Listening In

Passive Capture and Analysis of Oracle Network Traffic

Jonah H. Harris
myYearbook.com

About Me

- Sr. DBA at myYearbook.com
- Oracle DBA and developer since Oracle7
- Research Oracle Internals
- Speaker at IOUG, VOUG, NYOUG
- Technical Reviewer for IOUG SELECT
- Blog about Oracle technology

You’re asking… what’s in it for me?

- Learn how to detect network-related issues.
- Diagnose and solve network-related issues.
- Gain a better understanding of the Oracle Network Protocol

Disclaimer

- This is my hobby
- I’ve never been an Oracle insider
- The material in this presentation has been based on years of researching Oracle internals as well as analyzing network traffic and trace files.
- In addition to similar research from Ian Redfern, the majority of this paper is based primarily on my own personal research and discussions with Tanel Põder
- Do your own research!
- Use at your own risk!

A Common User Question

- Question
  - Why is the database sooooo slow?
- The sarcastic response you're considering…
  - The edition of Oracle we're using lacks the ALTER SYSTEM SPEEDUP DATABASE option.
- Answer
  - I don't know…
  - It's not the database, it's the application…
  - I'll look into it…
  - How do you know it's a database issue?

Troubleshooting the Issue

- Check Session Waits
- Check for a Long Running Query
- Check Session Performance Counters
- Check X, Y, Z…

Check Session Wait Events

- V$SESSION
- V$SESSION_WAIT
- V$SESSION_EVENT
- V$SESS_TIME_MODEL

Nope, nothing there…

Check Long Running Queries…

- V$SQL
- V$SESSION_LONGOPS

Hmm, looks like short queries…

Check Session Counters

- V$SESSTAT

Counters aren’t increasing, …

Troubleshooting the Issue

- Check Session Waits; zero.
- Check for a Long Running Query; zip.
- Check Session Performance Counters; zilch.
- Check X, Y, Z; nada.

All looks good from within Oracle… what next?

Check the Operating System

- Check Process CPU Usage and State
  - Determine whether it's doing anything…
- Dump Call Stack
  - Get a list of all the function calls made by Oracle as well as the call we're currently in…

Check the OS—CPU

- UNIX/Linux
  - nmon, top, glance, …
- Windows
  - More difficult due to threads-based model…

Examine the Process State

- Output from top

Check the OS—Call Stack

- UNIX/Linux
  - pstack, procstack, gdb, dbx, …
- Windows
  - More difficult…

Examine the Call Stack

- Represents the program's function calls
- Stack data structure
- Top entry is the current function

Top of the Stack—read(), write()

- Operating System calls (syscalls)
- Used to read/write data from a file descriptor (socket).

Trace Client/Server?

- SQLNET.ORA
  - CLIENT_TRACE_LEVEL
  - SERVER_TRACE_LEVEL
- **LISTENER.ORA**
  - TRACE_FILE_LISTENER

Trace Client/Server? Uh, no.

- Pros
  - It works [for some things]
- Cons
  - Is not passive; Tracing/Logging has overhead
  - Difficult to find valuable information
  - Difficult to use for more than one connection

Oracle Network Architecture

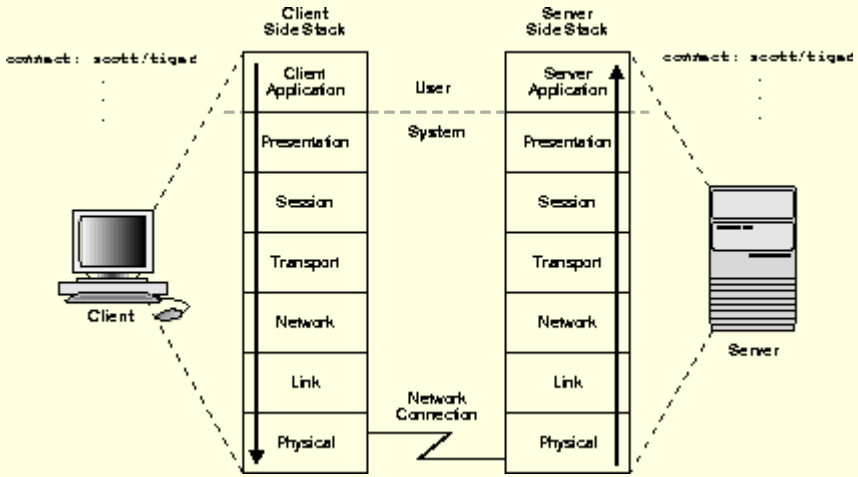

- Layered
  - Based on the Open Systems Interconnect (OSI) model

Oracle Net Components

- Oracle Protocol Support
- Oracle Net Foundation Layer
- Two Task Common Layer
- Application & RDBMS Layer

