TRAINING GUIDE

ArcGIS Server and Geodatabase Administration for 10.1

Revision: 2013

clarity through connections www.lucity.com 10561 Barkley, Suite 500 Overland Park, KS 66212 913.341.3105

Table of Contents

Setup 3
ArcGIS for Server Design 3
ArcGIS for Server Setup 5
Connect to ArcGIS for Server from ArcCatalog 10
Installing and Configuring ArcGIS for Server Web Adapter 12
SQL Server Specifications 16
SQL Instance Setup 17
Install SQL Server for ArcSDE for Workgroup 17
SQL Server Enterprise Setup 28
Enabling ArcSDE for Workgroup Databases 29
Install SQL Server Native Client 36
Connect and Create Geodatabase in Workgroup 36
Registering the Geodatabase in ArcGIS for Server 41
Create Enterprise Geodatabase (formerly ArcSDE) 47
Publishing and Configuring Map Services 48
Unlock Schema 52
Lucity Data Update SOE Configuration 56
Lucity Data Update SOE Logic 62
How to Create a Feature Service 63
Preparing Map for a Feature Service 63
Publishing a Feature Service 70
Enabling the Geometry Service 74
Creating a Geocoding Service 76
Securing a Map Service 78
Maintenance 85
Geodatabase Administration Tools 85
Create Role 87
Create Role within ArcCatalog 87
Change/Add Privileges 93
Analyze and Index 95
Compression 96
What is Compression? 96
Spatial and Attribute Indexes 99
Spatial Indexes 99
Attribute Indexes 100

Tune SQL Server space for Geodatabases 102
Backups 104
Automated Maintenance Scripts 107
SQL Server and ArcSDE Scripts 107
Automate Replica Synchronization using Python 108
Options 117
Versioning 117
What are A and D tables? 117
When Does Versioning Make Sense? 117
How to Create a Version 118
Replica 120

ArcGIS for Server v10.1 and Geodatabase Administration

Here at Lucity, we understand that our software is requiring more skills in ESRI technology and these Esri systems can be difficult to manage without training. In this workshop, we will explore the nature of ArcGIS for Server v10.1, touch on key principles and techniques, as well as explore best practices for integration into Lucity. We will cover some key administration best practices for: design, setup, maintenance, and optional editing workflow. Even though we won't be covering the full gamut of ESRI technologies, we will be focused on the administration of ArcGIS for Server, how this relates to the IT and GIS administrators and what are some basic principles for enabling ArcGIS for Server to be compatible with Lucity.

Setup

ArcGIS for Server Design

- 64 Bit Operating Systems from Windows XP (**depreciated in 2014**), Windows Vista (**ArcGIS Runtime will not support after 10.2**) and 7, Windows 2003 (**not supported after ArcGIS 10.2)** or higher. Recommend Windows Server 2008 R2 or 2012
- Server should have 2-4 cores minimum and 4GB of RAM per core. For a 4 core system, you should have 16GB of RAM.
- If you have virtual servers, make sure you dedicate resources and have the latest VMware or Hyper-V software available. Based on ESRI's testing, Virtual software is the key to many performance issues. Make sure in VMware environments that you subtract overhead of two cores for the Virtual Environment so ArcGIS for Server map services doesn't pool across the Virtual overhead cores. If you don't, you will see a significant performance decrease in ArcGIS for Server.
- Try to limit map services to utilize 4 pools per core maximum. Anything more than that, you could allow ArcGIS for Server crashes when many people are simultaneously hitting your server.
- If you have a dedicated system for ArcGIS for Server, try creating a RAID 1 SSD configuration for your hard drives. Also, use the latest version of SSD standards as this will increase performance and reliability. For all new SSD cards, it will take at least 40 years before they reach their read/write maximum. This benchmark test was performed on a 24/7 continuous read/write routine. If SAS drives, use RAID 10 with a good RAID controller.

