



IBM eServer™

Automated Domain Name Registration: Operation

@business on demand software

© 2007 IBM Corporation

A Sysplex ADNR "owns" the zones it maintains!

➤ **ADNR exclusively owns the zones defined in its configuration**

- A zone is owned by a Sysplex ADNR and is unique to a single z/OS® Sysplex
- Zones should not be shared between ADNRs in multiple z/OS Sysplexes
- Resources should not be manually added/updated/deleted in those zones

➤ **During start-up, ADNR does a zone transfer to learn which resources are currently in the zone (from a previous start of that Sysplex ADNR)**

- Based on that information and what ADNR learns from the z/OS LBA, it will add/delete/update resource records in the DNS
- This approach was chosen to avoid any windows where otherwise available resources would be seen as unavailable (not registered in the DNS)

➤ **To determine the contents of a zone that ADNR owns:**

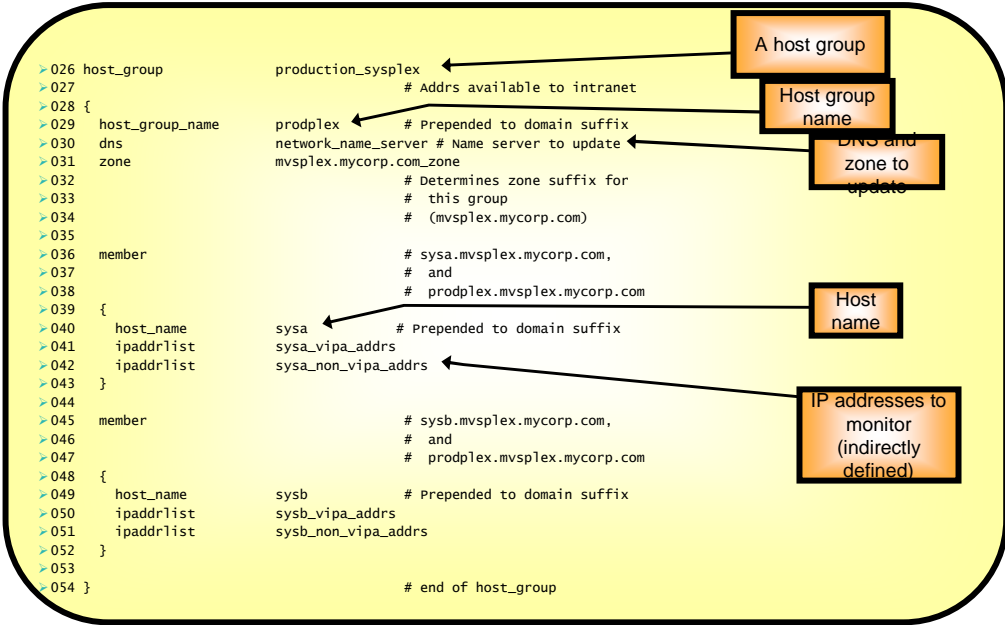
- Use the DNS dig client to manually display the contents of the zone. dig will require the transaction signature (TSIG) key for zones that have TSIG security protection.

Example ADNR config file - 1 of 4

```

>001 debug_level      7          # Error, Warning, Event
>002
>003 uuid             mycorp_sysplex_adnr ← Name ADNR registers with z/OS LBA
>004
>005 dns              network_name_server # Label used by other stmts
>006                  # and commands
>007 {
>008   dns_id          10.1.10.55 ← The DNS ADNR updates
>009                  # Network name server
>010   zone            mvsp1ex.mycorp.com_zone
>011                  # Label used by other stmts
>012                  # and commands
>013   {
>014     domain_suffix mvsp1ex.mycorp.com ← The zone ADNR updates
>015                  # Zone name in name server
>016   }
>017                  # end of zone
>018 }
>019                  # end of dns
>020 gwm              z/os_lba_advisor ← The z/OS LBA ADNR talks to
>021 {
>022   gwm_id           10.1.5.1.3860 ← LBA lb_connection_v4 address
>023   host_connection_addr 10.1.10.11 # Local address
>024 }
  
