Backup and Restore Strategy for sites without a DBA

Ron Privett
IBM

Session: D08
Tuesday, April 29, 2008 • 10:50 a.m. – 11:50 a.m.

2008 IIUG Informix Conference

www.iiug.org

Introduction

- This presentation will focus on specific backup methods and strategies that sites without a DBA would be able to implement. This presentation would appeal to end users looking for help with backup/recovery issues, or to consultants interested in archiving options for their clients.

- Goals…
  - For Novice IDS users:
    - Understand Backup/Restore options and be moved to implement a plan
  - for Experienced IDS users:
    - Use information presented to develop an automated backup/restore plan for DBA-less sites

Who is Ron Privett?

- Officially:
  - Ron is a member of the Advanced Support team, working with the Down Systems and Diagnostic group for IBM Informix engine products. He has worked with IDS for 11+ years, specializing in Replication issues a majority of that time. He has created customer tutorials, developed training materials, trained other engineers how to support ER, and presented at various user groups and user conferences in the past.
- Unofficially:
  - Ron is: a father to 3 boys (13, 9, and 7), the husband to a loving wife, a frustrated Do-It-Yourselfer, a huge U2 fan, and a Soccer fan.

Agenda

- Introduction - Why discuss Backup Strategy?
- Developing a Backup & Restore Strategy
  - Different Backup Methods
  - Creating Your Own Plan
  - Testing
- Ontape Backup to Dir
- Onbar/ISM Backup to files
- Dbexport
- Archecker
- Summary / Questions

Why Discuss Backup Strategy?

- As an IT professional - you should already know the importance of backups
- What some DBA's and Sys Admin's fail to do is to develop a plan, and test it (more than once)
- Do not be tempted to 'Set it and Forget it'
- Change happens, and even the best plans can break

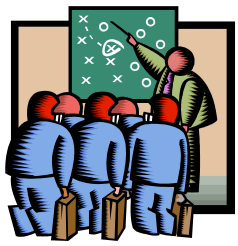

Seen this one before…

- As a member of the Down Systems group, I've seen my share of failed backup methods, and have heard many nightmare situations
- Some of the common themes are:
  - No company is safe - big or small
  - Lack of testing and resources
  - Lack of importance on Backup/Restore
  - Lack of understanding with processes

How I Lost My Job - Story #1

- Customer uses ontape via a script to backup to tape
- Instance grows over time
- Customer never checks the backup logs, or may fail to create one
- When they need to restore, ontape asks for the 2nd tape

What 2nd Tape?

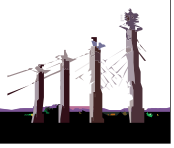

How I Lost My Job - Story #2

- Customer does a Level-0 once each month (Hmm)
- Customer archives logical logs between Level-0's
- Customer needs to restore near the end of the month
- After the 6 hour physical restore, they had approx. 1200 logical logs to restore
- Logical recovery completed after another 40+ hours

Total recovery took 46 hours!

How I Lost My Job - Story #3

- A developer dropped a production table
- Customer used onbar and decided to use the Point-In-Time option to restore the instance
- Customer started restore OVER the production environment
- Customer found out after the physical restore completed that LTAPEDEV was set to /dev/null

PIT restores *Require* logical logs

How I Lost My Job - Story #4

- Customer had been running Online 5 for years. Their DBA moved to another job, and was not replaced. A disk failure caused the staff to wonder who to call? (Informix 'the company' no longer existed). After numerous attempts to 'tbinit -iy' the instance, and even recreating the dbspaces - they still could not see their data.
- Oh… and there was no backup

oninit -iy with no backup

Agenda

- Introduction - Why discuss Backup Strategy?
- Developing a Backup & Restore Strategy
  - Different Backup Methods
  - Creating Your Own Plan
  - Testing
- Ontape Backup to Dir
- Onbar/ISM Backup to files
- Dbexport
- Archecker
- Summary / Questions

Different Backup Methods

- Full Backups
  - Usually considered a complete backup of all devices to a known point of consistency
  - Contains the physical image of device pages, logical logs, and any updates occurring during the backup
  - Might be run every day, or weekly depending upon needs

