

AppAssure Software Inc.

Written By: Kaustubh Pradhan

[REPLAY: PLANNING GUIDE]

This guide describes the considerations you should take into account when planning to deploy Replay technology.

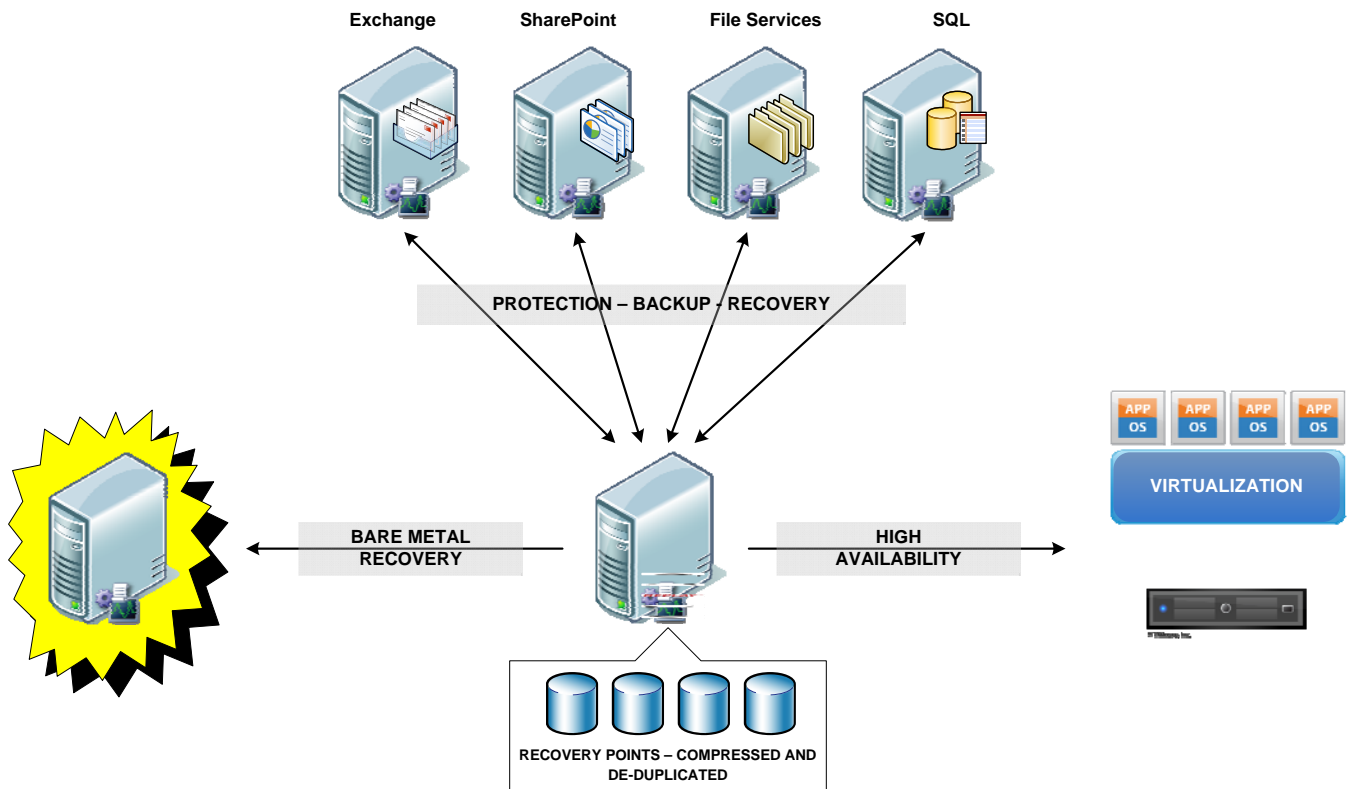
TABLE OF CONTENTS

TABLE OF CONTENTS	3
REPLAY	4
PRODUCT ARCHITECTURE	4
REPLAY UNIQUE FEATURES	4
REPLAY DEPLOYMENT PLANNING GUIDE	6
Architecture	6
Understanding Change Rate	6
REPLAY DATA DEDUPLICATION AND COMPRESSION	8
THE SERVER, OPERATING SYSTEM AND APPLICATIONS	10
Storage Subsystem	10
Processor and Memory	10
Supported Operating Systems	10
Protecting Mission Critical applications with Replay	10
ESTIMATING THE STORAGE REQUIRED	12
Size of Base Image	12
Size of Recovery Points	12
Temp Space	13
LOCAL AND WIDE AREA NETWORKS	15
NUMBER OF REPLAY SERVERS	17
Thumb Rules For Estimating Servers	17

REPLAY

PRODUCT ARCHITECTURE

Replay can be deployed as a single-server or a multi-server architecture. Single-server architecture consists of Replay Server and Agent running on the same server typically seen in small environments where as a multi-server architecture consists of a dedicated Replay server protecting one or multiple agents running different application workloads.



REPLAY UNIQUE FEATURES

CORRUPTION DETECTION

It is known that 20% of all outages in Exchange are due to data corruption.

When Microsoft Exchange 2003 and 2007 mailbox servers are protected with Replay, it performs a mountability check on all exchange databases after every incremental snapshot. This corruption detection feature delivers the ability to alert administrators before failure and provides a path for immediate recovery in the case of a failure. Replay also performs an off-host nightly attach to all the databases being protected on a SQL server

OFF-HOST PROCESSING

Backup applications consume memory and processor cycles application servers, therefore they have traditionally been run during off-hours. Also, as the data size continues to grow backup windows continue to shrink. Lastly, traditional backup applications are not equipped deliver effective post data processing such as compression and data de-duplication.

Replay allows ridding backup windows and alleviating server load caused by processor and memory hog backup programs by offloading all tasks to the Replay server. Also by performing continuous incremental snapshots, Replay eliminates backup windows and its optimized architecture allows effective compression and de-duplication of data helps you save costs associated with long term storage. Lastly, using a single pass approach Replay also delivers the ability to perform transaction log management on Exchange servers.

REPLAY LIVE

Replay enables you to bring your application service online even before the data restore is complete.

