Double-Take Pagefile Configuration

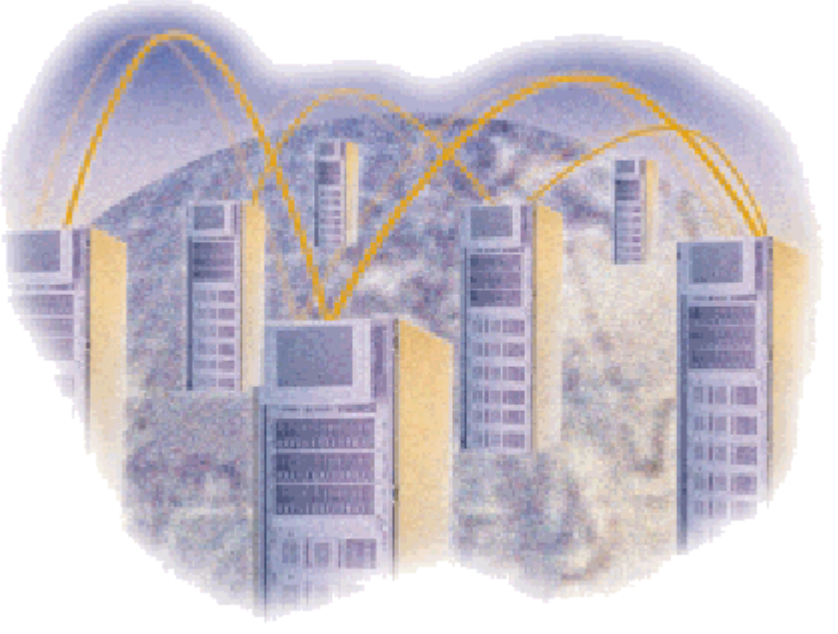

Double-Take Pagefile Configuration published August 2002

Table of Contents

Introduction

Double-Take is an asynchronous data replication application that utilizes a number of queues as it mirrors and replicates data on the source to the target. Most of the queues are transparent to system administrators and are optimized for the best performance. However, there are a number of settings related to the Double-Take pagefile that can impact the amount of memory Double-Take uses for queuing.

In order to provide maximum Double-Take queuing capacity and system performance, Double-Take servers should be configured with enough RAM to accommodate the maximum Double-Take pagefile (1GB) in addition to the server's other memory needs. When this amount of physical memory is not available, understanding of the following items is necessary to determine the appropriate configuration for the Double-Take pagefile on a server.

1. ***What is the Double-Take pagefile?***
2. ***What are the functions and conditions that impact Double-Take pagefile usage?***
3. ***What factors affect the Double-Take pagefile configuration?***
4. ***How does paging memory to disk impact performance?***
5. ***What does all of this mean to me?***

The following sections will discuss these points in detail and provide system administrators with the knowledge necessary to size the Double-Take pagefile for optimal functionality and performance.

1. What is the Double-Take pagefile?

When the Double-Take service starts, it reserves a pool of user-addressable memory equal to the Double-Take pagefile size. This reserved pool of memory is the Double-Take pagefile. Although Double-Take has the pool of memory reserved, it only uses what is necessary at any given time. The Double-Take pagefile plays an important role in the mirroring and replication processes, serving as a staging area through which mirror and replication operations pass on both the source and target systems.

When the Double-Take service starts, it creates or initializes a file, dtpagefile.sys, that is the size specified as the Double-Take pagefile size. Dtpagefile.sys is a Windows swap file, reserved for disk swapping operations from the Double-Take pagefile, and guarantees that Double-Take will always have access to the full amount of memory specified as the Double-Take pagefile size.

2. What are the functions and conditions that impact Double-Take pagefile usage?

The Double-Take pagefile is used to queue mirror and replication operations on both the source and target systems. There are a number of factors that determine how many operations will be queued at any one time, including, but not limited to, the following:

- Mirror operations being sent/received
- Replication operations being sent/received
- Double-Take pagefile configuration
- Unfinished operations retrying on the target
- Performance of network and system components

It is important to understand how the Double-Take pagefile is used during mirroring and replication on both the source and target. This is the key to anticipating how much of the Double-Take pagefile will be in use on your systems in various conditions.