Different Backup Methods

- Incremental Backups
  - A backup of any changes since the last full backup
  - Usually faster than full backups as only a portion of the instance is saved
  - Incremental backups will gradually get longer as more changes are made, and a full backup is not done

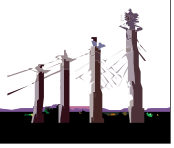

Specific IDS Backup Methods

- 2 different tools: OnBar and Ontape
- Archiving Modes: Online, Quiescent, External, Incremental
- Archives all data, or selected dbspaces (onbar)
- Continuous log backup and restore
- Restores all data, selected dbspaces, or tables
- Restore options: Cold, Warm, Mixed, Imported, Point-in-time, Redirected, Restartable
- Verification of archive (archecker)
- Ability to use a Storage Manager (onbar)
- Parallel archive and restore (onbar)
- Separate physical/logical restore (onbar)
- Archive to STDIO (onbar), or to Directory

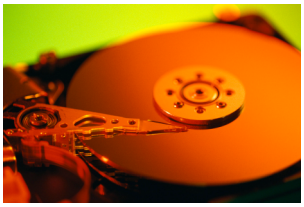

Archival Backups

- Many customers like to take an "archival" backup, one that can be kept for 5 or 10 years
  - With onbar - a parallelized whole system backup is perfect, because you don't need logical logs to restore this backup
  - With ontape - a backup to directory is ideal because you can also include other items to be archived together

Archival Backups - Requirements

1. Ontape: Level-0 archive, Onbar: Whole System Backup
2. Save a copy of the media (tar or cpio file) for this version of engine
3. Save copies of the following:
   1. onconfig file
   2. onstat -d output
   3. oncheck -pe
   4. dbschema -ss (for each database)
   5. sqlhosts file

Archival Backups - Requirements

3. Save copies of the following (cont):
   6. Environment variables of session used to create the backup
   7. output of the following SQL (for each database):

```
unload to 'systables.unl' select * from systables;
unload to 'sysfragments.unl' select * from sysfragments;
```

   8. ixbar file (onbar only)
   9. Storage Manager configuration (onbar only)
   10. Storage Manager software (onbar only)

Agenda

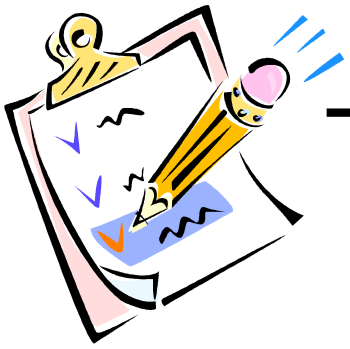

- Introduction - Why discuss Backup Strategy?
- Developing a Backup & Restore Strategy
  - Different Backup Methods
  - Creating Your Own Plan
  - Testing
- Ontape Backup to Dir
- Onbar/ISM Backup to files
- Dbexport
- Archecker
- Summary / Questions

What makes a Good Plan

- Backup and Restore Plans need the following components to be successful:
  - Documentation of the procedures
  - Communication of the details
  - Knowledge of Business Needs
  - Regular Testing of the backup and restore
  - Ability to Change
  - Archive Retention and Security

Creating Your Own Plan

- When do you start, or revise a Backup/Restore Plan?
  - Start at the beginning of System Design
  - Revisit the plan whenever changes are made
  - Test you Plan before migration
- Know the needs of the Business
  - Acceptable downtime window - Minutes, Hours, Days
  - Acceptable data loss - rows, tables, databases, dbspaces

Creating Your Own Plan (cont)

- Know your System:
  - Number of logical logs used during 24 hours
  - Time to backup / restore
  - Infrastructure
    - ontape or onbar
    - Space needed to complete full restore
    - Location of Backup Media at all times
    - Available Testing Systems
    - Available Emergency Systems

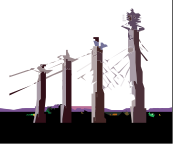

When to Backup?

- Very dependent upon your Business Needs
- Recommendations:
  - Schedule a Daily Backup
    - even if log volume is low, archive it to avoid data loss
  - Schedule a Weekly Backup
    - during a period of low activity on your system
    - use this time to run oncheck before the backup

What to Backup?