```

Example ADNR config file - 2 of 4



Example ADNR config file - 3 of 4

```
>056 ipaddrlist      sysa_vipa_addrs
>057 {
>058   ipaddr          10.1.10.3
>059 }                # end of ipaddrlist
>060
>061 ipaddrlist      sysa_non_vipa_addrs
>062 {
>063   ipaddr           10.1.10.22   # OSA on sysa
>064 }                # end of ipaddrlist
>065
>066 ipaddrlist      sysb_vipa_addrs
>067 {
>068   ipaddr           10.1.10.11
>069 }                # end of ipaddrlist
>070
>071 ipaddrlist      sysb_non_vipa_addrs
>072 {
>073   ipaddr           10.1.10.1   # OSA on sysb
>074 }                # end of ipaddrlist
```

IP Address lists
to be referenced
in other
definitions

Example ADNR config file - 4 of 4

```

>076 server_group      tn3270_group # TN3270 servers
>077 {
>078
>079   port              23 # TN3270 port
>080   protocol          TCP # Protocol for this port
>081   server_group_name ztelnet # Prepend to domain suffix
>082   dns               network_name_server # Name server to update
>083   zone              mvsp1ex.mycorp.com_zone
>084                  # Determines zone suffix for
>085                  # this group
>086                  # (mvsp1ex.mycorp.com)
>087   member           # telnetprimary.ztelnet.mvsp1ex.mycorp.com
>088                  # and,
>089                  # ztelnet.mvsp1ex.mycorp.com
>090   {
>091     server_name     telnetprimary # Prepend to
>092                   # server_group_name.domain
>093                   # suffix
>094     ipaddrlist      sysa_vipa_addr
>095   }
>096   member           # telnetsecondary.ztelnet.mvsp1ex.mycorp.com
>097                  # and,
>098                  # ztelnet.mvsp1ex.mycorp.com
>099   {
>100     server_name     telnetsecondary # Prepend to
>101                   # server_group_name.domain
>102                   # suffix
>103     ipaddrlist      sysb_vipa_addr
>104   }
>105 }
>106 } # end of server_group

```

A server group

Protocol and port

Server group name

DNS and zone to update

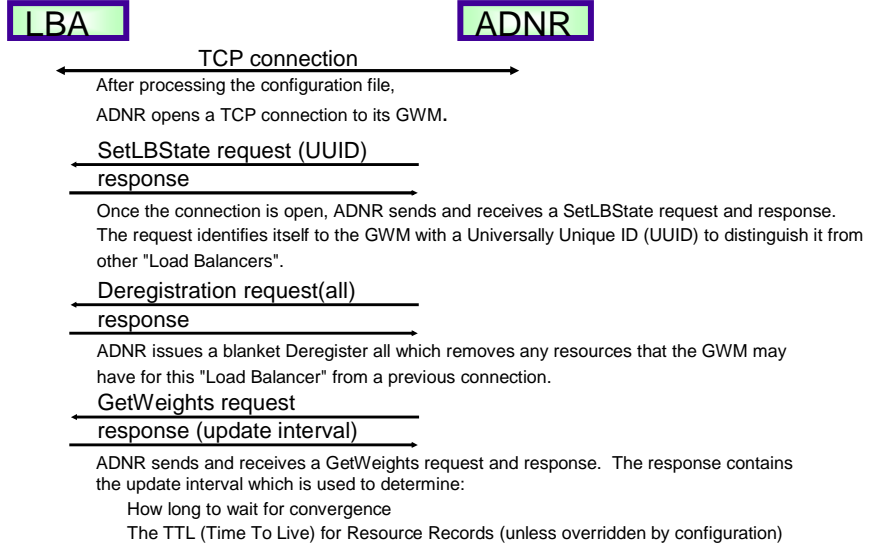
Member(s) of the server group

Member-specific name

IP addresses to monitor (indirectly defined)

