

2016CTAC

A7 Upgrade related

Infrastructure and Sizing decisions
Presenter: Subir Mukherjee

AMANDA 7 Upgrade Path

The following topics are pre-requisites for migration to A7.

- A.
 - Multi-browser support
 - Starting point of Amanda version
 - Databases supported by A7
 - Application server and Java SDK version
 - Operating system and Client System Requirements
 - New ANSI standards adopted for MS-SQL environments.
 - Usage of sequences for MS-SQL environments
- B.
 - Infrastructure and Sizing Decisions
 - Today we will cover this topic in detail

AMANDA 7 Upgrade Path – 101s

- Multi-browser support: IE 11, Edge, Chrome, Firefox and Safari
- Starting point of Amanda version: Need to upgrade to A6 at the very least.
- Databases supported by A7: Oracle 12c, SQL Server 2014
- Application server and Java SDK version: SDK 1.7
 - Tomcat 8.0
 - Oracle Web Logic 12c
 - IBM WebSphere 8.5.5 (with Java 7 add-on)
 - Jboss EAP 6.4.0 *
- Operating system and Client System Requirements:
 - Windows 2008 or 2012 Server 64-bit
 - Linux (Red Hat Enterprise Linux, SUSE Linux Enterprise Server, or Ubuntu Server)
 - Unix (Managed and supported by customer IT team)

SQL statements using ANSI standards

1. JOINS that have been deprecated in SQL Server broadly fall into 2 categories
 - INNER JOINS
 - OUTER JOINS for example

Join type	Proprietary	ANSI
Left join	Select * from Table1, Table2 where Table1.column1 *= Table2.column1	Select * from Table1 LEFT OUTER JOIN Table2 ON Table1.column1 = Table2.column1
Right join	Select * from Table1, Table2 where Table1.column1 =* Table2.column1	Select * from Table1 RIGHT OUTER JOIN Table2 ON Table1.column1 = Table2.column1

2. To ensure all code terminates with a semi-colon.
3. Compatibility mode 80 (SQL Server 2000) is deprecated.
4. Max+1 logic to get the next RSN will use Sequences going forward.

Scalability

AMANDA is designed to deliver high performance and unmatched versatility to an unlimited number of concurrent users.


- A7 platform could be scaled horizontally by adding more application or database server nodes to respective application or database server cluster(s).
- The architecture and design of the solution provides support for increased transactional load, increased data volume and increased concurrent users.

This approach offers a system capacity upgrade roadmap which is not limited by the capabilities of a single server. Its accomplished through the addition of more servers at each of the specific layers.

Scalability – cont'd

- The number of Application servers or Public Portal servers can be increased to support growth.
- However, scaling out the database servers must be done in accordance with the database engine capabilities.

Scalability – cont'd



Load Category

How to choose your Load Category

Firstly, we briefly define a few key terms

a)Users

b)Folder Units per Year, or

c)Peak Folder Units in 10 minutes

- "Folder Unit" corresponds roughly to the number of folders in your database.
- If you are using attachments, Executive Monitor/Dashboard, OpenTax, Mobile (with inspectors synchronizing each morning), or GIS Maps, then select the next-largest category.

Defining a Folder Unit of work

The "Folder Unit" unit of work is defined as:

- Logon,
- Adding a new "average" folder, entering data in all rows and most columns, including property, people, info,
- Attempts against 17 processes, and bills and payments
- Task List, with attempt against process
- Inspection Request
- Logoff

Criteria used to define a Folder Unit

- The load calculations contained in this document are based on data from 30 customers that was statistically analyzed to determine average and peak usage, and performance tests.
- The five most common folder types at each customer were identified, and then the data for each was statistically analyzed to find the average folder. It was determined that the average folder had:
 - 9 sub, 24 work
 - FolderProperty
 - FolderPeople
 - 30 FolderInfo (4 pick/choose, 8 alpha, 9 numeric, 8 boolean, 1 date)
 - 19 FolderProcess (16 with business-rule-stored procedures)
 - Fee procedures (1 new, 3 other), with corresponding fees, bills, and payments

Load Category – cont'd

How was the Test Data put together?

- AMANDA 7's application server (the browser html user interface) was load tested using **web application performance tool** (WAPT) varying users (human speed and no delay), CPUs, memory, and JDKs.
- The Search Engine was tested with varying memory sizes and disk drive types.
- The database server was driven using Folder Unit and other AMANDA 7 UI-required operations to 100% CPU and / or 100% Disk Utilization, including Attachment Blobs, Locking, and Disk Drive types.
- The observations were then applied against customer peak usage (user, annual Folder, 10 minute Folder/Process/Info) to create the tables in this document.

