

IBM Informix System Administration Foundations

Instance configuration, monitoring, storage and recovery

FA-0628 · 4 days

Mainframe & Enterprise Systems · DevOps, Cloud & Infrastructure · Data & Analytics

Configure and inspect a training Informix instance, review memory and storage, and practise a small backup/restore workflow.

Introduction

This four-day course introduces Informix server administration for participants with SQL and Linux experience. It covers connectivity, configuration, memory, virtual processors, storage, logs, monitoring, maintenance and troubleshooting.

The source was written for Informix 12.10. The course retains those foundational areas but checks commands, configuration parameters and prerequisites against the agreed training release. Legacy 12.10 support depends on its applicable lifecycle and agreement; it is not presented as the default current installation.

Hands-on work uses a disposable instance with a saved configuration baseline. Broad tuning, fault-tolerance and optional configuration topics use selected demonstrations or worked examples rather than a claim to optimise or administer any production system in four days.

Learning outcomes

The course teaches participants to:

- Identify server components and inspect release/platform prerequisites.
- Configure a simple training instance and verify client connectivity.
- Use sysmaster, onstat and appropriate inspection tools to examine server state.
- Explain memory, virtual processors, storage spaces, logs and checkpoints.

- Apply a selected configuration change in the training instance and compare evidence.
- Plan a small backup/log workflow and practise a supported restore exercise.
- Inspect sessions and event alarms and document a troubleshooting finding.
- Describe production review, capacity, retention, recovery and escalation needs.

Prerequisites

- Knowledge of Structured Query Language (SQL)
- Knowledge of dbaccess
- Experience using basic Linux functionality
- Experience using the vi editor

Access to a prepared disposable Informix/Linux instance with permitted administration rights, compatible client tools, backup space and saved baseline configuration.

Check the selected release's platform/kernel requirements, licensing, support, parameter availability and backup compatibility before delivery.

Training outline

1. Day 1 — Server Setup, Connectivity and Monitoring

Server Architecture

- Describe the major components of the Informix Database Server
- Define database server terms
- Describe major database server events

Configuration Preparation

- Review the contents of the release notes
- Identify applicable kernel settings from the selected release/platform documentation; do not apply a universal tuning recipe.
- Review required operating-system patches and compatibility for the chosen training environment.

Client Connectivity

- List the methods to connect to the database server
- Make appropriate entries into the sqlhosts file

- Set the environment variables necessary for a client application to connect to the database server

Instance Configuration

- Prepare designated disposable storage using a supported file/device method; no production-device reinitialisation.
- Set environment variables necessary to administer a database server
- Explain the purpose of the configuration file
- Identify basic configuration parameter settings required to bring up a simple database server
- Use a text editor to modify the configuration file
- Bring the database server online

Monitoring

- Describe the use of the SMI (System Monitoring Interface)
- Connect to the sysmaster database and run simple SQL queries
- Use onstat to monitor information about the database server
- Use oncheck to gather information about space usage and the status of indexes and tablespaces

2. Day 2 — Memory, Processors, Storage and Logs

Shared Memory and Modes

- List and describe the database server operating modes
- Use onmode and oninit to change the operating mode
- Describe the shared memory segments used by the database server
- Describe the configuration parameters that affect the size of shared memory
- Review RESIDENT or equivalent release/platform-specific memory-residency controls; evaluate applicability before changing them.
- Monitor the major components of shared memory

Virtual Processors and Threads

- Define a thread
- Describe the multithreaded architecture
- Describe how the virtual processors are implemented in Unix
- Use onstat to monitor VPs and threads
- List and explain the virtual processor classes

- Describe how network connections are handled by the server
- Set server configuration parameters related to VPs and threads
- Demonstrate permitted dynamic VP changes where supported and observe effects in the lab.

