



IBM Software Group

IBM WebSphere® Extended Deployment for z/OS® V6.0.1 *On Demand Router*



@business on demand.

© 2006 IBM Corporation
Updated May 8, 2006

This presentation will cover the on demand router component of WebSphere extended deployment V6.0.1 for z/OS.

Agenda

- On Demand Router in WebSphere Extended Deployment
- On Demand Router topologies
- On Demand Router in z/OS
- On Demand Router administration



This presentation will first discuss the On Demand Router and how it interacts with other WebSphere Extended Deployment and z/OS features. Then some common topologies for using the On Demand Router and how to create and administer an On Demand Router will be covered.

Section

On Demand Router in WebSphere Extended Deployment



This section will discuss briefly how the On Demand Router interoperates with Other WebSphere Extended Deployment features.

On Demand Router

- The On Demand Router is an intelligent HTTP proxy server
 - ▶ Enhanced version of the proxy server from WebSphere Network Deployment V6.0.2
 - ▶ Can replace or complement the HTTP server plug-in
 - ▶ Prioritizes requests and controls traffic flow according to operational policy
- World class edge of Sysplex distribution
 - ▶ Compliments Sysplex Distributor
 - ▶ Complements z/OS workload management



The On-Demand Router is an intelligent HTTP proxy server that is provided with WebSphere Extended Deployment. It is the point of entry into a WebSphere Extended Deployment environment, and is responsible for request prioritization, flow control, and distribution to application servers. It can momentarily queue requests for less important applications in order to allow requests for more important applications to be handled more quickly. It is aware of the current location of a dynamic cluster's server instances, so that requests can be routed to the correct endpoint. The On Demand Router can also dynamically adjust the amount of traffic sent to each individual server instance based on processor utilization and response times. The On Demand Router can detect and react to a "storm drain condition". A storm drain condition can occur when a backend subsystem or resource crashes. Until the resource becomes available again, the server can return error codes quite quickly. A simple router can falsely assume the server is processing requests very efficiently and start sending more requests to the affected server. The server can become swamped with work when it is able to satisfy requests again. The On Demand Router detects this condition and deals with the condition in cooperation with the health monitor. These and other advanced features distinguish the On Demand Router from the HTTP server plug-in, and give the On Demand Router the ability to ensure a more consistent quality of service for your enterprise applications. It can be used in place of, or in concert with the HTTP server plug-in, depending on your needs.

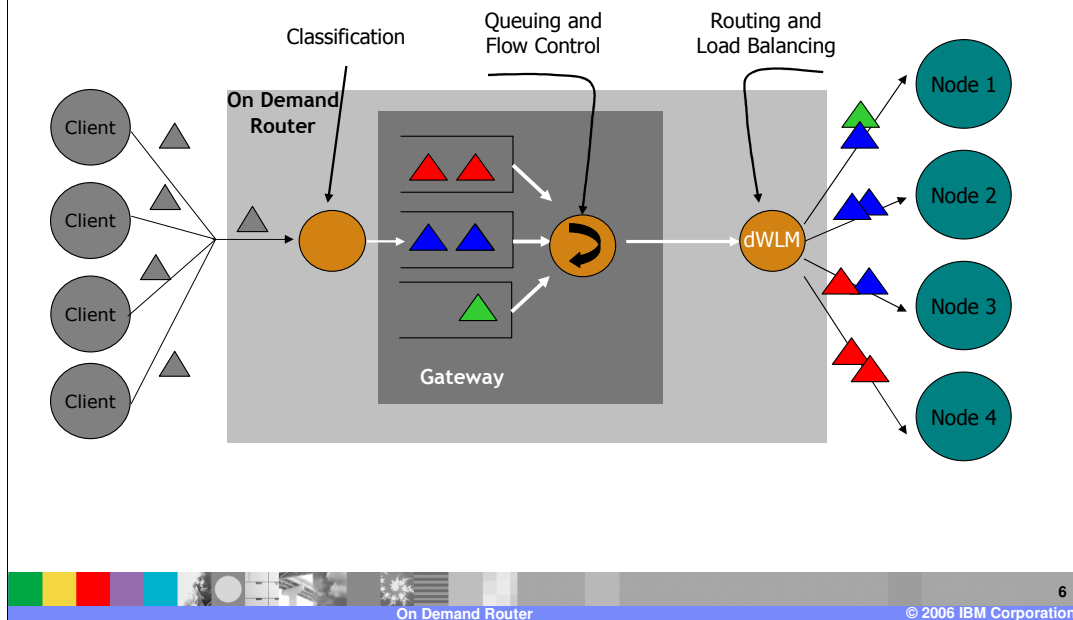
On Demand Router (cont.)

