Beginner's Guide To HDR

Hal Maner
M Systems International, Inc.

Session A06
April 28, 2008 • 04:40 – 05:40 p.m.

2008 IIUG Informix Conference

www.iiug.org

Presenter – Hal Maner

- 1993: Started using Informix as a 4GL developer.
- 1995: Liked it so much, started wearing the DBA hat too.
- 1996: Started M Systems International, Inc., a service/solution partner of Informix and now IBM, specializing primarily in Informix and Java/Web technologies.
- 1998: Informix Corporation stock, despite all its problems, allowed me to purchase a nice new car!
- 1999: Setup replication for the first time (not with IDS!).

Presentation Overview: High-Availability Data Replication (HDR)

- What is it?
- Why would I want to use it?
- Setup - planning
- Setup – step by step
- Monitoring
- Real World examples

HDR – What Is It?

- HDR stands for “High-Availability Data Replication” in IDS terminology. Some people read it as “High-Availability Disaster Recovery” which is appropriate…
- Replication’s common dictionary meaning is to “copy” or “reproduce”, so HDR is technology that allows you to copy your data to separate media.
- This, of course, sounds very simple and almost like duplicating a DVD…

HDR – What Is It?

- HDR allows all the data in an IDS instance to be copied *on an ongoing basis* to another server across a network connection.
- It is an “engine level” replication technology to provide copies of your Informix data on different servers.
- Does not have restrictions such as requiring primary keys in your tables, etc. Requires that your databases be logged, however, since its underlying technology is based on logical logs.

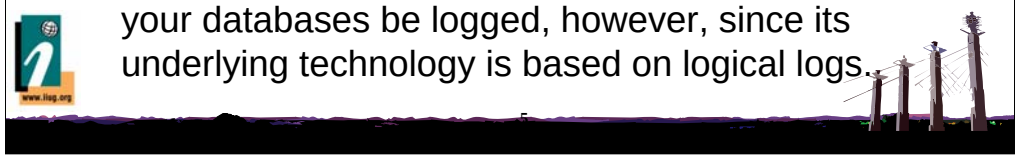

HDR does not care about how beautiful (or ugly) your database design is (however, your databases must use logging)!!! It is an “engine level” replication technology designed for disaster recovery, primarily. Some replication technologies, including IDS’ own Enterprise Replication (ER) and Microsoft SQL Server for example, require primary keys to exist in tables that are replicated (and unlike IDS that lets you choose the proper replication technology for your needs, i.e. HDR and ER, Microsoft SQL Server does not have anything like HDR – or at least it did not as of SQL Server 2000).

HDR – What Is It?

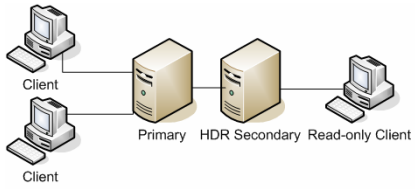

HDR Terminology:

- Primary: This is the server your users normally connect to. It allows updates to the data as usual. This is the source of the data copied to the secondary.
- Secondary: Read-only backup server that receives its updates from the Primary.

HDR – What Is It?

Additional terminology as of IDS 11 – where we used to have just a “secondary” before, now we have various types of secondary servers:

- HDR secondary server: this is the secondary as defined in previous IDS versions, i.e. a hot-standby that receives updates "synchronously".
- Remote Standalone (RS) secondary server ("RSS"): receives updates asyncronously and can be used with low speed or intermittent network connections.

HDR – What Is It?

Additional terminology as of IDS 11:

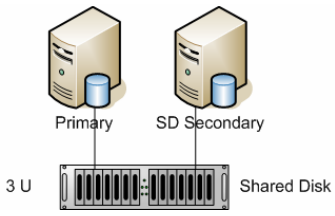

- Shared disk (SD) secondary server: a physical server that shares the same disk space as its primary. This secondary still only has read-only access to the data.

Note: HDR is currently available with the Workgroup and Enterprise Editions of IDS.

