



**IBM Tivoli Storage Manager for Virtual Environments  
Data Protection for VMware**

**Step by Step Guide To vStorage Backup Server (Proxy) Sizing  
20 May 2016**

*Authors: Dan Wolfe  
Art Roy*



## Tivoli Storage Manager for Virtual Environments Guide to Proxy Sizing

**Note:** Before using this information and the product it supports, read the information in the [Notices](#) section.



- 1. Introduction ..... 5
  - 1.1 Overview ..... 6
    - 1.1.1 Backup performance and incremental forever ..... 7
    - 1.1.2 Backup performance and full backups ..... 7
      - 1.1.2.1 Data deduplication ..... 7
      - 1.1.2.2 Data transfer/transport methods ..... 8
  - 1.3 Scope of this document ..... 9
    - 1.3.1 What is not covered by this document ..... 10
    - 1.3.2 External dependencies and assumptions ..... 10
    - 1.3.3 Proxy hardware configuration ..... 10
  - 1.4 Scheduling of backups ..... 10
- 2. Step-by-step proxy sizing ..... 12
  - 2.1 Assumptions ..... 12
  - 2.2 Example environment ..... 12
    - 2.2.1 Notes regarding example environment ..... 12
  - 2.3 Perform the estimate ..... 13
    - 2.3.1 Calculate steady state daily workload ..... 13
      - 2.3.1.1 Steady state daily workload assumptions ..... 13
      - 2.3.1.2 Steady state incremental backup workload calculation ..... 14
      - 2.3.1.3 Steady state full backup workload calculation ..... 14
      - 2.3.1.4 Steady state full restore workload calculation ..... 14

Note: Daily restore workload considers the VM image restores that occur during the normal backup window. Additional capacity is available for VM image restores outside of the backup window and not included in this calculation. .... 14
    - 2.3.2 Calculate initial full backup phase-in workload ..... 15
      - 2.3.2.1 Initial full backup phase-in calculation: assumptions ..... 15
      - 2.3.2.2 Calculate initial full backup phase-in workload ..... 15
    - 2.3.3 Calculate peak VM image restore workload ..... 16
      - 2.3.3.1 Calculate peak VM image restore workload: assumptions ..... 16
      - 2.3.3.2 Calculate peak VM image restore workload for critical VMs ..... 16
    - 2.3.4 Non-deduplicated data example with initial phase-in and peak restores ..... 16
      - 2.3.4.1 Calculate number of concurrent sessions required: non-deduplicated data ..... 17
      - 2.3.4.2 Calculate number of proxy hosts required: non-deduplicated data ..... 17
    - 2.3.5 Non-deduplicated data example excluding initial phase-in and peak restores ..... 18
      - 2.3.5.1 Calculate number of concurrent sessions required: non-deduplicated data ..... 19
      - 2.3.5.2 Calculate number of proxy hosts required: non-deduplicated data ..... 19
    - 2.3.6 Data deduplication example with initial phase-in and peak restores ..... 20

2.3.6.1 Calculate number of concurrent sessions required: data deduplication .....	20
2.3.6.2 Calculate number of proxy hosts required (data deduplication) .....	21
2.3.6.2.1 Calculate number of proxy hosts required by CPU requirements .....	21
2.3.6.2.2 Calculate number of proxy hosts required by throughput (data deduplication)...	21
2.3.7 Data deduplication example excluding initial phase-in and peak Restores .....	22
2.3.7.1 Calculate number of concurrent sessions required: data deduplication .....	22
.....	23
2.3.7.2 Calculate number of proxy hosts required by CPU (data deduplication) .....	23
2.3.7.3 Calculate number of proxy hosts required by throughput (data deduplication) .....	23
2.4 Constraint checking and architectural considerations .....	23
2.4.1 Check for constraints .....	24
2.4.2 Additional capacity requirements .....	24
2.4.3 Physical or virtual proxy? .....	25
2.4.3.1 Questions to ask when you decide between a physical and a virtual proxy .....	25
2.4.3.2 ESX clusters and distribution of virtual proxies .....	25
<b>3. Proxy host resource requirements .....</b>	<b>27</b>
3.1 Determining proxy resource requirements .....	27
3.1.1 Determining I/O resource requirements .....	27
3.1.2 Determining CPU requirements .....	27
3.1.3 Memory estimation .....	28
4. Simplified estimation .....	29
4.1 Non-deduplicated data example .....	29
4.1.1 Including initial phase-in and peak restore workloads .....	29
4.1.2 Excluding initial phase-in and peak restore workloads .....	30
4.2 Data deduplication example .....	30
4.2.1 Including initial phase-in and peak restore workloads .....	30
4.2.2 Excluding initial phase-in and peak restore workloads .....	31