- On Demand Router Contributes to:
 - ▶ Java™ 2 Enterprise Edition (J2EE) batch
- On Demand Router is Required for:
 - ▶ Edition management
 - ▶ Health monitoring
 - ▶ Visualization
 - ▶ Scale out
 - ▶ Dynamic operations



As this chart points out the On Demand Router is quite ubiquitous in WebSphere Extended Deployment. By itself the On Demand Router enhances both Sysplex distributor and z/OS workload management. At the same time, it is required for many of the other WebSphere Extended Deployment features to work.

On Demand Router functional components



This diagram illustrates the flow of an HTTP request through the On Demand Router. When a request enters the On Demand Router, it is classified according to rules that you have defined. These rules can be based on many criteria, including Uniform Resource Identifier, HTTP headers, cookies, or the client's IP address. When a rule is matched, the request is placed into a queue with all other requests that are in the same class of service. The On Demand Router can drain these queues at different speeds, to ensure that requests in each class are able to meet their goals. This can mean holding less important requests in the queues for longer than more important requests, so that more resources are available to serve the important requests. As requests are released from the queues, they are routed to application servers by the dynamic workload management component, which is constantly aware of the location of each dynamic cluster member. Each member has a dynamically assigned weight, which can increase or decrease the number of requests it receives based on how quickly it is responding to requests. This feature is labeled dynamic work load management or dWLM as opposed to the z/OS workload management.

Autonomic Request Flow Manager (ARFM)

- Dispatches work out of the queues based on class weights
- Dynamically adjusts weights
 - ▶ Achieves response time goals for different classes of requests
 - ▶ Responds to varying load conditions and request surges
 - Limits the number of requests on a given server
 - Overload protection



The Autonomic Request Flow Manager is a component that runs inside of the On Demand Router. It is responsible for queuing work by service policy as requests come into the On Demand Router, and for controlling the flow of requests exiting the queues. It dynamically adjusts the rates at which requests flow through to application servers to ensure that response time goals are met and that no application servers are sent more requests than can be handled.

Autonomic Request Flow Manager (cont.)

- ARFM Gateway queues and dispatches requests
 - ▶ Manages logical queues
 - ▶ Queues are dynamically created as service classes appear in the request flow
 - ▶ Acts on every request
- ARFM Controller sets gateway dispatch weights
 - ▶ Monitors request flow, run times and application placements
 - ▶ Sets service class dispatch weights
 - ▶ Operates at scheduled intervals



The Autonomic Request Flow Manager has two logical components, the gateway and the controller. The gateway is responsible for putting each incoming request into a queue that represents its service class, and later dispatching requests from the queues. The controller operates periodically, monitoring request flow, response times, and application placements, and dynamically sets the dispatch weights for each logical queue, dictating how quickly requests are released from each of them.


Section

On Demand Router topologies



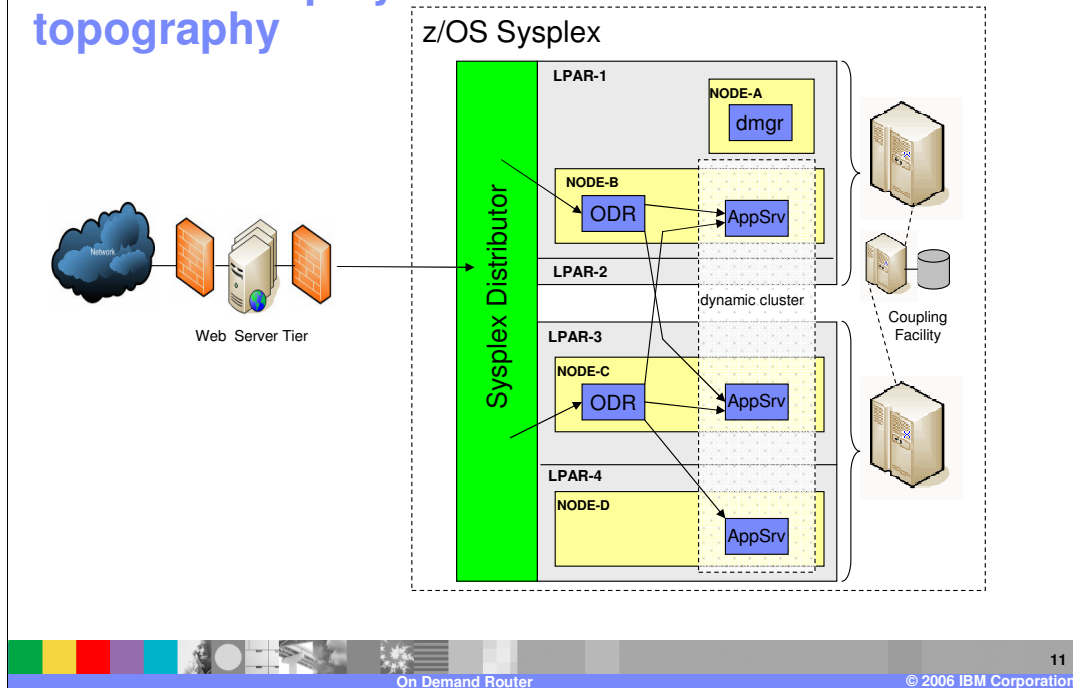
This section will present several possible topologies to use an On Demand Router.

