SYSTEM SPECIFICATIONS GUIDE

AD eDiscovery®

COLLECT, AUDIT & ANALYZE

Revision: April 2020

ACCESSDATA®

www.accessdata.com

Contents

AD eDiscovery Overview

AD eDiscovery is a fully integrated software platform designed to help organizations preserve, collect, process, review, and produce ESI (Electronically Stored Information) in accordance with EDRM model and FRCP rules regarding evidence preservation.

The solution is entirely web-based and scalable to meet even the most demanding e-discovery challenges. AD eDiscovery is comprised of a series of functional components that allow the solution to be customized to meet the customers' collection, processing and production requirements. All components can be installed on a single server or distributed in various ways across multiple servers depending on the organization's needs and desired workflow.

The following section contains a brief explanation of each of the AD eDiscovery components and its role within the solution:

- **Web Suite** – The Web Suite provides the interface through which users access AD eDiscovery.
- **Application Services**
    - **Windows Communication Foundation Services**—The Windows Communication Foundation Services ("WCF") manage the flow of data between the various AD eDiscovery components.
    - **Asynchronous Processing Services**—The Asynchronous Processing Services ("Async") are responsible for the execution of certain user actions such as bulk coding, searching, and load file import ingestion.
    - **Work Manager**—The Work Manager governs the flow of work to the Processing Engines and Site Servers, as well as performing collections of structured data sources such as Microsoft Exchange and Microsoft SharePoint.
- **Processing Engine**—The Processing Engine performs data processing tasks such as the expansion of archives (e.g., .PST, .NSF, and .ZIP files), indexing, de-duplication analysis, file identification, secondary culling and filtering, and the creation of production and export sets.
- **Site Server**—The Site Server is generally responsible for managing communication with from the Work Manger through to Agents and for performing collections from Network Share sources.
    - **Root Site Server**—The Root Site Server controls the overall flow of information from the Work Manager down to the appropriate collection endpoints and of collected data up from the appropriate collection endpoints to the Work Manager.
    - **Private/Private-Protected Site Server**—The Private Site Server manages the collection activity of specified endpoints within the local network which have been initiated by the Work Manager.
    - **Public Site Server**—The Public Site Server manages collection activity initiated by Agents located outside the local network.
    - **PostgreSQL**—PostgreSQL is an open-source object-relational database management system. Each Site Server utilizes a PostgreSQL database to store information about active collections temporarily.
- **Agent**—The Agent is a modular application that can be deployed to targeted computers and performs secure forensic-level access, analysis, and preservation of the target's static data.
- **SQL Database**—AD eDiscovery utilizes a Microsoft SQL Server instance to maintain databases containing file metadata, user data, and workflow information.
- **Project Data/Evidence/Collection Storage**—AD eDiscovery can leverage many types of local or external storage, including network attached storage (NAS), storage area network (SAN), and direct-attached storage (DAS), to host evidence and other case-related data.

General Considerations

Virtualization and Shared Servers

AccessData fully supports the use of either hardware or virtual machine systems as long as the virtual system is configured according to industry best practices. The support of any implementation which attempts to host one or more components on virtualized platforms is subject to the discretion of AccessData. AccessData reserves the right, during the troubleshooting of a support issue, to withdraw support on a specific issue if it is found to be induced by virtualization. Further information regarding virtualization of AccessData products can be found here.

The support of any implementation which attempts to host the SQL Database component on the same hardware platform as other enterprise applications is subject to the discretion of AccessData. Attempts to host the SQL database component in the same instance as other enterprise applications will not be supported.

AccessData forbids the installation of any of AD eDiscovery's components on any system that hosts a Microsoft Domain Controller.

Please contact your AccessData technical support representative for further information.

Service Account

AD eDiscovery requires the use of a single, dedicated, service account to operate properly. In a multi-server installation environment, a domain-level service account is required. Workgroup authentication is only supported for single-server installation environments. In either case, the service account must be a local administrator with "Logon as Service" and "Interactive Logon" system permissions. The password on the service account should not be set to expire at any interval.

The service account must be added to the Logins of the SQL Instance being used to host the SQL Database component.

Certificates

Communications between AD eDiscovery and the web-based end-user interface are protected by Secure Socket Layer encryption (SSL), which requires the use of a public certificate signed by a trusted certificate authority. Some implementations may require the purchase of a properly-configured certificate from a commercial Certificate Authority.

Communications between the Site Servers and the Agents are protected by Secure Socket Layer encryption (SSL), which requires the use of a public/private certificate pair. The certificate pair must contain the complete chain of trust, but does not require a signature from a trusted certificate authority and an internally-managed certificate is recommended.

AD eDiscovery will accept the following Certificate formats:

- Public[1]: .CER, .CRT, .P7B
- Private[2]: .PFX, .P12, .PEM, .ADP12

Additional Recommendations

Windows Indexing

AccessData strongly recommends that the Microsoft Windows Indexing Service be configured to either exclude the directories or drives containing case files, database files, temp/log files or disabled entirely.

Anti-Virus Scanning

AccessData strongly recommends that any anti-virus or anti-malware software on any each server hosting components of AD eDiscovery are configured to disable on-access scanning of the directories or drives containing case files, database files, or temp/log files. Additionally, should any full scans be scheduled, they should be monitored to ensure they are not interfering with the overall performance of the solution.

