



AUTODESK UNIVERSITY 2013

All Natural Autodesk Vault Enhancements

Irvin Hayes Jr., Autodesk, Inc.

Robert Stein, Autodesk, Inc.

PL3530 Do you have a check-in that lasts more than 4 hours? Trouble reaching peak performance in your Autodesk Vault software environment? Do you want to manage files all day long? Not as excited about files as you wish to be? Then this class is for you. Learn how to make check-ins complete as fast as possible. Learn how to keep your environment operating at its peak and ready to perform when it needs to.

Learning Objectives

At the end of this class, you will be able to:

- List and describe the different hardware configurations
- Find and use the Vault tools to monitor and maintain performance
- Make adjustments to make Vault perform better
- Assess and determine when to scale your environment

About the Speaker

Irvin Hayes Jr. is a product manager for the data management group at Autodesk in Novi, Michigan. Irvin has worked at Autodesk for eight years starting in product support and as a

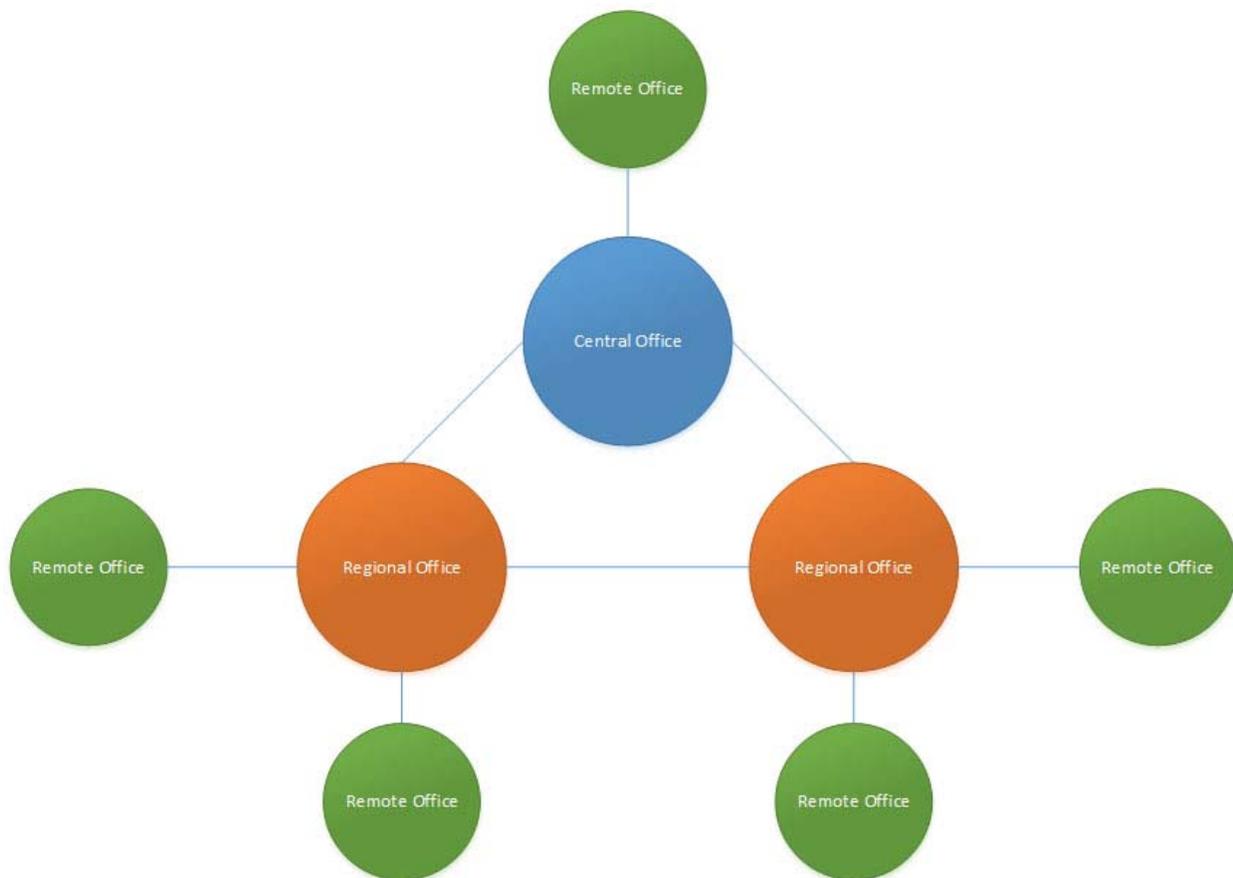
user experience designer. Irvin is a Microsoft® Certified Professional, and has been working in the information technology field for more than 19 years.

Introduction

This class is intended to help the Vault administrator to configure, monitor, and tweak the Vault environment so that it runs optimally. The tools and suggestions below should aid you in this endeavor.

