

Archive Shuttle 11.5
SQL Best Practices Guide



© 2024 Quest Software Inc. ALL RIGHTS RESERVED.

This guide contains proprietary information protected by copyright. The software described in this guide is furnished under a software license or nondisclosure agreement. This software may be used or copied only in accordance with the terms of the applicable agreement. No part of this guide may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording for any purpose other than the purchaser's personal use without the written permission of Quest Software Inc.

The information in this document is provided in connection with Quest Software products. No license, express or implied, by estoppel or otherwise, to any intellectual property right is granted by this document or in connection with the sale of Quest Software products. EXCEPT AS SET FORTH IN THE TERMS AND CONDITIONS AS SPECIFIED IN THE LICENSE AGREEMENT FOR THIS PRODUCT, QUEST SOFTWARE ASSUMES NO LIABILITY WHATSOEVER AND DISCLAIMS ANY EXPRESS, IMPLIED OR STATUTORY WARRANTY RELATING TO ITS PRODUCTS INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. IN NO EVENT SHALL QUEST SOFTWARE BE LIABLE FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL, PUNITIVE, SPECIAL OR INCIDENTAL DAMAGES (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF PROFITS, BUSINESS INTERRUPTION OR LOSS OF INFORMATION) ARISING OUT OF THE USE OR INABILITY TO USE THIS DOCUMENT, EVEN IF QUEST SOFTWARE HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Quest Software makes no representations or warranties with respect to the accuracy or completeness of the contents of this document and reserves the right to make changes to specifications and product descriptions at any time without notice. Quest Software does not make any commitment to update the information contained in this document.

If you have any questions regarding your potential use of this material, contact:

Quest Software Inc.
Attn: LEGAL Dept.
20 Enterprise, Suite 100
Aliso Viejo, CA 92656

Refer to our Web site (<https://www.quest.com>) for regional and international office information.

Patents

Quest Software is proud of our advanced technology. Patents and pending patents may apply to this product. For the most current information about applicable patents for this product, please visit our website at <https://www.quest.com/legal>.

Trademarks

Quest, the Quest logo, and Metalogix are trademarks and registered trademarks of Quest Software Inc. and its affiliates. For a complete list of Quest marks, visit <https://www.quest.com/legal/trademark-information.aspx>. All other trademarks and registered trademarks are the property of their respective owners.

Legend



CAUTION: A caution icon indicates potential damage to hardware or loss of data if instructions are not followed.



IMPORTANT, NOTE, TIP, MOBILE OR VIDEO: An information icon indicates supporting information.

Archive Shuttle
Updated October 2024
Version 11.5

Contents

Requirements overview	4
Hardware considerations	5
Configuration requirements for a small migration	5
Configuration requirements for a medium migration	6
Configuration recommendations for normal migrations	6
Configuration requirements for a large migration	7
Configuration recommendations for Journal Transformation migrations	8
Sharing	9
CPU considerations	9
Memory considerations	10
Network considerations	10
Storage considerations	12
Types of storage device	12
Sizing considerations	13
Virtualized infrastructure	14
Maintenance plans	15
Database backups	15
Index rebuilds and statistics updates	15
About Us	17
Contacting Quest	18

Requirements overview

To best support your SQL server and the needs of a migration, it's recommended that all production Archive Shuttle deployments run on a dedicated instance of Microsoft SQL Server.

Archive Shuttle is a feature-rich product and can therefore result in very diverse SQL Server load profiles between customer deployments. As Archive Shuttle scales, additional databases and the associated increase in load will require scaling of the SQL Servers to meet the load.

Although Enterprise Edition of Microsoft SQL Server is recommended, Standard Edition may be used if the SQL instance uses the recommended (not minimum) resources associated with the size of migration you are performing. Planning for additional time will be required to accommodate regularly required offline maintenance.

It may be necessary to separate out specific Archive Shuttle databases such as the Directory and Link databases to dedicated SQL Servers.

There are many factors which may influence the deployment, but an initial server sizing guideline would be to provision the CPU cores and RAM described in the information below arranged in as many servers as desired, evenly distributing the Archive Shuttle databases if deployed as multiple SQL servers.

Hardware considerations

Archive Shuttle requires a number of SQL databases:

- The Archive Shuttle Directory database holds the configuration information.
- Each Archive Shuttle Item database holds item information for a link. For every source link, there needs to be an item database.

The SQL Server that manages these databases must reside on a separate server than the Archive Shuttle core and modules. The only exception is for a non-production pilot/demonstration system. It's also recommended that it has a dedicated SQL Server instance that's not shared with other applications.

