DB2 Universal database

IBM

DB2 Connect for NT and the Microsoft Windows NT Load Balancing Service

Achieving Scalability and High Availability

Abstract

DB2 Connect Enterprise Edition for Windows NT provides fast and robust connectivity to IBM mainframe databases from desktop and web applications. Although DB2 Connect fully exploits the capabilities of the Intel platform and the NT operating system, multiple servers are often required to handle the workload. Running DB2 Connect in conjunction with Microsoft's new Windows NT Load Balancing Service provides the ability to expand capacity beyond the abilities of a single server. This document will explain the capabilities, advantages, and limitations of this approach.

INTRODUCTION

IBM's DB2 for OS/390, DB2 for AS/400, and DB2 for VSE & VM databases continue to be the systems of choice for managing critical data for the world's largest organizations. While these host and AS/400 databases manage the data, there is a great demand to integrate this data with applications running on PCs, UNIX workstations, and Apple Macintosh workstations. The DB2 Connect Enterprise Edition product is a host data access server that provides managed connectivity from desktop and web applications to these databases.

DB2 Connect for NT provides access to those databases in a cost-effective way by using a standard architecture for managing distributed data, known as Distributed Relational Database Architecture (DRDA). Use of DRDA allows your applications to establish a fast connection to host and AS/400 databases without expensive host components or proprietary gateways.

DB2 Connect products provide a run-time environment for database applications that are written in C, C++, Java, COBOL, FORTRAN, SmallTalk, REXX, and other programming languages. Programmers can use the latest tools, such as Microsoft Visual Studio, Borland Delphi, and many others, to develop database applications using the following Application Programming Interfaces (APIs):

Microsoft ODBC
DB2 Embedded SQL (both static and dynamic)
DB2 Call Level Interface
Java Support (JDBC)
Embedded SQL for Java (SQLJ)

This allows you to use a wide range of off-the-shelf or custom-developed database applications, and provides application programmers with a broad selection of tools and functions.

Many organizations choose to implement DB2 Connect on the Windows NT platform to take advantage of its ease of use and relatively inexpensive hardware. Although DB2 Connect fully exploits the capabilities of Windows NT and the Intel platform, through features like SMP support and multithreading, a single NT box cannot always meet the needs of the organization. Also, many organizations choose to implement DB2 Connect in phases and expand their capacity as the need arises.

Implementing DB2 Connect with Microsoft Windows NT Load Balancing Service (WLBS), available as part of the Windows NT Server Enterprise Edition, can help meet both of these needs. WLBS allows up to 32 computers (called nodes or hosts) to provide TCP/IP services via a single IP address (or set of addresses). It is completely transparent to both the client application and the target service, which in this case is DB2 Connect. WLBS statistically balances client requests across the machines in the cluster to achieve load balancing and scalability. WLBS also provides a degree of high availability by redistributing requests to other nodes in the event of a node failure.

HOW DOES IT WORK?

DB2 Connect

In a WLBS environment, DB2 Connect operates exactly like it would in any Windows NT environment. Options for host connectivity are unchanged; both TCP/IP and SNA are possible. DB2 Connect is installed and configured as if it were not a part of the WLBS cluster. As mentioned previously, DB2 Connect is completely unaware of WLBS. The only restriction is that all client requests must be TCP/IP. WLBS does not balance any protocol other than TCP/IP.

WLBS

To understand how WLBS distributes requests among multiple DB2 Connect servers, let's take a quick look at it's architecture. WLBS installs as an NT networking driver and will act as a virtual network interface card (NIC). Figure 1 shows how the components in a DB2 Connect/WLBS implementation interact.   In a typical configuration, one NIC will be dedicated to cluster traffic and one to other network traffic. In WLBS terminology, the adapter which is used for WLBS clustering is the *cluster adapter* and the other adapter is the *dedicated adapter*. In this example, DB2 Connect reaches the host via SNA (could also be TCP/IP) over the dedicated adapter and receives client requests via TCP/IP over the cluster adapter. Each cluster adapter is configured with the cluster's "virtual" IP address. WLBS uses Level 2 multicast to allow all cluster nodes to simultaneously listen to this "virtual" address. The WLBS driver resides between the NIC and the TCP/IP layer to filter the requests for a particular node.

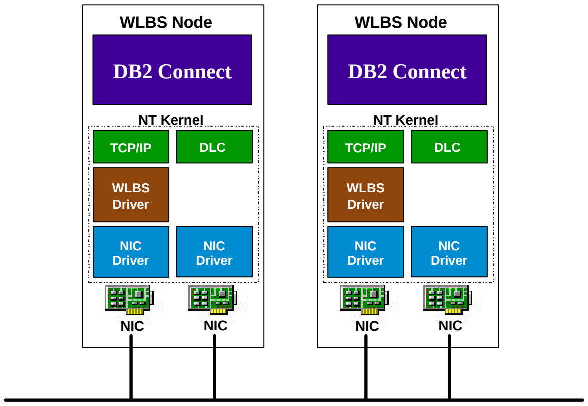

Figure 1: WLBS Architecture

Load Balancing

WLBS maps incoming client requests to a particular node based on the client IP address, port, and other information. Because database connections through DB2 Connect are conversational, it is critical that WLBS be configured to map all requests from a particular client to the same WLBS node. This is called *single host affinity*. What this means is that WLBS won't provide true work load balancing, since it is unaware of the true workload, but can merely attempt to evenly distribute the clients among the WLBS nodes. If each client has a similar workload, then the effect should be balanced activity across the servers. More specifically, if any single client (identified by an IP address) runs multiple connections thru DB2 Connect, this will be treated as a single connection by WLBS. Also, once a client has connected to a database through a particular node, all subsequent requests will go through that node unless the node fails or is removed from the cluster.

