www.realdolmen.com

REALDOLMEN

HADR/TSA Overview

REALDOLMEN

SESSION GOAL

- Overview of HADR functionality and advantages
- Description of Tivoli System Automation (TSA) software and integration with DB2
- Operational tasks and lessons learned

REALDOLMEN

WHY HADR?

- Failover solution available from v8.2
- Design choices:
  - Ultra-fast failover capability
  - Negligible impact on performance
  - Configurable degree of consistency
  - Low administrative costs
  - Avoid service interruption for fixpack installs and changes requiring instance restart
  - Transparent failover and failback for applications
  - Spread over different geographic locations
  - Standard hard- and software
  - Built in clustering software in DB2 9.5
- Cheap

BASIC PRINCIPLE HADR

- Database scope
  - Allows flexible configurations
- Primary database
  - Processes transactions
  - Send logs to standby
- Standby database
  - Applies transactions recieved from primary
  - Available for read-only workload (v9.7)
- In case of primary failure, standby database can take over primary role

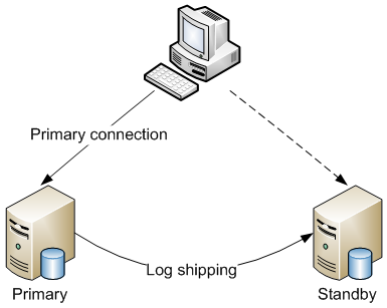

HADR Synchronisation Modes

- Configurable degree of consistency
  - SYNC, NEARSYNC, ASYNC

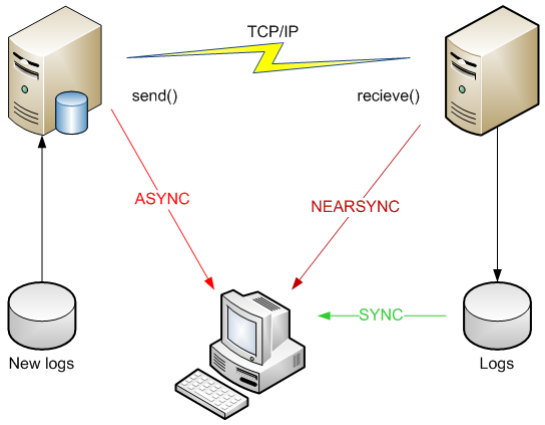

REALDOLMEN

Ready to implement?

- Network layout
  - High speed network
  - Private or public HADR communication?
  - Reserve HADR communication port
- Use identical OS & database version and hardware family
  - For automation with TSA, make sure DB2 > v9.5 FP5 or v9.7 FP1
- Use identical file system layout
- Keep clock synchronised
- Recommended:
  - Same amount of memory & hardware

REALDOLMEN

IMPLEMENTING

- Configure primary database and backup
- Restore on standby server
- Adjust db cfg on standby database
- Start HADR
- Configure Automatic Client Reroute

REALDOLMEN

CONFIGURE PRIMARY DATABASE

- Database configuration

```
db2 update db cfg for <DBNAME> using LOGARCHMETH1 DISK:/share/logarchs
db2 update db cfg for <DBNAME> using INDEXREC RESTART
db2 update db cfg for <DBNAME> using LOGINDEXBUILD ON
db2 update db cfg for <DBNAME> using HADR_TIMEOUT 30
db2 update db cfg for <DBNAME> using HADR_PEER_WINDOW 120
db2 update db cfg for <DBNAME> using HADR_LOCAL_HOST <PRIMARY_HOST>
db2 update db cfg for <DBNAME> using HADR_LOCAL_SVC <PRIMARY_HADR_PORT>
db2 update db cfg for <DBNAME> using HADR_REMOTE_HOST <STANDBY_HOST>
db2 update db cfg for <DBNAME> using HADR_REMOTE_SVC <STANDBY_HADR_PORT>
db2 update db cfg for <DBNAME> using HADR_REMOTE_INST <STANDBY_INST>
db2 update db cfg for <DBNAME> using HADR_SYNCMODE NEARSYNC
```

