

Best Practices: Getting Started With Informix Connection Manager

Thomas Beebe

Advanced DataTools Corporation

tom@advanceddatatools.com

Advanced DataTools



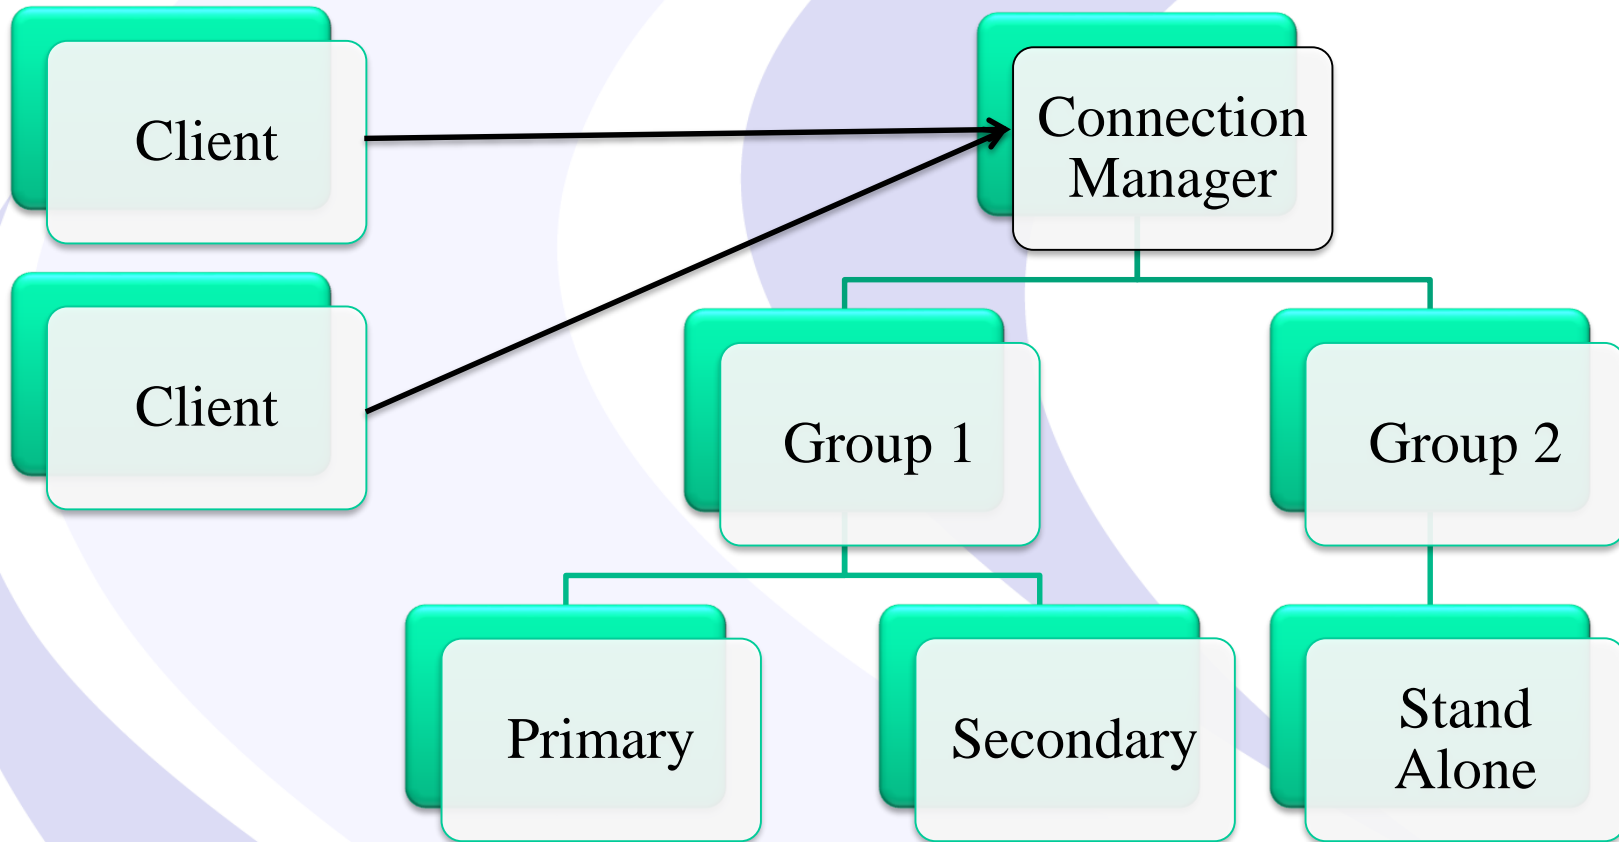
About This Talk

- Created to give an entry point for getting going with connection manager
- Supplement the rather sparse documentation on the topic
- Document some of the confusion I ran into configuring this for clients
- Assistance provided from others actively using connection manager in production environments