- For a full system backup - be sure to capture the following:
  - Onconfig file, oncheck -pe, dbschema -ss (from each database), environment variables used to take the backup
  - Output of the following SQL's (from each database):
    - unload to 'systab.unl' select * from systables;
    - unload to 'sysfrag.unl' select * from sysfragments;
  - Onbar only:
    - ixbar file, storage manager configuration, storage manager software

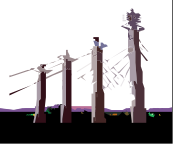

Bonus Backup's

- Multiple backup methods could really come in handy in a pinch.
- On smaller systems or systems with time, take another type of backup:
  - Dbexport
  - Onunload
  - Dbunload
  - HPL for critical tables
- Note - there are some locking issues with the above utilities

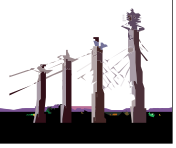

Agenda

- Introduction - Why discuss Backup Strategy?
- Developing a Backup & Restore Strategy
  - Different Backup Methods
  - Creating Your Own Plan
  - Testing
- Ontape Backup to Dir
- Onbar/ISM Backup to files
- Dbexport
- Archecker
- Summary / Questions

Why Test your Backup

- If you never try to restore it… how do you know it works?
- Testing your backup/restore procedures reduces your exposure to the unknown

- Make sure you have (and maintain) access to testing systems
- As you add space to your system, be sure to add space to your testing system

What to test?

- Test the restore completes without error
- Test the consistency of the restored instance
  - via oncheck -cr, ce, cc, cD, cI
- Test the storage manager, and any imported restore operations

- If a full test is not feasible, use Archecker to 'verify' the backup

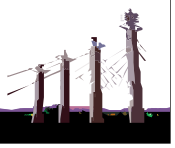

Agenda

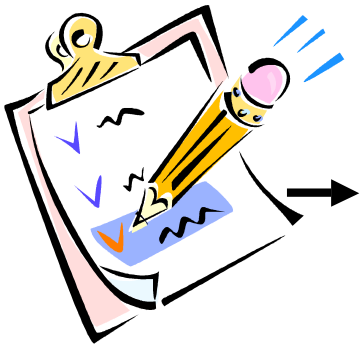

- Introduction - Why discuss Backup Strategy?
- Developing a Backup & Restore Strategy
  - Different Backup Methods
  - Creating Your Own Plan
  - Testing
- Ontape Backup to Dir
- Onbar/ISM Backup to files
- Dbexport
- Archecker
- Summary / Questions

ONTAPE Backup To Directory

- Finally!!! Support for backup to directories with ontape
  - Simple to setup and administer
  - Allows for automated archives
  - Automatically renames old archives when new ones are created

What has changed?

- ontape has been enhanced to allow TAPEDEV and LTAPEDEV to be set to a directory path
  - The directory can be a local directory or a mounted directory
  - Windows: TAPEDEV c:\IDS_BACKUP_DIR
  - Unix: TAPEDEV /informix/ids_backup_dir
    TAPEDEV /mounted_dir/ids_backup_dir
- Fully integrated with archecker verification and table level restore

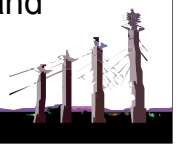

What happens during an archive?

- When this feature is used ontape will create a single file in the backup directory for each level of an archive, and for each log file backed up
- The naming convention used is as follows:
  - <hostname>_<servernum>_L<#>
    - <hostname> is the machine hostname for the instance.
    - <servernum> is the instance server number
    - L<#> is the level number of the archive
  - <hostname>_<servernum>_Log<##########>
    - <hostname> is the machine hostname for the instance
    - <servernum> is the instance server number
    - Log<##########> is the log unique id number

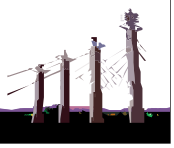

Backup to Directory - Example

- 1st Backup taken with TAPEDEV set to c:\IDS

```
C:\IDS>ontape –s –L 0
10 percent done.
20 percent done.
30 percent done.
40 percent done.
100 percent done.
File created: c:\IDS\IBM-65G8X7HJPJY_0_L0

Please label this tape as number 1 in the arc tape sequence.
This tape contains the following logical logs:

 38

Program over.
```

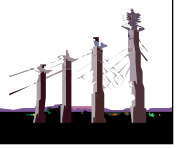

Backup to Directory - Example (cont.)

