



How to Setup Application Server to Access DB2 z/OS with High Availability

Data Management

Maryela Weihrauch,
DE DB2 z/OS, IBM Silicon Valley Lab,
weihrau@us.ibm.com

Important Disclaimer

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY.

WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED.

IN ADDITION, THIS INFORMATION IS BASED ON IBM'S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE.

IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION.

NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, OR SHALL HAVE THE EFFECT OF:

- **CREATING ANY WARRANTY OR REPRESENTATION FROM IBM (OR ITS AFFILIATES OR ITS OR THEIR SUPPLIERS AND/OR LICENSORS); OR**
- **ALTERING THE TERMS AND CONDITIONS OF THE APPLICABLE LICENSE AGREEMENT GOVERNING THE USE OF IBM SOFTWARE.**

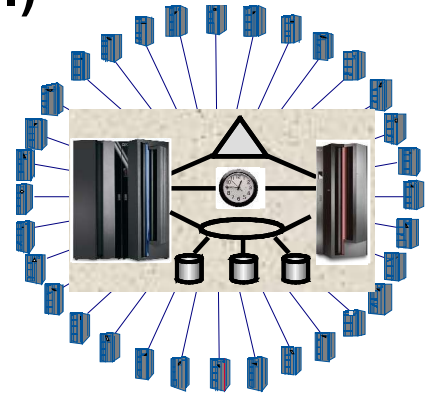
Agenda

- **DB2 z/OS concepts for availability**
- **Client configurations**
- **DB2 Sysplex Workload Balancing**
- **Application server considerations**

DB2 Data Sharing - Availability

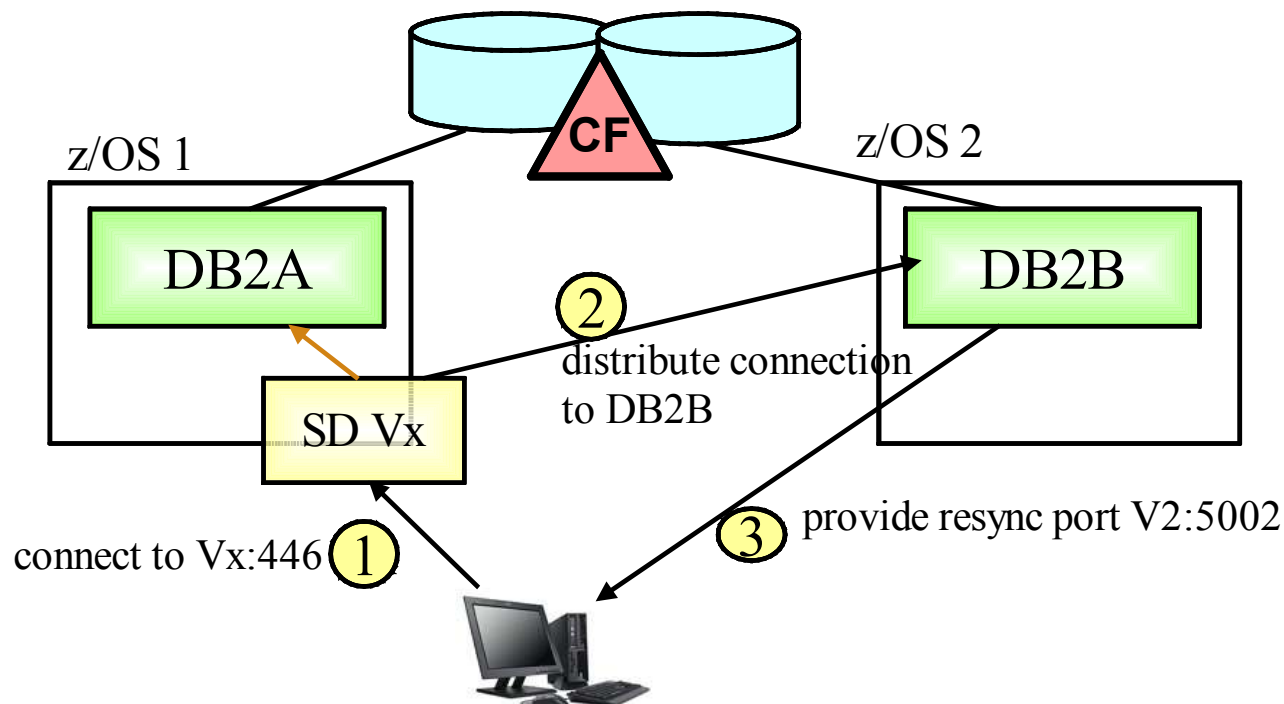
Goal: Continuous availability across any planned or unplanned outage across entire z HW/SW stack.

- **Elimination of single point of failure (DB2 subsystem, z/OS system, CPC, I/O paths,..)**
- **Remove all causes for planned outages**
 - rolling “on-line” upgrades and maintenance
 - online schema evolution
 - online utilities
- **On a failure:**
 - Isolate failure to lowest granularity possible
 - Automate recovery and recover fast



Challenge: the distributed application server needs to find the (best) available path to the data.

