

IBM Information Management

IBM Informix[®] 12.10

Enhancements in Storage Space Management

A Technical White Paper



Contents

- Executive Summary** 3
- Primary Storage Manager** 4
 - Features* 4
 - Design* 5
 - Benefits*..... 7
 - Summary* 7
- TimeSeries Rolling Window**..... 8
 - Features, Design, and Benefits*..... 8
- Purging Tables by Time and/or Size** 10
 - Features, Design, and Benefits*..... 10
- Summary** 12

Executive Summary

The IBM Informix® database server has a long tradition of being easy to maintain. The new Informix Primary Storage Manager (Informix PSM) continues this tradition by making it easier to create and maintain backups of critical data. It is simpler to use and more powerful than its predecessor, the Informix Storage Manager (ISM), and more fully featured than the ontape utility. Its feature set is useful in the general market, but is especially applicable to the embedded market, where low maintenance overhead is crucial.

Two smaller features also make the Informix server easier to manage: the ability to manage time series data in rolling window containers and autonomic storage management for rolling window tables. Although similar in name and concept, these are distinct and separate features. Both enable users to implement business rules to automatically remove outdated information from their system for purposes of performance and storage requirements.

Primary Storage Manager

The Informix Primary Storage Manager (Informix PSM) is intended to relieve the shortcomings of the Informix Storage Manager (ISM). Most customers prefer the simplicity of ontape, and especially appreciate the ease with which an ontape backup can be transported to another system and restored to create a clone of the original. However, they find it is not full-featured enough for incremental and partial backups. Customers who use the ON-Bar utility with ISM as the storage manager find it too cumbersome to configure and manage; Informix PSM is designed to be easy for you to configure and manage. Additionally, ISM has poor performance and is limited to four storage devices. The Informix PSM extends storage to multiple devices for backup and recovery, provides significantly upgraded performance, and is easier to configure.

Features

Informix PSM is designed to meet your needs by providing a simple storage manager that can back data up to file systems and support all the features of ON-Bar, like automatic log backup, unattended database backup, parallel backup, backup generation management with pruning, and others. You can perform a cold restore onto a computer that is not identical to the one on which the backup was performed by giving any chunk a new path name and offset during the restore. Informix PSM can also automatically compress and/or encrypt backup data using off-the-shelf algorithms.

The one feature that ISM has that Informix PSM lacks is the ability to support tape devices.

Feature	Ontape	ON-Bar with ISM	ON-Bar with Informix PSM
Back up all database server data	Yes	Yes	Yes
Back up logical-log files	Yes	Yes	Yes
Back up while the database server is online	Yes	Yes	Yes
Restore data to a specific point in time	No	Yes	Yes
Back up and restore different storage spaces in parallel	No	4 devices maximum	Yes
Works independently of an additional server process	Yes	No	Yes

Table 1: Feature comparisons between Ontape, ISM, and Informix PSM

Informix PSM lets you perform the following tasks:

- Create, modify, and delete storage devices
- Define and modify maximum sizes for devices
- Group devices into separate logical collections for log backups, space backups, and exporting/importing to/from the outside world
- Transfer backup information from one device to another within a pool
- Determine whether volumes, objects, and devices are locked or busy
- Release locked volumes, objects, and devices
- Verify volume names and labels
- Enable objects to be compressed and/or encrypted
- Expire and remove out of date backup information as directed by **onsmsync**

Design

Informix PSM is made up of three components: the **onpsm** utility, the XBSA shared library, and the Informix PSM catalogs (ISAM files). These components provide the essentials of secondary storage management. Informix PSM keeps track of all objects, devices, and pools in a set of ISAM tables called the Informix PSM catalog. Informix PSM provides command-line services to set up the Informix PSM catalog. Whenever ON-Bar requests retrieval or storage of data, Informix PSM coordinates the execution of the request at the device level.