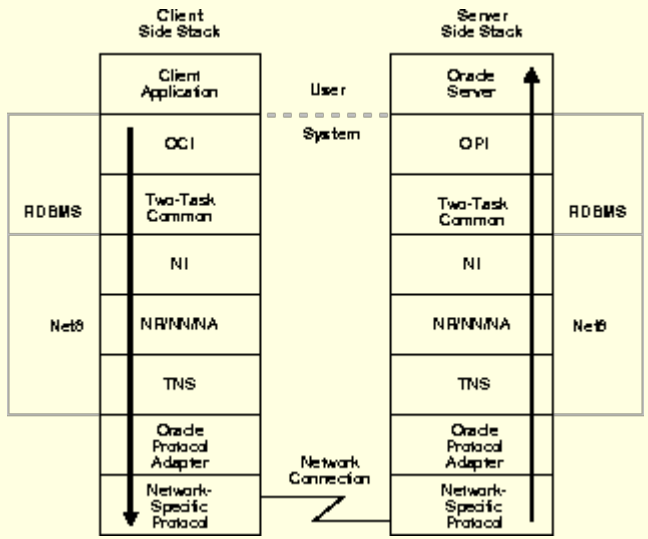

Oracle Protocol Support

- Maps TNS to underlying network transport

Oracle Net Foundation Layer

- Handles connections and messaging
- Transparent Network Substrate (TNS)

Two Task Common Layer

- Performs client/server character set conversion.

Application & RDBMS Layer

- Application (Client) Interface
  - UPI
  - OCI
  - JDBC
  - .NET
- RDBMS (Server) Interface
  - OPI (Oracle Programmatic Interface)

TNS Packets

- Transparent Network Substrate (TNS)
- Note:1007807.6, SQL*NET PACKET STRUCTURE: NS PACKET HEADER
- Every TNS packet has a header

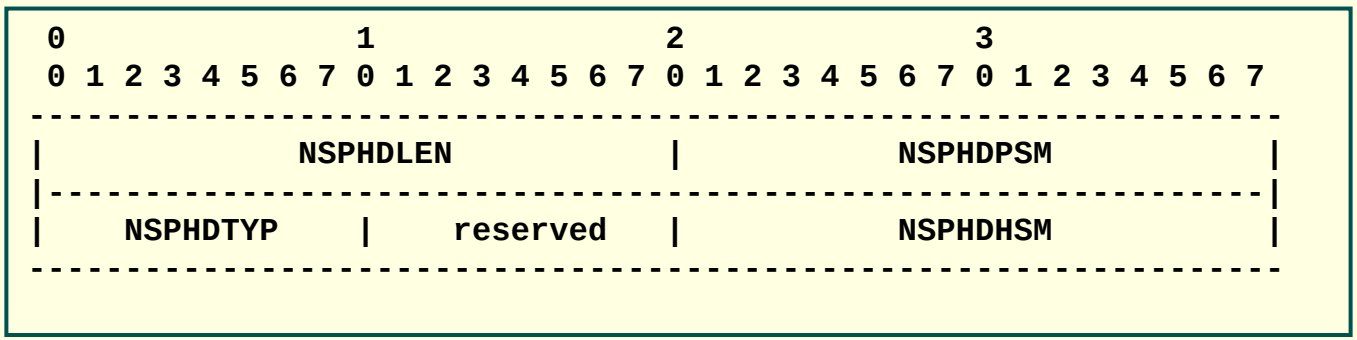

TNS Packet Types

- Connect
- Accept
- Acknowledge
- Refuse
- Redirect
- Data
- Null
- Abort
- Resend
- Marker
- Attention
- Control Information

