




IBM – Informix Dynamic Server 11

“MACH-11”

Continuous Availability and Scalability




Alexander Koerner, Channel Technical Sales – Information Management

Alexander.Koerner@de.ibm.com

Certified Senior IT-Specialist, IBM Germany (Munich)

Bitte füllen Sie die Feedbackbögen aus... 😊




IBM Informix Dynamic Server 11


The Next Generation in OLTP Data Server Technology

IBM Informix Dynamic Server (IDS) is the optimal data server for high-transaction OLTP and integrated environments. Start learning about IDS today with this guide to the latest features and their value to your business. IDS helps businesses achieve Information On Demand by providing unmatched high reliability, "nearly hands-free" scalable administration, blazing-fast OLTP, and extensive platform coverage.

- Download IDS 11 trial software to benefit from the latest features:
ibm.com/informix/ids
- Take your skills to the next level with *Administering Informix Dynamic Server, Building the Foundation* (MC Press, 2008), and IDS 11.
- Keep up to date on the available Information Management resources by visiting ibm.com/software/data/education.




CARLTON DOE had over 10 years of Informix experience as a DBA, Data Server Administrator and 4GL Developer before joining Informix in 2000. During this time, he was actively involved in the local Informix user group and was one of the five founders of the International Informix Users Group (IIUG). He served as IIUG President, Advocacy Director, and sat on the IIUG Board of Directors for several years. He is best known for having written *Informix-OnLine Dynamic Server Handbook* and *Administering Informix Dynamic Server on Windows NT* both of which are currently undergoing revisions for a second edition. He is the author of numerous IBM whitepapers and technical articles and co-authored several IBM Redbooks. Carlton currently works for IBM as the Informix Technical Specialist for the America's Geographies and is a Sr. Certified Consulting IT Specialist. He lives in Dallas, TX with his family and his neighbor's dog "puppy."



MC Press Online, LP
125 N. Woodland Trail
Lewisville, TX 75077

Price: \$14.95 US/\$18.95 CN

ISBN 978-1-58347-075-6




9781583470756

DOE
IBM Informix Dynamic Server 11

IBM Informix Dynamic Server 11

The Next Generation in OLTP Data Server Technology



- Unmatched continuous availability
- Blazing fast performance
- Legendary reliability
- Virtually zero administration

CARLTON DOE

Agenda

Resilient

Agile

Invisible

IBM Informix Dynamic Server 11

- Overview – What is MACH 11?
- Near-Line Standby (Continuous Log Restore)
- Remote Standalone Secondary Server
- Shared Disk Secondary Server **COOL**
- A Complex Failover Scenario
- Live Demo

What is “MACH-11”?

- Multi-instance Active Cluster for High Availability.
 - Extends HDR to support more than a single primary with a single secondary instance.
 - Three new types of secondary instances:
 - Continuous Log Restore (CLR) or “near-line” standby**
 - Remote Standalone Secondary (RSS)
 - Shared Disk Secondary (SDS)
 - HDR, RSS & SDS technology can be used in any combination.
- “MACH-11” is not:
 - Just 1-to-N HDR.
 - “MACH-11” is the treating of all three forms of secondary as a multi-tiered availability solution.



** not exactly a MACH-11 instance but included for the purposes of discussion


Supporting Infrastructure

- Two New Subcomponents Used to Support “MACH-11”
 - Server Multiplexer (SMX)
 - Used internally to establish network connections between the nodes
 - Supports multiple logical connections over a single TCP connection
 - Index Page Logging
 - Used to log the pages created as part of an index build
 - Requirement for RSS



Server Multiplexer (SMX)

- Multiplexed network connection.
- Uses full duplex protocol
 - Sends packets without waiting for “ack” back:
 - Does monitor to make sure “acks” come back.
 - HDR uses half-duplex so primary knows secondary received transaction before committing.
- Supports encryption.
- Automatically activated.
- Requires no configuration other than encryption.



