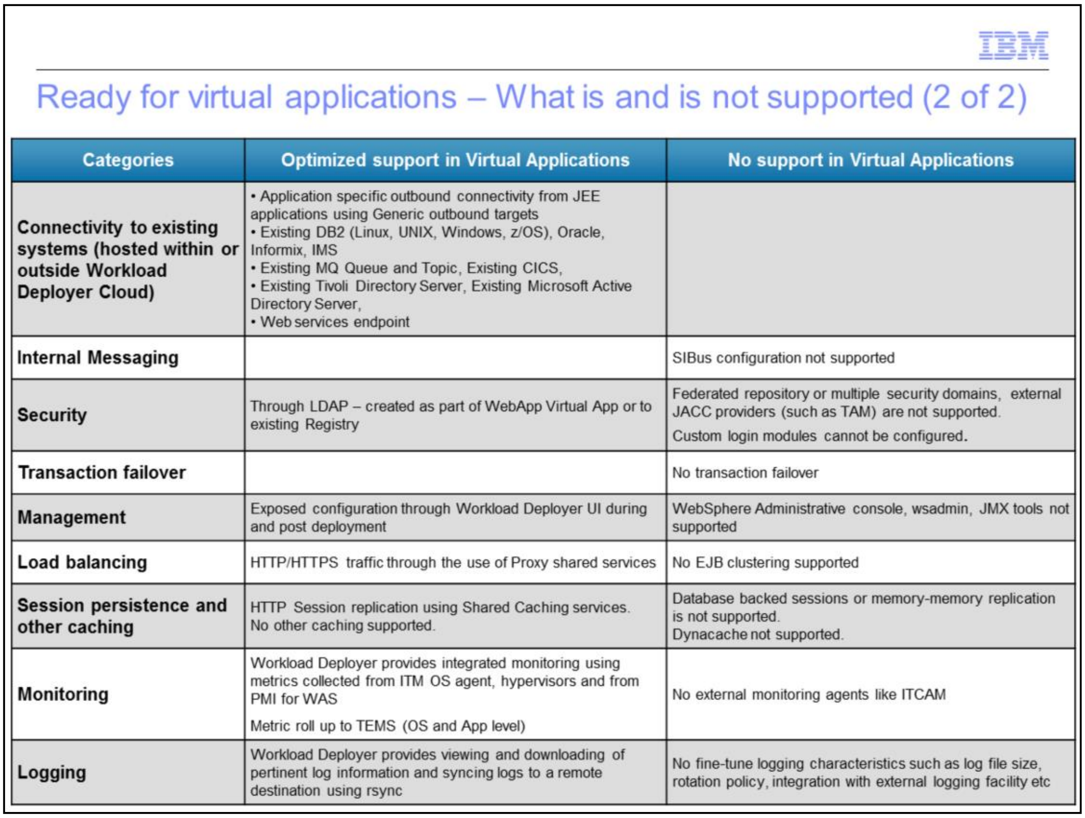

Ready for virtual applications – What is and is not supported (2 of 2)

Categories	Optimized support in Virtual Applications	No support in Virtual Applications
Connectivity to existing systems (hosted within or outside Workload Deployer Cloud)	• Application specific outbound connectivity from JEE applications using Generic outbound targets • Existing DB2 (Linux, UNIX, Windows, z/OS), Oracle, Informix, IMS • Existing MQ Queue and Topic, Existing CICS, • Existing Tivoli Directory Server, Existing Microsoft Active Directory Server, • Web services endpoint	
Internal Messaging		SIBus configuration not supported
Security	Through LDAP – created as part of WebApp Virtual App or to existing Registry	Federated repository or multiple security domains, external JACC providers (such as TAM) are not supported. Custom login modules cannot be configured.
Transaction failover		No transaction failover
Management	Exposed configuration through Workload Deployer UI during and post deployment	WebSphere Administrative console, wsadmin, JMX tools not supported
Load balancing	HTTP/HTTPS traffic through the use of Proxy shared services	No EJB clustering supported
Session persistence and other caching	HTTP Session replication using Shared Caching services. No other caching supported.	Database backed sessions or memory-memory replication is not supported. Dynacache not supported.
Monitoring	Workload Deployer provides integrated monitoring using metrics collected from ITM OS agent, hypervisors and from PMI for WAS Metric roll up to TEMS (OS and App level)	No external monitoring agents like ITCAM
Logging	Workload Deployer provides viewing and downloading of pertinent log information and syncing logs to a remote destination using rsync	No fine-tune logging characteristics such as log file size, rotation policy, integration with external logging facility etc

Virtual applications do not support the full range of services typically supported by WebSphere Application Server. This table shows what is and is not supported for virtual applications. For example, virtual applications do not support transaction failover, external JACC security providers, or external monitoring agents. Session persistence is not available through database persistence or memory-to-memory replication, but can be achieved using the shared caching service provided by IBM Workload Deployer.