Replay Live enables data to be instantly available during recovery – As matter of fact the data is available for use before it is recovered. We also perform Bare Metal Recovery to similar, dissimilar and virtual infrastructure that most competitors do not. Replay Live used with Bare Metal Recovery (both to similar and dissimilar hardware) can be coupled to recover application servers that face total outage

CONTINUOUS HIGH-AVAILABILITY

Native high-availability solutions are expensive, needy and localized. Customers looking at keeping the Recovery Point and Recovery Time to a minimum find management of these solutions to be very overwhelming.

The Standby features of Replay deliver continuous high availability to physical or virtual infrastructure using incremental approach – allows you to keep the Recovery Point and Recovery Time to a minimum. Replay Standby allows you to leverage physical or virtual infrastructure to keep the Recovery Point and Recovery Time to a minimum. Replay standby solutions do not require the acquisition of additional Licenses needed for OS or Application Servers.

REPLAY DEPLOYMENT PLANNING GUIDE

Welcome to the Replay Deployment Planning Guide. This document will help you plan an effective Replay deployment by answering the following questions

- How much Storage will I need?
- How much Bandwidth will I need for Replication?
- How many Replay Servers will I need?

ARCHITECTURE

Replay uses Microsoft Volume Shadow Copy Service technology to capture an online snapshot (point-in-time consistent copy) of server data. The captured data is rapidly transferred to and stored on the Replay Core. Using VSS for Backup does not render the application server in "backup mode" for an extended period of time because the length of time to perform the snapshot is seconds and not hours. Another benefit of using VSS for backups is that it allows snapshot of large quantities of data at one time since the snapshot works at the volume level.

There are two methods for creating shadow copies: making either a complete copy (a full copy) or copying only the changes to the volume (a copy-on-write). Each method results in two data images — the original volume and the shadow copy volume. The functional difference between the two is that the original volume maintains full read/write capabilities, whereas the shadow copy volume is read-only. This read-only status ensures that the shadow copy volume remains a point-in-time copy until its status is changed by the administrator for a specific purpose. A Full Copy is a copy of the original data on a volume whereas a Copy-on-Write method creates shadow copies that are incremental rather than full copies of the original data.