Agenda

Resilient

Agile

Invisible

IBM Informix Dynamic Server 11

- Overview – What is MACH 11?
- **Near-Line Standby (Continuous Log Restore)**
- Remote Standalone Secondary Server
- Shared Disk Secondary Server
- A Complex Failover Scenario
- Live Demo

Near-line Standby

- An enhancement to ontape and OnBar
 - **Full and backed-up** production logical logs copied to other servers supporting instances in logical roll-forward mode
 - Copied logs are applied on target instance(s) when ever the business wants
 - Maybe never, just want the logs saved somewhere
 - Maybe N times per day as set by DBScheduler
 - `ontape -l -C` (capital "C")
 - `onbar -r -l -C` (capital "C")
- Assumption is that these environments will be backed up to disk rather than tape



Setup of the near-line node



1. Create level 0 backup of primary instance
2. Copy backup to near-line server and execute a physical restore
`ontape -p`
3. As logical logs fill, they are backed up, copied to near-line server and renamed as needed
4. If near-line server is on, copied logs are applied on a regular basis

```
ontape -l -C / onbar -r -l -C
```

5. When needed,

```
ontape -l -X / onbar -r -l -X
```

stops roll forward and puts server in quiescent mode

Agenda

Resilient

Agile

Invisible

IBM Informix Dynamic Server 11


- Overview – What is MACH 11?
- Near-Line Standby (Continuous Log Restore)
- Remote Standalone Secondary Server
- Shared Disk Secondary Server
- A Complex Failover Scenario
- Live Demo

Remote Standalone Secondary (RSS)

- Similar to HDR
 - Maintains a disk copy of the database
 - Can be used for report processing
 - Created by performing a backup / restore of the instance

- Distinct from HDR
 - Uses Fully Duplexed Communications – better throughput over slower lines
 - Does not support SYNC mode, **not even for checkpoints**
 - Can **not** be ‘promoted’ to primary – but can be promoted to HDR secondary – FOR DISASTER RECOVERY, not HA
 - There can be any number of RSS nodes
 - Requires Index Page Logging

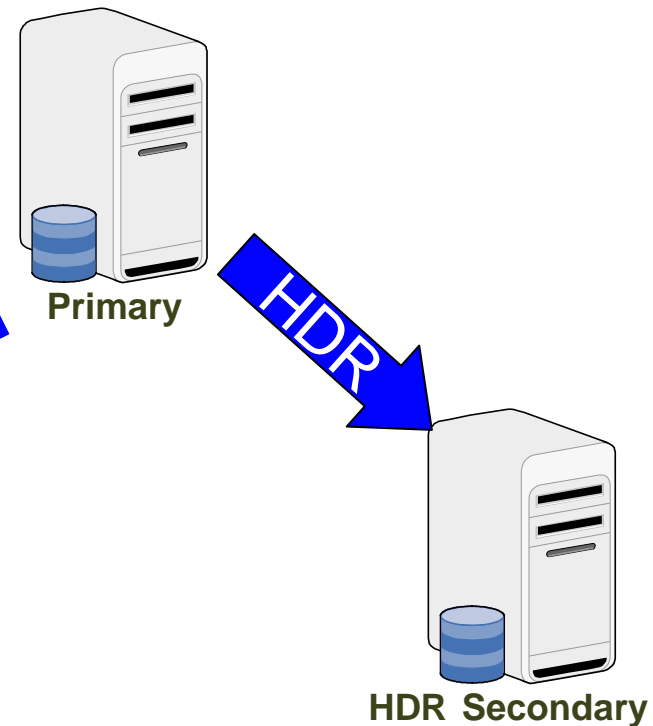
- RSS can be used in combination with HDR secondary
 - RSS can be converted into HDR secondary
 - HDR secondary can be converted into RSS




RSS – Bunker backup and more

Customer is currently using HDR for high availability but would like to have an additional backup of their system in the event of a disaster in which both the primary and the secondary should be lost.

Using HDR to provide High Availability is a proven choice. Additional disaster availability is provided by using RSS to replicate to a secure 'bunker'.



RSS – Backup to the Backup



If the primary fails, it is possible to convert the existing HDR secondary into the primary instance. If it appears that the original primary is going to be down for an extended period of time, then the RSS instance can be converted into an HDR secondary instance.