Archive Shuttle keeps track of every item it exports/imports for auditing purposes. It places a reliance on SQL Server for this task. An adequately prepared SQL Server that meets the recommended specification is essential for a fast migration. Ensuring that an adequately prepared SQL Server which meets the recommended specification is essential for a fast migration.

! | CAUTION:

- Performance of the migration will be impacted if the SQL Server environment is not sufficient for the needs of the migration.
- Archive Shuttle recommends that Archive Shuttle SQL servers run a single SQL Server instance only, with no other applications or services running on the servers. SQL Server needs to fully utilize the server resources, and another application may introduce contention issues that will result in undesirable performance.

Configuration requirements for a small migration

The table below indicates the requirements for SQL Server for a small Archive Shuttle migration (less than 5 Tb).

Component	Configuration
Processors (cores)	Minimum 4, Recommended 8
Memory	Minimum 16 GB, Recommended 32 GB

Component	Configuration
Hard disk	Disks should be configured as per Microsoft SQL Server best practices with TempDB, databases and logs on separate disks.
Virtualization	Supported, however performance may be impacted if resources mentioned above are not dedicated

Configuration requirements for a medium migration

The table below indicates the requirements for SQL Server for a medium Archive Shuttle migration (**less than 20 Tb**).

Component	Configuration
Processors (cores)	Minimum 8, Recommended 16
Memory	Minimum 32 GB, Recommended 64 GB
Hard disk	Disks should be configured as per Microsoft SQL Server best practices with TempDB, databases and logs on separate disks.
Virtualization	Supported, however performance may be impacted if resources mentioned above are not dedicated

Configuration recommendations for normal migrations

The following configuration enables better tuning of normal migrations, which allows Archive Shuttle to leverage the environment more effectively for superior performance. You may also need consider tuning the following recommendations within your own environment.

Normal mappings to Office 365 need to be configured as wide as possible:

- Try to use smaller batches
 - Office 365 Batch size Bytes - 5242880 (5 Mb)
 - Office 365 Batch size Items - 50

- Office 365 Timeout seconds - 1800
- Perform a wide migration
 - Office 365 Mailbox parallelism - 150
 - Office 365 Item Parallelism - 5
 - Office 365 Batch parallelism - 2
- Hidden settings configuration
 - Office365NumberOfItemsperContainerMapping (ID - 418)
 - The default value (1000) should be sufficient

Configuration requirements for a large migration

The table below indicates the requirements for SQL Server for a large Archive Shuttle migration (more than 20 Tb).

Component	Configuration
Processors (cores)	Minimum 16, Recommended 32
Memory	Minimum 64 GB, Recommended 128 GB
Hard disk	Disks should be configured as per Microsoft SQL Server best practices with TempDB, databases and logs on separate disks.
Virtualization	Not normally supported. Contact Quest if virtualization is the only possibility.

Archive Shuttle keeps track of every item it exports/imports for auditing purposes. It places a reliance on SQL Server for this task. An adequately prepared SQL Server that meets the recommended specification is essential for a fast migration. Ensuring that an adequately prepared SQL Server which meets the recommended specification is essential for a fast migration.

i | **NOTE:** Performance of the migration will be impacted if the SQL Server environment is not sufficient for the needs of the migration.

Configuration recommendations for Journal Transformation migrations

There are two separate cases that could be configured for Journal Transformation migrations:

1. Journal mapping is split into more mappings according to date filter

- a. Try to use smaller batches
 - i. Office 365 Batch size Bytes - 5242880 (5 Mb)
 - ii. Office 365 Batch size Items - 50
 - iii. Office 365 Timeout seconds - 1800
- b. Perform a wide migration
 - i. Office 365 Mailbox parallelism - 150
 - ii. Office 365 Item Parallelism - 5
 - iii. Office 365 Batch parallelism - 2
- c. Hidden settings configuration
 - i. Office365NumberOfItemsperContainerMapping (ID - 418)
 - 1. The default value 1000 should be enough

2. A big journal is migrated as one migration mapping

- a. Try to use smaller batches.
 - i. Office 365 Batch size Bytes - 1 048 576 (1 Mb)
 - ii. Office 365 Batch size Items - 5
 - iii. Office 365 Timeout seconds - 1800
- b. Perform a wide migration
 - i. Office 365 Mailbox parallelism - 150
 - ii. Office 365 Item Parallelism - 1
 - iii. Office 365 Batch parallelism - 10
- c. Hidden settings configuration
 - i. Office365NumberOfItemsperContainerMapping (ID - 418)
 - 1. Change the value to 100k. This will allow Archive Shuttle to have items spread across different mailboxes which should increase migration speed.