When a snapshot is created, the VSS writer on the target server is told to hold all writes to the disk, this takes only a few seconds. During that period of time, all disk IO operations are simply queued and will be resumed only after the snapshot is complete while the operations already in flight will be completed and all open files will be closed. This process of creating a shadow copy does not significantly impact the performance of the production system.

UNDERSTANDING CHANGE RATE

Once a server is added for protection, the process of protection commences by capturing an initial base image, thereafter the Replay Agent begins tracking all of the volume block changes in real-time. When the protection interval is reached, a VSS snapshot is taken and only the changed blocks are transferred to a Replay Core. The rate at which the blocks on the disk volumes change is known as the change rate. Change rates greatly vary depending upon the type of application, environment and the business.

It is easy to mistake Change Rate for Growth Rate. While Change Rate refers to the frequency at which the existing blocks change within a disk volume pool, Growth Rate refers to the frequency at which new blocks are created within a disk volume pool. Change Rate includes all the blocks associated with Growth Rate but the converse is not true.

A server can be transactional or non-transactional. Transactional systems have a higher change rate compared to non-transactional servers. A transactional server hosting Microsoft Exchange or SQL Server in a small business

will generate 1-3% change; on the other hand the same applications hosted in a Managed Service Provider will experience a change rate of up to 10 – 13%. Similarly, servers such as Domain Controllers, Web Front-end servers, BlackBerry Servers do not generate a lot of change on the disk and therefore are deemed as non-transactional servers. The following table lists how the change rate for a 1000 GB data affects the size of the backup

Change Rate		2%	6%	10%
Total Data		1000 GB	1000 GB	1000 GB
Daily Changed Data		20 GB	60 GB	100 GB
Snapshot interval	15 Minutes	0.21 GB	0.63 GB	1.04 GB
	60 Minutes	0.83 GB	2.50 GB	4.17 GB
	4 Hours	3.33 GB	10.00 GB	16.67 GB
	12 Hours	10 GB	30 GB	50 GB

A Replay Server can manage up to 4 TB of daily changed data.

REPLAY DATA DEDUPLICATION AND COMPRESSION

Using data compression and deduplication as a native part of the backup process eliminates redundant data by referencing duplicate copies of data to a single reference copy and by saving only data blocks that have changed since the last backup. Since only a small percentage of data changes every day, deduplication can result in huge – up to 90% - savings in storage capacity costs alone. The following factors must be considered when designing a data de-duplication solution

1. In-line or offline - Deduplication systems can be accomplished at the file level or block level. File level deduplication eliminates identical copies of whole files where as block level deduplication operates at a sub-file level. *Replay Core performs inline deduplication.*
2. File-level or block-level - Deduplication systems can be accomplished at the file level or block level. File level deduplication eliminates identical copies of whole files where as block level deduplication operates at a sub-file level. *Replay performs block-level deduplication.*
3. Source or Target – Defines where the deduplication is performed. Source-based deduplication provides for the elimination of the redundant data at the source where as target-based deduplication happens on the backup device or system. *For optimal performance and throughput of mission critical applications, Replay offers target based deduplication*

AppAssure Replay integrates backup, data compression, and data deduplication to eliminate data redundancy. The snapshot backups are sent to disk (on dedicated or virtual servers), rather than to tape which further reduces to overall storage footprint and increases the speed of backup operations.

Employing online deduplication, Replay maximizes nearline storage performance and avoids reading and writing backup files with additional storage. Storage capacity is minimized by combining incremental backups and data compression with fine-grain block level deduplication.

Rather than conducting one time backups – once every day, once every week – Replay incrementally backs up data. Backups might run as often as every 15 minutes. In that time frame, only a relatively very small amount of data has changed. Most of the data is redundant. That efficiency and Replay's engineered optimization makes backups run very fast with negligible, usually unnoticeable, performance impact. In fact, nearline storage runs at native storage speeds. The deduplicated and backed up data is then compressed for even greater storage savings.