8.3 Filenames

AccessData recommends disabling the creation of 8.3 character length filenames and updates to the last access timestamp on NTFS formatted volumes to improve performance in disk input/output operation.

Pagefile Configuration

AccessData recommends setting both the minimum and maximum sizes of the system pagefile to double the amount of RAM on the system. For optimal performance, the pagefile can be moved to a dedicated, low-latency (e.g., RAID 0 or SSD) disk space that meets the calculated capacity requirements. For further information, please read http://support.microsoft.com/kb/2860880 or contact your AccessData technical support representative for further information.

General Hardware Requirements

The overall performance of AD eDiscovery is dependent on the hardware employed to host its various components. Ideally, all implementations would employ the latest multi-threaded processors, large amounts of memory, and arrays of solid state disc drives. As this is often unrealistic due to budget restrictions, the following guidelines have been developed to assist in the creation of cost-effective environments that conform to the differing needs of a diverse client base.

Processor and Memory

The quality of the processors employed in the implementation environment will have a direct effect on the overall performance of AD eDiscovery. Sites such as cpubenchmark.net can be used to compare the relative performance of different processors. Additionally, some components use the number of logical processor cores on a system to calculate the total number of threads available to perform certain operations.

Minimum hardware recommendations for some of the components when deployed on their own servers in an enterprise environment can be found below in Table 1 and examples of some common configurations of AD eDiscovery are located in Appendix C. Please contact your AccessData technical support representative for further information and assistance.

Minimum Hardware Recommendations		
System Component	**CPU's**	**Memory**
Web Suite	4 logical cores	4GB RAM
Application Services	4 logical cores	16GB RAM
Processing Engine	8 logical cores	16GB RAM
Site Server	4 logical cores	8GB RAM
SQL Database	8 logical cores	32GB RAM

Table 1 - Minimum Hardware Recommendations

During certain operations, components in AD eDiscovery can leverage all available processor and memory resources available to the host system. Systems with insufficient memory resources can experience bottlenecks as certain operations may cause the system to start paging. The presence of any paging on a system will result in an associated reduction in the performance of the solution and severe paging—also known as "thrashing"—can lead to operational failure.

It is strongly recommended that any system involved in the implementation environment possess at least 2GB of RAM for each logical processor core (e.g., an 8-core system should have at least 16GB of RAM) to reduce the likelihood of paging. Additionally, it is recommended that the system hosting the SQL Database component possess at least 4GB of RAM for each logical processor core (e.g., an 8-core system should have at least 32GB of RAM).

Storage

The storage requirements of AD eDiscovery are dependent on a number of variables including the number of active projects, the volume of data involved in the projects and the workflow of the organization. Both the back-end storage hardware being employed and its configuration can greatly affect the overall performance of AD eDiscovery. Table 2 contains descriptions, characteristics, and recommendations on the various types of storage involved in AD eDiscovery.

Minimum Hardware Recommendations		
	Description	**Storage Characteristics**
Operating System and Applications	Local disk volume on any system hosting one or more components that provides storage for the operating system and application files.	The initial space requirements should include 40GB for the operating system and additional space sufficient to accommodate the components being hosted. Systems with more than 16GB of RAM will require additional space to accommodate the system pagefile. This storage should be fault-tolerant. **Recommendation: RAID 1.**
Staged Evidence	File share on either a local disk volume or network storage that provides storage for data that will be ingested as evidence or imported via loadfile (e.g., forensic images, native files, TIFF images, PDF images, OCR text files, and loadfiles).	The initial space requirements are dependent on the needs of organization, but can be significant. This storage should be fault-tolerant with low latency. **Recommendation: RAID 10 or RAID 5.**
Case Data	File share on either a local disk volume or network storage that provides storage for case-specific data, application-generated files, and internally-maintained copies of specific types of ingested data.	The initial space requirements for ingested evidence are roughly 33% of the space of the associated staged evidence and the initial space requirements for imported data are 100% of the space of the associated staged evidence. Additional space will be required to support ongoing workflow operations. This storage should be fault tolerant with low latency. **Recommendation: RAID 10 or RAID 5.**
Exported Data	File share on either a local disk volume or network storage that is used as a target for exported native files, TIFF images, PDF images, and loadfiles.	Exported data is separate from the associated records in a case and can be purged to reduce the requirements of this storage space. The space requirements and fault tolerance are entirely dependent on the organization's workflow. **Recommendation: None.**
SQL Databases	Local disk volume on the system hosting the SQL Database component that provides storage for the system and application database files.	The initial space requirements are roughly 33% of the space of the associated staged evidence. Additional space will be required to support ongoing workflow operations. This storage should be fault tolerant with low latency. Recommendation: RAID 10 or RAID 5.
SQL Logs	Local disk volume on the system hosting the SQL Database component that provides storage for the system and application database log files.	The initial space requirements are dependent on the size and number of databases and the frequency of database maintenance operations, but will be smaller than the space required for the SQL Databases. Additional space will be required to support ongoing workflow operations. This storage should be fault tolerant. **Recommendation: RAID 1.**
Temp DB	Local disk volume on the system hosting the SQL Database component that provides storage for the temporary database files.	The space requirements are dependent on the frequency of database maintenance operations. The speed of this space is important. This storage requires no fault tolerance. **Recommendation: RAID 0 or SSD.**
ADTemp	Local disc volume on any system hosting the Processing Engine and Site Server components that provides storage for ephemeral files generated by the Processing Engine and Site Server components.	At least 50GB of space is required, but a minimum of 500GB is recommended. The most important characteristic of this space is its speed. This storage requires no fault tolerance. **Recommendation: RAID 0 or SSD.**