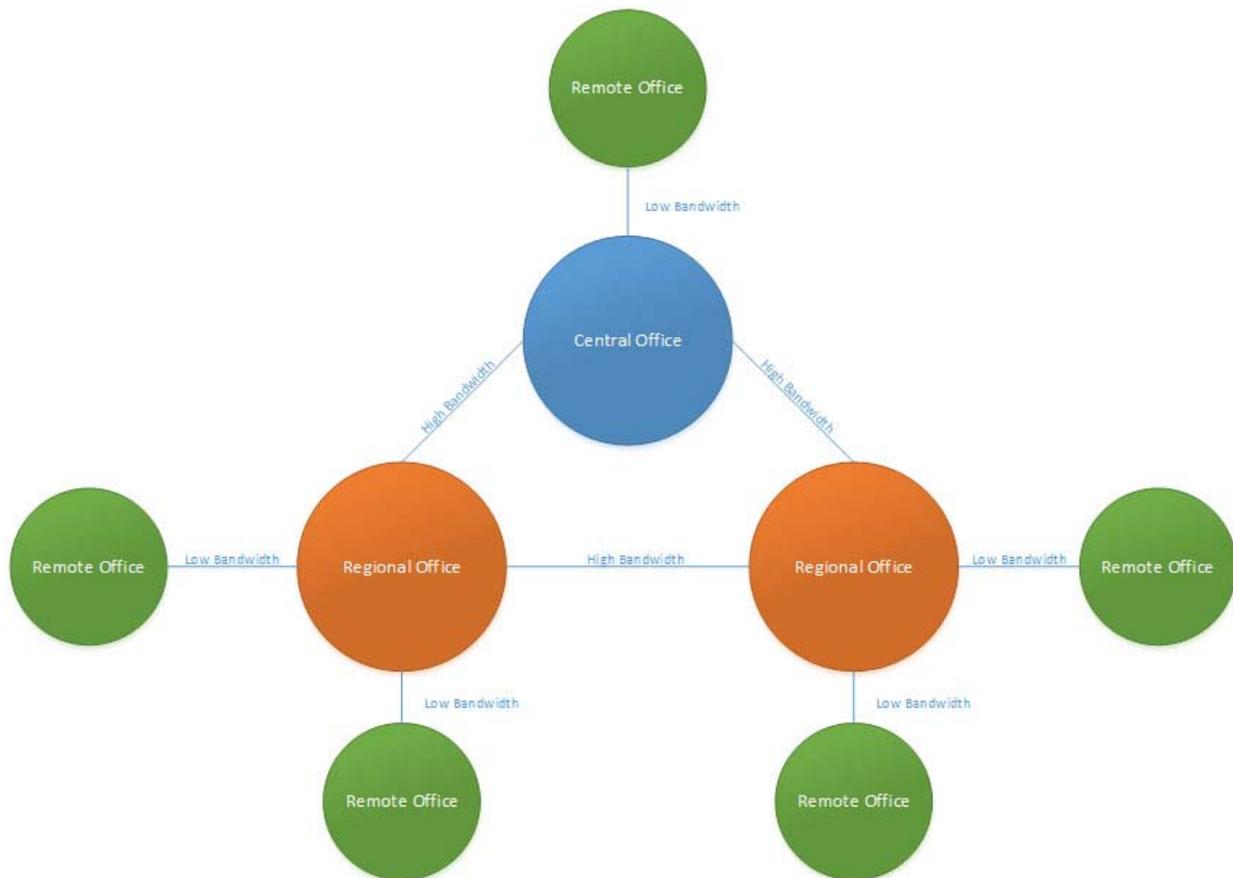
Environments

There is no one-size-fits-all when it comes to network environments. You need to have a very good understanding of your network and a good idea of how it may grow in the future. A good way to start is by building a complete picture of your current network. Examining your infrastructure can help you determine the best configuration for your implementation of Autodesk Vault. The image below is a good example of a high level overview of a network. While each office may not be in the implementation plan, documenting that they exist can help understand where other network traffic is come from.



Environment Configurations

While the previous image gives you a high level network overview there is more that needs to be examined. You will need to know the type of network connections between each site as well as the placement of routers and switches. A key to having successful communications is the understanding of the amount of bandwidth and latency between sites.