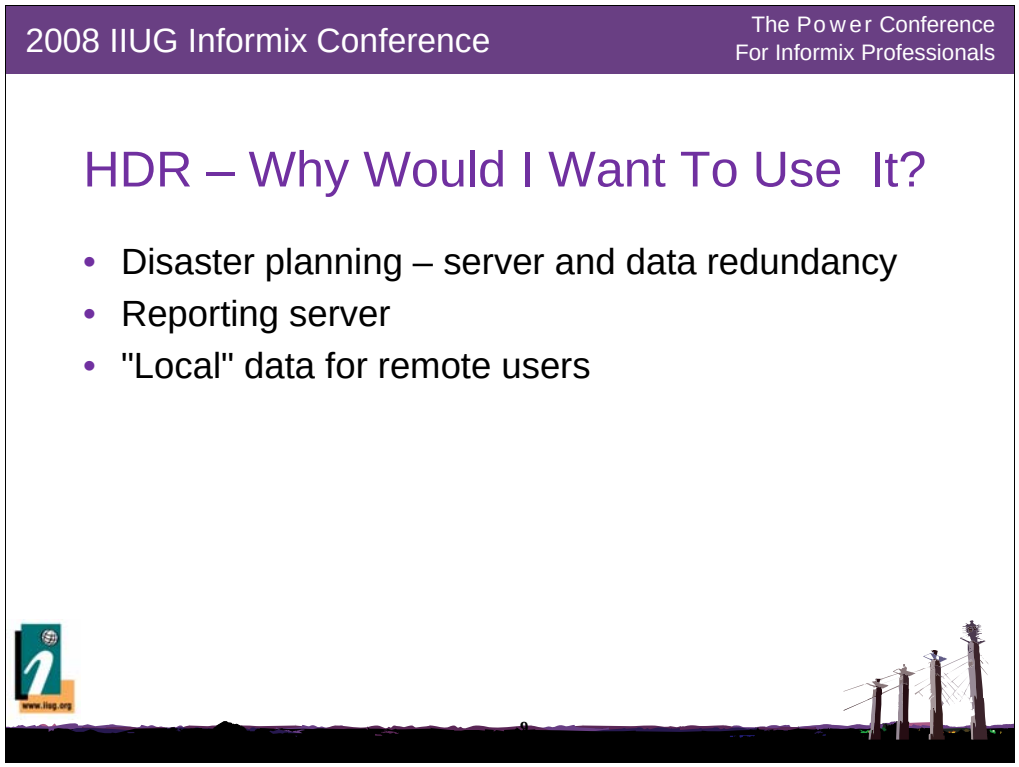

The president of one of our client companies (a company with several hundred employees, in the distribution business) contacted me a few years ago… He told me how a friend of his had a business in Detroit and a fire burnt down their whole facility including their data backups and the company went out of business within weeks…

This president wanted to make sure this would not happen to his company. He wanted an effective disaster recovery plan for one of their most important assets – their data.

Since they were using a supply chain management system running on IDS (7.31 at that time), the answer was pretty easy: use HDR.

HDR is a relatively simple and very effective way of having your whole database engine duplicated on another server preferably in a different geographical location from where your main server is. All this without any programming or other changes in your databases being required.

If you want disaster recovery preparedness for your databases, especially a “hot standby” server, then you should consider HDR.

If you have some read-only/reporting users that you wish you could offload to a separate box for reporting, you should consider HDR. Note: check with your sales rep on licensing requirements for this option.

There are various other side benefits of using HDR to provide “local” data to the users at the sites you are replicating to for example.

HDR – How Does It Work?

- Initial setup with backup&restore (except for SD Secondary)
- Kept in synch using logical log data – how?

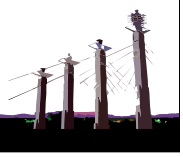

HDR – How Does It Work?

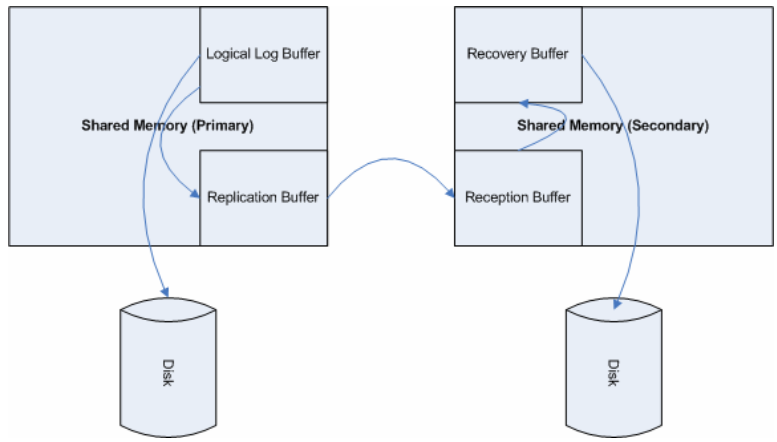

HDR requires a level zero backup&restore for the initial setup.

IDS makes use of the logical logs to replicate data from the primary server:

When the contents of the logical log buffer in the primary server's shared memory gets flushed to disk, IDS also copies this data to a data "replication buffer". The primary server sends these logical log records to the HDR secondary database server.