There are several factors that contribute to the final deduplication and compression ratio that one will experience. Some of these factors are intrinsic to the technology while the others are environmental.

1. **Type of data** is a good indicator of how well it will deduplicate and compress. Files created by office workers often contain redundant data and are frequently distributed or copied. At the other extreme, data derived from natural sources is typically unique. Estimating the benefits of data deduplication for a specific environment may be performed using a redundancy modeling or analysis tool which reflects the level of effort and accuracy desired. Note that company policy may influence the type of data that may

be stored on business systems; for example, if a company does not backup MP3 or JPEG data or prevents creation of PST files, such data would not be deduplicated. Lastly, if the data is already compressed, then further compressing backups might not reduce their size by much, if at all.

2. **Encryption.** Encrypted data compresses significantly less than equivalent unencrypted data. If transparent data encryption is used to encrypt an entire database, compressing backups might not reduce their size by much, if at all.
3. **Change Rate** impacts the likelihood that duplicate data will be created or detected. The less that data is modified the greater the chance that copies of that data will contain the same data as other copies of that data. Frequent update, copy, or append operations may also make it more difficult for some algorithms to detect duplicate data. The efficiency of deduplication and compression is inversely proportional to the change rate. This also implies a higher data deduplication ratio should be expected as the percentage of reference data to total active data increases because reference data is not changed.
4. **Retention Period** or the length of time that data is retained impacts data deduplication ratios in two ways. First, if more data is examined when deduplicating new data, the likelihood of finding duplicate data is increased and the space savings may increase. Secondly, if a data deduplication ratio is calculated over longer periods of time it may increase because the numerator tends to increase more rapidly than the denominator.

Deduplication and Compression Rate	50%	65%	75%
Total Data	1000 GB	1000 GB	1000 GB
Post Deduplication and Compression	500 GB	350 GB	250 GB

Depending upon the aforementioned factors Replay can deliver a deduplication and compression rates from 2:1 (50%) to 4:1 (75%).

THE SERVER, OPERATING SYSTEM AND APPLICATIONS

STORAGE SUBSYSTEM

Replay stores the recovery points of protected servers as epochs in a flat file format on a volume designated as the Replay repository. These repository volumes should preferably reside on a low latency device.

Replay Core supports the use of block-level and CIFS or SMB devices. Devices such as internal storage, direct attached or distant devices accessed via a storage area network (SAN) are classified as Block-Level devices. CIFS or SMB devices are network shares or NAS.

Replay Agent can protect volumes that reside on internal drives, direct attached storage (DAS) or on Storage Area Networks (SAN). Data residing on Network Attached Storage (NAS) cannot be protected by Replay

Type of Storage used does not have any impact on the Level of deduplication and compression.

PROCESSOR AND MEMORY

Replay Core and Agent can run as a 32-Bit or 64-Bit application. Minimum processor requirements are 3 GHz or multi-core 2 GHz (or greater). Use Multi-core CPUs. Two CPUs will not be as fast as one CPU that is twice as fast. In addition, larger L2 processor caches will provide better performance.

Minimum memory requirements are 4 GB for 32-Bit systems and at least 8 GB for 64-Bit systems.

SUPPORTED OPERATING SYSTEMS

Operating System	Platform Architecture
Windows® Server 2008 R2	64-bit
Windows® Server 2008	64-bit
Windows® Server 2003 R2	32-bit or 64-bit
Windows® Server 2003	32-bit or 64-bit
Windows® Small Business Server 2008	64-bit
Windows® Small Business Server 2003	32-bit
Windows® Vista	32-bit or 64-bit
Windows® 7	32-bit or 64-bit
Windows® XP	32-bit or 64-bit

PROTECTING MISSION CRITICAL APPLICATIONS WITH REPLAY

Protecting Exchange 2003 and Exchange 2007 with Replay for Exchange

Replay for Exchange support includes the ability to perform periodic mountability checks on all databases, nightly or weekly page-by-page integrity checksum checks on all databases and finally the ability to truncate storage group logs on the exchange server. These features are only exposed if the following criterion is met