---

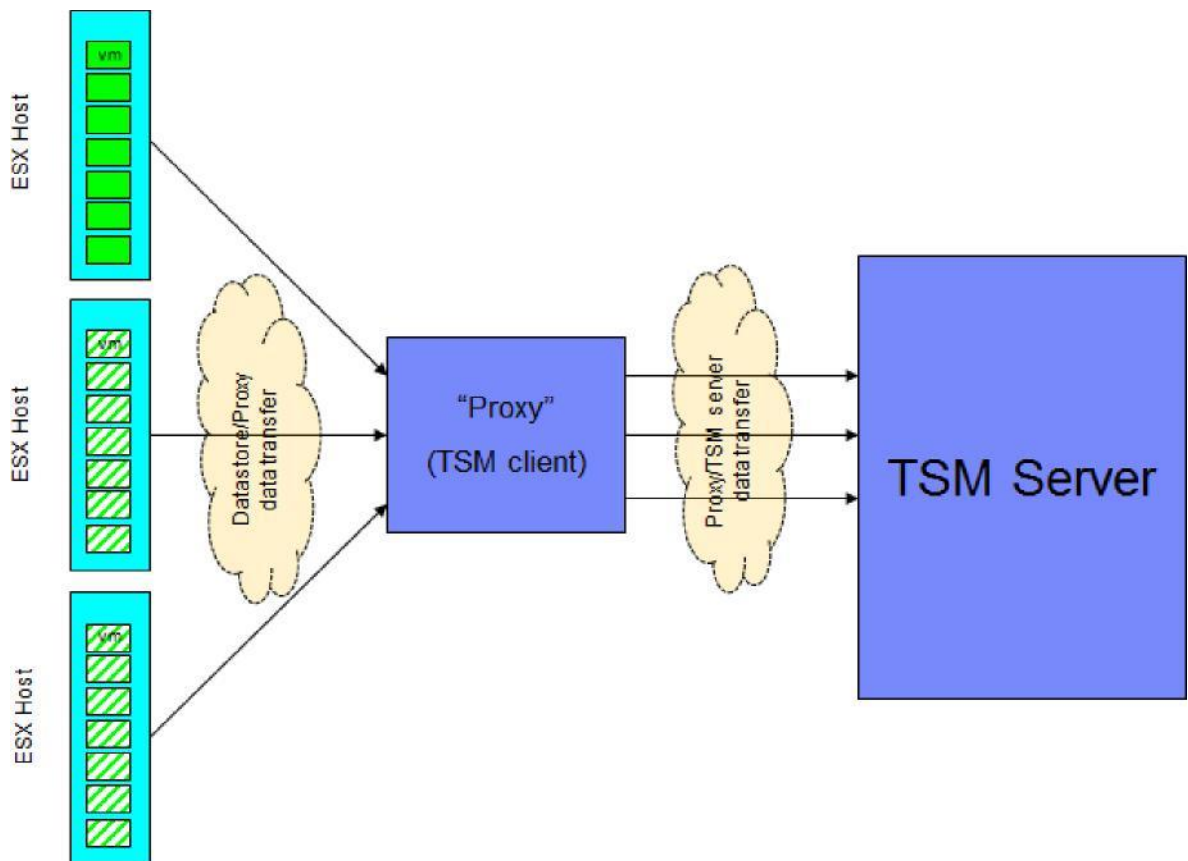
# 1. Introduction

Tivoli Storage Manager Data Protection for VMware is a feature of the Tivoli Storage Manager product family for backing up virtual machines in a VMware (vSphere) environment that provides for the ability to protect the data within VMware deployments (“DP for VMware”). DP for VMware uses the latest backup technology provided by VMware, called vStorage API (also known as “vStorage APIs for Data Protection” or “VADP”).

An essential component of Tivoli Storage Manager for Virtual Environments is the VStorage Backup Server which performs the data transfer from the VMware’s ESX datastores containing the virtual machine data, to the Tivoli Storage Manager server. The VStorage Backup Server is also referred to as the “VBS” and “proxy”. Throughout this document, the VStorage Backup Server will be referred to as the “proxy”. A proxy that is configured on a virtual machine is referred to as a “virtual proxy”, and if configured on a physical machine is referred to as a “physical proxy”. The advantage of the physical proxy is that it completely offloads the backup workload from the ESX server and acts as a proxy for a backup, whereas the virtual proxy exists within the ESX environment where the workload is retained. In either case the backup workload is removed from the virtual machine that is being backed up.

When you consider a backup solution using Tivoli Storage Manager for Virtual Environments, one of the frequently asked questions is how to estimate the number of proxies required for a specific environment. This document guides you through the estimation process.

The following diagram provides a simplified, high level overview of the components involved with DP for VMware image backup and restore:





## 1.1 Overview

The proxy estimation method described in this document is intended to help you plan a deployment of DP for VMware. A suggested approach is described. DP for VMware provides flexibility for deploying the proxies and selecting virtual, physical, or a combination of both proxies. . The intent of this document is to provide a starting point for quantifying the specification and quantity of proxy servers within the solution architecture.

When deciding between virtual and physical proxies the choice is primarily one of allocating the workload of taking and restoring backup copies of data. The secondary decision is how that data is transferred to and from the appropriate storage. If there is a desire to remove the workload from the ESX server then the proxy should be physical, if the ESX server has available resource then a virtual proxy may be appropriate. Irrespective of the workload, if it is the case that the backup copies of data should be sent directly to the disk and tape storage, without passing over the LAN or through the Tivoli Storage Manager backup server, then a physical proxy is required, if the data can flow over the LAN and through the backup server then again a virtual proxy may be acceptable.

It is imperative to understand that the proxy sizing is the same for virtual and physical proxies. The number and size of the proxies is the same in both cases and in both cases the system capability (hardware capability) governs the number and size of the proxies required.

Beginning with DP for VMware V6.4, the scheduling of VMware backups is simplified through two major enhancements:

**Multi-session datamover capability:** This feature allows a single DP for VMware schedule instance to create multiple, concurrent backup sessions. Using this feature significantly reduces the number of schedules that must be set up to backup an enterprise environment.

**Progressive Incremental backups:** This feature eliminates the need for periodic full VM backups. Following the initial full backup, all subsequent backups can be incremental. With this new feature, the amount of data that must be backed up on a daily basis is greatly reduced.

The use of the proven progressive incremental backup approach means that the amount of data that is backed up on a daily basis is relatively small; it corresponds directly to the daily change or growth in source data. With the small amount of data normally generated with incremental backups, it becomes more important to consider factors other than daily backups when sizing the proxy capacity, such as image restores, unplanned full backups, and new VM's that require full backups. The estimation process discussed in this document includes these other factors in the estimation. The proxy estimation process comprises the following steps which are described in detail in Chapter 2, Step by Step Proxy Sizing:

Calculate the workload for the steady-state daily processing. This includes daily incremental backups, occasional full backups due to "CBT resets", and expected daily image restores. CBT resets ("Change Block Tracking") result from the occasional loss of VMware maintained change tracking that can occur from various causes such as the deletion and restore of a VM. Incremental backups cannot be performed without the CBT information.

Calculate the workload for the initial phase-in of full backups. This applies to new environments that have not been backed up by DP for VMware previously. Although this factor is considered a onetime process, it is included in the estimate to provide a contingency. It may optionally be omitted from the estimate if a less conservative estimate is desired.

Calculate peak load restore workload. This workload provides a contingency for restore of critical VM images in the case of the loss of multiple VM's or ESXi hosts. It is suggested to include this workload in the estimate, although it may optionally be omitted.

Calculate the total number of concurrent sessions required to support the required workload.

Estimate the number of proxy hosts required based on CPU and throughput capabilities.

Check for any constraints in the environment based on the assumptions used in the estimate.



### 1.1.1 Backup performance and incremental forever

Estimating the number of proxies requires some assumptions about the performance characteristics of individual backup sessions. DP for VMware uses efficient disk block-level I/O for the backup process, and the data transfer portion of the backup process itself consumes minimal CPU and memory resources.

Beginning with DP for VMware V6.4 the factors involved in backup performance are changed significantly due to the use of incremental forever. With previous versions of DP for VMware, the backup workload was defined predominantly by the full backup component, even though it only occurred weekly. With DP for VMware V6.4 and later, the predominant factor in backups is the incremental workload.

With an incremental backup, the time required for transfer of the backup data can often be the same or less as the time required for the start and stop of the backup due to operations such as the VM snapshot creation and removal. As a result, the estimated throughput performance used for incremental backups is lower than that used for full backups. This is because the estimated throughput value for incremental backups is largely driven by the start/stop time of the backup and much less by the actual data transfer. Nevertheless, the use of incremental forever provides a substantial improvement in efficiency over backup methods that require a period full or rely exclusively on data deduplication to perform data reduction.