If there are too many items to migrate just from one big mailbox, consider changing the "Ingest users with most items first" setting:

- Enabled - gets the users with most items to import
- Disabled - get users alphabetically regardless of how many items they have to import

Sharing

Archive Shuttle recommends that Archive Shuttle SQL servers run a single SQL Server instance only, with no other applications or services running on the servers. SQL Server needs to fully utilize the server resources, and another application may introduce contention issues that will result in undesirable performance.

CPU considerations

The power of a server is not necessarily determined by the CPU speed in terms of cycles per second. Factors such as the server architecture and number and type of processors and cores can provide a far greater benefit over increasing CPU speed.

Hyper-threading technology is said to provide up to 30% improvement in performance. These processors contain two architectural states on a single processor core, making each physical processor act as two logical processors. However, the two logical processors must share the execution resources of the processor core, so performance gains may not be attained and in some circumstances can even lead to degradation of performance.

Multi-core technology provides similar performance to a comparable multi-CPU server. These processors contain multiple complete processor cores, which act as complete physical processors. Each physical core has its own architectural state and its own execution resources, so the performance gains are reliable.

With the ever-increasing number of processor core combinations and high clock speeds, the traditional x86 Front Side Bus architecture can start to become a bottleneck beyond eight processor cores. A popular and cost-effective method of scaling up the x86 architecture is to use an architecture that supports non-uniform memory access (NUMA). Processors and memory are grouped into nodes that have high-speed local access. However, access to memory co-located with other processor nodes is slower. Therefore, the operating system (and potentially application software) needs to be NUMA-aware and optimized to make the best use of processors, cores, and their associated resources. Windows server and SQL Server support NUMA.

i | **NOTE:** See the Microsoft TechNet article "How SQL Server supports NUMA".

The recommended number of processor cores can be composed of either physical CPUs or similar combination of multi-core CPUs, but the sizing should not be based on hyper-threaded logical cores.

! | **CAUTION:** Hyper-threading is not recommended to improve the database performance due to potential performance problems when the database places a load on the memory. See the Microsoft Knowledge Base article 322385 <http://support.microsoft.com/kb/322385> and the

MSDN article “Be aware: To Hyper or not to Hyper” for further information. If hyper-threading is to be used, particular attention should be paid to the MAXDOP setting as described in KB322385.

In most cases, the SQL Server instance should manage the CPU resources. Do not set the CPU affinity mask unless absolutely necessary, as this can significantly impact the performance. When running multiple SQL Server instances, the most common reason for setting the CPU affinity mask is to prevent an instance being starved of resources.

Memory considerations

The recommended memory should be available at each SQL Server instance to ensure the data manipulation does not cause excessive paging to disk both in the Archive Shuttle databases and tempdb, which will quickly degrade performance.

Install the appropriate edition of Windows Server and SQL Server to support the capacity of memory that is installed. See the Archive Shuttle compatibility charts for supported versions of SQL Server.

Under normal circumstances, SQL Server should be allowed to manage memory dynamically. It may be necessary to change the SQL Server minimum and maximum memory to ensure the memory is used appropriately between SQL Server instances, Reporting services or other co-located services.

Network considerations

Archive Shuttle recommends that you connect the Archive Shuttle SQL servers and Archive Shuttle servers via gigabit network technology. The SQL servers may require multiple network interface cards to support the anticipated loads.

Archive Shuttle also recommends that you disable the TCP Chimney Offload, TCP/IP Offload Engine (TOE) or TCP Segmentation Offload (TSO) to prevent network issues. For guidance in disabling these, see this Veritas technical article.

Minimum Requirements

Component	Configuration
Network	100 MBit

Recommended Requirements

Component	Configuration
Network	1 GBit

Storage considerations

Types of storage device

It is vital to ensure the storage does not become a bottleneck. By following Microsoft SQL Server best practices, it can be ensured that the SQL server is suitably sized. Avoid using network-based storage for the database files.

In most cases, RAID-based storage will be needed to achieve the storage requirements. To maintain performance and reliability, consider hardware-based RAID rather than software-based RAID. To achieve redundancy on striped arrays while maintaining performance, consider the RAID scheme carefully.