Table 2 - Storage

For optimal performance, initial consideration should be given to the seek time, latency, and data transfer rates of the storage. High disk activity can be expected during certain operations and is not necessarily indicative of a problem. Sustained rates of disk activity above 85% or persistent disc queues over 2 per disk during operations will result in a bottleneck effect and a corresponding reduction in the overall performance of the solution.

Note: Sustained periods of high disk use and persistent disk queues can be a symptom of insufficient memory resources. Please see the recommendations found in the "Processor and Memory" section on page 6.

Ongoing attention should also be paid to the space utilization and fragmentation of the storage which can themselves lead to a decrease in performance. There are a number of different methods by which disc queuing and fragmentation issues can be addressed including the use of high-RPM drives, RAID technologies[3], or solid-state drives (SSD).

Network

AD eDiscovery is a componentized, web-based platform. Communication between the various components is performed over Transmission Control Protocol (TCP) ports as depicted in Figure 1. A more comprehensive list of the ports used for communication can be found on the next page in Table 3 on page 9.

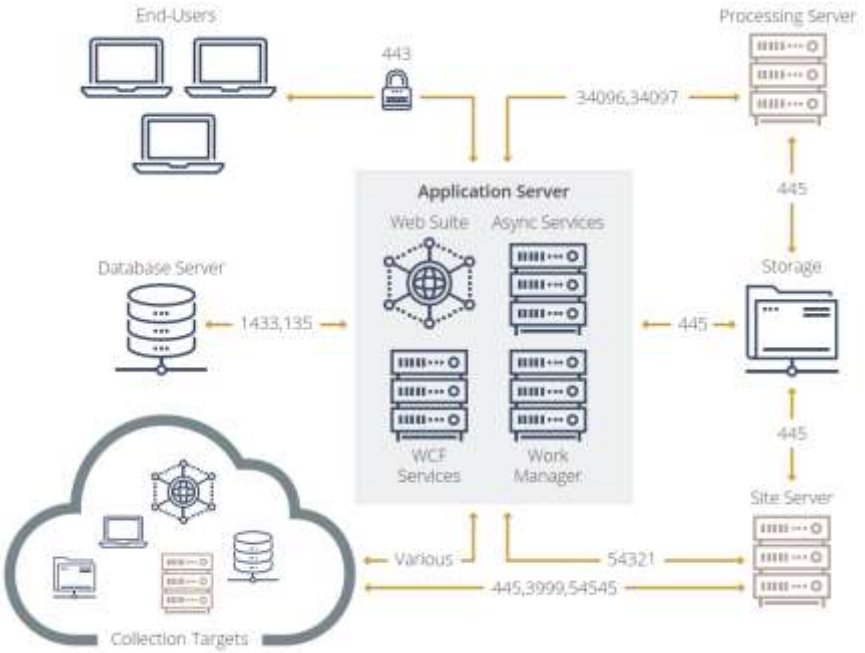

Figure 1 - Simplified AD eDiscovery Network Communication Diagram

It is important to note that some of the ports listed below are only used to negotiate the connection between two components. The actual communication taking place between the components will use ephemeral ports in the dynamic port ranges of the respective servers[4]. All specified ports must be open in both directions.

TCP Ports Employed By the Various Components		
Source Component	**Destination Component**	**Port**
Web Suite	Asynchronous Processing Services Case Data/Evidence Storage End-Users SQL Database Windows Communication Foundation Services Work Manager	80/8080 445 443 1433 9132 9132
Windows Communication Foundation Services	Case Data/Evidence Storage SQL Database Web Suite Work Manager	445 1433 9132 9132
Asynchronous Processing Services	Case Data/Evidence Storage SQL Database Web Suite	445 1433/135 80/8080
Processing Engine	Case Data/Evidence Storage SQL Database Work Manager	445 1433/135 34096/34097
Work Manager	Case Data/Evidence Storage Processing Engine SQL Database Web Suite Windows Communication Foundation Services Site Server	445 34096/34097 1433 9132 9132 54321
Site Server	Agents Network Shares (Collection) Project Data/Evidence Storage Site Server Work Manager	3999/54545 445 445 54545 54321
Agents	Site Server	3999/54545
SQL Database	Asynchronous Processing Services Processing Engine Web Suite Windows Communication Foundation Services Work Manager	1433/135 1433 1433 1433 1433
Project Data / Evidence / Collection Storage	Asynchronous Processing Services Processing Engine Web Suite Windows Communication Foundation Services Work Manager	445 445 445 445 445

Table 3 - TCP Ports Employed By the Various Components

General Software Requirements

AD eDiscovery has been designed to leverage Microsoft server technologies. The following table contains the software prerequisites for each of the components with such requirements.

