System Specifications Guide

LAB

July 24, 2014

Contents

AccessData LAB Overview

AccessData® enables computer forensics labs of all sizes, facing an array of challenges, to work more effectively. A single-person lab can radically speed up the processing of cases with the four-worker distributed processing available with FTK®. However, labs handling massive data sets, utilizing a distributed workforce, or looking to collaborate with attorneys, HR personnel or other non-forensic parties can step up to AccessData Lab. AD Lab adds powerful and intuitive web-based review functionality, expanded distributed processing capabilities with a centralized processing farm, and a centralized database infrastructure. This allows collaborative analysis among multiple forensic examiners, real-time task and case management, and secure, web-based collaboration with parties outside the lab. Regardless of the size, scope or mission of your computer forensics lab, AccessData has a solution that will meet your needs. The following contains a brief explanation of the AccessData LAB components and their role within the solution:

- **LAB** – This is the main interface where users login to create cases and add evidence.
- **Processing Engine** – Performs data processing tasks such as expansion or archives (e.g., .PST, .NSF and .ZIP), indexing, de-duplication analysis, file identification, secondary culling and filtering, etc. Component can be installed as a distributed processing engine for environments that process large amounts of data and want to have multiple engines processing the evidence at the same time. Distributed setups of processing engine utilize the Distributed Processing Manager (DPM) and the Distributed Processing Engine (DPE).
- **Database** – The AccessData LAB solution utilizes Oracle, PostgreSQL or Microsoft SQL Server to maintain databases containing file metadata, user data and workflow information. PostgreSQL is provided by AccessData with installation. If customer decides to use Oracle or Microsoft SQL then they must provide the license for the database of choice. Please note that if using the optional web review addition to LAB you must use Microsoft SQL as the database for the application.
- **Case Data/Evidence Storage** – The AccessData LAB solution 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 the evidence and other case related data.

Optional Components

- **Websuite** – This provides the interface through which users access the web review portion of the AccessData LAB solution.
- **Application Services:**
  - **Windows Communication Foundation Services** – WCF services manage the flow of data between the various components of the web review portion or the AccessData LAB solution.
  - **Asynchronous Processing Services** – Async is responsible for the execution of certain user actions such as bulk coding and searching.
  - **Work Manager** – Governs the flow of work to the processing engines.

General Considerations

AccessData strongly encourages the use of physical hardware platforms in any implementation of the AccessData LAB solution. 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.

If using Microsoft SQL server as the database for the LAB solution, AccessData strongly encourages the instance be dedicated and on its own hardware. The support of any implementation attempting to host the MS SQL database component of the LAB solution on the same hardware as other components is subject to the discretion of AccessData. Attempts to host SQL database component in same instance as another enterprise application will not be supported.

AccessData forbids the installation of any of the LAB components on a system that hosts a Microsoft Domain Controller.

Please contact your AccessData technical support representative for further information.

Service Account

If you will be using the optional Web Review portion of the AccessData LAB solution then it is required to setup a single, dedicated service account, otherwise the web portion will not operate properly. If the Web portion components will be installed in a multi-server environment, a domain level service account is required. Workgroup authentication is only supported for single server installations. In either case the service account must be a local administrator with the “Logon as Service”, “Logon as Batch” and “Interactive Logon” system permissions.

General Hardware Requirements

LAB UI / Evidence Processing Engine		
System Component	**Minimum Basic**	**Minimum Recommended**
OS	Server 2008 R2/ Server 2012	Server 2008 R2/ Server 2012
Processor	4 Logical Cores	8-32 Logical Cores
Memory	8GB RAM (2GB/Core)	16-64GB RAM (2GB/Core)
Storage	Separate physical disks for OS and ADTemp files 7200 RPM drives	Single disk for OS/Apps RAID 0 or SSD for AD Temp 10-15K RPM drives if not using SSD
Network Interface	1GbE NIC	10GbE NIC

