



Version 6.6 and Later

Platform Planning Guide

Contact Information

RSA Link at <https://community.rsa.com> contains a knowledgebase that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA Conference Logo, RSA, and other trademarks, are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to <https://www.rsa.com/en-us/company/rsa-trademarks>. Other trademarks are trademarks of their respective owners.

License Agreement.

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person. No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability. This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on RSA Link. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 -2021 RSA Security LLC or its affiliates. All Rights Reserved.

Contents

Revision History	6
Preface	7
About This Guide	7
Support and Service	7
Archer Documentation	8
Chapter 1: Archer Suite Architecture Overview	10
Platform Architecture	10
Environment Options	11
Single Server Environment	12
Multiple Server Environment	12
High Availability Multiple Server Environment	13
System Complexity	15
Platform Server Considerations	16
Web Servers	16
Services Servers	16
Database Servers	17
Cloud and Hosting Support	18
Chapter 2: Sizing Guidelines	20
Very Small Single-Host Environment	20
Small Environment	22
Medium Environment	24
Large Environment	26
Very Large Environment	29
Offline Access Laptop	29
Disk Sizing	30
Audit Logging	30
Advanced Workflow (AWF)	32
Workflow Complexity	32
Deployment Size	32

Appendix A: Sizing Quick Reference Chart	35
Appendix B: Performance Maintenance	37
Before You Begin: Software Recommendations	37
Advanced Workflow Service Maintenance	38
Browser Choice	38
Bulk Update Actions	39
Caching Maintenance	40
Calculated Cross-Reference	40
Data Feeds	40
Databases	41
Database Encryption	41
External Vendor Interfaces	42
Reverse Proxy	42
Direct Mapping / Single Environment	43
Custom Web Forms	45
Two Separate Environments	46
Fast Response Time	49
File Attachments	49
High Availability	49
Indexing Service (Elasticsearch)	51
IT Security Vulnerabilities Program	53
Job Engine Service	55
Job Engine Service Tasks to Check for All Environments	56
Job Engine Service Tasks to Check for Very Small to Medium Environments	57
Job Engine Service Tasks to Check for Large and Very Large Environments	58
Native SAML 2.0 for Single Sign-On	58
No Impact Factors	58
APIs	59
Inline Editing	59
IP Whitelist	59
Offline Access	60
Peak User Load	60
Performance Monitoring Counters	60
Archer	61

- Asynch Framework Counters61
- Queuing Service 61
- Performance Testing Guidelines 62
- Record Count62
- Relationship Visualization63
- Archer Mobile 63
- Transactional Density 64
- Use Case Design Best Practices64
- Virtualization65

Revision History

Revision Number	Date	Revision Notes
1	October 2019	<p>Added AWF-specific sizing recommendations. For more information, see Advanced Workflow (AWF).</p> <p>Updated High Availability best practices. For more information, see High Availability.</p>
2	March 2020	<p>Added performance analysis data for AWF bulk data actions. For more information, see Advanced Workflow Service Maintenance.</p> <p>Added performance analysis data for the Elasticsearch indexing service. For more information, see Indexing Service (Elasticsearch).</p> <p>Added performance analysis data for the Archer IT Security Vulnerabilities Program. For more information, see IT Security Vulnerabilities Program.</p>

Preface

About This Guide	7
Support and Service	7
Archer Documentation	8

About This Guide

Use this guide to plan your Archer environment, encompassing platform and all solutions.

This document is intended for system administrators who are responsible for planning and installing the Archer environment.

Support and Service

Open a Support case via the toll free phone number for your locale, or using the Case Management portal on RSA Link. Step by step instructions for opening a Support case and using the Case Management portal can be found here: <https://community.rsa.com/t5/support-information/how-to-contact-rsa-support/ta-p/563897>.

Other Resources

Resource	Description
Archer Community	Our public forum, on the RSA Link Community platform, brings together customers, partners and analysts to discuss risk and compliance as a practice. https://community.rsa.com/t5/archer/ct-p/archer
Archer Customer / Partner Community	Our private community is a powerful governance, risk and compliance online network that promotes collaboration among customers, partners, industry analysts, and product experts. https://community.rsa.com/t5/archer/ct-p/archer

Resource	Description
RSA Ready	<p>RSA's Technology Partner Program is where third parties gain access to RSA Software in order to develop an interoperability and have it documented and certified. RSA Ready certifications are posted to an online community and supported by RSA Support.</p> <p>https://community.rsa.com/t5/rsa-ready/ct-p/rsa-ready</p>
RSA Exchange	<p>The RSA Exchange offerings help you rapidly deploy adjacent or supporting risk business processes, quickly integrate new risk data sources, and implement administrative utilities to make the most out of their risk and compliance investment.</p> <p>https://community.rsa.com/t5/archer-exchange/ct-p/archer-exchange</p>

Archer Documentation

You can access Archer documentation on the RSA Link: <https://community.rsa.com/t5/archer-platform-documentation/tkb-p/archer-platform-documentation>.

The following table describes each document.

Document	Description
Release Notes	A list of issues fixed in the release and a list of issues known at the time of the release.
What's New Guide	An overview of the new, updated, and retired features in the current release. Available in PDF format.
Installation and Upgrade Guide	Instructions for installing and upgrading to the latest Archer release. Available in PDF format.

Document	Description
Online Documentation	Information for using Archer, including how to set up and maintain the Archer Platform, how to use the Archer Platform features, how to use the RESTful, Web and Content APIs, security configuration information, and how to install and use the solution use cases. Available from within the product.
Archer Control Panel (ACP) Help	Information for using the Archer Control Panel (ACP) module to manage the internal settings of the Platform, such as license keys, global paths and settings. Available from within the ACP module.
Planning Guide	Information about how to plan for your Archer installation. This document is intended for system administrators who are responsible for installing and managing Archer. Available in PDF format.
Qualified and Supported Environments	Information on the required software platforms for running Archer. Available on RSA Link at: https://community.rsa.com/t5/archer-platform-documentation/rsa-archer-qualified-and-supported-environments/tap/568750 .
Security Configuration Guide	Information about security configuration settings available in the Archer Platform and security best practices for using those settings to help ensure secure operation of Archer. Available in PDF format.

Chapter 1: Archer Suite Architecture Overview

This chapter provides an introduction to the Archer® Suite, information on how you can configure Archer to your business needs, and other details necessary to plan your ideal environment.

For more information on how to use Archer, see the Archer Online Documentation. For specific details on how to install Archer, see the *Archer Platform Installation and Upgrade Guide*. For more information on the Archer documentation set, see [Archer Documentation](#).

Platform Architecture	10
Environment Options	11
System Complexity	15
Platform Server Considerations	16
Cloud and Hosting Support	18

Platform Architecture

Archer is composed of multiple components that you can deploy on either a single server or among a set of related servers. These components include:

- Web Application
- Instance Database
- File Repository
- Configuration Database
- Services

Web Application

The Archer Platform uses a web-based user interface that runs on a Web Server. Manage the Web Application through an Application Pool using Microsoft® Internet Information Services (IIS).

Instance Database

An Archer instance is a single setup that includes unique content in a database, the connection to the database, the interface, and user credentials. For example, you may have individual instances for each office location or region, or for development, test, and production environments.

With the purchase of each Archer license, you receive permission to create a single environment for one production instance and two non-production instances.

Production instances contain live data that can be altered manually by users or through automated interactions.

File Repository

The File Repository serves many purposes for your Archer configuration. For example, the File Repository provides storage for the following services:

- Attachments uploaded as content
- Temporary files, such as packages, exports
- Data feeds
- Charts created by reports and searches

Configuration Database

Archer uses the Configuration Database to store data that is not instance specific; for example, client information and application information (including date and version).

Services

There are several services, listed with Microsoft Windows Services, that control various functions of Archer. They control features such as configuration data, job engine, and advanced workflow.

For full list of services, see the *Archer Platform Installation and Upgrade Guide*.

Environment Options

Environment refers to the way Archer is generally configured. Configurations refer to the arrangement of hardware and software to support a specific environment.

Organizations deploy Archer in a variety of configurations based on the expected user load, utilization, and availability requirements. As business needs evolve, you can adapt and scale the environment to meet the new demands.

This section explains the following three environment categories:

- [Single server environment](#)
- [Multiple server environment](#)
- [High availability multiple server environment](#)

Single Server Environment

A single server environment represents the simplest type of configuration to deploy and maintain. With a single server environment, all of the components exist on one server. It is recommended that you use a single server environment for quickly deploying a basic configuration for demonstration, training, development, or testing purposes.

Multiple Server Environment

Multiple server environments offer a wide range of options for combining or isolating the various Archer server components, thus providing numerous possibilities for flexibly scaling Archer.

This section provides some examples of configurations. Due to the highly configurable nature of Archer, it is unusual to see the same configuration more than once.

Before you start designing your configuration, here are a few tips to keep in mind:

- For enhanced security, review the *Archer Security Configuration Guide* for available options within a multiple server configuration.
- For production environments, host the Instance Database component separately on a dedicated server.
- All Web Servers require the RSA Archer Configuration service to be installed and enabled.
- The RSA Archer Advanced Workflow service runs on a Web Server.

Example 1: Simple Configuration

The following table shows an example of a multiple server configuration with two servers. All of the components, except for the database components, are on the same Web Server.

Web Server	Database Server
Web Application	Instance Database
RSA Archer Configuration service	Configuration Database
RSA Archer Advanced Workflow service	
Remaining services	

Example 2: Advanced Workflow Configuration

The following table shows an example of a multiple server configuration with four servers. The configuration has two Web Servers, one dedicated to RSA Archer Advanced Workflow service, a Services Server for the remaining services, and a Database Server for the database components.

Web Server	Web Server	Services Server	Database Server
Web Application	RSA Archer Advanced Workflow service	Remaining services	Instance Database
RSA Archer Configuration service	RSA Archer Configuration service		Configuration Database

Example 3: Alternative Advanced Workflow Configuration

The following table shows an example of a multiple server configuration with three servers. This configuration has one of each type of server: Services Server, Web Server, and Database Server. The Services Server hosts the services components. The Web Server hosts the Web Application, and RSA Archer Advanced Workflow. The Database Server hosts the database components.

Web Server	Services Server	Database Server
Web Application	Remaining services	Instance Database
RSA Archer Configuration service		Configuration Database
RSA Archer Advanced Workflow service		

High Availability Multiple Server Environment

High Availability systems provide consistent continuous operation for long periods of time. Accomplish this with Archer by providing redundant servers that take over main operations when another server fails.

For high availability environments, it is recommended that you set up the following minimum number of servers:

- 2 Web Servers
- 2 Services Servers
- 2 Database Servers
- Cache Server

The RSA Archer Advanced Workflow service can utilize its own dedicated server or be installed on both the Web Servers or the Services Servers. The final component in a high availability environment is a load balancer to act as a reverse proxy and distribute network or application traffic across a number of servers.

The following table shows an example of a high availability multiple server configuration. Each of the servers described here are duplicated to allow for high availability.

Web Server	Services Server	Database Server	Cache Server
Web Application	Remaining services	Instance Database	
RSA Archer Configuration service		Configuration Database	
RSA Archer Advanced Workflow service			
For third-party caching applications, do the following: <ul style="list-style-type: none"> • If ScaleOut StateServer is selected, the ScaleOut StateServer client is required. • If Redis is selected, no client is required. 			Third-party caching application

To improve availability in the event of a database failure, Archer supports a number of high availability options, such as failover clustering, availability groups, database mirroring, log shipping, and replication. For more information on high availability options, see the following article from MSDN: <http://msdn.microsoft.com/en-us/library/ms190202.aspx>.