Protected Application and Platform	and Replay Core runs on			
	Windows Server 2003 32-Bit	Windows Server 2003 64-Bit	Windows Server 2008 32-Bit	Windows Server 2008 64-Bit

If Replay Agent is installed on	Exchange Server 2003 on Windows Server (x86)	Yes	Yes	No	No
	On Exchange Server 2007 on Windows 2003 (x64)	No	Yes	No	Yes
	On Exchange Server 2007 on Windows 2008 (x64)	No	Yes	No	Yes

Protecting SQL 2000, 2005 or 2008 with Replay

Replay for SQL support includes the ability to perform an instance aware nightly attachability checks on all databases using a locally installed instance of SQL on the Replay Core. These features are only exposed the following criterion is met

		And Replay Core runs on			
Protected Application and Platform		Windows Server 2003 32-Bit*	Windows Server 2003 64-Bit*	Windows Server 2008 32-Bit*	Windows Server 2008 64-Bit*
If Replay Agent is installed on	SQL 2000 on Windows 2003 (x86)	Yes	Yes	Yes	Yes
	On SQL 2005 on Windows 2003 (x86)	Yes	Yes	Yes	Yes
	On SQL 2005 on Windows 2003 (x64)	Yes	Yes	Yes	Yes
	On SQL 2008 on Windows 2008 (x86)	Yes	Yes	Yes	Yes
	On SQL 2008 on Windows 2008 (x64)	Yes	Yes	Yes	Yes

Replay attachability test requires that equal or a newer version of SQL server be installed on the Replay Core server.

ESTIMATING THE STORAGE REQUIRED

Replay uses data-deduplication and compression to create and store the recovery points on nearline storage. The process of protecting a server starts with an initial base image of the targeted volumes followed by periodic incremental snapshot backups. In fact, these incremental snapshots capture only the changed disk blocks since the previous snapshot backup. More the change larger the size of the snapshot backup and vice versa. To calculate the storage required by Replay use the formula:

$$\text{Storage Required} = \text{Size of Base Image} + \text{Size of Recovery Points} + \text{Temp Space}$$

SIZE OF BASE IMAGE

Size of Base Image is the final deduplicated and compressed size of the initial snapshot backup. To estimate the size of the base image use the formula:

$$\frac{T * \left(1 - \frac{C}{100}\right)}{10^9}$$

Where:

- T = Total used space of data across all the disk volumes in bytes.
- C = Deduplication and Compression Ratio. This is expressed as a percentage between 1 and 100. So for example, if you set C to 65, indicating 65% compression, the compressed data is now 35 percent of its original size.
- R = Total estimated Change Rate. This is expressed as a percentage between 1 and 100. So for example if you set R to 10, it indicates that 10% of the data changed within a Recovery Point.

Use the following thumb rule, small business will generate 1-3% change; on the other hand the same applications hosted in a Managed Service Provider will experience a change rate of up to 10 – 13%.

Average sized businesses will see a change rate of around 5 – 7%.

SIZE OF RECOVERY POINTS

Size of the Recovery Points is the cumulative space used by the recovery points over the life of the retention period. Replay allows you to define a granular retention policy for the periodic snapshot backups. As the backups age the Replay Rollup process automatically rolls them to the next retention tier. The retention tiers are hourly, daily, weekly and monthly. The storage utilization is directly affected by the retention policy. To estimate the size of the recovery points use the formula:

$$\left[\frac{\left(T * \frac{R}{100}\right) * \left(1 - \frac{C}{100}\right)}{10^9} \right] * P$$

Where:

- T = Total used space of data across all the disk volumes in bytes.
- C = Deduplication and Compression Ratio. This is expressed as a percentage between 1 and 100. So for example, if you set C to 65, indicating 65% compression, the compressed data is now 35 percent of its original size.
- R = Total estimated Change Rate. This is expressed as a percentage between 1 and 100. So for example if you set R to 10, it indicates that 10% of the data changed within a Recovery Point.

