



Technical White Paper

Symantec Backup Exec 10d System Sizing

**“Best Practices”
For
Optimizing Performance
of the
Continuous Protection Server**

Table of Contents

Table of Contents	2
Executive Summary	3
System Sizing and Performance Considerations	4
Calculating Performance.....	5
Measurable Elements	5
Data Factors	7
Test Results	7
Initial Image	7
Performance Tests.....	8
Scaling Tests	10
Proper System Sizing: Best Practices	12
Appendix A – Sizing Process and Formulae	13
About Diogenes Analytical Laboratories, Inc.	15



Executive Summary

Continuous data protection (CDP) is an emerging technology that offers faster data restoration and more certain data protection than traditional backup and recovery (B/R) technologies. Traditional backup and recovery using magnetic tape technology generally is limited to a recovery time of 24 hours: If a failure occurs today, the most recent data available to restore was yesterday's. Thus, any data added or changed since yesterday will be lost. This would be referred to as a "24-hour recovery point objective (RPO)." Therefore, up to 24 hours of data can be at risk.

Data-replication methodologies have emerged to reduce the 24-hour loss window of risk associated with tape backup. One such approach, called mirroring, creates a complete replication of the data copied to a separate disk system. If the primary disk system fails, the secondary system can be mounted for continued processing. With a data mirror, complete data restoration can be accomplished in minutes rather than the hours that a tape system would require. But because mirrors double the amount of disk space required, most organizations will mirror the data no more than every eight hours. (Note that each mirror image incurs 100% disk space overhead.) With such a scheme, an eight hour RPO would be in place.

Data snapshots, another replication technology, are more efficient than data mirrors in that they capture only changed data. Thus, the overhead is substantially less, typically 3% to 20% depending upon the rate of data change. Most organizations will implement snapshots every four hours, although some will do so every hour. With this scheme, the RPO can be as little as an hour. However, snapshots have a key limiting factor: The snapshot is stored on the primary disk system. If the primary system fails, the snapshot will be lost, as well.

Continuous data protection does what its name implies: capture data changes as they occur. With CDP, RPOs can be reduced to an hour or less. Moreover, CDP data is stored on a disk system separate from the primary disk, so it will not be lost if the primary is lost as a snapshot would be. Because CDP captures only changed data, it does not incur the 100% disk space per copy overhead as mirror images do.

Because CDP operates fundamentally different from traditional B/R, mirror and snapshot technologies, the stress points that it places on system architecture are somewhat different. To characterize these points and to quantify best practices, Symantec commissioned Diogenes Analytical Laboratories, Inc., to conduct appropriate lab tests.¹ We installed Backup Exec 10d in a variety of configurations and scenarios in order to stress the three main points of a CDP system: system CPU, network and disk subsystem. We also wanted to find out if there is a limit to the number of file servers that can be attached to a single CDP server. This document does not discuss the operation of Backup Exec 10d in detail, and it is assumed that the reader is generally familiar with that product and its operation.

The principles of sizing a CDP system are similar to those of sizing a B/R system, except that tape drive speed is not a factor. From our testing, the key findings were:

¹ Diogenes is an independent lab that does not endorse products. This document should not be construed as an endorsement nor as a statement that any technology or product is suitable for a particular purpose.



1. A two-way Windows server as the CDP server offers ample capacity for the preponderance of situations. Scalability of CPU resources is unlikely to cause performance problems.
2. Either the network or disk subsystem is likely to be the limiting factor of performance.
3. Backup Exec 10d was able to back up 33 GB of data spread across 33 file servers in just 35 minutes without loss of functionality, even when 20 of the servers were connected using a 10/100 network.

The third point is important because it measures the efficiency of the application. The test placed just a 20.4% load on the CPU. While it is possible to undersize any component, we believe that following the guidelines in this document will assure satisfactory performance for the preponderance of environments.