2.1 How is the source pagefile used during mirroring?

During a mirror, operations are created and queued in the source's Double-Take pagefile. The queued operations contain pointers to the data and do not contain file data, so they are quite small (less than 600 bytes). When a mirror operation is processed, data is read from the disk and sent to the target without queuing the data in the source Double-Take pagefile.

The maximum number of operations that will be queued in the source's Double-Take pagefile is controlled by the Mirror Queue Limit. In the Management Console, this setting is called Maximum Pending Mirror Operations and can be found on the Server Properties Source tab. By default, this value is set to 1000 which results in less than .6 MB of the source pagefile being in use during a mirror. Increasing or decreasing this value is not likely to improve mirror performance.

2.2 How is the target pagefile used during mirroring?

When the target receives mirror operations, they are first queued in the Double-Take pagefile and then committed to disk. The maximum number of mirror operations that will be queued on the target is determined by the TG Mirror Capacity High setting. In the Management Console, this setting is called Target Mirror Capacity High and can be found on the Server Properties Target tab.

When the number of mirror operations queued on the target reaches this threshold, the target will pause mirroring from the source. While mirroring is paused, the target continues to commit operations to disk and the number of mirror operations queued in the target Double-Take pagefile decreases accordingly. When the number of queued operations reaches another threshold, TG Mirror Capacity Low, mirroring is resumed. In

the Management Console, this setting is called Target Mirror Capacity Low and can also be found on the Server Properties Target tab.

NOTE: These settings were renamed in version 4.2 to more accurately reflect their functions and changed to percentages so that they will be automatically adjusted when the Double-Take pagefile size is changed.

If you are using Double-Take version 4.1 or earlier, these options are called TG Open File Count High and TG Open File Count Low.

When Double-Take reaches the TG Open File Count High threshold, the entire connection from the source is paused, not just mirroring as it is in version 4.2 and later.

In many cases, mirror operations will accumulate in the target Double-Take pagefile because it receives operations faster than they can be committed to disk. Pause/resume cycles during mirroring are common and do not indicate a problem. When pause/resume cycles are occurring, the Transmit Mode statistic displayed in the right pane of the Management Console will cycle between Paused and Started. The pause/resume activity will also be reported in the Double-Take log file.

The primary goal when adjusting these values is to decrease the amount of memory used by the target during a mirror. Lowering the values will increase the number of pause/resume cycles while raising the values will decrease them. The overhead associated with pausing and resuming the connection is slight enough that there are no significant performance benefits to be gained by adjusting the values.

The amount of the pagefile a target will use for queuing during a mirror is calculated by [Pagefile size in MB] * [TG Mirror Capacity High]. By default, TG Mirror Capacity High is set to 20. Since the setting is a percentage, adjustments occur dynamically when the Double-Take pagefile is resized.

NOTE: If you are using Double-Take 4.1 or earlier, the amount of the pagefile a target will use for queuing during a mirror is calculated by [TG Open File Count High] * [Mirror Chunk Size] / 1048576.

The Target Open File Count High and Target Open File Count Low settings are not automatically changed when the Double-Take pagefile size is changed. Therefore, if the Double-Take pagefile size is changed, these settings should be manually adjusted using the following formulas.

Target Open File High = [Double-Take pagefile size in kilobytes] / 64
Target Open File Low = [Target Open File High] * .75

Mirroring will be paused when either the TG Mirror Capacity High threshold is reached or when the Drop Threshold is met. The Drop Threshold, discussed in the following sections relating to Double-Take pagefile usage during replication, is intended as a threshold for replication queuing. However, the connection will be paused if the Drop Threshold is reached, regardless of the type of operations queued in the Double-Take pagefile.

2.3 How is the source pagefile used during replication?

When a change to the file system results in the creation of a replication operation on the source system, Double-Take queues the operation in the source Double-Take pagefile and then transmits it to the target. During periods of normal replication load with adequate bandwidth, replication operations are typically removed from the source Double-Take pagefile without being swapped to disk and transmitted to the target immediately.