System Complexity

The size of your Archer environment depends upon the complexity of your system. Size your system resources appropriately to get the most out of Archer.

To determine the size of your environment:

1. Determine your largest factor from the following:
 - Number of concurrent users
 - Content records
 - Number of significant data feeds
 - Number of applications enabled for Advanced Workflow
2. Use the following chart to find your environment size.

Environment Size	Number of Concurrent Users	Content Records	Number of Significant Data Feeds	Number of Applications Enabled for Advanced Workflow
Small	Up to 100	Less than 100,000	1 to 2	1 to 4
Medium	100 to 250	100,000 to 250,000	3 to 5	5 to 9
Large	250 to 750	250,000 to 750,000	6 to 9	10 to 20
Very Large	More than 750	More than 750,000	10 or more	20 or more

Important: If you plan to use the Advanced Workflow feature, use the feature-specific sizing guidelines in [Advanced Workflow \(AWF\)](#). If your deployment meets the criteria of a very large deployment for any of the categories, contact your RSA sales representative for additional hardware sizing assistance.

API

The Web Services API, RESTful API, and Content API have minimal impact on server performance. APIs run on the same Web Servers that end users interact with. No additional hardware allocation is necessary.

Platform Server Considerations

It is recommended that you plan for these servers as part of your platform environment:

- [Web Servers](#)
- [Services Servers](#)
- [Database Servers](#)

Web Servers

The Web Server hosts the Web Application and RSA Archer Configuration service. A configuration can have multiple Web Servers depending upon the environment complexity.

When planning your configuration, be aware of the following:

- Know your system's complexity, specifically the number of concurrent users, the activities of those users, and the amount of data stored in Archer. To determine your system's complexity, see [System Complexity](#).
- The Web Server should have a RAID5 or RAID10 array spanning four or more 10K or 15K RPM spindles, or an attachment to a storage area network (SAN) logical unit number (LUN).
- Multiple Web Servers require a network share that is accessible to all Web and Services Servers to store keyword search indexes, document repository, and appearance files. For security purposes, it is recommended that the network share be hosted on a dedicated file server outside of a Web or Services Server.
- HTTP compression is enabled by default. If using a load balancer, it is recommended that you deactivate HTTP compression from the Web Servers and configuring HTTP compression to occur on the load balancer.
- More information is available on setting up Web Servers, see “Preparing Archer for Installation” in the *Archer Platform Installation and Upgrade Guide* on the Archer Community on RSA Link.
- Recommendations on disk sizing are available; see [Disk Sizing](#).
- For third-party caching applications, do the following:
 - If ScaleOut StateServer is selected, the ScaleOut StateServer client is required.
 - If Redis is selected, no client is required.

Services Servers

The Services Server hosts many of the services required to run Archer. A configuration can have multiple Services Servers depending upon the environment complexity.

When planning the configuration of the Services Server, consider the frequency, scope, and volume of the following:

- Data Feeds
- Transactions involving creation and deletion of both Content and Metadata
- Notifications
- Calculations
- Packages
- Advanced Workflow
- Search Indexes
- LDAP Synchronization
- Offline Access Synchronization
- Other asynchronous Job Engine activities

If the environment heavily uses services or shared storage, it is recommended that you dedicate more than one Services Server.

Note: For third-party caching applications, do the following:

- If ScaleOut StateServer is selected, the ScaleOut StateServer client is required.
- If Redis is selected, no client is required.

Database Servers

The Database Server stores database information, such as the Instance Database and the File Repository, for Archer. A configuration can have multiple Database Servers depending upon the environment complexity.

When planning the configuration of the Database Server, the following is recommended:

- Separate volumes to store SQL Server data, SQL Server Log, and TempDB components because SQL Server produces one writer thread for each physical disk. While it is preferable to also use separate spindle sets for each of these storage volumes, the separate writer thread improves performance even though these components are on the same spindle set.

Note: Create completely separate disks to allow SQL servers to produce multiple threads.

- RAID10 arrays with at least separate 15K RPM spindles, solid state drives (SSD), or SSD-cached drives (for example, EMC FAST Cache) for the SQL data, SQL log, and TempDB components.
- 64-bit versions of Windows® Server and SQL Server. The 64-bit editions access up to 2 TB locally and use that memory penalty-free for any operation chosen by the Database Management System. Even on servers with less memory (for example, 8 GB), the 64-bit editions provide access to all memory instead of just a portion.

Cloud and Hosting Support

Archer supports hosting in Microsoft Azure® and Amazon Web Services® (AWS) cloud environments. This section provides information to assess and plan for an installation using cloud environments.

When using cloud vendors, only using virtual machines to run your Archer environment is supported.

Use the same process for a cloud or hosted environment as outlined in [system complexity](#) to determine your environment size. Choose a product from your cloud provider that most closely matches your configuration requirements.

For example, consider the specifications for a [small environment](#). As of this publication, the details in this table are accurate based on current vendor specifications.

Element	Small Environment	AWS (m4.xlarge)	Azure (Standard_DS3 package)
Processor	Four cores	Four cores	Four cores
Memory	16 GB	16 GB	14 GB
Disk Space	50 GB HDD	100 GB SSD (Using Elastic Block Store)	100 GB HDD

Note: This table describes hardware requirements only. To understand all requirements for your configuration, see [Sizing Guidelines](#).

Other factors to consider when preparing your cloud-based configuration:

- Input/Output per second (IOPS) directly affects your Archer performance. If you find your performance is slow, consider choosing a vendor product with more IOPS per disk.
- Communication between your on-premises systems and your cloud vendor is key. Contact your vendor to select a method that works best for your environment.

For more information about the different vendor products offered, review the Azure and AWS documentation:

- For Azure, see <https://docs.microsoft.com/en-us/azure/>.
- For AWS, see <https://aws.amazon.com/>.

Chapter 2: Sizing Guidelines

This chapter describes minimum hardware recommendations based on the size of your Archer environment. Choose the configuration that suits your environment best. For more details on determining your environment size based upon the complexity of your environment, see [System Complexity](#).

Important: These sizing metrics are derived from performance tests conducted in the RSA lab environment. Results may vary based on the setup and configuration of your environment. If you plan to use the Advanced Workflow feature, use the feature-specific sizing guidelines in [Advanced Workflow \(AWF\)](#) instead of the general environment guidelines in the following sections.

Very Small Single-Host Environment	20
Small Environment	22
Medium Environment	24
Large Environment	26
Very Large Environment	29
Offline Access Laptop	29
Disk Sizing	30
Audit Logging	30
Advanced Workflow (AWF)	32

Very Small Single-Host Environment

A very small single-host environment is suitable for organizations with the following complexity factors:

- Up to ten concurrent users
- One Archer instance
- Does not require a high-performance or high-availability solution
- Fewer than 75,000 content records
- 10,000 or fewer new content records added per year

In this configuration, the Web, Services, and Database Servers exist on the same server. This environment is suitable for development work, basic functionality testing, or demonstration purposes. For a light-duty production environment, it is recommended to implement a separate Database Server, as described in [Small Environment](#).

The recommended system requirements can vary based on the number of concurrent users and the amount of data stored in Archer database.

The following table lists the minimum requirements for various hardware and software elements.

Element	Requirements
Processor	Two total processor cores Recommended: Four total processor cores Note: Hyper-threading technology is not equivalent to multiple dedicated cores.
Operating System	Microsoft Windows Server 2016 or 2019
Memory	8 GB RAM Recommended: 16 GB RAM
Disk Space	100 GB Recommended: 200 GB
Network Interface	1 Gbps Recommended: 10 Gbps
Video	SVGA compatible (1280x1024) (for installer and Archer Control Panel)
DB Version	Microsoft SQL Server 2017 (64-bit), SQL Server on Linux (Ubuntu), or 2019 (64-bit)
Dedicated System	No
Disk File System	Formatted to NTFS
SQL Server Mixed Mode Authentication	Mixed mode enabled Administrator account
SQL Server Memory Settings	One-half to two-thirds of system memory
Data Files (SQL)	If possible, install the SQL Server data files on spindle sets separate from the Operating System.

Element	Requirements
Other Requirements	Administrator-level account
Required Services	Microsoft SQL Server, Microsoft SQL Agent (to run the recommended scripts)

Small Environment

A small environment is suitable for organizations with the following complexity factors:

- Up to 100 concurrent users
- Fewer than 100,000 content records
- 10,000 or fewer new content records added per year
- Up to two significant data feeds per day
- Up to four applications with simple Advanced Workflow processes enabled

For a small environment, it is recommended that you set up one Database Server and one combined Web and Services Server.

Database Server (1)	Web and Services Server (1)
Instance Database	Web Application
Configuration Database	RSA Archer Configuration
	RSA Archer Advanced Workflow service
	Remaining services

Database Server

The following table lists the minimum required hardware and software elements to build a Database Server for a small environment.

Element	Requirements
Processor	At least four total processor cores (Hyper-threading technology is not equivalent to multiple dedicated cores)
Operating System	Microsoft Windows Server 2016 or 2019

Element	Requirements
Memory	16 GB RAM
DB Version	Microsoft SQL Server 2016 with the latest available SP (64-bit) or 2016 Enterprise Edition (64-bit) or 2017 (64-bit)
Disk Space	50 GB (Instance Database) 2 GB (Configuration Database)
Disks	15K spindles are the minimum for the DB tier, though it is strongly recommended to use a SAN with SSDs, either as regular disks, caching (for example, EMC FAST Cache), or for automated tiering. Distinct spindles for data, transaction logs, and tempdb are required if SSDs are not available.
Network Interface	Required: 1 Gbps Recommended: 10 Gbps

Important: For High Availability configurations, use a load balancer to route traffic equally between available web servers. A SQL Cluster with two database servers is recommended for database fault tolerance.

Combined Web and Services Server

The following table lists the minimum required hardware and software elements to build a combined Web and Services Server for a small environment. In this configuration, it is assumed that the Web Application and all Archer services run on the same server.

Element	Requirements
Processor	At least four total processor cores (Hyper-threading technology is not equivalent to multiple dedicated cores)
Operating System	Microsoft Windows Server 2016 or 2019
Memory	16 GB RAM
Disk Space	50 GB
Network Interface	1 Gbps

Important: For High Availability configurations, use a load balancer to route traffic equally between available web servers. A SQL Cluster with two database servers is recommended for database fault tolerance.

Medium Environment

A medium environment is suitable for organizations with the following complexity factors:

- Up to 250 concurrent users
- Fewer than 250,000 content records
- 20,000 or fewer new content records added per year
- Up to five significant data feeds per day
- Up to nine applications with semi-complex Advanced Workflow processes enabled

For a medium environment, it is recommended that you set up one Database Server, two Web Servers, and one Services Server.

Database Server (1)	Web Server (2)	Services Server (1)
Instance Database	Web Application RSA Archer Advanced Workflow service	Remaining services
Configuration Database	RSA Archer Configuration service	

Database Server

The following table lists the minimum required hardware and software elements to build a Database Server for a medium environment.