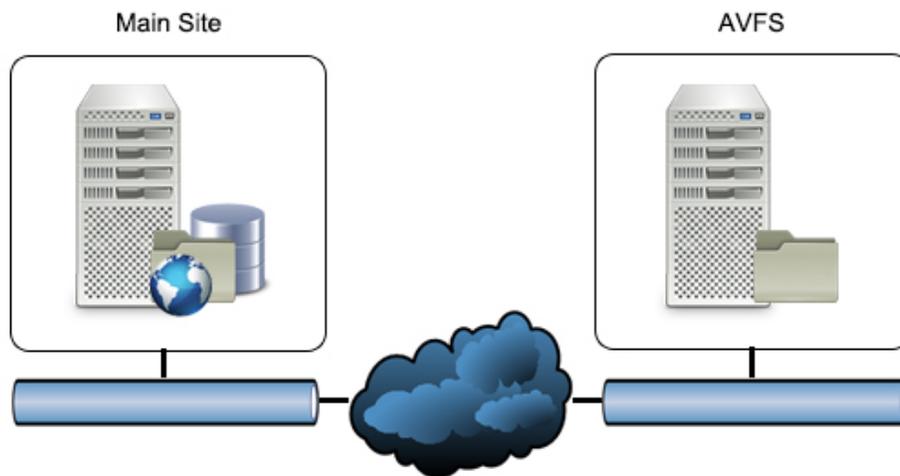


Bandwidth is the speed of transmission over a network connection in kilobits per second. Latency refers to the amount of time it takes in milliseconds to transfer data from one point to another. Both of these factors combine to determine the amount of data that can be transmitted in a certain amount of time over the network. The product of these two factors directly affects the performance of Vault and the user perception of performing certain actions.

Autodesk Vault supports distributed sites by offering 2 types of replication configurations, file replication and full replication. This information can help determine which type of Vault replication maybe deployed throughout your environment.

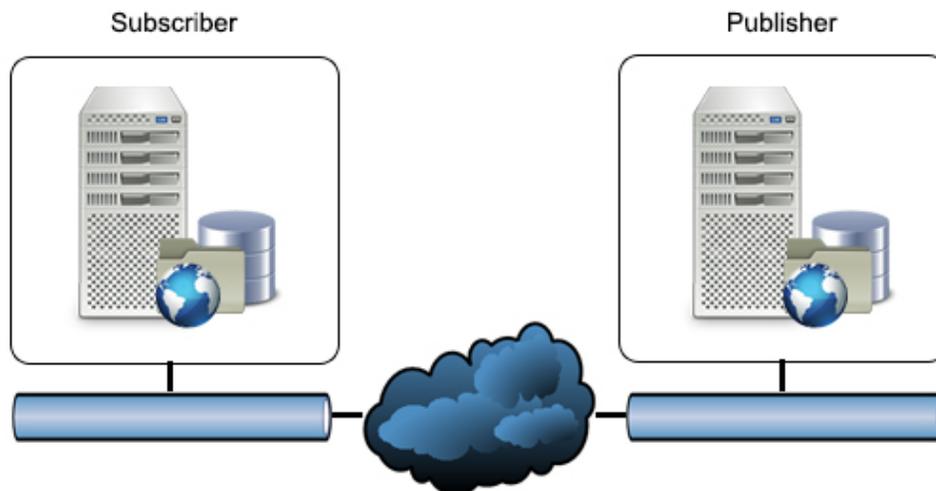
File Replication

With file replication, using the Autodesk Vault File Server, only files are replicated between sites. A central Vault server and database is located at the main site which all remote sites access for metadata information. You can use this configuration when you have low latency between sites. This gives the users at the remote site fast access to local files and good response from the remote metadata.



Full Replication

With full replication, the databases and files are both replicated between sites. The Microsoft SQL Servers are configured to use the merge replication model for the databases. This creates a Publisher at the main site and a Subscriber at the remote site(s). You will want to consider using this configuration when the latency between the sites is starting to exceed 120ms.



A full replication environment setup is more complicated than a file replicated environment. Consider several factors when deciding if you want to set up a multi-site environment:

- Network reliability.
- Network latency and Bandwidth between sites. (low latency and high bandwidth is ideal)
- Number of remote sites.
- Number of users per site.
- Size and number of files to replicate per day/per hour.
- Estimated number of concurrent users.

Hardware Considerations

Many factors have to be taken into consideration when selecting the right type of hardware and how to configure that hardware. These will include the number of users, the average number of files to be upload daily, the average size of the files being uploaded and performance of the components that make up Vault.

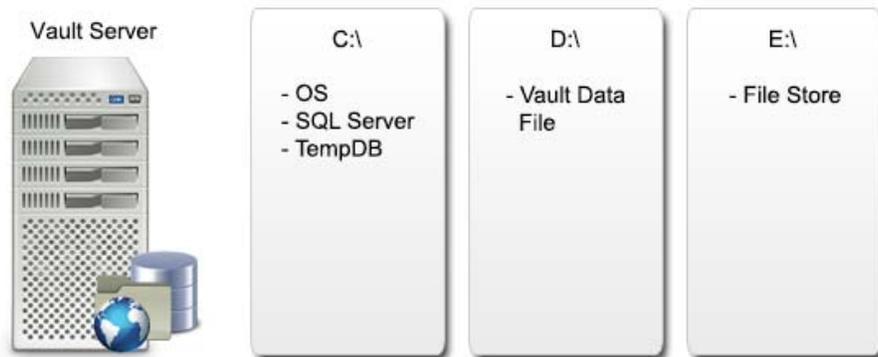
There are 4 components that are a part of the Vault server.

