



IBM Software Group

IBM® WebSphere® Application Server V6.0.2

Proxy Server



@business on demand.


© 2005 IBM Corporation
Updated July 7, 2005

This presentation will focus on the WebSphere Application Server V6.0.2 proxy server component.

Goals

- The goal is to understand the WebSphere Application Server V6.0.2 proxy server component

- Prerequisite:
 - ▶ None



The goal of this presentation is to explain the new proxy server component that is enabled with the V6.0.2 maintenance pack.

Agenda


- Proxy Server Big Picture
- Proxy Server Capabilities
 - ▶ Configuration
 - ▶ Routing
 - ▶ Caching
- Common Topologies
 - ▶ Basic Concerns
 - ▶ Including a Load Balancer
- Creating a Proxy Server



This presentation will begin by explaining the concepts behind a proxy server and the various proxy server solutions that are available for WebSphere Application Server. It will then go into detail on the capabilities provided by the new proxy server component. It will also discuss common topologies that would be appropriate for the new proxy server. And also explain how to create a proxy server after installing the 6.0.2 maintenance pack.

Section

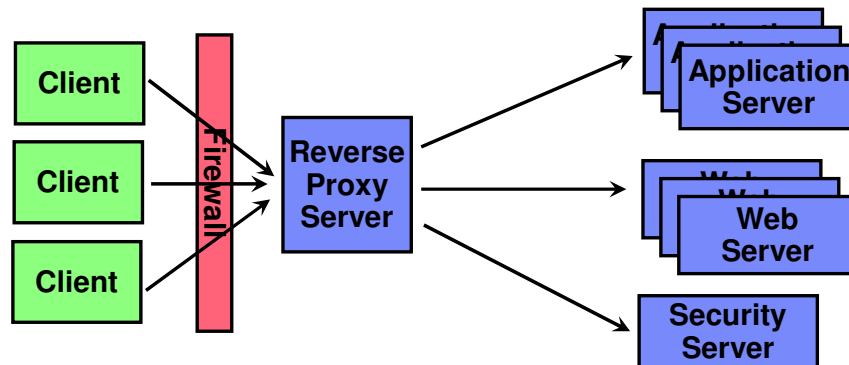
Proxy Server Big Picture



This section will detail the basic concepts of the proxy server.

Definition of Reverse Proxy Servers

- Reverse Proxy Servers (RPS) act as the initial point of entry, after a firewall, for requests entering an enterprise environment



A reverse proxy server acts as the initial point of entry for client requests entering your environment after a firewall. This picture illustrates this concept, showing client requests being sent through a firewall to a proxy server. The proxy server can then route those requests to a number of servers within an enterprise environment. This picture is somewhat incomplete, in that most situations will actually involve a second firewall on the other side of the proxy server, thus forming a secured area or DMZ. Often proxy server solutions are located within these DMZs.

Proxy Servers in V6.0.2

- WebSphere V6.0.2 adds the capability to define a Proxy Server within an ND environment
 - ▶ Create a managed node to be defined as a Proxy Server
- The Proxy Server provides enhanced functionality over a Web server plug-in
 - ▶ Route requests
 - ▶ Load balance
 - ▶ Improve response time for Web resources