Element	Requirements
Processor	At least eight total processor cores (Hyper-threading technology is not equivalent to multiple dedicated cores)
Operating System	Microsoft Windows Server 2016 or 2019
Memory	48 GB RAM

Element	Requirements
DB Version	Microsoft SQL Server 2016 with the latest available SP (64-bit) or 2016 Enterprise Edition (64-bit) or 2017 (64-bit)
Disk Space	250 GB (Instance Database) 2 GB (Configuration Database)
Disks	15K spindles should be the minimum for the DB tier, though it is recommended to use a SAN with SSDs, either as regular disks, caching (for example, EMC FAST Cache), or for automated tiering. If SSDs are not available, distinct spindles for data, transaction logs, and tempdb are necessary.
Network Interface	Required: 1 Gbps Recommended: 10 Gbps

Important: For High Availability configurations, use a load balancer to route traffic equally between available web servers. A SQL Cluster with two database servers is recommended for database fault tolerance.

Web Server

The following table lists the minimum required hardware and software elements to build the two Web Servers for a medium environment.

It is recommended that you run the Web Application and Advanced Workflow services together on both Web Servers and using a load balancer to route traffic equally between the servers.

Element	Requirements
Processor	At least four total processor cores (Hyper-threading technology is not equivalent to multiple dedicated cores)
Operating System	Microsoft Windows Server 2016 or 2019
Memory	16 GB RAM
Disk Space	50 GB
Network Interface	1 Gbps

Services Servers

The following table lists the minimum required hardware and software elements to build a Services Server for a medium environment.

Element	Requirements
Processor	At least four total processor cores (Hyper-threading technology is not equivalent to multiple dedicated cores)
Operating System	Microsoft Windows Server 2016 or 2019
Memory	16 GB RAM
Disk Space	50 GB
Network Interface	1 Gbps

Large Environment

A large environment is suitable for organizations with the following complexity factors:

- Up to 750 concurrent users
- Fewer than 750,000 content records
- 50,000 or fewer new content records added per year
- Up to nine significant data feeds per day
- Up to 20 applications with complex Advanced Workflow processes enabled

For a large environment, it is recommended that you set up four Web Servers running Web Application (IIS) and Advanced Workflow services, two Services Servers, one Database Server, and one or more Cache Servers.

Cache Server (1)	Web Servers (4)	Services Servers (2)	Database Server (1)
	Web Application RSA Archer Advanced Workflow service	Remaining services	Instance Database
	RSA Archer Configuration service		Configuration Database

Cache Server (1)	Web Servers (4)	Services Servers (2)	Database Server (1)
Third-party caching application	<ul style="list-style-type: none"> If ScaleOut StateServer is selected, the ScaleOut StateServer client is required. If Redis is selected, no client is required. 	<ul style="list-style-type: none"> If ScaleOut StateServer is selected, the ScaleOut StateServer client is required. If Redis is selected, no client is required. 	

Database Server

The following table lists the minimum required hardware and software elements to build a Database Server for a large environment.

Element	Requirements
Processor	At least 16 total processor cores (Hyper-threading technology is not equivalent to multiple dedicated cores)
Operating System	Microsoft Windows Server 2016 or 2019
Memory	96 GB RAM
DB Version	Microsoft SQL Server 2016 with the latest available SP (64-bit) or 2016 Enterprise Edition (64-bit) or 2017 (64-bit)
Disk Space	1 TB (for 600K to 750K content records) 800 GB (for 400K to 600K content records) 600 GB (for 250K to 400K content records) 2 GB (Configuration Database)
Disks	15K spindles should be the minimum for the DB tier, though it is strongly recommended that you use a SAN with SSDs, either as regular disks, caching (for example, EMC FAST Cache), or for automated tiering. Distinct spindles for data, transaction logs, and tempdb are a necessity if SSDs are not available.
Network Interface	Required: 1 Gbps Recommended: 10 Gbps

Important: For High Availability configurations, use a load balancer to route traffic equally between available web servers. A SQL Cluster with two database servers is recommended for database fault tolerance. You can use multiple cache servers to handle failover and high availability.

Web Servers

The following table lists the minimum required hardware and software elements to build the four Web Servers for a large environment. It is recommended that you run the Web Application and Advanced Workflow services together on all four Web Servers and using a load balancer to route traffic equally between the servers.

Element	Requirements
Processor	At least eight total processor cores (Hyper-threading technology is not equivalent to multiple dedicated cores)
Operating System	Microsoft Windows Server 2016 or 2019
Memory	24 GB RAM
Disk Space	50 GB
Network Interface	1 Gbps

Services Servers

The following table lists the minimum required hardware and software elements to build a Services Server for a large environment.

Element	Requirements
Processor	At least eight total processor cores (Hyper-threading technology is not equivalent to multiple dedicated cores)
Operating System	Microsoft Windows Server 2016 or 2019
Memory	24 GB RAM
Disk Space	50 GB
Network Interface	1 Gbps

Cache Servers

See the documentation from ScaleOut StateServer or Redis for the requirements of a large environment.

Very Large Environment

A very large environment is suitable for organizations with the following complexity factors:

- 750 concurrent users
- 750,000 content records
- Nine significant data feeds per day
- More than 20 applications with Advanced Workflow processes enabled

Contact your RSA sales representative for very large environment recommendations.

Offline Access Laptop

There are occasions where users, such as Auditors, need access to Archer while disconnected from the primary network. The Archer Offline Access feature allows these users to accomplish their tasks without disrupting day to day business.

Offline Access roughly parallels the architecture of Archer itself and requires a relatively robust machine. At a minimum, the system assumes there are less than 1,000 content records in the environment. If the number of content records in your environment is larger, the system requirements are proportionally higher.

The following table lists the minimum recommended specifications for various hardware and software components.

Element	Requirements
Processor	At least two total processor cores (Hyper-threading technology is not equivalent to multiple dedicated cores)
Operating System	Microsoft Windows® 10, 64-Bit
Memory	Minimum of 6 GB RAM
Disk Space	Required: 50 GB Recommended: 100 GB
Network Interface	1 Gbps
Video	SVGA compatible (1280x1024) (for installer and Archer Control Panel)
Dedicated System	No

Element	Requirements
Disk File System	Formatted to NTFS
Other Requirements	Administrator-level account

Disk Sizing

The following table provides general disk space sizing guidelines for all environment sizes.

Files Stored	Typical Size	High-Usage Size	Typical Expected Annual Growth
Appearance files	1 to 20 MB	60 MB	Zero to little
Search indexes	100 MB to 1 GB	4 GB	500 MB
Document repository	500 MB to 10 GB	10 to 50 GB	Varies greatly
Database (data files only)	5 to 25 GB	150 GB	Less than 25 GB
Database (with full recovery logging)	40 to 100 GB	400 GB	Varies greatly

Actual disk usage requirements vary depending upon your specific needs. In particular, the amount of space required by the File Repository is dependent upon an individual usage scenario. For example, a customer who uses large media files (such as high resolution photos) requires more disk space than a customer whose data content is mostly simple text.

With larger and more numerous files, the limitations of a slow disk subsystem becomes more apparent and can create a performance bottleneck.

Archer can be negatively impacted if any file storage area runs out of space.

Audit Logging

Audit Logging allows deep analysis and monitoring of activity within Archer and thus each user transaction generates multiple log items. This feature affects the environment sizing due to increased traffic over a TCP/IP or UDP port. The final destination for this traffic is a monitoring tool, such as RSA NetWitness.

The following table provides an estimate for the number of items logged based on typical user activity. The numbers are based on a single user.

Type of Transaction	Approximate Number of Monitoring Tool Entries
Data Feed / Data Import	Four times the number of items in source file
Logging into default workplace	240
Switching workplace	150
Display All	35
Advanced Search	40
View or Edit a Record	50
Lookup and Add a Cross Reference	55
Apply/Save	95
Delete a Record	60
Execute a Search or Report	25
Modify Report	10
Save Report	55
Export Report	40
Print Record	15
Access Quick Links	35
Administrative Functions	50-60*

*Average. Creating a simple questionnaire has a log entry rate of about 500, but most administrative functions are much smaller. Since few users typically perform administrative functions, their impact is limited.

Audit Logging can potentially generate a large amount of network traffic between the servers and the monitoring tool. Depending on retention policies, additional disk space for storing log data can be required. Consult the sizing documentation for the chosen monitoring tool to predict this storage requirement based on an analysis of expected activities. For example, RSA NetWitness requires between 0.5 kb and 2 kb storage for each syslog message.

Audit Logging has no impact on the Database Server, but increases CPU utilization on the Web and Services Servers. In a large environment with hundreds of concurrent users, CPU utilization could double. If Audit Logging is required for a deployment, factor this into the hardware design and acquire additional cores to ensure performance is acceptable. If the increased CPU usage does not cause contention (for example, when adequate cores are available), end-user page load times are not impacted.

Advanced Workflow (AWF)

The AWF feature manages the lifecycle of records in specific types of applications or questionnaires by visually modeling your business process and connecting steps in that process to actions in Archer. The following sections provide performance notes and deployment size recommendations according to the number and complexity of AWF scenarios in your deployment.

Workflow Complexity

When planning your Archer deployment, evaluate the complexity of your AWF scenarios according to the guidelines in the following table.

Complexity	Standard Nodes	Send Notification Nodes
Simple	5-10	0
Semi-Complex	10-30	1-2
Complex	30+	2+

Deployment Size

The following table provides deployment size recommendations based on the number and complexity of AWF scenarios in your deployment. Set up your environment according to the sample hardware configuration and performance notes for the recommended deployment size.

Size	Number of Scenarios	Complexity
Small Deployment	1-4	Simple
Medium Deployment	5-9	Semi-Complex
Large Deployment	10-20	Complex

Important: If your deployment exceeds the criteria of a large deployment for any of the categories, contact your RSA sales representative for additional hardware sizing assistance.

Small Deployment

Hardware Configuration

Server Type	Quantity	CPUs	Memory
Web Server	1	4	16GB
Database Server	1	4	16GB

Performance Analysis Notes

- Data imports of up to 1,000 new records to applications with simple workflow typically do not over-utilize hardware resources or cause errors.
- Data imports of up to 10,000 new records to applications with semi-complex workflow typically utilize moderate Web Server CPU resources, and sometimes cause errors.
- Data imports of any size for applications with complex workflow typically utilize near-maximum Web Server CPU resources, and frequently cause errors.
- This environment supports the creation of up to 17 new records per minute for applications with simple workflow. Further increasing the rate of record creation typically results in longer content save times.

Medium Deployment

Hardware Configuration

Server Type	Quantity	CPUs	Memory
Web Server (runs IIS and AWF services)	2 (load balanced)	4	16GB
Services Server	1	4	16GB
Database Server	1	8	48GB

Performance Analysis Notes

- Data imports of up to 1000 new records to an application with simple or semi-complex workflow typically do not over-utilize hardware resources, and rarely cause errors.

- Data imports of up to 10,000 new records for an application with semi-complex workflow typically utilize moderate hardware resources, and rarely cause errors.
- This environment supports the creation of up to 133 new records per minute for applications with simple workflow. Further increasing the rate of record creation typically results in longer content save times.

Large Deployment

Hardware Configuration

Server Type	Quantity	CPUs	Memory
Web Server (runs IIS and AWF services)	4 (load balanced)	8	24GB
Services Server	2	8	24GB
Database Server	1	16	96GB

Performance Analysis Notes

- Data imports of any size to applications with workflow of any complexity typically do not over-utilize hardware resources or cause errors.
- This environment supports the creation of up to 300 new records per minute for applications with simple workflow. Further increasing the rate of record creation typically results in longer content save times.