PostgreSQL Database		
System Component	**Minimum Basic**	**Minimum Recommended**
OS	Server 2008 R2/ Server 2012	Server 2008 R2/ Server 2012
Processor	4 Logical Cores	8-24 Logical Cores
Memory	8GB RAM	16-48GB RAM
Storage	Separate physical disks for OS and ADTemp files 7200 RPM drives	RAID 1 for OS/Apps RAID 10 for Database 10-15K RPM drives
Network Interface	1GbE NIC	10GbE NIC

Microsoft SQL Database (Must use if using optional web review portion of LAB)		
System Component	**Minimum Basic**	**Minimum Recommended**
OS	Server 2008 R2/ Server 2012	Server 2008 R2/ Server 2012
Processor	8 Logical Cores	16 Logical Cores
Memory	16GB RAM	32GB RAM
Storage	OS/Apps – 10-15K RPM drive DB/logs – 10-15K RPM drive	RAID 1 for OS/Apps RAID 10 for Database RAID 1 for DB logs RAID 0 or SSD for TempDB RAID 1 for TempDB Logs
Network Interface	1GbE NIC	10GbE NIC

Distributed Processing Engine		
System Component	**Minimum Basic**	**Minimum Recommended**
OS	Server 2008 R2/ Server 2012	Server 2008 R2/ Server 2012
Processor	4 Logical Cores	8-16 Logical Cores
Memory	8GB RAM	16-32GB RAM
Storage	OS/Apps – 10K RPM drive ADTemp – 15K RPM drive	RAID 1 for OS/Apps RAID 0 or SSD for ADTemp
Network Interface	1GbE NIC	10GbE NIC

Websuite (Optional)		
System Component	**Minimum Basic**	**Minimum Recommended**
OS	Server 2008 R2/ Server 2012	Server 2008 R2/ Server 2012
Processor	2 Logical Cores	8 Logical Cores
Memory	4GB RAM	8GB RAM
Storage	OS/Apps – 10-15K RPM drive	RAID 1 for OS/Apps
Network Interface	1GbE NIC	10GbE NIC

Application Services (Optional)		
System Component	**Minimum Basic**	**Minimum Recommended**
OS	Server 2008 R2/ Server 2012	Server 2008 R2/ Server 2012
Processor	4 Logical Cores	8-16 Logical Cores
Memory	8GB RAM	16-32GB RAM
Storage	OS/Apps – 10-15K RPM drive	RAID 1 for OS/Apps
Network Interface	1GbE NIC	10GbE NIC

Processor and Memory Considerations

The quality of the processors used in the implementation environment will have a direct effect on the overall performance of the AccessData LAB solution. 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.

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, AKA "thrashing", can lead to operational failure. To reduce the likelihood of paging, it is strongly recommended that any system involved in the implementation should possess at least 1GB RAM for each logical processor core and any system hosting a processing engine should possess at least 2GB RAM for each logical processor core.

NOTE: SUSTAINED PERIODS OF HIGH DISK USE AND PERSISTENT DISK QUEUES CAN BE A SYMPTOM OF INSUFFICIENT MEMORY RESOURCES. PLEASE SEE THE STORAGE CONSIDERATIONS SECTION OF THIS DOCUMENT FOR ADDITIONAL INFORMATION.

The AccessData LAB solution has been designed in such a way that components can be expected to leverage all available processor and memory resources available to the host system during certain operations. Please contact your AccessData technical support representative for further information.

Network Considerations

AccessData recommends the use of 1GbE or 10GbE network connections and discourages the use of iSCSI network connectors or link aggregation (i.e., NIC teaming) in any form. The AccessData LAB solution does not use IPv6 for communication and AccessData recommends disabling IPv6 if it is not otherwise required. Communications between the AccessData LAB solution 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. Please contact your AccessData technical support representative for further information.

Storage Considerations

Both the back-end storage hardware being employed and its configuration can greatly affect the overall performance of the AccessData LAB solution. 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 a sign 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. 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, or solid-state drives (SSD).