- Autodesk Vault Server – a standalone application which allows you to create and amange vaults and libraries.
- Internet Information Services (IIS) – Web services used for client to server communications. It uses standard HTTP methods, communicating over port 80 by default.
- Microsoft SQL Server – stores the metadata extracted from files when they are added to the vault.
- File store – the file system that stores the actual files checked in by the clients.



Multiple Drives

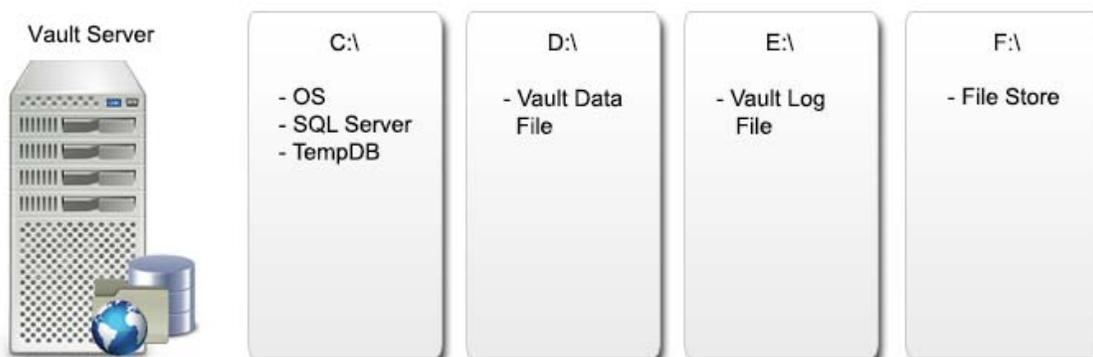
Companies with a small quantity, 5 – 15, of Vault users may be able to setup all of the components on the same physical machine. Even with one machine, multiple physical hard drives with different data on each would help optimize the server's performance.



With this configuration, you will increase the systems performance by spitting the drive I/O requests between the physical drives.

Separating Database files

Separate the Vault database files on to separate physical drives. Databases consist of 2 files, Master Database File (MDF) and Log Database File (LDF) files. The MDF file contains all of the recorded data. As information is being processed in and out of memory, data is stored in the LDF file for error management.



Splitting the Servers

Larger environments may require you to split the Vault components onto different physical computers such as the image below. With this configuration you are giving the components all the dedicated system resources to perform what is needed.



To help you decide what would be the best hardware configuration to use, I would recommend that you use the [IBM Sizing Guide](#). This guide produces recommended hardware and hardware configuration based a number of input parameters. These include:

- Total number of users
- Percentage of simultaneous users
- Amortization period of server
- New Vault implementation vs. migration from an older version
- Total size & number of files

Test Environment

A test environment provides a lot of different benefits. These includes verification that backups can be successfully restore, testing different product configuration and testing product upgrades. A test upgrade provides confidence that the backup is valid and identifies potential pitfalls with no impact on the production environment. Following is an overview of test environment benefits.

Test Machine Settings	Benefits
Similar CPU, memory, disk space	Provides an estimate of how long the migration can take. Gives insight into migration settings, such as time-out values to be changed. Exposes other configuration settings that may require changes.
Database	Test machines that need access to a vault that is 2 GB or more also need a similar version of full Microsoft SQL server installed.
Other applications that are installed on the production server	Test for compatibility with various proxy clients, firewalls, anti-virus programs, etc.
Operating System, Service Packs, firewalls, etc.	Compatibility testing.
Microsoft® Windows® Administration settings	Check for compatibility with Windows Domain policies, login scripts, permissions, and other settings.

If a test environment similar to production is not attainable, a low cost alternative is virtual machines like VMware or Virtual PC from Microsoft. Virtual machines lack many of the benefits outlined above but are far better than not performing a backup and migration validation

Tools & Documentation

In this section we will discuss what documentation should be used to record your environment as well as what tools could be used to monitor and report on your environment.

Documentation

These are some examples of documentation that you could use to record your Vault environment.