Appendix A: Sizing Quick Reference Chart

Size	Content Records	Concurrent Users	Description	Element
Very Small Single-Host Environment	Up to 75,000	Up to 10	Web, Services, and Database Servers all on the same server.	<ul style="list-style-type: none"> • 2 CPU/vCPU • 8 GB RAM
Small Environment	Up to 100,000	Up to 100	Combined Web and Services Server with a separate Database Server.	For all servers: <ul style="list-style-type: none"> • 8 CPU/vCPU • 24 GB RAM
Medium Environment	Up to 250,000	Up to 250	Two Web Servers, one for Advanced Workflow and one for Web Application, one Services Server, and one Database Server.	For Web and Services Servers: <ul style="list-style-type: none"> • 8 CPU/vCPU • 24 GB RAM For Database Servers: <ul style="list-style-type: none"> • 8 CPU/vCPU • 48 GB RAM
Large Environment	Up to 750,000	Up to 750	Four Web Servers running Web Application (IIS) and Advanced Workflow services, two Services Servers, and one Database Server.	For Web and Services Servers: <ul style="list-style-type: none"> • 8 CPU/vCPU • 24 GB RAM For Database Servers: <ul style="list-style-type: none"> • 16 CPU/vCPU • 96 GB RAM

Size	Content Records	Concurrent Users	Description	Element
Very Large Environment	More than 750,000	More than 750	Contact your RSA sales representative for very large environment recommendations.	<p>For Services Servers:</p> <ul style="list-style-type: none"> • 6 CPU/vCPU • 32 GB RAM <p>For Web Servers:</p> <ul style="list-style-type: none"> • 16 CPU/vCPU • 32 GB RAM <p>For Database Servers:</p> <ul style="list-style-type: none"> • 32 CPU/vCPU • 128 GB RAM
Offline Access Laptop	Up to 1,000	1	Standalone computer that can manual sync with Archer for offline use.	<ul style="list-style-type: none"> • 2 CPU/vCPU • 6 GB RAM

Appendix B: Performance Maintenance

The default Archer configuration provides optimal performance for most organizations. However, deployment, user, and administration factors can influence overall performance over time. Regular maintenance ensures continued and reliable performance of your Archer environment.

This section provides a high-level overview of factors that impact Archer system performance, and suggestions for making adjustments to get better performance from your Archer environment.

Note: System administrators responsible for Archer performance maintenance must have knowledge of Microsoft Internet Information Services, Microsoft SQL Server, Windows Servers, and have appropriate permissions to the infrastructure.

Before You Begin: Software Recommendations

Software recommendations vary based on the number of concurrent users and the amount of data stored in Archer. It is recommended to use a configuration that supports your environment's peak usage levels. For a complete list of configuration recommendations and characteristics, see [Architecture Overview](#).

Before beginning any performance maintenance, review the minimum required software elements for each area within your Archer configuration. For a complete list, see "Archer Qualified and Supported Environments" in the Archer Online Documentation.

Area	Recommended Software
Operating System	Microsoft Windows Server 2016 or 2019 Standard or Datacenter editions.
Database Server	Microsoft SQL Server 2017 (64-bit), SQL Server on Linux (Ubuntu), or 2019 (64-bit) SQL Express [®] is not supported.
Web Server	Microsoft Internet Information Services (included in Microsoft Windows Server 2016 or 2019) Microsoft Office [®] 2010 or 2013 Filter Packs [®] (to enable indexing of MS Office files) Requires Microsoft Filter Pack 2.0 or later. Microsoft .NET Framework 4.7.2

Area	Recommended Software
Services Server	Microsoft Sync Framework [®] 2.1 (for offline access)
Offline Access	Microsoft Sync Framework 2.1 Use a sizing heap of 32 GB to 48 GB.

Advanced Workflow Service Maintenance

The Advanced Workflow service administers the Archer Advanced Workflow feature, and must run constantly in environments that use workflows.

Tables within the Archer Instance Database support Advanced Workflow functionality. These tables contain metadata about workflow processes, and transactional details about enrolled content. Some of these tables accumulate large amounts of data, depending upon Advanced Workflow activity, and purge automatically after 30 days. Administrators can change the Advanced Workflow data retention period.

Initiating data feed or data import actions for applications enabled for Advanced Workflow can cause the Advanced Workflow transaction tables to rapidly accumulate data. In these cases, disable the target application's Advanced Workflow process before a bulk data action, and re-enable it upon completion.

Performance Analysis Notes for Advanced Workflow Bulk Data Actions

- Testing produced the following results in an environment with one web server and one database server, each running Windows Server 2012 R2 with 8 CPUs and 48GB RAM.
- Archer processes AWF bulk data actions at a rate of approximately 1000 records per minute.
- Average CPU utilization on both servers averages 5% while a bulk action job runs.
- Memory utilization on both servers is insignificant.
- Performance metrics are similar whether records wait on a User Action Node or a Wait For Content Update Node while a bulk action job runs.

Browser Choice

The Archer Platform uses a web-based user interface that users access through a web browser.

For information on which browsers are supported for each version of Archer, see the [Archer Qualified and Supported Environments Guide](#).

Bulk Update Actions

On-demand bulk actions enable you to perform actions on multiple records in bulk from a Search results page in an application. Update fields across multiple records, create new reference records from search result records, or delete records in bulk. Bulk update actions allow you to inline edit fields from a cross-reference grid in a record.

Testing produced the following results in an environment with a web server, a services server, a database server, and a caching server. Each server was equipped with eight cores and 24 GB memory, with the exception of the web server that had 24 cores and 24 GB memory.

During the first test, one million records were loaded into a simple application with 18 fields. No calculations or calculated cross-references were involved. Bulk Delete action was performed with different numbers of jobs running to perform the action to understand how this parameter affected the time required to complete the task.

- The time required to complete the Bulk Delete action increased linearly as the number of records to be deleted increased.
- The time required to complete the Bulk Delete action generally decreased as the number of jobs increased. This is not a linear relationship because adding more jobs also adds overhead. Improvement seen in completion time when increasing the number of jobs from one to two is much larger than increasing from 5 to 10.
- Simply adding more jobs does not guarantee that Bulk Delete will complete faster.

During the second test, one million records were loaded in a simple application with 19 fields (including one cross-reference field) and no calculations were involved. The time required to update 100,000 records in the application using a data feed was compared to the time required to complete the same task using the Bulk Update action.

- When all fields (including cross-reference) were updated with the data feed, it required approximately four hours to update 100,000 records. The Bulk Update action completed the same task in about two hours.
- When the Cross-Reference field was excluded from the updates, the data feed required about 90 minutes to update 100,000 records. The Bulk Update action completed the same task in about 10 minutes.
- Bulk Update action was generally faster for updates with synchronous calculations. The performance of Bulk Update action was better with five jobs for synchronous calculations compared to 10 jobs.

Caching Maintenance

Caching places frequently used data into high performance storage making it easier to access.

Large and Very Large environments that experience heavy Database Server traffic can benefit from implementing caching. Data caching software reduces Database Server traffic by reusing commonly-requested datasets. Data caching also reduces the need for higher-end hardware for the Database Server.

For more information on Large and Very Large environments, see [Large Environment](#) and [Very Large Environment](#).

Calculated Cross-Reference

The Cross-Reference field type enables you to create associations between records in the same application (internal references) or records in one or more different applications (external references).

Performance Analysis Notes for Calculated Cross-Reference

- Testing produced results in an environment with two load-balanced web servers, a services server, and a database server, each running Microsoft Windows Server 2012 with two Intel Xeon CPU E5-2690 v3 with 12 cores each.
- 10,000 records were imported into two apps where every record automatically related to each other, created 100,000 content records with 100 million content relationships.
- It took three hours of processing time for both data feeds to load applications.
- The database growth from creating the content relationships was approximately 13 times the original database size; most growth was from the table storing cross-reference relationships.
- Creating keyword search indexes took approximately 26 hours after starting the data feed, with the completed search index approximately 13 times the size of the original.
- For feeds that create large amounts of new content relationships, the database impact results in reduced application performance until the database indexes are reorganized.

Data Feeds

Data feeds allow you to build dynamic integrations with external enterprise systems and files, which run automatically on an on-going schedule.

It is recommended that you dedicate a server to the RSA Archer Job Engine service under either of the following conditions:

- Execution of eight or more significant data feeds on a frequent basis.
- Data feeds process more than 50,000 records a day.

Match this server to the configuration of a Web Server without the requirement to respond to web requests from users.

Databases

Proper Microsoft SQL Server maintenance and configuration settings are crucial. As data volume grows, an improperly configured or neglected Database Server can negatively impact the overall performance of Archer.

It is recommended that you follow these best practices for SQL Server maintenance within the SQL Server Management tool:

- Run a currently-supported version of SQL Server with the most recent service pack from Microsoft. To determine if your software is current, see [Before You Begin: Software Recommendations](#).
- Set Max Degree of Parallelism to 1. This is generally applicable to a Database Server dedicated to Archer.
- Configure the Database Server for one tempdb data file per physical CPU core, up to a maximum of eight tempdb data files.
- Set regular processes to back up both the Instance Database and the Configuration Database.
- Set regular processes to rebuild the Instance Database's indexes and update statistics.
- Implement regular processes to manage the transaction log properly when the recovery model for the Instance Database is not set to Simple.
- Set Auto Shrink for the Instance Database to False.
- Set Autogrowth to absolute amounts instead of percentages. This setting exists for data and log files in both the Instance Database and the temporary database.

Database Encryption

Database encryption uses an algorithm to transform data stored in a database into cipher text that is incomprehensible without decryption. Database encryption protects information stored in a database from being accessed by individuals with potentially malicious intentions.

Encrypt the database using one of the following methods:

- Microsoft Encrypting File System (EFS)
- Transparent Data Encryption (TDE)

The additional overhead of encrypting the database causes a 5 to 30 percent loss of Archer performance. Generally, the overhead associated with TDE is less than with EFS. It is recommended that you upgrade to processors with higher clock speeds, more cores, and sufficiently fast disks.

External Vendor Interfaces

Third party vendors use an external vendor interface to Archer to import data into a customer's environment, which open using one of the following methods:

- [Reverse Proxy](#)
- [Direct Mapping / Single Environment](#)
- [Custom Web Forms](#)
- [Two Separate Environments](#)

Each method has distinct advantages and disadvantages.