### 1.1.2 Backup performance and full backups

Relative to overall aggregate performance, the performance of full VM image backups are not affected as significantly by start/stop overhead, as is the case with incremental backups. Therefore when estimating the throughput of full backups, different throughput estimates are used. Although incremental backups are still affected by environment and system characteristics, these factors have the most impact on full backups. Following lists some of the most significant factors:

- I/O capabilities of the datastore storage arrays

- Back-end storage device used by the Tivoli Storage Manager server, for example, Virtual Tape Library (VTL) or disk

- Infrastructure connectivity, for example, storage area network (SAN) or local area network (LAN) bandwidth

- The use of Tivoli Storage Manager data deduplication.

It is suggested that you use benchmarking to refine the estimate of backup throughput specific to your environment.

#### 1.1.2.1 Data deduplication

Tivoli Storage Manager client-side (*inline*) data deduplication is highly effective with Tivoli Storage Manager for Virtual Environments and can substantially reduce back-end storage requirements as well as the proxy to Tivoli Storage Manager server connection bandwidth requirements. Client-side data deduplication requires additional processing (by the proxy host) that will slow the backup throughput. For a specific amount of data to backup, you may require more backup sessions to meet a given backup window when using data deduplication as compared with not using data deduplication. For estimation purposes, you can assume that backup throughput when you use client data deduplication is approximately 50% of the throughput without data deduplication. However, when using a properly configured proxy (e.g., sufficient CPU resources), the total number of proxies may be the same with or without data deduplication.

As an alternative to using client-side (*inline*) data deduplication, Tivoli Storage Manager server-side (*post-process*) data deduplication may be used if backup throughput requirements are the highest priority, and proxy to Tivoli Storage Manager server connection bandwidth is not constrained.

Tivoli Storage Manager Client-side compression can work with client-side data deduplication (since the duplicated chunks are identified before the client compresses the data). However, do not use client-side compression

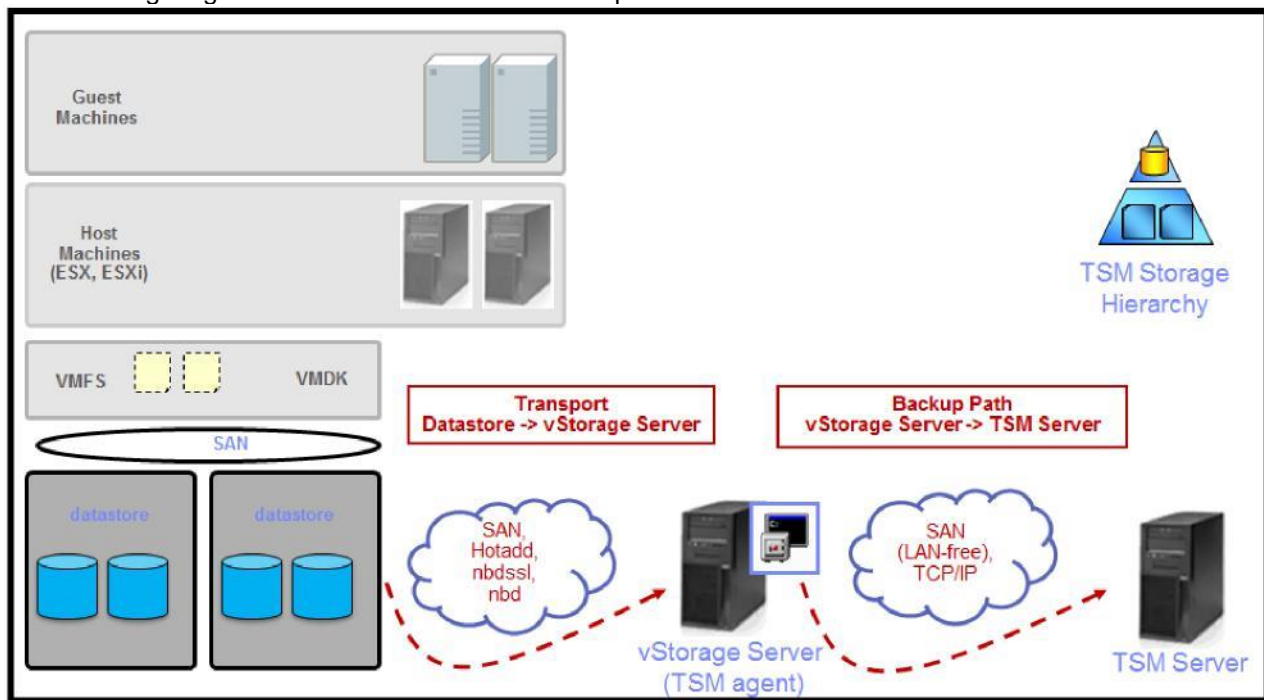
when performing server-side data deduplication as this will negatively impact the effectiveness of server-side data deduplication (client-side compression will obscure the data and make finding the duplicate chunks harder).

Remember that Tivoli Storage Manager native data deduplication requires the use of a Tivoli Storage Manager sequential disk pool (FILE device class). Do not attempt to use Tivoli Storage Manager native data deduplication when the target storage pool is virtual or physical tape. It is not suggested to use Tivoli Storage Manager client-side compression if the target device will try to compress or deduplicate the data.

### 1.1.2.2 Data transfer/transport methods

The methods used for data transfer from datastore to proxy and from proxy to the Tivoli Storage Manager server can have an impact on the per-session performance of Tivoli Storage Manager for Virtual Environments backups and restores.

The following diagram shows the data transfer/transport methods:



The following information about the methods available is listed here for reference. Tivoli Storage Manager user documentation should be referenced for more details on the methods available and how to configure.

Table 1

Data I/O from ESX datastore to proxy			
Transport method	Available to virtual proxy?	Available to physical proxy?	Comments
NBD	Yes	Yes	Uses LAN connection
NBDSSL	Yes	Yes	Uses LAN connection. Slower than NBD due the encryption processing.
SAN	No	Yes	Uses direct SAN connection to datastore (for SAN-attached datastores only).
HOTADD	Yes	No	Uses SAN connection (via ESX host) for SAN-attached volumes which is nearly as efficient as the "SAN" transport. For NFS datastores, provides more efficient transport than NBD.