What Is Connection Manager

- Introduced in 11.5 (revamped in 11.7)
- Bundled with the engine, also included with the csdk
- Standalone program that runs and passes connections to the correct server or group
- Can be used as a central broker for connections
- Can be run solo or with a group of CM servers working together
- Intended for HA, but very useful in any replicated environment
- OAT module to monitor Connection Manager

Topology



Setting Up Access

- Can be as simple as one server and one CM
 - Client -> Connection Manager -> Instance
- It can be configured as a replication set
 - Client -> Connection Manager -> Primary + HDR + RSS
- It can be configured with multiple groups each with their own rules
- Clients connect to the CM on the listener that will tell the manager where to redirect them and what rules to use
- If the connection manager server is not trusted to the Informix server, use an encrypted password to authenticate

Relevant Files

- **SQLHOSTS**
 - Specifies the groups, also controls the CM listeners
- **\$INFORMIXDIR/bin/oncmsm**
 - Connection Manager Binary
- **\$INFORMIXDIR/etc/cmsm.cfg.sample**
 - Sample config file, there are several versions
- **\$CMALARMPROGRAM**
 - Program, defined in configs that triggers on a failover

ONCONFIG Parameters

- DRAUTO – Controls failover processing, 3 means it will rely on the connection manager to initiate failovers.
- HA_FOC_ORDER – Default order (HDR,RSS,SDS) to fail over servers.

oncmsm

- **Startup:**
 - `$INFORMIXDIR/bin/oncmsm -c $INFORMIXDIR/etc/cmsm.cfg`
 - Optional environmental variable `$CMCONFIG`
- **Commands:**
 - `oncmsm -k -c <config file>`
 - Shutdown
 - `oncmsm -r -c <config file>`
 - Reload with updated config file

Basic Example - SQLHOSTS

#Instance Config port 9088

instance1_tcp	onsoctcp	host_ip	sqlexec
---------------	----------	---------	---------

#Connection Manager Listener

report_group	onsoctcp	host_ip	9090
--------------	----------	---------	------

Basic Example – cmsm.cfg

```
NAME samplecsm
```

```
LOGFILE
```

```
    ${INFORMIXDIR}/tmp/cmsm.log
```

```
CLUSTER samplecluster {
```

```
    INFORMIXSERVER    instance1_tcp
```

```
    SLA    report_group
```

```
    DBSERVERS=primary
```

```
}
```

Base CM Config Sample

NAME cm_1

LOGFILE

 \${INFORMIXDIR}/tmp/cmsm.log

LOG 1

CM_TIMEOUT 300

<Connection Info>

csmm.cfg Parameters

- **NAME** – Must be unique across the cluster
- **LOG** – log level, 1 is on
- **LOGFILE** – Path to CM log
- **CM_TIMEOUT** – Number of seconds to wait for a response before promoting the next highest ranked connection manager. (60 default)
- **EVENT_TIMEOUT** – Number of seconds to wait before failover occurs of Informix servers. If a secondary triggers 'primary offline' it will also trigger failover before the timeout (default 60)
- **SECONDARY_EVENT_TIMEOUT** – Seconds to wait before disconnecting from a secondary (Default 60)
- **SQLHOSTS** – If it should use a local, remote or both SQLhosts files to find instances. (Default local + remote)
- **LOCAL_IP** – Optional, can be used to tie CM to a specific IP address to listen for database status changes
- **MACRO** – Used to create variables to be used in other parts of the script

Connection Types

- CLUSTER – Group or selection of servers to connect to that support HDR failover
- GRID – ER Grid to connect to
- REPL_SET – ER replicate set to connect to
- SERVERSET – Unrelated servers that do not use failover

CLUSTER Example

```
CLUSTER cluster_1
{
    INFORMIXSERVER repl1_tcp
    SLA report_1 DBSERVERS=(PRI,HDR) \
        POLICY=WORKLOAD
    FOC ORDER=ENABLED \
        PRIORITY=1
    CMALARMPROGRAM $INFORMIXDIR/etc/CMALARMPROGRAM.sh
}
```