Server Information

Server Name:

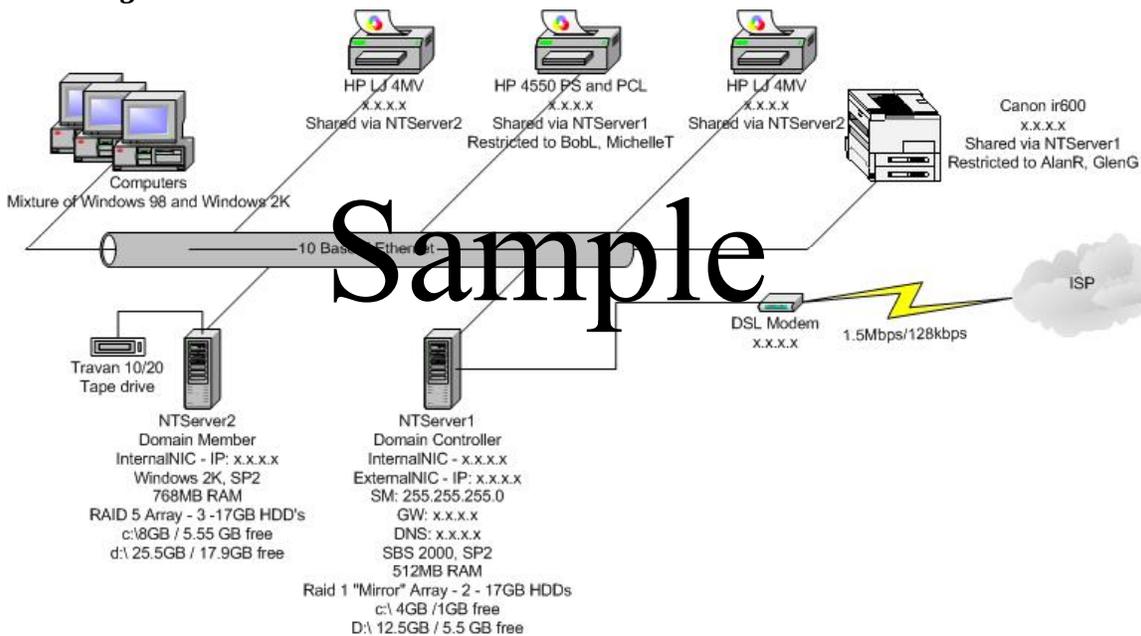
Hardware

Hardware Item	Description
Computer Model	
S/N	
CPU	
Hard disk(s)	
RAM	
NIC	
Other	

Windows Server Installation

Item	Description	
Windows OS		
Domain Name		
Item	Address	
	Internal	External
IP Address		
Subnet Mask		
Default Gateway		

Network Diagram



Vault Information***Install Information***

Item	Description
Vault Server Role	
Install Directory	
License Server Name	
Remote SQL (Yes/No)	
Default IIS Site (Yes/No)	
Impersonation Account	
Impersonation Password	

SQL Information

Item	Description
SQL Server Edition	
SQL Server Name	
SQL Admin Account	
SQL Admin Password	
Database File Path	
Log File Path	

IIS Configuration

Item	Description
IIS Web Site	
IIS Web Site Port	
IIS Application Pool	
IIS Connection Timeout	

Replication

Vaults Available by Site

Vault	Server Name / Location				
	Server 1	Server 2	Server 3	Server 4	Server 5
	USA	Germany	India	Japan	England
Vault 1	X			X	X
Vault 2	X		X		X
Vault 3	X	X		X	
Vault 4	X	X			

File Store Location

Vault	Location

Monitoring Tools

It is very important to monitor the Vault environment. SQL Server provides replication alerts and threshold settings that can be monitored and emailed to system administrators if an alert occurs or a threshold has been passed. You will need to use SQL's Database Mail component to setup the emailed alert.

SQL Database Mail

To start the Database Mail Configuration Wizard

1. From Object Explorer, connect to an instance of SQL Server.
2. Expand Management, right-click Database Mail, and select Configure Database Mail.
3. Choose the Set up Database Mail option to set up Database Mail for the first time.
4. Choose one of the other options for specific maintenance tasks:
 - a. Manage Database Mail accounts and profiles.
 - b. Manage profile security.
 - c. View or change system parameters.

Replication Alerts

Here are a few links which explain what to monitor in a replicated environment:

- [Monitoring Replication with Replication Monitor](#)
- [Setting Thresholds and Warnings in Replication Monitor](#)
- [How to: Set Thresholds and Warnings for Merge Publication](#)
- [Monitoring Performance with Replication Monitor](#)

Splunk

Another tool that can be used to monitor your Vault and server environments is Splunk. [Splunk](#) has the capability of indexing any machines data regardless of data format. Once the data is indexed, it could be used for troubleshooting, network monitoring, alerts and event changes.

Vault Log Files

Vault log files can be monitored and indexed for reporting through a report dashboard. This section explains how to configure Splunk to index the files and extract data to build dashboards and reports.