Table 2

Data I/O from proxy to Tivoli Storage Manager server			
Communication method	Available to virtual proxy?	Available to physical proxy?	Comments
LAN	Yes	Yes	Data transfers over LAN to Tivoli Storage Manager server
LAN-free	No	Yes	Data transfers over SAN to Tivoli Storage Manager server storage pool devices (Tape or Virtual Tape) <b>Notes:</b> <ul style="list-style-type: none"><li>• LAN-free with disk is possible using SANergy or GPFS.</li><li>• LAN-free cannot be used with client-side data deduplication.</li></ul>

## 1.2 A few definitions

Term	Definition
Proxy	The host that performs the offloaded backup. This host can be a virtual or physical machine. Also called “vStorage Backup Server” (VBS). The Tivoli Storage Manager Backup/Archive Client is installed on this host and provides the VMware backup function.
Datamover	An individual backup process that performs the VMware guest backups. Each datamover is associated with one or more Tivoli Storage Manager backup schedules. With DP for VMware V6.4 and later a single datamover instance can create multiple, concurrent backup sessions. Each proxy host will have at least one datamover configured.

## 1.3 Scope of this document

This document is intended to provide a first order estimate of the quantity of proxy hosts required for a specific Tivoli Storage Manager for Virtual Environments backup environment. Using these guidelines can help to provide a successful deployment by establishing a quantitative basis for determining the quantity, placement, and sizing of the proxy hosts. There are many assumptions made within this document and actual results can vary significantly depending upon the environment and infrastructure characteristics. Careful evaluation of the environment is necessary and benchmarking during the planning phase is strongly encouraged to characterize the capabilities of the environment.



### **1.3.1 What is not covered by this document**

The following topics are outside the scope of this document:

Tivoli Storage Manager Server sizing

Tivoli Storage Manager Server storage pool device selection and configuration

Performance analysis and troubleshooting

Tivoli Storage Manager DP for VMware solution architecture

Backup scheduling. This topic is covered in another document that can be found at:

<https://www.ibm.com/developerworks/mydeveloperworks/wikis/home?lang=en#/wiki/Tivoli%20Storage%20Manager/page/Recommendations%20for%20Scheduling%20with%20Tivoli%20Storage%20Manager%20for%20Virtual%20Environments>

### **1.3.2 External dependencies and assumptions**

The estimation process in this document is based on the assumption that no constraints exist in the environment, and that storage capacity per VM, ESX host, and cluster are consistent across the environment. If there are large disparities in storage capacity for individual VM's, for example, a separate sizing may need to be done for the largest VM's. The main assumptions are:

1. Datastore I/O capacity is sufficient to sustain the backup workloads.
2. LAN or SAN connectivity is sufficient to sustain the backup workloads.
3. Tivoli Storage Manager server and storage device capacity and the number of instances are designed and configured to sustain the backup workload provided by the proxies.

After completing an initial sizing estimate, it is important to validate the assumptions and other constraints that may exist in the environment. For example, more than one Tivoli Storage Manager server may be required to manage the workload to achieve the overall throughput requirement. Tivoli Storage Manager server sizing is beyond the scope of this document.

### **1.3.3 Proxy hardware configuration**

Although some guidelines are provided for proxy host resource requirements, it is not the intent of this document to provide specific guidance on hardware or system configurations of physical or virtual proxy hosts. Hardware configuration (or in the case of a virtual machine, resource allocation) should be defined by a qualified system engineer that is familiar with hardware capabilities, I/O throughput, and other system requirements.

IBM Techline provides a service for pre-sales configuration of Tivoli Storage Manager hardware including Tivoli Storage Manager for Virtual Environments proxy sizing. Consult your IBM Tivoli sales representative for more information.

## **1.4 Scheduling of backups**

The estimation technique described in this document is based on the use of incremental forever backup. The only exceptions to the use of incremental forever are for VM's that have not been backed up by DP for VMware, newly created virtual machines, and virtual machines that have had a reset of CBT (change block tracking). DP for VMware V6.4 and later uses a highly efficient technology for tracking and grouping changed blocks



without the requirement for synthetic reconstruction of a VM image backup version. Therefore there is generally no advantage to using any other backup technique.



## 2. Step-by-step proxy sizing

This section provides the steps for sizing a proxy for a specific deployment scenario across an entire data center. The same approach may be applied individually to separate environments that may have different characteristics, such as the average VM size.

### 2.1 Assumptions

Reasonably equal distribution (within +- 20%) of utilized virtual machine storage capacity (datastores) across all VM's.

Incremental backups are used for all backups except for initial phase-in of new VM's and a small percentage of VM's that require full backup due to CBT reset.

### 2.2 Example environment

The following example environment is used to illustrate the estimation process:

Table 3

Environment description		
Total number of virtual machines		1000
Average used storage per VM		50 GB
Total used storage	1000 * 50 GB	50,000 GB
Number of ESX hosts		25
Number of DRS clusters		5
Backup window		10 Hours
Assumed daily change rate		2%
For new environment: Initial full backup phase-in time (total days).		25 days (during normal daily backup window)
Peak image restore requirement: % of VM's that are critical.		25%
Peak image restore requirement: RTO for critical VM's		72 hours (continuous, 24x7)

#### 2.2.1 Notes regarding example environment

The example environment represents what is considered a small to medium sized VMware environment. The actual parameters selected are expected to vary depending upon the particular environment, but provide a realistic set of assumptions as a starting point. Following notes provide additional explanation for some of the parameters:



### Daily change rate

A value of 2% daily change rate is used in this example. This value is expected to vary depending upon the environment, and will of course not be exactly the same each day. The 2% represents the average over time. Typical change rates can vary from less than 1% to over 5%.

### Initial full backup phase-in time

For this example, a value of 25 days is used. This value will be driven by business needs. However it is suggested to always plan for a phase-in period that results in a backup workload (in terms of aggregate data transfer requirements) that is somewhat consistent with the workload of ongoing, daily backups. Planning such a phase-in period provides these advantages:

- Reduces the need to establish additional, temporary proxy hosts.
- Allows observation of actual throughput performance in the environment to validate the initial estimate, and allow for adjustments in the required number of backup sessions and proxy hosts. If the initial phase-in workload estimate is included in the proxy sizing, the full backup phase-in can occur concurrently with daily backups of VM's that have already had the initial full backup.

## 2.3 Perform the estimate

The following sections demonstrate the estimation process.

### 2.3.1 Calculate steady state daily workload

The steady-state workload is the basis for the proxy sizing and is intended to describe the normal, expected workload (both backup and restore). The steady-state workload consists of the following components:

- Daily incremental backups of all previously backed up VM's.
- Full backups of newly created VM's each day.
- Full backups of VM's impacted by CBT Reset. This refers to VM's that have been previously backed up but the VMware Change Block Tracking (CBT) data has been lost. A small percentage of the total VM's should be assumed to have this occur on a daily or weekly basis.
- Daily full VM image restores that are expected during the normal backup window. Note that since the estimate is based on the normal backup window this component only refers to the restores that occur within the backup window and does not include restores which could occur outside of the backup window. File-level restores are not included in this estimate since they can be done without the use of the proxy host. If it is desired to perform file-level restores from the proxy host, then additional resources may need to be configured on the proxy. This scenario is not covered in this document.