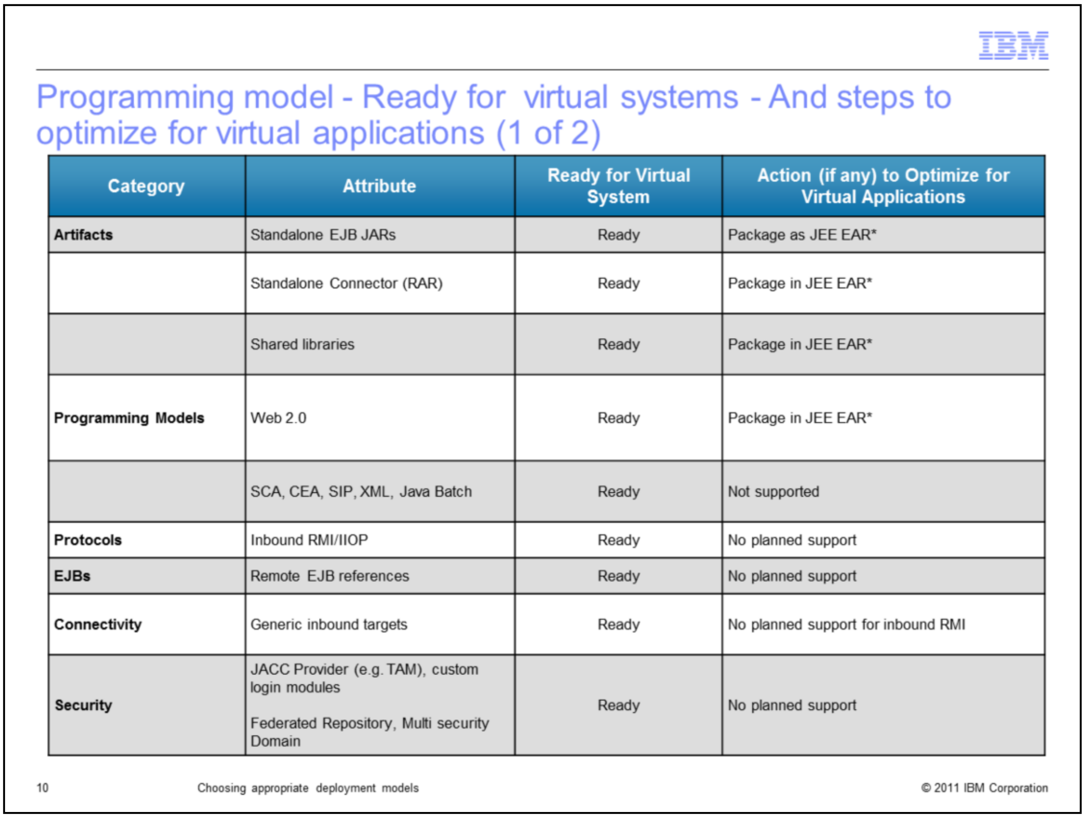

Programming model - Ready for virtual systems - And steps to optimize for virtual applications (1 of 2)

Category	Attribute	Ready for Virtual System	Action (if any) to Optimize for Virtual Applications
Artifacts	Standalone EJB JARs	Ready	Package as JEE EAR*
	Standalone Connector (RAR)	Ready	Package in JEE EAR*
	Shared libraries	Ready	Package in JEE EAR*
Programming Models	Web 2.0	Ready	Package in JEE EAR*
	SCA, CEA, SIP, XML, Java Batch	Ready	Not supported
Protocols	Inbound RMI/IIOP	Ready	No planned support
EJBs	Remote EJB references	Ready	No planned support
Connectivity	Generic inbound targets	Ready	No planned support for inbound RMI
Security	JACC Provider (e.g. TAM), custom login modules Federated Repository, Multi security Domain	Ready	No planned support

10 Choosing appropriate deployment models © 2011 IBM Corporation

There are several programming models that are not supported by the virtual application deployment model. Some programming models can be used by packaging the appropriate libraries in your EAR file. Others, are not supported at all, as shown on this chart.