The HDR secondary has a "reception buffer" in its shared memory that receives the logical log records. The server is in logical recovery mode in order to apply these records continuously.

HDR – How Does It Work?

- Synchronous vs Asynchronous: DRINTERVAL
- DRINTERVAL -1 (synchronous, i.e. primary waits for acknowledgement from secondary before declaring the logical-log buffer flush, i.e. commit, complete)
- DRINTERVAL other than -1 (asynchronous – e.g. 30 seconds)

Logical log records are copied to replication buffer which is then sent to reception buffer and then to the recovery buffer and then finally to disk on the secondary…

…but how often are these sent? When can I expect my newly inserted row to show up on the secondary… you may ask.

This is where you, the DBA, has control as this is done with IDS configuration. The contents of the replication buffer can be sent either synchronously (i.e. immediately) or asynchronously. The ONCONFIG parameter called DRINTERVAL controls this. If DRINTERVAL is set to -1, replication occurs synchronously.

Watch out: with synchronous replication, there is a dependency on the HDR secondary to complete its transaction commits before the rows on the primary can be committed. Primary waits for an acknowledgement from the HDR secondary. The disadvantage of this method is that there is a slight performance hit. The advantage is that you can be certain that your committed rows on the secondary will always match your primary in case of a failure that breaks the replication.

Asynchronous updating: If DRINTERVAL is set to any value other than -1, replication occurs asynchronously:

Primary server copies the logical log buffer contents to the replication buffer

Primary server flushes the logical log buffer to disk

Primary server sends the replication buffer contents to the secondary when the replication buffer becomes full, or, the time interval specified by DRINTERVAL (in seconds) has elapsed. The default value for DRINTERVAL is 30 seconds. This means that, in approximately 30 seconds, you should see the data replicated on the secondary.

This method has a performance advantage over the synchronous method, however the disadvantage here is that it is possible to have lost transactions.

If, after a transaction is committed on the primary, a failure occurs and the commit is never sent to the secondary, there will be a mismatch. In such cases, IDS makes use of a lost-and-found file specified in the DRLOSTFOUND parameter. These transactions are typically never applied by the database server.

Using unbuffered logging as you database logging mode reduces the likelihood of running into the lost transaction scenario because the transactions are transferred sooner due to the flushing to disk that goes with unbuffered logging.

*** Checkpoints are synchronous regardless of the DRINTERVAL – checkpoint on the secondary completes before checkpoint on the primary can complete ***

HDR – How Does It Work?

- HDR Threads
  - drprsend: (primary) sends buffer
  - drsecrecv: (secondary) receives buffer
  - drprping: (primary) checks connectivity
  - drsecping: (secondary) checks connectivity
  - drsecapply: (secondary) copy reception buffer to recovery buffer

drprsend: runs on the primary to send the replication buffer contents to the secondary.

drsecrcv: runs on the HDR secondary to receive the replicated data.

drprping and drsecping: run on the primary and the secondary, respectively, and ensure connectivity between the servers is up.

drsecapply: copies data from the reception buffer to the recovery buffer on the secondary

logrecvr: performs logical recovery on the secondary – updates the actual dbspaces from the logical log records. OFF_RECVRY_THREADS configuration parameter specifies how many logrecvr threads should be used.

There are some different threads used by the RS Secondary servers but we will not get into those in this presentation.