#### 2.3.1.1 Steady state daily workload assumptions

Table 4

Steady state daily workload assumptions	
<b>Note:</b> For illustrative purposes only. You should choose values appropriate to your operational environment.	
Daily Change Rate (from Table 3)	2% of total utilized disk
Full Backups of daily, newly created VM's	1% of total original VM's
Full Backups due to CBT reset	1% of total original VM's
Daily full VM image restores during backup window	1% of total original VM's



### 2.3.1.2 Steady state incremental backup workload calculation

Table 5

<b>Example</b>		
Daily Incremental Backup Workload	50000 GB * 2%	1000 GB
Aggregate Throughput requirement per backup window for Incrementals	1000 GB / 10 hours	100 GB/hour

### 2.3.1.3 Steady state full backup workload calculation

Table 6

<b>Example</b>		
Full Backups of new VM's	50000 GB * 1%	500 GB
Full Backups due to CBT Reset	50000 GB * 1%	500 GB
Total Steady State Daily Full Backup Workload		1000 GB
Aggregate Throughput requirement per backup window for full backups	1000 GB / 10 hours	100 GB/hour

### 2.3.1.4 Steady state full restore workload calculation

**Note:** Daily restore workload considers the VM image restores that occur during the normal backup window. Additional capacity is available for VM image restores outside of the backup window and not included in this calculation.

Table 7

<b>Example</b>		
Daily full VM image restores	50000 GB * 1%	500 GB
Aggregate Daily Restore Throughput requirement per backup window	500 GB / 10 hours	50 GB/hour



### 2.3.2 Calculate initial full backup phase-in workload

The initial full backup phase-in workload is intended to provide an estimate for data throughput requirements when backing up a VMware environment that had not been previously backed up by Tivoli Storage Manager. Although this workload can be considered part of an initial, one time only process, it is included in this estimate to provide additional contingency. This component may be omitted from the overall steady-state estimate if desired. Proxies may be temporarily provisioned to handle the initial phase-in process, particular if an aggressive phase-in schedule is desired.

#### 2.3.2.1 Initial full backup phase-in calculation: assumptions

Table 8

Initial full backup phase-in calculation assumptions	
Daily backup window during phase-in	25 days
Total days allowed for phase-in	10 hours

#### 2.3.2.2 Calculate initial full backup phase-in workload

Table 9

Steady state full backup workload calculation		
Amount of data to backup daily: full backups during the initial phase-in	50000 GB / 25 days	2000 GB/Day
Aggregate Throughput Requirement	2000 GB / 10 hours	200 GB/Hour



### 2.3.3 Calculate peak VM image restore workload

#### 2.3.3.1 Calculate peak VM image restore workload: assumptions

**Note:** The recovery time for this workload is assumed to run continuously, and thus will overlap the normal daily backup window.

Table 10

Example	
Percent of all VM's that are critical for immediate recovery	25%
Recovery Time Objective for critical VM's*	72 hours
*Assumes continuous, 24x7 recovery window	

#### 2.3.3.2 Calculate peak VM image restore workload for critical VMs

Table 11

Example		
Amount of data to restore, for the critical VM's	50000 * 25%	12500 GB
Aggregate Throughput Requirement	12500 GB / 72 hours	173 GB/hour

### 2.3.4 Non-deduplicated data example with initial phase-in and peak restores

Separate example calculations are provided for the case where Tivoli Storage Manager client-side data deduplication is not used, and for the case where it is used. Backup without Tivoli Storage Manager client-side data deduplication uses minimal CPU resources, whereas the use of client-side data deduplication uses CPU resources on the proxy to perform the data deduplication and compression (compression does not need to be used with client-side data deduplication, but is usually beneficial). The throughput characteristics are different between the non-data deduplication and client-side data deduplication cases. Note that Tivoli Storage Manager server-side data deduplication should be treated the same as the non-deduplicated data case for proxy estimation.

This section covers the non-deduplicated data example.





### 2.3.4.1 Calculate number of concurrent sessions required: non-deduplicated data

Table 12

<b>Example</b>			
<b>Component</b>	<b>Required aggregate throughput</b>	<b>*Estimated per-process throughput</b>	<b>Required number of sessions</b>
Steady-state Incr backup workload	100 GB/hour	20 GB/hour	5
Steady-state Full backup workload	100 GB/hour	40 GB/hour	2.5
Steady-state daily restores	50 GB/hour	40 GB/hour	1.3
Initial Full Backup phase-in	200 GB/hour	40 GB/hour	5
Peak VM Image Restore Workload	173 GB/hour	40 GB/hour	4.3 (rounded)
<b>Total</b>	<b>623 GB/Hour</b>		<b>18.1</b>

**Note:** Throughput numbers are provided for illustrative purposes for this example calculation, and depend upon a number of environment specific factors. Benchmarking is recommended to determine actual values to use for planning.

### 2.3.4.2 Calculate number of proxy hosts required: non-deduplicated data

The determination of the required number of proxy hosts is based primarily on two factors: CPU requirements and throughput requirements. When Tivoli Storage Manager data deduplication is not used, the required number of hosts is usually driven by the throughput requirement since the CPU requirement is minimal. The following assumptions are used for this calculation:

Proxy host is a 4 CPU core (e.g., single socket dual core or single socket quad core)

Although CPU requirements are minimal when data deduplication is not used, we use a guideline of 0.1 CPU core per session.

A conservative estimate of 350 GB/Hour maximum per proxy host is used. Greater throughput is possible with a properly configured host with appropriate adapter cards.



Table 13

<b>Calculate number of proxy hosts required by CPU (non-duplicated data)</b>		
Number of concurrent sessions required from Table 12		18.1
Estimated CPUs per session		0.1
Required number of CPU's	18.1 sessions * 0.1 CPU/session	1.81
CPU's per proxy		4
Number of proxy hosts required	1.81 CPU's / 4 CPU's per host	0.45 hosts

Table 14

<b>Calculate number of proxy hosts required by throughput (non-duplicated data)</b>		
Total aggregate throughput requirement from table 10		623 GB/Hour
Estimated maximum throughput per proxy host		350 GB/Hour
Number of proxy hosts required	623 GB/Hour / 350 GB/Hour/Host	2 (rounded up)

In this example, two proxy hosts are required, based on the throughput requirement. This is usually the case when client-side data deduplication is not used.

### **2.3.5 Non-duplicated data example excluding initial phase-in and peak restores**

The prior example included the contingency for the initial full backup phase-in and peak load VM image restores. This provides a more conservative estimate. A less conservative estimate may be made by excluding these components from the estimate. However, it is advisable to always consider how the initial phase-in will be accommodated as well as unexpected, mass restores or VM images that might be required as a result of the loss of multiple ESX hosts or other infrastructure.