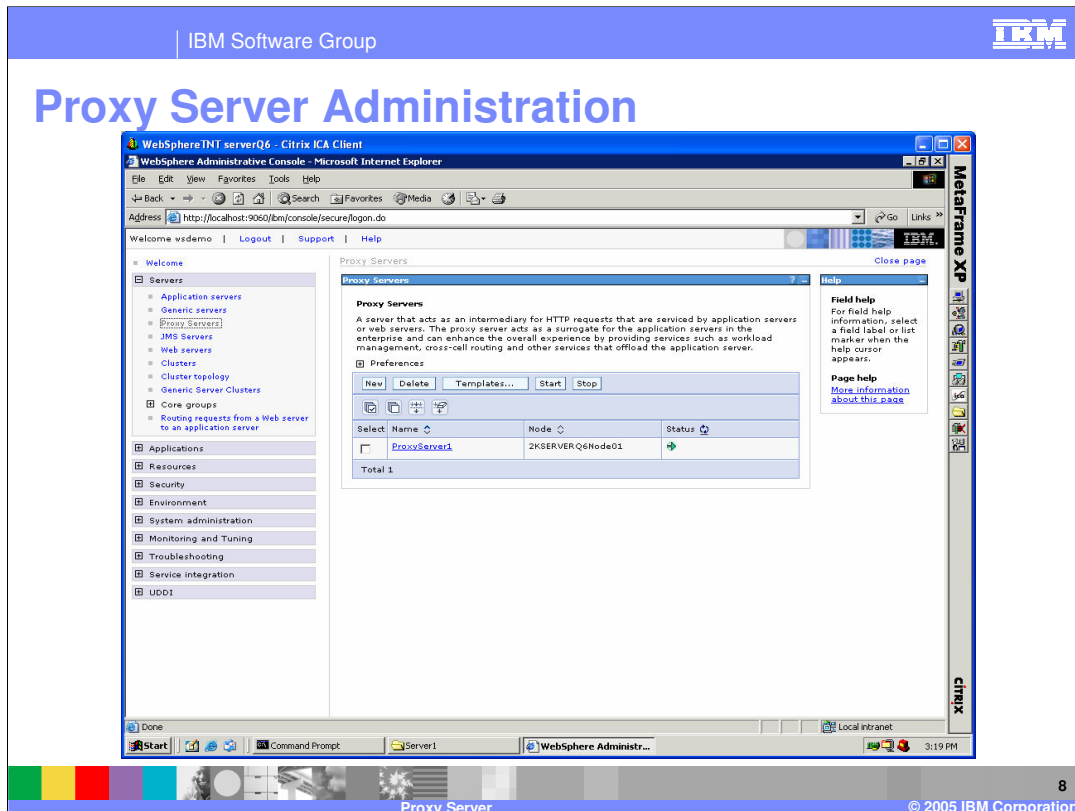
The WebSphere V6.0.2 maintenance pack enables users to define proxy servers within a ND topology. These proxy servers will be created as dedicated managed nodes. By using a managed node as a dedicated proxy server, a number of advanced features can be enabled. The proxy server allows for advanced features for routing requests into an environment, for load balancing requests, and for improving response time for web resources by utilizing the caching capabilities within WebSphere Application Server.

Proxy Server Options for WebSphere

- WebSphere XD On Demand Router (ODR)
 - ▶ An extension of the Proxy server that offers more advanced functionality
- Edge Component Caching Proxy (WTE)
 - ▶ Proxy Server is a strategic replacement with better performance and easier management
- Web server with plug-in
 - ▶ Proxy Server is a strategic replacement with better performance and easier management, as well as other advanced functions
- Tivoli® Access Manager WebSeal
 - ▶ Integrates with Proxy server to secure requests
- Other security proxies



There are a number of other proxy server solution that can be used with WebSphere. The On Demand Router (ODR) is an extension of the Proxy server and provides higher end capabilities such as SLA enforcement and integrates with the autonomic capabilities of WebSphere XD such as the dynamic clustering, provisioning and visualization. The Proxy server allows a seamless upgrade to ODR as the enterprise's need grow. The Proxy server is the strategic replacement and the next evolution of the Edge Component Caching Proxy (WTE) and provides a higher performing solution with most of the relevant reverse proxy capabilities of WTE in addition to dynamic routing capabilities and integrated Systems management. Many users also use the WebSphere plug-in. The proxy server provides a higher performing solution by employing a highly optimized intermediary architecture and provides capabilities such as dynamic work load management and caching over and above what is available in the WebSphere plug-in, including more complete reverse proxy function and better integration with WebSphere Systems management. The Tivoli Access Manager security proxy (WebSeal) can be used in conjunction with the Proxy server to secure requests for content in the enterprise. The proxy server provides better integration with WebSphere Application server. Existing security proxies can be used with the Proxy server to enable dynamic work load management and improve response times and throughput by utilizing services like persistent caching.



This screen shot shows the additions to the administration console to support the proxy server. This will only be added to the administration console by performing the profile augment command that is explained later in this presentation. An addition proxy servers option is now listed under the servers tab, this can be used to add and manage proxy servers for an environment.

Proxy Server Administration Cont.

