



IBM DB2 Server for z/OS

DB2 Data Sharing Best Practices: Configuring remote access for optimum availability

Jim Pickel
pickel@us.ibm.com
IBM DB2 for z/OS Development

February 2008

© 2008 IBM Corporation

Agenda

- **What is the network topology to obtain high availability when accessing a DB2 Data Sharing group?**
- **Why is DB2 Sysplex Workload Balancing so important?**
- **How to balance work to only a subset of subsystems in a data sharing group**
- **High Availability Update**

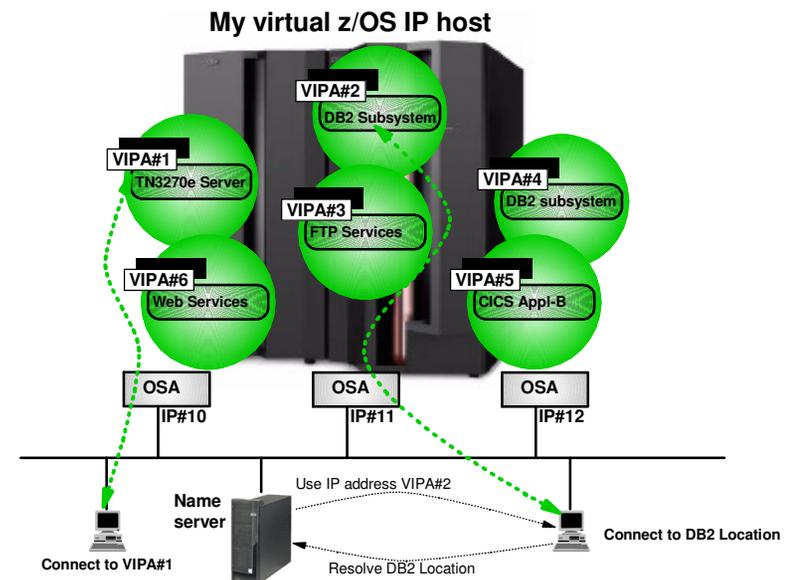
High Availability Configuration for Remote Access

- **TCP/IP Sysplex Distributor**
 - Needs to be configured at the network layer and used for the initial connections to the DB2 group

- **DB2 Sysplex Workload Balancing**
 - Needs to be enabled on the DB2 Connect Server when accessing a DB2 group

- **Load distribution managed by WLM**

The network view of a Parallel Sysplex
 • a single large server with many network interfaces and many services like a DB2 data sharing group