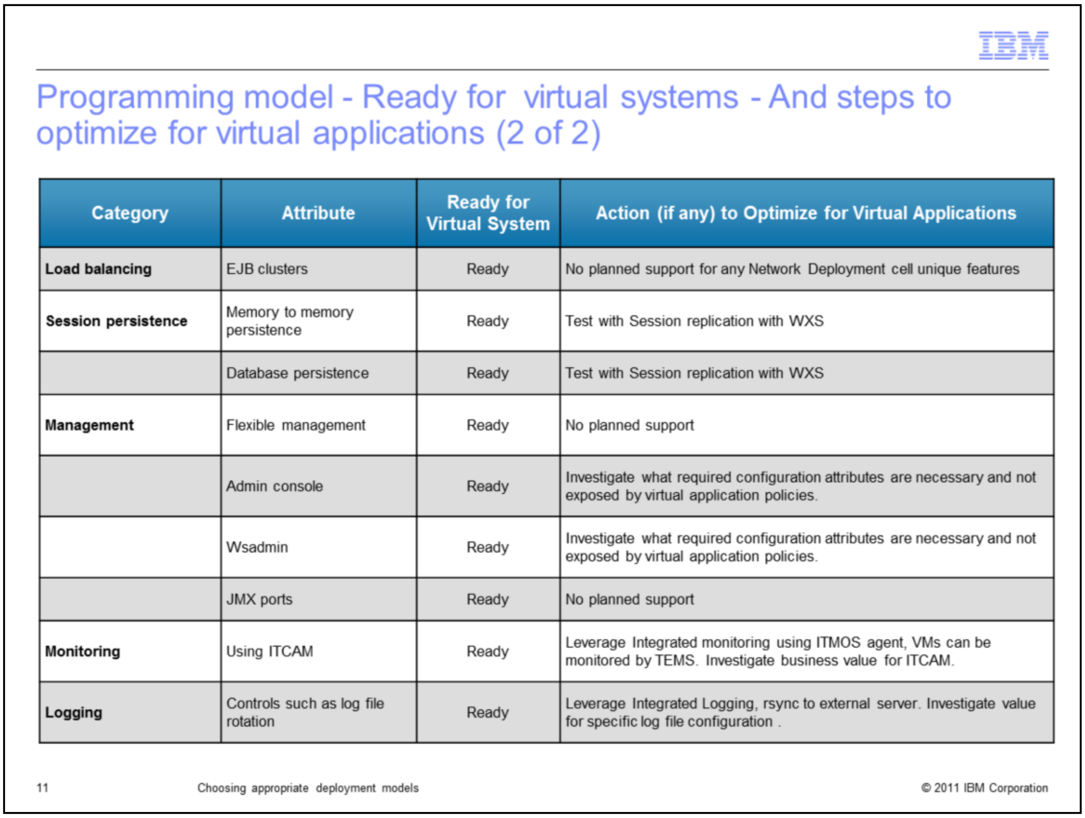

IBM

Programming model - Ready for virtual systems - And steps to optimize for virtual applications (2 of 2)

Category	Attribute	Ready for Virtual System	Action (if any) to Optimize for Virtual Applications
Load balancing	EJB clusters	Ready	No planned support for any Network Deployment cell unique features
Session persistence	Memory to memory persistence	Ready	Test with Session replication with WXS
	Database persistence	Ready	Test with Session replication with WXS
Management	Flexible management	Ready	No planned support
	Admin console	Ready	Investigate what required configuration attributes are necessary and not exposed by virtual application policies.
	Wsadmin	Ready	Investigate what required configuration attributes are necessary and not exposed by virtual application policies.
	JMX ports	Ready	No planned support
Monitoring	Using ITCAM	Ready	Leverage Integrated monitoring using ITMOS agent, VMs can be monitored by TEMS. Investigate business value for ITCAM.
Logging	Controls such as log file rotation	Ready	Leverage Integrated Logging, rsync to external server. Investigate value for specific log file configuration .

11 Choosing appropriate deployment models © 2011 IBM Corporation

Virtual applications also do not expose traditional management interfaces like the administrative console or wsadmin. You will need to determine if your application can be sufficiently configured using the settings and policies exposed by the Virtual Application Builder.