Availability

WLBS will automatically detect the failure of a cluster node and redistribute the requests for that node to other cluster nodes. WLBS nodes exchange messages to monitor the status of the cluster. WLBS provides protection from any failure that would prevent a node from exchanging these messages: node failure, WLBS failure, major OS failure, etc. It does not provide protection from failures like lost link to the host or DB2 Connect failure on a box.

When WLBS detects a failed node, it invokes a process called *convergence*. During this process, the remaining nodes determine the state of the cluster and redistribute the load from the failed node. From a DB2 Connect perspective, the connection to the host will be lost and the client application will receive an error (most likely SQL30081, which indicates a communication error, or SQL0900 which indicates an application state error). If the application immediately attempts to reconnect to the database, it will then get a successful connection through one of the remaining nodes.

CONFIGURING WLBS WITH DB2 CONNECT

For detailed WLBS installation and configuration consult the WLBS online help. For DB2 Connect, the multiple NIC configuration is recommended since it allows you to split the client traffic and host traffic across multiple adapters, thus improving performance. Perhaps, the trickiest part of configuring WLBS is the network bindings. Microsoft Knowledge Base Article Q197991 provides good step-by-step instructions for this task.

Configuring WLBS Parameters

Figure 2 show the WLBS setup screen. This is where you define the IP information for the cluster (in the *Cluster parameters* section), the information unique to each node (the *Host parameters)*, and which ports will be allowed to participate as balanced services, (the *Port rules* section). The cluster parameters should be the same on each node. The host parameters will contain information unique to the particular node, like the *Priority (ID)*, which acts as both a priority for choosing a default node and as a unique identifier for each node, and the *Dedicated IP address* and *Subnet mask*. In a multiple NIC environment, these refer to the IP information on the dedicated adpater, not the cluster adapter. This is used for remote administration and other tasks. Of primary importance are the Port rules. To enable DB2 Connect requests to be distributed, you must define a port rule (on all nodes) for the TCP/IP port that DB2 Connect is using. To define a rule for DB2 Connect, you must first verify that DB2 Connect on all cluster nodes is configured for the same port. To verify, open the DB2 Control Center,

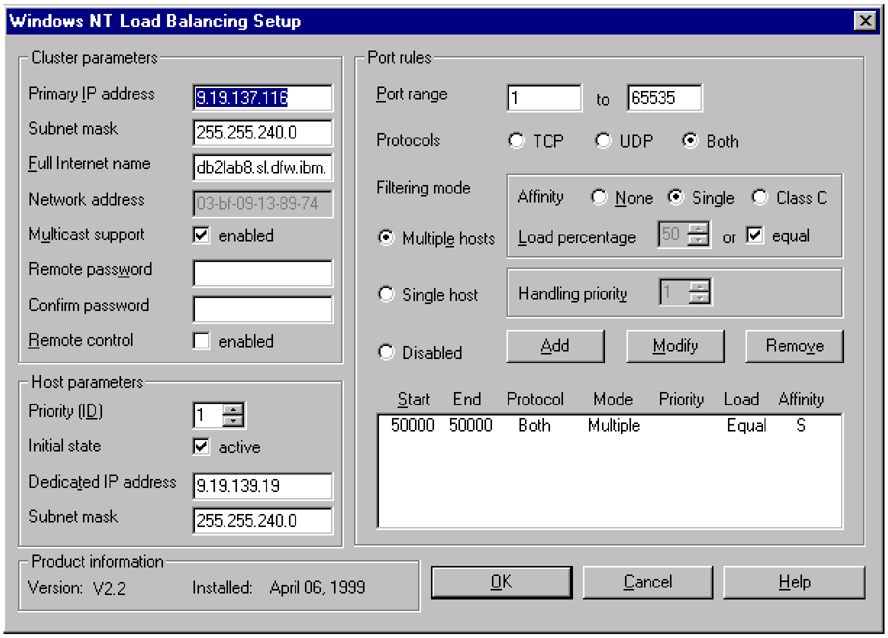

Cluster parameters

Primary IP address 9.19.137.116
Subnet mask 255.255.240.0
Full Internet name db2lab8.sl.dfw.ibm.
Network address 03-bf-09-13-89-74
Multicast support ☑ enabled
Remote password
Confirm password
Remote control □ enabled

Host parameters

Priority (ID) 1
Initial state ☑ active
Dedicated IP address 9.19.139.19
Subnet mask 255.255.240.0

Product information

Version: V2.2 Installed: April 06, 1999

Port rules

Port range 1 to 65535
Protocols ○ TCP ○ UDP ● Both
Filtering mode
● Multiple hosts Affinity ○ None ● Single ○ Class C
Load percentage 50 or ☑ equal
○ Single host Handling priority 1
○ Disabled

Start	End	Protocol	Mode	Priority	Load	Affinity
50000	50000	Both	Multiple		Equal	S

Figure 2: WLBS Setup Screen

right click on the DB2 Connect instance, select *Setup communications...*, and click the *Properties* button next to TCP/IP on the Setup Communications window. The *Service name* and *Port number* should be the same on all nodes. This is shown in Figure 3.