Disk Architecture

- Describe the Dynamic Server page and chunk structure
- Explain how pages are allocated within extents
- Explain the difference between table and index extents
- Describe the relationship between chunks and dbspaces
- Describe the overhead pages in various types of chunks
- Describe the structure of a dbspace, blobspace, and sbospace
- Use the oncheck command to display reports about space allocation, tblspace usage, and index usage
- Use onstat to display dbspace and chunk information

Storage Management

- Create/drop only designated disposable storage spaces after reviewing dependencies and recovery implications.
- Review mirroring-status behaviour through a prepared example where the training platform supports it.
- Monitor dbspaces, blobspaces, and sbospaces
- Use the storage pool and the APIs to administer storage spaces

Physical and Logical Logs

- Review physical-log placement and practise a bounded change only where supported by the disposable setup.
- Review and practise selected logical-log changes with a saved lab baseline and backup plan.
- Change the logging mode of a database
- Define a long transaction
- List and describe configuration parameters related to physical and logical logging
- Monitor the physical and logical logs

3. Day 3 — Checkpoints, Recovery and Temporary Objects

Buffer Pools and Checkpoints

- List the steps performed by the server to update a row on disk

- Describe the checkpoint process
- Describe the differences between a blocking and a nonblocking checkpoint
- Assess release-specific checkpoint parameters and evidence rather than promise an optimal fixed setting.
- Describe what events cause a checkpoint to occur
- Monitor checkpoints and buffer writes performed by the server

Fault Tolerance

- Describe the types of failures that could occur on the server system
- Describe the fast recovery process
- Examine the status of chunks and dbspaces
- Identify the configuration parameters associated with recovery

Backup and Restore Concepts

- Identify Informix backup and restore utilities
- Explain the backup and restore process
- Explain the incremental backup levels
- Explain the two different methods to backup logical logs
- Distinguish warm/cold restore conditions and relevant server modes; restoration requires suitable backups and logs, not merely a restart.

ontape Exercise

- Configure your server to use ontape for backup and restore
- Perform a server backup using ontape
- Backup logical log files using ontape
- Practise or demonstrate a supported warm restore of selected training storage spaces and verify consistency; use the documented mode and log requirements.

Temporary Objects

- Describe the temporary objects created by the database server
- Describe the configuration parameters and environment variables related to temporary objects
- Explain the importance of using multiple temporary dbspaces
- Explain the purpose of the WITH NO LOG clause when used to create temporary tables
- Use onstat and the SMI to monitor the space used to create temporary objects

4. Day 4 – Performance, Monitoring and Troubleshooting

Performance Features

- Review read-ahead options and measure an example where relevant to the selected release.
- Review recovery-thread settings and the limitations of the worked example.
- Review release-specific timing options where applicable; do not assume an old parameter exists in every release.
- Describe the PDQ operations available in the database server
- Describe configuration parameters that control PDQ resources
- Review SQL statement-cache configuration and workload trade-offs.

Session Monitoring

- Use onstat and the sysmaster database to monitor user activity on the database server
- Identify the resources held by a user session
- Inspect sample session SQL while respecting access permissions and sensitive-data handling.
- Identify which user sessions have created temporary tables
- Describe how to identify the type of locks held by user sessions
- Identify a disposable test session and explain termination effects; terminate only that designated lab session if the exercise requires it.
- Explain the purpose of the dynamic lock manager

Event Alarms

- Explain how to use the event alarm feature
- Interpret event severity codes, not "security codes".
- List the classes of alarm events in the server environment
- Explain the purpose of the system alarm program

Administrative SQL and Scheduler

- Perform a small supported administrative SQL task within the training instance.
- Review the built-in scheduler and a prepared bounded maintenance task; no production schedule is created.

Troubleshooting

- Identify and respond to common server problems
- Run the commands to check for consistency of the disk subsystem
- Identify the version and serial number of the database server
- Identify permission problems with devices and server files
- Describe the information required by Technical Support
- Describe the parts of a database version number
- Diagnose a prepared permission/storage-location problem, propose a justified correction and validate it in the disposable instance.

5. Optional Consolidation Case — Time Permitting

Optional Configuration Case

- Set required server environment variables
- Create an sqlhosts file entry to enable a connection
- Initialise only the explicitly disposable training instance after confirming baseline and restore materials.
- Review reinitialisation consequences; perform only designated lab changes using supported utilities or administrative SQL.
- Create dbspaces
- Create and delete logical logs
- Move the physical log