Extended Deployment inboard OLTP topology



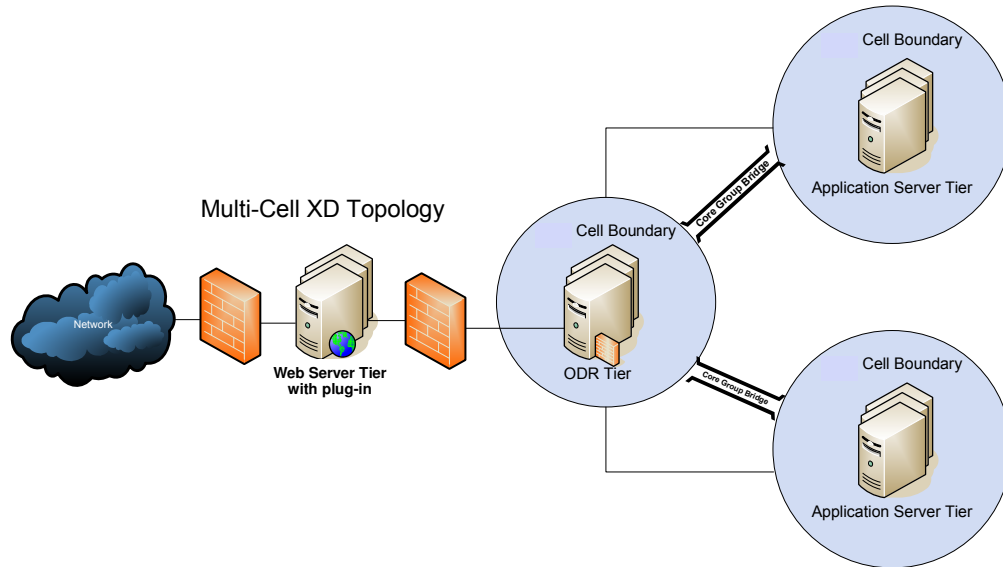
This is a high level view of a WebSphere Extended Deployment topology that uses an On Demand Router. It is also a very simple topology. Note that several On Demand Routers are present for redundancy and they are inboard or in the same cell as the application servers.

Extended Deployment z/OS inboard OLTP topography



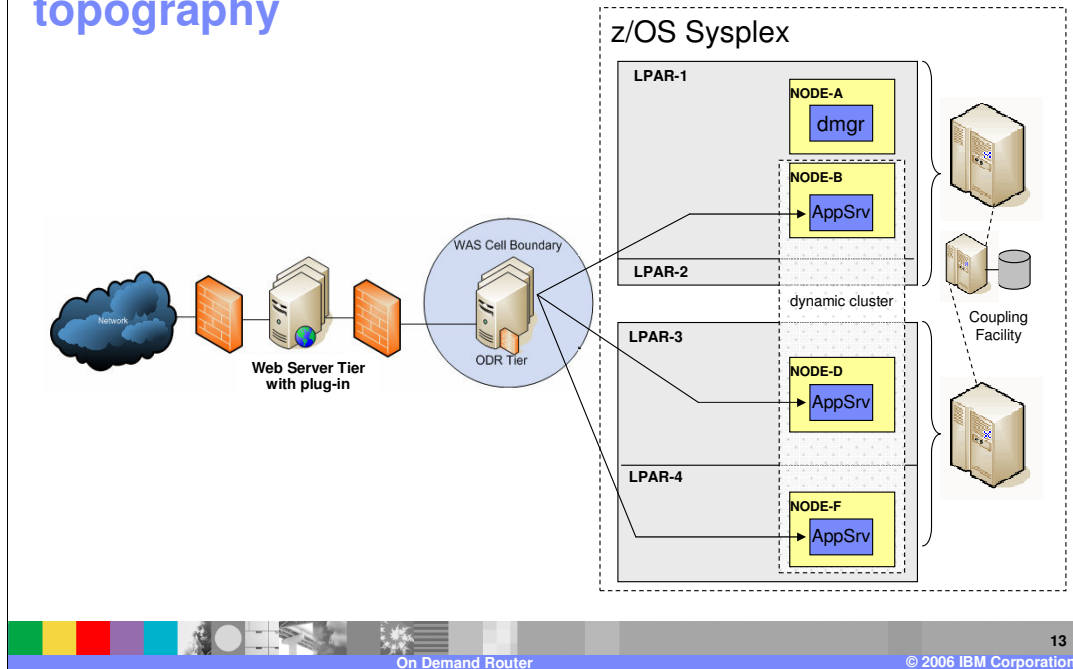
This slide presents the inboard On Demand Router configuration. Node A is shown as deployment manager or 'dmgr'. The other three nodes contain servers, both On Demand Routers, shown here as ODR's and application servers shown here as 'AppSrv'. The inboard On Demand Router configuration can take advantage of Sysplex distributor. Sysplex Distributor cluster addresses is a natural way to provide network transparency to the multiple On Demand Routers. Note that while the On Demand Routers are not clustered, they still provide high availability through redundancy; the multiple On Demand Routers work cooperatively as peers off of shared state data. In addition an On Demand Router can start additional servers in a dynamic cluster to accommodate increased load, for example if the load increases the On Demand Router running on NODE-C can start a server on NODE-D. This topic is discussed more in the section on dynamic operations. This type of configuration with the On Demand Router taking advantage of Sysplex distributor is the recommended kind of topology. However user constraints might require the On Demand Router to be outboard.