- Take a backup

REALDOLMEN

CONFIGURE STANDBY SERVER

- Restore backup
  - Keep name
- Database configuration

```
db2 update db cfg for <DBNAME> using HADR_LOCAL_HOST <STANDBY_HOST>
db2 update db cfg for <DBNAME> using HADR_LOCAL_SVC <STANDBY_HADR_PORT>
db2 update db cfg for <DBNAME> using HADR_REMOTE_HOST <PRIMARY_HOST>
db2 update db cfg for <DBNAME> using HADR_REMOTE_SVC <PRIMARY_HADR_PORT>
db2 update db cfg for <DBNAME> using HADR_REMOTE_INST <PRIMARY_INST>
```


REALDOLMEN

START HADR

- Standby

  `db2 start hadr on db <DBNAME> as standby`

  - Replay local log files
- Primary

  `db2 start hadr on db <DBNAME> as primary`

  - Replay primary archived logs
  - Replay primary active logs
  - Replay in-memory log buffer

  -> Peer state

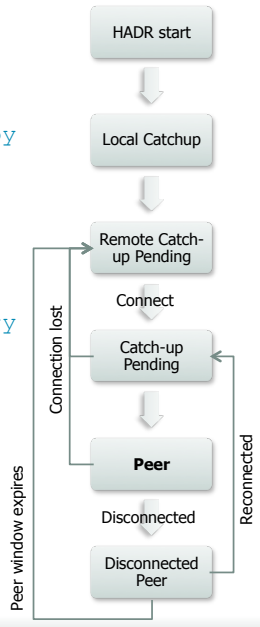

REALDOLMEN

HADR RESULT

- db2pd –hadr –db ORC599

```
Database Partition 0 -- Database ORC599 -- Active -- Up 79 days 06:28:50
HADR Information:
Role    State              SyncMode HeartBeatsMissed   LogGapRunAvg (bytes)
Primary Peer               Nearsync 0                  0

ConnectStatus ConnectTime                            Timeout
Connected     Fri May 21 18:59:35 2010 (1274461175) 120
PeerWindowEnd                          PeerWindow
Fri May 28 16:09:34 2010 (1275055774) 120

LocalHost                                 LocalService
node1                                     50225
RemoteHost                                RemoteService      RemoteInstance
node2                                     50225              db2itst

PrimaryFile  PrimaryPg  PrimaryLSN
S0000413.LOG 786        0x0000000069E9A899
StandByFile  StandByPg  StandByLSN
S0000413.LOG 786        0x0000000069E9A899
```


REALDOLMEN

BASIC OPERATIONS

- Takeover

  `db2 takeover hadr on db <dbname> [by force [peer window only]]`

- Stop database

  `db2 deactivate db <dbname> && db2stop`

  ✋ "`stop hadr`" will turn db to standalone db

- Actually, that's about it…

REALDOLMEN

AUTOMATIC CLIENT REROUTE

- DB2 client reroutes connection to alternative server in case of communication error

```
db2 update alternate server for <dbname> using hostname <host>
  port <svcename>
```

- On connection failure:

```
SQL30108N A connection failed but has been reestablished
```

- HADR is not required

- JDBC URL

```
jdbc:db2://<dbhost>:<port>/
  <db>:clientRerouteAlternateServerName=<althost>;clientRerou
  teAlternatePortNumber=<altport>;
```

- Configure on client and server

THINGS TO KEEP IN MIND

- What gets replicated? Everything that's logged.
- What not?
  - DB CFG
  - History file
  - Not Logged Initially
  - LOB columns that are not logged (>1G)
- Common operations
  - LOAD: COPY YES
  - REORG: consider reorg online or deactivate standby to avoid blocking primary

ARCHIVAL LOGS

- Only primary server archives logs
- Store on shared file system
  - Faster catch-up process
  - Avoid manually copying logs when restore on standby
- Attention when using AUTO_DEL_REC_OBJ

REALDOLMEN

BLOCKNONLOGGED

- Introduced in v9.5 FP4 / v9.7
- Blocks non-logged operations
  - Not logged initially
  - Not logged colums (BLOB and CLOB > 1G)
- Simple measure to protect from human error

REALDOLMEN

MONITORING

- db2pd -hadr
- DB2 Health Monitor
- db2diag.log – check for messages from d2hadrp and db2hadrs

REALDOLMEN

WHERE DOES TSA FIT IN?