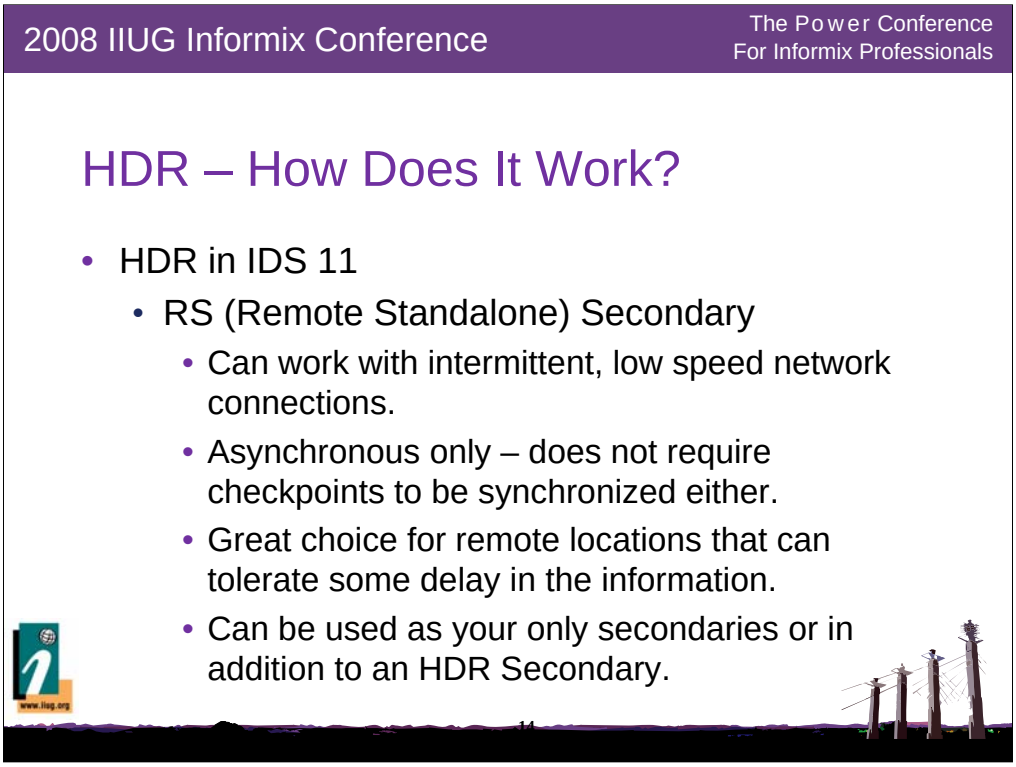

Before IDS 11, HDR was limited to having two servers: one primary and one secondary. You needed a relatively high speed and dependable network connection between the two.

IDS 11 introduced new types of secondaries and made HDR a lot more powerful and flexible by also allowing multiple secondaries to be used:

RS (Remote Standalone) Secondary Server

One of the issues with HDR prior to IDS 11 is that the network connection between the primary and secondary needs to be very reliable. Network connectivity issues can, fairly frequently, cause the replication to register a failure which causes the replication to stop and require re-synching.

To address this issue, RS secondary servers can be used. These secondaries are designed to work even with less dependable (i.e. slower or erratic) network connections with the primary. They do not require checkpoints to be synchronized with the primary either.

Since IDS 11 also allows multiple secondaries to be configured, if you do not have a reliable connection to your remote location, you can, for example, configure an HDR Secondary in the same physical location as your primary and configure one or more RS Secondary servers at remote locations.

RS Secondary servers can be promoted to HDR secondary if needed.

SD (Shared Disk) Secondary server:

This new type of secondary, as its name implies, uses the same disk storage as its primary.

This configuration allows “cluster” type setups where you can provide computer server hardware redundancy. You would also need to provide disk storage redundancy by using RAID and/or SAN technology.

IDS 11 enhancements have now opened the door for virtually limitless possibilities on planning and configuring your disaster recovery alternatives.

HDR Setup: Planning

- Planning for HDR is a must!
- Business requirements (e.g. disaster recovery planning, additional reporting capacity, "local" data) drive the technical configuration.
- Operating system and IDS versions must be identical.
- Chunk layouts of primary and secondary must be identical.

Planning is a very important step in getting ready to use HDR.

At the business level:

Your HDR configuration specifics will depend on the business requirements. You must clearly define your goals (such as disaster recovery preparedness, additional reporting capacity) and plan your infrastructure accordingly.

At the server level:

HDR requires that the same version of IDS and operating system are used on all hardware. It is possible to use different hardware, however the operating system must be the same.

At the database level:

Your databases must use transaction logging. The advantages of unbuffered logging also apply in HDR.

At the network level:

Depending on the business requirements, you may need to make network connectivity improvements/changes. IDS 11 provides you with a lot more flexibility as compared with earlier versions, however you should make sure you have the appropriate network connectivity

HDR Setup: Planning (onconfig)

DRAUTO

0 (OFF) means do not automatically switch the type of the server.

1 (RETAIN_TYPE) means automatically switch secondary to standard if replication fails, back to secondary upon HDR restart.

2 (REVERSE_TYPE) means automatically switch secondary to standard on failure, to primary upon HDR restart.

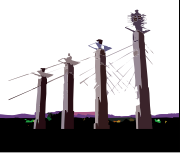

16

Why is zero (OFF) in bold? This is the one I recommend you use unless you have some very specific requirements in your environment. Automatically switching type can cause unwanted consequences because HDR may go down because of a temporary network glitch, for example.

HDR Setup: Planning (onconfig)

DRIDXAUTO

- 0 (off) means no automatic replication of indexes to the secondary in case of corruption.
- 1 (on) means replicate an index to the secondary automatically if a corrupted index is detected. The "on" setting is required for RSS.