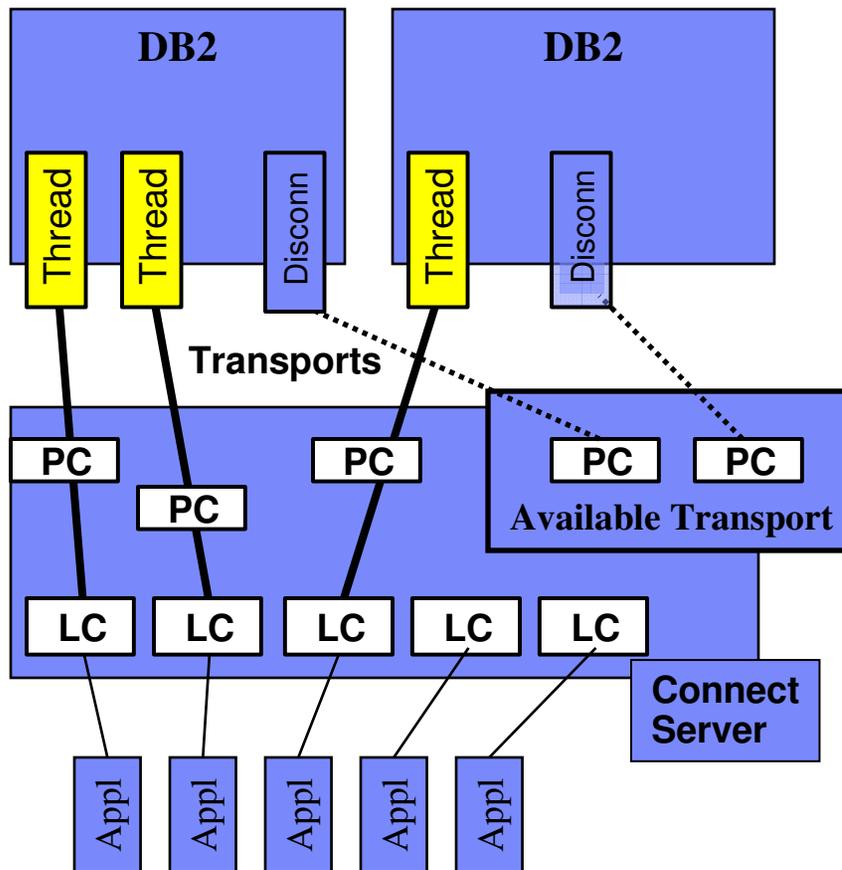
Balancing connections across group

- **The TCP/IP Sysplex Distributor is used to establish the initial connection to the DB2 group**
 - IP address used by clients to access DB2 is configured as a distributed dynamic VIPA
- **Provides connection load distribution**
- **Ensures highest availability possible**
- **But...**
 - Typical connections used for SQL have a long lifespan. For example, an application server that has utilizes connection pooling...

Balancing transactions across group

- **DB2 Sysplex Workload Balancing provides transaction level load distribution**
 - Each DB2 member is configured with a unique dynamic VIPA and automatic VIPA takeover
- **DB2 Sysplex Workload Balancing is enabled on DB2 Connect Server**
- **WLM ensures work is routed to the member with the most capacity and best health.**
- **Superior utilization of group resources**