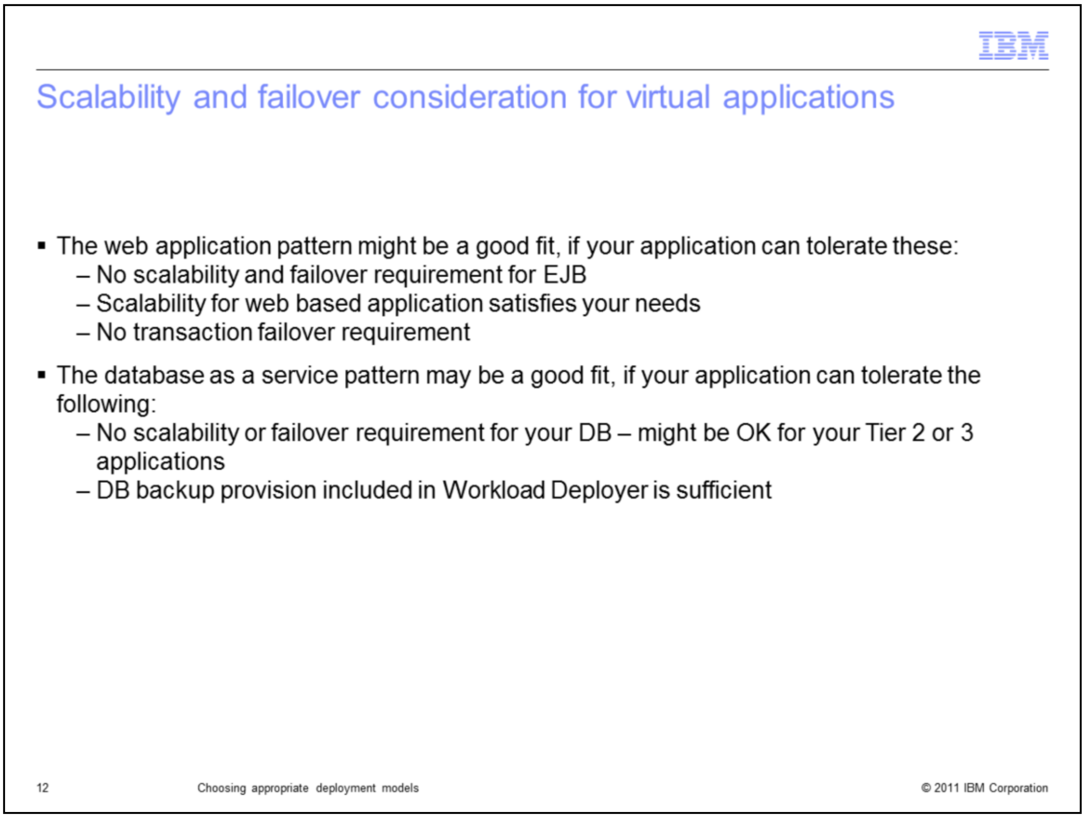

Scalability and failover consideration for virtual applications

- The web application pattern might be a good fit, if your application can tolerate these:
  - No scalability and failover requirement for EJB
  - Scalability for web based application satisfies your needs
  - No transaction failover requirement
- The database as a service pattern may be a good fit, if your application can tolerate the following:
  - No scalability or failover requirement for your DB – might be OK for your Tier 2 or 3 applications
  - DB backup provision included in Workload Deployer is sufficient

Scalability and failover requirements are important considerations for using the Virtual Application pattern types. The IBM web Application pattern type does not support scalability or failover for EJBs, and does not support transaction failover. As mentioned previously, sessions can be retained during server failure by sharing them using the shared cache service. The database as a service pattern type does not support scalability or failover, so database availability is not guaranteed. IBM Workload Deployer does support backup and restore of databases in this pattern type, and if that is sufficient, you might want to consider the database as a service pattern type. For web applications that require better database availability, you can configure your virtual application to use an externally configured database server, though this introduces additional complexity, so consider the trade-offs in determining the needs of your application.

Section

Virtual systems deployment challenges

13 Choosing appropriate deployment models

This section covers the challenges related to deploying virtual systems.