Extended Deployment outboard OLTP topology



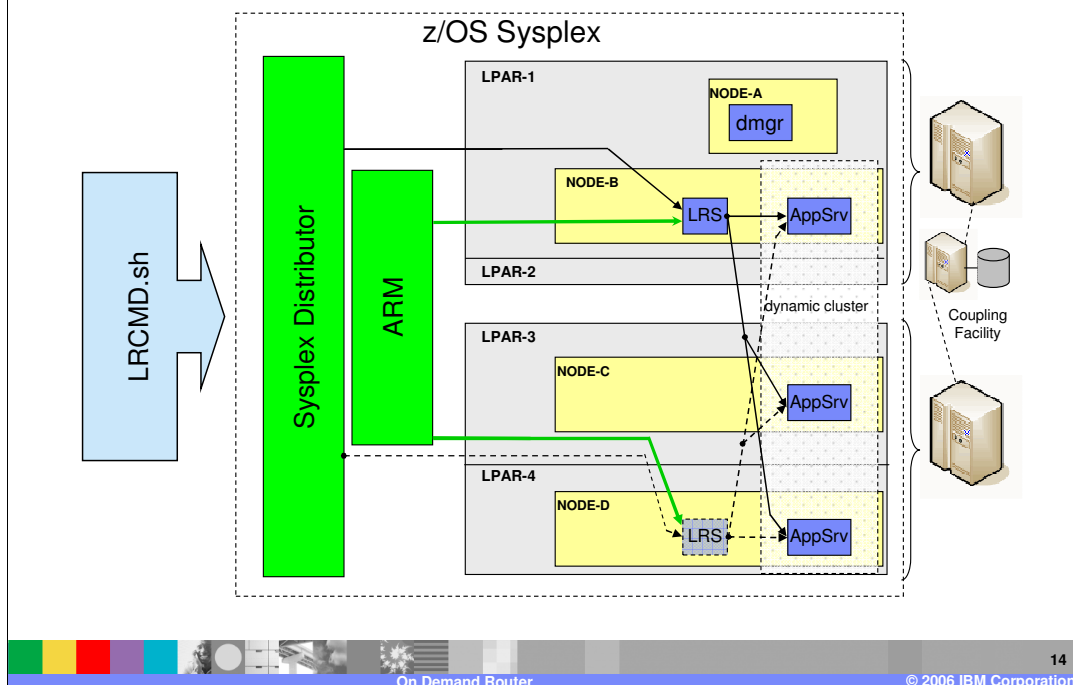
This slide shows the On Demand Router isolated from other WebSphere Extended Deployment cells, communicating with them through core group bridges. Here the On Demand Router could be on a z/OS platform and still take advantage of Sysplex distributor or alternately the On Demand Router could be on a Distributed platform. In either case the On Demand Router interacts with z/OS workload management and carries on the two way communication as will be discussing shortly, even when the On Demand Router is not on a z/OS platform.

