



IBM Software Group

IBM® WebSphere® Application Server V6

Configuration Archives



@business on demand.

© 2005 IBM Corporation
Updated April 28, 2005

This presentation will focus on Configuration Archives, a new feature in WebSphere Application Server V6 that introduces a declarative approach to modify the configuration of WebSphere Application Server.

Agenda

- Definition of Configuration Archives
- Configuration Archive Life Cycle
- More Details....
 - ▶ Configuration Archive Export
 - ▶ Configuration Archive Import
- Example of using Configuration Archives



In this presentation you will learn the definition of WebSphere configuration archives, the life cycle of WebSphere configuration archives, some additional detail on the export and import operations, and an example of using WebSphere configuration archives.

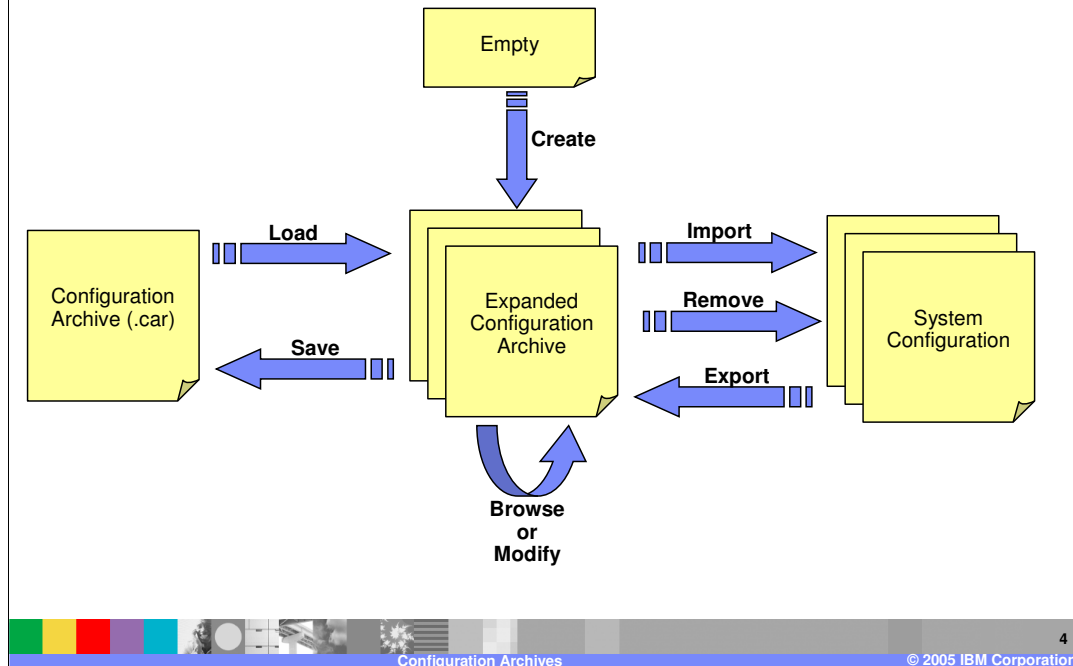
Configuration Archive

- It's a set of complete or subset of WebSphere Application Server configuration
 - ▶ Basically the same as the real WebSphere Application Server's configuration except they may be partial
 - ▶ The configuration information is virtualized to make it portable
 - Removes any install specific information, like the host name
- Used to import/export configurations into or from WebSphere Application Server
 - ▶ Allows declarative way to modify WebSphere Application Server configurations



WebSphere configuration archive introduces a declarative approach to modifying configurations. Instead of having you modify the system configuration through operations, you can now define a configuration in an archive and modify the configuration by importing an archive into the system. A configuration archive is an archive that consists of a complete or a subset of WebSphere configurations. The layout of the configuration archive is basically the same as the real WebSphere configuration except that it may be incomplete. You may export some configurations in the system to a configuration archive. Exporting configurations to a configuration archive virtualizes the configuration and disassociates it from the real physical topology of the system. This makes the configuration portable across systems. Later on, users can import the configuration archive to the same system or other similar systems. Importing a configuration archive creates the configurations defined in the configuration archive to the system. This allows users to set up configurations quickly and easily without learning any configuration operation syntax and semantics.

Configuration Archive: Life Cycle



The picture above illustrates the life cycle of a configuration archive and operations that can be applied on it. A configuration archive is basically a zipped file. You may use the extension of “car” for archive files, but it is not required. Expanded format for configuration archives is also supported. While an archive package is easier to be carried over across systems, it is more efficient to operate on the configuration archive if it is in an expanded format. Thus, the configuration archive is loaded into a directory in an expanded format before import operation. An empty configuration archive can be created in a directory. Loading a configuration archive into directory basically expands the archive while saving a configuration archive is the opposite process which packages the expanded configuration archive in the directory into a single file archive. Once a configuration archive is loaded into a directory, users may browse or modify the configuration archive as they wish. Configurations in the system can be exported into configuration archive while configuration archives can be imported into the system.