Translating TNS Packets to Code

- Packet Header

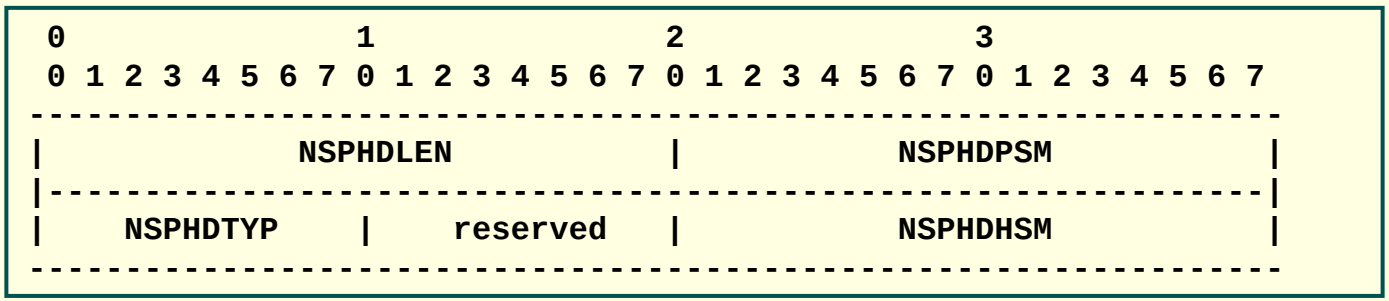

- Code

```
struct nsphd
{
    ub2     nsphdlen;                   /* Packet Length (in bytes) */
    ub2     nsphdpsm;                           /* Packet Checksum */
    ub1     nsphdtyp;                              /* Packet Type */
    ub1     nsphdrsv;                  /* Reserved for Future Use? */
    ub2     nsphdhsm;                    /* Packet Header Checksum */
};
```

TNS Connect Packet

- Performs a connection to an Oracle server.

```
struct nspcn
{
    ub2     nspcnvsn;                                   /* Packet Version */
    ub2     nspcnlov;                           /* Lowest Compatible Version */
    ub2     nspcnopt;                       /* Supports Global Service Options */
    ub2     nspcnsdu;                     /* Session Data Unit Size (in bytes) */
    ub2     nspcntdu;                 /* Transport Data Unit Size (in bytes) */
    ub2     nspcnntc;                         /* NT Protocol Characteristics */
    ub2     nspcntna;                             /* Line Turnaround Value */
    ub2     nspcnone;                     /* The number 1 in Host Byte Order */
    ub2     nspcnlen;                   /* Length of Connect Data (in bytes) */
    ub2     nspcnoff;                         /* Byte Offset to Connect Data */
    ub4     nspcnmxc;                              /* Maximum Connect Data */
    ub1     nspcnfl0;                                   /* Connect Flags 0 */
    ub1     nspcnfl1;                                   /* Connect Flags 1 */
    ub4     nspcncf1;                              /* cross facility item 1 */
    ub4     nspcncf2;                              /* cross facility item 2 */
    text    nspcncid[8];                            /* unique connection id */
    text    nspcncix[8];                            /* unique connection id */
    text   *nspcndat;                                     /* Connect Data */
};
```

TNS Accept Packet

- Server's response to a connection request.

```
struct nspac
{
    ub2     nspacvsn;          /* Version that this connection is to run at */
    ub2     nspacopt;                          /* Global service options */
    ub2     nspacsdu;                /* Session Data Unit Size (in bytes) */
    ub2     nspactdu;              /* Transport Data Unit Size (in bytes) */
    ub2     nspacone;                 /* The value '1' in host byte order */
    ub2     nspaclen;                          /* Length of connect data */
    ub2     nspacoff;                 /* Offset to start of connect data */
    ub1     nspacfl0;                                /* Connect Flags 0 */
    ub1     nspacfl1;                                /* Connect Flags 1 */
    text   *nspacdat;                                   /* Connect Data */
};
```

TNS Refuse Packet

- A denied connection request by the server.

```
struct nsprf
{
    ub1     nsprfurs;              /* User (application) reason for refusal */
    ub1     nsprfsrs;                     /* System (NS) reason for refusal */
    ub2     nsprflen;                              /* Length of refuse data */
    text   *nsprfdat;                              /* Start of connect data */
};
```

TNS Redirect Packet

- Server asks client to performs connection redirection.

```
struct nsprd
{
    ub2     nsprdlen;                          /* Length of redirect data */
    text   *nsprddat;                            /* Start of connect data */
};
```

TNS Resend Packet

- Server requests the client to resend the packet.
- Packet is the standard TNS header with packet type NSPTRS.