During high usage, replication operations may be created faster than they can be sent to the target and may accumulate in the source pagefile. Concurrent mirrors, lack of adequate network bandwidth to transmit the replication operations immediately, and paused connections are also conditions that will cause replication operations to accumulate in the source pagefile.

When the amount of Double-Take pagefile in use on the source reaches the Drop Threshold, the connection will automatically be disconnected. This is called an auto-disconnect. When an auto-disconnect occurs, the source will immediately flush all operations from its pagefile and attempt to reconnect and and remirror (a difference mirror by default). In the Management Console, these settings, Drop Threshold and the remirror, are Auto-Disconnect Threshold and Perform Remirror after Auto-reconnect, respectively. They are located on the Server Properties Memory and Startup tabs. Along with enabling or disabling the remirror, you can specify the type of remirror, full or partial, to be performed.

The HP Queue Ratio is the ratio of replication packets to one mirror packet. This setting is on the Server Properties Source tab and provides control over the priority that queued replication will receive over mirror packets. When replication and mirroring are occurring simultaneously, replication operations may remain in the Double-Take pagefile for longer periods of time, particularly if pause/resume cycles are occurring. The number of queued replication operations can be monitored in the Management Console and if they continually increase during a mirror, it may be necessary to increase the HP Queue Ratio from the default of 5. Usually, this setting is adjusted when replication is traversing a low bandwidth link and/or a heavy load of replication is occurring during a mirror.

2.4 How is the target pagefile used during replication?

The Double-Take target processes operations as quickly as possible in order to commit them from memory to disk. However, if the target receives replication operations faster than it can commit them to disk, operations will begin to accumulate in the target pagefile. If the target pagefile utilization reaches the Drop Threshold, the target will pause all connections established to it. The target will resume the connections when the pagefile utilization drops to 50% of the Drop Threshold.

NOTE: Both the source and target modules of Double-Take use the Drop Threshold. A Double-Take source will auto-disconnect all of its source connections when the threshold is reached, while a Double-Take target will pause all inbound connections when the threshold is met.

2.5 How does the target handle unfinished operations?

On Double-Take targets, unfinished operations occur when Double-Take cannot commit an operation to disk. For example, sharing violation (if another application, like a backup, has the target file open with a lock) and disk full conditions are common causes of unfinished operations. Double-Take will retry the operation at three second intervals (by default) and will not commit any more operations until the one being retried is committed to disk.

During the time retries are occurring, the target will continue to accept mirror and replication operations from the source and queue them in its target pagefile. Accordingly, when unfinished operations are occurring, target Double-Take pagefile usage will grow. If the target reaches the Drop Threshold, it will pause any connections. When the retry condition is resolved, the target will commit operations in the pagefile to disk. When the pagefile utilization reaches 50% of the Drop Threshold, Double-Take will resume the connections.

2.6 How does the performance of network and system components impact the Double-Take pagefile?

The speed of the network can have a significant impact on how much of the Double-Take pagefile is utilized. If the source and target are communicating and a file is copied into the replication set on the source, many of the resulting replication operations will remain queued in the source pagefile for a period of time as the source transmits the data. If the servers are connected via a gigabit link, the operations will, of course, remain in the source pagefile for much less time.

The network speed will also affect how many operations are queued on the target system under various conditions. A slow network link will serve to throttle the delivery of operations, which will then be less likely to accumulate on the target since it is able to commit them to disk faster than they arrive.

Various system components can have an impact on how quickly replication and mirror operations move into and out of the Double-Take pagefile. For example, a target system writing data to a RAID 0 array will be able to commit operations to disk faster than if it were writing to a RAID 5 array.

Although most system components cannot be configured to improve Double-Take throughput, the target's disk subsystem configuration can be optimized so that it can move operations from memory to disk as quickly as possible. The RAID level, the amount of cache available for write operations, and the write policy can all affect the target's throughput. Often times, a RAID 0 array configured with a write-back write policy and 100% cache dedicated to write operations will provide the best write performance.