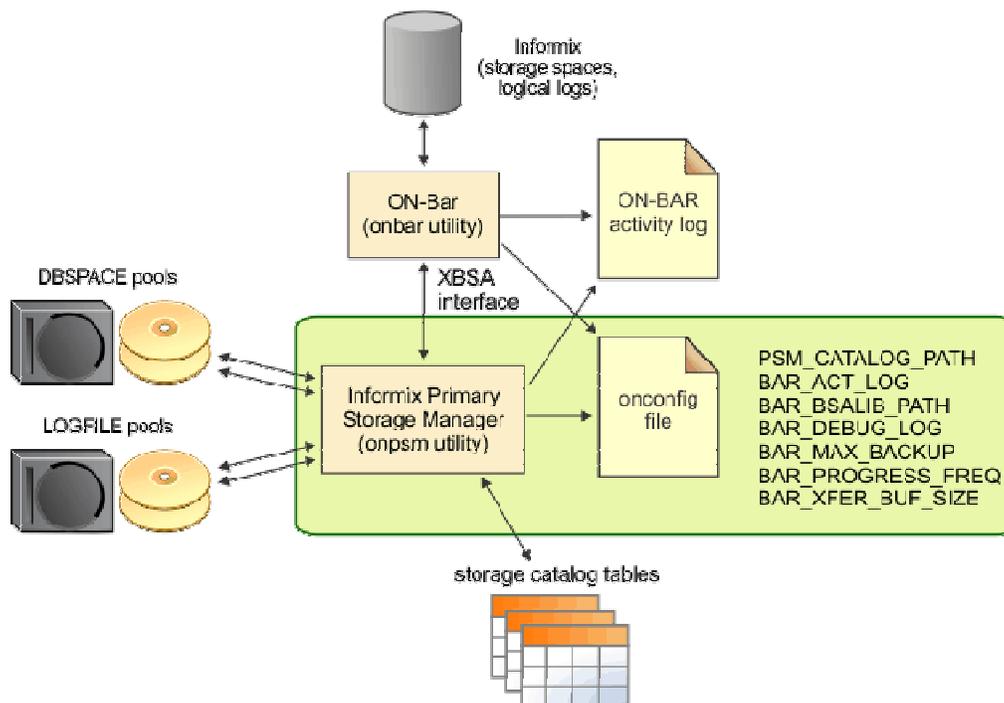


Figure 1: Components of Informix Primary Storage Manager

A pool is a named group of devices that you define in Informix PSM as a repository for a particular class of objects. Pools make it possible for Informix PSM to select from multiple devices when it services a request, and to move automatically from one device to another when one becomes full or if a device fails.

Three pools are always available; they cannot be deleted as they are system pools:

- DBSPool – Holds online backups of dbspaces, blobspaces, and smart blobspaces.
- LOGPool – Holds online backups of logical logs.
- EXTPool – Serves as a staging area for exporting a backup set to a single large logical object or importing a backup set back in from such an object. This lets you back up on one machine and restore onto a different machine, as ontape allows. Files in the EXTERNAL pool are “offline” (not visible to ON-Bar and not tracked by Informix PSM).

Pools are maintained by adding devices that are used for storing Informix backup data. You can add or delete devices from pools.

Informix PSM places no limit on the number of devices that can be defined in a pool. If there is more than one device defined in a pool, then a backup object can be split across those devices. When writing an object, if the target device runs out of space, Informix PSM will automatically continue the object on the next available device with an incremented <piece_number> suffix for the object filename.

Informix PSM is designed for a multiple server environment. The DBSERVERNAME differentiates between servers and allows Informix PSM to service multiple servers simultaneously. Informix PSM can be configured once for all servers wishing to use it, as long as those servers all use the same Informix PSM catalog location (PSM_CATALOG_PATH configuration parameter). If the servers share an \$INFORMIXDIR, by default they will all share the same Informix PSM catalog location, as this defaults to a fixed location under \$INFORMIXDIR.

Informix PSM supports parallelism against file devices limited only by the hardware in use. When a device is needed to write an object, the devices within a pool are tested for availability in the order they were added to the pool. File devices only become unavailable for writes once they are full.

When performing a parallel backup with ON-Bar, the maximum number of backup objects that can be written in parallel is controlled by the ON-Bar parameter `BAR_MAX_BACKUP` in the `onconfig` file. There is not much benefit to be gained by setting this larger than the number of physical devices.

Benefits

The larger buffers and the ability to write to any number of devices at the same time relieve performance bottlenecks that hinder ISM. While ISM is limited to four devices, Informix PSM can simultaneously access as many devices as you need for backup and recovery. The larger buffers should lead to performance enhancements for both reads and writes on each device. Informix PSM is also different from ISM in that it requires little configuration and requires no daemons.

Major Differences with ISM

- Larger internal buffers
- No limit on number of storage devices
- No limit on the number of devices that can be written in parallel
- No need for a background process

Summary

The new Informix Primary Storage Manager represents an improvement in design that enables it to provide the same or better set of features that the Informix Storage Manager currently provides, with better performance, and it requires less administrative overhead.

More information on using this feature can be found in the documentation, under the section, *Informix Primary Storage Manager*.

TimeSeries Rolling Window

Time series data tends to accumulate in massive amounts. This accumulation can place extensive demands for storage space, query operations, and data analysis on your database system. Depending on business requirements, it may not be necessary to keep detailed measurement records beyond a certain age. One way of removing the old information is to delete it, record by record, in a traditional manner. A faster, more efficient way to remove the old information is in blocks of data at a time. This is made possible by storing the data in time series rolling window containers. These containers allow entire blocks of data to be removed in one operation.

Features, Design, and Benefits

You can control the amount of time series data that is stored in containers by specifying criteria on when to delete obsolete data. This is done by creating a rolling window container that has multiple partitions that can optionally be stored in multiple dbspaces. A rolling window container can be configured to define the time interval for each partition and how many partitions are allowed. For example, 12 partitions that each store a month of data can be specified. After data is inserted for a new month, a new partition is created, and if the number of partitions exceeds the maximum that is allowed, the oldest partition becomes dormant. You specify when to destroy dormant partitions. This process of expiring information can be fully automated.