Storage Types

The storage requirements of the AccessData LAB solution 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. The following table contains information on the various types of storage.

	Description	Storage Characteristics
OS and Apps	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 component that provides storage for ephemeral files generated by the Processing Engine component.	At least 50GB of space is required, but 100GB to 500GB is recommended. The most important characteristic of this space is its speed. This storage requires no fault tolerance. **Recommendation: RAID 0 or SSD.**

Ports used for component communication

The table below lists the Transmission Control Protocol (TCP) ports the specified components use to communicate with each other. 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.

Source Component	Destination Component	Port
LAB UI	Case Data/Evidence	445
	Database (SQL, PG, Oracle)	1433, 5432, 1521
	Processing Engine (Only with DPM, DPE)	34096/34097
Processing Engine	Case Data/Evidence	445
	LAB UI (Only with DPM, DPE)	34096/34097
	Work Manager (Only with DPM, DPE)	34096/34097
Database (MS SQL, PostgreSQL, Oracle)	Async	1433,5432,1521/135
	LAB UI	1433,5432,1521
	WCF	1433,5432,1521
	Work Manager	1433,5432,1521
	Websuite	1433,5432,1521
Websuite	Async	80/808
	Case Data/Evidence	445
	End Users	443
	Database (SQL, PG, Oracle)	1433,5432,1521
	WCF Services	9132
	Work Manager	9132
WCF Services	Case Data/Evidence	445
	Database (SQL, PG, Oracle)	1433,5432,1521
	Websuite	9132
	Work Manager	9132
Async Processing Services	Case Data/Evidence	445
	Database (SQL, PG, Oracle)	1433,5432,1521/135
	Websuite	80/808
Work Manager	Case Data/Evidence	445
	Processing Engine (Only with DPM, DPE)	34096/34097
	Database (SQL, PG, Oracle)	1433,5432,1521
	Websuite	9132
	WCF Services	9132
Case Data/Evidence Storage	Async	445
	Processing Engine	445
	Websuite	445
	WCF Services	445
	Work Manager	445
	LAB UI	445

SQL Server Considerations

The SQL Database component is the heart of the AccessData LAB solution 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 the AccessData LAB solution, 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.

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.

NOTE: IF USING MICROSOFT SQL EXPRESS AS THE UNDERLYING DATABASE, MONITORING THE SIZES OF THE DATABASES IS CRUCIAL. MICROSOFT SQL EXPRESS CAPS THE MAXIMUM SIZE OF DATABASES TO 10GB. IF THIS LIMIT IS EXCEEDED, THE DATABASES MAY NOT BE RECOVERABLE.

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.

Microsoft SQL Server requires ongoing maintenance to maintain its performance. For more information on SQL database maintenance, please see http://technet.microsoft.com/en-us/magazine/2008.08.database.aspx or contact your AccessData technical support representative for further information.

Additional Recommendations

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

AccessData strongly recommends that any anti-virus or anti-malware software on any each server hosting components of the AccessData LAB solution 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.

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.

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.

Appendix A: Sample Environments

Single Server Environment		
Server	**Components**	**Hardware Specifications**
Single Server	LAB UI Processing Engine Database (SQL, PG, Oracle) Websuite (SQL Only) WCF Services (SQL Only) Async (SQL Only) Work Manager (SQL Only) Case Data/Evidence	Server 2008 R2/ Server 2012 16-24 Logical Cores 32-48GB RAM 1GbE NIC Drive Sets: - RAID 1- 150GB- 10K RPM for OS/Apps - RAID 5/10- 500GB to 1TB (depends on case load) for Database - RAID 0– TempDB – 150GB 15K RPM (Only needed for SQL) - No RAID – ADTemp – 300GB 15K or SSD - RAID 5- 500GB to 1TB (depends on case load) for Case Data and Evidence. (SAN, NAS or DAS is also ok just make sure whatever you setup you can expand when needed)