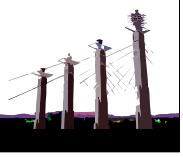

HDR Setup: Planning (onconfig)

DRINTERVAL

- Maximum interval, in seconds, between the flushing of the replication buffer.
- Accepts -1, 0, positive integer values.
- Default is 30 seconds which means if may take up to 30 seconds for the changes on your primary to be replicated to the HDR secondary.
- -1 means sychronous setting.

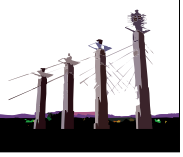

HDR Setup: Planning (onconfig)

DRLOSTFOUND

Pathname to the lost&found file for replication. When asynchronous replication is used and in the case of failure, it is possible that some transactions are committed on the primary but not on the secondary. These transactions are saved in a lost&found file.

HDR Setup: Planning (onconfig)

DRTIMEOUT

- Length of time, in seconds, that the primary waits for a receipt acknowledgement from the HDR secondary.
- After this length of time is elapsed with no acknowledgement, IDS assumes that a failure has occurred.
- Default is 30 seconds.

The actual wait time for the server is documented as four times the value set in DRTIMEOUT, i.e. 120 seconds when the setting is 30.

HDR Setup: Step By Step (Primary-HDR Secondary)

1. Planning stage complete – two identical servers ready and databases are logged.
2. TCP/IP steps:
   - Ensure both servers are in each other's /etc/hosts files with the correct IP address.
   - Update /etc/hosts.equiv on each server with the other's name.
   - Ensure /etc/services file on each server has the proper IDS service.

Planning is a very important step in getting ready to use HDR.

At the business level:

Your HDR configuration specifics will depend on the business requirements. You must clearly define your goals (such as disaster recovery preparedness, additional reporting capacity) and plan your infrastructure accordingly.

At the server level:

HDR requires that the same version of IDS and operating system are used on all hardware. It is possible to use different hardware, however the operating system must be the same.

At the database level:

Your databases must use transaction logging. The advantages of unbuffered logging also apply in HDR.

At the network level:

Depending on the business requirements, you may need to make network connectivity improvements/changes. IDS 11 provides you with a lot more flexibility as compared with earlier versions, however you should make sure you have the appropriate network connectivity

HDR Setup: Step By Step (Primary-HDR Secondary)

3. IDS sqlhosts file: each server should have the other's entry in it. You must use a TCP/IP connection and not shared memory connection for HDR configuration (your engines can still have shared memory connections configured as well).

Example:

```
ids_shm_tsunami  onipcshm       tsunami        ipcshm
ids_tsunami      ontlitcp       tsunami          ids
ids_shm_tsunami2 onipcshm       avalanche     ipcshm
ids_tsunami2     ontlitcp       avalanche       ids2
```

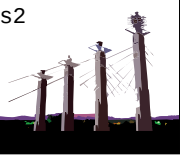

HDR Setup: Step By Step (Primary-HDR Secondary)

4. Onconfig files: the following parameter values must be identical on the primary and the secondary:

   ROOT

   PHYSLOG

   DRAUTO

   DRINTERVAL

   DRLOSTFOUND

   DRTIMEOUT

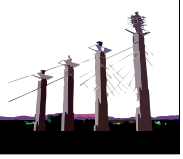

HDR Setup: Step By Step (Primary-HDR Secondary)

5. Take a level zero archive on the (soon to be) primary:

   ontape –s –L 0

6. Set the server to be primary:

   onmode –d primary ids_tsunami2

IDS secondary name

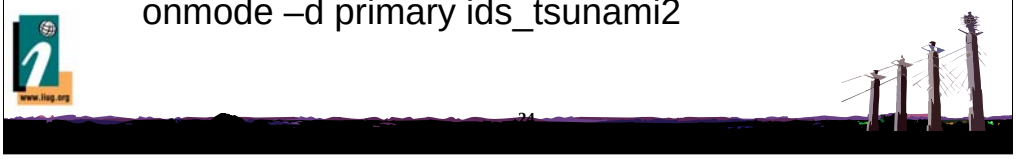

HDR Setup: Step By Step (Primary-HDR Secondary)

7. Restore the level zero backup as a physical restore on the secondary:

   ontape –p

8. Set the server to be secondary:

   onmode –d secondary ids_tsunami

IDS primary name

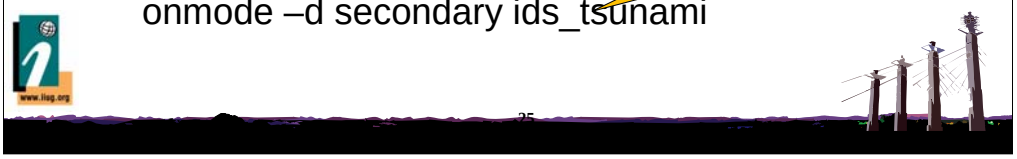

HDR Setup: Step By Step (Primary-HDR Secondary)

