

Best Practices for System Administration

LabKey User Conference and Workshop 2016

Brian Connolly

Senior Systems Engineer, LabKey Software



Agenda:



- Your hardware and environments
- Prerequisite software
- Overview of a LabKey Server installation
- Overview of the Admin Console
- Troubleshooting a LabKey Server
- Staging Servers
- Monitoring and backups
- Q & A


Who am I



Senior Systems Engineer at LabKey

- What do I do at LabKey?
 - Run LabKey Servers at AWS and at Customer's sites
 - Help our Customers run LabKey Server
- Worked at LabKey for last 9 years
 - Previously worked at Microsoft, RFA systems, UW and a few startups

Outside of LabKey I spend lots of time in the mountains





Consider Your Starting Points, Pre-existing Systems

- Hardware (Physical or Virtual)
- Operating System
- Database(s)
- Using a Staging or Test Environment

Pre-requisite Software



- Oracle Java
- Apache Tomcat
- Database Software (PostgreSQL or MS SQL)
- Third Party Tools

Documented on our Support Technologies page at
<https://www.labkey.org/wiki/home/Documentation/page.view?name=supported>

Pre-requisite Software: Java



Recommended Version:

- 1.8.x - Latest Version
- Distribution to use: **ServerJRE**
- Download URL:

<http://www.oracle.com/technetwork/java/javase/downloads/index.html>

Updates:

- Oracle releases Critical Patch Updates (CPU) once a quarter.
- Recommendation: *Always install updates during next maintenance window*

Pre-requisite Software: Tomcat



Recommended Version:

- Latest version of 7.0 or 8.0

Download URL:

- <https://tomcat.apache.org/download-80.cgi>
- Linux/Mac OSX: use “tar.gz” dist
- Windows: use “32-bit/64-bit Windows Service Installer”

Updates:

- Follow your standard patching policy

Pre-requisite Software: Tomcat



Configuration:

- Follow the guidelines in LabKey's documentation (#7 to #10)
 - <https://www.labkey.org/wiki/home/Documentation/page.view?name=configTomcat>
- server.xml: Use LabKey's sample configuration
 - <https://github.com/LabKey/samples/tree/master/ops/config-examples>
 - These are the **configuration settings we use**.
- Tomcat memory
 - Recommended: *Startup = 512MB | Max Size = 4096MB or greater*
 - <https://www.labkey.org/wiki/home/Documentation/page.view?name=configWebappMemory>



Configuration (continued)

- Windows Service Logon
 - Do not use Local System

Pre-requisite Software: PostgreSQL



Recommended Version: 9.5 or later

- Version 9.2, 9.3, 9.4 are still supported and regularly tested.

Download URL:

- <http://www.postgresql.org/download/>

Updates:

- Follow your standard patching policy
- LabKey will notify you if update contains **critical hot-fix or security vulnerability**



Configuration

- User and Privileges:
 - LabKey Server assumes that the user is a “superuser”
- Configuration Guidance (for “large” server)
 - Effective Cache Size: 75% of memory
 - Shared Buffers:
 - Linux: 25% of memory
 - Windows: 64MB to 512MB
 - Work mem: 20MB | Maintenance Work mem: 1024MB | Autovacuum Work mem: 512MB
 - Checkpoint Segments: 10
 - Checkpoint Timeout: 15
 - Random Page Cost: 1.4 | Join collapse limit: 10

Pre-requisite Software: MS SQL Server



- Recommended Version: SQL Server 2014
 - Versions 2008 and 2012 is supported and regularly tested
- Updates:
 - Follow your standard patching policy



Configuration

- User and Privileges:
 - LabKey Server assumes that the user is a member of the **sysadmin** role
 - Can be installed without user having **sysadmin** role.
 - DBA will create new database for LabKey Server
 - User requires **db_owner** role for database
 - After installation, DBA must manually install GROUP_CONCAT
 - <https://www.labkey.org/wiki/home/Documentation/page.view?name=groupconcatinstall>

Pre-requisite Software: Third Party Software



See

[https://www.labkey.org/wiki/home/Documentation/
page.view?name=thirdPartyCode](https://www.labkey.org/wiki/home/Documentation/page.view?name=thirdPartyCode)

LabKey Server Administration: Agenda



- Your hardware and environments
- Pre-requisite software
- **Overview of a LabKey Server installation**
- Overview of the Admin Console
- Troubleshooting a LabKey Server
- Staging Servers
- Monitoring and backups
- Q & A

Overview of LabKey Server Installation



Installation and Upgrade

- LabKey does not have *official* installation or upgrade scripts
- Sample scripts are **available** in our Samples repo
 - <https://github.com/LabKey/samples/tree/master/ops>