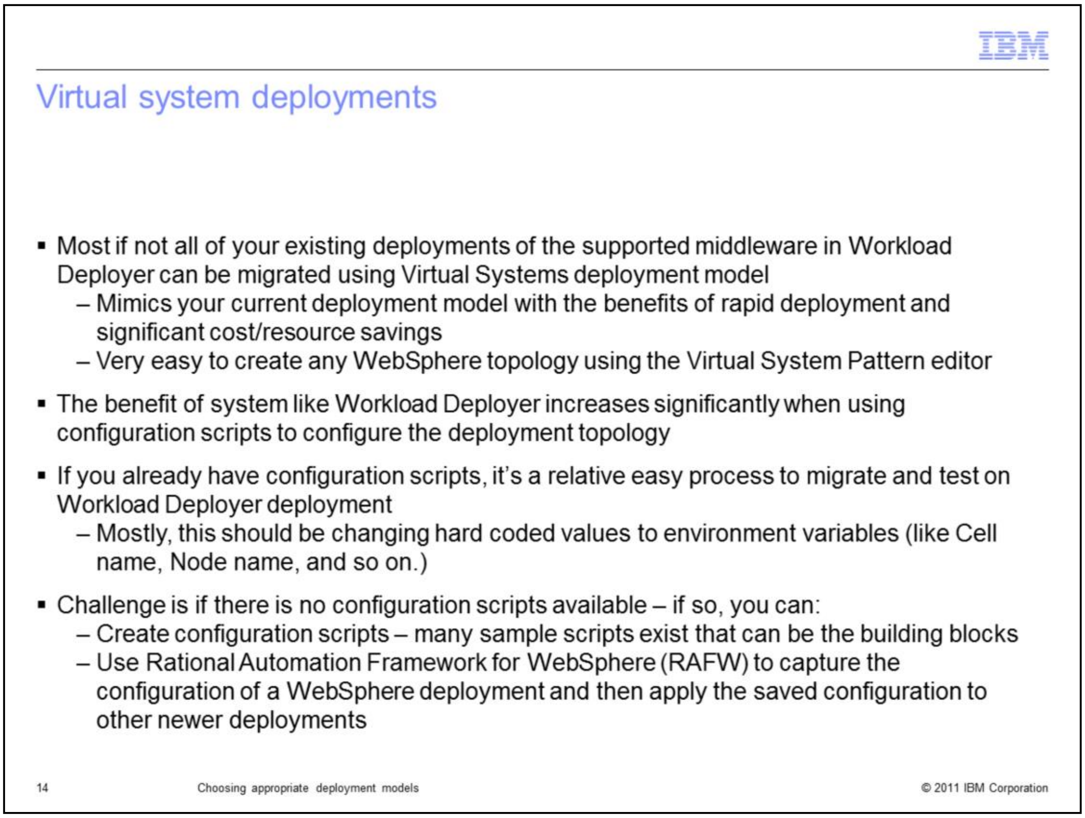

Virtual system deployments

- Most if not all of your existing deployments of the supported middleware in Workload Deployer can be migrated using Virtual Systems deployment model
  - Mimics your current deployment model with the benefits of rapid deployment and significant cost/resource savings
  - Very easy to create any WebSphere topology using the Virtual System Pattern editor
- The benefit of system like Workload Deployer increases significantly when using configuration scripts to configure the deployment topology
- If you already have configuration scripts, it’s a relative easy process to migrate and test on Workload Deployer deployment
  - Mostly, this should be changing hard coded values to environment variables (like Cell name, Node name, and so on.)
- Challenge is if there is no configuration scripts available – if so, you can:
  - Create configuration scripts – many sample scripts exist that can be the building blocks
  - Use Rational Automation Framework for WebSphere (RAFW) to capture the configuration of a WebSphere deployment and then apply the saved configuration to other newer deployments

Most WebSphere Application Server topologies can be replicated in a Virtual System deployment. If your systems are already configured using scripting, the amount of work required to migrate them to a Virtual System environment is minimal. If not, you will want to strongly consider scripting your configuration, since scripting will enable you to take advantage of rapid provisioning and deployment.

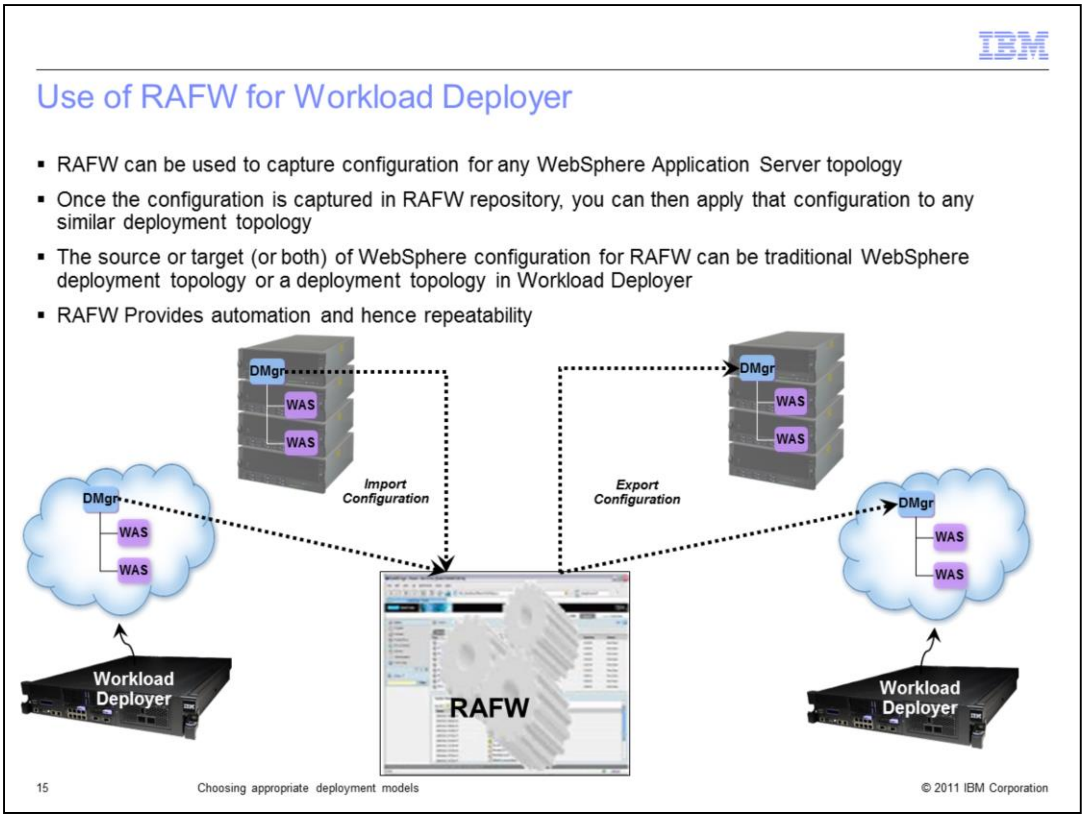

Use of RAFW for Workload Deployer

- RAFW can be used to capture configuration for any WebSphere Application Server topology
- Once the configuration is captured in RAFW repository, you can then apply that configuration to any similar deployment topology
- The source or target (or both) of WebSphere configuration for RAFW can be traditional WebSphere deployment topology or a deployment topology in Workload Deployer
- RAFW Provides automation and hence repeatability

