



# **DALNET Wide Area Network (WAN) Connectivity Assessment Pilot Program**

April 16<sup>th</sup>, 2002

**A report on WAN connectivity assessment for DALNET  
Members and Horizon application performance  
Executive Summary**

Prepared by:



**123 S. Main St., Suite 210  
Royal Oak, MI 48067  
(248) 837-1400**

# DALNET Wide Area Network (WAN) Connectivity Assessment Pilot Program Final Report

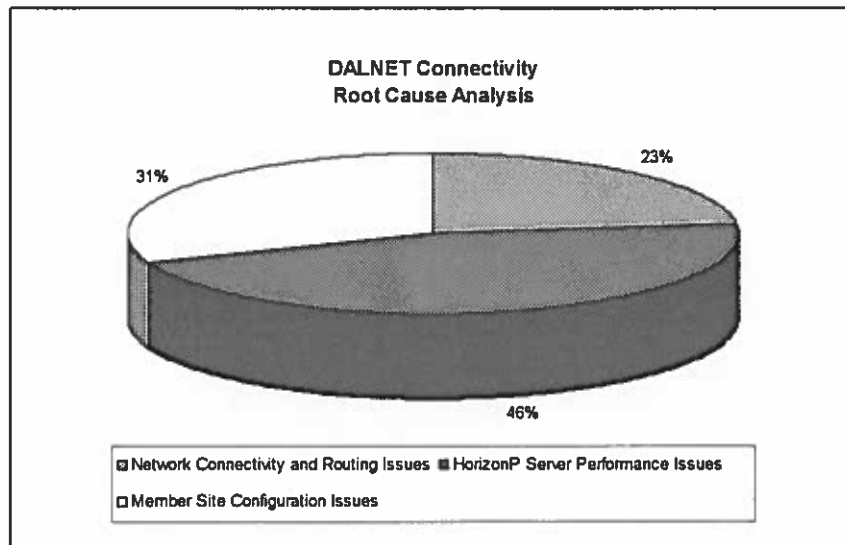
## Background and Scope

DALNET retained Intelligent Connections to perform a Wide Area Network (WAN) Connectivity Assessment Pilot Program of select DALNET member sites. The purpose was to determine the source of connectivity problems and to propose solutions and alternative frame relay providers or alternative connectivity methods. This final report concludes Intelligent Connections, obligations in regards to this project.

## Summary

Various technical root causes for the member connectivity problems have been identified. These root causes are classified into three major categories:

<u>Description</u>	<u>#</u>	<u>Impact</u>
1. Network Connectivity and Routing Issues	3	High
2. Horizon Production Server Performance Issues	6	High
3. Member Site Configuration Issues	4	Medium
Total Root Causes	13	



Six member sites participated in the pilot program. You can see from the population study above, some sites are classified into more than one root cause category. All six of the sites involved during the pilot were observed to experience Category 2 issues.

These categories are our best effort to categorize the root causes for the purpose of priority handling and impact analysis to the DALNET Wide Area Network (WAN). These categories are intended to provide a summary view and are not the definitive root cause of individual member site connectivity issues. Additional analysis and findings throughout this report identify individual member connectivity issues in greater detail.

## Definitions

Network Connectivity and Routing issues are concerned primarily with connectivity of DALNET routing equipment and IP based routing problems at the WSU Data Center and also include problems related to routing equipment located on the member premise.

Horizon Production Server Performance issues are specific performance issues concerning Horizon Production Sybase SQL Server errors. These issues were addressed in a previous report 'Horizon Library Server Audit', dated Jan 29<sup>th</sup>, 2002. Further analysis during this study determined that client server response times are well beyond acceptable levels and immediate action should be taken to implement the recommendations in our prior 'Horizon Library Server Audit', dated Jan 29<sup>th</sup>, 2002.

Member Site Configuration issues are concerned with configuration issues on IT systems owned and operated by member sites. Solutions to these issues will require member sites to reconfigure their systems.

## Observations

During our post-mortem data and root cause analysis we discovered that the DALNET FR head end router connection Ethernet 0/0 141.217.1.116 that should connect to the WSU Campus Backbone network 141.217.1.0/24 has been disconnected. This disconnected router interface is the root cause of many of the member connectivity issues.

Our network protocol analysis and network performance monitoring clearly show longer than expected server response times for all client connections to the Horizon production server. The normalized average response time is four (4) times the expected response time. The normalized maximum response time is ten (10) times the expected response time. See Appendix A for graph of Horizon Server response time.

During our site visits we discovered that member sites using the Internet as their connectivity method were prone to experiencing TCP timeout and Network Address Translation (NAT) issues due to the improper configuration of their on-site IT systems.