Configuration Archive: Export

- Export converts the real configuration into templates
 - ▶ Breaks any relationship of configuration in the archive and the system configuration
- Clusters
 - ▶ When an individual cluster member is exported, the reference to cluster is stripped off
 - ▶ When a cluster is exported, the relationship between the cluster and its members is preserved
- Configuration may also refer to system resources outside scope: database driver library, database name, etc.
- Limited types of export targets
 - ▶ server and wasprofile
- Examples
 - ▶ `$AdminTask exportServer { -archive c:\myServer.car -nodeName node1 -serverName server1}`
 - ▶ `$AdminTask exportWasprofile {-archive c:\myCell.car}`



When you export configurations from a real WebSphere configuration repository, the configuration archive feature virtualizes the configuration and will disassociate the configuration from a real WebSphere configuration to make the exported configuration portable from system to system. At a high level, the export function converts the real WebSphere configuration into template that breaks up any relation of the configuration stored in the archive and the system configuration. There are a couple of things to consider when you are working with WebSphere configuration archives in a cluster environment. The first is when you export an individual cluster member the relationship of the cluster member to the cluster is eliminated. However, when a complete cluster is exported its relationship to its cluster members is preserved. Another thing to point out is references to resources outside in the system configuration are also saved in the configuration archive so you would need to go through the manual update of defining those resources.

In terms of the export operation there are only two target types that are supported. One is a server and the other is WebSphere profile. This means that you can either export a particular WebSphere server instance, like server1, or a complete WebSphere profile that could contain more than one server instance.

Provided in the last bullet of this chart is some wsadmin operational tasks to perform the export function on a server and a WebSphere profile.

Configuration Archive: Import

- Importing a configuration archive is the process of updating/creating system configuration with contents in configuration archive
- Target configuration object type either matches source configuration object or is a parent
 - ▶ server to server (same type) – configuration merged
 - ▶ server to node (parent) – new object type created for configuration archive
 - In this case, configuration must be complete as new server created
- Import operation also compares node properties between source and target
 - ▶ Example: If configuration archive is for base, it may not work for WebSphere Business Integration Server Foundation
- Examples
 - ▶ `$AdminTask importServer {- archive c:\myServer.car [-nodeInArchive node1] [-serverInArchive server1] [--nodeName node1] [--serverName server1] }`
 - ▶ `$AdminTask importWasprofile {- archive c:\myCell.ca}`



When you import a configuration archive a process is run that lays out the configuration defined in the archive to the real WebSphere configuration repository. When you use the import function, you select a configuration object defined in the configuration archive as the source and select a configuration object on the system as the target. The target configuration object can either match the source configuration object or be its parent. For example is the source and target configuration object are both servers the configuration is merged. This would be the case if you exported a server object type configuration, updated or performed other administrative tasks to it, than imported it back into the server. An example of the parent style would be if you imported a configuration archive that represents a server object type into a node object type. This would create a new server instance for the configuration archive you are importing. One other thing done when importing a WebSphere configuration archive is the node properties are compared. For example, trying to import a base server configuration into a WebSphere Business Integration Server Foundation package would not work. In the last bullet you are provided a couple of examples of importing a server and WebSphere profile archive.

Configuration Archive: Example

- Example:
 - ▶ Users export the Configuration Archive
 - ▶ Modify the configurations
 - ▶ The exported Configuration Archive can be used to import on the same or other systems
- Allows users to set up configurations quickly and easily without learning any configuration operation syntax and semantics
- Provides powerful way to manage the configurations of a large number of similar systems
 - ▶ Within the cell boundary or across cell boundary, such as a server farm or retail scenario



Here you can see an example of using WebSphere configuration archives. First, you could export a real WebSphere configuration to a configuration archive. Configuration tasks can then be done on that archive to meet a specific situation. Remember that exporting the configuration will virtualized the configuration so it has no association to the physical system. Once the configuration tasks are complete, that same archive can then be imported and used as a real WebSphere configuration. This allows users the ability to set up configurations quickly and easily without the overhead of learning any specific configuration operation syntax or semantics. WebSphere configuration archives provides a powerful way manage and exchange WebSphere configuration across different systems.

Section

Summary and Reference

Summary

- Configuration Archives are a set of complete or subset of WebSphere Application Server configuration
- Can be used to create or modify configurations



In this presentation you learned what are WebSphere configuration archive and the benefits they provide to the WebSphere administrator. An explanation of the overall life cycle, or operation flow of WebSphere configuration archives was discussed. You also learned about both the export and import functions when working with WebSphere configuration archives. The presentation was wrapped up with an example of using WebSphere configuration archives.

Trademarks, Copyrights, and Disclaimers

The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:

IBM	CICS	IMS	MQSeries	Tivoli
IBM (logo)	Cloudscape	Informix	OS/390	WebSphere
eIogo business	DB2	Series	OS/400	xSeries
AIX	DB2 Universal Database	Lotus	pSeries	zSeries

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

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

Intel, ActionMedia, LANDesk, MMX, Pentium and ProShare are trademarks of Intel 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.

Linux is a registered trademark of Linus Torvalds.

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

Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. This document could include technical inaccuracies or typographical errors. IBM may make improvements and/or changes in the product(s) and/or program(s) described herein at any time without notice. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectual property rights, may be used instead.

Information is provided "AS IS" without warranty of any kind. THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted, if at all, according to the terms and conditions of the agreements (e.g., IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products in connection with this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. IBM makes no representations or warranties, express or implied, regarding non-IBM products and services.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprocessing in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

© Copyright International Business Machines Corporation 2004. All rights reserved.

Note to U.S. Government Users - Documentation related to restricted rights-Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract and IBM Corp.