Extended Deployment z/OS outboard OLTP topography



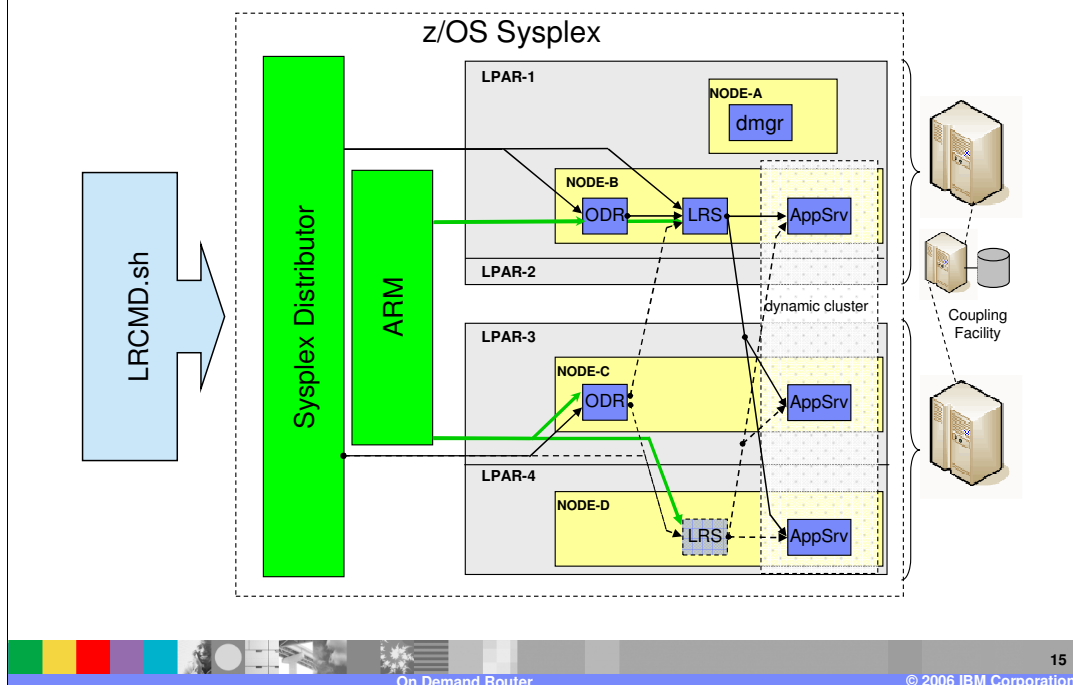
When the On Demand Router is not configured on a z/OS platform it cannot take advantage of Sysplex distributor. But the On Demand Routers still provide high availability through redundancy and they can start additional servers in a dynamic cluster to accommodate increased load and most importantly they can carry out the same bi-directional communication with the z/OS work load manager.

J2EE batch without On Demand Routers



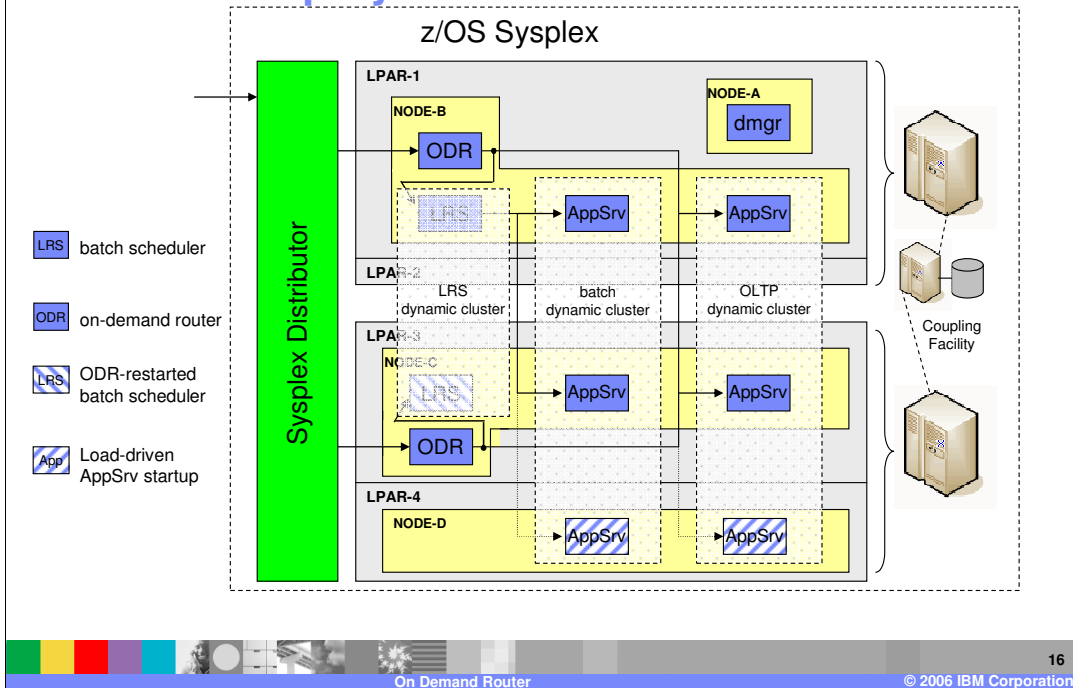
Because the Long Running Scheduler, or LRS, is an application running on an application server, an On Demand Router can improve the availability of J2EE batch. In particular, without an On Demand Router present, you will not have continuous availability of the Batch Application. At the start of this picture the Long Running Scheduler is running in NODE-B. The Long Running Scheduler is an application and the rule is there can only be one running Long Running Scheduler cell. Here a user enters batch commands through LRCMD.sh. These flow through Sysplex distributor to the Long Running Scheduler running in NODE-B. The control from Long Running Scheduler then follows the solid arrows to the application servers running the J2EE batch steps and all is fine. If for some reason the server running the Long Running Scheduler crashes, the solid arrow paths are lost. Here, ARM restarts an Long Running Scheduler on NODE-D in the picture as represented by the lower green arrow. Once the new Long Running Scheduler “comes up”, the flow is similar to before except following the dashed arrows. Between the time the first Long Running Scheduler crashed and the second Long Running Scheduler is ready, all communication is lost.