DVIPA and Sysplex Distributor – the Concept



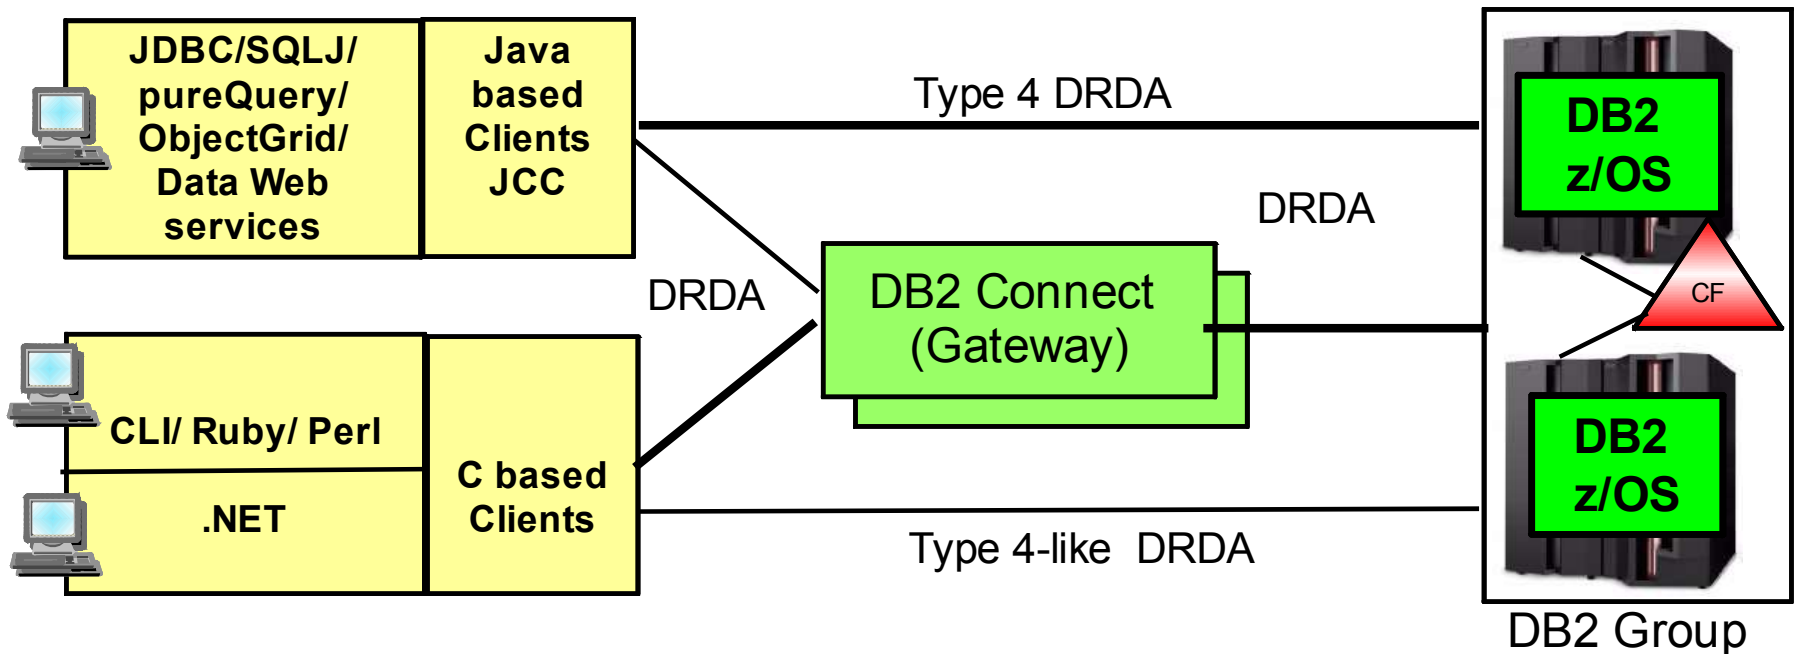
- Group DVIPA provides a virtual TCPIP address into the Data Sharing Group
- Sysplex Distributor routes the connection request to the most available member based on WLM recommendation

See also: IOD1336 - DB2 for z/OS Distribute Processing Best

DVIPA and Sysplex Distributor – Consideration

- **Benefit:**
 - **Connections are successful as long as one member is up**
 - Connection level workload balancing between members
 - Setup is isolated to z/OS environment
- **Drawbacks**
 - SD on one lpar may route to a member on a different lpar, which results into slightly higher response time compared to direct member access.
 - Information about availability of data sharing members is only considered at creation of a "new" connection but application server typically maintain long-running connections.
- ✓ **DVIPA and Sysplex Workload Balancing at the application server need to be enabled to ensure the highest availability characteristics in a distr. environment.**