Setup of the RSS node

1. Turn on Index Page Logging

```
LOG_INDEX_BUILDS          1 ($ONCONFIG) or
```

```
onmode -wf LOG_INDEX_BUILDS=1
```

onconfig parameter must be set to activate index page logging

2. Define RSS node on Source Server

```
onmode -d add RSS <RSS Server Name> <optional password>
```

Creates an entry in the SYSHA database. Used for verification that the primary node is expecting a remote node to connect. Password is an optional entry used to authenticate the connection.

Setup of the RSS node (cont)

3. Perform a level zero backup on the primary node
4. Perform a physical restore on the RSS node (ontape -p)
5. Start HDR on RSS node

```
onmode -d RSS <source node> <optional password>
```



Agenda

Resilient

Agile


Invisible

IBM Informix Dynamic Server 11

- Overview – What is MACH 11?
- Near-Line Standby (Continuous Log Restore)
- Remote Standalone Secondary Server
- Shared Disk Secondary Server **COOL**
- A Complex Failover Scenario
- Live Demo

Shared Disk Secondary (SDS) instances

- HDR on top of a shared disk subsystem.
- Primary transmits the current Log Sequence Number (LSN) as it is flushing logs.
- SDS instance(s) receives the LSN from the primary and reads the logs from the shared disks.
- SDS instance(s) applies log changes to its buffer cache.
- SDS instance(s) resynch processed LSN to primary.



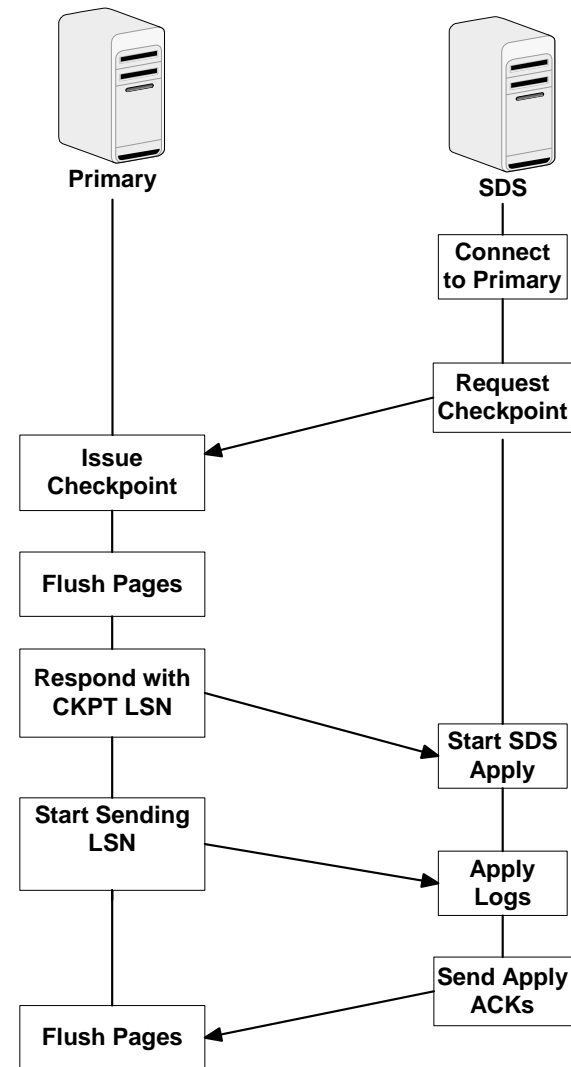
Fundamental Rule of SDS instances

- SDS instance will not write to the shared disk chunks:
 - Never flushes to shared disk chunks, even during a checkpoint.
 - If the SDS needs to remove buffer cache pages, they are placed in a temporary 'paging file' until the next checkpoint.