9. Verify that you have a successful HDR connection:

   onstat –m

   onstat –g dri

   sysmaster:sysdri table

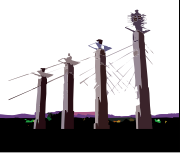

HDR Setup: Step By Step (Primary-HDR Secondary)

Level zero archive

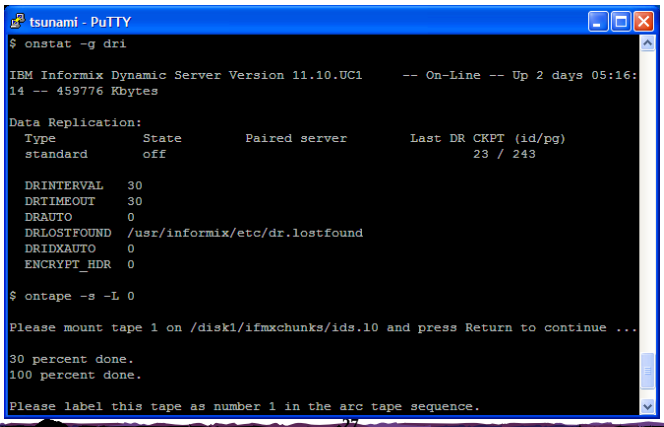

HDR Setup: Step By Step (Primary-HDR Secondary)

After onmode –d primary command execution

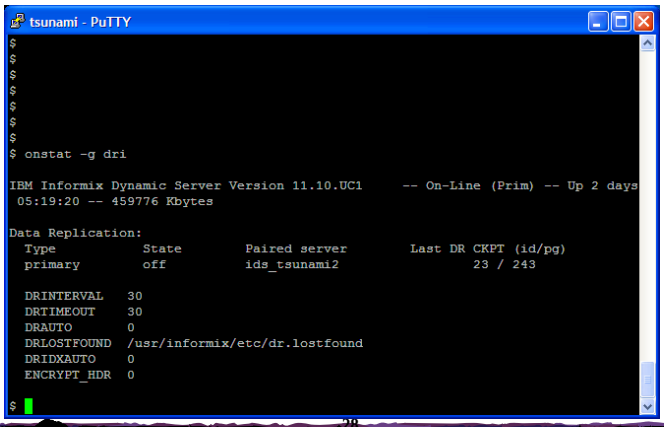

HDR Setup: Step By Step (Primary-HDR Secondary)

Physical restore on the secondary

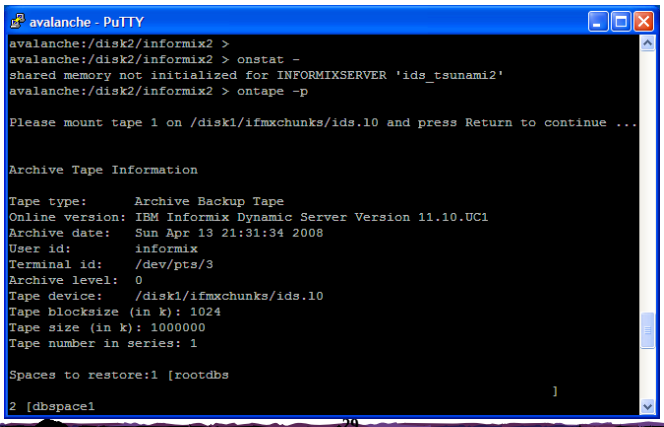

HDR Setup: Step By Step (Primary-HDR Secondary)

After onmode –d secondary command execution

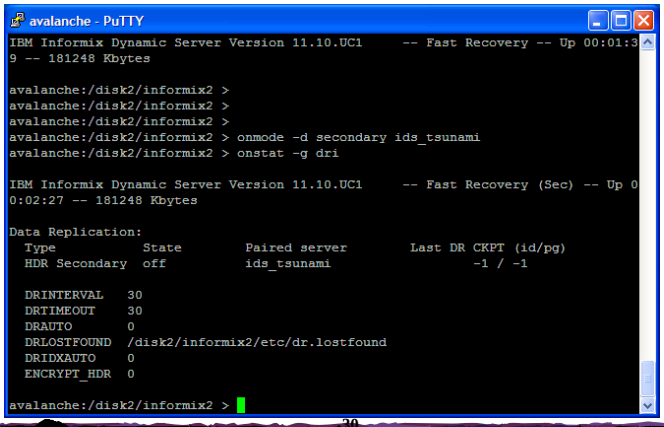

HDR Setup: Step By Step (Primary-HDR Secondary)