System Sizing and Performance Considerations

When an IT organization implements a Backup Exec 10d environment, two possible configurations are most common. As a best practice, CDP should be implemented as a supplement to B/R; no CDP product is generally considered a B/R replacement. IT organizations will continue to back data up to tape either for offsite vaulting or long-term data retention. Therefore, a media (backup) server is still necessary. One key question that we sought to answer in our lab tests is when to separate the media server from the CDP server. Figure 1 illustrates a consolidated configuration, in which the media server and the CDP server are running on the same server. For simplicity, we have not shown the connection to other systems, such as tape libraries, as they do not factor into our calculations.

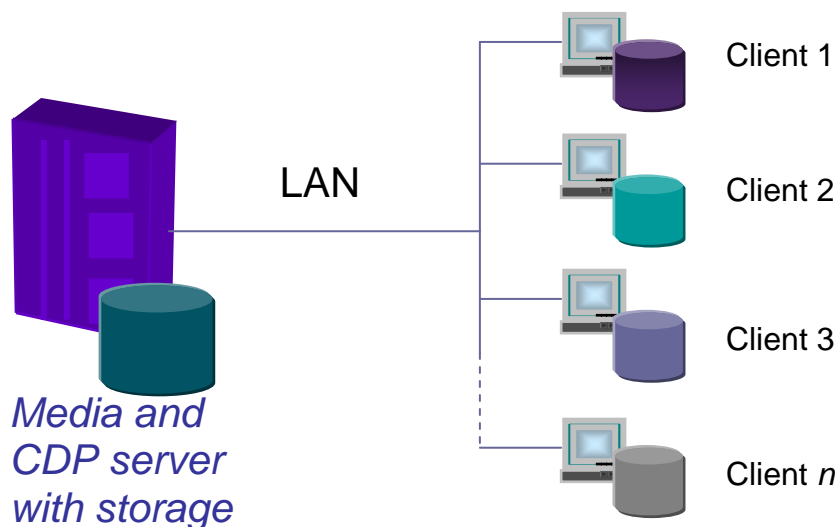


Figure 1 – Consolidated configuration

The three potential bottlenecks of this configuration are the LAN bandwidth, the media/CDP server processing power, and the CDP storage I/O throughput.



Figure 2 illustrates an extended configuration in which the media server and CDP server are separated.




Figure 2 – Extended configuration

Obviously, the consolidated configuration yields a lower cost of ownership and is therefore preferable all other things being equal. When do best practices dictate that the two functions be separated? We will answer this in the succeeding sections.

Calculating Performance

Measurable Elements

The elements that must be calculated for performance purposes include:

1. Total data to be protected per hour
2. CPU utilization
3. Network bandwidth
4. Disk I/O throughput

Backup and recovery operations are often very CPU-intensive, consuming 60% - 80% of the CPU during peak times. In most organizations, backup activity is most intense during the evening and night hours. Figure 3 illustrates common CPU utilization over a normal backup cycle using traditional backup methods.

The line in Figure 3 could also represent aggregate resource utilization, including the CPU, network and disk I/O. When IT organizations assess this total resource line, they should consider that free resources above the utilization line are available for other applications such



as CDP. Clearly, the total resources or any single resource needed cannot exceed 100% without causing some jobs to fail. That is, the slowest resource becomes the “weak link” of the entire system. If the combined system requirements exceed 100% at the peak time, then additional resources are needed. If the CPU resources are exceeded, then separating the media server from the CDP server is the next step.