15 Choosing appropriate deployment models © 2011 IBM Corporation

Rational Automation Framework for WebSphere (RAFW) can be used to capture configuration information for an existing WebSphere Application Server topology. Once the configuration information has been captured, RAFW can be used to create scripts to automate configuration of the environment. Those scripts can be used to create script packages for deploying the topology in a repeatable and automated way as virtual systems.

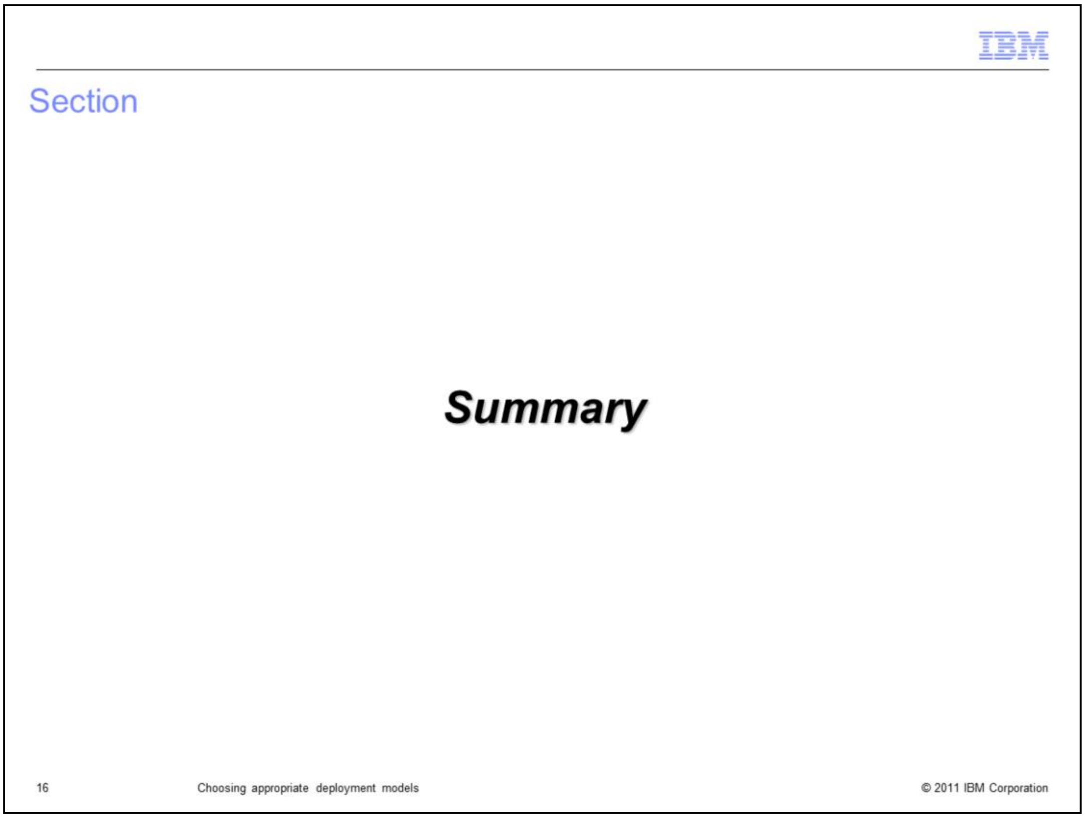

Section

Summary

This section will summarize the presentation.