## Methodologies

Intelligent Connections (IC) assigned and dedicated two network engineers to the pilot program. Each engineer was dispatched to his/her respective testing location on the scheduled date. One IC engineer went to the member site to conduct the site survey and interview the Library and IT Staff. The other IC engineer went to the WSU Data Center where the DALNET Horizon Library servers reside. Each engineer had a protocol analyzer to capture all network packets related to the session(s) being tested for that member site.

In addition, we worked together with the member site IT Staff to gather information from other pertinent sources such as:

1. Horizon workstation TCP/IP netstat statistics
2. Firewall Log Entries
3. Firewall Configuration Settings
4. MS Proxy Server TCP/IP netstat statistics
5. MS Proxy Server configuration settings.

Each engineer recorded copious notes regarding each step of the testing process and documented data from the above sources as well as field observations. These notes, and the protocol analyzer captures, constitute the source of our empirical data.

We also spoke to representatives from various Internet Service Providers servicing member sites as well as SBC/Ameritech representatives to review configuration and performance data provide by SBC/Ameritech in regards to the DALNET Frame Relay Network.

During the course of the field testing our IC network engineers made numerous field observations many of these were communicated directly to DALNET personnel involved in the field testing. Upon completion of the field testing the IC engineers began the data analysis and data correlation process.

### Findings and Conclusions

DALNET WAN routing problems are traced to several root causes. First, the DALNET Frame Relay router at WSU head end has a disconnected Ethernet interface through which all default routing occurs. Frame Relay clients attempting to access WSU resources that are not on the 141.217.3.0/24 network do not have an available route and therefore the connection appears dead.

*why wouldnt SBC know?*

The reason the DALNET FR router interface is down is unknown. But design analysis indicated that since the router is configured with Routing Information Protocol (RIP) and RIP is also enabled for pass-thru on the DALNET/WSU Engineering and Sciences Library Firewall this configuration would cause a routing loop. It is assumed that WSU Network staff disconnected this interface to eliminate the routing loop.

Future design changes and discussions regarding routing decisions should be made involving both WSU Network Ops and SBC/Ameritech DataComm engineers. Intelligent Connections would also be available to participate in these design discussions. Immediate action should be taken to determine why this connection was disconnected and remedial action should be taken to correct the situation.

The root cause for Horizon Server response time was previously reported in our Horizon Library Server Audit', dated Jan 29<sup>th</sup>, 2002. In the report we concluded that a lack of physical memory was causing an exhaustion of resources available for additional connections. During our testing we observed active client server connections in the range of 300-500. We also observed that a single workstation creates approximately 12 connections for a single logon and search session.

During the pilot every member site we analyzed experienced poor server response times. Appendices A and B show relevant data in support of this finding. This is the single most common and pervasive cause of server responsiveness and SQL Server connections failures.

The root cause, for member site configuration issues are numerous and varied. This description is in fact a catch all for a myriad of specific technical issues that need to be addressed on a case by case basis. It is our conclusion that members who connect via the Internet tend to have more configuration problems and TCP timeout issues due to the mis-configuration of intermediate member site IT systems and security components at each member site. However, member sites which use the internet for connectivity to DALNET experience better response times than Frame-relay connected sites. This adds to the complexity and administrative burden for each member site IT support staff.

The next few pages provide a summary for each member site. Appendix D contains the individual member site Audit Template with site specific findings and recommendations.

The root cause for the majority of the dropped TCP sessions was due to TCP timeouts caused by member site intermediate IT systems that are improperly configured to support the necessary TCP timeouts required by the Horizon Application. However, requiring these sites to reconfigure their IT systems with longer TCP timeouts could in fact have adverse impact on other Applications the member sites may be supporting.

Some member sites reported the training site was unreachable <http://training.dalnet.lib.mi.us> Our DNS lookup shows that this site resolves to 141.217.3.136 and cannot be reached via the internet. This may indicate that there is a DNS record error or Firewall configuration problem. Our investigation indicates that training is now linked off of <http://www.dalnet.lib.mi.us> at 141.217.3.121 and this should to be communicated to DALNET member sites.

We also observed member sites who were routing web traffic (http) over the DALNET Frame Relay network. Most sites have small FR Committed Information Rates (CIR) bandwidth provisioning of 56/64 or 128 kbps. These CIR allocations do not adequately support both the web traffic (http) and the primary DALNET application for Horizon.

During our on-site visits at WSU Data Center we observed several occasions when minor changes to the production environment were made during production hours. However, even these minor changes tend to have adverse impact on the Horizon Application due to its TCP sensitivity.