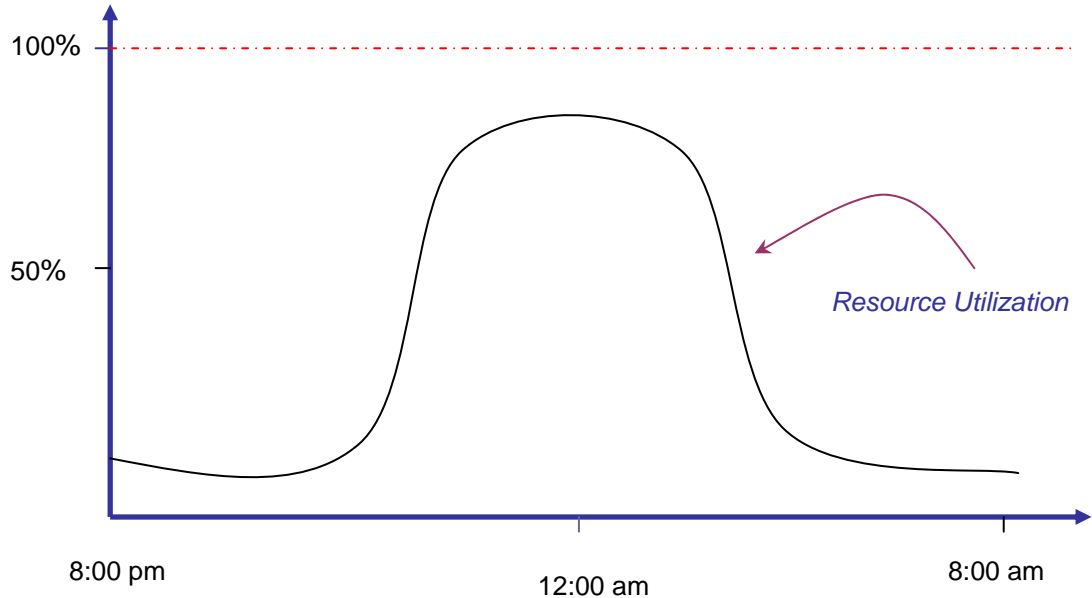


Figure 3 – Resource utilization over time in a typical B/R cycle

When calculating resource availability, storage administrators must also keep restore operations in mind. Sufficient resources should be reserved to accommodate data restore operations. We recommend that data protection consume no more than 80% of the total available resources to allow for unexpected restore requirements.

The first task in properly sizing a Backup Exec 10d implementation is to quantify the existing resource utilization. Thus far, our discussion has focused on system resources. However, the key factor is the amount of data that must be moved. After the initial image is created on the CDP server, Backup Exec 10d moves only new and changed data from the primary server. The calculation for the total amount of data is simple, as follows:

$$\text{Data change rate (GB) per hour} + \text{new data rate (GB) per hour} = \text{total data/hr}$$

Readers should note that Backup Exec 10d sends only changed data (blocks) except for small files. If a file is smaller than 1 MB, then the entire file will be transmitted.

Network metrics are easier to determine. Some sample network throughput values are as follows:

10/100 Ethernet: 1.0 MB/sec – 10 MB/sec per file server

Gigabit Ethernet (GbE): 100 MB/sec per file server



These rates are theoretical maximums, and the actual throughput can be substantially less. Dedicated subnets will yield throughput closer to the maximum than networks that mix user data and backup data. A safe estimate is 70% of the theoretical maximum.

As we noted earlier, B/R can consume a significant portion of the media server CPU. The actual amount will vary as a result of a number of factors, but those factors are not really material. It is better for an IT organization to measure the actual CPU utilization over a period of time to determine peak periods usage. Remember, it is not the average that is important but the peak value. CPU utilization is easily captured using Windows Performance Monitor (Perfmon). Perfmon will capture and graph CPU and memory utilization to help determine average and peak usage.

Also, determining disk system throughput must be accomplished through actual monitoring and testing. With the wide variety of disk subsystems on the market, a theoretical calculation is not beneficial. Fortunately, open source tools are available to perform this task. We recommend using IOmeter (www.iometer.org).