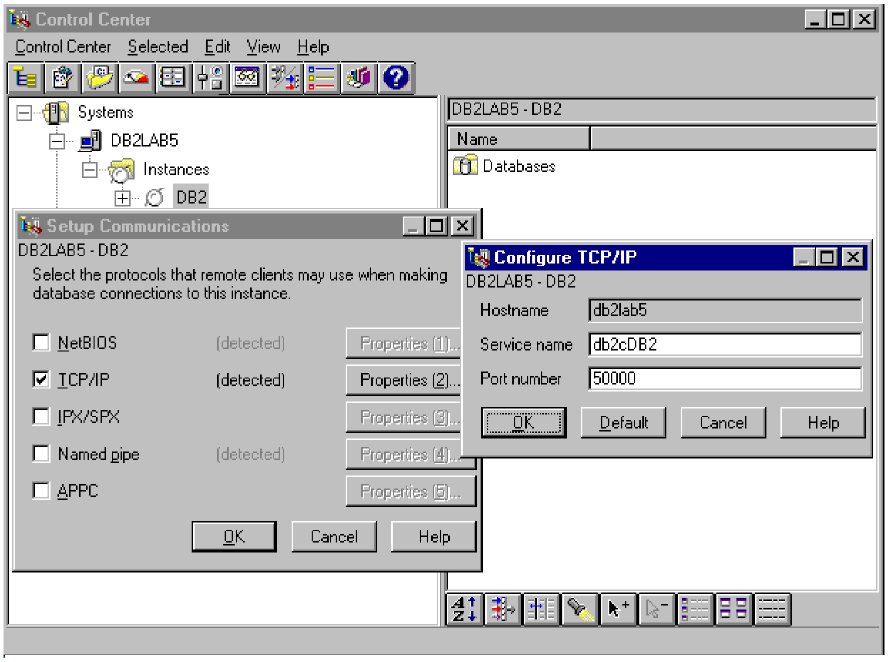

DB2LAB5 - DB2

Select the protocols that remote clients may use when making database connections to this instance.

☐ NetBIOS (detected)
☑ TCP/IP (detected)
☐ IPX/SPX
☐ Named pipe (detected)
☐ APPC

DB2LAB5 - DB2

Hostname: db2lab5
Service name: db2cDB2
Port number: 50000

Figure 3: DB2 Connect Port Information

To create the port rule in the WLBS setup window, enter the DB2 connect port number, 50000 by default, in both the start and end fields for the port range. This signifies the rule applies only to this one port. Choose either *TCP* or *Both* for protocol. Select *Multiple hosts* for filtering mode to indicate that this service will run on more than one cluster node. Select *Single* node Affinity. This ensures that all requests from a client go to the same node. Choose whether you want equal balancing or if you want a relative percentage of the workload for the specific node. Select *Add*. The rule should now show up in the list box in the lower right corner as shown in Figure 2.

If you want to enable the Administration Server, create a similar rule for it’s port 523. It is recommended that you access the Administration Server through the dedicated adapter, not the clustered services, since any changes would have to be done to each node anyway. Microsoft Knowledge Base Article Q198496 provides more information on adding port rules.

ADDITIONAL NOTES

Performance

WLBS’s implementation allows it to operate without placing much additional burden on the CPU of the cluster node. Because its status messages are broadcast on the network and not on a private network, it adds addtional traffic. If the traffic is excessive, you can tune WLBS to help minimize the impact. Consult Microsoft Knowledge Base Article Q193601 for details on WLBS parameters. Unfortunately, WLBS itself does not provide any performance counters for monitoring its performance.

Troubleshooting

In addition to logging convergence and other status messages in the NT Event Log, WLBS provides a command interface, the *WLBS* command, to display status, configuration, event messages, and IP configuration information. This can be very helpful in resolving problems. This command also allows starting, stopping, suspending, resuming, etc of the WLBS services.

Limitations

WLBS is only available as part of Window NT Enterprise Edition and is only supported on Ethernet or FDDI LANs.

FOR MORE INFORMATION

IBM DB2 Connect Home Page:
Http://www.software.ibm.com/data/db2/db2connect/

IBM Data Management Home Page:
Http://www.software.ibm.com/data/

IBM DB2 Product and Service Technical Library
Http://www.software.ibm.com/data/db2/library/

Microsoft WLBS Home Page (and download site):
www.microsoft.com/ntserver/ntserverenterprise/exec/feature/wlbs