## Recommendations

### **DALNET Frame Relay Connectivity**

DALNET members appear to have outgrown the Frame Relay (FR) solution currently provided by SBC/Ameritech. Bandwidth requirements for the DALNET Horizon Application can no longer be met with the existing FR solution without significant incremental costs to DALNET member sites.

Several sites were also found to be routing web traffic (http) in addition to the Horizon application traffic over these FR circuits. The original network design was never intended for that purpose. Web traffic (http) tends to burst bandwidth consumption due to the protocol overhead and added web graphics.

In comparison to Internet connectivity solutions FR is far more expensive solution based on cost per Bandwidth \$/MB. All of Member sites that participated in the pilot already have some form of Internet connectivity and could easily migrate from FR to Internet for access to the Horizon Application.

### **Alternative Connectivity Solutions**

With the price of Internet Connectivity and the adoption rate within DALNET member sites, it is our recommendation that existing Frame Relay DALNET members' transition to Internet based connectivity. Member sites should experience increased response times, but also experience increased administrative and complexity issues as previously notes.

### **Encryption and Security**

There isn't a current requirement for encryption of the Horizon SQL queries for updates. However, we have concern with the current DALNET network architecture which places public web sites on the same network segment with the Horizon Production and Horizon test servers.

*can't*

We strongly suggest creating a Public DMZ and placing the public facing Web Search engines on this segment and firewall them from the Horizon SQL server and database. This follows current best practice methodologies.

### **Network Connectivity and Routing Issues**

Immediate action should be taken to resolve the disconnected DALNET router interface and/or re-route the backbone traffic through the WSU/DALNET Horizon Library Firewall.

#### **Horizon Application Modifications**

It is our recommendation that DALNET work with Epixtech to modify the Horizon Client application to support application layer TCP keep-alive to stop the TCP sessions from expiring during idle periods.

#### **Horizon Production Server Performance**

All server audit and network performance analysis indicates that the Horizon Production server is seriously lacking in host memory resources to adequately serve the user community. It is our recommendation to purchase and install an additional 4GB of system memory as soon as possible.

#### **Member Site Configurations Solution(s):**

1. Increase member Firewall Global TCP timeout from 3600 seconds to 7200 seconds.
2. Increase member Firewall Service TCP timeout for a specific service by editing init.def
3. Horizon client application should be modified to send keep-alive.

# Appendix F

## Sniffer™ Expert Diagnosis Terms and Definitions

### Slow Sybase/Microsoft SQL Server

The Expert generates this alarm when the average transaction time observed for a Sybase or Microsoft TDS server exceeds either the Slow TDS Server Time w/Cursors threshold (for transactions using cursors) or the Slow TDS Server Time threshold (for transactions that do not use cursors) in the Alarms tab of the Expert Properties dialog.

At the Service layer, the Expert observes transactions for individual servers (in this case, a TDS server). For each transaction the server takes part in, the Expert measures the time it takes to complete the transaction. Transactions using cursors are measured from the start of an Open Cursor statement to the cursor's close. Transactions that do not use cursors are measured from the first frame of a request to the last frame of the corresponding response. When the average time of either type of transaction observed for a single TDS server exceeds the corresponding threshold (Slow TDS Server Time w/Cursors for transactions using cursors; Slow TDS Server Time for transactions that do not use cursors), the Expert generates this alarm.

Possible causes:

- 1 The problem is network related.
  - a. Determine if your network is overloaded. Increase your network's bandwidth.
  - b. Determine if your network is optimally configured. Minimize frame fragmentation by using larger frame sizes or more efficient protocols.
  - c. Reduce polling frequency of polled devices in your network.
- 2 The problem is in the Multiprotocol Interchange.
  - a. Configure your MPI for better data throughputs rather than for connections.
  - b. Decrease the number of pumps per MPI.
  - c. Add more MPIs.
- 3 The problem is in the server.
  - a. Assess the current utilization of your server. Determine if you have enough CPU power to support the number of users in this server.
  - b. Optimize the file distribution of your database. Place the tables and their associated indexes in separate disk drives. Tables that are most frequently referenced should be placed in separate disk drives.
  - c. Determine if your server has enough memory to support the applications and services that have to run concurrently.
  - d. Install faster disk drives.
- 4 The problem is in the client workstation.
  - a. Use optimized query statements. The "explain plan" command will show the access paths to the data.
  - b. Remove all unnecessary "lock table" statements. Do not use "for update" clause in select statements if there is no need to update immediately.
  - c. Issue commit statements immediately after "update", "delete", and "insert" statements to release locked resources.
  - d. Use explicit cursors in all of your PL/SQL blocks.
  - e. Identical query statements in your application must match character-by-character to use the same shared pool area.
- 5 Check for runaway (orphan) processes and clean up after them. The problem is in the database. Perform database tuning suggested by experts in this area. For example:
  - a. Make sure that the shared pool is large enough to reduce reparsing.
  - b. Check if indexed clusters can speed up your access to tables that are commonly joined together on a standard set of matching columns.