DB2 Client Configuration to Access DB2 z/OS

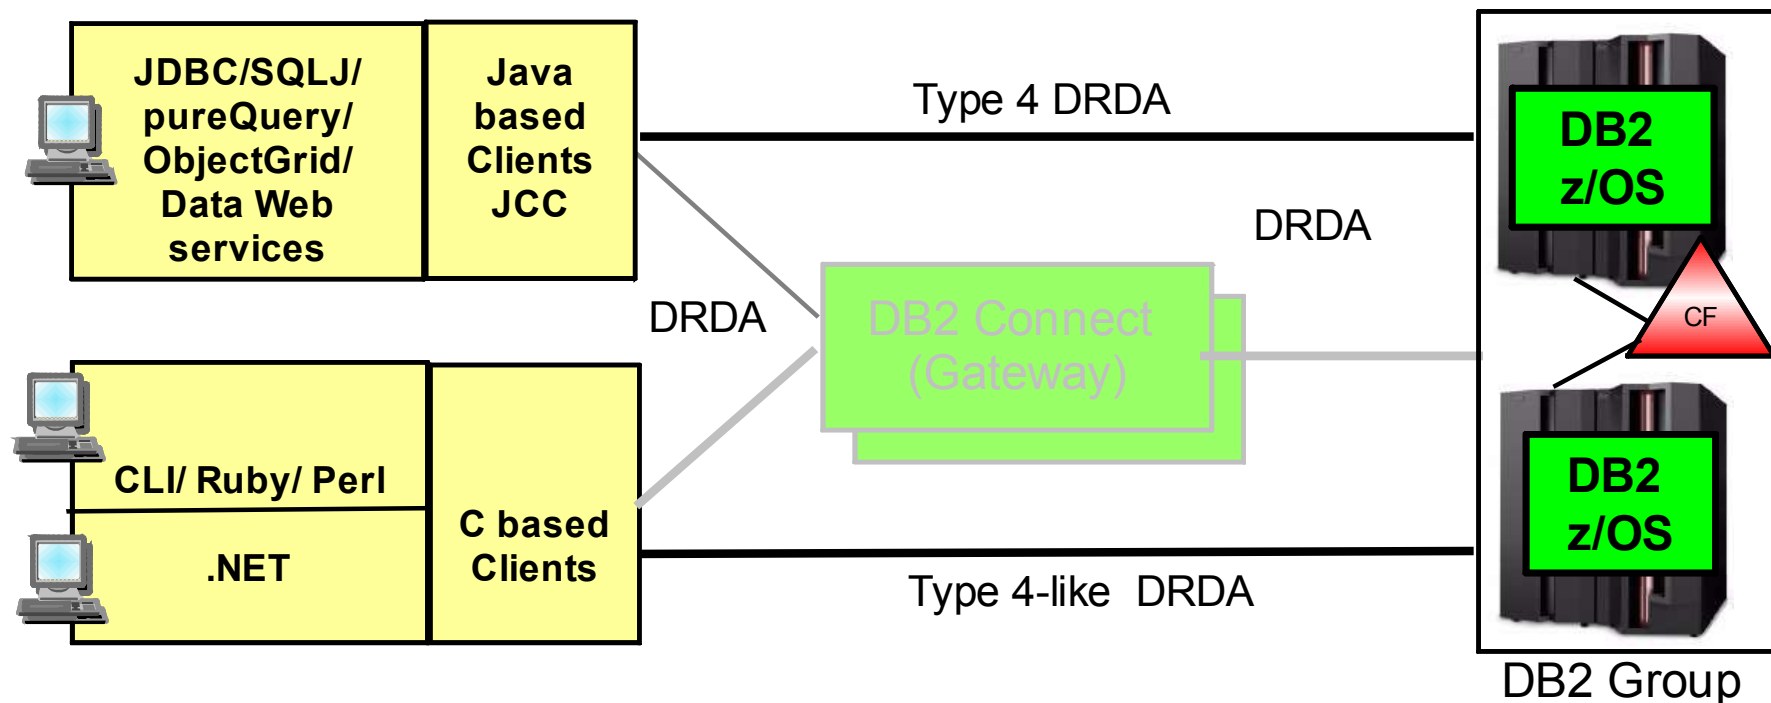


- Until DB2 Connect V7, DB2 Connect converted DB2RA (distr. protocol for DB2 LUW) into DRDA (distr. protocol for DB2 z/OS, DB2 iSeries)
- All C based application connect via DB2 Connect (client or gateway) or type 4 – like with IBM Data Server Driver 9.5
- Java based application connect directly via JCC type 4 or via DB2 Connect Gateway.

IBM Data Server Driver 9.5 for CLI and .NET App

- **Smaller footprint and stand alone download for driver and development component (11M+30M)**
- **run DB2 Connect EE or PE license manager to install license entitlement for connectivity to DB2 z/OS**
- **Sysplex Workload Balancing support in Data Server Driver 9.5 FP3**
- **More detailed instructions**
<http://www.ibm.com/developerworks/wikis/display/DB2/DB2%20and%20.NET%20FAQ>

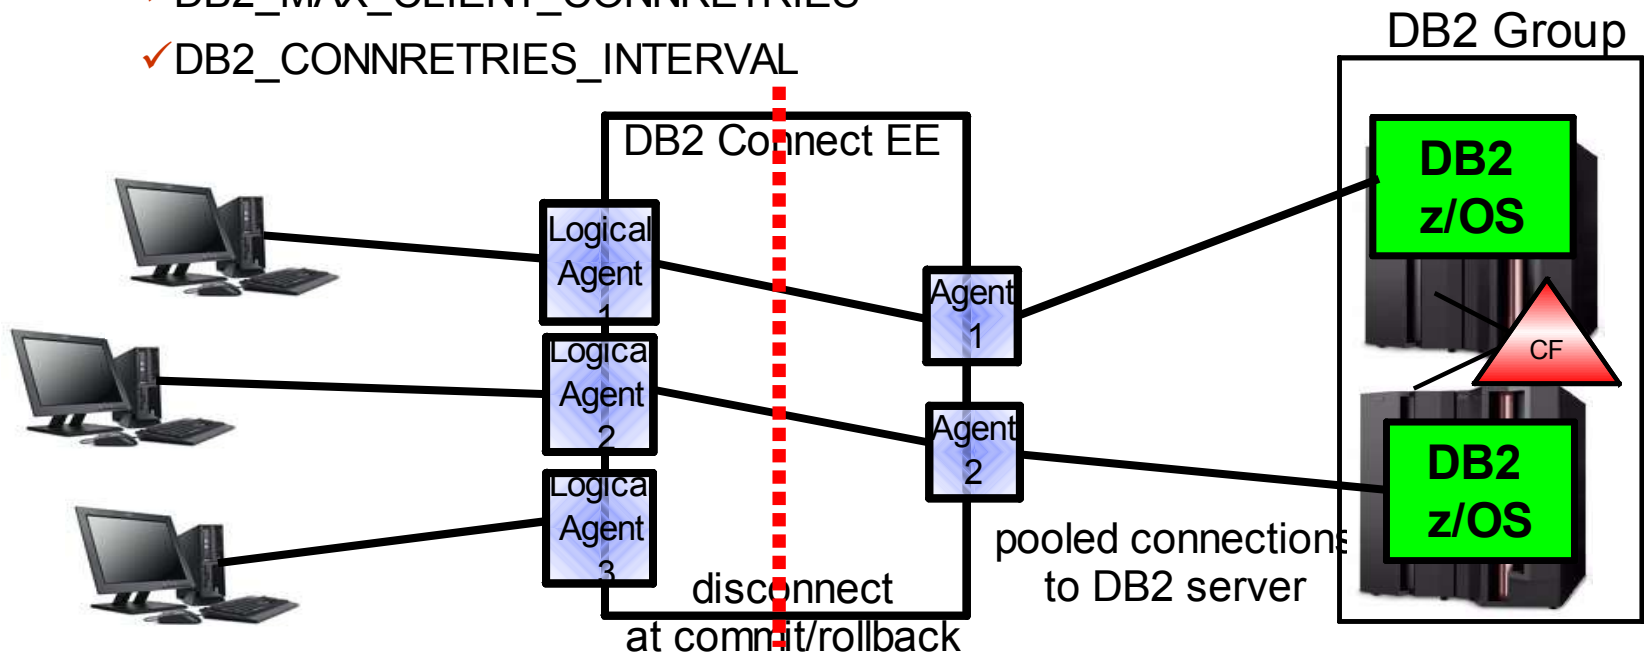
DB2 Client Configuration to Access DB2 z/OS – Future Directions



- **Clients - especially application server - access DB2 z/OS directly (will not change licensing model)**
- **Sysplex Workload Balancing is supported by DB2 Connect gateway, JCC type 4, and Data Server Driver 9.5 FP3 for CLI and .NET application**