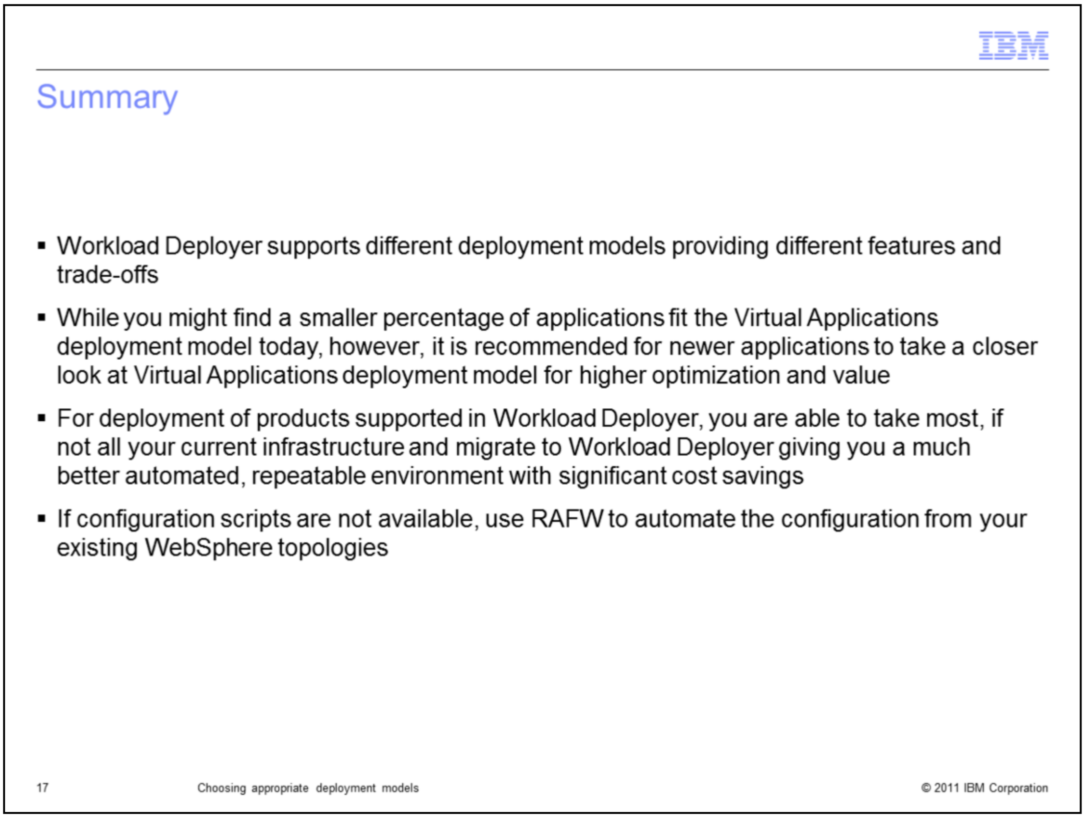

Summary

- Workload Deployer supports different deployment models providing different features and trade-offs
- While you might find a smaller percentage of applications fit the Virtual Applications deployment model today, however, it is recommended for newer applications to take a closer look at Virtual Applications deployment model for higher optimization and value
- For deployment of products supported in Workload Deployer, you are able to take most, if not all your current infrastructure and migrate to Workload Deployer giving you a much better automated, repeatable environment with significant cost savings
- If configuration scripts are not available, use RAFW to automate the configuration from your existing WebSphere topologies

17 Choosing appropriate deployment models © 2011 IBM Corporation

The deployment models supported by IBM Workload Deployer allow for several types of solutions to be deployed in the virtualized environment. Virtual applications offer the lowest total cost of ownership, and can take the most advantage of automation in IBM Workload Deployer. Not all applications are suited to deployment as virtual applications, since virtual applications do not support the full range of features supported by WebSphere Application Server. Those applications can typically be deployed as virtual systems, which focus more on the topology than the application. Virtual systems require configuration scripts to be available to configure the application and related resources. If those scripts do not already exist, you can create them using Rational Automation Framework for WebSphere. If you are developing new applications, it is recommended to see if you can build your application to fit the virtual application deployment model, since they offer the most value in a virtualized environment.

