



Frequently Asked Questions about RTD

NOTE: All references to RTD apply to both RTD/zOS and RTD/DB2.

- What is the RTD Suite?
- What functions does RTD perform?
- How does RTD work?
- Why is RTD better than DFDSS?
- When should you consider implementing RTD?
- What types of datasets can RTD process?
- Does RTD process multivolume datasets?
- What does RTD do with database volumes?
- Can RTD process DB2 datasets?
- Is RTD effective for volumes with hardware defragmentation?
- How much CPU time does RTD use?
- How is RTD controlled?
- How can one measure RTD's success?
- What improvements can be expected from RTD?
- How often does RTD process a volume?
- How does RTD handle datasets which are allocated?
- How much manpower is required to install and maintain RTD?
- Does RTD support DFSMS?
- Can RTD handle large volumes of DASD?
- Does RTD reduce the Fragmentation Index to zero?
- How can RTD be tuned?
- Where should RTD run in a multi-CPU environment?
- Are there special considerations for RTD under Parallel Sysplex?
- Does RTD support FlashCopy, PPRC and Preserve Mirror?



What is the RTD Suite?

The RTD Suite currently includes the following products:

- RealTime Defrag for z/OS (RTD/zOS)
RTD/zOS performs space management for non-DB2 DASD volumes on a z/OS system.
- RealTime Defrag for DB2 (RTD/DB2)
RTD/DB2 performs space management for DB2 DASD volumes on a z/OS system.

What functions does RTD perform?

RTD performs the following three functions:

- DEFRAG - Defragmentation of DASD volumes.
- RELEASE - Release dataset allocated but unused space.
- COMBINE - Combine multiple dataset extents into one extent.

How does RTD work?

RTD runs in the background as a Started Task (STC) working 24 hours a day, 7 days a week. The RTD Administrator invokes ISPF Panels to define the cycles and volumes to process for RTD. During each cycle, RTD processes a portion of the datasets on a volume so that only a minimal amount of resources (CPU and Channel Activity) is utilized. After a number of iterations, RTD will have groomed the volumes into an optimal state.

In contrast to other Defrag utilities, RTD's goal is not to achieve a Fragmentation Index value of 0. Since any allocation must be satisfied in 5 extents (or less), it is RTD's goal to increase the size of the largest 5 free extents. This provides the best environment for dataset allocations.

Why is RTD better than DFDSS?

DFDSS must allocate the entire DASD Volume for the full amount of time required to perform a DEFRAG. RTD only allocates those datasets which it must move and only for the time required to move the dataset. Datasets which are in use will be bypassed for the current interval and processed during a subsequent interval. Since RTD runs continuously, the volumes are always in an optimal state. With DFDSS the volumes are optimal once a day, once a week or less. RTD does not require a Batch window to perform its functions. It works constantly keeping your DASD environment in an optimal, stable state.

When should you consider implementing RTD?

You should consider RTD if:

1. The Batch window is too small to process all the Defrag jobs.
2. DASD volumes are underutilized.
3. Jobs are terminated due to "out-of-space" conditions.

What types of datasets can RTD process?

RTD processes all types of datasets (non-VSAM and VSAM).

Does RTD process multivolume datasets?

RTD processes multivolume datasets as follows:

- DEFRAG – extents on all volumes are processed.
- COMBINE – extents on all volume are processed by default. The RTD Administrator can specify that only extents on the last volume of a multivolume dataset be processed.
- RELEASE - multivolume datasets are not processed.

What does RTD do with database volumes?

RTD processes database volumes the same as other volumes. However, one must realize that RTD will not process allocated datasets. Hence, individual databases will only be processed when they are not allocated.

Can RTD process DB2 datasets?

RTD/DB2 is specifically designed to process DB2 datasets.

RTD/DB2 can process DB2 datasets which are currently not allocated by DB2 as well as DB datasets that are allocated by DB2 but not currently in use by a transaction.

Is RTD effective for volumes with hardware defragmentation?

Yes. Hardware defragmentation reduces space wastage, but only physically. The restrictions imposed by MVS logic (allocation routines) and control blocks (e.g., the VTOC) still exist for such volumes. RTD alleviates these restrictions and allows for better utilization of all DASD, including those with hardware defragmentation.

How much CPU time does RTD use?

RTD was designed to minimize resource utilization. RTD user experience has shown that RTD's CPU utilization is hardly measurable. This is due to the fact that RTD does not attempt to process the entire DASD volume each interval. RTD's intelligent processing achieves maximum results with minimum resource utilization.