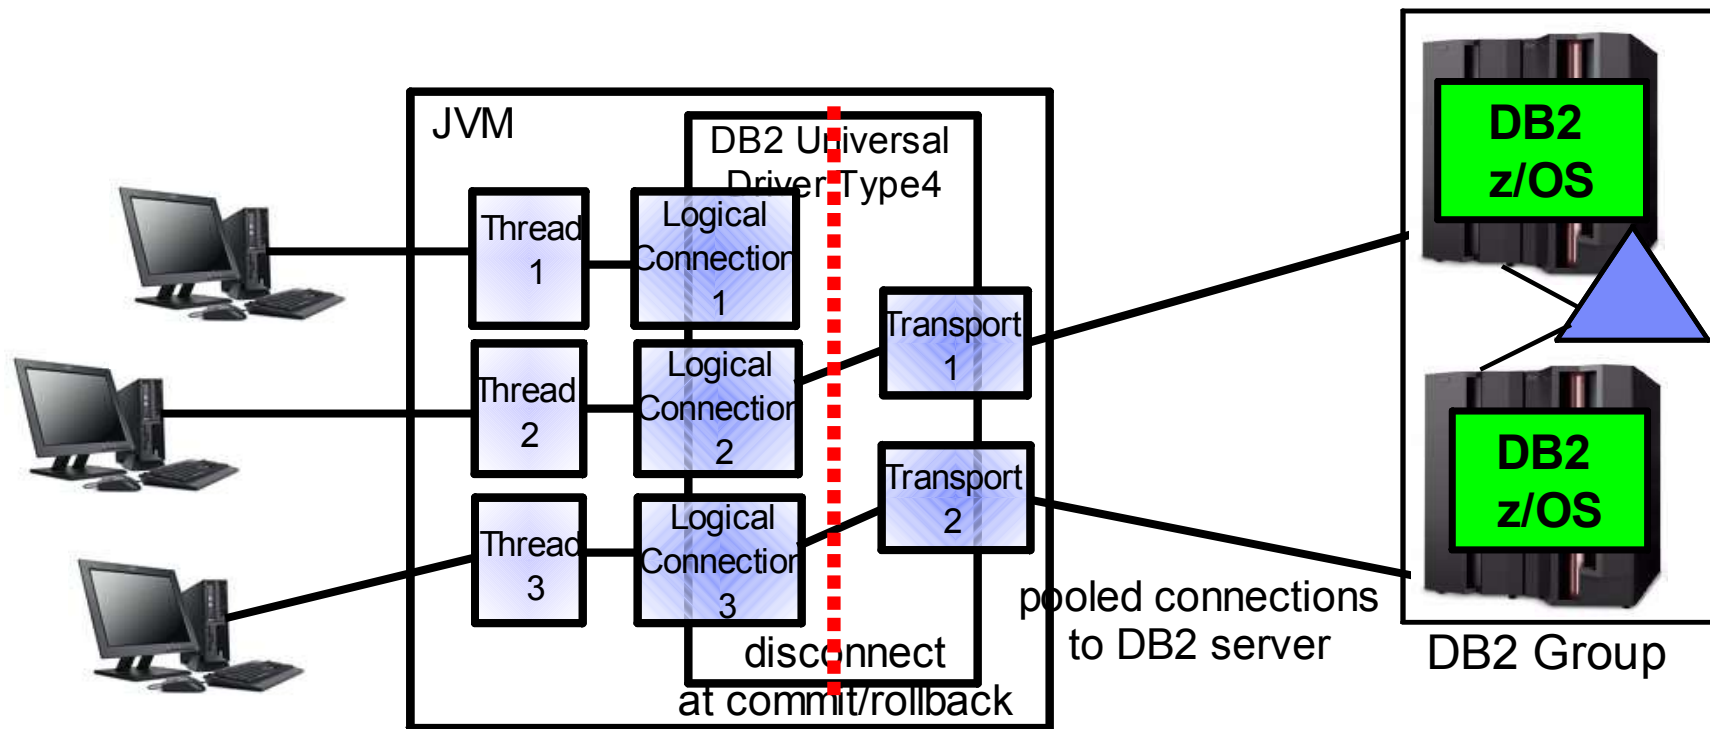
DB2 Connect EE Connection Concentrator

- ✓ MAXAGENTS - max. # connections to DB2 server
- ✓ MAX_CONNECTIONS - max. # of concurrent client connections to DB2 Connect (default is equal to MAXAGENTS - concentrator disabled)
- ✓ **MAX_CONNECTIONS > MAXAGENTS to enable sysplex workload balancing**
 - ✓ DB2 Connect 8.5 introduced 2 new options
 - ✓ DB2_MAX_CLIENT_CONNRETRIES
 - ✓ DB2_CONNRETRIES_INTERVAL



JCC Type 4 Sysplex Workload Balancing

- JCC type 4 supports Sysplex Workload Balancing
 - ✓ JDBC 2.0 datasource since DB2 Connect V8 FP10 – >JCC 2.7.xx (e.g. in WebSphere Application Server)
 - ✓ JDBC 1.2 DriverManager since DB2 Connect 9.5 – >JCC 3.50.xx (e.g. in Weblogic)



JCC Type 4 Sysplex Workload Balancing . . .

➤ Typical DataSource Properties

☞ **enableSysplexWLB=YES** - enables Sysplex Workload Balancing

☞ default is false, disabled

☞ **maxTransportObjects** - max # of connections to DB2 server from this DataSource.

