



A World of Difference in
Database Management



Best Practices for Managing Highly-available MANMAN (CODASYL DBMS) Databases (what I've learned in the last 25 years)

Software Concepts International, LLC
402 Amherst Street, Suite 300
Nashua, NH 03063, USA

Phone: 603-879-9022
e-mail: holland@sciinc.com
www.sciinc.com

Presentation Overview

- About Software Concepts International, LLC
- Database Administration Requirements
- MANMAN/DBMS Database Configuration
- Baseline system recommendations
- Ongoing maintenance
- Special tasks
- Summary
- Questions and Feedback



About Software Concepts

- Located in Nashua, NH (USA)
- Celebrating our 21st year anniversary
- International reputation
 - Leading provider of remote OpenVMS management and Database Administration services for Rdb and DBMS databases
- Proven track record
 - Actively managing over 100 databases
 - Remote DBA service since 1995 (still supporting many of the same sites)



What do MANMAN/DBMS sites have in common?

- MANMAN is critical to the business
- Large investment in technology
- High performance requirements
- Need for disaster tolerance
- High availability requirements
(few, if any opportunities for downtime)



MANMAN/DBMS Administration

Simple requirement:

“To protect the integrity, availability, reliability, performance and security of the database”

...24 x 7

Key Issues:

- Database Reliability/Integrity
- Database Performance
- Database Availability



Preparation

(for Integrity, Availability, Reliability, Performance and Security)

The database is only as strong as its weakest component.

To begin supporting a production DBMS database, an audit and review must be performed of the following:

- Database Configuration
- System Configuration
- Storage Configuration
- ...and the application!



What not to do

The following are discoveries found during pre-support audits of production database at real customer sites

- AIJs were not enabled
- RUJs and database backups shared disks
- AIJ & database shared same disk
- RUJ & AIJ files shared same disk
- Databases had never been verified
- Databases had never been analyzed (and extended many times)
- Snapshot files grew endlessly (x time area size)
- High-water marking was enabled
- SPAM thresholds had default values
- Database configured with default buffering (10 buffers, 10 blocks)
- Fail to schedule backups – not all databases were backed up.
- Delete prior backup files before verifying if they have been backed up
- Keep all data forever (never archive)

Yes, these were found in critical production databases!



Database Configuration

- Database design
- File placement
- AIJ configuration
- RUJ configuration (placement is critical)
- *Configure & enable record caching
- Number of nodes (does more really mean better?)
- Enable operator notification
- Buffering
- Snapshots (enabled, deferred)
- Fast Incremental Backups (disable?)

*performance option



Database/Schema design

- In the MANMAN environment, you pretty much have to live with what you have been given by the vendor...
- Consider data volatility/loading patterns when creating/choosing
- Intelligent storage area design.
- ...and much more



File Placement

Proper file placement is critical to both the performance and recoverability of the database


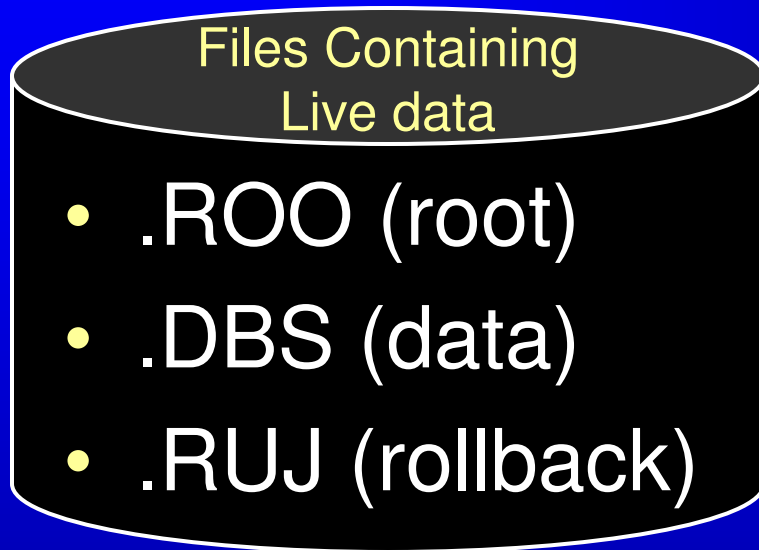
Database files can be grouped into two categories:

1. Files that contain current/live data
2. Files needed to restore/recover your database

Files from *different* groups must NOT be placed on the same physical disk!