Software Requirements	
Source Component	**Destination Component**
Web Suite	Microsoft Windows Server 2012 R2, or 2016[5] Microsoft Internet Information Services 7.5 Microsoft Distributed Transaction Coordinator *Microsoft .NET Framework 4.0*[6] *Microsoft Visual C++ 2010 x64 Redistributable*
Application Services	Microsoft Windows Server 2012 R2, or 2016 Microsoft Distributed Transaction Coordinator *Microsoft .NET Framework 4.0* *Microsoft SQL Server 2008 R2 Management Objects (x64)* *Microsoft SQL Server System CLR Types (x64)* *Microsoft Visual C++ 2010 x64 Redistributable* *Microsoft Visual C++ 2010 x86 Redistributable* Microsoft Outlook 32-bit[7]
Processing Engine	Microsoft Windows Server 2012 R2, or 2016 Microsoft Distributed Transaction Coordinator *Microsoft .NET Framework 4.0* *Microsoft Visual C++ 2008 Redistributable* *Microsoft Visual C++ 2010 x64 Redistributable*
SQL Database	Microsoft Windows Server 2012 R2, or 2016 Microsoft SQL Server 2012 or 2014 or 2016 or 2017
End-User Computers	Microsoft Internet Explorer 9+ Microsoft Silverlight 5.1 Adobe Flash 11.7 AccessData NearNative Viewer AccessData Bulk Print Local

SQL Database

The SQL Database component is the heart of AD eDiscovery and its performance is crucial to the overall performance of the application. Microsoft SQL Server operates under the assumption that the server hosting it exists solely to host its databases. Understanding this behavior and the reasoning behind it is important to the performance of AD eDiscovery, especially in implementation environments in which the SQL Database component is sharing a server with additional components. AccessData recommends that a qualified Database Administrator assist in both the initial configuration and ongoing maintenance of the SQL Database component. More recommendations on the configuration of SQL Server 2016/2017 are found here.

Microsoft SQL Server will cache the data it reads from storage in memory to improve its performance and will cache entire databases if it has the resources available to do so. The benefit of this behavior is that adding memory to the server hosting the SQL Database component can be expected to improve its performance. The drawback of this behavior is that Microsoft SQL Server's default settings allow it to claim up to 2 petabytes of memory. AccessData recommends that the Maximum Server Memory setting in Microsoft SQL Server be set to reduce the likelihood of the SQL component claiming all of the server's available memory.

The storage used by the SQL database component also plays an important role in the application's overall performance. AccessData recommends that the SQL data files, the SQL transaction log files, and the TempDB database are physically segregated from each other and from the operating system. Ideally, SQL data files should be located on storage with high read-write performance and redundancy; SQL transaction log files should be located on storage with high write performance and redundancy; and the TempDB should be located on storage with the fastest possible read-write performance, but does not require any redundancy. For more information, please see http://technet.microsoft.com/en-us/library/cc966534.aspx or contact your AccessData technical support representative.

AccessData software will create a multiple Logins/Users (used internally by our software) within the SQL database. Because of this, Microsoft SQL Server must be licensed based on CPU Core count rather than by User count (CAL).

SQL Server Requirements

The support of any implementation which attempts to host the SQL Database component on the same hardware platform as other enterprise applications is subject to the discretion of AccessData. The use of a single SQL instance to host the SQL database component and any other enterprise application is not supported.

AccessData requires that the SQL instance being used to host the SQL Database component is created using the Default US Collation, "SQL_Latin1_General_CP1_CI_AS."

AccessData requires that the SQL Instance being used to host the SQL Database component must have Mixed Mode Authentication enabled and the Service Account must be added as a "sysadmin" to the instance.

AD eDiscovery will create a separate database for each project, as well as associated SQL logins. To ensure the proper operation of the platform, no modifications should be made to any logins generated in this manner.

AccessData requires that the Network Configuration of the SQL Instance being used to host the SQL Database component must have the TCP/IP and Named Pipes Protocols enabled.

Database Maintenance

Databases require ongoing maintenance to prevent poor application performance, system downtime, and data loss.[8] There is no one-size-fits-all solution to database maintenance and regular attention must be given to ensure the continued successful operation of any maintenance plan, but the implementation of a basic maintenance plan is a relatively simple undertaking. The following guidelines can be used to assist in the development of a basic Microsoft SQL Server maintenance plan.

For additional information on SQL database maintenance, please see the footnotes included within this section or contact your AccessData technical support representative. Additional assistance with database maintenance, including the development of customized plans and professionally-staffed on-going maintenance, is available through AccessData's Support Services department. Please contact your AccessData salesperson for additional information and pricing.

Database Recovery Model

The selection of a database recovery model[9] is the first decision that must be made when developing a SQL maintenance plan. The recovery models provided by Microsoft SQL Server are meant to address varying levels of resource availability and acceptable data loss.

AccessData recommends the use of the Full Recovery model with user databases, but supports the use of either the Full Recovery model or the Simple Recovery model.