- 2nd Backup taken with TAPEDEV set to c:\IDS

```
C:\IDS>ontape –s –L 0
10 percent done.
20 percent done.
30 percent done.
40 percent done.
100 percent done.
File created: c:\IDS\IBM-65G8X7HJPJY_0_L0

Please label this tape as number 1 in the arc tape sequence.
This tape contains the following logical logs:

 38

Program over.
```

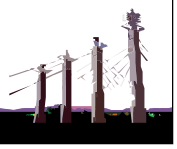

Backup to Directory - File Naming

- Naming of the backups

```
Directory of C:\IDS

04/06/2007  03:22 PM    <DIR>          .
04/06/2007  03:22  PM    <DIR>          ..
04/06/2007  02:49 PM        16,941,056 IBM-
   65G8X7HJPJY_0_20070406_144941_L0
04/06/2007  03:22 PM        16,941,056 IBM-65G8X7HJPJY_0_L0
               2 File(s)     33,882,112 bytes
               2 Dir(s)  50,500,202,496 bytes free
```

- First L0 backup was renamed when 2nd backup taken.

  The backup file <hostname>_<servernum>_L0 was renamed to <hostname>_<servernum**>_<YYYY-MM-DD_HHMMSS>**_L0

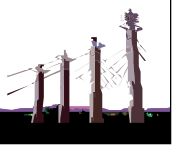

Backup to Directory - Log Example

- Log backup taken with LTAPEDEV set to c:\IDS

```
C:\IDS>ontape -a -d
Your evaluation license will expire on 2007-10-01 00:00:00

Performing automatic backup of logical logs.

File created: c:\IDS\IBM-65G8X7HJPJY_0_Log0000000041
Do you want to back up the current logical log? (y/n) n

Program over.
```

- Current log **NOT** backed up

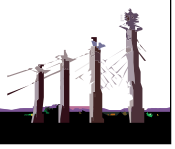

New Ontape option

- Non-Interactive mode -d
  - -d can be used where ontape would normally be interactive
    - Log backup – prompt to backup current log
    - Restore – prompt to do log salvage
  - Useful for ALARMPROGRAM or scripting ontape to perform unattended backups

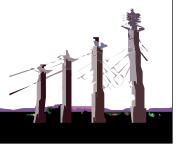

New Environment Variable

- IFX_ONTAPE_FILE_PREFIX
  - An environment variable that ontape will use to change the prefix portion of the backup file name
- Export IFX_ONTAPE_FILE_PREFIX=RCP
  - RCP_L<#>
  - RCP_Log<##########>

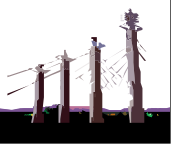

IFX_ONTAPE_FILE_PREFIX Example

- Backup taken with TAPEDEV set to c:\IDS and IFX_ONTAPE_FILE_PREFIX = RCP

```
C:\IDS>ontape -s -L 0
10 percent done.
20 percent done.
30 percent done.
40 percent done.
100 percent done.
File created: c:\IDS\RCP_L0

Please label this tape as number 1 in the arc tape sequence.
This tape contains the following logical logs:

 42

Program over.
```

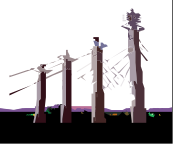

IFX_ONTAPE_FILE_PREFIX Anomaly

- Log backups taken without the environment variable will have a naming convention different than log backups taken with the environment variable on
- It is not recommended to mix log backups with different naming conventions
- You will not be able to restore all your log files unless you rename them with a single naming convention