- TSA provides automatic monitoring and availability management of “resources”
- TSA helps you when one of the cluster components fails
- But…TSA will also be there when you are
  - Performing maintenance activities
  - Troubleshooting
- Basic knowledge of TSA and how it interacts with DB2 HADR is essential to perform your tasks as a DBA succesfully

REALDOLMEN

GOAL

- Explain components and terminology
- Show TSA domain configuration
- Operations

REALDOLMEN

TSA TERMINOLOGY

- Cluster domain
- Resource
  - Anything for which stop/start/monitor scripts can be written
  - HADR role, DB2 instance, virtual IP, file system, tiebreaker,…
- Resource group
  - All resources in group are online on one and the same physical node
- Quorum
- Tiebreaker
- Equivalency

- Let's see that in a diagram…

REALDOLMEN

CLUSTER DOMAIN OVERVIEW

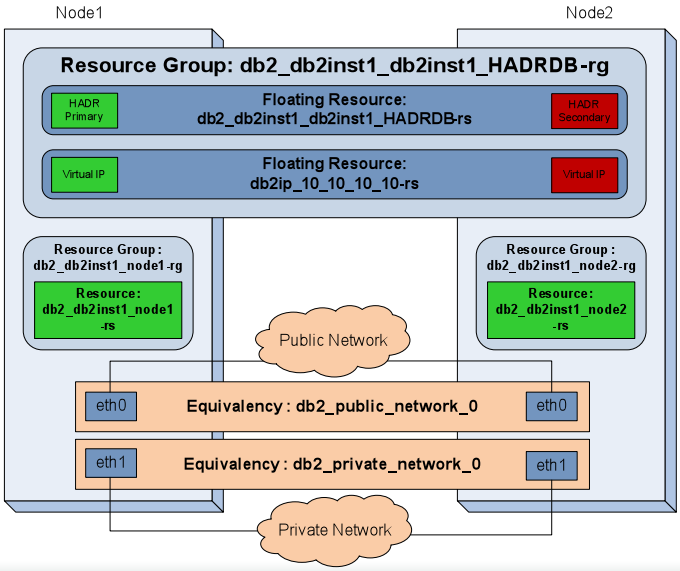

REALDOLMEN

IBM.APPLICATION CLASS – EXAMPLE IS OF A DB2 INSTANCE RESOURCE

```
# lsrsrc -s "Name = 'db2_db2inst1_node1_0-rs'" -Ab IBM.Application

Name            = "db2_db2inst1_node1_0-rs"
StartCommand    = "/usr/sbin/rsct/sapolicies/db2/db2V95_start.ksh
    db2inst1 0"
StopCommand     = "/usr/sbin/rsct/sapolicies/db2/db2V95_stop.ksh
    db2inst1 0"
MonitorCommand  = "/usr/sbin/rsct/sapolicies/db2/db2V95_monitor.ksh
    db2inst1 0"
MonitorCommandPeriod    = 10
MonitorCommandTimeout   = 120
StartCommandTimeout     = 330
StopCommandTimeout      = 140
UserName                = "root"
RunCommandsSync         = 1
ProtectionMode          = 1
ActivePeerDomain        = hadr_domain
NodeNameList            = {"node1"}
OpState                 = 1
```


REALDOLMEN

IBM.APPLICATION CLASS – EXAMPLE IS OF A DB2 HADR RESOURCE

```
# lsrsrc -s "Name = 'db2_db2inst1_db2inst1_HADRDB-rs' " -Ab
    IBM.Application

Name            = "db2hadr_hadrdb-rs"
StartCommand    = "/usr/sbin/rsct/sapolicies/db2/hadrV95_start.ksh
    db2inst1 db2inst1 HADRDB"
StopCommand     = "usr/sbin/rsct/sapolicies/db2/hadrV95_stop.ksh
    db2inst1 db2inst1 HADRDB"
MonitorCommand  = "usr/sbin/rsct/sapolicies/db2/hadrV95_monitor.ksh
    db2inst1 db2inst1 HADRDB"
MonitorCommandPeriod    = 21
MonitorCommandTimeout   = 29
StartCommandTimeout     = 330
StopCommandTimeout      = 140
UserName                = "root"
RunCommandsSync         = 1
ProtectionMode          = 1
ActivePeerDomain        = hadr_domain
NodeNameList            = {"node1","node2"}
OpState                 = 1
```