Cluster Example

CLUSTER cluster_1 – Unique name of cluster, needs to be identical on other connection managers

```
{  
  INFORMIXSERVER repl1_tcp -- The sqlhost entry (group or server) the  
    connection manager will listen on  
  SLA report_1 \ -- the SLA is for the report_1 group in sqlhosts  
    DBSERVERS=(PRI,HDR) \ -- This is the order it will  
    maintain  
    POLICY=WORKLOAD – This is the type of SLA policy it users  
  FOC ORDER=ENABLED \ -- Says to use the failover order above  
    PRIORITY=1 – Says this connection manager is the first  
    one to handle failover for this SLA  
  CMALARMPROGRAM $INFORMIXDIR/etc/CMALARMPROGRAM.sh –  
    If failover fails after 8 attempts it calls this program.  
}
```

INFORMIXSERVER

- Works with all 4 types of connection
- Specify the group of servers or standalone server this SLA should service
- This is what the connection manager will connect to when it comes up to establish the replication status

SLA

- Service Level Agreement
- This is the directive of how a connection manager should treat a particular group of servers
- Any linked connection managers should have similar settings and the same name for a group
- Each SLA will have its own listener port

SLA - DBSERVERS

- List of servers to connect to, and the order to connect to them in
- Can use server names, group names, server aliases, server types (HDR, SDS, ANY)


DBSERVERS - Cluster Keywords

- PRI, PRIMARY
- HDR – Secondary
- SDS – Shared disk secondary
- RSS
- ANY


SLA - MODE

- Redirect – (Default) this will redirect the client directly to the server, only works with versions later then CSDK 3.0 and JDBC 3.5.1
- PROXY – Will pass all data through the connection manager directly, allows for older clients to be supported. Also use this if the client cannot directly access the Informix server.

Topology - Redirect



Topology - PROXY



SLA - USEALIASES

- On – Default, this will add any entries in DBSERVERALIASES into the mix.
- OFF – Only will use DBSERVERNAME none of the aliases

SLA - WORKERS

- Numerical Value – Default is 4.
- Used to define how many worker threads the CM is allocated

SLA - POLICY

- **WORKLOAD** – (Default) – Assigns the work to the least busy server at the time
- **ROUNDROBIN** – Rotates between all of the available servers
- **FAILURE** – Requests pointed to the server with the fewest apply failures. (Replset and GRID only)
- **LATENCY** – Redirects to the server with the lowest transaction latency (Replset and GRID only)
- **SECAPPLYBACKLOG**:<num of pages> -- Stops sending requests to the secondary after it exceeds the number of pages in a backlog. CLUSTER only. (version 12.10xc2 or 11.70xc8 required)

FOC

- Is used as a stand alone element inside a connection type to define how failover occurs
- Also used to specify the priority between servers

FOC - ORDER

- If not defined, the primary server's HA_FOC_ORDER parameter is used
- Default if neither are set is SDS, HDR, RSS
- If enabled it will use the order defined by the DBSERVERS in the SLA
- **ENABLED** – Means the connection manager will allow failover

FOC - PRIORITY

- Defines the priority between connection managers.
- Must be a positive number, the lower the number the higher the priority
- Required for CLUSTER types

FOC - TIMEOUT

- Additional time before a failover occurs
- Adds to the value of EVENT_TIMEOUT
- Defaults to 0

SQLHOSTS

- Connection manager will use SQLHOSTS like any other Informix tool
- By default will read the local SQLHOSTS and if a server is not found it will probe the remote sqlhosts for other relevant hosts
- Will use the INFORMIXSERVER directive in the SLA section to determine the primary server or group to connect to
- The SLA name will be the connection name for the connection manager listener for that SLA.
- Best practices is to always use groups when doing replication rather than individual servers

Sample Config

```
NAME connection_manager_1
```

Name Of This Connection Manager

```
LOG 1
```

```
LOGFILE $INFORMIXDIR/tmp/my_cm1_log.log
```

```
EVENT_TIMEOUT 20
```

```
CLUSTER primary_cluster
```

Cluster Name, Shared Among Connection Managers

```
{
```

```
INFORMIXSERVER cluster1
```

Informix Group/Server To connect to

```
SLA report_1 DBSERVERS=(PRI,HDR) \
```

```
    POLICY=WORKLOAD
```

Connection Manager Listener

```
FOC ORDER=ENABLED \
```

```
    PRIORITY=1
```

```
}
```

SQLHOSTS - Example

#HDR Pair of Servers

cluster_1	group	-	-	c=1,e=repl2_tcp
repl1_tcp	onsoctcp	server1	9088	g=cluster_1
repl2_tcp	onsoctcp	server2	9088	g=cluster_1