☞ Can not be bigger than db2.jcc.maxTransportObjects

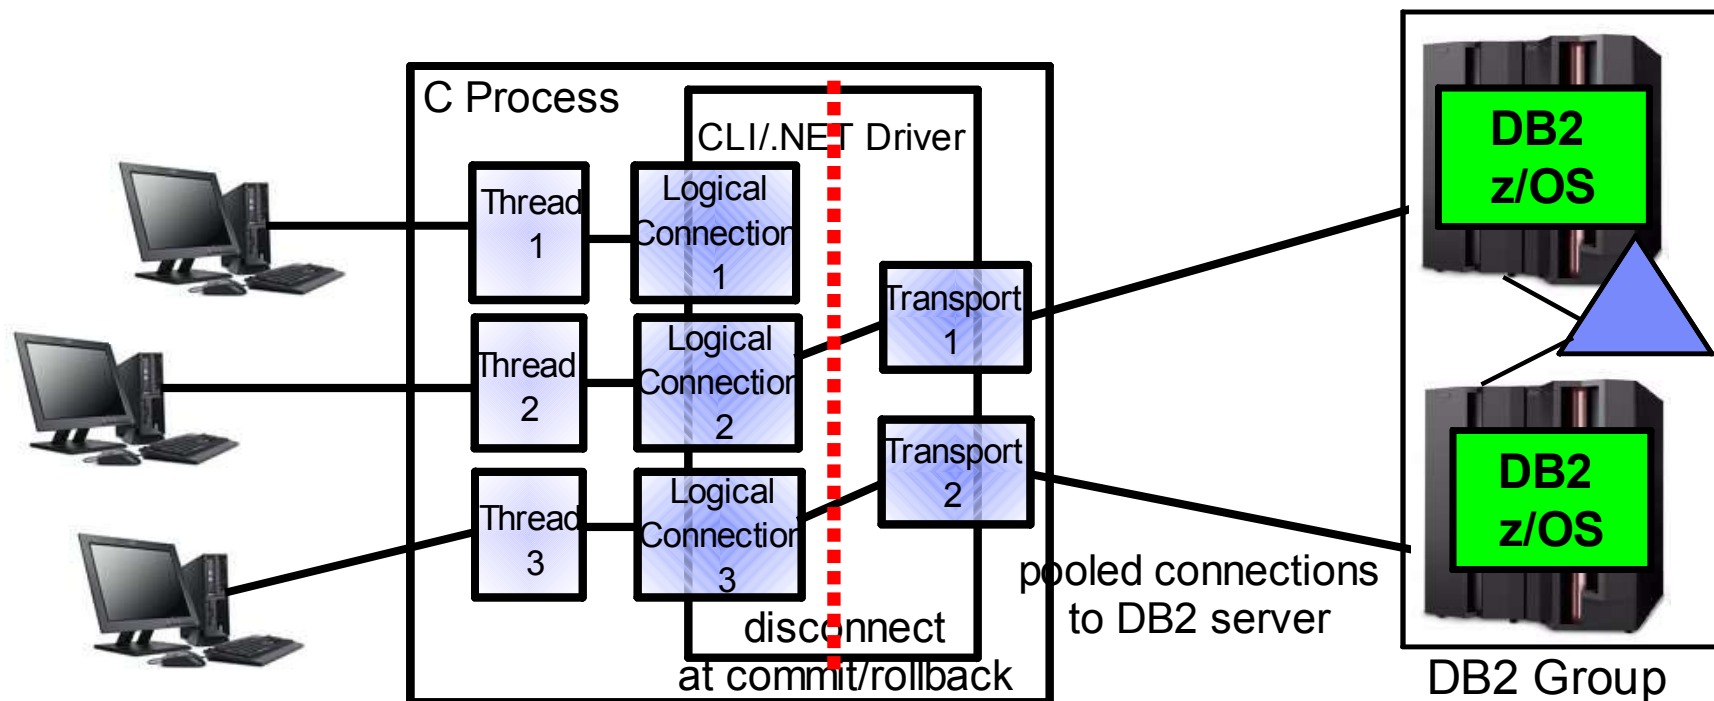
☞ default value is -1, meaning no limit

JCC Type 4 Sysplex Workload Balancing . . .

- **Global properties defined in Global Properties File:**
 - ☞ **db2.jcc.maxTransportObjects** - max # of connections to DB2 server across all datasources (default value is -1, meaning no limit)
 - ☞ **db2.jcc.maxTransportObjectIdleTime** -time in sec., a connection stays idle in the pool before it is closed,
 - ☞ default value is 60 sec
 - ☞ **db2.jcc.maxTransportObjectWaitTime** - if maxTransportObjects is reached - time in sec., an application waits to get a connection before throwing a SQLException
 - ☞ default value is 5 sec
 - ☞ **blockingReadConnectionTimeout** –
 - ☞ socket timeout value
 - ☞ should be set to a value greater by few seconds than the time it takes to execute the longest query in the application
 - ☞ **loginTimeout** –
 - ☞ only used for new connection


C-based Client Sysplex Workload Balancing

- Data Server Driver 9.5 FP3 for CLI and .NET supports Sysplex Workload Balancing
- Limited XA Support



XML Config. File – db2dsdriver.cfg

```
< configuration >  
  < DSN_Collection >  
    < dsn alias=, name=, host=, port= >  
    < parameter name=, value= />  
  < /dsn >  
< /DSN_Collection >  
< databases >  
  < database name=, host=, port=,>  
    < parameter name=, value= />  
    < WLB >  
    < parameter name=, value= />  
    < /WLB >  
    < ACR >  
    < parameter name=, value= />  
    < /ACR >  
  < /database >  
< /databases >  
< parameters >  
  < parameter name=, value= />  
< /parameters >  
< /configuration >
```



WLB Configuration

```
<databases>
  <database name="STLEC1" host="9.30.30.5" port="446">
    <WLB>
      <parameter name="enableWLB" value="true"/>
      <parameter name="maxTransports" value="100"/>
      <parameter name="maxTransportIdleTime" value="600"/>
      <parameter name="maxTransportWaitTime" value="30"/>
      <parameter name="maxRefreshInterval" value="30"/>
    </WLB>
    <ACR>
      <parameter name="enableACR" value="true"/>
      <parameter name="enableSeamlessACR"
value="true"/>
    </ACR>
  </database>
</databases>
```


Fault Tolerant Sysplex Support

- **Initial connection to DB2 data sharing group url**
 - recommended to be a distributed DVIPA. This virtual IP address resolves to any available member based on WLM
 - or a DNS name that resolves into a list of IP address of each member that is used round robin
 - It always works if at least 1 member is up.
- **The initial connection (or any ACCRDB responds) returns a member IP list with WLM weights**
- **Any following connection: uses client sysplex workload balancing algorithm to determine which member to use if reuse is OK:**
 - No open WITH HOLD cursor
 - No declared global temporary tables must exist
 - (used declared global temp tables must be explicitly or implicitly dropped)
 - No reference to packages bound with KEEP DYNAMIC YES
 - Rollback always leaves a connection in a reusable state

The server list

The server list can be explored with db2pd -sysplex

Sysplex List:

Count: 3

IP Address	Port	Priority	Connections	Status
9.12.6.70	38320	53	0	0
9.12.4.202	38320	53	0	0
9.12.6.9	38320	21	0	0