J2EE batch with On Demand Routers



With On Demand Routers present, you will have continuous availability of the Batch Application. This is the same picture as before except for the addition of two On Demand Routers. The main difference is to reroute the communication path from the Sysplex distributor from the running Long Running Scheduler to the On Demand Routers. ARM restart an On Demand Router if it goes down. Since multiple On Demand Routers can be running at the same time availability is assured. When an Long Running Scheduler goes down in this picture, the On Demand Router restarts a copy and while the Long Running Scheduler is restarting the On Demand Routers queues up requests. In this scenario, the On Demand Router has made an already great system better.

Extended Deployment z/OS Business Grid



This picture is collage of the onboard On Demand Router topographic options as discussed.






















Section

On Demand Router in z/OS



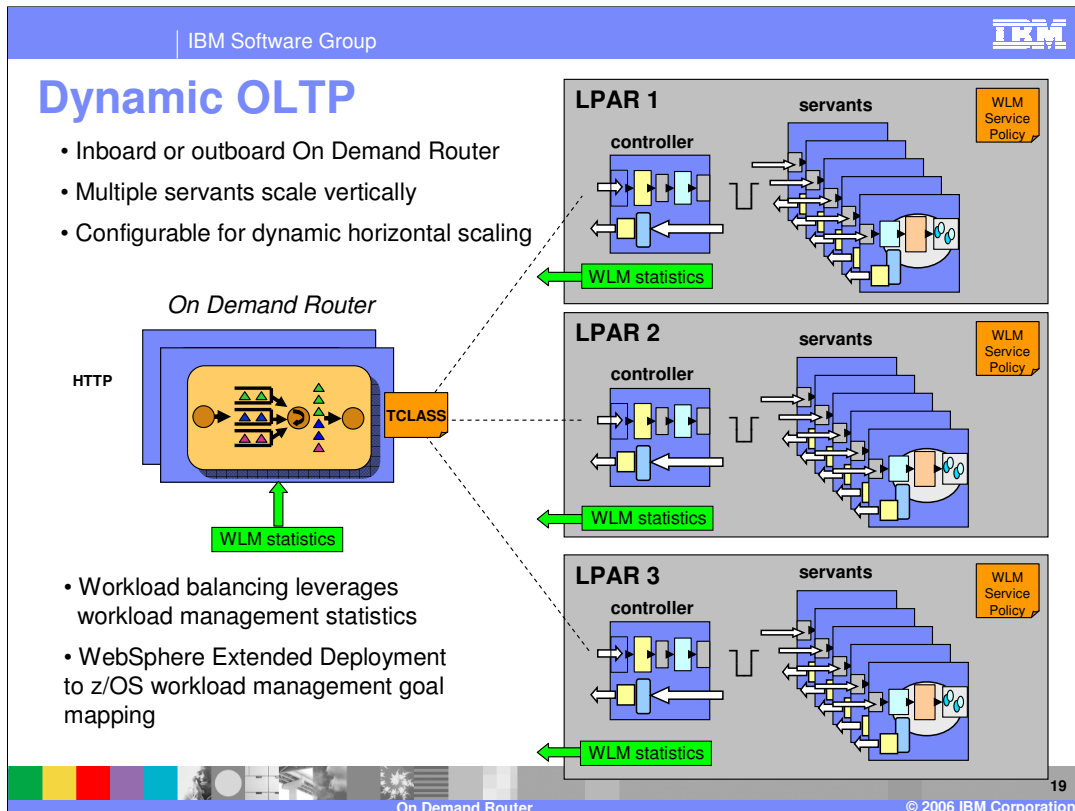
This section will discuss briefly how the On Demand Router interoperates with z/OS.