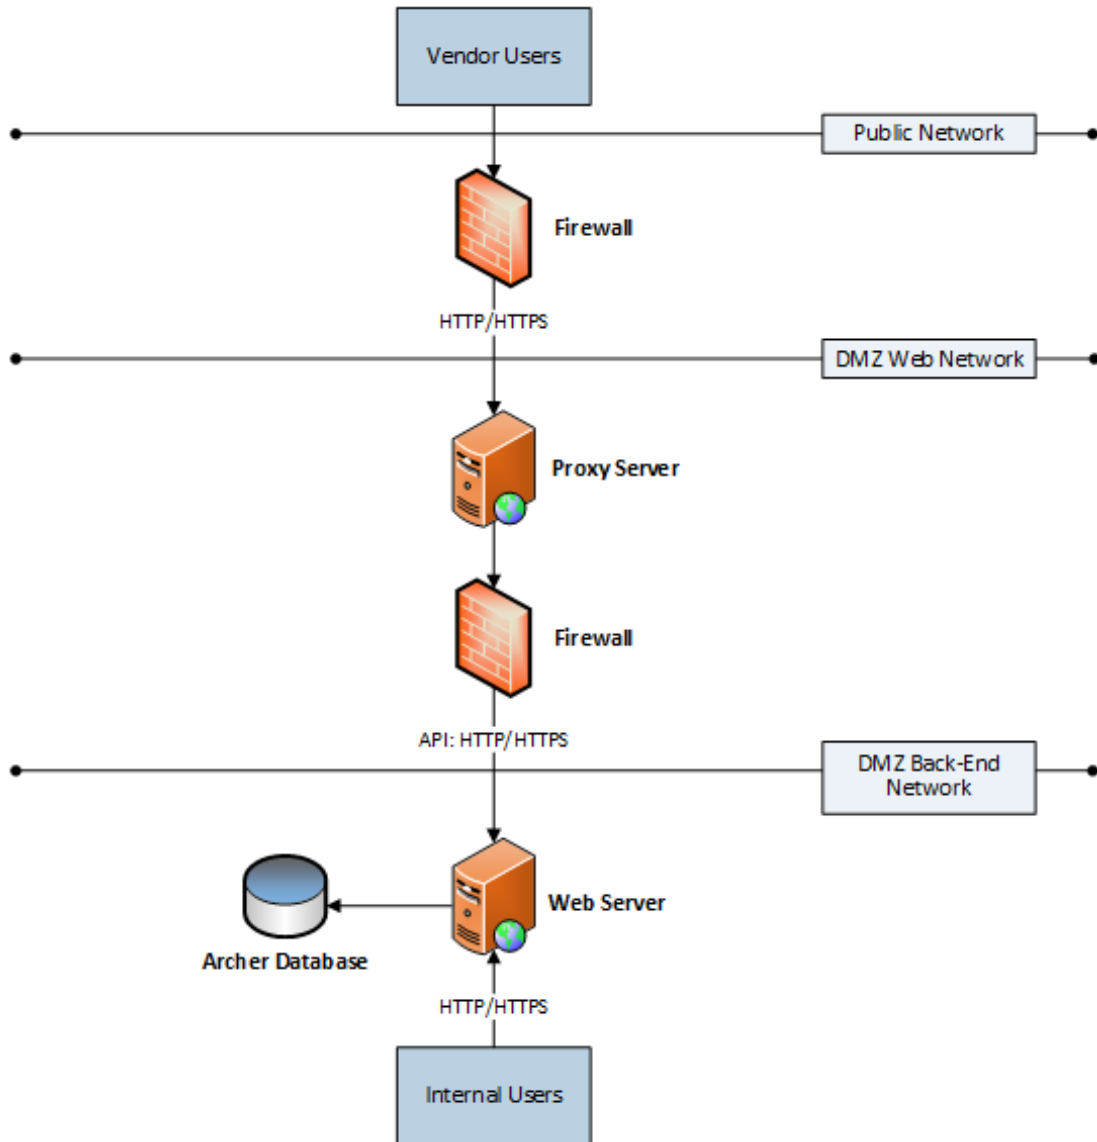
For preventative maintenance, it is recommended to implement the [reverse proxy](#) option because it provides the cleanest, least-complicated solution, and requires only HTTP/HTTPS traffic to traverse the firewall.

Reverse Proxy

A reverse proxy server sits behind the firewall in a private network and directs client requests to the appropriate back end server. The reverse proxy method provides an external interface for users outside the trusted network, and ensures the smooth flow of network traffic between clients and servers.

Advantages	Disadvantages
<ul style="list-style-type: none">• Lightweight• Does not allow direct access to Archer servers• Benefits from all flexibility and features present in Archer• Integrated access control and authentication mechanisms	<ul style="list-style-type: none">• Requires a firewall rule for HTTP/HTTPS traffic to pass from the external proxy server to the inside server• May require some further security configuration changes if using Single Sign-On

The proxy then routes those requests to the internal server environment as shown in the following diagram.

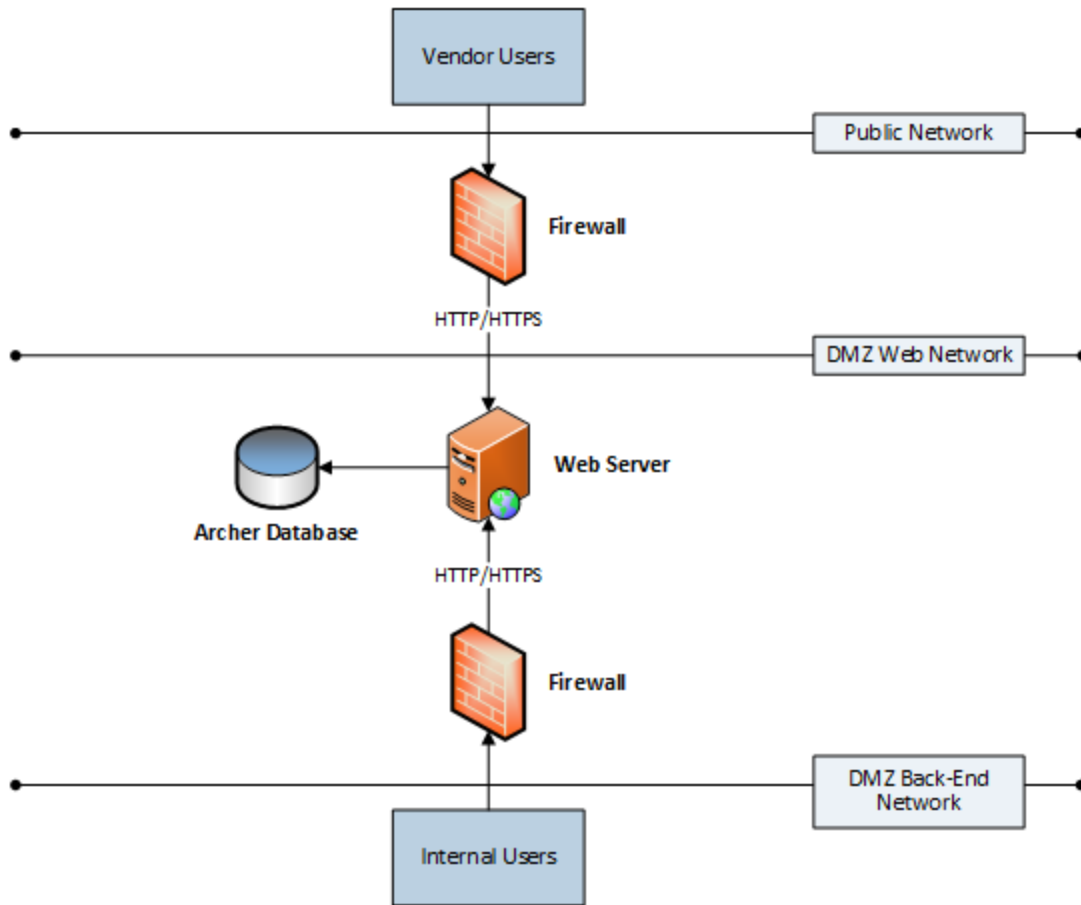


Direct Mapping / Single Environment

The direct mapping/single environment method configures Archer to be accessible to both external and internal users, either in an internal, trusted environment, or in the Demilitarized Zone (DMZ).

Advantages	Disadvantages
<ul style="list-style-type: none"> • Lightweight • Better availability, since establishing an HA environment applies to all users, regardless of source • Benefits from all flexibility and features present in Archer • Integrated access control and authentication mechanisms 	<ul style="list-style-type: none"> • Security concerns because vendors have direct access to environment

The following diagram illustrates an external vendor interface in a direct mapping in a single environment.

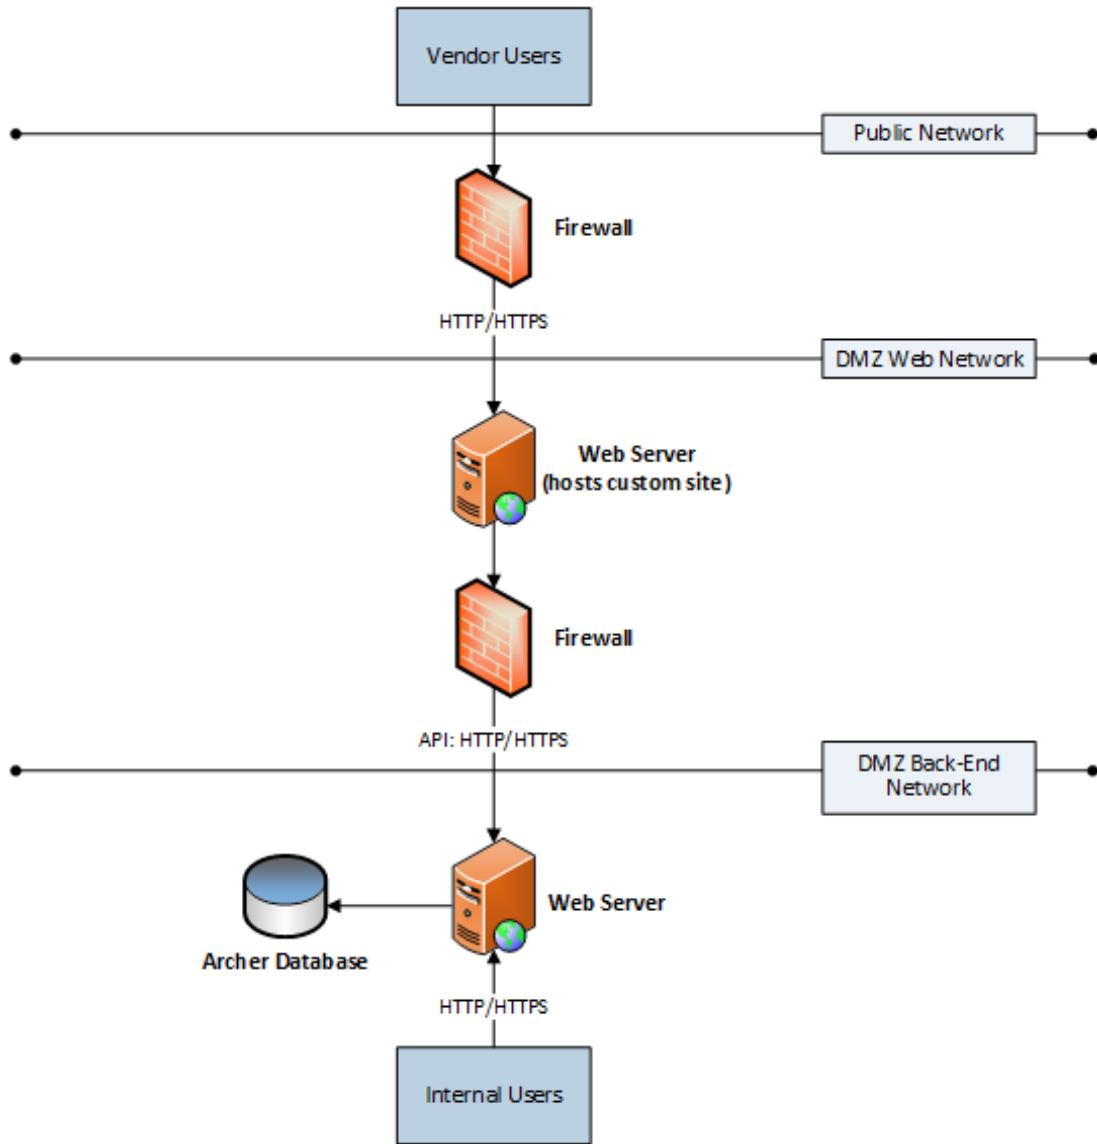


Custom Web Forms

The custom web forms method creates a custom web form that runs on an external Web Server. There is no database back end. When a vendor submits the form, Web Services API sends the data to the internal instance.