Data Factors

Another factor that impact system sizing and performance is the nature of the data. Currently, Backup Exec 10d Continuous Protection Server supports file system data, so it is not necessary to consider Exchange or SQL Server data for these tests. Within file systems, data can be characterized as predominantly small file (less than 1 MB), large file (greater than 1 MB) or mixed. From a CDP perspective, Backup Exec 10d indexes files as they are transferred to the CDP server to facilitate faster restore. This indexing requires both processing cycles as well as disk space.

Given the measurable factors and the data factors, our tests set out to establish guidelines that IT organizations can use in sizing their Backup Exec 10d environment.

Test Results

Initial Image

After installing the Backup Exec 10d software, our first task was to create an initial image of the primary data on the CDP server. Our first test, then, was to determine how much disk space should be allocated for this initial image. Because of the indexing, we did not expect a 1:1 ratio but needed to determine what the guidelines are. For our tests, we used three scenarios:

- Scenario 1: Many small files, deep tree structure
- Scenario 2: Mix of large and small files, moderate tree structure
- Scenario 3: Large files, flat tree structure

Figure 4 indicates the results of our tests of these scenarios. Table 1 indicates how the different scenarios were created.

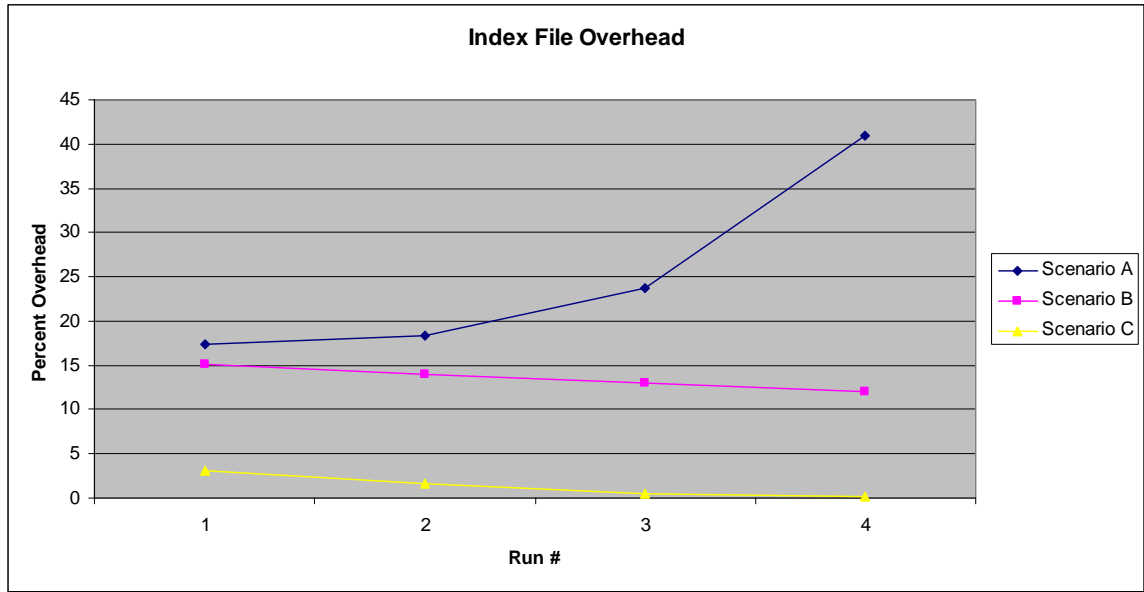


Figure 4 – Disk space overhead resulting from indexing

	Run No. 1	Run No. 2	Run No. 3	Run No. 4
Scenario A	9 files per directory each 2KB 10 directories per folder 3 folders deep	9 files per directory each 2KB 10 directories per folder 4 folders deep	9 files per directory each 2KB 10 directories per folder 5 folders deep	9 files per directory each 2KB 10 directories per folder 6 folder
Scenario B	4,095 files per directory ranging from 2KB to 4MB 3 directories per folder 2 folders deep	4,095 files per directory ranging from 2KB to 4MB 4 directories per folder 3 folders deep	4,095 files per directory ranging from 2KB to 4MB 5 directories per folder 4 folders deep	4,095 files per directory ranging from 2KB to 4MB 6 directories per folder 5 folder
Scenario C	30 files per directory each 4MB 0 directories per folder 0 folders deep	65 files per directory each 4MB 0 directories per folder 0 folders deep	200 files per directory each 4MB 0 directories per folder 0 folders deep	1,000 files per directory each 4MB 0 directories per folder 0 folders

Table 1 – Description of how the test scenarios were created.