DB2 transaction level load distribution

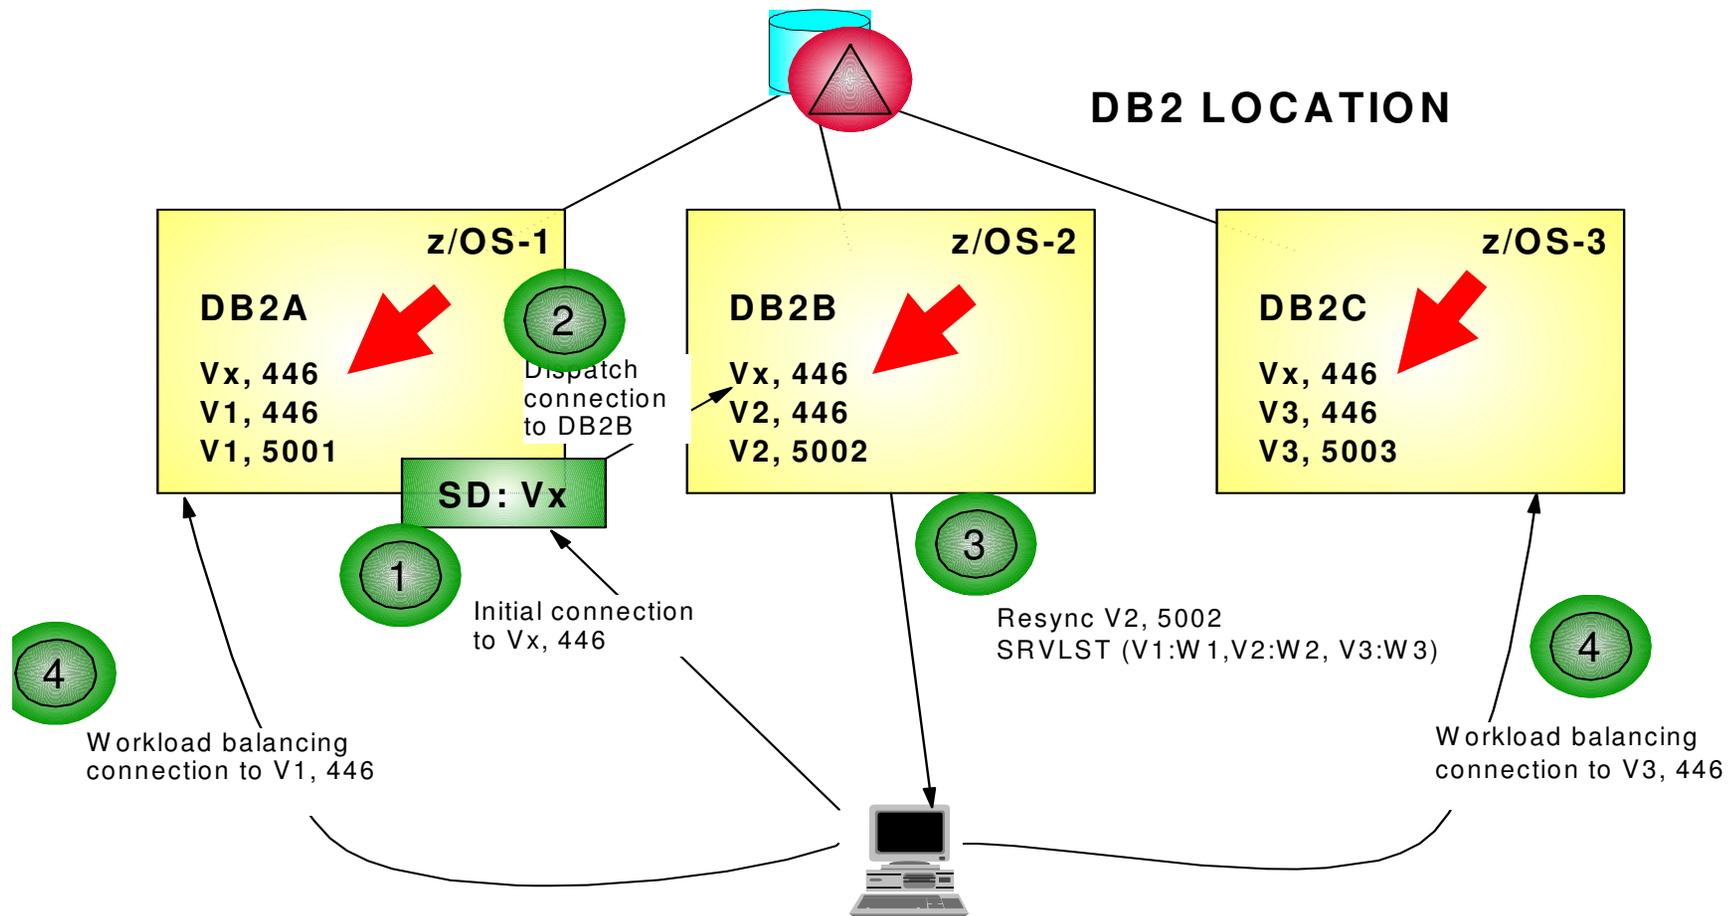


- DB2 Connect Server maintains pool of physical connections (PC) to members called transports to DB2 subsystems based on WLM recommendations
- Applications establish logical connections (LC) to DB2 Connect Server
- Connections not associated with a DB2 thread called disconnected connections by DB2
- Connections are associated with a subsystem transport on transaction boundaries based on WLM recommendations

Configure network routing to the DB2 group

- **Configure a group distributed DVIPA for the DB2 data sharing group.**
 - All members listen to this IP address for the SQL port
 - Connections are distributed across all members
 - Connections are successful as long one member is up
 - Used by clients to access the group providing a single image
- **Configure a unique member IP address for each subsystem**
 - This member DVIPA is not distributed since used to route connections to a specific member
 - Allows routing even if a member fails over to another LPAR using VIPA takeover
 - WLM weight and member DVIPA are provided to DB2 Connect Server for all subsystems registered