Use the following thumb rule, small business will generate 1-3% change; on the other hand the same applications hosted in a Managed Service Provider will experience a change rate of up to 10 – 13%. Average sized businesses will see a change rate of around 5 – 7%.

- P = Retention Policy

$$P=N+H+D+W+M$$

Where:

- M = Number of Monthly snapshot backups. Default is 1
- W = Number of Weekly snapshot backups Default is 1
- D = Number of Daily snapshot backups. Default is 3
- H = Number of Hourly snapshot backups. Default is 2
- N = Number of New snapshot backups. Default is 1

TEMP SPACE

Replay requires some upfront storage capacity for post-process task known as Rollup. Rollup requires additional capacity to temporarily store backup data to enforce retention policies that control the recovery point retention time on a daily, weekly, and monthly basis. The rollups occur automatically on a scheduled basis and begin when the nightly job is scheduled. To estimate the upfront capacity use the size of the largest single contiguous volume to be protected.

For Example, to estimate the total storage required to protect a Server with 1000 GB of total data with an approximate change rate of 2%, Compression Rate of 65% and protection interval of 1 hour using the default retention policy

Sizing Factor		Value
Total Data		1000 GB
Daily Change Rate		2%

[REPLAY: PLANNING GUIDE]

Compression		65%
Snapshot Interval		1 Hour
Daily Changed Data		20 GB
Daily Changed Data After Compression		7 GB
Size of Base Image		350 GB
Total Size of Monthly Recovery Point	7 GB * 30	210 GB
Total Size of Weekly Recovery Point	7 GB * 7	49 GB
Total Size of Daily Recovery Point	7 GB * 3	27 GB
Total Size of Hourly Recovery Point	7 GB * 2	14 GB
Total Size of New Recovery Point	7 GB * 1	7
Upfront Capacity		350 GB
Space Required for Monthly Retention		1007 GB

LOCAL AND WIDE AREA NETWORKS

Network performance is critical to Replay as it uses the network connection to transfer the snapshots from the protected server to the Replay server. Therefore, all components of the network subsystem need to perform optimally and ensure maximum performance.

Use adapters with the Gigabit bandwidth available for best performance. Note that increasing bandwidth increases the number of transmissions that are taking place and in turn makes more work for your system, including more interrupts being generated. Remove unused network adapters to reduce overhead.

Replay delivers a consistent throughput of over 5 GB/minute on GbE networks.

Replay has the ability to replicate recovery points of a protected server to a peer Replay server. This replication gives you the ability to provide off-site backups and off-site disaster recovery. The process simply transfers the compressed and de-duplicated recovery to the remote peer server that can exist across a LAN or a WAN. The replication can be setup in the following configurations

- 1:1 – This is a point to point configuration where a Replay core can replicate uni- or bi-directionally with another peer Replay core.
- MANY:1 – This is a configuration where one Replay core will serve as a single target for multiple Replay cores

To plan Replay replication, consider the following

1. Total size of the initial seed. The initial seed consists of the Base Image and additional snapshots. Replay offers an online and an offline method to transfer these snapshots to the remote location. When using the online method ensure that you have enough bandwidth to replicate the data to the remote location. Alternatively, Replay offers a Copy-Consume approach which allows you to perform an offline copy of the recovery points to the remote location. Replay automatically resumes from the last snapshot within the offline copy. To validate if you have sufficient bandwidth, you can use the bandwidth tools available on <http://www.bandwidth.com/tools/calc.html> or <http://www.ibeast.com/content/tools/band-calc.asp> to estimate the time required to copy the initial base image and snapshot seed to the remote location.
2. Bandwidth capacity for replication of snapshots. To validate if you have enough bandwidth, you can use the bandwidth tools available on <http://www.bandwidth.com/tools/calc.html> or <http://www.ibeast.com/content/tools/band-calc.asp> to estimate the time required to copy the incremental snapshots to the remote location.