As the graph indicates, additional disk space allocated for small file systems in a very deep tree structure can be as much as 42%, whereas 15% is sufficient for mixed file and just 5% for large files environments.

Performance Tests

To provide guidance to Backup Exec 10d users, we created five different scenarios. These included a variety of CPU, network and client configurations. Each of these scenarios and the results are described below. Readers should note that neither the total amount of data in the file system nor the percent of data changed is important; only the total amount of new and changed data impacts a CDP system. In the first two scenarios, we wanted to determine



the limiting factor of the environment. Thus, our variable factors were the elements of the infrastructure.

Scenario 1:

Server: Dual processor 3.06 Ghz Xeon, 2 GB memory

Clients: 3 file servers with 10 GB of changed data each (30 GB total)

Network: 10/100 Ethernet

Theoretical 10 MB/sec throughput (30 GB in 50 mins.)

Disk subsystem: 1 Gbs Fibre Channel, RAID 0, 500 GB

Theoretical 100 MB/sec throughput (30 GB in 5 mins.)

Results:

CPU Utilization: 13%

Memory Utilization: 17%

Total elapsed time: 2 hrs 34 mins.

Network throughput: 3.8 MB/sec.

From these results, we can see that the network throughput was substantially lower than the theoretical throughput. To determine the cause, we re-ran the test with just a single client. In this test, we achieved 10 MB/sec, the maximum. So, we ran the test a third time with two clients. In this case, network throughput was just 6.4 MB per second. Thus, it became clear that network performance deteriorated significantly with additional clients, probably because of network contention. In subsequent tests, we were able to connect 20 servers using a 10/100 network (see Scenario 5 below). If IT organizations experience performance degradation on a 10/100 network, it may be advisable to upgrade to newer networking infrastructure such as GbE. We did not experience any network degradation problems in later tests using GbE. The low CPU and memory utilization indicate that this configuration could be quite adequately served by a single-processor system.

Scenario 2:

Server: Quad processor 1.5 Ghz Xeon, 3.5 GB memory

Clients: 6 file servers with 15 GB of changed data each (90 GB total)

Network: Gigabit Ethernet

Theoretical 100 MB/sec throughput

Disk subsystem: SCSI Ultra 160, 700 GB RAID 1

Theoretical 160 MB/sec throughput

Results:

CPU Utilization: 13.3%



Memory Utilization: 8%
Disk throughput: 22.8 MB/sec
Network throughput: 24 MB/sec
Total elapsed time: 1 hrs 9 mins.

In this scenario, the disk subsystem was our limiting factor. Although Ultra SCSI 160 supports 160 GB/sec channel speed, the combination of RAID, disk allocation and controller speeds results in substantially lower actual throughput. Running at 22.8MB/s the 90GB should have taken 1 hr and 6 minutes, while the actual time was 1 hour 9 minutes. To improve performance, we would have needed to use a faster disk system. It is also worth noting that the CPU utilization was only a small fraction of its capacity. A dual-processor system, and perhaps even a single processor system, would deliver similar results at a fraction of the cost of a quad processor.

Scaling Tests

After completing our performance tests, it was clear that we could use Backup Exec 10d to the maximum capacity of the disk or the network. So we next wanted to know to what extent the addition of more file servers to the environment would impact total performance. Adverse impact would manifest as either degraded disk performance as a result of disk contention (thrashing), or increased CPU/memory utilization.