Overview of LabKey Server Installation



Important directories

- LabKey Server Installation Directory
- Site-Wide Fileroot
- Tomcat TMP directory

Overview of LabKey Server Installation



LabKey Server Install Directory: (`LABKEY_HOME`)

Contains LabKey Server Software

- `./modules`: LabKey modules installed here
- `./labkeywebapp`: all files used by web application
- `./pipeline-lib`: libraries
- `./bin`: compiled 3rd party tools
- `./files`: Site-Wide file root (default location)
- `./threadDumpRequest` & `./heapDumpRequest`:
- Optional:
 - `./externalModules`: put custom modules here
 - `./extraWebapp`: use this for splash pages or robots.txt, etc

We use `c:\labkey\labkey (/labkey/labkey)`

Overview of LabKey Server Installation



Site-Wide Fileroot:

- Contains
 - Default location **for all files** associated with every folder
- By default, located in Installation directory.
 - Location can be changed via the Files Admin Console
- Directory structure **matches** the folder layout of your LabKey Server
 - Example: if your “home” project has two sub-folders named bob and alice Sitewide Fileroot directory structure will be
 - `./`
 - `./home`
 - `./home/bob`
 - `./home/alice`

Overview of LabKey Server Installation



Tomcat TEMP directory:

- Contains
 - Temporary files such as report contents, thumbnails, etc
 - Default location of Full Text Search index
- Must be **secured**. Only admins and the user running the Tomcat server need access
- Location can be changed via TOMCAT service configuration tool
 - Default location is `TOMCAT_INSTALL_DIR/temp`
 - Recommend:
 - Do **not** use the default location (we use `c:\labkey\tomcat-tmp`)
 - Place on local disk (not on network storage)

Overview of LabKey Server Installation



LabKey Server configuration file: (`labkey.xml` file)

- Installed in `TOMCAT_INSTALL_DIR/conf/Catalina/localhost`
- What does it do? Tells the Tomcat server
 - How to **connect** to the LabKey database
 - **Specifies** location of the labkeywebapp directory
 - How to **connect** to the SMTP server
 - **Sets** various configuration parameters

Overview of LabKey Server Installation



LabKey Server configuration file: (continued)

- Security
 - This file **contains** passwords and sensitive information
 - **Secure** this file!
- Documentation at <https://www.labkey.org/wiki/home/Documentation/page.view?name=cpasxml>

Overview of LabKey Server Installation



LabKey Server configuration file: Important settings

- Context Path:
 - Default is `labkey`:
 - This means your URL will be something like <http://host.example.org/labkey>
 - Controlled by **name** of the file
 - If you want your URL to be <http://host.example.org>
 - **change** name of file to `ROOT.xml`

Overview of LabKey Server Installation



LabKey Server configuration file: Important settings

```
<Context docBase="@@appDocBase@" debug="0" reloadable="true"  
crossContext="true">
```

- `docBase`: is the location of `labkeywebapp` directory.

Overview of LabKey Server Installation - Postgres



LabKey Server configuration file: Important settings

Database Configuration:

```
<Resource name="jdbc/pgDataSource" auth="Container"  
  type="javax.sql.DataSource"  
  username="USERNAME"  
  password="PASSWORD"  
  driverClassName="org.postgresql.Driver"  
  url="jdbc:postgresql://localhost:5432/test"  
  maxActive="20"  
  maxIdle="10"  
  accessToUnderlyingConnectionAllowed="true"/>
```



LabKey Server configuration file: Important settings

Database Configuration:

```
<Resource name="jdbc/mssqlDataSource" auth="Container"
  type="javax.sql.DataSource"
  username="USERNAME"
  password="PASSWORD"
  driverClassName="net.sourceforge.jtds.jdbc.Driver"
  url="jdbc:jtds:sqlserver://localhost:1433/DATABASE_NAME"
  maxActive="20"
  maxIdle="10"
  accessToUnderlyingConnectionAllowed="true"
  validationQuery="SELECT 1"/>
```

Overview of LabKey Server Installation



LabKey Server configuration file: Important settings

SMTP Configuration:

```
<Resource name="mail/Session" auth="Container"  
  type="javax.mail.Session"  
  mail.smtp.host="@smtpHost@"  
  mail.smtp.user="@smtpUser@"  
  mail.smtp.port="@smtpPort@" />
```

- Configuration supports TLS and other options
 - <https://www.labkey.org/wiki/home/Documentation/page.view?name=cpasxml>



