



IBM Software Group

WebSphere® Extended Deployment V6.1

WebSphere Virtual Enterprise

Formerly Operations Optimization

Middleware descriptors and discovery



@business on demand.

© 2008 IBM Corporation
Updated June 16, 2008

WebSphere Extended Deployment version 6.1 provides automation that will discover selected middleware server installations and generate server representations within the administrative console.

Middleware descriptors

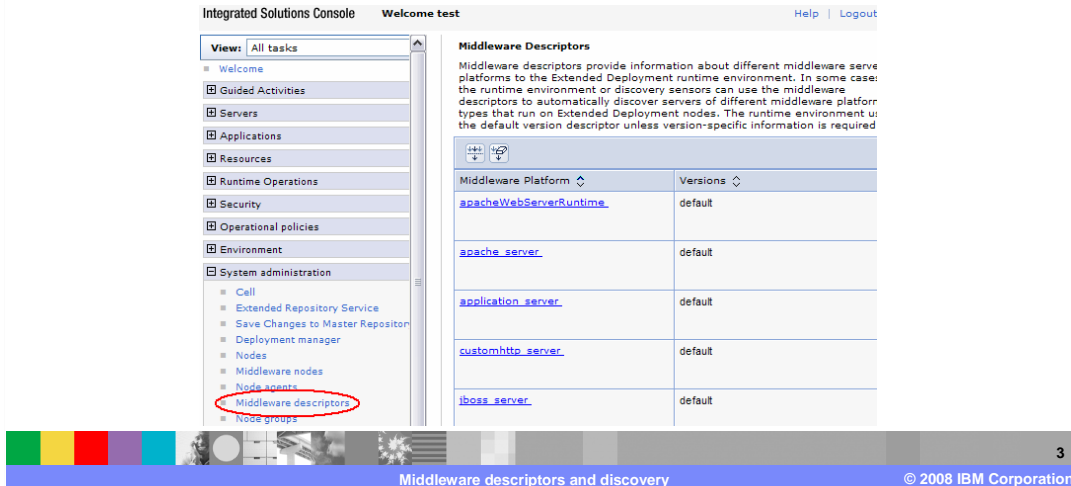
- Provide information about different middleware server platforms
 - ▶ Installation locations
 - ▶ Session affinity identifiers



WebSphere Extended Deployment V6.1.0 introduces middleware descriptors, which provide high level information about different middleware platform types. This information is used for automated discovery support which enables Extended Deployment to locate installations of various middleware servers. Once the server is located, the middleware descriptor is primarily used by the on demand router for routing and session affinity.

Middleware descriptors collection page

- For assisted and complete life cycle servers
 - ▶ Provides default settings for server installations
 - ▶ Allows user to customize settings for specific or non-default installations



Integrated Solutions Console Welcome test Help | Logout

View: All tasks

- Welcome
- Guided Activities
- Servers
- Applications
- Resources
- Runtime Operations
- Security
- Operational policies
- Environment
- System administration
 - Cell
 - Extended Repository Service
 - Save Changes to Master Repository
 - Deployment manager
 - Nodes
 - Middleware nodes
 - Node agents
 - **Middleware descriptors**
 - Node groups

Middleware Descriptors

Middleware descriptors provide information about different middleware server platforms to the Extended Deployment runtime environment. In some cases, the runtime environment or discovery sensors can use the middleware descriptors to automatically discover servers of different middleware platform types that run on Extended Deployment nodes. The runtime environment uses the default version descriptor unless version-specific information is required.

Middleware Platform	Versions
apacheWebServerRuntime	default
apache_server	default
application_server	default
customhttp_server	default
iboss_server	default

3

Middleware descriptors and discovery © 2008 IBM Corporation

WebSphere Extended Deployment V6.1 provides access to middleware descriptors through the administrative console. The Middleware descriptor collection page provides default settings and installation information for supported middleware server platforms. It also allows the administrator to customize these settings for specific or non-default middleware server installations. The information on this page enables Extended Deployment to locate installations and to provide session affinity definitions for the various middleware server platforms.