HTTP distribution scorecard

	Session Affinity Routing	Dynamic Workload Balancing	Autonomic Capacity Management	Storm Drain Avoidance	Application Lazy Start
Web Server and Plug-in					
Sysplex Distributor					
Web Server, Plug-in, Sysplex Distributor					
On-demand Router					

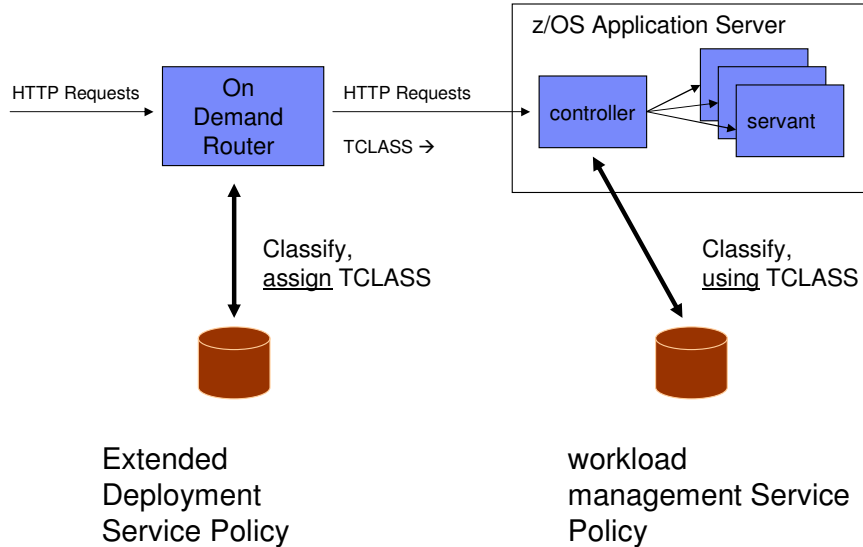


The On Demand Router works with and enhances Sysplex distributor. Here the Web Server with plug-in handles session affinity quite well, but does not perform dynamic workload balancing well. On the other hand Sysplex distributor handles dynamic workload balancing extremely well but has no sense of session affinity. Working together, Sysplex distributor and Web Servers with plug-in handle both dynamic workload balancing and session affinity quite well. The On Demand Router not only performs dynamic workload balancing and session affinity extremely well, the On Demand Router also handles the other items shown quite well. However an On Demand Router exists at the boundary of a Sysplex. As such it does not replace Sysplex distributor, instead On Demand Routers working in tandem with Sysplex distributor provide state of the art request delivery to WebSphere servers.



Here is the promised picture on dynamic OLTP. Again the On Demand Router is required for dynamic OLTP and it can be either inboard with Sysplex distributor or the On Demand Router can be outboard. The basic WebSphere application servers provide vertical scaling by deploying multiple servant regions through z/OS workload management. The On Demand Router adds horizontal scaling to WebSphere. In particular the On Demand Router can start or stop additional servers as needed for horizontal scaling. The On Demand Router accepts z/OS workload management just as Sysplex distributor does and feed output back to z/OS workload management through the TCLASS parameters and WebSphere Extended Deployment to z/OS workload management goal mapping. This topic will be discussed in greater detail shortly. As has been shown is the previous slides, the On Demand Router provides autonomic computing in that the On Demand Router can restart failed servers. This occurs on either on the same node or a different node if an entire node is lost for some reason.

OLTP: Extended Deployment workload management service policy mapping



Here is a little more detail of the feed back from the On Demand Router back to z/OS workload management. As already seen, the On Demand Router performs a classification and a prioritization of incoming OLTP requests. The result of the classification is to assign a TCLASS value to the incoming request. This TCLASS value is then passed along with the request to a server to process. The request's first stop in a server is to the controller region for the server. The next step is configurable, but by default, the controller region uses the TCLASS value to assign a z/OS workload management service policy to the request. The request is then queued to z/OS workload management and the flow follows the normal WebSphere for z/OS processing. If the default value is overridden, the controller region performs its own classification and selects a service policy. Both the On Demand Router and a controller region have stats from z/OS workload management to make decisions with. However, the On Demand Router also has the view of the complete WebSphere offered load and current server response times. In addition, managing the mapping of requests to TCLASS values is much easier with the On Demand Routers graphical interface. The standard WebSphere uses a flat file that has to be hand edited and the server restarted.

Powerful Request-to-TCLASS Match Engine

Rule Builder

Use this panel to create conditions for a match rule.