Database Backup Strategy

A database backup strategy[10] is generally the focus of any maintenance plan. While primarily meant to protect against data loss, database backups may also be necessary to address other significant maintenance requirements. Microsoft SQL Server supports three primary database backup methods: Full, Differential, and Log:

- The Full backup method creates a complete record of a database. A Full backup record provides the ability to restore a database to a single point-in-time.
- The Differential backup method requires the existence of a Full backup and creates a record of any extents that have been modified since the Full backup was created. A Differential backup record in combination with its associated Full backup record provides the ability to restore a database to a single point-in-time.
- The Log backup method creates a record of all transactions made in a database since the last Log backup. A Log backup record in combination with its associated Full and Differential backup records provides the ability to restore to any point from the time of the Full backup record to, contingent on the success of a tail log backup, the most recent transaction in the database.

The Full and Differential backup methods are available to both the Simple and Full recovery models. The Log backup method is unavailable under the Simple recovery model. The output of any backup method should be directed to a location that is not being used to store active SQL database files (e.g., MDF, NDF, or LDF files).

Note: Log backups must be taken for any database using the Full recovery model; the file containing the database's transaction log will otherwise continue to grow indefinitely.

AccessData strongly recommends that, at minimum, full backups of both system and user databases are made regularly. Additional complexity and scheduling will be dictated by criteria such as acceptable work-loss exposure, the speed and volume of storage available for both the data files themselves and the backup records, and the maintenance's impact on the overall performance of the application.

Database Index Optimization

Microsoft SQL server uses indexes[11] to quickly query commonly used data and improve operation. Rebuilding and reorganizing these indexes is important to the overall performance of the application.

As modifications are made to the tables within a database, the associated indexes will be subject to internal (i.e., excessive, unused memory allocated to the indexes) and external fragmentation (i.e., indexes that are stored non-sequentially) which can degrade performance. Regular reorganization (i.e., reordering an index within its existing allocated memory) and rebuilding (i.e., reordering the index into freshly-allocated contiguous memory) of fragmented indexes is necessary to counteract the results of this activity.

AD eDiscovery performs database index optimization following certain activities, but AccessData recommends performing scheduled index optimization regularly in concert with Full backups[12].

Database Integrity Checks

Database integrity checks[13] are a method by which any logical or physical issues in a database can be identified. Depending on the severity of an identified issue, a database can either be repaired or restored to a point prior to the genesis of the damage.

AccessData recommends that integrity checks be run prior to any Full backup to help ensure the integrity of the database contained within the backup.

Maintenance Cleanup

Maintenance cleanup[14] is a necessary piece of any maintenance plan which must be explicitly run to remove old backup files and other unnecessary maintenance records.

AccessData recommends establishing a regular maintenance cleanup schedule based on the Full backup schedule and organizational backup retention policies.

Appendix A: Pre-implementation Checklist

The following checklist should be used to document the prerequisites necessary to ensure the successful implementation of AD eDiscovery and should be completed prior to product implementation by an AccessData technician.

1. Hardware Information

- ☐ 1.1. The servers that have been designated for component configuration are available
- ☐ 1.2. The servers' operating systems have been installed and are fully-patched.
- ☐ 1.3. The servers' storage volumes have been properly provisioned and formatted.
- ☐ 1.4. The Microsoft DTC feature on each server with a cloned operating system has been uninstalled and reinstalled (i.e., ensure each server has a unique CID).

2. Network Configuration

- ☐ 2.1. The appropriate ports are open between the servers (see page 9).

3. Service Account

- ☐ 3.1. A dedicated service account named_____has been created.
- ☐ 3.2. The service account has been added to the local Administrators group on all servers in the environment.
- ☐ 3.3. The service account has been provided with the Interactive Logon permission.
- ☐ 3.4. The service account has been provided with the Logon As Service permission.
- ☐ 3.5. The service account's password options have been set to Password Never Expires and User Cannot Change Password.

4. SQL Server Configuration

- ☐ 4.1. Microsoft SQL server has been installed and fully patched.
- ☐ 4.2. The SQL instance name is_____(default: "Default").
- ☐ 4.3. The SQL instance is configured to use port_____(default: 1443).
- ☐ 4.4. The SQL instance is configured to use "SQL_Latin1_General_CP1_CI_AS" coalition.
- ☐ 4.5. The SQL instance has Mixed Mode authentication enabled.
- ☐ 4.6. The Service Account has been added to the SQL instance as a user and has been given sysadmin rights.
- ☐ 4.7. Microsoft DTC is enabled.
- ☐ 4.8. Named Pipes have been enabled for the instance.

5. Software Licensing

- ☐ 5.1. The license dongle is accessible and has been properly stocked with the appropriate licenses.

6. Software Installation Media

- ☐ 6.1. The AccessData technician has provided the FTP credentials to retrieve the latest software ISO.
- ☐ 6.2. The latest software ISO has been downloaded and copied to the servers.
- ☐ 6.3. Software capable of mounting an ISO (e.g., WinCDEMU) or extracting from an ISO (e.g., 7-ZIP) has been installed on at least one of the servers.

7. Certificates

- ☐ 7.1. A certificate has been created for use with IIS with this common name:_____.
- ☐ 7.2. A valid certificate pair has been created for use with the Site Server[15].

8. Collection Connector Configuration Information

- ☐ 8.1. The configuration information for each of the desired connectors has been identified and documented (see page 15).

9. Authentication Configuration

- ☐ 9.1. eDiscovery will be configured to use "Forms" / "Active Directory[16]" for user authentication (pick one).