### 2.3.5.1 Calculate number of concurrent sessions required: non-deduplicated data

Table 15

<b>Example</b>			
<b>Component</b>	<b>Required aggregate throughput</b>	<b>Estimated per-session throughput</b>	<b>Required number of sessions</b>
Steady-state Incr backup workload	100 GB/hour	20 GB/hour	5
Steady-state Full backup workload	100 GB/hour	40 GB/hour	2.5
Steady-state daily restores	50 GB/hour	40 GB/hour	1.3
<b>Total</b>	<b>250 GB/Hour</b>		<b>8.8</b>

### 2.3.5.2 Calculate number of proxy hosts required: non-deduplicated data

Table 16

<b>Calculate number of proxy hosts required by CPU (non-deduplicated data)</b>		
Number of concurrent sessions required from Table 15		8.8
Estimated CPUs per session		0.1
Required number of CPU's	8.1 sessions * 0.1 CPU/session	0.8
CPU's per proxy		4
Number of proxy hosts required	0.8 CPU's / 4 CPU's per host	0.2 hosts

Table 17

<b>Calculate number of proxy hosts required by throughput (non-deduplicated data)</b>		
Total aggregate throughput requirement from Table 15		250 GB/Hour
Estimated maximum throughput per proxy host		350 GB/Hour
Number of proxy hosts required	250 GB/Hour / 350 GB/Hour/Host	1 (rounded up)



In this example, one proxy host is required. However, you should consider configuring a minimum of two proxies and distribute the VM backups across the proxies to reduce the risk associated with a single point of failure for backup operations.

### 2.3.6 Data deduplication example with initial phase-in and peak restores

This section considers the case where Tivoli Storage Manager client-side data deduplication is used. For the data deduplication scenario the following assumptions apply:

- Proxy host has 8 CPU cores (e.g., dual socket quad core). Since client-side data deduplication requires CPU resources, it is advisable to configure more CPU resources than with the non-deduplicated data case. Alternatively, fewer CPU cores can be configured and more proxy hosts can be used.
- When using Tivoli Storage Manager client-side data deduplication, the CPU requirements for the proxy is roughly proportional to throughput. Therefore, the CPU requirement for each session performing an incremental backup would be less than a session performing a full backup since the throughput of the incremental session will be less. For the purpose of the estimate, use the following values:
- Backup with client-side data deduplication uses approximately 0.025 CPU cores (2.2Ghz processor) per GB/Hour of throughput. For example, a single, 10GB/hour backup session will consume  $10 * 0.025 = 0.25$  CPU cores. The CPU estimation is based on throughput regardless of whether the backup is incremental or full.
- Although the restore process throughput does not strongly correlate to CPU utilization, we need to factor the restore workload on the CPU requirement. We use an estimate of 0.01 CPU core per GB/hour of throughput. For example, a single, 20GB/hour restore session will consume  $20 * 0.01 = 0.1$  CPU cores.
- A conservative estimate of 350 GB/Hour maximum per proxy host is used. Much greater throughput is possible with a properly configured host with appropriate adapter cards.

#### 2.3.6.1 Calculate number of concurrent sessions required: data deduplication

Table 18

Calculate number of concurrent sessions required: data deduplication example			
Component	Required aggregate throughput	Estimated per-session throughput*	Required number of sessions
Steady-state incremental backup workload	100 GB/hour	10 GB/hour	10
Steady-state full backup workload	100 GB/hour	20 GB/hour	5
Steady-state daily restores	50 GB/hour	20 GB/hour	2.5
Initial full backup phase-in	200 GB/hour	20 GB/hour	10
Peak VM image restore workload	173 GB/hour	20 GB/hour	8.7 (rounded)
<b>Total</b>	<b>623 GB/Hour</b>		<b>36.2</b>

**Note:** Throughput numbers are provided for illustrative purposes for this example calculation, and depend upon a number of environment specific factors. Benchmarking is recommended to determine actual values to use for planning.



### 2.3.6.2 Calculate number of proxy hosts required (data deduplication)

The number of proxies is calculated based on both factors of CPU and throughput. When using client-side data deduplication, CPU resources are critical factor in determining the number of proxies.

#### 2.3.6.2.1 Calculate number of proxy hosts required by CPU requirements

Table 19

Calculate number of proxy hosts required by CPU (data deduplication)				
Component	Required aggregate throughput	CPU cores required per GB/Hour	CPU cores required calculation	Required number of CPU cores
Steady-state incremental backup workload	100 GB/hour	0.025	100 * 0.025	2.5
Steady-state full backup workload	100 GB/hour	0.025	100 * 0.025	2.5
Steady-state daily restores	50 GB/hour	0.01	50 * 0.01	0.5
Initial full backup phase-in	200 GB/hour	0.025	200 * 0.025	5.0
Peak VM image restore workload	173 GB/hour	0.01	173 * 0.01	1.7
<b>Total</b>				<b>12.2</b>

If we assume an 8-core (e.g., dual socket 4-core) configuration, in this example we require two proxies based on the CPU requirement of 12.2 cores ( $12.2/8 = 2$  rounded up).

#### 2.3.6.2.2 Calculate number of proxy hosts required by throughput (data deduplication)

Table 20

Example		
Total aggregate throughput requirement from Table 18		623 GB/Hour
Estimated maximum throughput per proxy host		350 GB/Hour
Number of proxy hosts required	623 GB/Hour / 350 GB/Hour/Host	2 (rounded up)

In this example, two proxies are required by both the CPU and throughput calculations.



### 2.3.7 Data deduplication example excluding initial phase-in and peak Restores

#### 2.3.7.1 Calculate number of concurrent sessions required: data deduplication

Table 21

<b>Example</b>			
<b>Component</b>	<b>Required aggregate throughput</b>	<b>Estimated per-process throughput</b>	<b>Required number of sessions</b>
Steady-state incr backup workload	100 GB/hour	10 GB/hour	10
Steady-state full backup workload	100 GB/hour	20 GB/hour	5
Steady-state daily restores	50 GB/hour	20 GB/hour	2.5
<b>Total</b>	250 GB/Hour		17.5



### 2.3.7.2 Calculate number of proxy hosts required by CPU (data deduplication)

Table 22

<b>Example</b>				
<b>Component</b>	<b>Required aggregate throughput</b>	<b>CPU cores required per GB/Hour</b>	<b>CPU cores required calculation</b>	<b>Required number of CPU cores</b>
Steady-state incr backup workload	100 GB/hour	0.025	100 * 0.025	2.5
Steady-state full backup workload	100 GB/hour	0.025	100 * 0.025	2.5
Steady-state daily restores	50 GB/hour	0.01	50 * 0.01	0.5
<b>Total</b>				<b>5.5</b>

Based on CPU requirements, one proxy is sufficient for the workload.