REALDOLMEN

IBM.SERVICEIP CLASS – VIRTUAL IP ADDRESSES

```
# lsrsrc -Ab IBM.ServiceIP

Name              = "db2ip_10_20_30_42-rs"
IPAddress         = "10.20.30.42"
NetMask           = "255.255.255.0"
ProtectionMode    = 1
ActivePeerDomain  = "hadr_dom"
NodeNameList      = {"node1","node2"}
OpState           = 1
```


REALDOLMEN

MONITOR SCRIPT EXAMPLE

```
#!/bin/ksh -p
#-------------------------------------------------------------------------
# (C) COPYRIGHT International Business Machines Corp. 2001-2009
# All Rights Reserved
#
# US Government Users Restricted Rights - Use, duplication or
# disclosure restricted by GSA ADP Schedule Contract with IBM Corp.
#
# INPUT:  hadrV95_monitor.ksh db2instp db2insts db2hadrdb [verbose]
#
# OUTPUT: 1 if online, 2 if offline, 0 if not known
#
# NOTES: Can only be used in the context of a TSA environment configured
#        with the db2 HA Integrated Configuration Utility (db2haicu)
#
#-------------------------------------------------------------------------
```


REALDOLMEN

QUORUM AND TIEBREAKER

- Quorum: the number of nodes in a cluster that are required to control the resources
- Example: in a 3 node cluster, 2 surviving nodes have the quorum
- Majority with only 2 HADR nodes? Tiebreaker

- Network Tiebreaker
  - only supported type
  - A pingable IP address (default gateway is a good choice)
- Without quorum (+ tiebreaker), no failover will occur !

REALDOLMEN

CONFIGURING FOR HA

- Install TSA + DB2 (at least 9.5.5 or 9.7.1)
  - Who would've guessed…
- Configure HADR (Peer State)
  - We already know how to do that
- Disable db2fm
- Prepare nodes
  - `#preprpnode node1 node2`
- Run DB2 High Availability Instance Configuration Utility (db2haicu)
  - Interactively
  - XML

REALDOLMEN

DB2HAICU XML FILE EXAMPLE

```
<DB2Cluster xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:noNamespaceSchemaLocation="db2ha.xsd" clusterManagerName="TSA"
    version="1.0">
<ClusterDomain domainName="hadr_linux_domain">
    <Quorum quorumDeviceProtocol="network"
    quorumDeviceName="9.26.4.5"/>
    <PhysicalNetwork physicalNetworkName="db2_public_network_0"
    physicalNetworkProtocol="ip">
        <Interface interfaceName="eth0" clusterNodeName="node1">
                <IPAddress baseAddress="9.26.124.30"
                subnetMask="255.255.255.0"
                networkName="db2_public_network_0"/>
        </Interface>
        <Interface interfaceName="eth0" clusterNodeName="node2">
                <IPAddress baseAddress="9.26.124.31"
                subnetMask="255.255.255.0"
                networkName="db2_public_network_0"/>
        </Interface>
    </PhysicalNetwork>
```


REALDOLMEN

DB2HAICU XML FILE EXAMPLE (CONTINUED)

```
        <ClusterNode clusterNodeName="node1"/>
        <ClusterNode clusterNodeName="node2"/>
    </ClusterDomain>
    <FailoverPolicy>
        <HADRFailover></HADRFailover>
    </FailoverPolicy>
    <DB2PartitionSet>
        <DB2Partition dbpartitionnum="0" instanceName="db2inst1"/>
    </DB2PartitionSet>
    <HADRDBSet>
        <HADRDB databaseName="HADRDB" localInstance="db2inst1"
        remoteInstance="db2inst1" localHost="node1"
        remoteHost="node2" />
        <VirtualIPAddress baseAddress="9.26.124.22"
        subnetMask="255.255.245.0"
        networkName="db2_public_network_0"/>
    </HADRDBSet>
</DB2Cluster>
```


REALDOLMEN

DB2HAICU

- Start on standby instance
  - Don't mix primary/standby databases in the same instance

```
$db2haicu -f ha_cfg.xml
```

- Repeat for primary instance
- Done

REALDOLMEN

DB2HAICU STANDBY OUTPUT

db2haicu determined the current DB2 database manager instance is db2inst1. The cluster configuration that follows will apply to this instance.