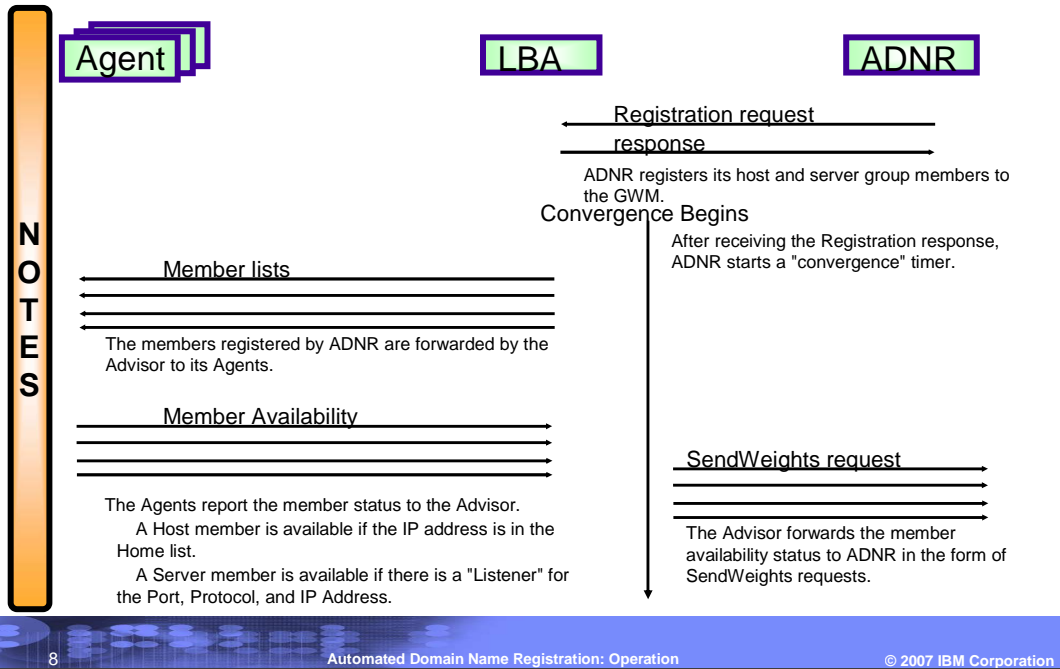
Overview of Flows - Bringup

ADNR communicates with the GWM as a Load Balancer using architected SASP flows:



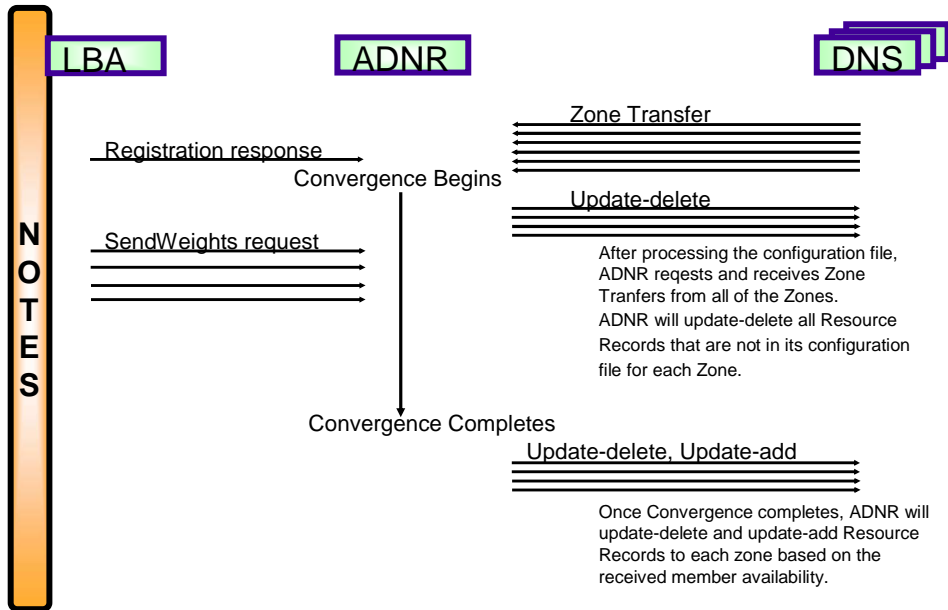
NOTES

Overview of Flows - Bringup..

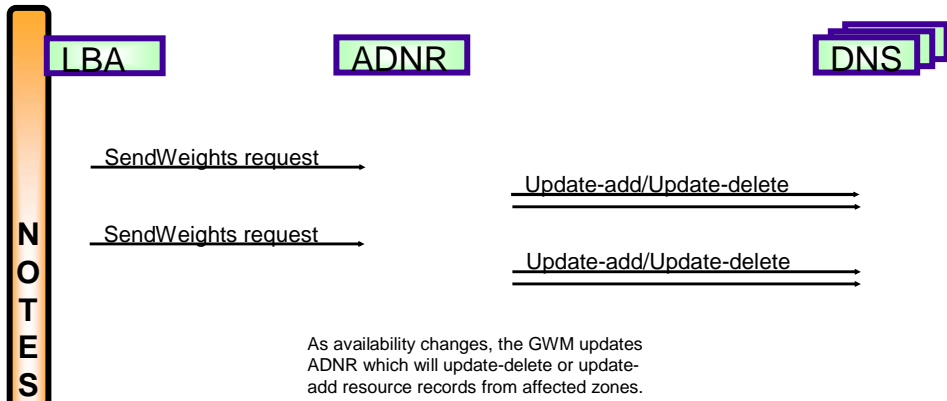


NOTES

Overview of Flows - Bringup..