Advantages	Disadvantages
<ul style="list-style-type: none"> • Lightweight • Secure; only one open port necessary to allow the external server to communicate with the Web Services API on the internal instance 	<ul style="list-style-type: none"> • Lacks benefits from the Archer access control • Lacks benefits of the Archer feature set • Challenging authentication • Requires custom integration code • Intermediate saves not allowed before final form submission • Requires code modifications to change external static form. • Requires special provisions for file attachments

The following diagram illustrates an external vendor interface in a custom web forms environment.



Two Separate Environments

The two separate environments method runs two independent Archer environments with separate SQL databases. Vendors interact only with the external instance, and the instance transfers that information to the internal instance of Archer.

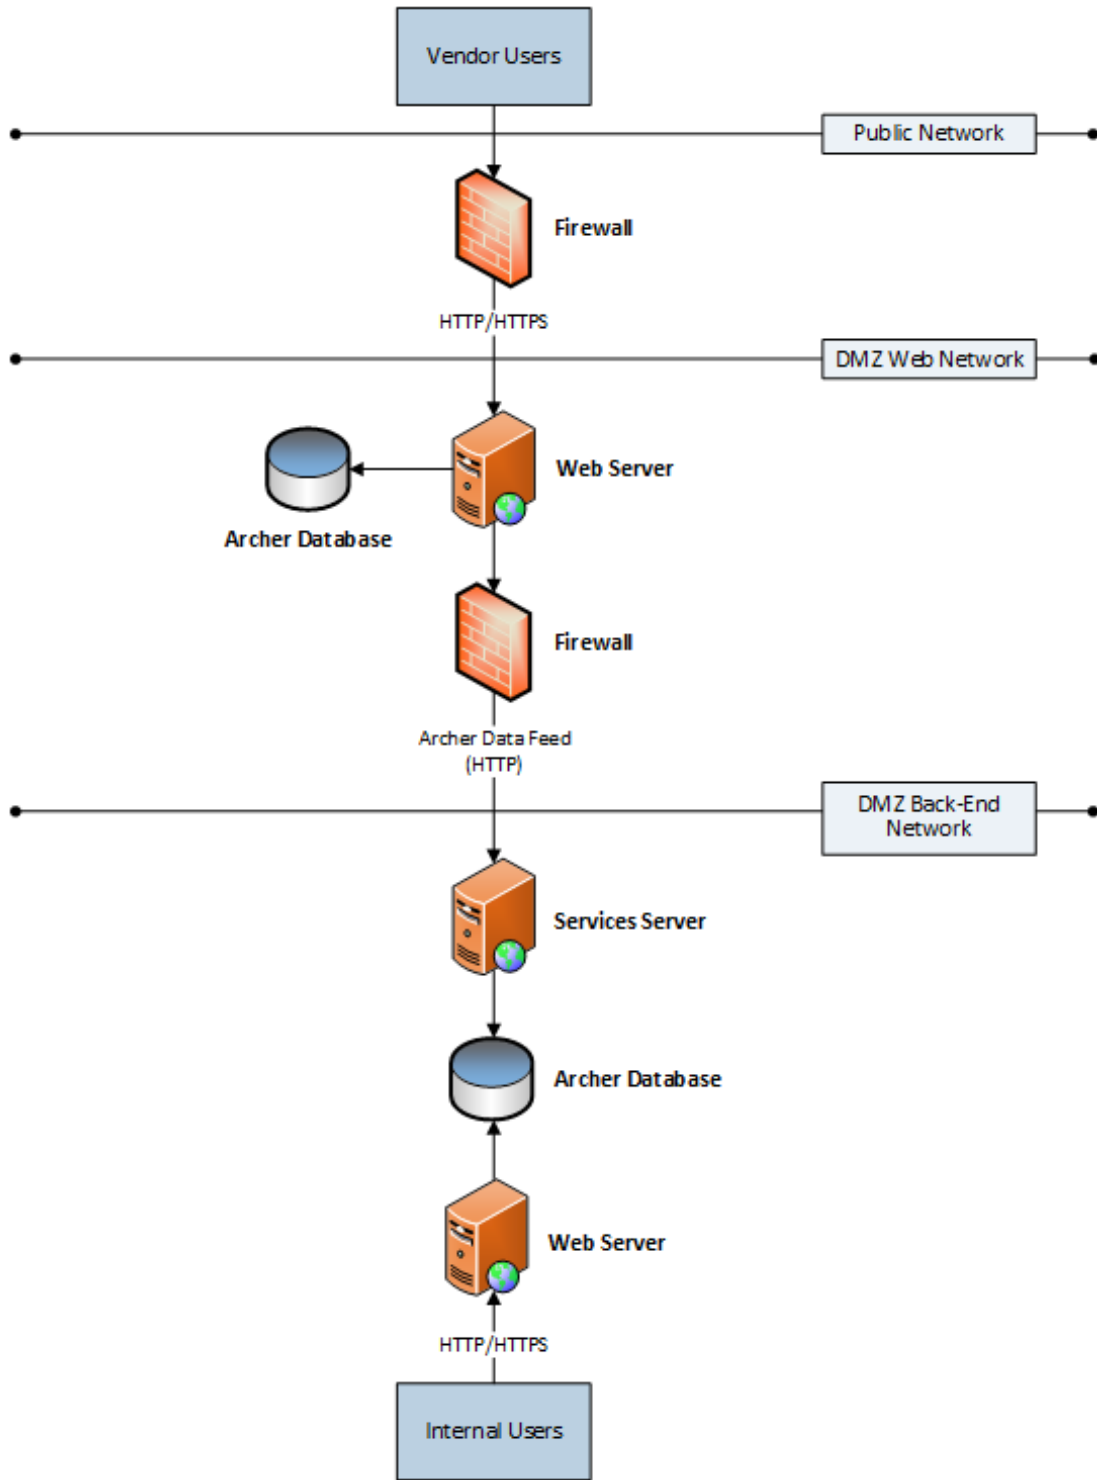
Data transfer methods include:

- Manual data export, import, or both
- Automated submission upon record save
- Scheduled synchronizations

Note: Both automated submission upon record save and scheduled synchronizations require an RSA Professional Services engagement to create the custom data transportation component.

Advantages	Disadvantages
<ul style="list-style-type: none"> • Highly secure (when using manual data export/import) to reasonably secure (automated submission and scheduled synchronizations require only a single port) • Draft stages available • Benefits from all Archer features • Integrated access control and authentication mechanisms 	<ul style="list-style-type: none"> • Requires custom code for the sync component or a user to manually export/import data • Requires more robust external hardware and requires a SQL database • Requires special provisions for file attachments

The following diagram illustrates an external vendor interface in two separate environments.



Fast Response Time

Some actions require more time than others. For example, viewing a record should be quicker than executing a large search or report, saving a record, or viewing a workspace with numerous report-based iViews.

Increase response time by using higher-end current-generation processors. Disk performance can also play a large role. To improve response times, use separate 15K RPM spindles, solid state drives (SSD), or SSD-cached drives such as FAST Cache.

File Attachments

The number and types of file attachments can impact sizing guidance for the document repository storage location, as indicated in the Archer Control Panel. For more information, see "Designating the File Repository Path for an Instance" in the *Archer Control Panel Help*.

It is recommended to perform the following two tasks at regular intervals:

- Estimate growth for the next one to three years.
- Utilize multiple Service Servers if there are other factors that also suggest a high services load. The following file attachments indexed for keyword search can induce more load on the Services Server:
 - Text
 - HTML
 - Word files
 - PDF files

High Availability

High Availability (HA) refers to a system or component that is continuously operational in case of hardware failure. Use HA as a preventative measure against down time in the event of failure. To meet HA requirements, Archer can:

- Deploy on multiple servers
- Use load balancing technology
- Utilize shared Unified Naming Convention (UNC)-accessible storage locations
- Utilize Network Attached Storage (NAS) appliance or clustered file server
- Create shares on one of the participating Services Servers

In addition, you can specify shared, UNC-accessible storage locations for:

- Archer document repository
- Keyword search index files
- Appearance files usually stored in file repositories

The following best practices are recommended to ensure HA performance efficiency:

- Deploy duplicate web servers running identical services
- If your environment leverages the Advanced Workflow (AWF) feature, run the AWF service on each web server. AWF uses port 8000 for HTTP and 8443 for HTTPS
- Configure a load balancer to manage network traffic between web servers

Note: The Archer Job Engine is self-balancing, and does not rely on a load balancer to distribute workload. Each Job Engine process periodically queries the database and starts the next job in the queue automatically

- If available, use a preexisting, highly available Network Attached Storage (NAS) appliance or clustered file server. Store the Archer File Repository and Company Files in this shared location
- Store search indexes locally, as keyword search index performance slows down if search indexes are stored on the NAS. Instead, create shares on one of the participating Services Servers

Note: Most components of Archer failover from one server to another automatically when properly load balanced, but flat file search index does not. If you do not use the Elasticsearch feature available in Archer 6.5 or later, you must manually select a new server to run the search index by its DNS name. This option might not meet the requirements for HA. If you choose this option, scale your server vertically.

- For the database tier, use an HA-clustered configuration. HA clusters are groups of computers that support server applications, and are reliable with a minimum amount of down-time. It is recommended that you configure HA clusters due to:
 - Industry prevalence
 - Front-end configuration simplicity
 - Fast automatic fail-over

You can also use other enterprise-class HA techniques for database recovery, such as activity groups, log shipping, native mirroring, or regular SQL backup and restore. In most situations, you can independently implement SQL Server HA practices in preexisting environments.


For more information, see "High Availability Multiple Server Environment" in [Environment Options](#).