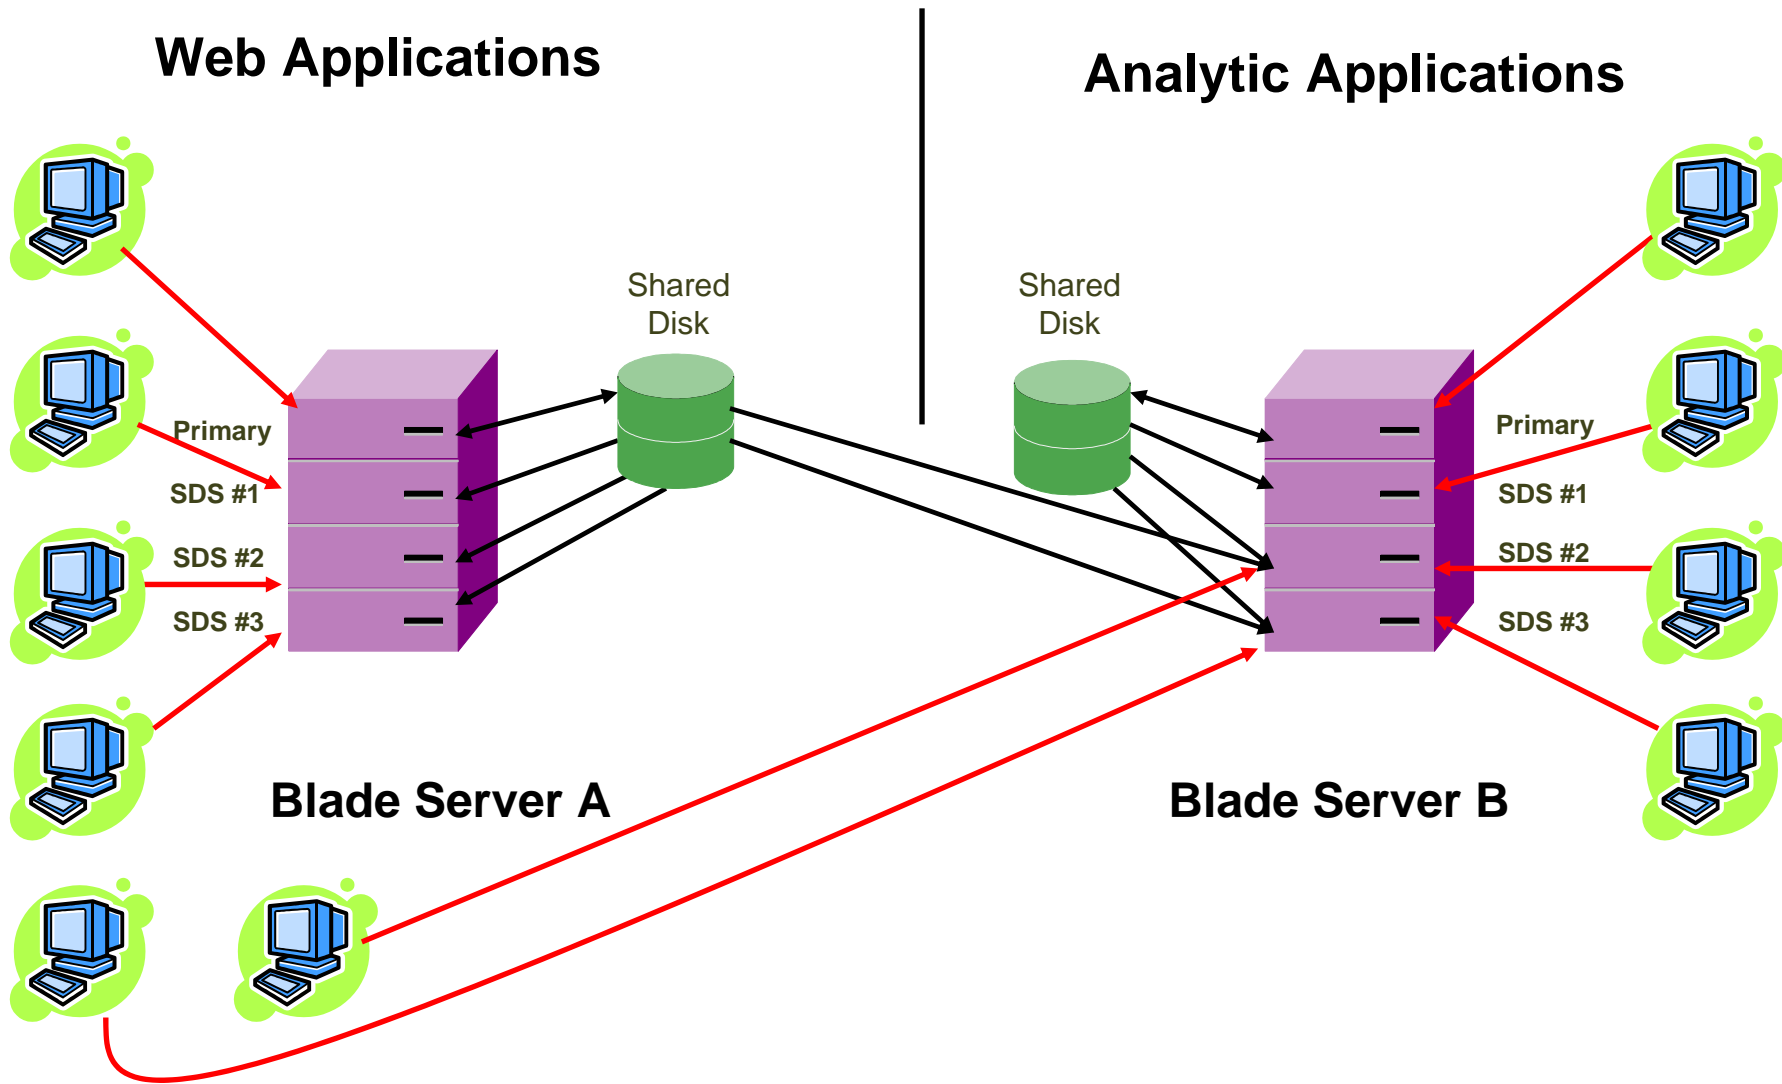
- Primary will not flush a page to disk until it is certain that the SDS instances will not need the previous version of the page:
 - Primary keeps track of the LSN (log unique ID + logpos) of every page as it is first dirtied on the primary.
 - The dirtied pages is not flushed until the SDS instances have progressed past that LSN.