Replication synching up for the first time…

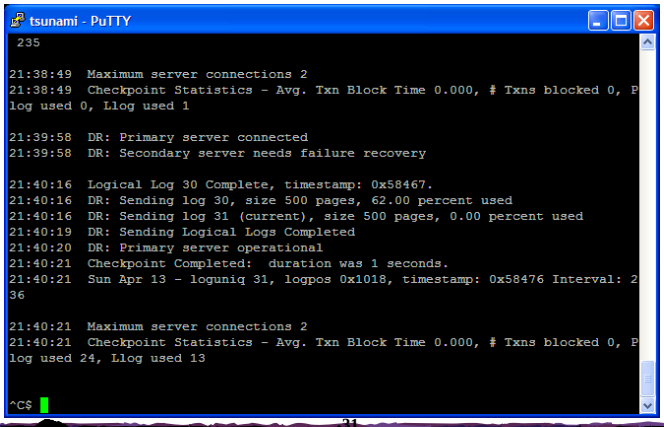

HDR Setup: Step By Step (Primary-HDR Secondary)

Replication is up (view from the primary)

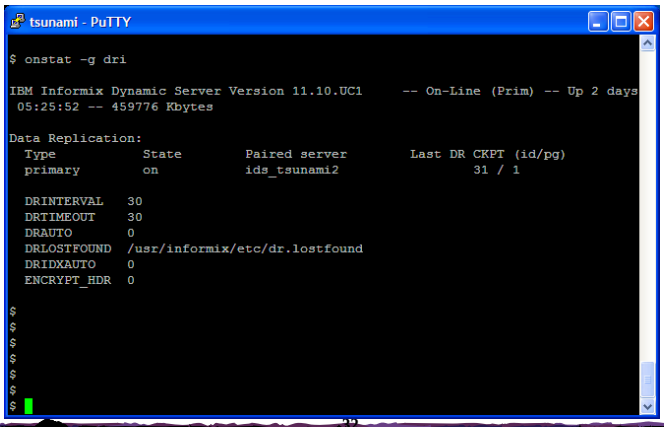

HDR Setup: Step By Step (Primary-HDR Secondary)

Replication is up (view from the HDR Secondary)