LabKey Server configuration file: Important settings

```
<!-- Encryption key for encrypted property store -->  
<Parameter name="MasterEncryptionKey" value="@@masterEncryptionKey@" />
```

- Master Encryption Key:
 - This key is used to **encrypt credentials** used to access remote services
 - Store this key in **safe place**
 - If using a Staging server: key must be **same** on staging and production

Overview of LabKey Server Installation



LabKey Server **Installation** has two steps:

1. Install the binary distribution **files**
2. Install the LabKey **database** schema

Overview of LabKey Server Installation



Step 1: Install the binary distribution files

- What happens during this step:
 - a. Files from distribution are **copied** to LabKey Installation directory
 - b. LabKey Web Service is **started**

Overview of LabKey Server Installation



Step 2: Install the LabKey database schema

- What happens during this step:
 - a. LabKey Web Server starts up
 - b. Attempts to connect to database instance
 - If database does not exist, then **creates** it
 - If database exists, but is empty goes to next step
 - c. **Installs** labkey and core schemas
 - d. Asks the installer to create the first user account
 - This account will be member of **Site Admin** group
 - e. **Installs** all other schemas in database 31

Overview of LabKey Server Installation



LabKey Server **Upgrade** has two steps:

1. Install the binary distribution **files**

2. Upgrade the LabKey **database** schema

(upgrades are a little more involved)

Overview of LabKey Server Installation



Step 1: Install the binary distribution files

- What happens during this step:
 - a. LabKey Web Server service is **stopped**
 - b. Files from the distribution directory are **copied** to LabKey installation directory
 - c. LabKey Web Server service is **started**

How is this done:

- **Using** `upgrade-windows-manual.bat` **script**
 - *use `upgrade-windows-manual.sh` for linux*³³

Overview of LabKey Server Installation



Step 2: Upgrade the LabKey database schema

- What happens during this step:
 - a. LabKey Web Server Server starts up
 - b. Connects to database instance
 - c. Compares version of newly installed **files** with information in the **database** **

** Version information is **stored both** in the database and in the files/software on disk)

Overview of LabKey Server Installation



Step 2: Schema upgrade not required

- Two scenarios where this can happen:
 - a. If the software is **not** newer than the database
 - b. If the software is **newer** than the database, but **no database changes** are required.

Overview of LabKey Server Installation



Step 2: Schema upgrade is required

- This occurs if the software is **newer** than the database and **database changes are required**.
- So what happens :
 - a. Web server will only allow Site Admins to login while the upgrade is running
 - b. After upgrade is finished all users are allowed to login

Overview of LabKey Server Installation



Typical workflow for installer/upgrader:

1. Run upgrade script
 - a. **Review** messages printed to screen for errors
 - b. *Note: LabKey Web Server is started at end of script*
2. Open `labkey.log` and **verify** the server has started
3. Open browser and login
 - a. If required, perform database upgrade

Overview of LabKey Server Installation



Typical workflow for installer/upgrader: (cont.)

4. When upgrade is complete:

- a. Review `labkey.log` and verify there were no errors
- b. Open Admin Console and **verify** version of LabKey Server running
- c. Start **acceptance testing**.

LabKey Server Administration: Agenda



- Your hardware and environments
- Pre-requisite software
- Overview of a LabKey Server installation
- **Overview of the Admin Console**
- Troubleshooting a LabKey Server
- Staging servers
- Monitoring and backups
- Q & A

Overview of Admin Console



The LabKey Server Admin Console contains lots of **good information**.

- Management pages for server
- Version: LabKey, JAVA, PostgreSQL, TOMCAT
- Diagnostic Info: Links to log files, Memory usage, etc
- Audit logs

Let me give you run-down....

LabKey Server Administration: Agenda



- Your hardware and environments
- Pre-requisite software
- Overview of a LabKey Server installation
- Overview of the Admin Console
- **Troubleshooting a LabKey Server**
- Staging Servers
- Monitoring and backups
- Q & A

Troubleshooting a LabKey Server



LabKey Server writes **lots** of log files. Here is an **overview**.

Troubleshooting a LabKey Server



Log files are located in CATALINA_HOME/logs

- `labkey.log`: 99% of the **important logs** are here
 - Contains debug, informational, warning and error
 - File is rotated when reaches 10MB.
 - Four previous versions retained
- `labkey-errors.log`: **Same as** `labkey.log` **but only ERROR** messages
 - File is rotated with each server restart
 - Four previous versions retained
- `catalina.YYYY-MM-DD.log`: (or `catalina.out`)
 - Contains **TOMCAT specific** logs
 - New file is created daily (as long as activity) and retained forever

Troubleshooting a LabKey Server



For the expert:

- `localhost_access_log.*`:
 - Contains access logs (ie info about each request to the server)
 - Contains URL, referrer, user account making requesting, etc
 - New file is created daily (as long as activity) and retained forever
- `commons-daemon.*.log`:
 - Contains logs from the Windows Service. If service does not start, look here.
 - New file is created daily (as long as activity) and retained forever



When an ETL runs a log file is created.