- Space depends on caching. If you plan to cache imagery, please account for large image files. Remember, each scale in the cache will have separate tiles for each scale. If you have 9 scales within a cache, you will have nine separate tiles of cache. Smaller scale images will be smaller in size and the larger scale images will be larger in size.
- Suggest HTTP and HTTPS Secure Design when setting up ArcGIS for Server security Configuration
- If intranet only, use HTTP only unless otherwise specified by your IT administration

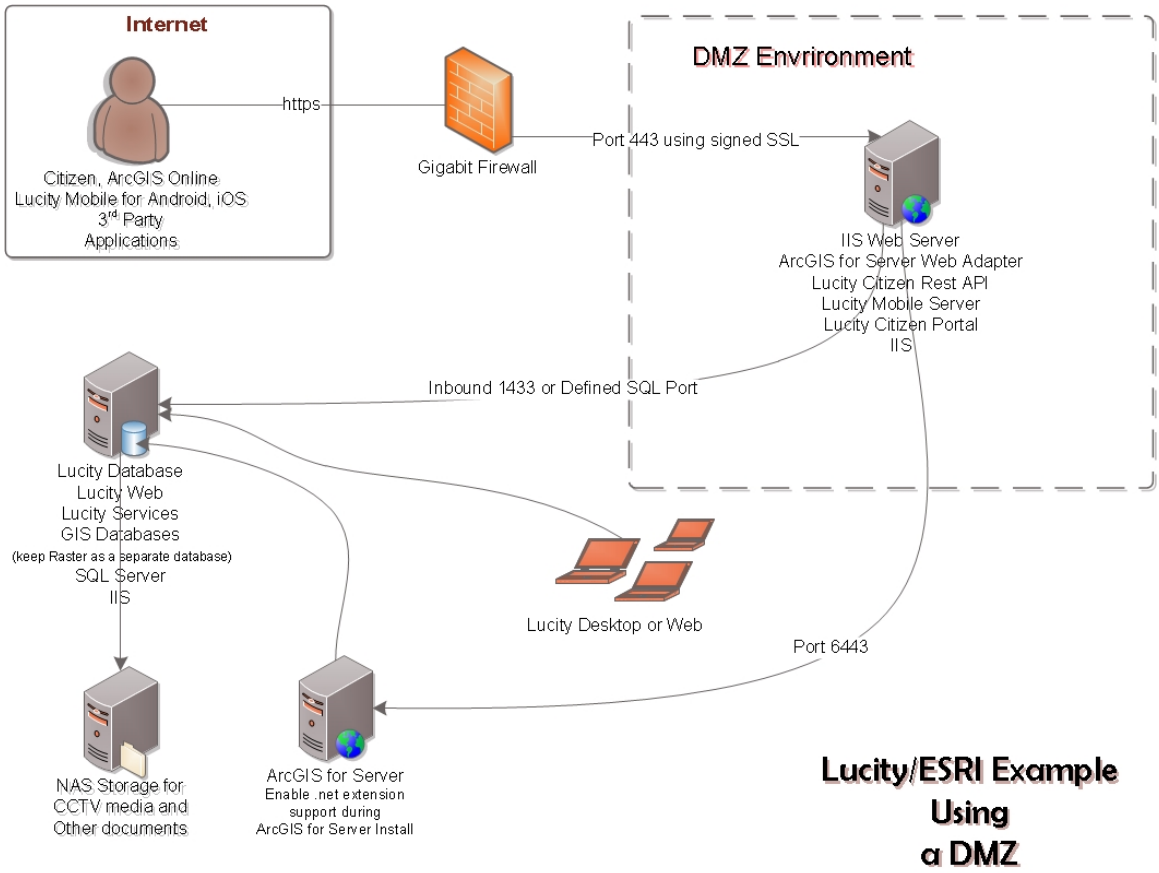

ArcGIS for Server Setup

1. First, run the upgrade utility on existing system. If this is a new system, bypass the utility. Remember, ArcGIS for Server needs to be 64bit but the web adapter (optional) can be on either 32bit or 64bit.