Load Category – cont'd

- In table below, select the table row based on whichever one of the following gives you the largest Load Category

Load Category	Users	Folder Units per Year	Peak Folder Units in 10 Minutes
A	0 - 24	0 - 9,999	10
B	25 - 49	10,000 - 19,999	25
C	50 - 74	20,000 - 29,999	35
D	75 - 99	30,000 - 39,999	50
E	100 - 199	40,000 - 79,999	100
F	200 - 299	80,000 - 119,999	150
G	300 - 399	120,000 - 159,999	175
H	400 - 499	160,000 - 199,999	250
I	500 - 999	200,000 - 399,999	500
J	1000 - 1999	400,000 - 800,000	1000

Load Category – suggested hardware choice

AMANDA 7 Application Server (Optional Cluster) for AMANDA, Public Portal, EAI, and Mobile Application Server

Table 1-2: AMANDA 7 Application Server

Item	Description									
Components	Back Office, Administration, and Search Engine. (All on one server)									
Operating System	<ul style="list-style-type: none"> Windows 2008 or 2012 Linux (Red Hat Enterprise Linux, SUSE Linux Enterprise Server, or Ubuntu Server) Unix (Managed and supported by customer IT team) 									
Application Server Software Suite	<ul style="list-style-type: none"> Tomcat 8.0 Oracle Web Logic 12c IBM WebSphere 8.5.5 (with Java 7 Add On) JBoss EAP 6.4.0 <p>Note: If clustering or high-availability deployment is required, we recommend WebLogic, WebSphere, or JBoss.</p>									
Network Interface Card	1 Gbps									
Disk Storage	100GB (OS + Paging + Tomcat + AMANDA7 WARs/etc. + Logs + Search Engine data)									
Load Category	A	B	C	D	E	F	G	H	I	J
Server Processor	1000	2000	3000	4000	8000	12000	16000	20000	40000	80000
Server RAM	8GB	8GB	10GB	12GB	14GB	14GB	14GB	16GB	20GB	24GB
Http Requests/min.	530	1100	1600	2100	4300	6400	8500	11000	21000	43000

Load Category – hardware choice cont'd

Table 1-2: AMANDA 7 Application Server (Continued)

Notes:

- Http Requests can be used for Load Balancer performance requirements.
- RAM includes OS + Tomcat + Search Engine. e.g., 2GB + 6GB heap + 4GB.
- Processor units are "PassMark 'CPU Mark' " scores. You can translate these into corresponding CPU(s) at <https://www.cpubenchmark.net/>

For CPU Marks less than 4000, a portion of a virtual server may be utilized. For example you might have a computer with "Intel Xeon E5-2650L v3 @ 1.80GHz" (Passmark Score 12311, 12 cores (2 logical cores/threads/hyperthreads per physical core)), and you could virtually allocate 4 cores to the AMANDA virtual server giving that virtual server with the equivalent passmark score of 4103, which is large enough for Capacity Category D.

Since AMANDA7 scales almost linearly as more CPU power is added, from a performance perspective, scaling vertically (single computer more CPU) is recommended over scaling horizontally (multiple computers each with less CPU).

A single computer with two "Intel Xeon E5-2695 v3 @ 2.30GHz" (each with Passmark Score 20766, or 41532 total for a single computer with two CPUs), would be large enough for Capacity Category I. Today, no single computer with two CPUs could handle Capacity Category J, so a four CPU computer or multiple computers would be required. CPUs should have a Single Thread Rating of 1800 or higher, 2000 or higher preferred.

- Hyperthreading gives an increase in performance of about 6%.
- Java 8 runs AMANDA7 6% faster than Java 7.
- Tomcat java heap should be 4GB+.
- Search Engine uses ~500MB RAM, but may grab up to 4GB for Full Import.
- Search Engine runs 20% faster on SSD drives, vs spinning hard disks.
- Servers with CPU Marks less than 4000 may be tedious to support (software installation, zipping logs, etc.)