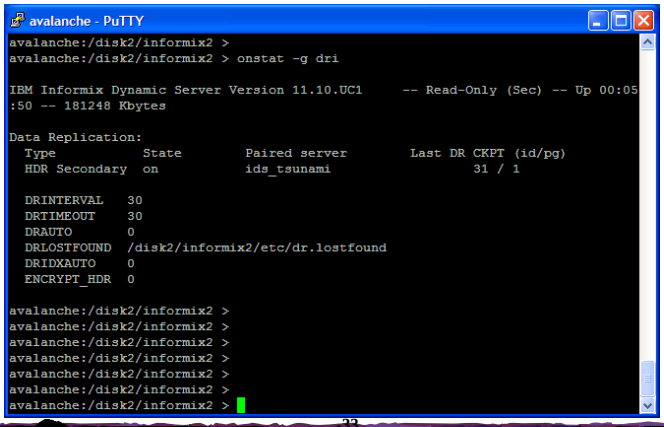

HDR Monitoring

onstat –m

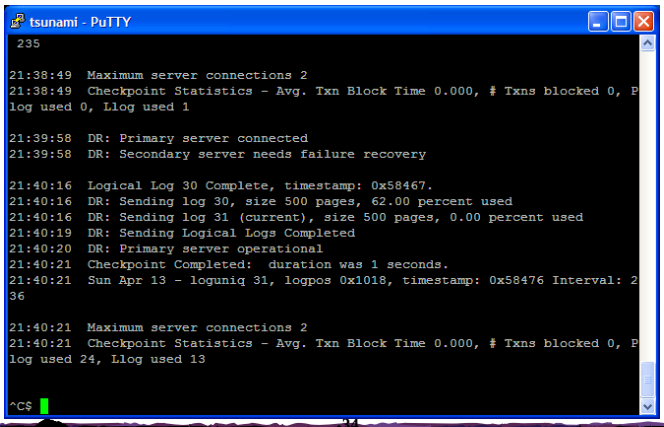

HDR Monitoring

onstat –g dri

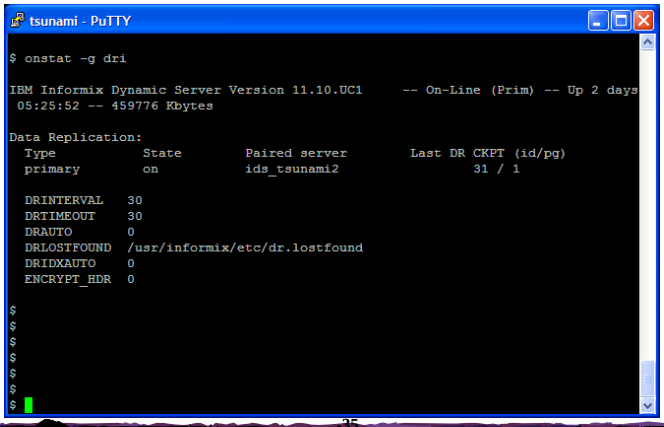

HDR Monitoring

sysmaster:sysdri table

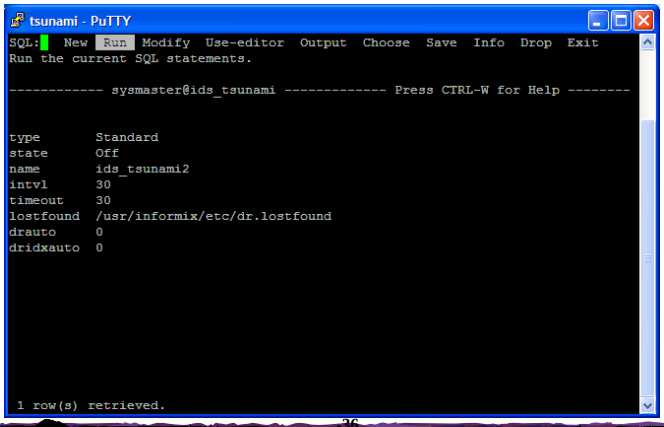

HDR Examples From The Real World

Industrial equipment distributor, Virginia

- Informix-4GL based green screen supply chain management application
- Running on IDS 7.31
- Approximately 300 users located in multiple locations in the United States
- Setup HDR in 2003 with the primary goal of disaster recovery preparedness

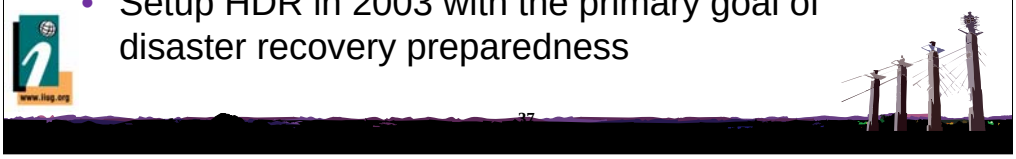

HDR Examples From The Real World

Industrial equipment distributor, Virginia

- HDR Primary located in the headquarters building IBM RS/6000 B80 running AIX 5.2 and IDS 7.31
- HDR secondary located in a co-location facility 10 miles away – identical hardware, operating system, IDS.
- Connection between primary and secondary: T-1.
- Unbuffered logging and DRINTERVAL of 30 seconds.

HDR Examples From The Real World

Industrial equipment distributor, Virginia

- HDR interruptions: four times in the past five years due to network connectivity issues.
- Reversed the roles of the servers two years ago (co-location facility better equipped).

Replication interrupted three times in the past five years due to network connectivity issues.

A couple of years ago, the business decided to reverse roles of the machines as the co-location facility was better equipped to house a mission critical server (corporate headquarters location was prone to power outages and did not have an adequate server room). Since then, the primary is in the co-location facility and the secondary is in the company headquarters. Reversing the roles involved breaking the replication and then re-synching the servers using the regular initial synching procedure.

HDR Examples From The Real World

Gift Card Transaction Processing Company, Tennessee, Kentucky

- ESQL/C based, extremely mission critical (24x7x365) back end transaction processing application.
- Processing gift cards for major retailers.
- High availability/minimizing downtime and disaster recovery preparedness are main goals.
- Primary in Tennessee, secondary in Kentucky

HDR in IDS 11.5 (Cheetah 2)

There are even more exciting replication enhancements in the upcoming 11.5 version of IDS (Cheetah 2).

These are being announced at this conference – be sure to catch these announcements if you are interested in replication technology.

These may be discussed in more detail if time (and Madison:o) allow.

Session A06

Beginner's Guide To HDR

Hal Maner
M Systems International, Inc.
hmaner@msystemsintl.com