#Group of connection managers that service the report
SLA

report	group	-	-	c=1,e=report_2
report_1	onsoctcp	server1	10088	g=report
report_2	onsoctcp	server2	10088	g=report

Application Set Up

SQLHOSTS:

```
report      group      -      -  
      c=1,e=report_2
```

```
report_1    onsoctcp    server1  10088  
      g=report
```

```
report_2    onsoctcp    server2  10088  
      g=report
```

- Set up your application to connect to report_1

Oncmsm Log - Startup

22:45:41 listener report initializing
22:45:41 listener report_rr initializing
22:45:41 listener current_rss initializing
22:45:41 listener proxy_rss initializing
22:45:41 Listener report_rr DBSERVERS=(HDR,RSS) POLICY=ROUNDROBIN is active with 4 worker threads
22:45:41 Listener current_rss DBSERVERS=RSS POLICY=SECAPPLYBACKLOG:5500+WORKLOAD is active with 4 worker threads
22:45:41 Listener report DBSERVERS=(HDR,RSS) POLICY=WORKLOAD is active with 4 worker threads
22:45:41 Listener primary_cm DBSERVERS=primary is active with 4 worker threads
22:45:41 Listener proxy_rss DBSERVERS=RSS POLICY=WORKLOAD MODE=PROXY is active with 4 worker threads
22:45:42 Connection Manager successfully connected to maytcp
22:45:42 The server type of cluster aos_cluster server furytcp is Primary.
22:45:48 The server type of cluster aos_cluster server fury is Primary.
22:46:04 Connection Manager started successfully

Oncmsm Log – Startup Cont

22:46:04 Connection Manager successfully connected to coulson tcp
22:46:04 Cluster aos_cluster Arbitrator FOC ORDER=ENABLED PRIORITY=1
22:46:04 Connection Manager successfully connected to fury nosql
22:45:42 The server type of cluster aos_cluster server fury tcp is Primary.
22:45:48 The server type of cluster aos_cluster server fury is Primary.
22:46:04 Connection Manager started successfully
22:46:04 Connection Manager successfully connected to coulson tcp
22:46:04 Cluster aos_cluster Arbitrator FOC ORDER=ENABLED PRIORITY=1
22:46:04 Connection Manager successfully connected to fury nosql
22:46:04 Connection Manager successfully connected to fury tcp
22:46:04 Connection Manager successfully connected to fury rest
22:46:35 CM cm_1 arbitrator for aos_cluster is active
22:46:35 Cluster aos_cluster Arbitrator FOC ORDER=SDS,HDR,RSS PRIORITY=1
TIMEOUT=0

Oncmsm log - Connections

23:04:23 SLA report_rr redirect SQLI client from
10.10.20.60 to maytcp may.10088

23:30:25 SLA report_rr redirect SQLI client from
10.10.20.71 to furynosql fury.10098

23:30:25 SLA primary_cm redirect SQLI client from
10.10.20.63 to furynosql fury.10098

onstat -g cmsm

```
informix@fury:~$ onstat -g cmsm
```

```
IBM Informix Dynamic Server Version 12.10.UC4DE -- On-Line (Prim) -- Up 00:23:45 -- 154032 Kbytes  
Unified Connection Manager: cm_1                               Hostname: fitz
```

```
CLUSTER      aos_cluster      LOCAL  
Informix Servers: shield_group  
SLA           Connections  Service/Protocol  Rule  
primary_cm    1      20100/onsoctcp    DBSERVERS=primary  
report        9      20101/onsoctcp    DBSERVERS=(HDR,RSS) POLICY=WORKLOAD  
report_rr     25     20102/onsoctcp    DBSERVERS=(HDR,RSS) POLICY=ROUNDROBIN  
current_rss   0      20103/onsoctcp    DBSERVERS=RSS POLICY=SECAPPLYBACKLOG:5500+WORKLOAD  
proxy_hdr     5      20104/onsoctcp    DBSERVERS=HDR POLICY=WORKLOAD MODE=PROXY  
  
Failover Arbitrator: Active Arbitrator, Primary is up  
ORDER=SDS,HDR,RSS PRIORITY=1 TIMEOUT=0
```

```
informix@fury:~$ █
```

Best Practices

- Use groups rather than individual servers
- Run more than one connection manager
- Make sure the connection managers are on different servers from the instances.
- If running PROXY mode make sure to have the resources on the CM to handle the data.
- Make sure to set up applications to use the group of connection managers, or at least failover.
- Make sure applications reconnect with at least a short delay.