Load Category – hardware choice cont'd

AMANDA Database Server (Optional Cluster)

Table 1-3: AMANDA Database Server

Item	Description
Components	Database
Operating System	<ul style="list-style-type: none"> Windows 2008 or 2012 Linux (Red Hat Enterprise Linux, SUSE Linux Enterprise Server, or Ubuntu Server) Unix (Managed and supported by customer IT team)
Database Software	<ul style="list-style-type: none"> Oracle 11g or 12c; Enterprise Edition; Enterprise Edition required for Oracle's Real Application Cluster option and Enterprise Edition required for RLS (Row Level Security) option. MS-SQL Server 2012 or 2014 Standard or Enterprise Edition; Enterprise Edition required for Clustering.
Network Interface Card	1 Gbps

Item	Description									
Load Category	A	B	C	D	E	F	G	H	I	J
Server Processor	1000	1000	1000	1000	1500	2250	2600	3700	7500	15000
Server RAM	6GB	6GB	6GB	6GB	8GB	12GB	16GB	20GB	32GB	64GB
RDBMS Software GB	4	4	4	4	4	4	4	4	4	4
Database Files GB	35	35	35	35	35	35	48	60	120	230
AMANDA Data + Index	4	8	10	15	30	45	60	75	150	300
AMANDA Attachment	50	100	150	200	400	600	800	1000	2000	4000
Total Disk GB	93	147	199	254	469	684	912	1139	2274	4534
Disk Bandwidth B/sec	40K	100K	130K	200K	385K	575K	700K	1M	2M	4M



Load Category – hardware choice cont'd

Notes:

- To find your "Load Category", see Appendix B, "Load Category" on page 14.
- System performance is dependent upon disk subsystem, processors, memory.
- Data + Index is calculated at 38KB per Folder Unit, over 10 years.
- Attachment is calculated at 2x256KB per Folder Unit, over 10 years.
- Database Files, for Oracle, would include: REDO, CONTROL, SYSAUX, SYSTEM, TEMP, UNDO, USERS, OSSYS, OSSTATE, etc.
- Server Processor estimates are based upon Folder Unit processing, not reporting.
- Peak 8k random IOPs minimum is 2,000 (~5 mirrored HDD).
Peak 8k random IOPs preferred is 50,000+ (2+ SSD).
Bandwidth divided by 8k (the IOs per second) is considerably lower than these values, but AMANDA performance would be abysmal. Attaching AMANDA DB to a SAN to take advantage of its fast access times (IOPS 50,000+), and using a tiny percentage of its IOs (575KB/s in 8k random 50%read/50%write), is common.
- Putting datafiles AMANDA_DATA, AMANDA_INDEX, AMANDA_ATTACHMENT, UNDOTBS1, REDO on separate physical disks is an easy performance optimization. Storage optimizations beyond this are city specific and should be determined by a DBA.
- OpenTAX's Tax Bill and Bill Print are remarkably database bound and thus beyond the scope of this document which is "Folder Unit" centric. Please consult with CSDC re existing OpenTAX customers for hardware examples.

Note: As the AMANDA Database Server is the core of the AMANDA System, CSDC recommends making DB server(s) as powerful as possible (disk subsystem, processors, memory) to ensure overall system performance.

Browser Response Times

Description	Average	Minimum	Maximum
Logon: <ul style="list-style-type: none"> Choose Database Enter userid/password Logoff other users See first menu 	1.30	0.78	5.60
Permit Auto Add 1 - Permit: <ul style="list-style-type: none"> Add Permit Enter FolderType, Sub, Work 	0.58	0.33	2.23
Permit Auto Add 2 - Enter Permit Details: <ul style="list-style-type: none"> Enter FolderDescription, FolderConditions, etc {insert Folder} 	0.31	0.14	0.83
Permit Auto Add 3 - FolderPeople	1.87	0.05	5.43
Permit Auto Add 4 - Set Info (4 choose, 8 alpha, 9 numeric, 8 boolean)	2.60	1.06	7.83
Permit Auto Add 5 - enter Attempt on first Process	0.71	0.25	1.61
Permit Auto Add 6 - enter Attempts on next 16 Processes	11.10	4.40	17.60
Permit Auto Add 7 - Fee, Create Bill, Add Payment	1.06	0.27	2.57
Permit Auto Add 8 - Task List (for this test this always returns 2 rows): <ul style="list-style-type: none"> See list Enter Attempt on first Process 	3.40	0.47	9.20
Permit Auto Add 9 - Inspection Request: <ul style="list-style-type: none"> Query Schedule 	0.68	0.44	1.34
Permit Auto Add 10 - close tabs	0.31	0.17	0.73
Logoff	0.19	0.03	0.58

2016CTAC

Thank you!