db2haicu is collecting information on your current setup. This step may take some time as
db2haicu will need to activate all databases for the instance to discover all paths ...
Creating domain hadr_linux_domain in the cluster ...
Creating domain hadr_linux_domain in the cluster was successful.
Configuring quorum device for domain hadr_linux_domain ...
Configuring quorum device for domain hadr_linux_domain was successful.
Adding network interface card eth0 on cluster node node1 to the network
db2_public_network_0 ...
Adding network interface card eth0 on cluster node node1 to the network
db2_public_network_0 was successful.
Adding network interface card eth0 on cluster node node2 to the network
db2_public_network_0 ...
Adding network interface card eth0 on cluster node node2 to the network
db2_public_network_0 was successful.
Adding DB2 database partition 0 to the cluster ...
Adding DB2 database partition 0 to the cluster was successful.
The HADR database HADRDB has been determined to be valid for high availability. However, the
database cannot be added to the cluster from this node because db2haicu detected this node is
the standby for the HADR database HADRDB. Run db2haicu on the primary for the HADR
database HADRDB to configure the database for automated failover.
All cluster configurations have been completed successfully. db2haicu exiting ...

REALDOLMEN

DB2HAICU PRIMARY OUTPUT

db2haicu determined the current DB2 database manager instance is db2inst1. The cluster configuration that follows will apply to this instance.

db2haicu is collecting information on your current setup. This step may take some time as db2haicu will need to activate all databases for the instance to discover all paths ...

Configuring quorum device for domain hadr_linux_domain ...
Configuring quorum device for domain hadr_linux_domain was successful.
The network adapter eth0 on node node1 is already defined in network db2_public_network_0
and cannot be added to another network until it is removed from its current network.
The network adapter eth0 on node node2 is already defined in network db2_public_network_0
and cannot be added to another network until it is removed from its current network.
Adding DB2 database partition 0 to the cluster ...
Adding DB2 database partition 0 to the cluster was successful.
Adding HADR database HADRDB to the domain ...
Adding HADR database HADRDB to the domain was successful.
All cluster configurations have been completed successfully. db2haicu exiting ...

REALDOLMEN

RESULT

```
~# lssam
Online IBM.ResourceGroup:db2_db2inst1_node1_0-rg Nominal=Online
   '- Online IBM.Application:db2_db2inst1_node1_0-rs
      '- Online IBM.Application:db2_db2inst1_node1_0-rs:node1
Online IBM.ResourceGroup:db2_db2inst1_node2_0-rg Nominal=Online
   '- Online IBM.Application:db2_db2inst1_node2_0-rs
      '- Online IBM.Application:db2_db2inst1_node2_0-rs:node2
Online IBM.ResourceGroup:db2_db2inst1_db2inst1_HADRDB-rg Nominal=Online
   |- Online IBM.Application:db2_db2inst1_db2inst1_HADRDB-rs
      |- Online IBM.Application:db2_db2inst1_db2inst1_HADRDB-rs:node1
      '- Offline IBM.Application:db2_db2inst1_db2inst1_HADRDB-rs:node2
   '- Online IBM.ServiceIP:db2ip_9_26_124_22-rs
      |- Online IBM.ServiceIP:db2ip_9_26_124_22-rs:node1
      '- Offline IBM.ServiceIP:db2ip_9_26_124_22-rs:node2
~> db2 get dbm cfg | grep CLUSTER_MGR
Cluster manager                         (CLUSTER_MGR) = TSA
```


REALDOLMEN

OPERATIONS

- Use regular commands whenever possible:
  - Takeover, db2stop, db2start,…
- Takeover example
  - $ db2 takeover hadr on db <DBNAME>

REALDOLMEN

TAKEOVER EXAMPLE

```
Online IBM.ResourceGroup:db2_db2icom_db2icom_CAT054-rg Nominal=Online
      '- Online IBM.Application:db2_db2icom_db2icom_CAT054-rs
            |- Offline IBM.Application:db2_db2icom_db2icom_CAT054-rs:s999jclnx14
            '- Online IBM.Application:db2_db2icom_db2icom_CAT054-rs:s999jclnx17
Online IBM.ResourceGroup:db2_db2icom_db2icom_ORC054-rg Nominal=Online
      '- Online IBM.Application:db2_db2icom_db2icom_ORC054-rs
            |- Offline IBM.Application:db2_db2icom_db2icom_ORC054-rs:s999jclnx14
            '- Online IBM.Application:db2_db2icom_db2icom_ORC054-rs:s999jclnx17
Online IBM.ResourceGroup:db2_db2icom_s999jclnx14.concept.gfdi.be_0-rg Nominal=Online
      '- Online IBM.Application:db2_db2icom_s999jclnx14.concept.gfdi.be_0-rs
            '- Online IBM.Application:db2_db2icom_s999jclnx14.concept.gfdi.be_0-rs:s999jclnx14
Online IBM.ResourceGroup:db2_db2icom_s999jclnx17.concept.gfdi.be_0-rg Nominal=Online
      '- Online IBM.Application:db2_db2icom_s999jclnx17.concept.gfdi.be_0-rs
            '- Online IBM.Application:db2_db2icom_s999jclnx17.concept.gfdi.be_0-rs:s999jclnx17
```