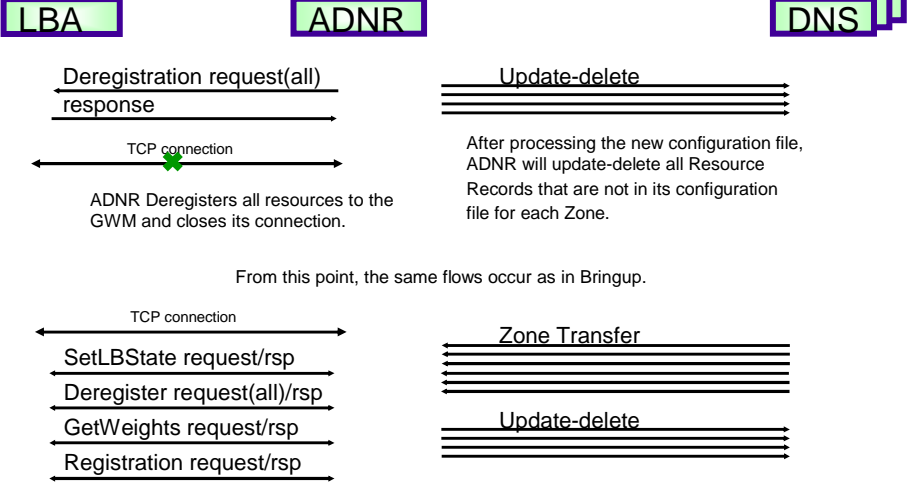
Overview of Flows - Steady State



Overview of Flows - Reconfiguration

ADNR can be reconfigured by changing its configuration file and issuing a "Refresh" command.

NOTES



From this point, the same flows occur as in Bringup.

Zone synchronization - convergence

- **Once convergence completes, ADNR will attempt to update all of the nameserver's zones with the appropriate resource records.**
 - Resource records that are no longer available will be update-deleted from each zone.
 - Resource records that are available will be update-added to each zone (if they are not already present).

- **Once this has occurred the Zone will be in the SYNCHRONIZED state**
 - As ADNR's GWM state becomes GWM_ACTIVE, each Zone will change state from RESYNCH_RECONCILE_PENDING
 - Assuming that there are no errors, each Zone will transiently enter RESYNCH_ZONE_UPDATE_PENDING state while the updates are in progress
 - Finally the Zone will enter the SYNCHRONIZED state
 - Issue a MODIFY procname,DISP,ZONES command to verify the zone state

- **Zones which are not responding (NOT_RESPONSIVE_*) will not become synchronized until the zone becomes responsive**

Automatic restart of ADNR

- **Ensure that automation is in place to restart ADNR on the same/other system in the sysplex in cases of failures (of ADNR or the System)**
- **Can be accomplished with Automatic Restart Manager (ARM) policy or other automation**
 - ↳ Note that while TCP/IP's AUTOLOG can be used to start ADNR, it can NOT be used to monitor the availability of ADNR after initial startup.
- **To enable ADNR to restart and operate properly on another TCP/IP stack (CINET) or on another sysplex system, ADNR must be configured with a unique application-instance DVIPA on the host_connection_addr keyword of the gwm statement.**
- **ADNR may be (re)started while the Advisor is already running or vice versa**
- **ADNR may be (re)started while its name servers are already running or vice versa**
- **Because ADNR relies on the LBA Advisor and Agent, those applications also need to be similarly configured for high availability. See next page.**

Advisor (GWM) outage

➤ **Mechanisms should be in place to keep the Advisor highly available in order to ensure ADNR continues to receive near real-time data.**

- Use of DVIPAs
- Automatic Restart Manager (ARM)
- TCP/IP Autolog
- Other automation

➤ **ADNR periodically attempts to reconnect to the Advisor when the connection to it is lost.**

➤ **When the connection to the Advisor is reestablished, ADNR resynchronizes the availability status of the sysplex resources with its name servers**