File Placement (Live vs Recovery)



Keep on separate
Physical disks

Never place “recovery” files on the same device as “Live data” files

.SNP files intentionally omitted (they can be rebuilt)



AIJ Configuration

- **Enable AIJs**
(I can't believe I actually had to say this)
- **Create multiple AIJs**
(ideally, enough to hold transactions between backups during peak load)
- **Place AIJ files on multiple disks**
- **Explicitly define backup files for each AIJ**
- **Place AIJ backup files on multiple disks.**
- **Create extra AIJ journal slots**
(allows AIJs to be created online, if needed)
- **Make each AIJ “relatively large”**
(but not so large that the AIJ initialization causes excessive stalls)



AIJ Configuration

- Perform manual AIJ backups, but...
- Enable the AIJ backup server
- Enable AIJ operator notification
- Enable the AIJ log server
(allows for “emergency” AIJ creation – see next)
- Enable “emergency AIJ” creation

```
Define/system DBM$BIND_ALS_CREATE_AIJ 1
```

```
Define/system DBM$BIND_AIJ_EMERGENCY_DIR <device:[directory]>
```




RUJ Configuration

By default, RUJ files are created on the same device as SYS\$LOGIN

Users may be on different devices -- so you never know where the RUJs are!

- Explicitly specify the device and directory for RUJ files
(that is not on the same device as any AIJ files or backup files)



Record Cache

Performance feature

- Cluster_nodes = 1
- Database open mode = “manual”
- Enable fast commit
- Reserve & create record caches
- Enable RCS and DBR log files (DBM\$BIND_RCL_LOG_FILE and DBM\$BIND_DBR_LOG_FILE). Limit log file size.
- Use “reasonable” RCS and FC checkpoint intervals (no more than 15-30 minutes).
- Force manual checkpoints on a regular basis (15 minutes?) (dbo/checkpoint/wait)
- Long-running read/write transactions may block the RCS from checkpointing.
- Row cache will not help if:
 - Records are typically insert-only
 - Sequential scans



Miscellaneous

- Specify location for DBMS bugcheck files
(define/sys DBM\$BUGCHECK_DIR)
- Create & enable Snapshots
(preferably deferred)
- Disable Fast Incremental Backups
(unless you actually perform incremental backups)
- Configure Adjustable Locking
- Consider hot-standby options
- Versions...stay “relatively” current
(read release notes)



Logging DBMS server processes

DBMS uses detached processes to perform many functions that are critical to DBMS performance and functionality.

Knowing what these processes are doing is invaluable for analyzing problems and improving availability.

DBMS allows you to enable logging for the following types of processes:

1. After Image Journaling (ABS, ALS)
2. Database Recovery (DBR)
3. Hot-standby (LCS, ALS, LRS)
4. Row Cache (RCS)



DBS server logging logicals

- \$ DEFINE/SYSTEM DBM\$BIND_ABS_LOG_FILE
DBM\$BUGCHECK_DIR:ABS_PID.LOG
- \$ DEFINE/SYSTEM DBM\$BIND_ALS_OUTPUT_FILE
DBM\$BUGCHECK_DIR:ALS_PID.LOG
- \$ DEFINE/SYSTEM DBM\$BIND_DBR_LOG_FILE
DBM\$BUGCHECK_DIR:DBR_PID.LOG
- \$ DEFINE/SYSTEM DBM\$BIND_LCS_OUTPUT_FILE
DBM\$BUGCHECK_DIR:LCS_PID.LOG
- \$ DEFINE/SYSTEM DBM\$BIND_LRS_OUTPUT_FILE
DBM\$BUGCHECK_DIR:LRS_PID.LOG
- \$ DEFINE/SYSTEM DBM\$BIND_RCS_LOG_FILE
DBM\$BUGCHECK_DIR:RCS_PID.LOG



System Configuration Issues:

Architecture

- HP Integrity (based on Itanium architecture)
- Alpha (based on AXP chip)
- VAX
- HP Proliant (using VAX or Alpha emulator software)
(Yes, you can run MANMAN on a PC server!)