Generated Rule

Add Client Host (clienthost) Remove Clear All Cancel OK

- Client Host (clienthost)
- Client IPV4 (clientipv4)
- Client IPV6 (clientipv6)
- Cookie Header Name (cookies\$)
- Group Id (gid)
- Request Header Name (headers\$)
- HTTP Method (HTTPMethod)
- MIME Type (MIMEType)
- Port (port)
- Protocol (protocol)
- Request Query Parameter Name (queryparm\$)
- Server Host (serverhost)
- Server IPV4 (serveripv4)
- Server IPV6 (serveripv6)
- User Id (uid)

Match Rule — Transaction Class — Service Policy

21
© 2006 IBM Corporation

Here is the administrative interface just mentioned. It is called the rule builder and you can see that the user can build rules based on the values of a large number of fields. The results of the rule builder accumulate on a routing rules page of the administration console. Once you get used to how the rules are built, you can enter the rules directly into the routing rules page.

Classification Operands for HTTP Requests

Operand	Protocols	Description
clientHost, serverHost, clientIPv4, serverIPv4, clientIPv6, serverIPv6	HTTP, SOAP	Client and server host names and IP addresses
protocol	HTTP, SOAP	The request protocol
header, queryParm, cookie	HTTP, SOAP	HTTP header, query parameters, and cookies
HTTPMethod	HTTP, SOAP	HTTP method
MIMEType	HTTP, SOAP	MIME type
webService	SOAP	Web service name
webOperation	SOAP	Web service operation
virtualportal	HTTP	Virtual Portal name

Here is a listing of the operands you can use for routing rules.

Section

On Demand Router administration

This section will introduce creation and administration of an on demand router. You can also view the demonstration “Administering an On Demand Router”.

IBM Software Group IBM

Creating a new On Demand Router

Welcome jimip | [Logout](#) | [Support](#) | [Help](#)

- Welcome
- Guided Activities
- Servers
 - Application servers
 - Generic servers
 - Proxy Servers
 - On Demand Routers
 - JMS Servers
 - Web servers
 - Clusters
 - Cluster topology
 - Dynamic Clusters
 - Generic Server Clusters
- Core groups
- Applications
- Resources
- Runtime Operations
- Security
- Operational Policies
- Environment
- System administration

On Demand Routers

On Demand Routers

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-call routing and other services that offload the application server. The On Demand Router is an extension to the Proxy server that provides value added services such request flow management to enforce SLA guarantees.

Preferences

Select	Name	Node	Status
Total 0			

■ Administrative Console or wsadmin “createodr.jacl”

24

On Demand Router © 2006 IBM Corporation

You can create a new On Demand Router from the wsadmin.sh script “createodr.jacl” or you can use the Administrative Console. The administration page for On Demand Routers can be found under Servers in the left pane of the Administrative Console. To create a new On Demand Router, click on new and follow the wizard instructions.

Administering an On Demand Router

On Demand Routers > c3odr1

A server that acts as an intermediary for HTTP requests that are serviced by application servers or web servers. surrogate for the application servers in the enterprise and can enhance the overall experience by providing serv management, cross-call routing and other services that offload the application server. The On Demand Router is Proxy server that provides value added services such request flow management to enforce SLA guarantees.

Configuration

General Properties

Name: c3odr1

*short Name: CL3ODR1

Unique Id: 8E44AF4733301F8500000FC0000001092A76EE

Run in development mode

Parallel start

Apply OK Reset Cancel

On Demand Router Settings

- On Demand Router Properties

Container Settings

- Web Container Settings
- EJB Container Settings
- Container Services
- Business Process Services

Server Infrastructure

- Java and Process Management
 - [Class loader](#)
 - [Process Definition](#)
 - [Environment Entries](#)
 - [Process Execution](#)
 - [Monitoring Policy](#)
 - [Server Instance](#)
- Administration
 - [Custom Properties](#)
 - [Administration Services](#)
 - [Server Components](#)
 - [Custom Services](#)

Communications

- Ports
- Messaging

Selecting a an On Demand Router from the On Demand Routers list produces the main administration page for an On Demand Router. Here you can set the short name. From the Server Instance you can select how many copies of this server to run. From the Custom Properties you can set cluster transition name and from the ports selection you can set the On Demand Routers ports. For more detail on this, refer to the demonstration named, "Administering an On Demand Router".

Summary

- On Demand Router is critical to many WebSphere Extended Deployment features
- Workload management integration
 - ▶ Workload management metric exploitation by On Demand Router
 - ▶ Extended Deployment Workload management service policy mapping for OLTP/HTTP



The on demand router offers world class edge of Sysplex routing. The on demand router is required by many WebSphere Extended Deployment features and enhances all OLTP traffic. The on demand router is best deployed with and enhances Sysplex distributor. And the on demand router interoperates effectively with z/OS work load manager.

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 2005,2006. 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.