Two Server Environment		
Server	**Components**	**Hardware Specifications**
Application Server	LAB UI Processing Engine Websuite (SQL Only) WCF Services (SQL Only) Async (SQL Only) Work Manager (SQL Only) Case Data/Evidence	Server 2008 R2/ Server 2012 16 Logical Cores 32GB RAM 1GbE NIC Drive Sets: - RAID 1- 150GB- 10K RPM for OS/Apps - No RAID – ADTemp – 300GB 15K or SSD - RAID 5- 500GB to 1TB (depends on case load) for Case Data and Evidence. (SAN, NAS or DAS is also ok just make sure whatever you setup you can expand when needed)
Database	Database (SQL, PG, Oracle)	Server 2008 R2/ Server 2012 16 Logical Cores 32GB RAM 1GbE NIC Drive Sets: - RAID 1- 150GB- 10K RPM for OS/Apps - RAID 5/10- 500GB to 1TB (depends on case load) for Database - RAID 0– TempDB – 150GB 15K RPM (Only needed for SQL)

Three Server Environment		
Server	**Components**	**Hardware Specifications**
Application Server	Websuite (SQL Only) WCF Services (SQL Only) Async (SQL Only) Work Manager (SQL Only) Case Data/Evidence DPE	Server 2008 R2/ Server 2012 8 Logical Cores 16GB RAM 1GbE NIC Drive Sets: - RAID 1- 150GB- 10K RPM for OS/Apps - No RAID – ADTemp – 300GB 15K or SSD - RAID 5- 500GB to 1TB (depends on case load) for Case Data and Evidence. (SAN, NAS or DAS is also ok just make sure whatever you setup you can expand when needed)
Database	Database (SQL, PG, Oracle)	Server 2008 R2/ Server 2012 16 Logical Cores 32GB RAM 1GbE NIC Drive Sets: - RAID 1- 150GB- 10K RPM for OS/Apps - RAID 5/10- 500GB to 1TB (depends on case load) for Database - RAID 0– TempDB – 150GB 15K RPM (Only needed for SQL)
LAB	LAB UI DPM DPE	Server 2008 R2/ Server 2012 8 Logical Cores 16GB RAM 1GbE NIC Drive Sets: - RAID 1- 150GB- 10K RPM for OS/Apps - No RAID – ADTemp – 300GB 15K or SSD

Four Server Environment		
Server	**Components**	**Hardware Specifications**
Application Server	Websuite (SQL Only) WCF Services (SQL Only) Async (SQL Only) Work Manager (SQL Only) Case Data/Evidence DPE	Server 2008 R2/ Server 2012 8 Logical Cores 16GB RAM 1GbE NIC Drive Sets: - RAID 1- 150GB- 10K RPM for OS/Apps - No RAID – ADTemp – 300GB 15K or SSD - RAID 5- 500GB to 1TB (depends on case load) for Case Data and Evidence. (SAN, NAS or DAS is also ok just make sure whatever you setup you can expand when needed)
Database	Database (SQL, PG, Oracle)	Server 2008 R2/ Server 2012 16 Logical Cores 32GB RAM 1GbE NIC Drive Sets: - RAID 1- 150GB- 10K RPM for OS/Apps - RAID 5/10- 500GB to 1TB (depends on case load) for Database - RAID 0– TempDB – 150GB 15K RPM (Only needed for SQL)
LAB	LAB UI DPM DPE	Server 2008 R2/ Server 2012 8 Logical Cores 16GB RAM 1GbE NIC Drive Sets: - RAID 1- 150GB- 10K RPM for OS/Apps - No RAID – ADTemp – 300GB 15K or SSD
Distributed Engine	DPE	Server 2008 R2/ Server 2012 8 Logical Cores 16GB RAM 1GbE NIC Drive Sets: - RAID 1- 150GB- 10K RPM for OS/Apps - No RAID – ADTemp – 300GB 15K or SSD