NOTE:	Keep in mind that 100% cache dedicated to write operations may not be acceptable if the machine is actively being used for other uses. For example, if the target is also acting as a source in a chained Double-Take configuration or if the target is monitoring the source for failover and may stand in if the source fails. In these cases, performance on the target will be negatively impacted by the lack of memory available for reading.

3. What factors affect the Double-Take pagefile configuration?

Since Double-Take queuing requires the use of system memory, the amount of system memory available and the priority of Double-Take queuing in relation to the other memory needs of the system should be considered when sizing the Double-Take pagefile. In some scenarios, such as replicating a large database over a slow WAN link, queuing capacity on the source is critical to reduce the occurrence of auto-disconnects. If the WAN link goes down, or a large amount of data is changed in a short period of time, the source's pagefile must be large enough to buffer the replication data until it can be transmitted to the target in order to avoid an auto-disconnect.

The appropriate Double-Take pagefile size relies primarily on four factors:

- **Time to remirror**—Some data, extremely large database files in particular, may have an opportunity to be in an invalid state on the target while a mirror is in progress. This is because a mirror moves files in 32K chunks (by default), and it may take several minutes or hours for a the database file to be mirrored. During that time, the database file and/or log on the target may be inconsistent and attempts to mount the database before a mirror is complete may be unsuccessful. Therefore, Double-Take queuing capacity is critical to minimizing auto-disconnects and the resulting remirrors.

- **Network and system resources used to remirror**—The use of network bandwidth and system resources to remirror is another factor when determining the Double-Take pagefile size. These factors become more important in WAN environments where limited bandwidth is available or when the source and/or target normally operate at a heavy load. Since mirrors require additional bandwidth, a larger Double-Take pagefile can potentially conserve bandwidth by reducing the number of auto-disconnects and remirrors.

- **Importance of data**—The relative importance of replicated data versus the performance of the source server should be considered when determining the Double-Take queueing capacity. A company's invoicing data is probably much more critical than a database for a social activities website on an intranet web server, and the amount of Double-Take queuing capacity should be adjusted to eliminate periods where a potential auto-disconnect could put the target data at risk.

- **Available physical memory and memory being paged to disk**—Relative to other applications, Double-Take pagefile configuration is optimized when both the source and target have a 1 GB Double-Take pagefiles as well as 1 GB of free physical RAM for Double-Take to use. This allows Double-Take to queue the maximum amount of replication without the server performance being affected by paging memory to disk. If the Double-Take pagefile is larger than the amount of free physical memory, Double-Take queuing may result in the system paging memory to disk, which will affect system performance.

Other factors, including service level agreements, hardware redundancy, system stability, and company policies may influence the priority of Double-Take queuing but will not be discussed here.

4. How does paging memory to disk impact performance?

Paging memory to disk impacts both system performance and Double-Take performance. These impacts are detailed separately in the next two sections.

4.1 How does paging memory to disk impact system performance?

The use of disk-based memory allows the operating system to provide more memory to the system than the amount of installed RAM. Keep in mind, since disk access is slower, system performance may be impacted when the contents of disk-based memory must be written to disk. The Windows operating system documentation and Windows Resource Kits are valuable resources to learn more about memory management. See Microsoft Knowledge Base article Q146005 to learn how to determine if excessive paging of memory to disk is occurring, causing a memory bottleneck.

The operating system controls memory-paging operations, and the memory management of the Double-Take pagefile is no exception to this rule. Double-Take interfaces with the operating system to use memory and is not responsible for determining whether the memory in use in the Double-Take pagefile is kept in physical memory or paged to a swap file.

The Double-Take pagefile size can be configured anywhere between 32 MB and 1 GB. The Double-Take installation program prompts for the Double-Take pagefile size, which defaults to a value of 256 MB. Regardless of its size, if the system does not have that amount of free physical memory, Double-Take's use of memory for queuing operations may result in the system paging memory to disk. If a large Double-Take pagefile is required, physical memory can be added to the system to accommodate it without the risk of memory being paged to disk and degrading system performance.