10. Litigation Hold Feature Configuration Information

- ☐ 10.1. The SMTP configuration information for the local email server has been documented[17].

Appendix B: AD eDiscovery Connectors

The section below contains information on each of the connectors available in AD eDiscovery.

Site Server Collections			
Source	**Technology**	**Permissions/Considerations**	**Version Support**
Microsoft Windows	AccessData Agent	Administrator credentials only needed for initial Agent deployment.	Windows Server 2003-2016, Windows XP, 7, 8, 8.1, and 10
Apple Macintosh	AccessData Agent	Administrator credentials only needed for initial Agent deployment.	OSX 10.5.8, 10.6.8, 10.8, 10.9, 10.10, 10.11, 10.12
Linux	AccessData Agent	Administrator credentials only needed for initial Agent deployment.	RedHat 3, 4, 5, 6.5 (32-bit and 64-bit), SUSE 10, 11 (32-bit and 64-bit), Ubuntu 9, 10, 12 (32-bit and 64-bit), CentOS 5 (32-biut and 64-bit), Solaris 8, 9, 10 (32-bit and 64-bit)[18]
Network Share	CIFS	Credential with “Read” permission on the target.	Not Applicable

Work Manager Collections			
Source	**Technology**	**Permissions/Considerations**	**Version Support**
Box.com	HTTPS	Administrative privileges capable of creating a custom Box Application which supports OAuth2 authentication via an API key.	
CMIS	Web Services	Valid credentials on the target	
Cloud Mail	HTTP/HTTPS	Valid credentials on the target	
DocuShare	HTTP	Administrative privileges (or equivalent) to target cabinet	DocuShare 6.5, 6.6
Documentum	DFS	Administrative privileges (or equivalent) to target cabinet	Documentum 6.5, 6.6
Domino	Notes Client	“Replicate or Copy Documents”, “Create Shared Folders/Views”, Write/Read Public Documents”, and “Edit NAMES.NSF file” permissions on the target	Domino 6.5 – 6.8
Druva	WebDAV	Credential with “Read” permission on the target.	Druva InSync 5.4
Enterprise Vault	AccessData Service	The AccessData Integration Service must be running as an account with Administrative access to the target	Enterprise Vault 8 – 10.0.3
FileNet	Web Services	Administrative privileges (or equivalent) to target cabinet	FileNet 5.1
Gmail	HTTP/HTTPS	Administrative privileges (or equivalent) to target cabinet	
Google Docs	HTTP/HTTPS	Administrative privileges (or equivalent) to target cabinet	
MS Exchange EWS	Outlook/EWS	“Receive As” and “View Information Store Status”	Exchange 2010 SP1+, Exchange 2013, Exchange 2016, Exchange 365
MS Exchange MAPI	Outlook/MAPI	“Receive As” and “View Information Store Status”	Exchange 2003, Exchange 2007, Exchange 2010
MS Sharepoint	Web Services	“Web Application Policy” and “SP FrontPage” extensions enabled and administrative privileges (or equivalent) on the target	SharePoint 2007 – 2013 Sharepoint 365
OpenText ECM	Web Services	Administrative privileges (or equivalent) to target cabinet	OpenText 10.0
Oracle URM	Web Services	Administrative privileges (or equivalent) to target cabinet	Oracle URM 10 – 11
Web Crawler	HTTP/HTTPS	Website must be publicly-accessible	

Appendix C: Sample Environments

The section below contains a series of hypothetical hardware configurations that illustrate some of the more common methods used to implement AD eDiscovery.

NOTE: These examples are for demonstrative purposes only and should not be solely relied upon as they may not be appropriate for your environment.

Example 1: Single Server Environment		
Server	**Components**	**Hardware Specifications**
Single Server	• Web Suite • Windows Communication Foundation Service • Asynchronous Processing Services • Distributed Processing Manager • Work Manager • Site Server (Root) • SQL Database • Case Data/Evidence/Collection Storage	Logical Cores: 16 RAM: 32GB Network Connection: 1GbE NIC Drive Sets: • OS/Apps – 150GB RAID 1 SAS • DB Files/Logs – 1TB RAID 1/5/10 DAS/NAS/SAN • TempDB – 100GB RAID 0 SAS or SSD • ADTemp – 500GB RAID 0 SAS or SSD • Case Data/Evidence/Collection Storage – 1TB RAID 1/5/10 NAS/SAN/DAS

Example 2: Three Server Environment (Processing-focused)		
Server	**Components**	**Hardware Specifications**
Web / Application / Collection Server	• Web Suite • Windows Communication Foundation Service • Asynchronous Processing Services • Distributed Processing Manager • Work Manager • Site Server (Root)	Logical Cores: 16 RAM: 32GB Network Connection: 1GbE NIC Drive Sets: • OS/Apps – 150GB RAID 1 SAS • Case Data/Evidence/Collection Storage – 1TB RAID 5/10 NAS/SAN/DAS
Processing Server	• Distributed Processing Engine • Case Data/Evidence/Collection Storage	Logical Cores: 16 RAM: 48GB Network Connection: 1GbE NIC Drive Sets: • OS/Apps – 150GB RAID 1 SAS • ADTemp – 500GB RAID 0 SAS or SSD
Database Server	• SQL Database	Logical Cores: 16 RAM: 32GB Network Connection: 1GbE NIC Drive Sets: • OS/Apps – 150GB RAID 1 SAS • DB Files – 1TB RAID 1/5/10 DAS/NAS/SAN • DB Logs – 500GB RAID 1 SAS • TempDB – 100GB RAID 0 SAS or SSD