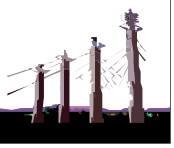

Agenda

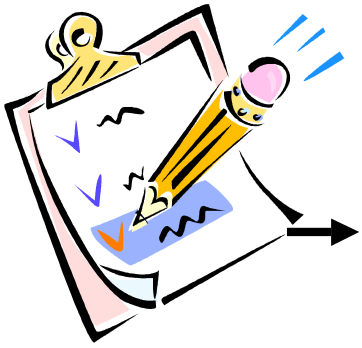

- Introduction - Why discuss Backup Strategy?
- Developing a Backup & Restore Strategy
  - Different Backup Methods
  - Creating Your Own Plan
  - Testing
- Ontape Backup to Dir
- Onbar/ISM Backup to files
- Dbexport
- Archecker
- Summary / Questions

Onbar/ISM Backup to files

- For smaller systems, or systems that are contained - using onbar with ISM can provide many benefits over ontape.
- While there is more setup involved, it could be argued the features available with onbar are enough to make it worth the effort.

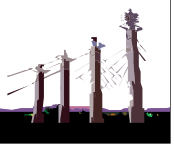

Onbar/ISM Advantages over ontape

- Archive of selected dbspaces
- Table Level Restore
- Point-in-time Recovery
- Restartable restore
- Ability to use a Storage Manager
- Parallel archives and restores
- Archive to STDIO

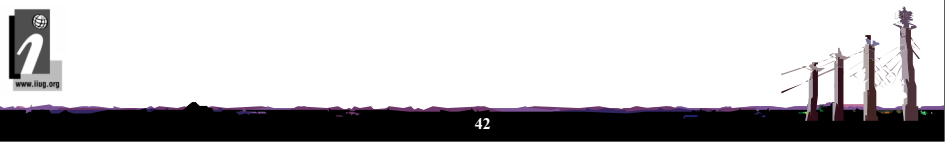

Pro/Cons of using Storage Manager

Pros	Cons
• Separates the Archive and Storage functions	• Media Management can be a problem
• Automatic management of media	• Multiple uses of backup devices
• Storage Managers usually have faster devices, and several of them	• Support of Storage Manager issues may require additional resources

Setting up ISM for Backup to files

- Refer to this document for details:
  http://www-1.ibm.com/support/docview.wss?uid=swg21105298
- Basic Steps are:
  - Create directories to store files
  - Set onconfig parameters ISM_DATA_POOL, ISM_LOG_POOL
  - Set LTAPEDEV to anything other than /dev/null
  - Run '`ismstartup -init`'
  - Add admin user, devices for Data and Log
  - Label, then mount the devices

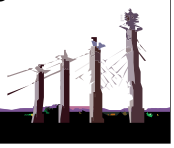

Agenda

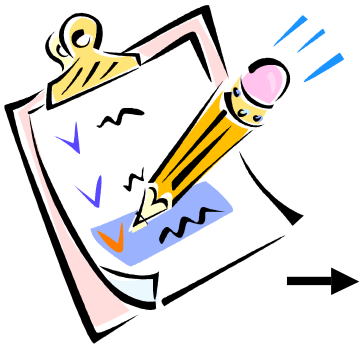

- Introduction - Why discuss Backup Strategy?
- Developing a Backup & Restore Strategy
  - Different Backup Methods
  - Creating Your Own Plan
  - Testing
- Ontape Backup to Dir
- Onbar/ISM Backup to files
- Dbexport
- Archecker
- Summary / Questions

How can dbexport help?

- What if…
  - a table is dropped on accident
  - someone deletes rows from a table by mistake
  - a table becomes corrupt
  - a bad update occurs, and you want to restore the old version of the row
- A dbexport of the database can help in any of the above situations

5 reasons to love dbexport

- Quick and easy to run (and automate)
- Unloads data to ASCII files
- Supports unloads of BLOB and CLOB columns
- Unloads the schema of all tables
- Unloads all SPL text and UDR functions

Using dbexport

- There are several options with dbexport

```
Usage:
 dbexport <database> [-X] [-c] [-q] [-d] [-ss [-si]]
[{ -o <dir> |
   -t <tapedev> -b <blksz> -s <tapesz>
  [-f <sql-command-file>] }]
NOTE: arguments to dbexport are order independent.
```

- Suggested options are:
  - -ss = server specific syntax
  - -o <dir> = directory where <database>.exp is created
  - -X = If you have byte data