APAR PK80474 adds display server list to -DIS DDF DETAIL

```
-D9C1 DIS DDF DET
DSNL080I  -D9C1 DSNLTDDF DISPLAY DDF REPORT FOLLOWS:
...
DSNL081I  STATUS=STARTD
DSNL100I  LOCATION SERVER LIST:
DSNL101I  WT IPADDR          IPADDR
DSNL102I  45 ::9.12.4.105
DSNL102I  42 ::9.12.4.103
DSNL102I  18 ::9.12.4.104
DSNL099I  DSNLTDDF DISPLAY DDF REPORT COMPLETE
```

Behavior at Connection Errors

- **DB2 returns indicator as part of COMMIT responds if transport can be reused and returns SET statements to replay connection state at reuse**
- **If first SQL stmt in transaction fails and reuse OK**
 - No errors reported back to application
 - SET statements associated with the logical connection are replayed with first SQL on another transport
- **If subsequent SQL fails and reuse OK,**
 - -30108 reuse error returned to application (transaction is rolled back and reconnected).
 - SET statements are replayed to recover connection state
 - Up to application to retry transaction
- **If subsequent SQL and reuse not OK**
 - -30081 connection failed error returned to application.
 - Connection returned to initial (default) state
 - application needs to reestablish connection state and retry transaction
- **If all members in the member list are tried and none seems to be available, the initial data source url (DVIPA/DNS name) is retried to make sure that really no member is available.**

WebSphere Connection Pooling

WebSphere Administrative Console - Microsoft Internet Explorer

Address: <http://carlsr31:9090/admin/preferenceAction.do?checkbox1=on&node1=System%2Fworkspace%23auto-refr>

WebSphere Application Server Administrative Console Version 5

Home | Save | Preferences | Logout | Help

Resource Adapters > J2CResourceAdapter > J2C Connection factories > J2C3 >

Connection Pools

Connection pool properties that can be modified to change the behavior of the J2C connection pool manager. Default values are provided for non-production use. Reviewing and possible modification of these configuration values is recommended. [?](#)

Configuration

General Properties	
Scope	cells:carlsr31Network:nodes:carlsr31
Connection Timeout	180 seconds
Max Connections	10 connections
Min Connections	1 connections
Reap Time	180 seconds
Unused Timeout	1800 seconds
Aged Timeout	0 seconds
Purge Policy	EntirePool

Apply OK Reset Cancel

WebSphere Status [?](#) < Previous Next > February 13, 2003 2:28:17 PM EST

Done Local intranet

WebSphere Connection Pool Properties

- **Max Connections**
 - max connections from JVM instance
- **Min Connections**
 - lazy minimum number of connections in pool
- **Reap Time**
 - How often cleanup of pool is scheduled in seconds
- **Unused Timeout**
 - How long to let a connection sit in the pool unused
- **Aged Timeout**
 - How long to let a connection live before recycling
- **Purge Policy**
 - After StaleConnection, does the entire pool get purged or only individual connection