The current container maintains a single partition. The ability to use a window container gives you the option to define a container over multiple partitions. Instead of a single partition, the window container is essentially a list of interval partitions ordered over time by using a regular time series value of *tsPartitionDesc_t* values. The interval of the designated Calendar used in this time series controls the datetime range of each partition. You can define window containers with a list of dbspaces that the Informix server will use in a round-robin allocation scheme when adding new datetime range partitions. A window container may limit the number of partitions in an active window so that, as new partitions are added, the older partitions become dormant. Dormant partitions are placed in a dormant area and are no longer considered an active part of the time series. Dormant partitions may be automatically or manually purged.

Each partition is known as an interval partition. There is an active set of partitions and a dormant set of partitions. Active partitions may be detached to become dormant partitions, and

dormant partitions may be attached to become active partitions again. You can delete dormant partitions by deleting the physical partition without deleting each element stored in that partition. This enables purge operations to be very quick. This feature also provides an active window size to control the number of active partitions and a dormant window size to control the number of dormant partitions. When the number of active partitions exceeds the active window size, they will be automatically detached to become dormant. If the active window size is 0, then active partitions are not automatically detached. Similarly, when the number of dormant partitions is greater than the dormant window size, they are automatically dropped. If the dormant window size is 0, then partitions are not automatically deleted.

Partitions are attached and detached in chronological order. There is no provision to leave gaps in the sequence of active partitions; otherwise, incorrect values would be returned especially in the case of irregular time series values. User-defined routines are provided to manage the active and dormant partitions and to generate usage statistics for the window containers.

Since the purge operation deletes a partition as a unit of work, it will be very efficient. Operations typically access elements in a time range, so organizing the partitions in a time series also allows operations to quickly locate the partitions with elements of interest.

More information on using this feature can be found in the documentation, under the section, *Manage time series data in rolling window containers*.

For tables that use a RANGE INTERVAL distributed storage strategy, the database server already creates new fragments automatically as needed, when an inserted record has a fragment key value outside the range of any existing fragment. You can now use new SQL syntax to define a purge policy for archiving or deleting interval fragments based on either a user-specified limit on its allocated storage space size, or on the number of its interval fragments. A new built-in Scheduler task runs periodically to enforce the purge policies on qualifying rolling window tables.

In order to avoid accidental deletion of data, besides dropping fragments outright, the feature offers the option to detach fragments without deleting them. Aside from providing a lifeline to recover from incorrectly specified purge policies, this implementation allows for more complex purging activities, whereby purged fragments can be attached (or the data otherwise moved) to archival tables.

For applications that deal with large amounts of information that tends to expire after some time, being able to configure the automated removal of information reduces the administrative overhead associated with the Informix server.

More information on using this feature can be found in the documentation, under the section, *Autonomic storage management for rolling window tables*.

Summary

Informix 12.10 includes three new features related to the space and storage management that make it easier to maintain, and can also produce substantial performance benefits. The Informix Primary Storage Manager (PSM) opens up the implementation of parallelism fully utilizing available hardware in backup and recovery tasks, and removes the need to configure and run a separate server process in order to automate the management of physical storage devices. The time series rolling windows allow for the automatic removal of old information from the active set, and facilitates the archiving of the information, while also delivering large performance gains for doing the same. The rolling windows feature completes the automation of space management of tables partitioned by interval by allowing for data in old interval partitions to be removed from active use.

For more information

To learn more about the Informix features, contact your IBM representative or IBM Business Partner, or visit ibm.com/software/data/informix



<p>IBM Informix Storage Optimization A Technical White Paper</p> <p>March 2013 Chris Golledge</p> <p>© Copyright 2013 IBM Corporation</p> <p>IBM Corporation Software Group Route 100 Somers, NY 10589 U.S.A.</p> <p><i>The information contained in this publication is provided for informational purposes only. While efforts were made to verify the completeness and accuracy of the information contained in this publication, it is provided AS IS without warranty of any kind, express or implied. In addition, this information is based on IBM's current product plans and strategy, which are subject to change by IBM without notice.</i></p>	<p><i>IBM shall not be responsible for any damages arising out of the use of, or otherwise related to, this publication or any other materials. Nothing contained in this publication is intended to, nor shall have the effect of, creating any warranties or representations from IBM or its suppliers or licensors, or altering the terms and conditions of the applicable license agreement governing the use of IBM software.</i></p> <p><i>References in this publication to IBM products, programs, or services do not imply that they will be available in all countries in which IBM operates. Product release dates and/or capabilities referenced in this presentation may change at any time at IBM's sole discretion based on market opportunities or other factors, and are not intended to be a commitment to future product or feature availability in any way. Nothing contained in these materials is intended to, nor shall have the effect of, stating or implying that any activities undertaken by you will result in any specific sales, revenue growth, savings or other results.</i></p> <p><i>IBM, the IBM logo, ibm.com, and Informix are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide.</i></p>
--	--