For these tests, we applied a mix of file types and sizes. The total amount of new and changed data ranged from 1 GB per file server to 24 GB per server. In most production situations, new and changed data are applied gradually during the course of business. However, applying large amounts of data (e.g., 24 GB) in a single event is referred to as a “data bomb.” Data bombs are a worst case scenario for CDP systems. And, that’s exactly what we did to Backup Exec 10d.

As a result of the data bomb, Backup Exec 10d’s journaling system operation temporarily suspended because of journal system overflow. It is important to note that the jobs did not fail and that normal replication continued. However, rather than reduce functionality we chose to run the product at full capability and throttled the disk throughput to a maximum of 11.4 MB/sec. We should also note that we believe the data volumes we applied to the test exceed those that will be typical in a small –to medium enterprise environment. IT organizations should evaluate their data volumes, as noted in our best practice recommendations at the end of this document, and determine whether to throttle disk performance or allow journaling to be temporarily suspended during peak periods.

For the scaling tests, our infrastructure became the constant and the number of file servers became the variable. Our infrastructure included:

Server: Quad processor 1.5 Ghz Xeon, 3.5 GB memory

Network: Gigabit Ethernet

Theoretical 10 MB/sec throughput

Disk subsystem: Ultra SCSI 160, 700 GB RAID 1

Set to 11.4 MB/sec throughput



Scenario 3:

8 file servers connected via GbE
22 GB of new/changed data per server
Theoretical best time of 4 hrs. 8 mins.

Results:

CPU Utilization: 15.9%
Memory Utilization: 8.4%
Total elapsed time: 4 hrs 10 mins.

The results do not indicate any adverse results from eight file servers.

Scenario 4:

13 file servers connected via GbE
24 GB of new/changed data per server
Theoretical best time of 7 hrs. 57 mins.

Results:

CPU Utilization: 16.0%
Memory Utilization: 8.3%
Total elapsed time: 7 hrs 59 mins.

The results do not indicate any adverse results from 13 servers.

Scenario 5:

13 file servers with connected via GbE
20 file servers with connected via 10/100 Ethernet
1 GB of new/changed data per server
Theoretical best time of 32 mins.

Results:

CPU Utilization: 20.4%
Memory Utilization: 8.5%
Total elapsed time: 35 mins.

Again, the results do not indicate any adverse results from 33 servers. The key issue in this scenario was possible processing overhead from simply connecting to and managing a larger number of file servers. Clearly, Backup Exec 10d easily scaled to handle 33 file servers.



Proper System Sizing: Best Practices

From the tests that we conducted, we recommend the following steps to properly sizing a Backup Exec 10d environment:

1. Calculate (or estimate if necessary) the combined peak data growth and change rate (referred to as total peak data).
 - a. If the total amount of data exceeds 40 GB of data per hour, then consider adding an additional CDP server. [*Based on our tests, 33 GB were backed up in 35 minutes. However, performance was not linear, so this data rate cannot be extrapolated with absolute certainty.]
2. Measure existing infrastructure elements to determine the following metrics:
 - a. Server CPU and memory during B/R operations (if a consolidated environment is anticipated)
 - b. Network bandwidth (GB/hr)
 - c. Disk I/O throughput
3. Divide the total peak data by the network bandwidth.
 - a. If the total peak data cannot be moved across the network within 80% of the Recovery Point Objective, i.e. if the RPO is 1 hr, then the data should be able to be transferred within 48 minutes (in order to allow time for restores, resource contention, etc.), consider upgrading network resources.
 - b. If your CDP data is operating on a 10/100 network and performance is insufficient, then consider upgrading to GbE.
4. Divide the total peak data by the disk I/O throughput
 - a. If the total peak data cannot be moved from the disk within 80% of the RPO, consider upgrading the disk subsystem (e.g., faster disk array, spread the data across multiple arrays)
5. Separate the media server from the CDP server if:
 - a. Peak CPU utilization during backup operations exceeds 60%
 - b. Average CPU utilization during backup operations exceeds 50%
 - c. Restore requirements are highly variable or unknown and media server availability must be assured
6. If any infrastructure element is insufficient, then it must be improved to the proper level.
7. We did not achieve an upper limit to the number of files servers that could be attached to a Backup Exec CDP server based on CPU overhead. The limiting factors are more likely to be data volume and its impact on network throughput and disk I/O throughput.
8. A dual processor CDP server should be sufficient for all but the largest configurations.