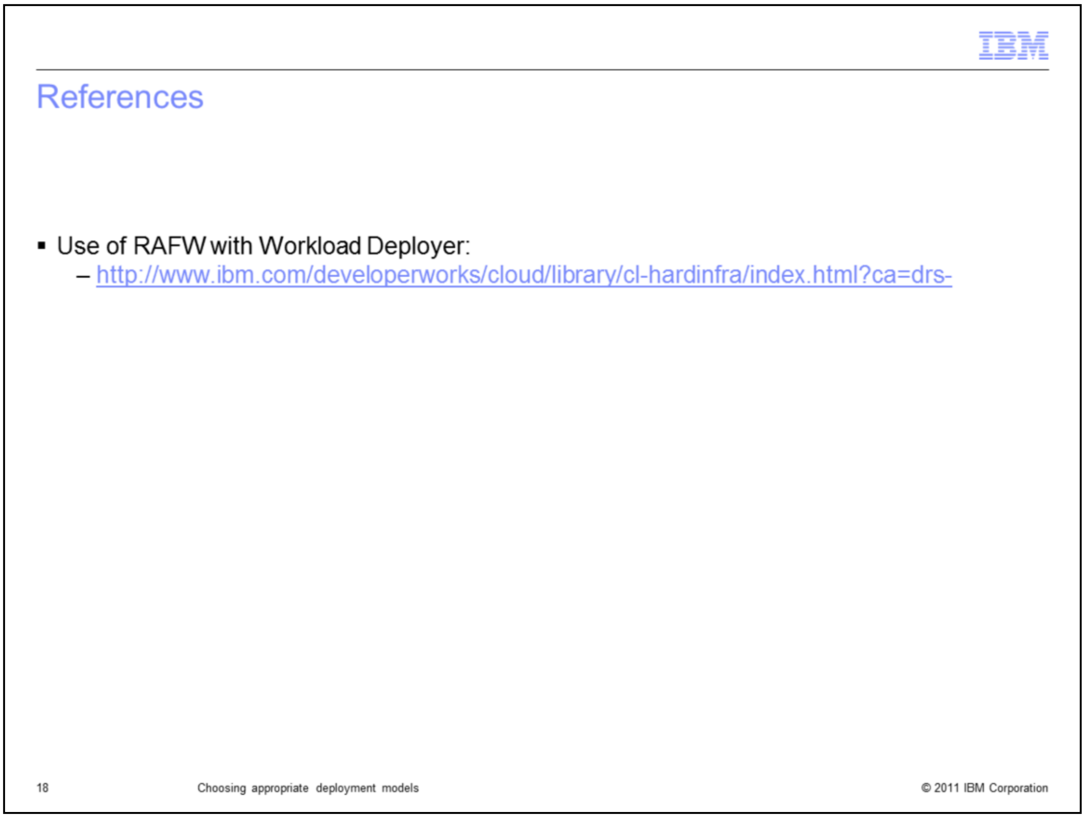

References

- Use of RAFW with Workload Deployer:
  - http://www.ibm.com/developerworks/cloud/library/cl-hardinfra/index.html?ca=drs-

18 Choosing appropriate deployment models © 2011 IBM Corporation

This link details using Rational Automation Framework for WebSphere with IBM Workload Deployer.

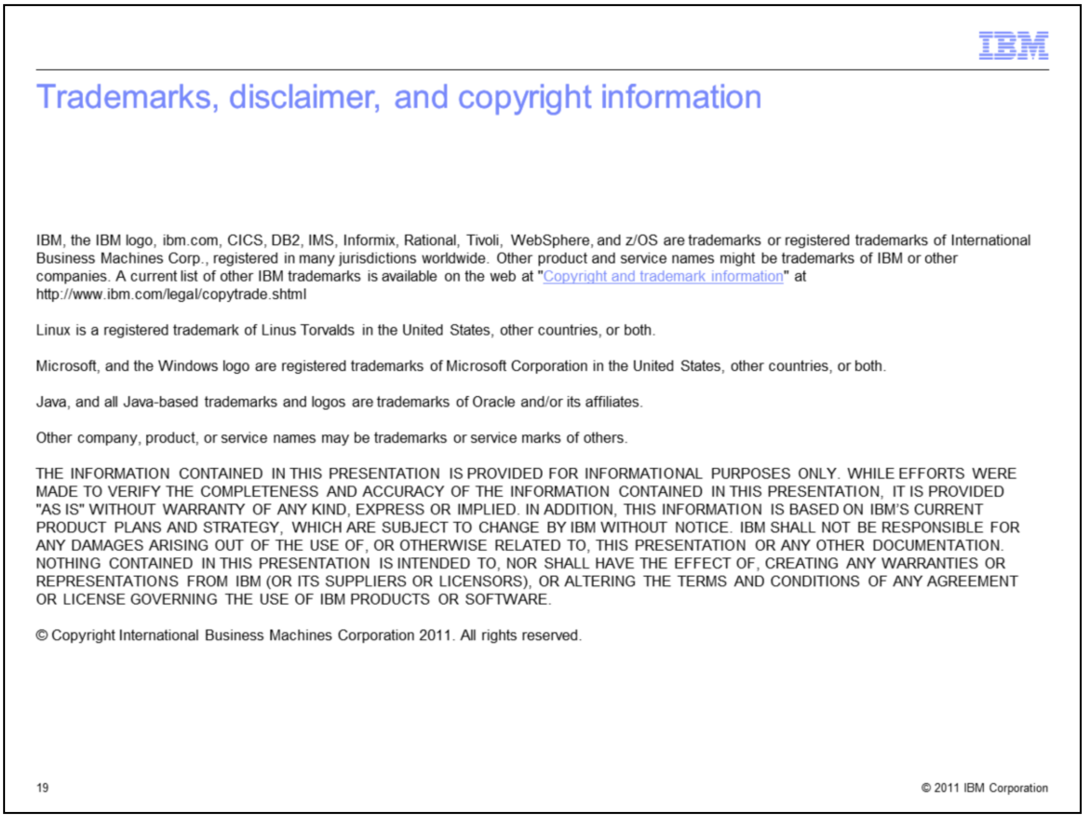

Trademarks, disclaimer, and copyright information

IBM, the IBM logo, ibm.com, CICS, DB2, IMS, Informix, Rational, Tivoli, WebSphere, and z/OS 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" 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, and the Windows logo are registered trademarks of Microsoft Corporation in the United States, other countries, or both.

Java, and all Java-based trademarks and logos are trademarks of Oracle and/or its affiliates.

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.