OpenVMS Versions:

- v6.1 or higher
- v7.2 for DBMS v7.1+ (on Alpha)
- v8.2 for DBMS v7.2+ (on Alpha)
- v8.3 for Integrity

Disk High-water marking

- Disable, except for DoD environments.

XFC file Cache

- Buy lots of memory, and allocate to XFC liberally



System Configuration Issues:

Redundancy

- Mirror everything
- Redundant I/O paths

SAN considerations...

- SAN environments hide the physical disk configurations that are beneath the “disks” we see at the OpenVMS level. This is good...and bad!
- Despite the “impossibility” of any failures with your new storage array, this “magic” is done in software.
- Configure your SAN environment to allow physical separation of live DB files from their recovery files.



Plan your recovery!

Specify your recovery requirements – *then* design a backup strategy to meet those requirements.

~~Backup~~ Recovery Strategies

- A successful recovery requires:
 - A valid full database backup
 - The most recent incremental backup (if any)
 - ALL AIJ files since the last DB backup (full or incr)
- Hot-standby database
- Do you KNOW when you will need to recover...
- If you don't have ALL of these when you need them, you can't buy them.

Unfortunately, this says that we have to perform backups that we hope we will never use



Recovery design

What's important to you?

- Confidence in successful recovery?
- Ease of restore?
- Minimize time to recover?
- Minimize ongoing impact to production?
- Site disaster protection?



Backup

- Use DBMS backup utility
 - Do not use VMS backup or file-level backups
 - Do not split shadow sets
- Perform AIJ then DB backups
(additional AIJ backups can be performed independently)



Backup decisions?

- Online or offline?
 - Offline (database is shutdown)
 - Online
 - Requires snapshots
 - Likely to impact performance
- Quiet_point versus noquiet_point
 - Quiet_point (known starting point)
 - Requires “quietpoint lock” – which may block new user transactions
 - Noquiet_point (may require prior AIJ files for recovery)

SCI's “Zero-Impact” backup strategy provides online backups w/out the need for snapshots or quiet points!

Ongoing Maintenance...

- Define an appropriate recovery strategy
(then develop a backup plan)
- Execute, test and monitor your backup plan!
- Perform scheduled AIJ backups
- Perform scheduled database backups
- Perform scheduled verifications/consistency checks
- Scan for Corrupt Page Table entries
- Search for and analyze all bugchecks
(this may be an early corruption indicator)



...Ongoing Maintenance...

- Monitor the database for parameter changes
- Monitor the database for schema changes
- Monitor database attaches for old TSNs
(remove processes with ancient TSNs)
- Reopen DBMS monitor log files daily
- Perform real-time monitoring of DBMS monitor log files for early problem notification
- Analyze DBMS monitor log files for excessive abnormal process terminations



...Ongoing Maintenance

- Monitor hot-standby (if enabled)
- Perform scheduled usage analysis
(use historical data to perform trend analysis and forecasting)
- Collect run-time performance statistics with
DBO/SHOW STATISTICS
- Automate notification/diagnostic data
collection/resolution of long database stalls
- Perform and review database audits of DB
activity.



Control Procedures

The best routines and procedures are only effective
if they are working and monitored

- All failures must be trapped and integrated with trouble-ticket reporting system
- Critical failures must alert appropriate support staff immediately & integrate with Voice Response Systems
- A checklist is required to validate the successful completion of scheduled tasks
 - Reporting is required for missing/late tasks.



Special Tasks

“Other tasks as necessary...”

- Periodic database restructuring
- Disaster recovery
- Corruption repairs



Summary

When properly configured *and* managed,
MANMAN and CODASYL DBMS
on OpenVMS provides a
highly reliable, high-performance
database for mission-critical
environments.





A World of Difference in
Database Management



Questions?