Example 3: Three Server Environment (Collections-focused)		
Server	**Components**	**Hardware Specifications**
Web / Application / Processing Server	• Web Suite • Windows Communication Foundation Service • Asynchronous Processing Services • Processing Engine • Work Manager (Processing/Export) • Case Data/Evidence/Collection Storage	Logical Cores: 16 RAM: 48GB Network Connection: 1GbE NIC Drive Sets: • OS/Apps—150GB RAID 1 SAS • ADTemp—500GB RAID 0 SAS or SSD • Case Data/Evidence/Collection Storage—1TB RAID 5/10 NAS/SAN/DAS
Collection Server	• Work Manager (Collection) • Site Server (Root)	Logical Cores: 8 RAM: 16GB Network Connection: 1GbE NIC Drive Sets: • OS/Apps—150GB RAID 1 SAS • Collection Staging—500GB RAID 5 SAS
Database Server	• SQL Database Engine	Logical Cores: 16 RAM: 32GB Network Connection: 1GbE NIC Drive Sets: • OS/Apps—150GB RAID 1 SAS • DB Files—1TB RAID 5/10 DAS/NAS/SAN • DB Logs—500GB RAID 1 SAS • TempDB—100GB RAID 0 SAS or SSD

Example 4: Four Server Environment		
Server	**Components**	**Hardware Specifications**
Web / Application Server	• Web Suite • Windows Communication Foundation Service • Asynchronous Processing Services • Distributed Processing Manager • Processing Engine • Work Manager (Processing/Export) • Case Data/Evidence/Collection Storage	Logical Cores: 16 RAM: 32GB Network Connection: 1GbE NIC Drive Sets: • OS/Apps – 150GB RAID 1 SAS • Case Data/Evidence/Collection Storage – 1TB RAID 5/10 NAS/SAN/DAS
Processing Server	• Distributed Processing Manager	Logical Cores: 16 RAM: 48GB Network Connection: 1GbE NIC Drive Sets: • OS/Apps – 150GB RAID 1 SAS • ADTemp – 500GB RAID 0 SAS or SSD
Collection Server	• Work Manager (Collection) • Site Server (Root)	Logical Cores: 16 RAM: 32GB Network Connection: 1GbE NIC Drive Sets: • OS/Apps – 150GB RAID 1 SAS • Collection Staging – 500GB RAID 5 SAS
Database Server	• SQL Database	Logical Cores: 32 RAM: 64GB Network Connection: 1GbE NIC Drive Sets: • OS/Apps – 150GB RAID 1 SAS • DB Files – 1TB RAID 5/10 DAS/NAS/SAN • DB Logs – 500GB RAID 1 SAS • TempDB – 100GB RAID 0 SAS or SSD

Example 5: Six Server Environment		
Server	**Components**	**Hardware Specifications**
Web Server	• Web Suite	Logical Cores: 8 RAM: 16GB Network Connection: 1GbE NIC Drive Sets: • OS/Apps – 150GB RAID 1 SAS
Application Server	• Windows Communication Foundation Service • Asynchronous Processing Services • Work Manager (Processing/Export) • Case Data/Evidence/Collection Storage	Logical Cores: 32 RAM: 64GB Network Connection: 1GbE NIC Drive Sets: • OS/Apps – 150GB RAID 1 SAS • Case Data/Evidence/Collection Storage – 10TB RAID 5/10 NAS/SAN/DAS
Processing Server	Distributed Processing Engine	Logical Cores: 16 RAM: 48GB Network Connection: 1GbE NIC Drive Sets: • OS/Apps – 150GB RAID 1 SAS • ADTemp – 500GB RAID 0 SAS or SSD
Collection Server	• Work Manager (Collection) • Site Server (Root)	Logical Cores: 16 RAM: 32GB Network Connection: 1GbE NIC Drive Sets: • OS/Apps – 150GB RAID 1 SAS • Collection Staging – 1TB RAID 5 SAS
Child Site Server	• Site Server (Private)	Logical Cores: 8 RAM: 16GB Network Connection: 1GbE NIC Drive Sets: • OS/Apps – 150GB RAID 1 SAS • Collection Staging – 1TB RAID 5 SAS
Database Server	• SQL Database	Logical Cores: 32 RAM: 64GB Network Connection: 1GbE NIC Drive Sets: • OS/Apps – 150GB RAID 1 SAS • DB Files – 5TB RAID 5/10 DAS/NAS/SAN • DB Logs – 2TB RAID 1 SAS • TempDB – 100GB RAID 0 SAS or SSD

Appendix D: Sample SQL Maintenance Plans

The section below contains a pair of hypothetical SQL maintenance plans.

Note: The maintenance tasks outlined below are for demonstrative purposes only and should not be solely relied upon as they may not be appropriate for your environment.

Simple Recovery Model Maintenance Plan

Job #1: Full Backup (System Databases)

Description: Performs an integrity check and full backup on all system databases.

Schedule: Occurs every day at 12:00:00 a.m.