Log File Setup

To configure Splunk to index the log files, perform the following steps:

1. Select Manager from the upper right-hand corner of Splunk Web.
2. In the Data section of the Manager page, click Data Inputs.
3. Click Files & Directories.
4. Click the New button to add an input.
5. In the Preview Data page, navigate to the location of the Vault log files. If you have installed Splunk on the same server as the Vault server, you can use “C:\programdata\autodesk\vaultserver\filestore\” for the location. You must also add the name of a specific log file in that folder to display sample data from, for example vlog-20130823.txt. If you have installed Splunk on a different server than the Vault server you will have to share the location of the Vault log files and use a UNC path at this step.
6. Click the Continue button.
7. Once the log is displayed, you will need to create a Source Type. Click the Continue button in the Set source type dialog box and click the Continue button on the Data Preview page.
8. Name the source type it Vault Logs in the Review settings dialog and click Save source type.
9. Click Save.

Field Extractions

When monitor Vault log files, information can be extracted for charting and reporting. Create extraction fields for the logs to get this information.

1. Select Manager from the upper right-hand corner of Splunk Web.
2. Click Fields in the Knowledge section.
3. Click Field extractions on the Fields page.
4. Click the New button.
5. Use the chart below to create the field extractions.

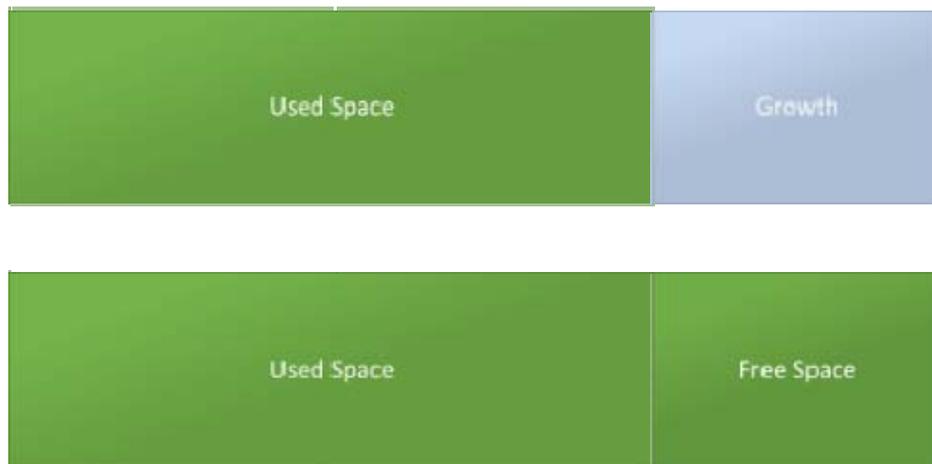
Name	named	Type	Extraction/Transform
Exception	Vault Logs	Inline	(?i).*?: (?P<Exception>\w+\s+\[\d+\])
Host	Vault Logs	Inline	(?i) host(((?P<Host>[^\s])++)
LicenseID	Vault Logs	Inline	(?i) .*?\((?P<LicenseID>\d+)(?=))
Log_Type	Vault Logs	Inline	(?i)^(?:[^\s]*){3}(?P<Log_Type>[^\s]+)
Username	Vault Logs	Inline	(?i) username\(((?P<Username>[^\s])++)



Turning the Dials

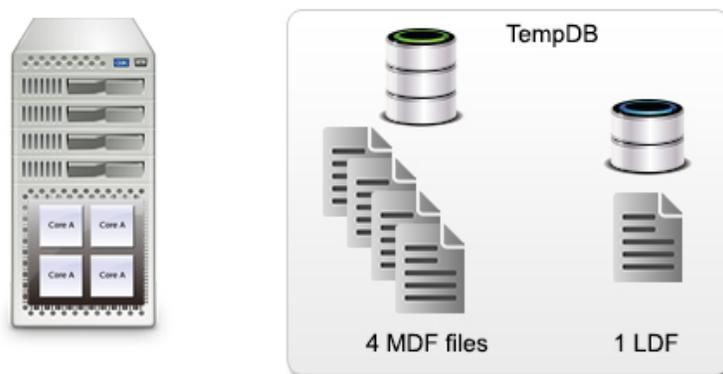
SQL Data File Size

The Vault and TempDB SQL database are set to an initial small size and auto growth size. During normal operations if the database has to grow, Vault performance may be effected during the growth process. It is better to increase the growth size so that each time the database has to grow, the database will have enough free space before it has to grow again. A 500 MB growth size is a good place to start.



Multiple TempDB Data Files

The TempDB MDF database file is one file by default. The TempDB is used for global or local temporary tables, temporary stored procedures, table variables or other objects. A method for increasing performance of the TempDB and its operations is to split it into multiple data files. It is recommended to create one file per CPU core as shown in the below example.



Adjusting SQL Replication and the ADMS Merge Agent:

The purpose of this section is to provide guidelines and suggestions for adjusting the ADMS Merge Agent for Connected Workgroup SQL Replication Performance.