Most server applications, such as Microsoft Exchange Server, include a parameter to control how much memory the application will use. By default, they may be configured to use all available physical memory. In some cases, it may be necessary to configure the application to ensure that it does not use physical memory that you intend to be available for Double-Take queuing. See the application's documentation or contact the software publisher's technical support to find out how to configure its memory usage.

4.2 How does paging memory to disk impact Double-Take performance?

Degraded system performance due to excessive use of disk-based memory on a target server can extend the time it takes Double-Take to move operations from the Double-Take pagefile to disk, resulting in connections being paused for longer periods of time. This can impact Double-Take throughput to such an extent that a smaller target Double-Take pagefile (or at least lower threshold values) would actually be better and result in less queued replication on the source. If physical memory on the target becomes a bottleneck for Double-Take throughput, it can cause more operations to be queued for longer periods of time on the source.

On the source system, the use of disk-based memory will not only affect Double-Take's performance, it will also affect the performance of the applications creating the replication operations. If limited physical memory is available, the priority of Double-Take queuing capacity in relation to the priority of system performance will be a factor in determining the appropriate Double-Take pagefile size.

5. What does all of this mean to me?

In order to provide maximum Double-Take queuing capacity and system performance, Double-Take servers should be configured with enough RAM to accommodate the Double-Take pagefile in addition to other memory needs that are required of the server

If adequate physical memory is not available and cannot be added to the source system, the requirement of Double-Take queuing must be weighed against the requirement of system performance. If the requirement for Double-Take queuing is greater, then the Double-Take pagefile size should be increased or maximized even though increased memory usage for queuing may result in system performance degradation due to memory being paged to disk. If the source's system performance is more critical than Double-Take queuing, then the Double-Take pagefile should be set to slightly less than the amount of free physical memory to ensure that the system rarely pages memory to disk.

In almost all cases, the Double-Take pagefile on the target should be sized to slightly less than the amount of free physical memory. The use of disk-based memory on the target typically results in performance degradation to the extent that the larger Double-Take pagefile is counterproductive if use of it causes the system to page memory to disk.

The following table summarizes the Pagefile related settings. Use this table to help you understand which settings you may need to modify.

Setting	Management Console Name (Server Properties Tab)	Description
Drop Threshold	Auto-Disconnect Threshold (Memory)	Percentage of the Double-Take pagefile memory pool that must be reached to trigger an auto-disconnect on the source ◆ **Default**—75% ◆ **Increasing the value**—Increases the amount of queuing in the Double-Take pagefile ◆ **Decreasing the value**—Decreases the amount of the queuing in the Double-Take pagefile ◆ **Reason to modify**—Change this setting if you want to change the amount of memory available to Double-Take for queuing operations without stopping and starting the Double-Take service.
HP Queue Ratio	Replication Packets to 1 Mirror Packet (Source)	Ratio of replication packets sent for every one mirror packet ◆ **Default**—5 ◆ **Increasing the value**—Increases the priority of replication packets over mirror packets which results in slower mirrors during high replication periods but keeps the replication queue from growing so quickly ◆ **Decreasing the value**—Decreases the priority of replication packets over mirror packets which results in faster mirrors during high replication periods but allows the replication queue to grow quickly ◆ **Reason to modify**—Increase the priority of replication packets if the number of queued replication bytes increases too quickly during mirrors, which is most common in high load and/or low bandwidth environments. Increasing the priority may reduce or eliminate auto-disconnects that occur during mirrors.