- ETL is run as Pipeline Job
- ETL log file is referred to as **pipeline job log** file
- Contains informational and error messages
- If the ETL fails look in this log for errors or other debug info.

Troubleshooting a LabKey Server



Heap and thread dumps:

- HeapDump:
 - What: “Dumps” **contents of memory** to a file
 - How: “touch” file
`LABKEY_HOME\heapDumpRequest`
 - Where: File is written to `LABKEY_HOME`
- ThreadDump:
 - What: “Dumps” **running threads** to `labkey.log` file
 - How: “touch” file
`LABKEY_HOME\threadDumpRequest`
 - Where: File is written to `LABKEY_HOME`⁴⁶

LabKey Server Administration: Agenda



- Your hardware and environments
- Pre-requisite software
- Overview of a LabKey Server installation
- Overview of the Admin Console
- Troubleshooting a LabKey Server
- **Staging Servers**
- Monitoring and backups
- Q & A



Guidance for setting up a Staging Server

- **Change** the Server GUID
- Write a SQL script to **automatically** change Site Settings and Look & Feel

Full documentation at

[https://www.labkey.org/wiki/home/Documentation/
page.view?name=stagingServerTips](https://www.labkey.org/wiki/home/Documentation/page.view?name=stagingServerTips)

Staging Environment: My Design



My design for building a Staging Env.

- Create exact duplicate of Production Server
 - Same OS and version of pre-req software
 - Tomcat, PostgreSQL/MSSQL, JAVA
 - Other software your server may use (ie R, python)
 - Same configuration
 - FileRoots, Tomcat TMP directories in same file locations
 - HTTPS (if configured on Prod)
- Change some settings
 - Set Server GUID
 - Set Master Encryption Key to be same as Prod
 - Change name of LabKey Database
 - ie. if Prod is `labkey`, then use `labkey-staging`



My design for building a Staging Env. (cont)

- Data Periodically “refreshed” from Production
 - When Production is to be upgraded
 - 1. Data in Staging Env will be wiped
 - 2. Data in Production Env will be copy to Staging

Staging Env: Using the design (High-Level)



1. On Production Env.

- Backup the “Data” in Prod Env
- The “Data” is stored in the LabKey Database and in SiteWide FileRoot
 - (and any other FileRoots you may be using)

2. On Staging Env.

- Shutdown LabKey Server
- Refresh Database and FileRoot “Data”
- Change Look and Feel (and other settings)
- Start LabKey Server
- Test

Staging Env: Using the design (Details)



On Production Env.

- Backup of the LabKey Database options
 - **Full Backup**: Write to file and copy to Staging Env.
 - **Other** methods are available.
- Backup of FileRoot(s) options
 - **Full**: “zip” up files to archive and copy to Staging Env.
 - **Incremental**: Use a tool like Rsync to transfer files directly to Staging Env.

Staging Env: Using the design (Details)



On Staging Env.

- Shutdown LabKey Server
- Restore of the LabKey Database options
 - Full Backup:
 - i. Drop existing LabKey Database
 - ii. Restore Production Backup
 - iii. *[Best Practice: Use different name for LabKey Database on Staging]*
 - Other methods are available.
- Change Look and Feel/Site Settings
 - See <https://www.labkey.org/wiki/home/Documentation/page.view?name=stagingServerTips>

Staging Env: Using the design (Details)



On Staging Env.

- Restore of FileRoot(s) options
 - Full:
 - i. “unzip” files to archive into proper directories
 - ii. Ensure OS permission of restored files is correct
 - Incremental:
 - i. Ensure OS permission of copied files are correct

Now you are ready to test

- New LabKey Server release
- OS patches
- Tomcat, PostgreSQL, JAVA etc upgrades



Tell me about your plans for

- Monitoring
- Backups



Monitoring: **What to monitor**

- Tomcat: Up/Down
- Disk space on `FileRoot`, `PipelineRoot`,
`Tomcat TMP dir`
- Database: Up/Down
- CPU/Memory usage
- Other...



Backup: Important to backup

- Files:
 - FileRoot (Site-wide and all custom)
 - PipelineRoot(if used)
 - Tomcat TMP directory
 - LabKey, Tomcat, database log files
 - System Files (Operating System, software, etc)
- Database

Questions?



Questions?