Step 1. Check the integrity of the system databases.

Step 2. Perform a Full backup of the system databases.

Job #2: Full Backup (User Databases)

Description: Performs an index optimization, integrity check, and full backup on all user databases.

Schedule: Occurs every day at 12:00:00 a.m.

Step 1. Defragment the indexes and update the statistics of the user databases.

Step 2. Check the integrity of the user databases.

Step 3. Perform a Full backup of the user databases.

Job #3: Differential Backup (User Databases)

Description: Performs a differential backup on all user databases.

Schedule: Occurs every day every 6 hours between 6:00:00 a.m. and 11:59:59 p.m.

Step 1. Perform a Differential backup of the user databases.

Job #4: Cleanup

Description: Deletes all backup and job history records that are older than 30 days.

Schedule: Occurs every week on Sunday at 12:00:00 a.m.

Step 1. Execute sp_delete_backuphistory.

Step 2. Execute sp_purge_jobhistory.

Full Recovery Model Maintenance Plan

Job #1: Full Backup (System Databases)

Description: Performs an integrity check and full backup on all system databases.

Schedule: Occurs every day at 1:00:00 a.m.

Step 1. Check the integrity of the system databases.

Step 2. Perform a Full backup of the system databases.

Job #2: Full Backup (User Databases)

Description: Performs an index optimization, integrity check, and full backup on all user databases.

Schedule: Occurs every week on Saturday at 1:00:00 a.m.

Step 1. Defragment the indexes and update the statistics of the user databases.

Step 2. Check the integrity of the user databases.

Step 3. Perform a Full backup of the user databases.

Job #3: Differential Backup (User Databases)

Description: Performs a differential backup on all user databases.

Schedule: Occurs every week on Monday, Tuesday, Wednesday, Thursday, Friday, and Sunday at 1:00:00 a.m.

Step 1. Perform a Differential backup of the user databases.

Job #4: Transaction Log Backup (User Databases)

Description: Performs a transaction log backup on all user databases.

Schedule: Occurs every day every 4 hours between 12:00:00 a.m. and 11:59:59 p.m.

Step 1. Perform a Log backup of the user databases.

Job #5: Cleanup

Description: Deletes all backup and job history records that are older than 30 days.

Schedule: Occurs every week on Sunday at 12:00:00 a.m.

Step 3. Execute sp_delete_backuphistory.

Step 4. Execute sp_purge_jobhistory.

Footnotes

[1] CRT and CER formatted public certificates are only acceptable when using self-signed certificates created with AccessData's Certman utility. P7B formatted certificates must contain the full certification chain.

[2] The .ADP12 file format is an AccessData-generated protected and encrypted P12 certificate format. The Site Server component will automatically generate an ADP12 file if supplied with a PFX, PEM, or P12.

[3] AccessData recommends a minimum 512MB cache on any RAID controller.

[4] Beginning with Windows Server 2008, the default dynamic port range of Windows Server is 49152 through 65535. Please see Microsoft's Knowledge Base (http://support.microsoft.com/kb/929851) for more details.

[5] The edition of Windows Server and SQL Server (i.e., Standard, Enterprise, etc.) will be dependent on the amount of memory being installed in the server.

[6] Italicized requirements are only listed for informational purposes and do not need to be in place prior to implementation, as they can be installed as part of the component installation process.

[7] Microsoft Outlook 32-bit is required for building PSTs during email collections and to export processed items to PST.

[8] For a summary of best-practice maintenance recommendations, see http://technet.microsoft.com/en-us/magazine/2008.08.database.aspx#id0230078.

[9] For detailed information on Microsoft SQL recovery models, see http://technet.microsoft.com/en-us/library/ms189275(v=sql.105).aspx.

[10] For detailed information on Microsoft SQL backup methods and strategies, see http://technet.microsoft.com/en-us/library/ms187048(v=sql.105).aspx.

[11] For detailed information on Microsoft SQL index optimization, see http://technet.microsoft.com/en-us/library/ms190910(v=sql.105).aspx.

[12] Scheduled index optimization should be performed prior to the Full backup; scheduled index optimization performed after a Full backup can dramatically increase the size of the subsequent differential and log backup records.

[13] For detailed information on Microsoft SQL database integrity checks, see http://msdn.microsoft.com/en-us/library/ms176064(v=sql.105).aspx.

[14] For detailed information on Microsoft SQL maintenance cleanup, see http://msdn.microsoft.com/en-us/library/ms345177(v=sql.105).aspx.

[15] Please have the certificate password available for use.

[16] If using Active Directory authentication, knowledge of the name of the Active Directory server to be used to authenticate users and credentials capable of querying Active Directory will be necessary.

[17] If a dedicated "sent by" account is being employed, please ensure it has been created and that its credentials have been documented.

[18] This list only includes "officially supported" versions that have been tested by AccessData's Quality Assurance department; additional versions may be compatible but are excluded here, as they have not been fully tested.

Visit us online:
www.accessdata.com

Global Headquarters
+1 801 377 5410
603 East Timpanogos
Circle Suite 2300
Orem, UT 84097

North American Sales
+1 800 574 5199
Fax: +1 801 765 4370
sales@accessdata.com

International Sales
+44 20 7010 7800
internationalsales@accessdata.com