Setting	Management Console Name (Server Properties Tab)	Description
Mirror Chunk Size	Size of Mirror Packets (Source)	Block size used in the mirroring process ◆ **Default**—32768 bytes ◆ **Increasing the value**—Increases the maximum size of mirror packets, which may improve mirror performance slightly due to a potential decrease of overhead since fewer operations are required to complete the mirror [a] ◆ **Decreasing the value**—Decreases the maximum size of mirror packets, which may degrade mirror performance slightly due to a potential increase of overhead since more operations are required to complete the mirror [b] ◆ **Reason to modify**—Increasing this value may improve mirror throughput marginally. This depends on file size. Larger files with a larger mirror chunk size may see improvement while small files would not see much improvement at all. Typically, the default setting is the optimal setting.
Mirror Queue Limit	Maximum Pending Mirror Operations (Source)	Maximum number of mirror operations that can be queued on the source machine ◆ **Default**—1000 ◆ **Increasing the value**—Allows the source system to queue more mirror operations causing more memory to be used (600 bytes per operation). This may improve mirror performance only if Double-Take is emptying the mirror queue regularly during a mirror. ◆ **Decreasing the value**—Allows the source system to queue fewer mirror operations causing less memory to be used (600 bytes per operation). This may decrease mirror performance if the lower value results in Double-Take emptying the mirror queue regularly. ◆ **Reason to modify**—Increasing this value may improve mirror throughput marginally, although the default setting is optimal in almost all cases.
Pagefile Size	Pagefile Size (Memory)	Size of the Double-Take pagefile ◆ **Default**—256 MB ◆ **Increasing the value**—Allows Double-Take to use more memory to queue operations and data ◆ **Decreasing the value**—Allows Double-Take to use less memory to queue operations and data. ◆ **Reason to modify**—The source pagefile size should be increased if auto-disconnects are occurring frequently and the system has memory that is not being utilized. This would allow Double-Take to use more memory to queue operations and data. The pagefile size can be reduced to limit Double-Take's use of memory if the system is paging memory to disk and affecting system performance.
TG Mirror Capacity High[c]	Target Mirror Capacity High (Target)	Maximum percentage of pagefile that can contain mirror data before the target signals the source to pause the sending of mirror operations ◆ **Default**—20% ◆ **Increasing the value**—Allows the target to queue more mirror operations before pausing the connection ◆ **Decreasing the value**—Allows the target to queue fewer mirror operations before pausing the connection ◆ **Reason to modify**—If a Double-Take target has available memory, this value can be increased to allow Double-Take to use more memory during mirror operations. If Double-Take is using more memory during a mirror than is desired, lower the value. The default setting is optimal in most cases. [d]

Setting	Management Console Name (Server Properties Tab)	Description
TG Mirror Capacity Low[e]	Target Mirror Capacity Low (Target)	Minimum percentage of pagefile that can contain mirror data before the target signals the source to resume the sending of mirror operations. ◆ **Default**—15% ◆ **Increasing the value**—Causes the pause/resume cycle during a mirror to be shorter ◆ **Decreasing the value**—Causes the pause/resume cycle during a mirror to be longer ◆ **Reason to modify**—This value should be adjusted in relation to the TG Mirror Capacity High value. Although the values can differ by as little as one, a difference of at least five is recommended to avoid excessive pause/resume cycles.
TG Retry Locked	Retry Delay for Incomplete Operations (Target)	Minimum number of seconds to wait before retrying a failed operation on a target ◆ **Default**—3 seconds ◆ **Increasing the value**—Increases the interval between retries of unfinished operations ◆ **Decreasing the value**—Decreases the interval between retries of unfinished operations ◆ **Reason to modify**—If unfinished operations are occuring, decreasing the retry interval will allow the target to commit the operation sooner when the condition causing the unfinished operation is resolved. The default setting of 3 is optimal in almost all cases.

a. In Double-Take version 4.1 and earlier, this will cause more bytes to be queued in the target Double-Take pagefile during a mirror since TG Open File Count High specifies the number of operations that can be queued before pausing.

b. In Double-Take version 4.1 and earlier, this will cause fewer bytes to be queued in the target Double-Take pagefile during a mirror since TG Open File Count High specifies the number of operations that can be queued before pausing.

c. This option is called TG Open File Count High in Double-Take version 4.1 and earlier and if using the default 256 MB Double-Take pagefile defaults to 4096.

d. If you are using a many-to-one Double-Take configuration, you may need to lower the setting to allow for the target to better manage its memory for each connection.

e. This option is called TG Open File Count Low in Double-Take version 4.1 and earlier and if using the default 256 MB Double-Take pagefile defaults to 3072.