How is RTD controlled?

Control parameters for RTD are entered via a user-friendly ISPF interface. The RTD Administrator defines which volumes (or SMS Storage Groups) will be processed, which functions will be performed, and the time interval for performing these functions. The control parameters may be changed at any time and become effective immediately.



How can one measure RTD's success?

RTD users tell us that RTD pays for itself within 6 to 12 months. Savings result from:

1. Increased DASD space utilization. Prior to RTD this is about 70-75%; with RTD this increases to 80-95%. Purchasing new DASD can be delayed.
2. Reduced number of Job abends due to space-not-available.
3. Increased stability of the production environment.
4. Reduced resources needed to perform Storage Management.

What improvements can be expected from RTD?

RTD users see DASD space utilization increase from 70-75% to 80-95%. The Fragmentation Index is reduced by 200-500 points.

How often does RTD process a volume?

The RTD processing interval can be defined on a storage group and/or volume basis and function-by-function basis. Volume processing intervals can be set from 1 to 1440 minutes (24 hours).

How does RTD handle datasets which are allocated?

- RTD/zOS uses standard MVS ENQs to allocate datasets which it would like to process. If the RTD/zOS ENQ request fails, RTD/zOS will bypass the dataset. During the next processing interval RTD/zOS may again attempt to allocate the dataset. If it is known that certain datasets are always allocated, the user can, if desired, easily exclude these from RTD/zOS processing
- RTD/DB2 uses DB2 services to obtain exclusive serialization of datasets it would like to process. If the RTD/DB2 serialization request fails, RTD/DB2 will bypass the dataset. During the next processing interval RTD/DB2 may again attempt to allocate the dataset. If it is known that certain datasets are always in use by DB2 transactions, the user can, if desired, easily exclude these from RTD/DB2 processing.

How much manpower is required to install and maintain RTD?

Installation takes from 30 to 60 minutes. Defining the global parameters and the first volumes for RTD processing takes about 15 minutes. Other volumes are then added as required. Since RTD supports easy-to-use masking facilities, a few control parameters are sufficient for the entire DASD configuration. Once the parameters have been set, RTD users look at the RTDLOG every few weeks just to see how RTD is doing. That's all there is to it.

Does RTD support DFSMS?

RTD processes both SMS and non-SMS managed volumes. RTD works independently of SMS. That is, RTD processing depends only on the parameters defined in the RTD control dataset. Additionally, SMS class names may be used for defining selection criteria to RTD.

Can RTD handle large volumes of DASD?

Yes. Many RTD users have large DASD configurations. RTD typically processes thousands of DASD volumes, which may be of any size including model 54s and EAVs. Since RTD supports multi-tasking and requires a minimum of resources, up to 28 volumes may be processed concurrently.

Does RTD reduce the Fragmentation Index to zero?

No, this is not necessary. The critical factor for dataset allocation is the size of the 5 largest extents. All allocations (be they primary or secondary) must be contained in a maximum of 5 extents on a single volume. If this is not possible, the allocation (and the job) will fail. RTD's goal is to ensure that the 5 largest extents are as large as possible. RTD reduces the possibility of job abends due to space not available.

How can RTD be tuned?

RTD was designed to use a minimum of system resources. Hence, there is very little to do in the way of tuning. However, knowledge of the environment can be used to further reduce RTD resource utilization. For example, if it is known that certain volumes contain many datasets that are allocated for a certain time frame, one can tell RTD to only process these volumes outside of this time frame.

Where should RTD run in a multi-CPU environment?

If the all of the DASD volumes are shared (or accessible from a particular CPU) then RTD should run on only one system which has access to all of the volumes. If the volumes are not shared, RTD must run on each CPU with exclusive access to the volumes.

Are there special considerations for RTD under Parallel Sysplex?

No. Since RTD only runs on one system, there are no special considerations for a Parallel Sysplex environment.

Does RTD support FlashCopy, PPRC and Preserve Mirror?

RTD provides complete support for FlashCopy, PPRC and Preserve Mirror (Remote Pair FlashCopy®) volumes. A comprehensive array of parameters allows the selection of and processing options for volumes based on their current status in such environments.

INTERCHIP AG

Elektrastrasse 6
D-81925 München
Telefon +49 - 89 - 99 14 99 0

Email: info@interchip.de
<http://www.interchip.de>