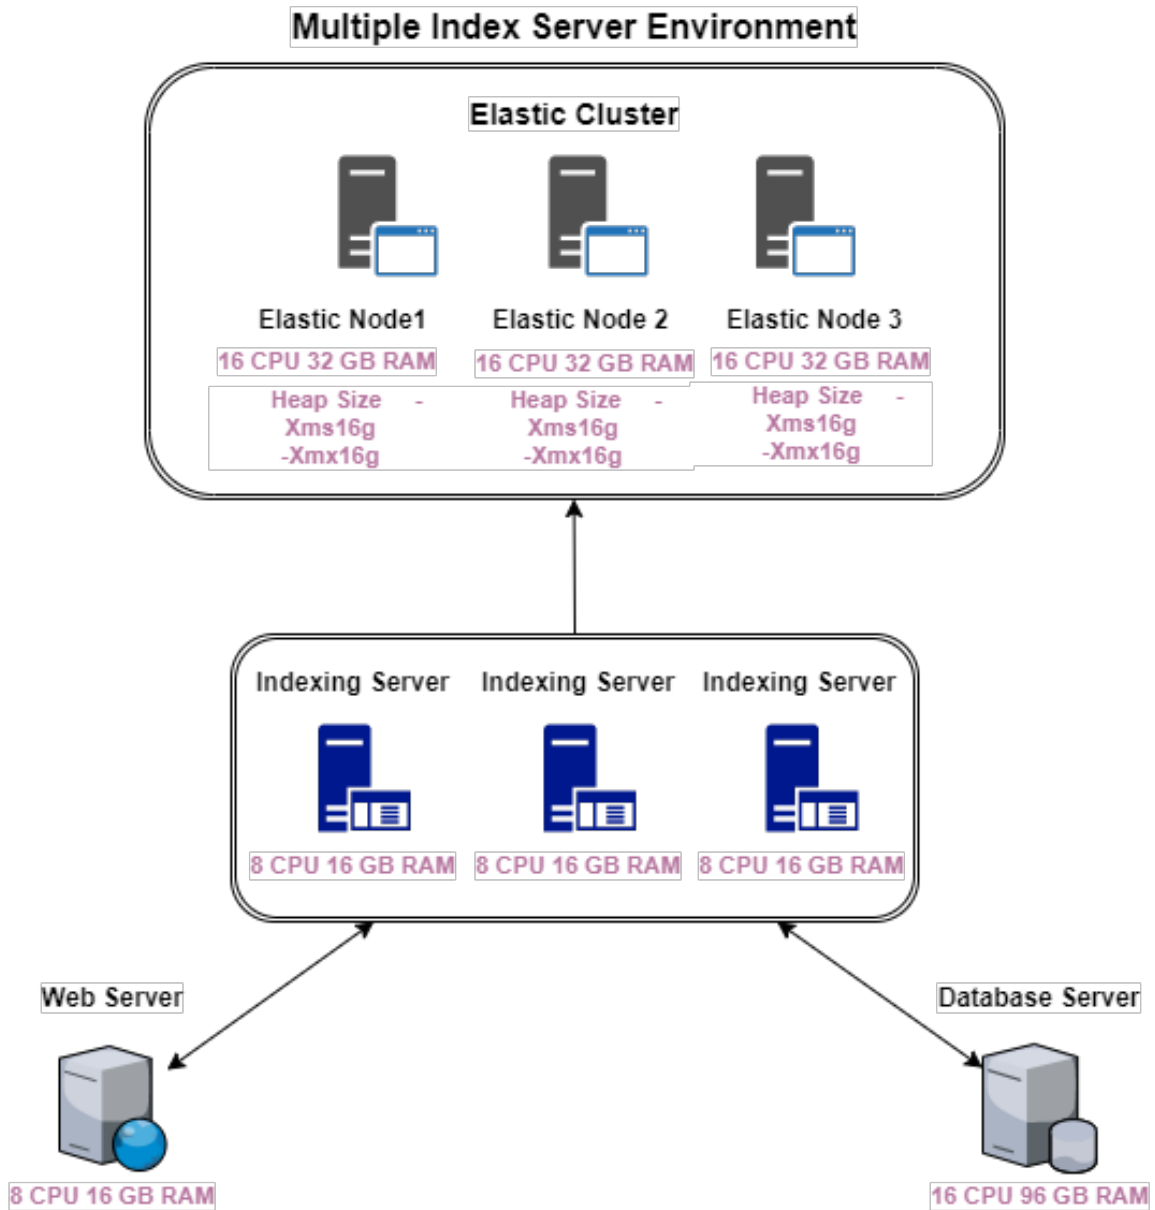
Indexing Service (Elasticsearch)

Elasticsearch is a third-party, scalable search indexing tool that you can configure to work with Archer. If your deployment uses Elasticsearch, refer to the following performance analysis notes to optimize your configuration for faster search database re-indexing.

Performance Analysis Notes for Elasticsearch Database Re-Indexing

- Testing produced the results below in the following hardware environments:





- Average transaction response times using Elasticsearch were approximately one second faster than the default DBSearch service.
- The following parameters produced the fastest index re-build times:
 - IndexRebuildContentBatchSize = 1000
 - ContentBatchMessageProcessorCount (Index Rebuild Batch Indexer Count) = 5
 - BatchPollSize = 50 (for environments with multiple index servers)

- Index re-build time depends on the number of content records and the complexity of the records, as estimated by the number of cross-references on any field in a record. Tests in an environment with a single indexing server produced the following index re-build times:
 - Two minutes for a database with 20,000 content records featuring 235,000 cross-references. The maximum number of cross-references on any particular field was 254.
 - One hour for a database with 5,400,000 content records featuring 18,500,000 cross-references. The maximum number of cross-references on any particular field was 93,000.
- Index re-build times decrease dramatically in environments with multiple indexing servers, but improvements become less significant with the addition of each subsequent server.

IT Security Vulnerabilities Program

The Archer IT Security Vulnerabilities Program helps you proactively manage security risks for your organization's IT assets by combining business context, actionable threat intelligence, vulnerability assessment results, and comprehensive workflows in one place. If your deployment uses this feature, refer to the following performance analysis notes to understand how the scope and complexity of your IT Security Vulnerabilities Program configuration can affect load times for applications, dashboards, calculations, and notifications.

Performance Analysis Notes for the IT Security Vulnerabilities Program

Hardware Environment

- Testing produced the results below in an environment with the following hardware configuration:

Server Type	Services	CPUs (Intel Xenon E5-2690)	Memory
Web Server	Web Application (IIS), Configuration Service, Advanced Workflow Service	16	128GB
Services Server	Job Engine, Queuing Service, Indexing Service, Configuration Service	8	64GB
Database Server	Database	16	128GB

Note: All servers used the Microsoft Windows Server 2012 R2 Standard operating system.

- CPU and memory resources remained below maximum capacity during performance analysis at

all levels of load.

- A single user accessing the Archer test environment measured all test results.

VSR Load

For the purpose of this performance analysis, load is measured by the number of records in the Vulnerability Scan Results (VSR) application in an Archer deployment that returns several hundred vulnerability detections for several hundred IT Security Vulnerabilities Program-enabled devices. VSR load is divided into the following size categories:

- Small: 5 million records
- Medium: 10 million records
- Large: 15 million records

Applications

The VSR load in your Archer deployment significantly affects performance of the following applications:

- Vulnerability Scan Results (VSR)
- Vulnerability Library
- Devices

Note: It is recommended to limit the number of calculated fields in records for these applications. Excessive use of calculated fields can further affect application performance.

Application Load Time

- Load time for the VSR application increases at a linear rate proportional to the number of VSR records (up to approximately 90 seconds for a large load).
- The Faceted Search Panel loads slower than the main application page (up to 150 seconds for a small load), and does not support medium or large loads in the VSR application.
- Other applications with fewer than 200,000 records typically load in under one second (including the Faceted Search Panel).

Record Creation and Save Time

- Interface load time when adding a new record in the VSR application increases at a linear rate proportional to the number of VSR records (up to approximately 85 seconds for a large load).
- Interface load time when saving content in the VSR application remains constant for small and medium loads (approximately 12 seconds), but increases dramatically for a large load (up to approximately 90 seconds).
- Load time for creating new records and saving content in other applications with fewer than 200,000 records is relatively consistent for all loads (typically less than 20 seconds).

Advanced Search and Faceted Search

- For all levels of load, response times for Advanced Search and Faceted Search increase dramatically when you select fields associated with Global Value Lists (such as Potential Operating System Vulnerabilities, Hardware Vulnerabilities, and Application Vulnerabilities) in the search parameters.
- For all levels of load, response times for Faceted Search in the Devices and VSR applications are much longer than those for Advanced Search.
- Faceted Search is not supported for medium and large loads due to timeout issues.
- Session size for searches in the VSR application increase at a linear rate proportional to the number of VSR records (up to approximately 84 megabytes for a large load).

Dashboards

- For all levels of load, CPU utilization on the Web Server and Services Server remains minimal when loading dashboards.
- Dashboard load times increase at a linear rate proportional to the number of VSR records (up to approximately 122 seconds for the IT Security Vulnerabilities Program Business and Executive Management dashboard with a large load).
- iView request and response sizes remain constant for all loads.

Data Feeds (Qualys, NVD, Tenable Security Center)

- It is recommended to run data feeds with "Optimize related calculations after data feed completes" enabled.
- Server CPU and memory utilization are minimal when running data feeds for all levels of load.
- It is recommended to schedule data feeds to run outside of peak business hours.

Calculations

- In deployments with 10,000 or more devices, jobs that modify fields (such as Business Unit Manager) in records associated with the devices may require hours to complete.
- It is recommended to schedule jobs that update large numbers of device records (10,000 or more) to run outside of peak business hours.

Job Engine Service

The RSA Archer Job Engine service administers all asynchronous job processing. This is the workhorse of your Archer configuration and runs on the Services Server.

Asynchronous jobs run two tasks that initiate and execute independently.

The Job Engine service processes all asynchronous jobs, including:

- Data feeds
- Findings generation
- Notifications
- Recalculations
- System jobs

Some asynchronous jobs are purely administrative, whereas others are generated by user activity. Use cases that frequently use notifications or calculations generate more asynchronous jobs as the number of users increases.

The following sections describe tasks that you can review on a regular basis to ensure your system is operating at peak potential. These best practices are based upon the size of your environment.

- [For All Environments](#)
- [For Extra Small to Medium Environments](#)
- [For Large and Very Large Environments](#)

Job Engine Service Tasks to Check for All Environments

As your system grows and you need to adjust your Job Server settings, it is recommended that you follow these general practices regardless of your environment size:

- Add more Memory (RAM), available CPU cores, or both, to improve the performance of:
 - Data Feeds
 - Calculation jobs
 - Large numbers of individual notifications
- Make gradual changes to Job Engine service threads. A thread is the smallest sequence of programmed instructions that can be managed independently by a scheduler. Adding threads can improve system performance. For example, a Job Engine service set for 10 threads can benefit from a higher setting of 12. Observe system performance, and make additional adjustments following the same approach.

Important: Do not change the threads setting from 10 to 20 without first observing the effects of the interim settings.

- Allocate Job Engine service thread settings properly:
 - Thread over-allocation occurs when an increased number of configured Job Engine threads negatively affects job completion and job queue lengths. If average CPU utilization during Job Engine activity is above 80%, decrease the number of configured threads, and consider dedicating another Service Server to the Job Engine service.
 - Thread under-allocation occurs when a very low thread setting creates a sequence of backlogged jobs, causing significant delays until completion. Gradually increase the number of threads, observe system performance, and make additional adjustments following the same approach.
- Dedicate additional Service Servers to the Job Engine to reduce pressure on available resources. You can dedicate a Service Server to a particular type (or several types) of Job Engine activity:
 - For environments dedicating more than one Service Server to the Job Engine service, the Discontinue Job Processing option within the Archer Control Panel Job Engine Plug-in affects all configured Service Servers in the environment, not just the currently selected one. For more information, see "Managing Asynchronous Job Processing" in the Archer Control Panel Help.
 - The use of virtualized Service Servers for the Job Engine service is acceptable, so long as users configure servers according to best practices, and properly allocate CPU and memory resources. For more information, see [Virtualization](#).

Job Engine Service Tasks to Check for Very Small to Medium Environments

For very small to medium environments, the default Job Engine configuration settings provide an adequate degree of asynchronous job throughput and acceptable job queue length. It is recommended that you follow these best practices to further improve performance:

- Configure smaller Archer environments with less intensive workloads with the Job Engine service running on the same server as the Web platform components.
- Adjust Job Engine threads:
 - For unconstrained Job Engine CPU and Memory resources, increase the Job Engine threads.
 - For overwhelmed Job Engine CPU and Memory resources, decrease the Job Engine threads

Job Engine Service Tasks to Check for Large and Very Large Environments

For large scale Archer environments, improve performance by modifying the default Job Engine service settings:

- Separate the Job Engine service from the Web Application component.
- Give the Job Engine service a dedicated Services Server.
- Dedicate one Services Server to particular Job Engine activities, such as offline access functionality.