Starting a Shared Disk Secondary

1. SDS instance creates SMX connection to the primary.
2. SDS instance requests a checkpoint on the primary.
3. Primary responds with the checkpoint LSN.
4. SDS instance activates recovery as of the checkpoint.
5. Primary starts periodically sending the current LSN.
6. SDS instance starts periodically sending the applied LSN.



SDS Usage: Capacity as Needed



Instantiating an SDS instance

1. Setup shared disk environment.

2. Set SDS_ENABLE on primary:

```
onmode -d set SDS primary <instance_name_of_SDS_primary>
```

Log entry (primary)

```
00:49:28 SDS Source instance Alias Changed from (<null>) to (hdr_primary)
```

3. On SDS instance, configure SDS_ENABLE, SDS_PAGING, and SDS_TEMPDBS.

4. On SDS instance, map the following to match the primary:

ROOTNAME	ROOTPATH
ROOTOFFSET	ROOTSIZE
PHYSDBS	PHYSFILE
LOGFILES	LOGSIZE



Leave other parameters the same, with the exception of
DBSERVERALIASES, DBSERVERNAME, and SERVERNUM



Instantiating an SDS instance

5. Start the SDS instance (oninit).

Log entry (SDS instance):

```
00:58:19 IBM Informix Dynamic Server Initialized -- Shared Memory Initialized.
00:58:19 Initializing SDS Temp dbspace sds1temp
00:58:19 Requesting logs from 7:2d
00:58:21 Space 'sds1temp' added.
00:58:21 Onconfig parameter MSGPATH modified from
/opt/IBM/informix/logs/hdr_primary.log to /opt/IBM/informix/logs/sds1.log.
00:58:21 Onconfig parameter DBSERVERNAME modified from hdr_primary to sds1.
00:58:21 Onconfig parameter SERVERNUM modified from 2 to 5.
00:58:21 Dataskip is now OFF for all dbspaces
00:58:21 Restartable Restore has been ENABLED
00:58:21 Recovery Mode
00:58:21 Logical Recovery Started.
00:58:21 10 recovery worker threads will be started.
00:58:21 Start Logical Recovery - Start Log 7, End Log ?
00:58:21 Starting Log Position - 7 0x2d018
00:58:22 DR: SDS secondary server operational
```