### 2.3.7.3 Calculate number of proxy hosts required by throughput (data deduplication)

Table 23

<b>Example</b>		
Total aggregate throughput requirement from Table 21		250 GB/Hour
Estimated maximum throughput per proxy host		350 GB/Hour
Number of proxy hosts required	250 GB/Hour / 350 GB/Hour/Host	1 (rounded up)

In this example, one proxy is required, based on both CPU and throughput calculation. However, you should consider configuring a minimum of two proxies and distribute the VM backups across the proxies to reduce the risk associated with a single point of failure for backup operations.

## 2.4 Constraint checking and architectural considerations

Now that we have an estimate of the number of proxies required to achieve the daily backup workload, we must now consider whether this makes sense in practical terms. DP for VMware provides a great deal of flexibility in deployment options, so we need to determine which options makes the most sense. We will consider other factors and determine if adjustments are required.



### 2.4.1 Check for constraints

The estimation method for proxy sizing will provide an arithmetically accurate result based on the accuracy of the underlying assumptions of the capacity of data to back up, assumed throughput, and even distribution of utilized storage capacity across all virtual machines and hosts. However, there are several constraints that should be validated to ensure that the resulting estimate is practical and feasible for the environment. The following table lists some of the constraint conditions that should be validated, although there may be other conditions unique to the environment that are not listed below.

Constraint	Validation
Proxy i/o throughput: Can each individual proxy sustain the required i/o throughput?	Ensure that the proxy can be configured with sufficient adapter cards (NICs and HBAs) to support the required throughput.
Tivoli Storage Manager server throughput: Can the Tivoli Storage Manager server support the required aggregate i/o throughput for all of the proxies?	Ensure that the Tivoli Storage Manager server is configured to support the required aggregate throughput. Multiple Tivoli Storage Manager servers may be required in some cases.
Tivoli Storage Manager server sessions: Can the Tivoli Storage Manager server support the required number of concurrent backup sessions from all of the datamover processes?	Ensure that the Tivoli Storage Manager server is configured to support the required number of concurrent backup sessions.
Infrastructure bandwidth: Can the LAN or SAN accommodate the aggregate workload required for all of the backup processes?	Ensure that the LAN and SAN networks have sufficient bandwidth to accommodate the backup (and restore) bandwidth requirements.
Datastore I/O Capacity: Can the ESX datastore accommodate the i/o required to support the required backup throughput?	Ensure that the Datastore I/O devices are capable of supporting the I/O data transfer rates required for the backup processes. The assumption for the per-process backup throughput may need to be adjusted.
Special case VM's with large virtual disks: Are there specific VM's with exceptionally large disks or change rate that would affect the ability to backup within the required window?	Determine if there are special applications or VM's that cannot be backed up as part of the overall schedule that may need their own dedicated backup process.

### 2.4.2 Additional capacity requirements

VM Image restores have been factored into the proxy sizing estimate. You should validate whether additional image restores (or possibly file-level restores) should be included in the estimate.

Additionally, projected growth should be factored into the estimate. If the estimate is based on the current state, then the number of proxies estimated may be insufficient. You should ensure that the estimate is based on near term growth projections.





### 2.4.3 Physical or virtual proxy?

The first consideration is to determine whether physical or virtual proxies will be used. Note that you don't need to decide all one or the other. It is entirely possible to use virtual proxies for part of an environment, and physical proxies for another where it makes sense.

For simplicity, we assume that the number of data movers will be the same whether you use a physical or virtual proxy host. This assumption is based on equivalent resources and connectivity when comparing the physical and virtual hosts. In practice, a virtual proxy will share resources (e.g., i/o adapters, CPU, etc.) with other VM's whereas a physical proxy will have dedicated, static resources.

#### 2.4.3.1 Questions to ask when you decide between a physical and a virtual proxy

Following is a list of questions you should consider when deciding between physical and virtual machines showing which type of proxy would be preferred. If the answer to the question is "Yes" then preference should be given to the type of proxy in the "Yes" column. . If the answer to the question is "No" then preference should be given to the type of proxy indicated in the "No" column.

Question	Yes	No
Do you require backup traffic to flow over the SAN as much as possible? <b>Note:</b> Virtual machine proxies can take advantage of Hotadd data transfers from a SAN datastore to the proxy which primarily uses SAN I/O via the ESX host HBA. However, a virtual machine proxy cannot take advantage of LAN-free data transfers from the proxy to the Tivoli Storage Manager server.	Physical	Virtual
Does your LAN (IP Network) have sufficient bandwidth to accommodate the backup traffic?	Virtual	Physical
Do you want to use LAN-free data transfers from the proxy to the Tivoli Storage Manager server? <b>Note:</b> LAN-free is usually only used with Tape or Virtual Tape backup storage devices.	Physical	Virtual
Do you prefer or require that all new hosts are virtual and not physical machines?	Virtual	Either
Do you want to minimize the number of proxy hosts? <b>Note:</b> The preference is based on the assumption that you will dedicate more resources to a physical proxy than a virtual proxy.	Physical	Virtual
Do you use NFS attached datastores?	Virtual	Either
Is 10Gbit Ethernet connectivity available from the virtual proxy to the Tivoli Storage Manager server?	Virtual	Either

#### 2.4.3.2 ESX clusters and distribution of virtual proxies

When using virtual proxies, it is important to consider the distribution of virtual proxy hosts within the infrastructure. If you are considering the use of virtual proxies, it may be desirable to place a virtual proxy in each ESX cluster. This is because a virtual proxy within an ESX cluster can access datastore storage via Hotadd. Hotadd provides an efficient (and low overhead) data transfer method. For SAN attached storage, the proxy transfers data directly through the ESX host's fibrechannel adapter.



In the example environment, you can use 10 proxies to cover all 10 ESX clusters. You can reduce the number of data movers per proxy to distribute the workload over a greater number of proxies and reduce the resource requirements for each proxy. Using 10 virtual proxies will also provide additional “reserve” capacity for restores which occur during the backup window, in addition to the contingency capacity that may have already been included.



## 3. Proxy host resource requirements

Resource requirements for a proxy are driven by the following key factors:

- I/O data transfer capacity
- CPU capacity

When using Tivoli Storage Manager client-side data deduplication, CPU resource utilization tends to be the dominant factor in the proxy host resource requirements. Without the use of client-side data deduplication, the I/O capabilities of the proxy host will be the dominant factor.

---

### 3.1 Determining proxy resource requirements

For simplicity, it is advisable to select a reference configuration for the proxy host and use the required quantity of the standard configuration to achieve the desired backup window. This section describes the factors to consider when determining the resource requirements.