- Update-deletes any RRs representing resources which became unavailable during the outage
- Update-adds any RRs representing resources which became available during the outage

Name server availability

➤ Name servers must also be highly available

- To clients requiring name resolution
- To ADNR for near real-time dynamic updates

➤ Typical DNS setups have multiple authoritative name servers per zone

- One primary name server
- One or more secondary name servers
 - Secondary name servers replicate data from the primary name server
- Avoids single point of failure for client queries

➤ Single point of failure possibility still exists for the name server ADNR updates

- No architected schemes to avoid this
- Possible solutions:
 - z/OS name server with VIPA takeover and shared file system
 - z/OS name server with automation for recovery
 - Network name server with automation for recovery

Name server outage/config mismatch

➤ **ADNR issues an immediate action console message when it needs to update a name server and...**

- Name server is down
- Name server is unreachable
- Name server and ADNR are incompatibly configured

➤ **Examples of incompatible configurations:**

- Incorrect IP address of name server configured in ADNR
- Name server specified is not actually authoritative for the zone being updated
- Name server is not configured to allow ADNR to perform zone transfers
- Name server is not configured to allow dynamic updates from ADNR
- Zone transfer key mismatch between ADNR and name server
- Zone update key mismatch between ADNR and name server

➤ **ADNR periodically sends health probes to name server it cannot update**

- 1) dynamic update-add + update-delete
- 2) zone transfer
- ADNR resynchronizes the availability status of the sysplex resources with the name server when both health probes are successful
 - Update-deletes any RRs representing resources that became unavailable during the outage
 - Update-adds any RRs representing resources which became available during the outage

Things to think about

Supported name servers

➤ **Supported name servers**

- Any name server which supports RFC 2136 (DNS Update)
- Includes BIND 8 and BIND 9
 - Includes z/OS BIND 9
- Microsoft® Windows® 2000 and 2003 Servers

➤ **Does NOT include z/OS BIND 4.9.3**

Zone issues

- **ADNR will examine the contents of zones specified on the dns statement during zone resynchronization. Address (A and AAAA) RRs other than those representing active ADNR managed resources will be update-deleted. Glue records will be ignored. All other RRs will not be deleted.**
- **Zones managed by ADNR must be updated exclusively by ADNR. The zones and zone files should not be edited by hand nor should anyone or any other program perform dynamic updates to the zones.**
- **Orphaned resource records (RRs)**
 - RRs in a name server which ADNR can no longer remove
 - May cause clients to attempt to connect to unavailable servers
 - Causes:
 - ADNR unable to delete RRs during a F REFRESH operation due to network or name server problems
 - Improper F REFRESH procedures (like removing a dns statement from the config file while ADNR is down)
 - Solutions:
 - Remove name server from service
 - "Flush" the name server or zone using ADNR. Consists of adding a dns or zone statement to ADNR, but not referencing the zone(s) from any ADNR groups, followed by a F REFRESH when ADNR can communicate with the name server or zone.

LBA compatibility

- **ADNR is compatible with all versions of the z/OS Load Balancing Advisor (V1R4 and higher).**
- **If using a down-level Advisor or Agent, obtain the latest maintenance for the LBA component.**

Migrating from DNS/WLM

- **ADNR provides one of the two main functions of DNS/WLM**
 - ADNR supports automated DNS registration of application-specific hostnames
 - ADNR does not support load balancing using Sysplex-awareness

- **ADNR can be used to replace DNS/WLM load balancing functions if:**
 - A hostname resolution approach to load balancing is sufficient for your environment (such as no client caching issues)
 - An alternative solution like SD or external load balancing solution based on the LBA can not be deployed
 - Round-robin distribution is sufficient for your sysplex load balancing needs
 - Note: Some resolvers, including the z/OS resolver, may re-sort the returned addresses
 - ADNR merely registers when servers and addresses are available, it does not provide WLM-based load-balancing recommendations

- **ADNR has many advantages over DNS/WLM**
 - Uses the updated BIND 9 name server instead of the older BIND 4.9.3 name server
 - Name server need not reside on z/OS
 - No need for target servers or TCP/IP stacks to register with WLM. This means any z/OS server can participate
 - No limit to the number of IP addresses supported per TCP/IP stack. DNS/WLM has a limit of 15.
 - Supports IPv6.

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 z/OS

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

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 2007. 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.