


Syncsort™ PipeSort



What is Syncsort™ PipeSort

Syncsort™ PipeSort is a key component of the Syncsort™ MFX Suite and offers the fastest way of sorting the same file many times with different key sequences. It can simultaneously execute up to eight differently sequenced sorts from a single pass, without having to enter data multiple times. Using cutting-edge parallel sorting technology, Syncsort™ PipeSort can cut the total elapsed time by more than 50% compared to running separate sorts, while also significantly reducing the I/O overhead.

As a licensable feature of Syncsort™ MFX, Syncsort™ PipeSort is the perfect utility for helping customers to meet batch windows' constraints with especially large files and multiple sort key sequences. It can be a powerful solution to help save time, computing resources and mainframe costs.

Elapsed time savings

Syncsort™ PipeSort exploits the rich multi-programming potential of MVS, OS/390, and z/OS to cut total elapsed time by replacing serial with parallel sort execution.

Technical specifications


- Do you routinely run applications that sort the same large data set in different ways?
- Would it be a significant benefit if you could cut the total elapsed time for those sorts by 50 percent and more?

If your answer is yes to both questions, PipeSort may be the tool you need.


Benchmark tests on a one-gigabyte DASD file sorted in three different sequences compared consecutive executions versus Syncsort™ PipeSort. The consecutive executions required 1 hr 9 minutes total elapsed time. Syncsort™ PipeSort needed only 31 minutes, a savings of 55%. Even greater savings can be achieved when Syncsort™ PipeSort replaces more individual sorts.

Syncsort™ PipeSort reads the input file only once and distributes the input records to multiple simultaneous Syncsort™ executions. This method enables Syncsort™ PipeSort to eliminate the considerable I/O overhead involved in reading the same input data for every sort.

Consecutive Sorts



Syncsort™ PipeSort



Resource Requirement

You must have adequate resources to run Syncsort™ PipeSort successfully. Syncsort™ PipeSort requires all the resources (DASD space and central storage) that would have been spread over the individual sort executions. For example, if Syncsort™ PipeSort combines five sorts that each require 50 cylinders of DASD, Syncsort™ PipeSort will need 250 cylinders. Also, CPU time may increase compared to the sum of CPU times for individual sorts.

Compatibility and Implementation

Syncsort™ PipeSort requires Syncsort™ and is easily installed. Syncsort™ MFX control statements in existing applications can be used without change.

Implementation is through simple JCL that includes DD statements specific to PipeSort; for example, SRTnCNTL DD statements replace SYSIN DD statements. The following example shows Syncsort™ PipeSort JCL for two sorts:

```
//SORT      EXEC PGM=PIPESORT,
//          REGION=3072K
//SORTIN    DD DSN=input.file,....
//SYSOUT    DD SYSOUT=*
//*
//SRT1CNTL DD * sort #1 cntl statements
           SORT FIELDS=(.....),DYNALLOC=SYSDA
//SYSOUT1   DD SYSOUT=* sort #1 SYSOUT
//SRT1OUT   DD DSN=output.file.#1,
//          DCB=RECFM=recfm,....
//*
//SRT2CNTL DD * sort #2 cntl statements
           SORT FIELDS=(.....),DYNALLOC=SYSDA
//SYSOUT2   DD SYSOUT=* sort #2 SYSOUT
//SRT2OUT   DD DSN=output.file.#2,
//          DCB=RECFM=recfm,....
//
```

Resource Requirement

Syncsort™ PipeSort may be the tool you need to expedite time-value applications such as billing, inventory updates, or regulatory reporting. Whatever your applications, if you frequently run large sorts on the same input data and can benefit from reducing their total elapsed time, Syncsort™ PipeSort can make a significant contribution to your system throughput.

To learn more about Syncsort™ PipeSort, contact Precisely at 877- 700-0970