#### 3.1.1 Determining I/O resource requirements

The examples used in this document establish a throughput capability of 350 GB/Hour per proxy host. This requires that the proxy host be configured with sufficient i/o adapter capacity (virtual or physical fibrechannel and/or network adapters) to sustain this data throughput. Infrastructure must also be able to support this throughput, such as available network bandwidth. Higher throughput configurations can be achieved by proper selection of host adapters.

In the case of LAN data transfers the critical I/O resources for the proxy are the network adapter cards (NICs). In the case of SAN data transfers (which includes SAN transport from the datastore to the proxy and LAN-free from the proxy to the Tivoli Storage Manager server), the FC adapters (HBA's) are the critical I/O resource. When both LAN and SAN transfers are used, then both NICs and HBAs are critical to support the required data transfer rates. An example of this is when LAN transport is used between the datastore and proxy, and LAN-free is used between the proxy and the Tivoli Storage Manager server.

For virtual proxies, the resource requirements apply to virtual adapters and the requirements (number and speed of adapters) will remain the same as a physical proxy. However, dedicating shared resources across an ESX hypervisor may require additional planning and configuration for a virtual machine to ensure that sufficient resources are available to other VM's hosted by the same ESX host.

The NIC and HBA adapters should be sized to ensure adequate capacity to handle the expected i/o data rates through the IP network and SAN, respectively.

Machine backplane i/o capacity must also be considered, but generally is not an issue for a properly configured system.

#### 3.1.2 Determining CPU requirements

For the purpose of this section, "CPU processor core" refers to a 2.2Ghz physical or virtual processor core.

**Note:** *Processor core* refers to a processing unit. For example, a quad-core socket has 4 processor cores.



An advantage of considering virtual proxies is that you can distribute a larger number of proxies with smaller resources to accommodate the same backup workload as fewer, physical proxies, that require more resources each. With virtual proxies, you can also easily deploy and remove proxies to serve specific purposes such as mass image restores or initial full backup phase-in.

You should consider using the following proxy host resources. You can alternatively reduce the resources per proxy host and increase the quantity of proxies to achieve the same result:

Without data deduplication:

2 to 4 CPU cores per proxy

With Tivoli Storage Manager client-side data deduplication:

8 to 16 CPU cores per proxy

### **3.1.3 Memory estimation**

The memory requirements for Tivoli Storage Manager for Virtual Environments backup processes are small and generally not a constraint. The memory should be sized adequately for the proxy host operating system (for example, Windows 2008R2). A minimum of 4GB of RAM should be considered.



## 4. Simplified estimation

This section proposes a simplified estimation process, based on the assumptions and calculations described in the previous sections of this document. The simplified process can be used for an initial estimate prior to a more in depth analysis.

**Note:** The use of the simplified estimation process does not eliminate the need for a more in depth analysis to plan and validate the required number of proxies. It is the responsibility of the system architect, administrator, or user, and reader of this document to validate the assumptions by actual benchmarking, select the appropriate parameters, and perform the required calculations to ensure the accuracy of the estimate.

Using the parameters for the example environment and the calculation methods described earlier in this document, we can determine the number proxies required for various change rates and total storage. Based on the results, a simple guideline can be established as a starting point for proxy sizing.

*As previously stated, it is advisable to always have a minimum of two proxies, even if one proxy is sufficient based on the capacity requirements. This will mitigate the risk to backup operations in the case of a failure of one of the proxies.*

---

### 4.1 Non-deduplicated data example

This example provides a simplified estimate when either Tivoli Storage Manager data deduplication is not used, or Tivoli Storage Manager server-side data deduplication is used. The proxy host for non-deduplicated data is assumed to have 4 CPU cores.

#### 4.1.1 Including initial phase-in and peak restore workloads

Table 24

Proxy sizing including phase-in and peak image restores				
	Daily change rate			
Total capacity in GB	0.02	0.04	0.06	0.08
20000	1	1	1	2
40000	2	2	2	3
60000	3	3	3	4
80000	3	4	4	5
100000	4	5	5	6

Based on the sample calculations and assumptions previously described, table 24 shows the estimated number of proxies for various change rates and storage capacities. If we assume a 4% to 6% daily change rate, we may estimate one proxy required per 20000 GB of source data when including initial phase-in and peak restore workloads.



### 4.1.2 Excluding initial phase-in and peak restore workloads

Table 25

Excluding phase-in and peak image restores				
	Daily change rate			
Total capacity in GB	0.02	0.04	0.06	0.08
35000	1	1	1	2
70000	2	2	2	3
105000	2	3	3	4
140000	3	3	4	5
175000	3	4	5	6

Based on the sample calculations and assumptions previously described, table 25 shows the estimated number of proxies for various change rates and storage capacities. If we assume a 4% to 6% daily change rate, we may estimate one proxy required per 35000 GB of source data when excluding initial phase-in and peak restore workloads.

## 4.2 Data deduplication example

This example provides a simplified estimate when Tivoli Storage Manager client-side data deduplication is used. The proxy host for client-side data deduplication is assumed to have 8 CPU cores

### 4.2.1 Including initial phase-in and peak restore workloads

Table 26

Including phase-in and peak image restores				
	Daily change rate			
Total capacity in GB	0.02	0.04	0.06	0.08
20000	1	1	1	2
40000	2	2	2	3
60000	3	3	3	4
80000	3	4	4	5
100000	4	5	5	6

Based on the sample calculations and assumptions previously described, table 26 shows the estimated number of proxies for various change rates and storage capacities. If we assume a 4% to 6% daily change rate, we may estimate one proxy required per 20000 GB of source data when including initial phase-in and peak restore workloads.



### 4.2.2 Excluding initial phase-in and peak restore workloads

Table 27

Excluding phase-in and peak image restores				
	Daily change rate			
Total capacity in GB	0.02	0.04	0.06	0.08
35000	1	1	1	2
70000	2	2	2	3
105000	2	3	3	4
140000	3	3	4	5
175000	3	4	5	6

Based on the sample calculations and assumptions previously described, table 27 shows the estimated number of proxies for various change rates and storage capacities. If we assume a 4% to 6% daily change rate, we may estimate one proxy required per 35000 GB of source data when excluding initial phase-in and peak restore workloads.



### Notices

This information was developed for products and services offered in the U.S.A. IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing  
IBM Corporation  
North Castle Drive  
Armonk, NY 10504-1785  
U.S.A.

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing  
Legal and Intellectual Property Law  
IBM Japan Ltd  
19-21, Nihonbashi-Hakozakicho, Chuo-ku  
Tokyo 103-8510 Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you. Licensees of this program who want to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation  
2Z4A/101  
11400 Burnet Road  
Austin, TX 78758  
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results





obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

#### COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample liable for any damages arising out of your use of the sample programs. Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows: © (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp  
.\_enter the year or years\_.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.



## Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide.

Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at “Copyright and trademark information” at <http://www.ibm.com/legal/copytrade.shtml>.

Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.