TNS Null Packet

- Generally used for keep-alive.
- Packet is the standard TNS header with packet type NSPTNL.

TNS Data Packet

- The most commonly used packet.
- Encapsulates TTI/TTC Subpackets

```
struct nspda
{
    ub2     nspdaflg;                                   /* Data Flags */
    text   *nspdadat;                                 /* start of data */
};
```

Issue #1: Slow Connections

- Connecting to server is slow

Establishing a Connection

- Connect to Server
- Negotiate Additional Network Options
- Negotiate Protocol Version
- Negotiate Data Types
- Authenticate

Code (Upon Request)

- oconnect.c

Making the Connection

- Client requests a connection to TNS entry ORCL.
- Network Naming attempts each directory_path method in order
- Network Naming parses TNSNAMES.ORA looking for the ORCL entry.
- Client builds and sends a TNS Connect Packet (NSPTCN) to the appropriate listener.
- The Listener responds with a TNS Resend (NSPTRS) or Redirect [to another port] (NSPTRD) packet.
- Client responds accordingly.
- Server responds with an Accept (NSPTAC) or Refuse (NSPTRF) packet

Negotiating Additional Options

- Additional Services
  - Authentication
  - Encryption
  - Data Integrity
  - Supervisor

Negotiating Protocol Versions

- TTI Protocol Packet
  - What versions are acceptable?
    - 4, 5, 6
  - Server replies with info

Negotiating Data Types

- TTI Data Type Packet
  - Client character conversions
- Server replies with its own representations

Authentication

- Client sends server basic information
  - User Name
  - Terminal Name
  - Machine Name
  - Program Name
  - …
- Server responds with challenge/response…

Why can it be slow?

- Parsing large TNSNAMES.ORA files
- Network Connectivity (Firewall/VPN)
- Overloaded Server

Large TNSNAMES.ORA

- Reason
  - Increases parse time
- How to detect
  - Oracle Net Tracing
    - Time between loading TNSNAMES.ORA and finding the appropriate entry.
- How to fix
  - Use different TNSNAMES.ORA files
  - Use EZCONNECT or a different directory method

Network Connectivity

- Reason
  - Network latency slows down packet transfer.
- How to detect
  - Inside Oracle
    - SQL*Net message to client with high wait times
  - Wire-level Monitoring
    - Look for delays in TCP ACK
  - Client/Server Oracle Net Tracing (least-preferred)
    - Look at time between send/recv from NS and NT
- How to fix
  - Fix the network.

Overloaded Server

- Reason
  - Server is CPU-bound and has a high load
- How to detect
  - Wire-level Monitoring (Server-side)
    - Calculate time between receiving a TNS packet and sending a response.
  - Server-side Listener Tracing (least-preferred)
- How to fix
  - Find cause of load
    - Buy new hardware
    - Tune queries

Issue #2: Slow Queries

- Simple queries take a long time to return…

Querying the Database

- Open a cursor
- Parse the query
- Execute the query
- Fetch the data
- Cancel the cursor
- Close the cursor

Code (Upon Request)

- oquery.c

Open a Cursor

- TTC Open (OOPEN)
  - Opens a statement
- Protocol
  - Client requests OOPEN
  - Server replies with cursor #

Parse/Execute the Query

- TTC (OALL7/8)
  - Supports several options
    - Parse
    - Bind
    - Execute
    - Fetch
    - Cancel
    - Commit
    - Exact Fetch
    - Send Vector
    - No PL/SQL
- Protocol
  - Client sends OALL packet with cursor #, SQL statement, and flags (Parse, No PL/SQL, Execute), then requesting column/data type info.
  - Server replies with Success/Fail

Fetch Data

- OALL7/8 Packet
  - Options
    - No PL/SQL
    - Fetch
- Protocol
  - Client sends OALL packet with No PL/SQL and Fetch options
  - Server replies with data.

Why can it be slow?