Link Connection	Link Capacity	Time Required to Transfer Deduplicated and Compressed Snapshots and/or Base Images		
		100 MB	100 GB	100 TB

[REPLAY: PLANNING GUIDE]

Link Connection	Link Capacity	Time Required to Transfer Deduplicated and Compressed Snapshots and/or Base Images		
		100 MB	100 GB	100 TB
33.6 K (Modem)	33,600 bps	41 mins, 36 secs	29 days, 14 hrs, 8 mins, 48 secs	83 years, 4 days, 14 hrs, 13 mins, 48 secs
56 K (Modem)	56,000 bps	24 mins, 57 secs	17 days, 18 hrs, 5 mins, 16 secs	49 years, 294 days, 18 hrs, 8 mins, 16 secs
64 K (DS-0)	64,000 bps	21 mins, 50 secs	15 days, 12 hrs, 49 mins, 37 secs	43 years, 212 days, 6 hrs, 52 mins, 14 secs
128 K (ISDN)	128,000 bps	10 mins, 55 secs	7 days, 18 hrs, 24 mins, 48 secs	21 years, 288 days, 15 hrs, 26 mins, 7 secs
256 K (DSL)	256,000 bps	5 mins, 27 secs	3 days, 21 hrs, 12 mins, 24 secs	10 years, 326 days, 19 hrs, 43 mins, 3 secs
640 K (DSL/Cable)	640,000 bps	2 mins, 11 secs	1 day, 13 hrs, 16 mins, 57 secs	4 years, 130 days, 17 hrs, 29 mins, 13 secs
768 K (DSL/Cable)	768,000 bps	1 min, 49 secs	1 day, 7 hrs, 4 mins, 8 secs	3 years, 230 days, 14 hrs, 34 mins, 21 secs
T1, DS-1	1.544 Mbps	54 secs	15 hrs, 27 mins, 14 secs	1 year, 294 days, 8 hrs, 51 mins, 54 secs
T3, DS-3	44.736 Mbps	1 sec	32 mins	22 days, 18 hrs, 10 mins, 22 secs
OC-1	51.840 Mbps	1 sec	27 mins, 37 secs	19 days, 15 hrs, 19 mins, 37 secs
OC-3	155.520 Mbps	Less than a second.	9 mins, 12 secs	6 days, 13 hrs, 6 mins, 32 secs
OC-12	622.520 Mbps	Less than a second.	2 mins, 18 secs	1 day, 15 hrs, 16 mins, 38 secs
OC-48	2.488 Gbps	Less than a second.	34 secs	9 hrs, 49 mins, 14 secs
OC-192	10Gbps	Less than a second.	8 secs	2 hrs, 26 mins, 36 secs

NUMBER OF REPLAY SERVERS

The Replay deployment plan should specify the number of Replay Core servers necessary to protect your data and where you plan to locate each Replay Core server within your network. Replay Core servers can perform additional post-processing tasks such as performing detailed page by page integrity checks of Exchange Databases and Nightly attachability tests on SQL databases. These post processing tasks need additional processing and therefore need must be considered when developing a deployment plan.

As you consider the number of Replay Core servers that your organization requires, keep in mind that there is no precise formula for determining the number of Replay Core servers. In practice, the number of servers and amount of data that a single Replay Core server can protect varies based on the following factors:

1. Total of Exchange and SQL Servers
2. Total data on the source
3. Type of server – Transactional or Non-transactional.
4. Change Rate of the data sources to be protected – low (2%), medium (6%) or high (10%)
5. Frequency of Snapshot Backups – frequent or intermittent

THUMB RULES FOR ESTIMATING SERVERS

1. Backup non-transactional Servers every 12 or 24 hours. There will be no additional benefit by using a more frequent protection interval
2. Backup transactional Servers every 1 hour. When protecting Exchange and SQL Servers with an added application protection options, use a protection of 2 hours.
3. Create a mix of transactional and non-transactional servers when protecting a large numbers of workloads
4. Intermittent snapshots means more protected servers per Replay Core server, whereas, frequent snapshots mean less protected server per Replay Core server.