Things To Be Cautious Of

- Split Brain
- Listeners missing info
- Alias issues
- Missing Trusted status
- Make sure DBSERVERNAME is the TCP Port




Split Brain

- Two Primary Servers on the same network
- Need to restore one of the servers and re-establish HDR
- Reduce the chance of this situation by having a reliable network connection
- Can use the Connection Manager Alarm Program to shutdown the Primary server if the Secondary cannot be reached AND can't get to the network – See article on IBM developerWorks:

https://www.ibm.com/developerworks/community/blogs/informix_admins_blog/entry/preventing_hdr_split_brain_scenario

Auto Failover and Network Loss

- What if the PRIMARY server loses the network connection?
- Connection Manager will promote the secondary server, incorrectly assuming that the primary is down
- When network connectivity is restored, there will be TWO primary servers – **Split Brain**



Listeners Missing Info

- Will get a -930 error
- Can be caused due to missing entries from sqlhosts on the connection manager server
- Can also get errors if the port is in use, each SLA needs a unique port

Alias Issues

- By default will probe all of the systems in the cluster for DBSERVERALIAS values
- Can be disabled by USEALIAS
- If it sees tcp connections that appear valid but cannot be reached by client, such as a private replication interface, it will still try to offer them for REDIRECT connections

Missing Trusted Sources

- If one of your servers does not trust the connection manager server it will not be able to connect successfully.
- This can be resolved by using the password management process.

Encrypted Password

- Set up a file defining the servers and what passwords to use
- Use onpassword to encrypt the file
- Stored as
\$INFORMIXDIR/etc/passwd_file
- Uses a key to encrypt/decrypt to make changes in the future

More Resources

- Connection Manager Manual Examples

https://www.ibm.com/support/knowledgecenter/SSGU8G_12.1.0/com.ibm.admin.doc/ids_admin_1437.htm

- Informix Replication Technologies:

<http://www.ibm.com/developerworks/data/library/techarticle/dm-0807geib/>

- Andrew Ford's Blog – Setting up HDR:

<http://www.informix-dba.com/2010/08/informix-hdr-will-save-your-butt.html>

- Connection Manager Alert List

https://www.ibm.com/support/knowledgecenter/SSGU8G_11.70.0/com.ibm.adref.doc/ids_adr_1148.htm

Questions?



Send follow-up questions to
tom@advancedatatools.com

Next Webcast

Informix Best Practices

- Informix Auditing by Mike Walker
 - Thursday, July 27, 2017 at 2:00pm EST
- Stay tuned for upcoming announcements about the next set of webcasts

Please register for each webcast here at:

<http://advanceddatatools.com/Informix/NextWebcast.html>

Upcoming WAIUG Meeting

- Aug 8th 2017 – Tysons Corner VA
- 8:30-1pm
- What HCL Brings to the Informix World – Stuart Litel
- Enterprise Replication – Shawn Moe
- Free
- <http://www.waiug.org> to RSVP



Informix Training in 2017

September 11-14, 2017

Advanced Informix Performance Tuning

Sept 18-21, 2017

Informix for Database Administrators

- All courses can be taken online on the web from your desk or at our training center in Virginia
- We guarantee to *NEVER* cancel a course and will teach a course as long as one student is registered



Informix Support and Training from the Informix Champions

Advanced DataTools is an Advanced Level IBM Informix Data Management Partner, and has been an authorized Informix partner since 1993. We have a long-term relationship with IBM, we have priority access to high-level support staff, technical information, and Beta programs. Our team has been working with Informix since its inception, and includes 8 Senior Informix Database Consultants, 4 IBM Champions, 2 IIUG Director's Award winners, and an IBM Gold Consultant. We have Informix specialists Lester Knutsen and Art Kagel available to support your Informix performance tuning and monitoring requirements!

- ***Informix Remote DBA Support Monitoring***
- ***Informix Performance Tuning***
- ***Informix Training***
- ***Informix Consulting***
- ***Informix Development***

Free Informix Performance Tuning Webcast replays at:

<http://advanceddatatools.com/Informix/Webcasts.html>

Call: (800) 807-6732 x101 or email: info@advanceddatatools.com

web: <http://www.advanceddatatools.com>



Advanced DataTools



Thank You

Thomas Beebe
Advanced DataTools Corporation

tom@advanceddatatools.com

For more information:
<http://www.advanceddatatools.com>

Advanced DataTools