RAID levels 5 and 6 are popular, cost-effective methods of achieving redundancy while maintaining striped disk read performance. However, writing incurs a cost of four to six physical operations per write. A poorly sized RAID-5 or 6 implementation can significantly reduce the performance of write-intensive activity. Correctly sizing a RAID-5 or 6 implementation to maintain write performance may become more costly than RAID-1+0, and therefore a RAID-1+0 scheme should be considered.

i | **NOTE:** Consider a RAID 1+0 scheme for maximum performance.

In the case of local or direct attached storage, use multiple controllers supporting multiple channels to distribute the load between the multiple storage locations and provide sufficient throughput. The controllers should also provide a battery-backed read and write cache to aid performance. A minimum of 512 MB controller cache is recommended for local or direct attached storage.

Before using partitions on a storage area network (SAN), consider the I/O load together with any other applications that are already using the SAN to ensure that the performance can be maintained. Ideally, discuss the implementation with the SAN hardware vendor to ensure that it can achieve optimum performance. Typically, LUNs should be created across as many suitable disks as possible, using entire disks rather than partial disks to prevent multiple I/O-intensive applications from using the same disks. When the HBA is configured on the host, ensure that the Queue Depth is set to an optimal value. This should be discussed with the storage vendor.

! | **CAUTION:** Be wary of using partitions on SANs without discussing requirements with the appropriate storage vendor.

When creating a basic NTFS volume on a storage device, it is very important to align the volume with the device sector or stripe unit boundaries to prevent unnecessary disk operations that can significantly impact performance. (Dynamic volumes cannot be aligned at time of publication). See the TechNet article “SQL Server best practices” for more information and using the diskpart

tool to create and align volumes. This article also recommends that both log and data partitions are formatted with 64 KB allocation unit sizes.

In most cases, you should create a single volume on each disk array to avoid contention at the disks between the partitions.

Each database requires the disks to be arranged for two different purposes; the database data files and the transaction log files. The data files require good random access, and therefore a striped array of many disks should be used. The log files require good sequential write performance, so each log file should be placed on its own high-speed array with good transfer rates.

i | **NOTE:** To achieve redundancy on the sequential write-intensive disks (log), use a RAID-1 or RAID-1+0 scheme with high-speed, 15k rpm disks.

Arrange the SQL server storage to accommodate the different types of data, distributing the load as appropriate. The following arrangements of storage might be considered for each data requirement:

Partition	RAID array
System drive	RAID-1 array
Tempdb log file	RAID-1 or 1+0 array
Tempdb data files	RAID-1 or 1+0 array
Archive Shuttle Directory data file	RAID-1 or 1+0 array
Archive Shuttle Directory log file	RAID-1 or 1+0 array
Each link Database data file	RAID-1 or 1+0 array
Each link Database log file	RAID-1 or 1+0 array

If multiple database files are located on one partition, it may require regular file defragmentation to maintain performance.

Sizing considerations

Archive Shuttle uses a single Directory Database, and potentially multiple Item databases (depending on the number of source environment Vault Stores that are configured to migrate).

The Directory Database size requirement is 500 MB. However, to allow for temporary transaction log growth, it is recommended to ensure at least 2 GB is available for the database and logs.

Each item database has an initial storage requirement of 2 GB; 1 GB for the data file, and 1 GB for the transaction log.

The Archive Shuttle Item databases will grow depending on the items that are being migrated. A basic sizing guide for each item database is 1024 – 1500 bytes for each item collected plus 1 GB for static data, transaction logs and temporary data fluctuations.

Component	Configuration
Directory Database	500 MB
Each Item Database	2 GB

In an environment when many source environment Vault Stores are considered for migration, Archive Shuttle recommends that you configure the SQL Server setting relating to the default database and transaction log location to a drive or folder with sufficient storage space. Archive Shuttle also recommends that all of the Archive Shuttle databases (Directory, and Item databases) have the SQL autogrowth feature enabled. Further, in this type of environment, Archive Shuttle recommends that you move the databases and transaction logs for the Item Databases to separate storage areas.

Virtualized infrastructure

There are important aspects to consider when installing SQL Server in a virtualized infrastructure. Follow the recommendations of the hypervisor vendor and Microsoft when sizing and configuring the environment.

The primary objective is to ensure that the resource requirements described above are dedicated to the virtual machine to ensure minimum impact to the performance from intermediate layers or co-existing guests.

The hypervisor should be type-1 (system-provided) to ensure the minimum impact on hardware resource requirements.