Increase Leveling Threshold:

On the Publisher SQL instance, run the following scripts to increase the generation leveling threshold and validate the change:

1. Open the SQL Server Management Studio (**SSMS**) and connected to the Publisher SQL Instance.
2. Select the **New Query** button and copy / paste the following Query in to the resulting window. This can be run in three segments in the event an error occurs. Lines 1–7 for the first, Lines 9–12 second, Lines 14–17 for the third query.

****Note** that the value is red below will need to reflect your replicated Vault database name. Example, if your replicated vault's name is Production, the `@publication='ADMS_%Vault%'` would be `@publication='ADMS_Production'`

```

/** Change the generation_leveling_threshold to 100000 *****/
/** Replace %Vault% with the name of the replicated Vault database *****/
USE Production
GO

/***** see value before change *****/
SELECT generation_leveling_threshold from sysmergepublications

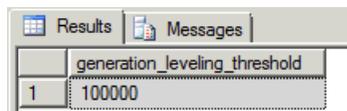
sp_changemergepublication @publication='ADMS_Production',
@property= 'generation_leveling_threshold',
@value= '100000'
GO

/***** verify value after change *****/
SELECT generation_leveling_threshold from sysmergepublications

/***** End of the script *****/

```

The results of the script will be as follows:



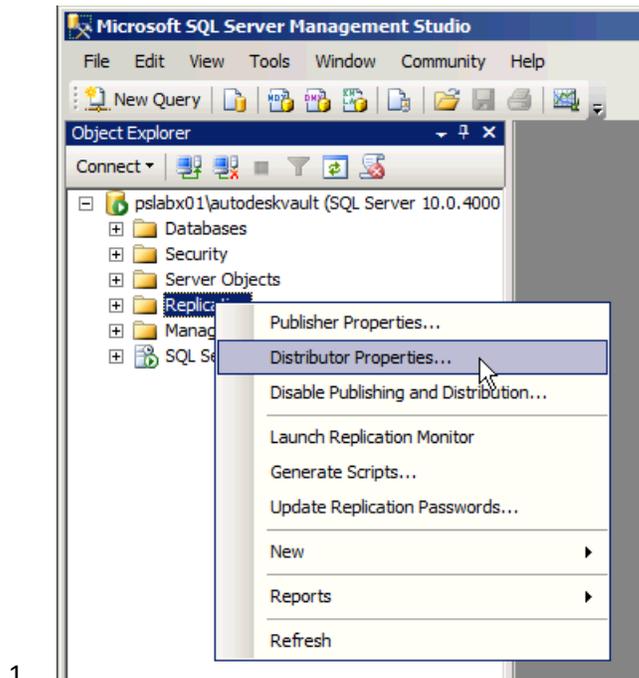
	generation_leveling_threshold
1	100000

Increase Merge Agent Properties:

On the Publisher SQL Instance, modify the ADMS Merge Subscription Agent Values:

- Increase **-QueryTimeout** to 1200 (*if your value is already greater than 1200, this can be left alone*)
- Increase **-DownloadReadChangesPerBatch** to 500
- Increase **-DownloadWriteChangesPerBatch** to 500
- Increase **-DownloadGenerationsPerBatch** to 500
- Increase **-UploadReadChangesPerBatch** to 500
- Increase **-UploadWriteChangesPerBatch** to 500
- Increase **-UploadGenerationsPerBatch** to 500

This can be performed through the SSMS by performing the following clicks:



2.

Distributor Properties - pslabx01\autodeskvault

Select a page

- General
- Publishers

Script Help

Distribution databases

Databases:

Name	Transaction Retention	History Retention	
AutodeskReplication	0 - 72 hours	14 days	...

New... Delete

To set the security for a Queue Reader Agent, click the properties button (...) for the distribution database associated with the agent.

Agent Profiles

Set the default profiles to use when replication agents are created.

Profile Defaults...

Connection

Server: pslabx01\autodeskvault

Connection: sa

[View connection properties](#)

Progress

Ready

3.