Agenda

Resilient

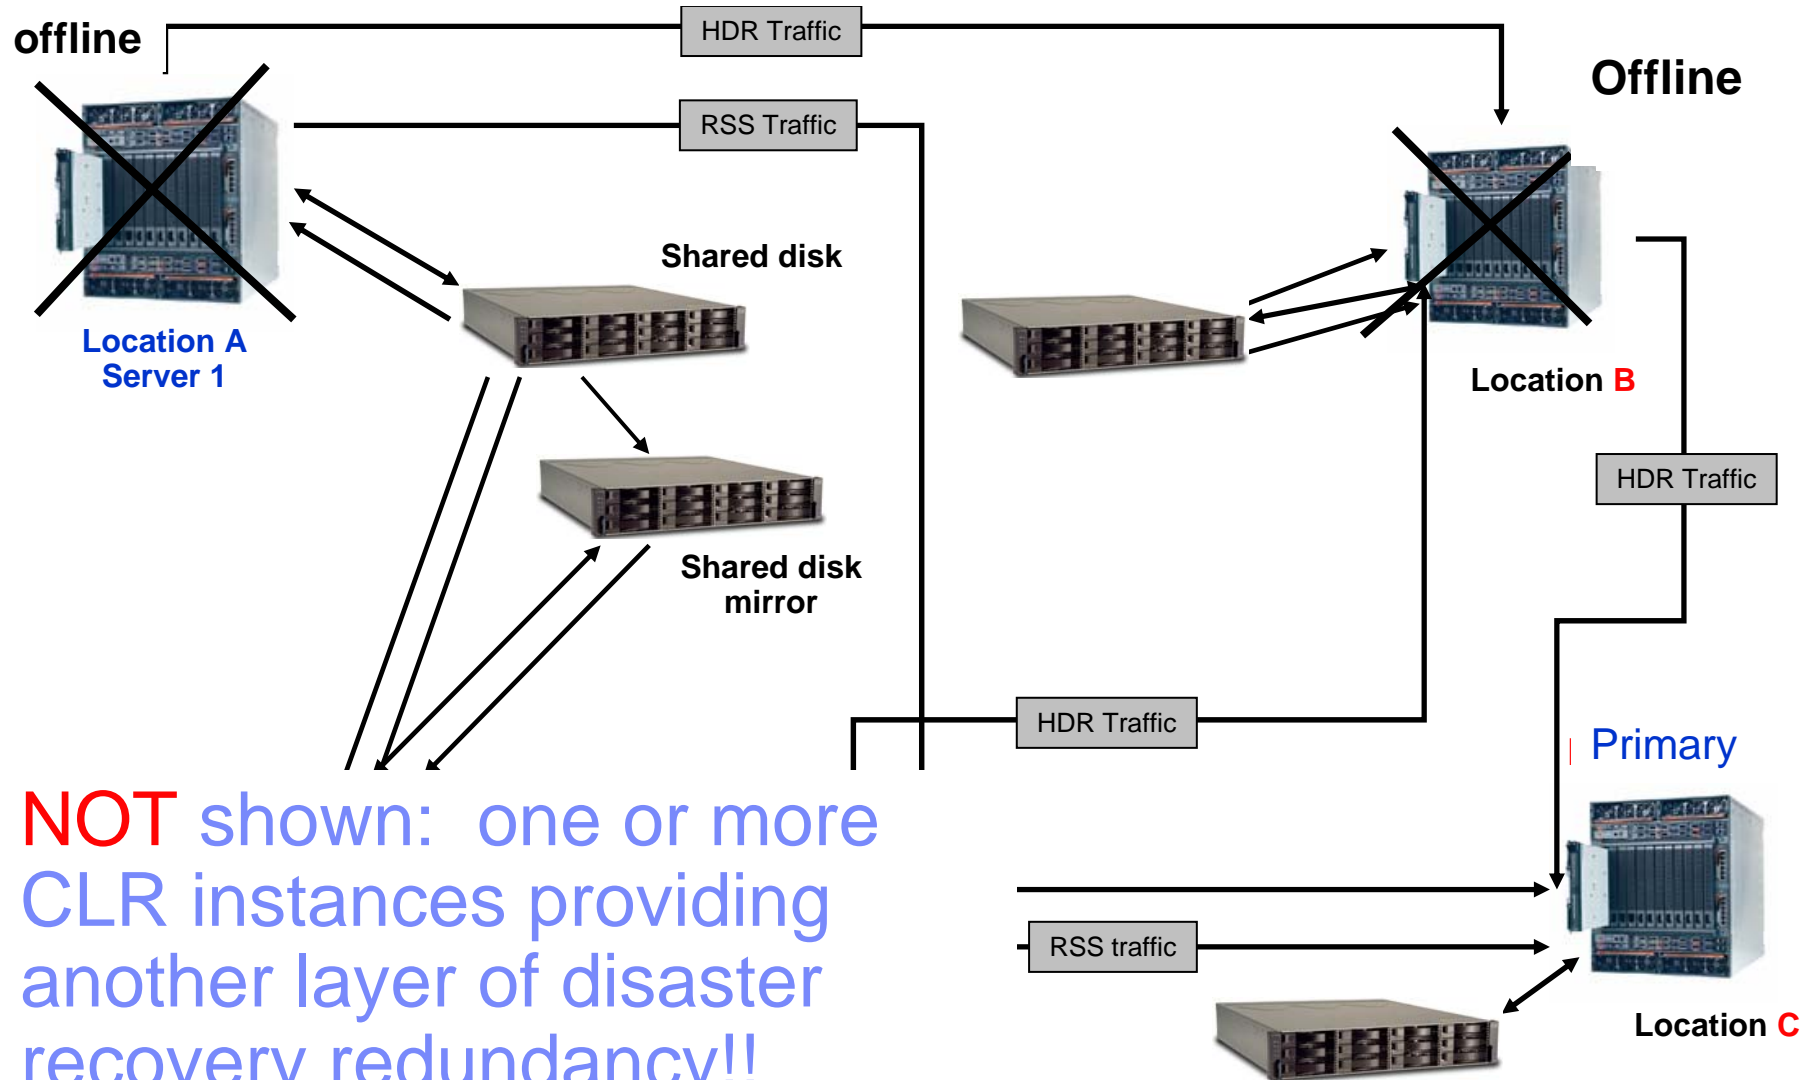
Agile

Invisible

IBM Informix Dynamic Server 11

- Overview – What is MACH 11?
- Near-Line Standby (Continuous Log Restore)
- Remote Standalone Secondary Server
- Shared Disk Secondary Server
- **A Complex Failover Scenario**
- Live Demo


The server at location B fails



NOT shown: one or more CLR instances providing another layer of disaster recovery redundancy!!

Failover

- SDS can be used with HDR and RSS.
- SDS can be promoted directly to primary:
`onmode -d set primary SDS <sds_instance_to_promote>`



When the primary is changed, other instances 'follow the change'.

COOL

- Order of failover should be:
 1. To an SDS instance.
 2. To the HDR secondary.
 3. To an RSS instance.

Not a good idea to use DRAUTO, any SDS instance will be shutdown as part of the HDR switch.

“MACH-11” with ER

- Any instance within the ER replication domain can also be in a “MACH-11” cluster.
- ER requires that all instances within the “MACH-11” cluster be in an \$SQLHOSTS ER group.
- ER will automatically ‘discover’ where the primary is.

Agenda

Resilient

Agile

Invisible


IBM Informix Dynamic Server 11

- Overview – What is MACH 11?
- Near-Line Standby (Continuous Log Restore)
- Remote Standalone Secondary Server
- Shared Disk Secondary Server
- A Complex Failover Scenario
- **Live Demo**

Questions ?


Coming Soon: The IDS 11 Redbook!

- Covers all aspects of IDS 11
 - IDS 11 Essentials / Architecture
 - Implementation / Configuration
 - SQL Language / Query Drill-Down
 - Data Privacy / High Availability
 - Functional Extensions to IDS
 - Dev. Tools, APIs and SOA
- Lots of step-by-step examples
- Availability: Q3 2007
- Download **for free** at:
<http://www.ibm.com/redbooks>
- Redbook ID: SG24-7465



I am

Ready for



Informix®
Dynamic
Server

data server software

You too?