REALDOLMEN

TAKEOVER EXAMPLE

```
Online IBM.ResourceGroup:db2_db2icom_db2icom_CAT054-rg Nominal=Online
      '- Online IBM.Application:db2_db2icom_db2icom_CAT054-rs
            |- Offline IBM.Application:db2_db2icom_db2icom_CAT054-rs:s999jclnx14
            '- Online IBM.Application:db2_db2icom_db2icom_CAT054-rs:s999jclnx17
Pending offline IBM.ResourceGroup:db2_db2icom_db2icom_ORC054-rg Nominal=Online
      '- Pending offline IBM.Application:db2_db2icom_db2icom_ORC054-rs
            |- Offline IBM.Application:db2_db2icom_db2icom_ORC054-rs:s999jclnx14
            '- Pending offline IBM.Application:db2_db2icom_db2icom_ORC054-rs:s999jclnx17
Online IBM.ResourceGroup:db2_db2icom_s999jclnx14.concept.gfdi.be_0-rg Nominal=Online
      '- Online IBM.Application:db2_db2icom_s999jclnx14.concept.gfdi.be_0-rs
            '- Online IBM.Application:db2_db2icom_s999jclnx14.concept.gfdi.be_0-rs:s999jclnx14
Online IBM.ResourceGroup:db2_db2icom_s999jclnx17.concept.gfdi.be_0-rg Nominal=Online
      '- Online IBM.Application:db2_db2icom_s999jclnx17.concept.gfdi.be_0-rs
            '- Online IBM.Application:db2_db2icom_s999jclnx17.concept.gfdi.be_0-rs:s999jclnx17
```


REALDOLMEN

TAKEOVER EXAMPLE

```
Online IBM.ResourceGroup:db2_db2icom_db2icom_CAT054-rg Nominal=Online
      '- Online IBM.Application:db2_db2icom_db2icom_CAT054-rs
            |- Offline IBM.Application:db2_db2icom_db2icom_CAT054-rs:s999jclnx14
            '- Online IBM.Application:db2_db2icom_db2icom_CAT054-rs:s999jclnx17
Pending online IBM.ResourceGroup:db2_db2icom_db2icom_ORC054-rg Nominal=Online
      '- Pending online IBM.Application:db2_db2icom_db2icom_ORC054-rs
            |- Pending online IBM.Application:db2_db2icom_db2icom_ORC054-rs:s999jclnx14
            '- Offline IBM.Application:db2_db2icom_db2icom_ORC054-rs:s999jclnx17
Online IBM.ResourceGroup:db2_db2icom_s999jclnx14.concept.gfdi.be_0-rg Nominal=Online
      '- Online IBM.Application:db2_db2icom_s999jclnx14.concept.gfdi.be_0-rs
            '- Online IBM.Application:db2_db2icom_s999jclnx14.concept.gfdi.be_0-rs:s999jclnx14
Online IBM.ResourceGroup:db2_db2icom_s999jclnx17.concept.gfdi.be_0-rg Nominal=Online
      '- Online IBM.Application:db2_db2icom_s999jclnx17.concept.gfdi.be_0-rs
            '- Online IBM.Application:db2_db2icom_s999jclnx17.concept.gfdi.be_0-rs:s999jclnx17
```