Agenda

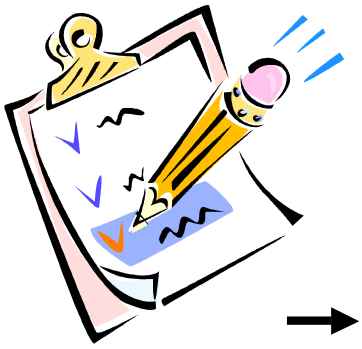

- Introduction - Why discuss Backup Strategy?
- Developing a Backup & Restore Strategy
  - Different Backup Methods
  - Creating Your Own Plan
  - Testing
- Ontape Backup to Dir
- Onbar/ISM Backup to files
- Dbexport
- Archecker
- Summary / Questions

Why use Archecker?

- Intergraded into ontape -v and onbar -v
- Helps to 'verify' or 'check' that all the items needed to restore a backup, exist within the backup
- Helps you see a *potential* problem before it becomes a *real* problem
- Could also be used to restore a table, or a dbspace (with help from Tech Support)

What is Checked?

- Format of each page on the archive is checked (similar to oncheck -cd)
- Tape control pages are sanity checked
- Table partition pages, data pages, and index pages are checked ensuring all pages exist in the archive
- Table extents are checked for overlaps
- Reserve page format is validated
- Each chunk free list is verified

Configuration files/settings

- Configuration file is $INFORMIXDIR/etc/ac_condig.std
- Important parameters are:
  - AC_MSGPATH # archecker message log
  - AC_STORAGE # Directory used for temp storage
  - AC_VERBOSE # 1 verbose, 0 terse messages
  - AC_TAPEDEV # archive file (or tape device)
  - AC_TAPEBLOCK # Tape block size (Kbytes)

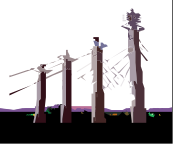

Output produced

- Normal output would resemble:

```
Tape type:      Archive Backup Tape
OnLine version: IBM Informix Dynamic Server Version
  11.10.TC2
Archive date:   Thu Mar 06 08:53:11 2008
Archive level:  0
Tape blocksize:  16384
Tape size:  20480
Tape number in series:  1

Scan PASSED
Control page checks PASSED
Reserve page validation PASSED
Table checks PASSED
Tables/Fragments validated:       374
Archive Validation PASSED.
```

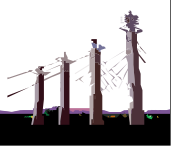

Files produced

- Archecker creates 3 directories in the AC_STORAGE location
  - CHUNK_BM
    - Contains chunk bitmap pages
  - INFO
    - Information for support, oncheck -pr, oncheck -ce broken down by dbspace, tape information, etc (All ASCII text files and readable)
  - SAVE
    - Contains binary pages for specific page types

Agenda

- Introduction - Why discuss Backup Strategy?
- Developing a Backup & Restore Strategy
  - Different Backup Methods
  - Creating Your Own Plan
  - Testing
- Ontape Backup to Dir
- Onbar/ISM Backup to files
- Dbexport
- Archecker
- Summary / Questions

Summary

- Developing a Backup & Restore Strategy
  - Different Backup Methods
  - Creating Your Own Plan
  - Testing
- Ontape Backup to Dir
- Onbar/ISM Backup to files
- Dbexport
- Archecker

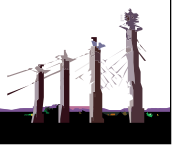

Questions?

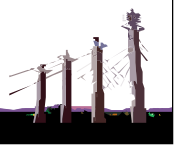

Thank You!

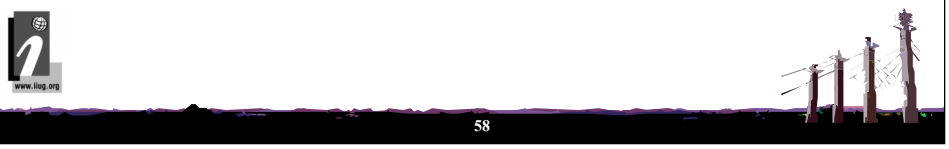

Contact Info:

- Session = D08
- Title = Backup and Restore Strategy for Sites Without a DBA
- Speaker = Ron Privett
- Company = IBM
- e-mail = rprivett@us.ibm.com