### DB Slow Server Response

Database slow server response(s). The time it took for a server to respond to one or more SQL commands exceeded the DB Slow Server Response Time threshold.

Possible causes: Threshold set to 50ms.

- 1 The problem is network related.
  - a. Determine if your network is overloaded. Increase your network's bandwidth.
  - b. Determine if your network is optimally configured. Minimize frame fragmentation by using larger frame sizes or more efficient protocols.
  - c. Reduce polling frequency of polled devices in your network.
- 2 The problem is in the Multiprotocol Interchange.
  - a. Configure your MPI for better data throughputs rather than for connections.
  - b. Decrease the number of pumps per MPI.
  - c. Add more MPIs.
- 3 The problem is in the server.
  - a. Assess the current utilization of your server. Determine if you have enough CPU power to support the number of users in this server.
  - b. Optimize the file distribution of your database. Place the tables and their associated indexes in separate disk drives. Tables that are most frequently referenced should be placed in separate disk drives.
  - c. Determine if your server has enough memory to support the applications and services that have to run concurrently.
  - d. Install faster disk drives.
- 4 The problem is in the client workstation.
  - a. Use optimized query statements. The "explain plan" command will show the access paths to the data.
  - b. Remove all unnecessary "lock table" statements. Do not use "for update" clause in select statements if there is no need to update immediately.
  - c. Issue commit statements immediately after "update", "delete", and "insert" statements to release locked resources.
  - d. Use explicit cursors in all of your PL/SQL blocks.
  - e. Identical query statements in your application must match character-by-character to use the same shared pool area.
- 5 Check for runaway (orphan) processes and clean up after them. The problem is in the database. Perform database tuning suggested by experts in this area. For example:
  - a. Make sure that the shared pool is large enough to reduce reparsing.
  - b. Check if indexed clusters can speed up your access to tables that are commonly joined together on a standard set of matching columns.

### DB Slow Connect

SQL Server slow connect(s). The time it took a SQL Server to respond to a request to login to a SQL Server exceeded the DB Slow Connect Time threshold. This may or may not be normal, depending upon the application and the setting of the DB Slow Connect Time threshold.

Threshold set to 30 ms.

Possible causes:

- 1 The server is too busy. If this occurs frequently in a server with several users already active, consider taking performance tuning steps to improve the server response.
- 2 There is a lot of traffic on your network.
- 3 Other users may be trying to connect to the server at the same time, thereby reducing the response time.



### Idle Too Long

This TCP or NFS connection has been idle (no frames in either direction) for longer than the Idle Time threshold. If the idle connection is still alive, it is reserving system resources at both ends of the connection without using them. Threshold set to 3m 0s 0ms.

Possible causes:

- 1 A station is inoperative.
- 2 A router has lost the connection.
- 3 The transport software has hung.
- 4 An application has hung.
- 5 An application is open but not in use.

### Window Frozen

In a windowing protocol (such as DECnet or TCP), the window size of this station has been stuck at some number of bytes for longer than the Window Frozen Time threshold. Threshold set to 5s 0ms.

When a window is stuck, data flow will not be as efficient because the transmitter will not send data in excess of the current window size.

Possible causes:

- 1 There is insufficient buffer space, because the recipient's memory allocation process is not providing enough buffer memory.
- 2 There is insufficient buffer space, because there are new open connections, so that the buffer space must be shared.

### Repeat Ack

The acknowledgement number is less than its previous value.

Possible cause:

Some implementations do this purposely to keep the circuit "alive". However, this is not good practice because it causes unnecessary network traffic.

### Ack Too Long

The time that it has taken to acknowledge data exceeds the Long Ack Time threshold plus three times the average acknowledgement time for this connection. Refer to the detail statistics screen for the value of this average acknowledgement time. Threshold set to 100ms.

Possible causes:

- 1 The recipient of the original data frame was temporarily busy, and could not process the frame as quickly as usual.
- 2 The ACK arrived late because a server had to look up and/or process data before responding with an ACK.
- 3 The path changed in a way that increased the time between the request and its acknowledgement.
- 4 There were multiple paths between the two stations, and the time to acknowledgement was longer for some paths than for others.