Native SAML 2.0 for Single Sign-On

Before the Archer 6.8 platform, no native SAML 2.0 support existed. This feature provides native SAML 2.0 support for single sign-on (SSO) in the Archer platform. Previously, with Federation Single Sign-On option, Microsoft ADFS service was also required when using a SAML-capable Identity Provider. With native SAML 2.0 support, choosing the SAML option in the Archer Control Panel allows Archer to communicate directly with Identity Providers that support SAML 2.0. Administrators continue to use the Federation option for SSO when using ADFS as the Identity Provider. Switch to SAML mode for ADFS configurations and utilize the Signing and Encryption certificate option under SAML mode for enhanced security of authentication assertions issued by ADFS.

Performance Analysis Notes for Native SAML 2.0 for Single Sign-On

No significant time difference was observed when establishing an Archer session through manual or SAML login.

When “Enable User Update” was turned on in the Archer Control Panel, with fields such as FirstName, LastName, Company, Title, PhoneNumber, FullAddress etc., updated through the SAML login process, no significant overhead was observed in the login process. The login time with and without profile updates remained approximately the same.

No Impact Factors

The following factors have no impact on Archer performance:

- [APIs](#)
- [Inline Editing](#)
- [IP Whitelist](#)

APIs

An Application Programming Interface (API) is a mechanism for loading data into Archer. Archer employs a few different APIs:

Web Services API

The Web Services API is ideal for operating on smaller hardware, and designed for transactional integrations with other systems. However, during a single bulk load, using a data feed is recommended for 30% increased performance speed.

The BCM Mobile application communicates with the main Archer environment through the Web Services API. BCM Mobile has no Job Engine jobs associated with it.

RESTful API

The RESTful API call loads records one at a time, but loops these calls in such a way as to load very large quantities of data quickly. The RESTful API is called through any Archer Web Servers in a given deployment. It has a minimal effect on other users accessing the system at the same time, save for normal behavior around recalculations.

Content API

The Content API acts as a translation layer that converts Archer Metadata and Content into logical entities so that those entities can be consumed in the same way that they are by users in the Archer UI.

Inline Editing

Inline editing enables users to make changes to data content directly from search pages. Inline editing can have a minor to moderate impact on Archer performance, depending upon how many fields underlying search pages return, and whether Expand All is enabled.

IP Whitelist

The IP Whitelist feature enables the Archer administrator to limit access to Archer to a specific list of known IP addresses.

Offline Access

The Offline Access feature can have a moderate to heavy performance impact on data synchronization between Archer and remote devices. The degree of impact for this feature depends on:

- The complexity of the applications subject to synchronization
- The volume of content data within the synced applications
- The number of concurrent users engaged in the synchronization process

Offline Access uses the RESTful API and Archer Job Engine service for incoming and outgoing synchronization processes. Environments experiencing heavy Offline Access usage benefit from using a Services Server dedicated to the Job Engine service.

Peak User Load

User load is critical to properly sizing the environment. Size your deployment with the expected peak user loads in mind.

Sizing metrics provided in this guide assume that all users are logged in with an active session, and performing basic actions such as:

- Adding new content
- Updating existing content
- Viewing dashboards and iViews
- Browsing reports

Certain user actions are more intensive than others. For example, view, edit, and save activities cause less impact on system resources than browsing workspaces and dashboards.

If you encounter performance issues, monitor the performance of Archer during the peak user loads, and consider adding system resources to the Services Server.

Performance Monitoring Counters

The Service accounts used to run the Archer Services need to be in the Power Users group for the Performance Monitoring Counters to properly work.

Note: More details about each counter are available in the description of the counter itself.

The following Windows Server Performance Monitoring Counters are available:

Archer

- ContentGet:Average
- Content Get Rate
- ContentSave:Insert Rate
- ContentSave:Insert Success Rate
- ContentSave:Total Average
- ContentSave:Total Rate
- ContentSave:Total Success Average
- ContentSave:Total Success Rate
- ContentSave:Update Rate
- ContentSave:Update Success Rate

Asynch Framework Counters

- % Job Threads Active]
- Job Threads Running
- Job Enqueued/Sec
- Jobs Started/Sec
- Jobs/Succeeded/Sec

Queuing Service

- Content Records/Second
- Content Records Queued
- Content Threads Running
- Document Objects Pending Flush
- Document Objects Pending Processing
- Index Messages Queued
- Pending Optimizations
- Repository Records/Second
- Repository Records Queued

Performance Testing Guidelines

Using the following guidelines is recommended to construct effective performance tests for the Archer Platform:

- When adding new content to an application, consider the anticipated number of records that will be added each year in your production environment. Performance tests can generate more activity in an application over several minutes than would typically occur in a year of use in production.
- Consider how a real user will work when using an application. Specifically:
 - Users do not typically log in, perform a single transaction, and then immediately log out. Tests should simulate multiple actions during each login session, over a set period of time.
 - All users in a production environment do not typically log in simultaneously. Tests should simulate multiple user logins spread across a realistic period of time.
 - Users typically require time to think between the sub-steps of an operation. Tests should simulate realistic delays between steps in each operation, such as between a lookup action and the selection of a returned value.
 - Users typically do not perform a Save operation after modifying each individual field in an application or a questionnaire. Tests should simulate Save operations spread across a realistic period of time.
- Consider the flow of the application when designing a load test. For example, executing a Save and Close operation for a content record will return the user to the search page from which the record was opened or added, and execute the search again. Tests should simulate whether users would typically execute another search from the page, or continue to work with the current search results.
- Consider whether references to related records in an application will be spread across a range of records, or whether they will all reference the same record. Tests should simulate references to related records as they will likely exist in your production deployment.

Record Count

Record count refers to the amount of data in an environment. Record count is critical to consider when sizing the environment.

It is recommended that you check your record count at regular intervals. Keep the short-term and long-term growth in perspective. If you anticipate a high influx of data, consider scaling up Archer Databases and Archer Web Servers. Environments with more records than planned run into processing issues.

If a tie occurs in the results, tie resolution considers only the first sort criteria—returning all matching values for the last spot and potentially exceeding the defined limit. Tie resolution works against numeric, date, text, and values list field types.

For more information on determining environment size based on record counts, see [Chapter 2: Sizing Guidelines](#).

For more information on configuring record count, see "Running Searches in Applications and Questionnaires" in the Archer Online Documentation

Relationship Visualization

The Relationship Visualization feature provides the ability to display content relationships by means of an image or graphic. This feature can have a minor impact on Web Application performance, and a minor to moderate impact on overall configuration performance.

Retrieval and render time for visualized data depends on the amount of related data associated with the element being visualized.

For the best user experience, use an up-to-date and fully supported web browser. For more information on proper browser usage, see [Browser Choice](#).

Archer Mobile

Important: This section does not apply to the Archer Suite mobile application.

The Archer Mobile application communicates with the main Archer environment through the Archer RESTful API, which is accessed by any of the Web Servers in a deployment. This enables a significant number of users for mobile off the main web server without any noticeable impact on the other users of the system.

Mobile-enabled users syncing content with their devices can cause load on the system in two places:

- Communicating with the Web Server over the API
- Processing of tasks related to preparing the data for downward sync in the RSA Archer Job Engine service.

It is recommended that you complete the following tasks at regular intervals:

- During the busiest times of user activity, check how many users are active. Performance degradation occurs when over 100 users simultaneously sync data. To increase performance, add more Web Servers.
- During large rollout of the mobile application, allocate a separate Web Server to act as the API endpoint.

Environments with very busy Service Servers dedicated to the Job Engine should assess whether this may negatively impact their user's experience, due to the dependency for downward sync on the Job Engine. To increase performance, add more Services Servers dedicated to the Job Engine.

Transactional Density

Transactional density is a variable factor that affects the performance of your Archer deployment. The average number of concurrent users divided by the average delay between the actions each user performs in the Archer web application equals the transactional density for your environment. While services and browser components also affect performance measurements, most user interactions with the Archer Platform take place over an HTTP GET/POST paradigm, so it is important to consider the implications of this paradigm when simulating load during performance testing.

With some exceptions, a user who is logged into the user interface but not actively performing actions does not generate load on the web server. For example, there is no performance impact while a user reads a page that has already loaded or while the user types into the fields on a page. Creating user sessions and reserving memory for each session generates some overhead on the web server. However, this load applies concurrently to the Archer Platform only when users send traffic over the network, such as by clicking Save to commit changes on a page or by entering data in a field that triggers a Data Driven Event (DDE).

When testing performance in your existing Archer deployment, or determining hardware requirements when planning a new deployment, it is recommended that you estimate transactional density using predictable values, such as a set number of users that perform a series of defined tasks repeatedly and at a specific frequency. This method creates an artificially high transactional density, because real world users typically pause while inputting data or take breaks between performing tasks.

Using this recommended evaluation method, the performance metrics captured for an environment with 200 concurrent users and a 10 second average delay between actions are approximately equivalent to metrics for 1000 concurrent users with a 50 second average delay.

Use Case Design Best Practices

When designing applications in Archer, administrators should be aware of the impact design decisions can have on the end-user experience. Rather than limit the user, Archer can support highly complicated use case designs, keeping in mind there may be a trade-off in terms of end-user experience.

The following choices may lead to degraded performance:

- Allowing N-tier searches over five or more levels
- Returning or exporting over 10,000 records
- Using more than seven filters in a report
- Displaying more than 500 fields in an application or questionnaire
- Displaying more than 10 root tabs, or 20 child tabs in an application or questionnaire

- Implementing more than 100 Data Driven Events (or more than 25 Apply Conditional Layout Data Driven Events) in an application or questionnaire
- Using more than 10 record permissions fields in an application or using more than 20 inherited record permission rules
- Allowing cross-reference or history log display grids to display every record slows load times in certain records
- Setting cross-reference fields on records to return more than 20 fields or more than 500 records slows the record lookup process
- Avoiding placing a history log field on the cross-reference side of any relationship containing large numbers of records; a better choice is on the related record side
- Adding Text Area Fields to application or questionnaire slows page load times

In isolation, any of these options may be the right approach, but administrators should be aware of the implications.

Virtualization

Virtualization creates a virtual, rather than actual, version of an operating system, a server, a storage device, or a network resource.

Completing the following tasks at regular intervals is recommended:

- Follow accepted IT industry best practices for configuring, monitoring, and maintaining virtual servers.
- Keep the allocation of your virtual machine (VM) host below 100% of its physical capacity.
- Leave at least 20% of the underlying host machine's physical capacity available to absorb usage spikes.
- Use at least 24 GB RAM on dedicated VM host hardware for a large environment configuration. If your virtual environment configuration does not support 24 GB RAM, consider adding more VMs to compensate the physical system configuration requirements.
- Set resource reservations for processor and memory to ensure availability for the VMs running Archer. Adjust these reservations later to tune performance.

Completing the following tasks is recommended if your environment is failing:

- Employ multiple processor cores for the Web Server. Microsoft Hyper-V and VMware vSphere products can support four or more virtual processors in guest mode, but default to one. Adding at least one virtual processor to the default (for a total of two or more) is recommended due to multithreading in Archer and thread pooling by running under ASP.NET. For more information on threading, see [Job Engine Service](#).
- Dedicate spindles or a storage area network (SAN) to lessen the likelihood of disk or input/output contention issues.