The screenshot displays the WebSphere Administrative Console interface within a Citrix ICA Client window. The browser address bar shows the URL `http://localhost:9060/ibm/console/secure/login.do`. The left-hand navigation pane lists various server categories, with 'Proxy Servers' selected. The main content area is titled 'Proxy Servers > ProxyServer1' and contains a description: 'A server that acts as an intermediary for HTTP requests that are serviced by application servers or web servers. The proxy server acts as a surrogate for the application servers in the enterprise and can enhance the overall experience by providing services such as workload management, cross-cell routing and other services that offload the application server.' Below this is a 'Configuration' section with several expandable areas: 'General Properties' (Name: ProxyServer1, Run in development mode: unchecked, Parallel start: checked), 'Proxy Settings' (HTTP Proxy Server Settings, Proxy settings, Routing rules, Static cache rules, Proxy server transports, Proxy Cache instance config), 'Container Settings' (Web Container Settings, EJB Container Settings, Container Services), 'Server Infrastructure' (Java and Process Management, Administration), and 'Communications' (Ports, Messaging). A 'Field help' section on the right explains how to use field help markers. The bottom of the screenshot shows the Windows taskbar with the Start button, several application icons, and the system tray displaying 'Local intranet' and the time '3:22 PM'. The footer of the slide includes the text 'Proxy Server' and '© 2005 IBM Corporation'.

By selecting an available proxy server in the servers list, a user will be taken to this panel. From this panel you can further configure the settings for a proxy server. In particular, the options listed under the proxy settings area are of interest. The proxy settings option can be used to set proxy specific configuration information, the routing rules section is where a user can define specific rules for routing behavior. The static cache rules is where a user can enable the capabilities of the proxy servers caching solution.

Section

Proxy Server Functions



The next section will explain the technical functions supported by the proxy server.

Proxy Server Configuration

- In the past, plug-in configuration file generation and distribution was required to route HTTP traffic into application servers
- The Proxy Server can perform retrieval and distribution of back end routing (virtual hosts, URIs) and end point data (app servers hosts and ports)
- Each Proxy Server can be configured with one or more cell destination endpoints
 - ▶ Each Proxy Server will automatically update its routing tables as applications are installed or removed from the cells
 - ▶ In the event that the DMGR is down or unavailable, the Proxy Server continues to use its current configuration data



The proxy server has advanced configuration capabilities over the Web Server plug-in. The Web Server plug-in requires the generation and distribution of plug-in configuration files when changes are made to the WebSphere configuration. The plug-in configuration file contains information used by a Web Server to route requests into a WebSphere environment. With the proxy server being a managed node, it can automatically receive routing information from deployment managers through normal WebSphere communications. When a cell's configuration is changed the deployment manager will update any proxy servers that have been configured with the new configuration information. Each proxy server can be configured to route requests to one or more back-end cells, the deployment managers for each of these cells will update the proxy server when changes occur. A proxy server can still route requests when a deployment manager is unavailable, but will not receive any information about changes to the configuration.

Proxy Server Routing

- Dynamically routes requests to WebSphere Application Servers in backend cells
- Performs cross cell routing and load balancing with minimal configuration required
- Enables routing to any HTTP Server through the use of generic server clusters
 - ▶ Generic server clusters are groups of remote servers that the Proxy Server needs to route to



The proxy server can route requests to any WebSphere Application servers contained within back-end cells. The proxy server is capable of routing across cells, as well as performing basic load balancing of request. The proxy server can route to other types of servers through the use of generic server clusters. Other server types can be configured within WebSphere as generic server clusters, and the proxy server can route requests to these clusters.

Proxy Server Caching

- The Proxy Server can utilize the WebSphere dynamic cache feature to cache web content
 - ▶ Also supports static caching
- Using Edge Side Include (ESI) extensions, caching data can be included with a response
 - ▶ The Proxy Server can read this data to make decisions about caching



The proxy server can also perform caching of data contained within responses based on information within Edge Side Include extensions. This is different from normal caching, since the response is not being generated by the proxy server. Rather the application server that generates a response will also include information in the response within Edge Side Includes. A proxy server can read the data contained within edge side includes and use it to make decisions to cache portions or a response within the proxy servers cache. Based on data stored within it's cache the proxy server may be able to directly respond to a clients request, rather than forward the request to an application server.

HTTP Session Affinity

- The Proxy Server supports standard WebSphere session affinity using the jsessionid cookie/query parameter
 - ▶ The clone id from jsessionid is extracted and used to make the routing decision for the request
- If the selected clone id is down, then the weighted round-robin algorithm will select a new clone to dispatch the request



The proxy server fully supports the concept of session affinity. Session affinity is used to route requests back to the same server that was used for earlier work. This can make for better performance since objects will not need to be serialized and sent to other application servers within an environment. When the proxy server receives a request that is using session affinity, it will read the information contained within the jsession cookie, and route the request to the appropriate server. If the selected server is down, the proxy server will send the request to a different server.