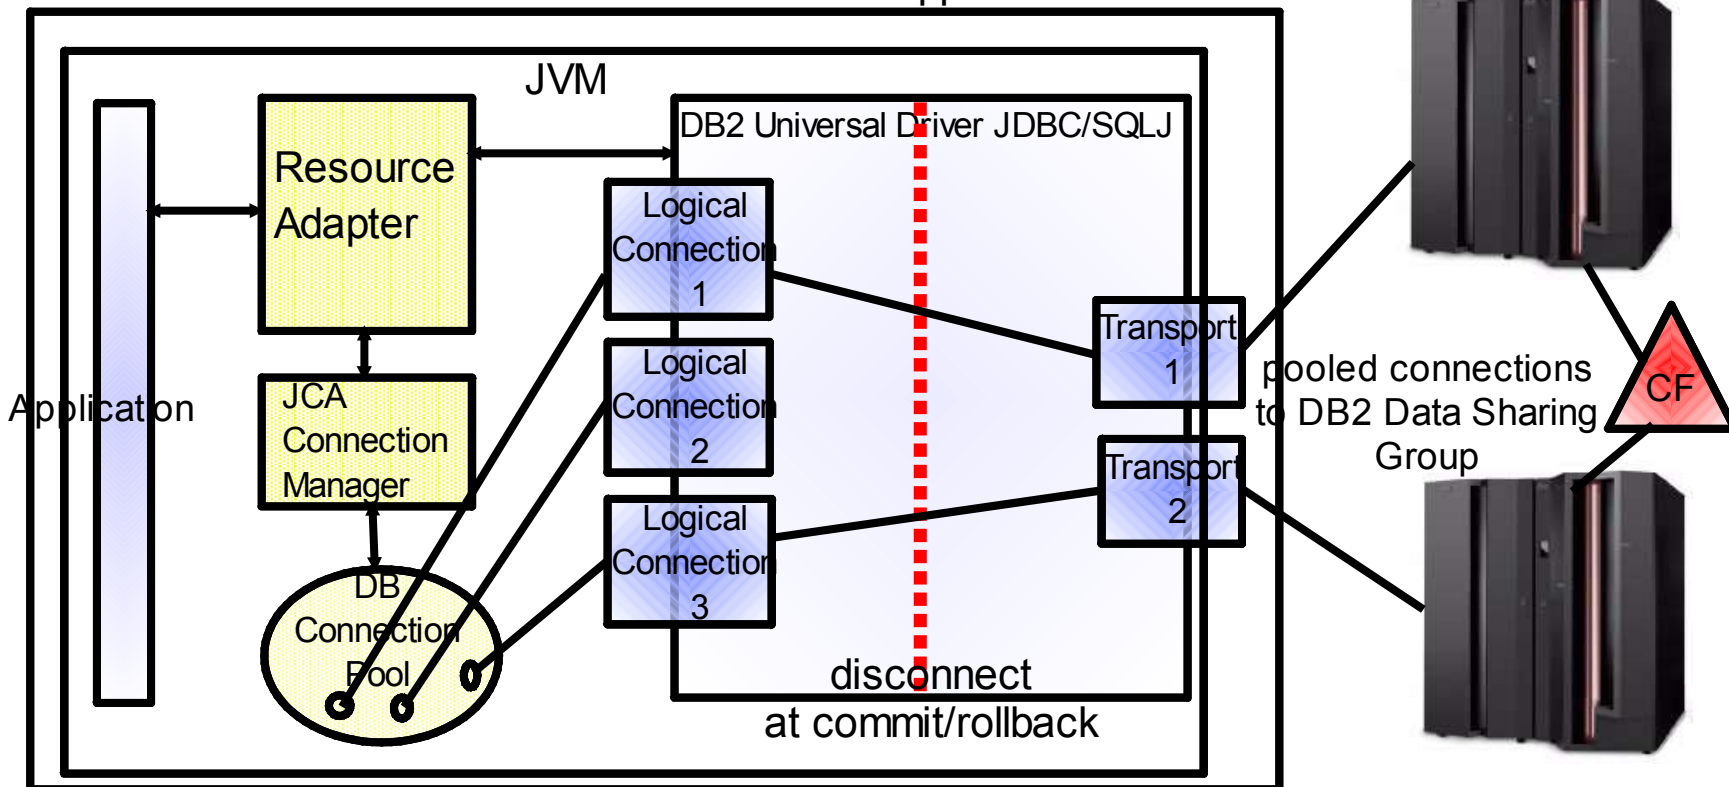
WebSphere Connection Pooling

- ▶ Connection object pool currently maintained by WebSphere
 - ☞ saves creating/destroying Connection objects which is relatively expensive
 - ☞ part of the Connection object creation is creating a physical connection to DB2
- ▶ Connection object holds references to other Java objects like preparedStatement objects that would be destroyed with the Connection object

WebSphere Application Server - End-to-end Example

- ▶ Always use connection pool closest to application
- ▶ Enable Sysplex Workload Balancing for best availability

Application Server



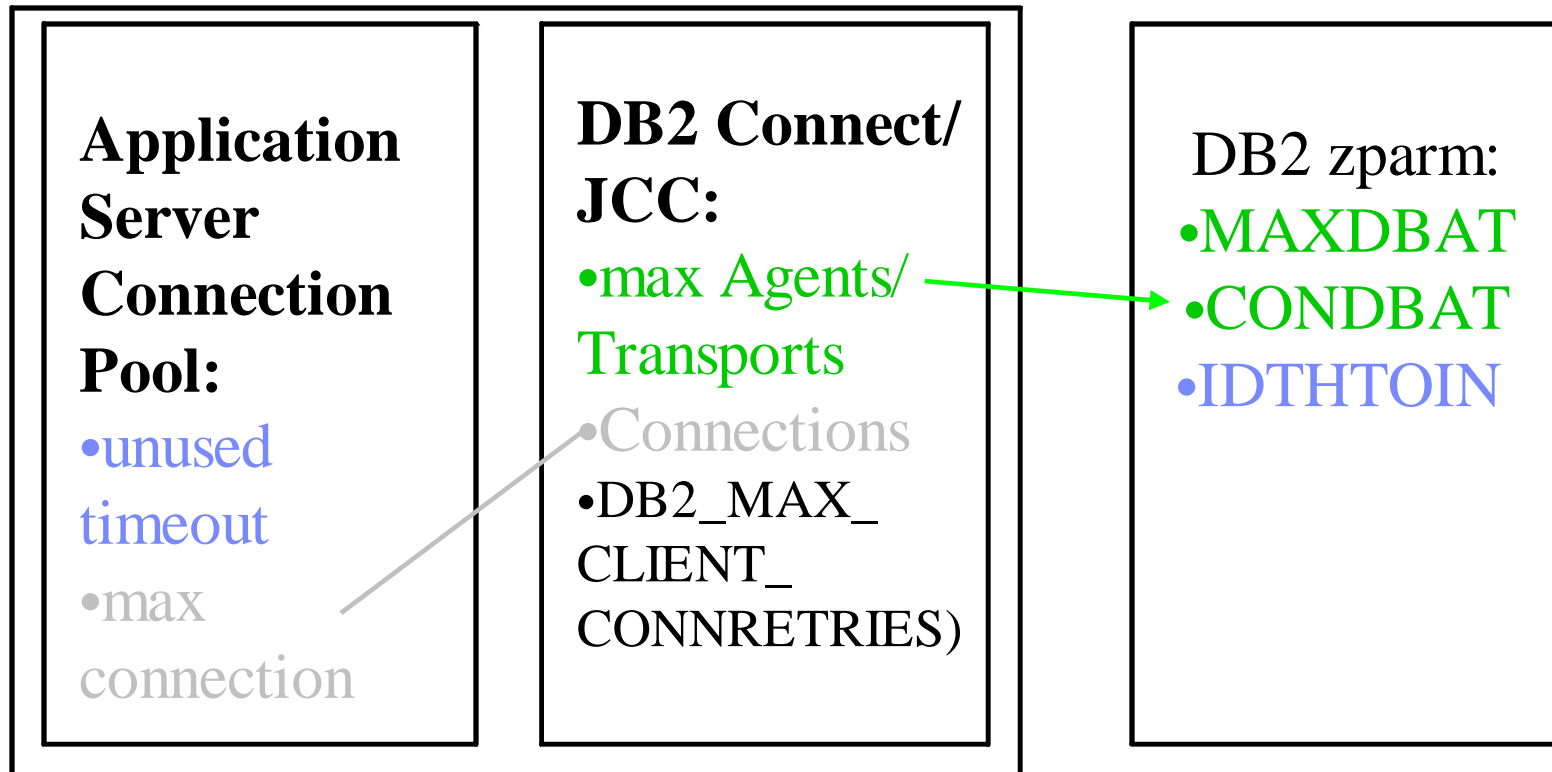
WAS/DB2 **Active** Thread - Tuning Considerations

- **WAS connections in connection pool that keep an active thread in DB2 are target of the “idle thread timeout”**
 - Type 2 on z/OS (RRSAF)
 - Type 4 connection that holds on to resources, e.g. WITH HOLD cursor, KEEP DYNAMIC YES.
- **Set WAS "connection unused time" to a smaller value than DB2 "idle thread timeout" to avoid stale connection conditions.**
- **Consider setting min connections to 0 (zero) and connection unused time to not higher than 10 - 15 min to free up unused resources in DB2 in a controlled way and to reduce the exposure of long living threads**