Data Sharing Fault Tolerant Network Topology



Application enablement

- **Enable inactive thread support for the DB2 group**
- **Minimize the use of KeepDynamic**
 - Static SQL most efficient use of resources
- **Close result sets and held cursors when done**
- **Drop global resources if not needed anymore**
- **Timely commits**
- **Allows for more efficient use of group resources**

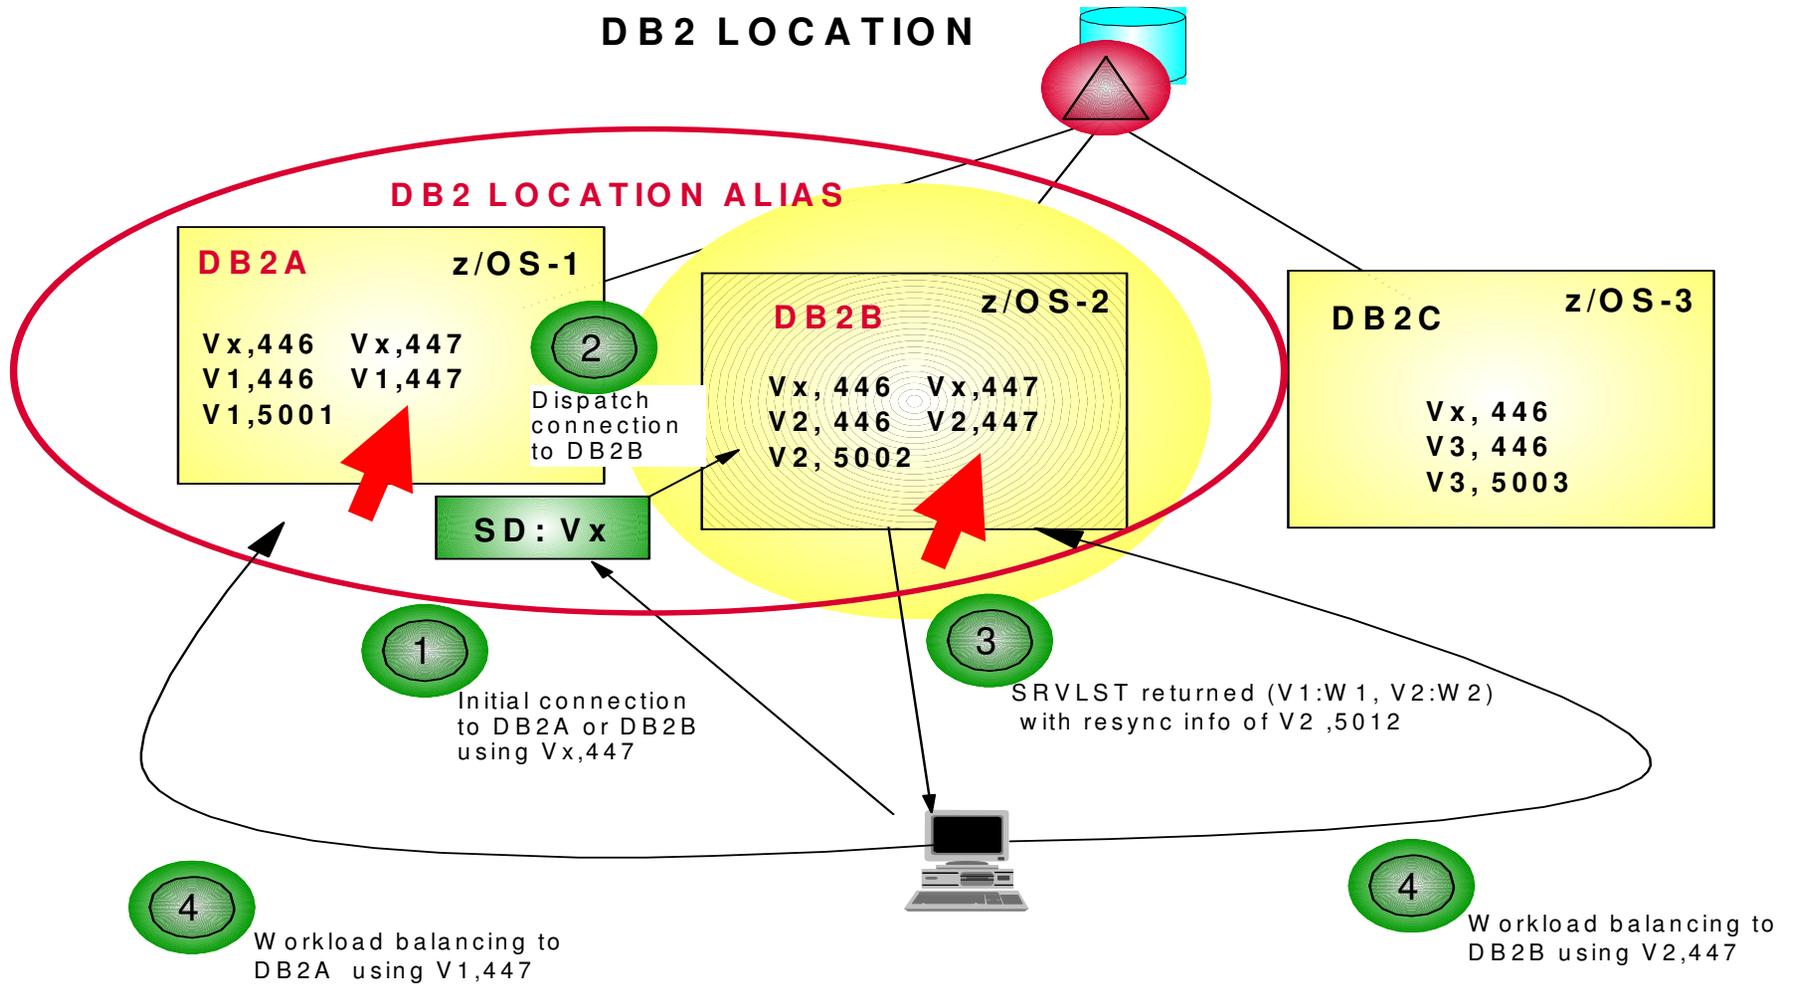
Workload balancing based on WLM weights

- **Weights provided to DB2 Connect Server via a server list**
 - Returned whenever a connection is established or reused
- **DB2 uses the WLM IWMSRSRS service to determine where to route work using the specific function**
 - List of servers associated with the DB2 location or alias are returned along with a relative weighting for each member
 - Server drops off list when no capacity
- **WLM Weights based**
 - Performance index that indicates the achievement of the WLM-defined goals of the member
 - Queue delays waiting for a DB2 thread to process an SQL request is considered
 - In V9, the member's health is considered. DB2 reports storage utilization to WLM based on storage thresholds

Balance to a subset of subsystems

- **Managed load distribution across a subset of members**
- **Again, recommend the use of Sysplex Workload Balancing in conjunction with Sysplex Distributor**
 - Configured with dynamic VIPA and automatic VIPA takeover
 - Connection successful if one member of subset started
 - A subset port is defined to distinguish between group and subset
- **Connections and load distribution occurs only to the members listening on the subset port**
- **Connection failures can occur even when other members outside of the subset are up**
- **Utilizes only a subset of the group resources**

DDF Subset Configuration



High Availability Update

- **Improved workload balancing algorithm in Connect Server**
- **Seamless reroute when a DB2 subsystem is quiesced**
 - Applications see no exceptions
- **Seamless reroute when a DB2 subsystem fails**
 - Applications on transaction boundary see no exceptions
- **zIIP capacity to be included in the WLM recommendations**

Thank You for Joining Us today!

Go to www.ibm.com/software/systemz to:

- ▶ Replay this teleconference
- ▶ Replay previously broadcast teleconferences
- ▶ Register for upcoming events