Advanced Proxy Server Capabilities

- Configure routing to HTTP servers based on URI and virtual host
 - ▶ Allows redirection of a request to a more appropriate or secure URI
 - ▶ To allow failure of certain requests for a specific URI
- The ability to terminate SSL connections at the Proxy Server
- Generate custom error pages based on information within the response attributes
- Logging of requests, cached and generated responses



The proxy server has other more advanced routing features supported. You can configure routing rules to route requests based on the URI and virtual host information. This can be used to route a request to a more appropriate or secure server, and to allow failure for certain requests. The proxy server can also be used to terminate SSL connections and generate custom error pages for requests. The proxy server also fully supports the logging of information about requests and responses that flow through it.

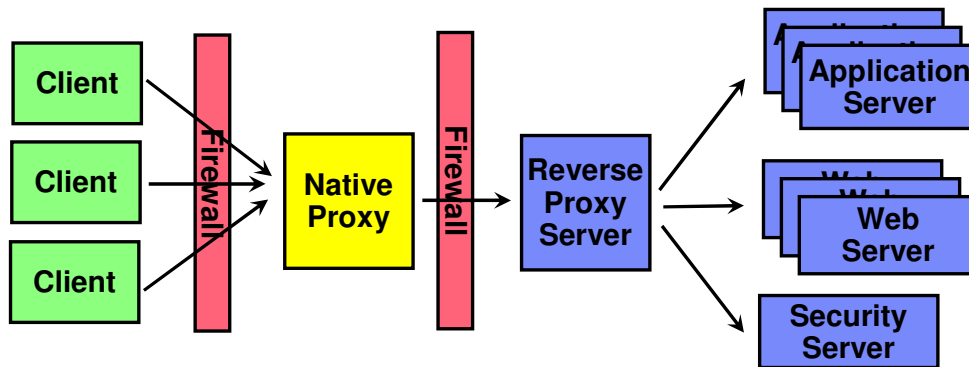
Section

Common Topologies

The next section will detail common topologies for a proxy server.

Topologies: Example 1


- RPS behind firewall and native proxy within firewall



Since the proxy server is a managed node, most users would not consider placing it within a DMZ for security reasons. Instead what is considered more appropriate in most situations is to place a proxy server behind a native proxy solution such as a web server plug-in. This offers a more secure solution, and enables the capabilities provided by the proxy server. However it also introduces a network hop, which can cause performance concerns.

Topologies: Example 2

- Integration with a Web Server farm
 - IP Sprayer -> Web Server -> RPS



In order to use multiple proxy servers, you should front the with an IP Sprayer, such as shown in this picture. IP sprayers are also often used to front multiple web servers contained within a web server farm. In this topology an additional IP sprayer could be placed in front of the proxy servers.

Section

Creating a Proxy Server



The next section will explain the additional steps required to create a proxy server.

Creating a Proxy Server

- Upgrade to V6.0.2
 - ▶ Upgrade DMGR and any nodes
 - ▶ DMGR must be stopped to upgrade
 - ▶ `<INSTALL_ROOT>\bin\wasprofile -augment -profileName PROFILE_NAME -templatePath <INSTALL_ROOT>\profileTemplates\proxy_augment`
- Within the administration console create a new proxy server
- Configure Proxy Server
 - ▶ Currently can not be administered as clusters



The first step of creating a proxy server is to install the V6.0.2 maintenance pack. After this has been done, a profile augment command needs to be run to add information into a profiles configuration to support a proxy server. This should be done on any deployment managers that will need to be configured for a proxy server, and on any managed nodes that will be used as a proxy server. The command to be used is shown in the example in the slide. When run on a deployment manager, the deployment manager must be stopped. The command will update the administrative console and this can only be done if the deployment manager is stopped. Once the profile augment is complete, proxy servers should be configured with in the updated administrative console.

Section

Summary

The next section is the summary for the presentation.

Summary

- Discussed the capabilities provided by the Proxy server in WebSphere 6.0.2
 - ▶ Configuration
 - ▶ Caching
 - ▶ HTTP session affinity
- Explained the common topology questions involved with Proxy Servers



This presentation explained the basic capabilities of the new proxy server component support by the WebSphere Application Server V6.0.2 maintenance pack. It explained other proxy server solutions that can be used with WebSphere and explained the topologies that are appropriate for this proxy server.

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
e(logo)/business	DB2	iSeries	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 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 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.