- Network Connectivity (Firewall/VPN)
- Overloaded Server
- Underlying Network Protocol Settings
- Fetching method used by an application

Underlying Network Protocol

- Reason
  - Oracle NT/NS relies on fast UNP
- How to detect
  - Wire-level Monitoring (UNP Fragmentation)
  - Inside Oracle (SDU/UNP Issues)
    - SQL*Net message to client/more data to client
- How to fix
  - Tune OS-level protocols using best practices
  - Make sure Oracle Net matches UNP (SDU==MTU) in listener.ora and tnsnames.ora

Fetch Method

- Reason
  - One-at-a-time fetching means one-roundtrip-per-fetch.
- How to detect
  - Wire-level Monitoring (lots of fetch packets)
  - Inside Oracle
    - Look for a lack of SQL*Net more data to client
- How to fix
  - Use Array-based Fetching
    - OCI (prefetching)
  - Tune Explicit Fetch Sizes

Wire-level Monitoring Tools for Oracle

- Oracle Itself
- Wireshark
- WireCache SQL Query Analyzer
- SCAPE for Oracle (SCAPE4O)

Oracle Network Monitoring

- Pros
  - The simplest method
- Cons
  - Not manageable for large-scale systems

Wireshark

- Pros
  - Free
  - Dissects TNS packets
  - Records Conversations
  - Can be used with stored packet captures
- Cons
  - Does not decode TTI/TTC packets (yet)

WireCache SQL Query Analyzer

- Pros
  - Dissects TNS packets
  - Dissects TTI/TTC Packets
  - Records Conversations
  - Can be used with stored packet captures
- Cons
  - Commercial Product

WireCache SQL Query Analyzer

```
SQL Summary
=-=-=-=-=-=-=-=-=-=-=-=-=-=-=-=-=-=-=-=-=-=-=-=-=-=-=-=-=-=-=-=-=-=-=-=-
             TIME                COUNT
%TIME     FIRST FOLLOW     FIRST FOLLOW STMT
----- --------- --------- ------ ------ ----------------------------------
100.0  18826.34 123.25    1460527 175967 <= TOTAL
 84.8  15983.94 80.36     1353843 154003 -- SELECT
 11.5   2171.89 0.00       14800 0      -- UPDATE
 2.7     459.78 42.89      26977 21964  -- INSERT
 0.6     104.81 0.00       33130 0      -- COMMIT/ROLLBACK
 0.5      86.71 0.00        2823 0      -- DELETE
 0.1      18.77 0.00       28108 0      -- PL-SQL
 0.0       0.43 0.00         846 0      -- ALTER/CREATE
----- --------- --------- ------ ------ ----------------------------------
 28.4   5385.38 0.00        3386 0      SELECT CC_FIRST_CASE_DT FROM PS_RF_INST
 10.4   1972.57 0.00         626 0      UPDATE PS_RF_INST_PROD SET CC_FIRST_CAS
  8.2   1544.53 0.00        1104 0      SELECT INST_PROD_ID FROM PS_RF_INST_PRO
  7.5   1427.12 0.00       48816 0      SELECT DISTINCT SC.CC_CAUSE_ID, C.DESCR
  4.0    749.77 0.00         474 0      SELECT INST.PRODUCT_ID,INST.CC_RECALL_F
  3.6    673.29 0.00         474 0      SELECT INST.CC_POP_FLG, PROD.CC_POP_GRA
  3.5    667.75 0.00        1498 0      SELECT PRODUCT_ID FROM PS_RF_INST_PROD
  2.7    516.45 0.00          84 0      SELECT DISTINCT BO_ID,BO_TYPE_ID,BO_NAM
  2.2    416.79 0.00         154 0      SELECT TO_CHAR(BO_ID) FROM PS_BO WHERE
  1.9    368.44 0.00          10 0      SELECT CM_ID, COUNTRY_CODE, PHONE, EXTE
```

SCAPE for Oracle (SCAPE4O)

- SQL Capture and Analysis by Passive Evaluation
- Pros
  - Free
  - Dissects TNS packets
  - Dissects TTI/TTC Packets
  - Records Conversations
  - Can be used with stored packet captures
- Cons
  - Alpha-quality

SCAPE for Oracle (SCAPE4O)

```
SCAPE4O 0.7.2 – CLIENT (jhh-laptop)
================================================================================
```

Items Learned in this Session

- How to detect network-related issues.
- How to diagnose and solve network-related issues.
- Gained a better understanding of the Oracle Network Protocol

Questions?

Thank You

- Fill out evaluation
  - Jonah H. Harris
  - Listening In: Passive Capture and Analysis of Oracle Network Traffic
  - Session 381
- Further Information
  - jonah.harris@gmail.com
  - http://www.oracle-internals.com/