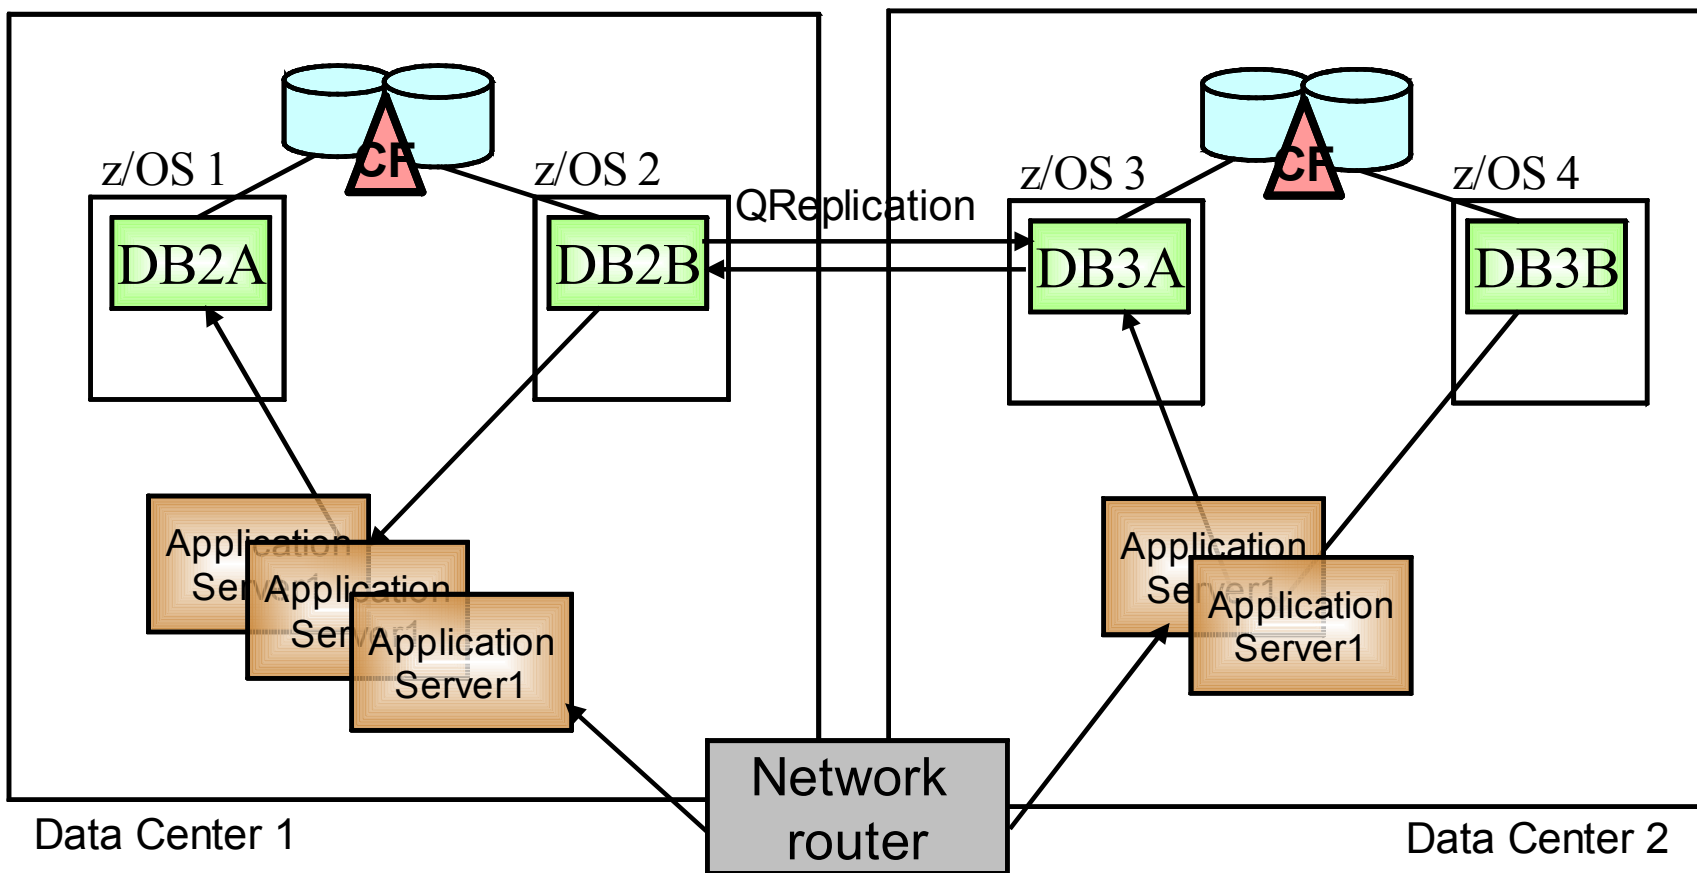
Important DB2 zparm Values

- **CONDBAT** – Max. # of distributed connection into DB2 system
 - includes inactive and active connections, may be large
 - DB2 queues DBAT requests to become active up to CONDBAT
- **MAXDBAT** - Max # database access threads (DBATs) that can be active concurrently.
 - In many installations, max. value determined by available storage in DBM1 (check IFCID 225)
 - Set this value conservatively
- **CMTSTAT INACTIVE** – make a thread inactive after it successfully commits or rolls back and thread does not hold resources
 - prerequisite for sysplex workload balancing
 - inactive connections use less storage and free up DBM1 resources
- **IDTHTOIN** – time in sec an active server thread remain idle before it is canceled
 - inactive connections are not subject to idle thread timeout
 - Strongly recommended to not set to 0 – disable, default works well

WAS/DB2 Inactive Thread - Tuning Considerations



Continues Availability Across Sysplexes



- ▶ Use Sysplex Workload Balancing for availability within a data center
- ▶ In case of Data Center outage, route the work away from application servers via network router

Reference

- **DB2 Version 9.1 for z/OS Data Sharing: Planning and Administration (SC18-9845)**
http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/index.jsp?topic=/com.ibm.db29.doc.dshare/db2z_dshare.htm
- **DB2 9 for z/OS Data Sharing: Distributed Load Balancing and Fault Tolerant Configuration**
<http://www.redbooks.ibm.com/abstracts/redp4449.html?Open>
- **DB2 9 for z/OS: Distributed Functions**
<http://www.redbooks.ibm.com/abstracts/sg246952.html?Open>