Agent Profiles

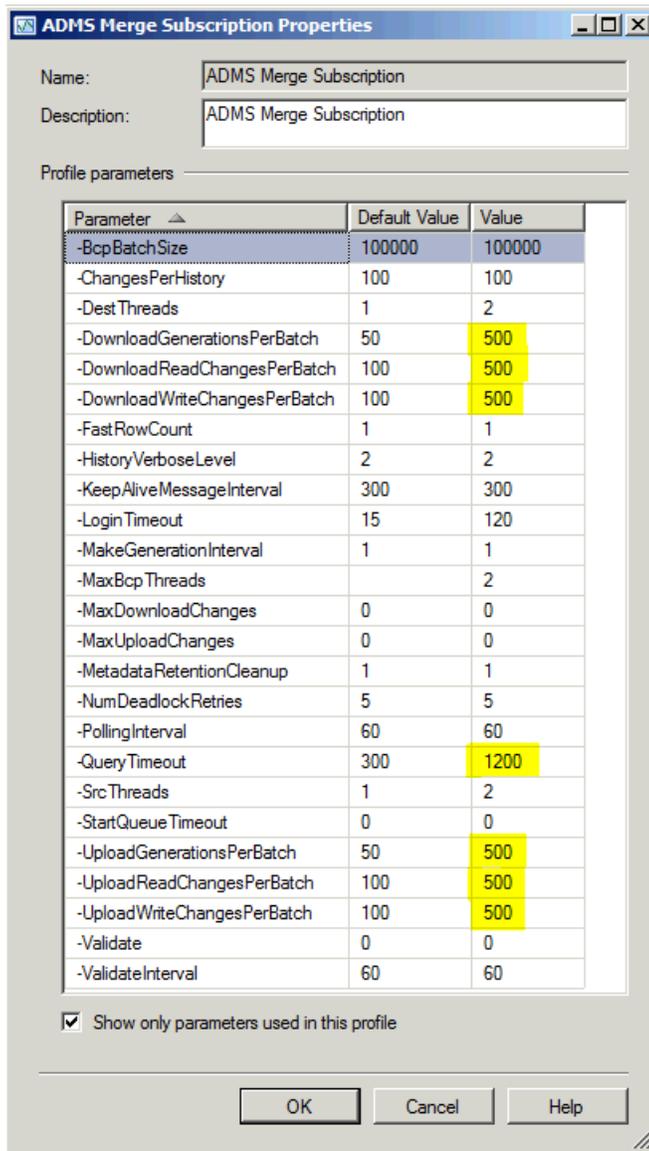
Select a page

- Distribution Agents
- Merge Agents
- Snapshot Agents
- Log Reader Agents
- Queue Reader Agents

Script Help

Agent profiles:

Default for New	Name	Type	
<input type="checkbox"/>	ADMS Merge Subscription	User	...
<input checked="" type="checkbox"/>	Default agent profile	System	...
<input type="checkbox"/>	High Volume Server-to-Server Profile	System	...
<input type="checkbox"/>	Rowcount and checksum validation profile.	System	...
<input type="checkbox"/>	Rowcount validation profile.	System	...
<input type="checkbox"/>	Slow link agent profile.	System	...
<input type="checkbox"/>	Verbose history agent profile.	System	...
<input type="checkbox"/>	Windows Synchronization Manager profile	System	...



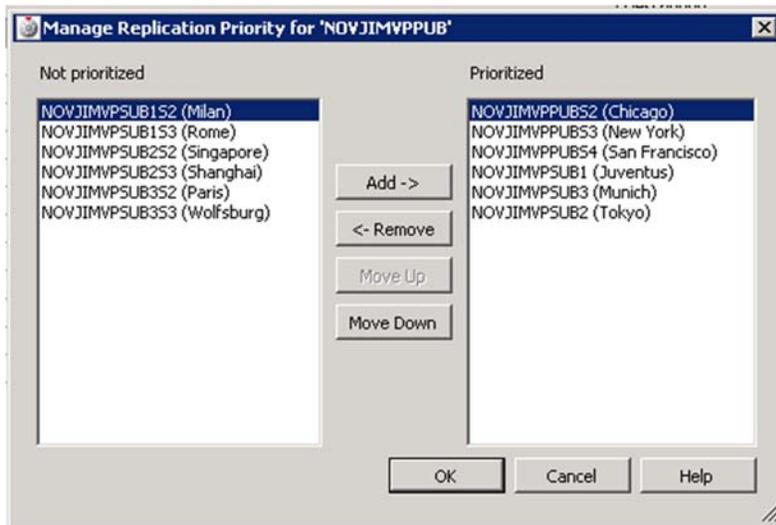
4.

Replication Priority

In a replicated environment files are copied from site to site in random order by default. Leaving this random or unprioritized setting a site could pull files from another site with a less than optimum network connection. Setting the priority between sites with an optimum network connection can improve the performance when files have to be replicated.

Configure Replication Priority

1. On the Server Console, select a Site by expanding the Workgroups node and then by expanding the workgroup the Site is in.
2. Right-click on the site and select Manage Replication Priority.
3. On the Manage Replication Priority dialog, add Sites to the Prioritized section by using the Add button.
4. Prioritize Sites in the Prioritized section by using the Move Up and Move Down buttons.
5. Click OK to save your changes.



Conclusion

Now that I have provided you with enough information to help you determine the hardware configuration you need and provided good tools which will help you monitor your environment. Take these suggesting and monitor and tweak your environment to keep Vault performing at its best.