Appendix A – Sizing Process and Formulae

Similar to sizing a backup environment, properly sizing a CDP system starts with determining the slowest performing component. With backup, the slowest component is typically either the tape drive or the network. CDP is similar in that the slowest component is most likely to be the network or the disk drive. The total amount of data that can be backed up is no more than the slowest component. The formula for calculating how much data can be backed up in a day is:

$$\text{Amount of Data (GB)} = 24 \text{ hours} * (X \text{ in GB/hr})$$

where X is the transfer rate of either the disk drive or network – whichever is the SLOWEST.

However, using 24 hours is not realistic, as no time is allocated for restores. So an advisable 80% should be used instead or about 19 hours. Determining the speed is simple; IOMeter can be used to measure the transfer rate of both disk drives and network.

However, if the CDP server is functioning as the tape backup (media) server, then other considerations must be factored. For this calculation, we are assuming that the CDP server will not be running at the same time as the backups. Readers should note that running both operations in parallel could impact performance significantly as the CDP server would be trying to write data to disk while the backup software is trying to read, resulting in disk I/O contention.

$$\text{Starting with: } T_c + T_b = 19 \text{ hrs (remember the 80\% rule above)}$$

Where T_c is the time spent doing CDP and T_b is the time spent doing backups.

$$T_c = \text{Amount of Data (GB)} / X \text{ (GB/hr)}$$

Where device speed is the same as above, the slower of the two transfer rates for disk or network and:

$$T_b = \text{Amount of Data (GB)} / Y \text{ (GB/hr)}$$

where Y is the slower of the two transfer rates for tape or network. Plugging these into our original equation, we get:



$$\text{Amount of Data (GB)} / X(\text{GB/hr}) + \text{Amount of Data(GB)} / Y(\text{GB/hr}) = 19\text{hrs}$$

Now, solving this equation for the amount of data we get:

$$\text{Amount of Data (GB)} = \frac{19 \text{ hrs}}{(1 / X (\text{GB/hr}) + 1 / Y (\text{GB/hr}))}$$

Again, IOmeter can be used to determine network xfer, disk xfer. For tape, 80% of the theoretical maximum is often a good rule of thumb. It is important to note that the amount of data is not the amount of storage the server has but rather the amount of data that is either new or changed.



About Diogenes Analytical Laboratories, Inc.

Diogenes Analytical Laboratories, Inc. is an independent organization dedicated to helping information technology buyers reduce the inherent risk and uncertainty associated with technology purchases. Our goal is to create an informed IT consumer and provide the complete information needed to make smart purchase decisions. This report is based on Diogenes Analytical Laboratories' actual lab testing experiences. The opinions expressed in this report are those of Diogenes Analytical Laboratories, Inc.

For more information, go to www.diogeneslab.com or call 303-402-9900.

© 2006 Diogenes Analytical Laboratories, Inc. All rights reserved. Neither this report nor any portion of it may be reproduced in any form without the express written consent of Diogenes Analytical Laboratories, Inc. The use of this report is governed by a separate Product and Services Use Agreement which may be viewed at www.diogeneslab.com/termsfuse.htm.

Trademark notices:

The Diogenes logo and flame are trademarks of Diogenes Analytical Laboratories, Inc.
Backup Exec 10d is a trademark of Symantec Corp.