REALDOLMEN

TAKEOVER EXAMPLE

```
Online IBM.ResourceGroup:db2_db2icom_db2icom_CAT054-rg Nominal=Online
      '- Online IBM.Application:db2_db2icom_db2icom_CAT054-rs
            |- Offline IBM.Application:db2_db2icom_db2icom_CAT054-rs:s999jclnx14
            '- Online IBM.Application:db2_db2icom_db2icom_CAT054-rs:s999jclnx17
Online IBM.ResourceGroup:db2_db2icom_db2icom_ORC054-rg Nominal=Online
      '- Online IBM.Application:db2_db2icom_db2icom_ORC054-rs
            |- Online IBM.Application:db2_db2icom_db2icom_ORC054-rs:s999jclnx14
            '- Offline IBM.Application:db2_db2icom_db2icom_ORC054-rs:s999jclnx17
Online IBM.ResourceGroup:db2_db2icom_s999jclnx14.concept.gfdi.be_0-rg Nominal=Online
      '- Online IBM.Application:db2_db2icom_s999jclnx14.concept.gfdi.be_0-rs
            '- Online IBM.Application:db2_db2icom_s999jclnx14.concept.gfdi.be_0-rs:s999jclnx14
Online IBM.ResourceGroup:db2_db2icom_s999jclnx17.concept.gfdi.be_0-rg Nominal=Online
      '- Online IBM.Application:db2_db2icom_s999jclnx17.concept.gfdi.be_0-rs
            '- Online IBM.Application:db2_db2icom_s999jclnx17.concept.gfdi.be_0-rs:s999jclnx17
```


REALDOLMEN

OPERATIONS

- TSA commands:
  - stoprpnode, startrpnode, stoprpdomain, startrpdomain
  - lssam, lsrel,
  - chrg, lsrg, rgreq, rmrg
  - lsrsrc, chrsrc, rmrsrc
- Most useful for troubleshooting:
  - `rgreq -o unlock <resource_group_name>`
  - `rgreq -o move <resource_group_name>`
  - `resetrsrc -s "Name like 'db2_db2inst1_db2inst1_HADRDB-rs'" IBM.Application`

REALDOLMEN

REGAIN CONTROL

- Disable TSA
  - `$ db2haicu -disable`
- Remove DB2 from TSA config
  - `$ db2haicu -delete`
- Change TSAMP to manual/automatic mode
  - `# samctrl -m T/F`
- Remove domain
  - `# rmrpdomain -f <domainname>`

REALDOLMEN

MONITORING

- db2pd –ha
- Lssam

- Syslog(/var/log/messages)
  - Output from automation scripts
  - Error messages from TSA
- /var/ct/<DOMAIN>/log/mc/IBM.<DAEMON>RM
  - Daemon log file directory, with <DAEMON> = Recovery, GblRes, …
- Getsadata script

REALDOLMEN

SYSLOG EXAMPLE

```
/usr/sbin/rsct/sapolicies/db2/hadrV95_monitor.ksh[31113]: Entering    : db2icom db2icom CAT054
/usr/sbin/rsct/sapolicies/db2/hadrV95_monitor.ksh[31307]: Returning 1 : db2icom db2icom CAT054
db2V95_start.ksh[677]: Entered /usr/sbin/rsct/sapolicies/db2/db2V95_start.ksh, db2icom, 0
/usr/sbin/rsct/sapolicies/db2/hadrV95_monitor.ksh[1375] : Entering    : db2icom db2icom CAT054
/usr/sbin/rsct/sapolicies/db2/hadrV95_monitor.ksh[1654] : Returning 2 : db2icom db2icom CAT054
/usr/sbin/rsct/sapolicies/db2/hadrV95_start.ksh[1665]   : Entering    : db2icom db2icom CAT054
db2V95_start.ksh[2927]: Returning 0 from /usr/sbin/rsct/sapolicies/db2/db2V95_start.ksh ( db2icom, 0)
/usr/sbin/rsct/sapolicies/db2/hadrV95_monitor.ksh[4695] : Entering    : db2icom db2icom CAT054
/usr/sbin/rsct/sapolicies/db2/hadrV95_monitor.ksh[5952] : Returning 1 : db2icom db2icom CAT054
/usr/sbin/rsct/sapolicies/db2/hadrV95_start.ksh[6198]: Returning 0    : db2icom db2icom CAT054
/usr/sbin/rsct/sapolicies/db2/hadrV95_monitor.ksh[6201] : Entering    : db2icom db2icom CAT054
/usr/sbin/rsct/sapolicies/db2/hadrV95_monitor.ksh[6680] : Returning 2 : db2icom db2icom CAT054
/usr/sbin/rsct/sapolicies/db2/hadrV95_start.ksh[6707]   : Entering    : db2icom db2icom CAT054
db2V95_start.ksh[6752]: Entered /usr/sbin/rsct/sapolicies/db2/db2V95_start.ksh, db2icom, 0
/usr/sbin/rsct/sapolicies/db2/hadrV95_monitor.ksh[9969] : Entering    : db2icom db2icom CAT054
/usr/sbin/rsct/sapolicies/db2/hadrV95_monitor.ksh[10650]: Returning 1 : db2icom db2icom CAT054
db2V95_start.ksh[10654]: Returning 0 from /usr/sbin/rsct/sapolicies/db2/db2V95_start.ksh ( db2icom, 0)
/usr/sbin/rsct/sapolicies/db2/hadrV95_monitor.ksh[10984]: Returning 1 : db2icom db2icom ORC054
/usr/sbin/rsct/sapolicies/db2/hadrV95_start.ksh[12895]  : Returning 0 : db2icom db2icom CAT054
/usr/sbin/rsct/sapolicies/db2/hadrV95_monitor.ksh[12925]: Entering    : db2icom db2icom CAT054
/usr/sbin/rsct/sapolicies/db2/hadrV95_monitor.ksh[13610]: Returning 1 : db2icom db2icom CAT054
```