Note the following general guidelines:

- In a typical virtualized infrastructure, local disks might be used for the hypervisor and SAN-based storage for the guest operating system images and data file locations. The operating system and data storage partitions should be independent dedicated locations, as described above.
- Disk partitions should be aligned with the device sector or stripe unit boundaries to prevent unnecessary disk operations that can significantly impact performance.
- The disk partitions to be used for the database log files should be created as recommended by the hypervisor vendor for sequential access (possibly raw hard disks).
- The disk partitions to be used for the database data files should be created as recommended by the hypervisor vendor for random access (most likely virtual hard disks).
- Virtual hard disks should be created as fixed size and not dynamic.
- Avoid the use of hyper-threading by the hypervisor.

- Avoid the use of virtual machine snapshots, which can impact performance.
- The memory requirements recommended above should be dedicated and prioritized to the virtual machine to prevent dynamic allocation or sharing.
- The number of processor cores as recommended above should be exclusively dedicated to the virtual machine, and the processor priority and bandwidth set to provide the virtual machine with full utilization of the selected CPUs.

Maintenance plans

There are three aspects of SQL Server Maintenance Plans to consider:

- Database backups
- Index rebuilds and Updating Statistics

i | **NOTE:** The information provided is for guidance only. Often a SQL DBA team will have their own plans and procedures in place for applications relying on databases.

Database backups

Archive Shuttle recommends recommend that you perform regular backups of the Archive Shuttle Directory Database and the Item Databases. These databases contain important data relating to the configuration of the migration as well as the progress, down to the individual item level.

i | **NOTE:** It is a best practice to perform nightly full backups of each of the databases. This will ensure the shortest recovery time.

Other backup types are supported; however, it may lengthen the recovery time if a database restore is needed.

Details for configuring database backups for specific versions of SQL Server is outside the scope of this document.

Index rebuilds and statistics updates

Archive Shuttle recommends that you perform weekly index rebuilds on the Archive Shuttle Directory Database and each Item Database. This is particularly important when 'Item Gathering' is being performed where a large amount of data is being added to the Item Databases.

Archive Shuttle contains a new 'Index Fragmentation' page in the Admin Interface. This page shows the current levels of Index Fragmentation in the Archive Shuttle databases. This page also color-codes the rows to help highlight any potential problems, as follows:

Yellow – Page count > 1000, average fragmentation 10-30%

Red – Page count > 1000, average fragmentation >30%

i | **NOTE:** Perform online Index Rebuilds in SQL Server weekly.

Details for configuring Index rebuilds for specific versions of SQL Server is outside the scope of this document. It is important to note that there are editions of SQL Server where the Index rebuild operation can be performed with the index online, and there are editions of SQL Server where the Index rebuild will cause the index to be made offline while the operation is performed. According to Microsoft documentation (referenced below), online index operations can be performed in SQL Server Developer Evaluation and Enterprise Editions.

[http://technet.microsoft.com/en-us/library/ms186880\(v=sql.105\).aspx](http://technet.microsoft.com/en-us/library/ms186880(v=sql.105).aspx)

If the edition of SQL Server in use for the migration does not support online index rebuilds, make sure to schedule the rebuild operations when there is a low level of activity from Archive Shuttle.

It is also recommended to do the 'Update Statistics' maintenance task, following an index rebuild.

i | **NOTE:** In many environments performance has been seen to dramatically increase if index rebuild and statistics updates are performed daily. Check the index fragmentation page frequently to see if this will help in a specific environment.

About Us

Quest creates software solutions that make the benefits of new technology real in an increasingly complex IT landscape. From database and systems management, to Active Directory and Office 365 management, and cyber security resilience, Quest helps customers solve their next IT challenge now. Around the globe, more than 130,000 companies and 95% of the Fortune 500 count on Quest to deliver proactive management and monitoring for the next enterprise initiative, find the next solution for complex Microsoft challenges and stay ahead of the next threat. Quest Software. Where next meets now. For more information, visit www.quest.com.

Contacting Quest

For sales or other inquiries, visit www.quest.com/contact.

Technical Support Resources

Technical support is available to Quest customers with a valid maintenance contract and customers who have trial versions. You can access the Quest Support Portal at <https://support.quest.com>. The Support Portal provides self-help tools you can use to solve problems quickly and independently, 24 hours a day, 365 days a year. The Support Portal enables you to:

- Submit and manage a Service Request
- View Knowledge Base articles
- Sign up for product notifications
- Download software and technical documentation
- View how-to videos
- Engage in community discussions
- Chat with support engineers online
- View services to assist you with your product