Middleware descriptor properties

- Used by on demand routers
- sessionID
 - ▶ Cookie
 - ▶ URL rewrite
 - ▶ Clone information
- Affinity mode
 - ▶ Passive
 - ▶ Active-conditional

The screenshot shows a configuration window titled "Configuration" with a tab for "General Properties". Under "Session Affinity Descriptor", the following properties are visible:

- Version: default
- Learn clone IDs: true
- Cookie names: WSJSESSIONID, JSESSIONID, SSLJSESSIONID
- URL rewrite: jsessionid
- Clone ID separator: :
- Alternate clone ID separator: +
- Affinity mode: passive



The middleware platform descriptors can be viewed by clicking on the middleware descriptor name from the previous administrative console panel.

When the on demand router (ODR) processes a request, it obtains the session affinity descriptor policy of the cluster to which the server belongs in order to properly route the request.

The session affinity descriptor properties for a middleware server are defined here.

The on demand router has the ability to parse clone IDs from the Set-Cookie header in a response as it flows back to the servlet. The “learn clone IDs” property must be set to true if you are routing to servers with session cookies that contain server clone information.

The cookie names field tells the on demand router where the middleware server stores the session ID in the HTTP reply. If the client browser has disabled cookies, the URL rewrite field can be used to find the session ID in the URL (Uniform Resource Locator). Both the cookie name and URL rewrite fields can contain multiple entries separated by commas.

The “clone ID separator” specifies the delimiter that is used by the on demand router when extracting clone IDs from the session information. For servers that support more than one delimiter character use the “alternate clone ID separator” field.

“Affinity mode” works with the “**Learn clone IDs**” value to determine how the session affinity descriptors are used. When **Passive** affinity mode is specified, the server clone information is obtained from the JSESSIONID cookie. When **Active-conditional** affinity mode is specified with a **Learn clone IDs** value of *false*, the server clone identification is generated internally and is set in the WSJSESSIONID cookie.

(For discovered servers, this information is read from the server configuration.)

Middleware discovery service

- **Discovery support**
 - ▶ Runs in deployment manager and middleware agent
 - Depends on “Middleware descriptors” to function properly
 - Runs at configured intervals
 - Can be turned off
 - ▶ Automatically discover installed middleware servers
 - ▶ Create server representations
 - Apache HTTP server
 - PHP
 - **WebSphere Application Server Community Edition** 
 - ▶ Not currently supported for other middleware server platforms



WebSphere Extended Deployment node agents and middleware agents include a discovery service that uses information from the middleware descriptors to automatically discover installed instances for certain middleware servers. The discovery service runs periodically based on the “Discovery interval” configured on the “Middleware descriptors” panel. Once the discovery service finds a known server installation, it will create an assisted life cycle server representation within the Extended Deployment cell configuration.

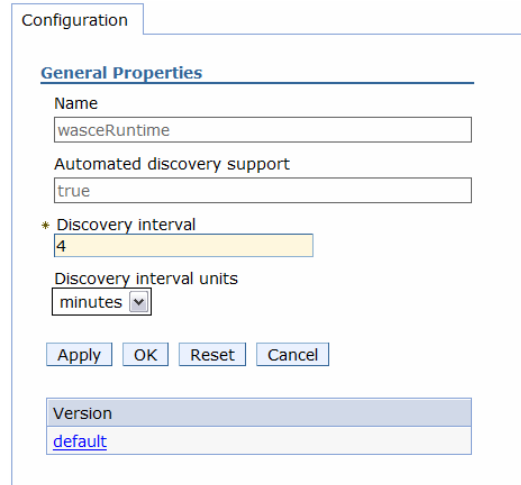
As an “assisted life cycle” server, deleting the representation does not actually delete the server from the remote installation. In fact, if discovery remains enabled, it will rebuild the server representation when the discovery service next runs.