REALDOLMEN

TIPS &TRICKS

- Use TSA to improve availability of standalone databases
  - Replace db2fm
  - Use HADR failover without specifying HADR database
  - Instance start-script will also activate databases
- Keep the # of databases / instance low
- Be consistent in configuration parameters (hostname, case, …)
- Sampolicy

Case Study

- Healthcare company
- ~3K users; ~2M customers

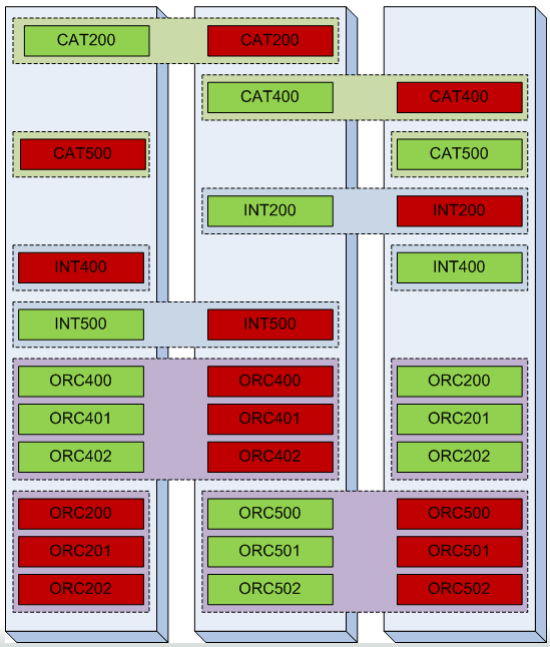

REALDOLMEN

DESIGN CHOICES

- Bonded NIC's -> HADR communication via public network
- Nearsync
  - Risk of both servers failing simultaneously is very small
- Archival logs on NFS share
- Fixed INSTANCE_MEMORY
- Automatic storage

REALDOLMEN

LIBRARY & REFERENCES

- DB2 Information Center
- TSAMP documentation
  - http://publib.boulder.ibm.com/tividd/td/IBMTivoliSystemAutomationforMultiplatforms2.2.html
- Automated Cluster Controlled HADR Configuration Setup using the IBM DB2 High Availability Instance Utility
  - ftp://ftp.software.ibm.com/software/data/pubs/papers/HADR_db2haicu.pdf
- Automating DB2 HADR Failover on Linux using Tivoli System Automation for Multiplatforms
  - ftp://ftp.software.ibm.com/software/data/pubs/papers/hadr_tsa.pdf
- Integrating TSAMP with DB2 HADR Configuration
  - http://www-01.ibm.com/support/docview.wss?uid=swg27016774&aid=1
- High Availability and Disaster Recovery Options for DB2 on Linux, UNIX, and Windows
  - www.redbooks.ibm.com/redbooks/pdfs/sg247363.pdf
- DB2 Best Practices – HADR
  - http://www.ibm.com/developerworks/data/bestpractices/

Questions?