WebSphere Extended Deployment V6.1 can automatically discover PHP servers and Apache Web Servers. V6.1.0.1 adds discovery for existing WebSphere Application Server Community Edition V2 and above servers.

Note that in V6.1.0.1 you can also define full life cycle WebSphere Community Edition servers within the administrative console. This means that you can both create and remove WebSphere Community Edition server instances using the WebSphere Extended Deployment administrative console. But server representations that result from the “discovery” of a pre-existing server, or that result from the manual definition of a pre-existing server, function as assisted life cycle servers.

Configuring discovery service

- Discovery support indicator
- Discovery interval



Configuration

General Properties

Name
wasceRuntime

Automated discovery support
true

* Discovery interval
4

Discovery interval units
minutes

Apply OK Reset Cancel

Version
default

The runtime environment discovery sensors can use middleware descriptors to automatically discover servers of different middleware platform types that run on Extended Deployment nodes.


Here are some additional middleware descriptor properties that can be set to facilitate discovery.

In version 6.1 the name and “Automated discovery support” fields are not editable. If Automated discover support is “true” then you can set the discovery interval. This value specifies how often the middleware agent runs discovery for this platform type. If the value is “-1” the discovery is disabled for this platform. If the value is zero the discovery service will only run when the agent starts. For any value greater than zero the discovery service will run at agent startup and at the specified interval. The specified discovery interval can be seconds, minutes, hours, or days.

The runtime environment uses the default version descriptor unless version-specific information is required.

Configuring discovery

- Extended Deployment must know the root location of WebSphere Community Edition installation
- Provided by “default” setting for wasceRuntime within Middleware descriptors page



Discovery can only function if the information in the Middleware descriptors for the associated server product is accurate. For example, for WebSphere Community Edition servers, you must ensure that the default settings associated with the WebSphere Community Edition Runtime entry within the Middleware descriptors page accurately reflects the possible installation locations for WebSphere Application Server Community Edition on the remote nodes. Default values are provided but can be modified for customized installations, such as a non-default directory location. The “root location” can specify multiple directory locations to check. For Windows installations the directories are separated by semi-colons; for UNIX systems the directories are separated by colons.

A typical example of the use the Middleware descriptors page is shown here for a WebSphere Application Server Community Edition installation. Using the Windows installation location as an example, the default setting must be modified within WebSphere Extended Deployment because the WebSphere Community Edition was not installed in the default location on the target server platform.

If the correct installation location is not specified, discovery will not be able to find pre-existing servers and definitions of new server instances on the WebSphere Community Edition target platform will not operate properly.

Summary

- WebSphere Extended Deployment V6.1 provides enhanced middleware support
 - ▶ Middleware descriptors
 - ▶ Automated discovery support



WebSphere Extended Deployment version 6.1 extends middleware support to various types of middleware platform servers. To do this, middleware descriptors provide information about different middleware platform types including middleware discovery intervals, enablement, and installation information. When enabled, the automated discovery service can run for Apache, WebSphere Application Server V5.1, V6.0, or Community Edition, and PHP runtimes that you have configured. These types of middleware servers and applications can then be represented within Extended Deployment so that many of the Extended Deployment control, routing, traffic shaping, and monitoring facilities can be used to efficiently manage a more diverse system environment.

Feedback

Your feedback is valuable

You can help improve the quality of IBM Education Assistant content to better meet your needs by providing feedback.

- Did you find this module useful?
- Did it help you solve a problem or answer a question?
- Do you have suggestions for improvements?

Click to send e-mail feedback:

mailto:iea@us.ibm.com?subject=Feedback_about_XD61_MiddlewareDiscovery.ppt

This module is also available in PDF format at: ../XD61_MiddlewareDiscovery.pdf



You can help improve the quality of IBM Education Assistant content by providing feedback.

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 WebSphere

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 or changes in the products or programs 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 (for example, 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 multiprogramming 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 2008. 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.