Note: When installing ArcGIS for Server on a separate server than your SQL Server 2012 instance with ArcSDE v10.1, you need the SQL Server Native Client 64 bit. Also, these versions need to match. For example, SQL Server 2012 original install media is a different version than the latest updated SQL Server 2012 from Microsoft updates. If the native client doesn't match the patched version of SQL Server 2012, it won't work. Esri recommends either the SQL Server 2012 matching version or the sql server 2008 R2 native client. Latest SQL Server 2012 Native 64 bit Client download link: http://go.microsoft.com/fwlink/?LinkID=239648&clcid=0x409

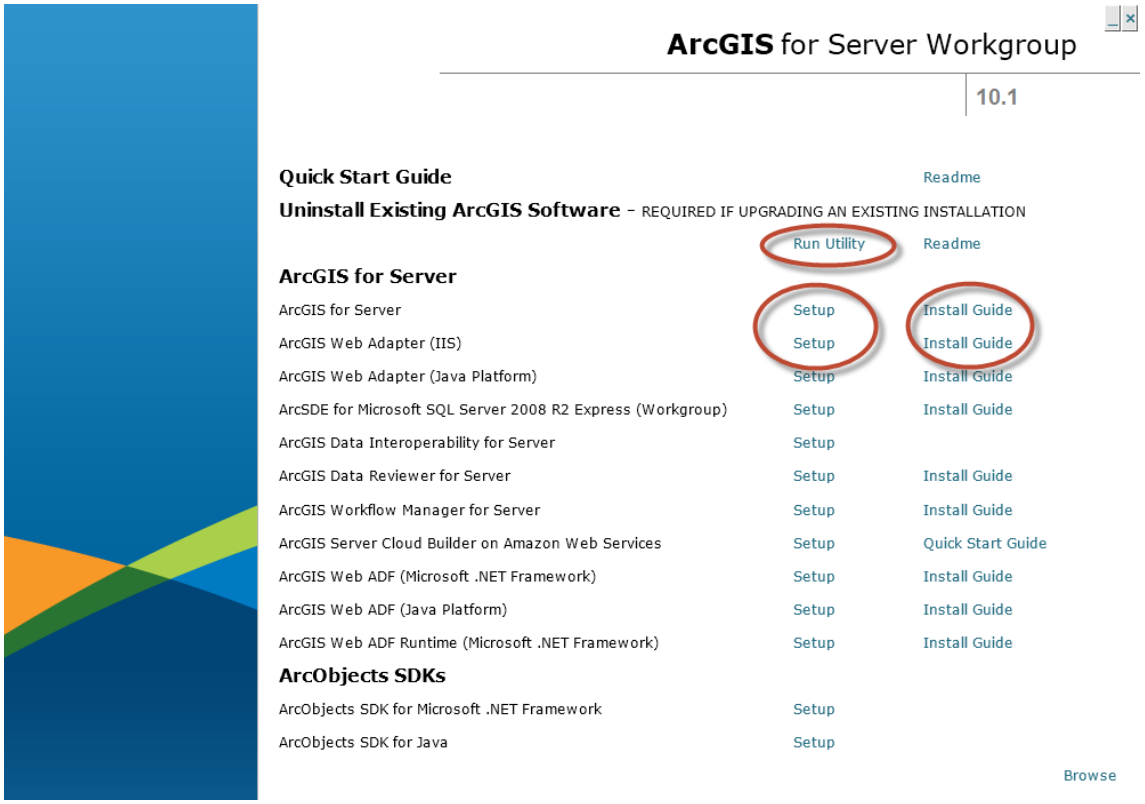

2. Click on ArcGIS for Server Setup
3. If you receive a warning for one or more versions of v10, please uninstall v10 products.

4. Once installation starts click next.

5. In the next dialog, click on 'I Accept' and click next.
6. In the 'Select Feature' dialog, click next and accept defaults.
7. For python, this is needed for geoprocessing and click next.
8. In the 'Specify ArcGIS for Server Account, create a username and passwords for arcgis. Recommend to add arcgis as a user within Active Directory and add this to ArcGIS for Server Account information <domain>\<user>.

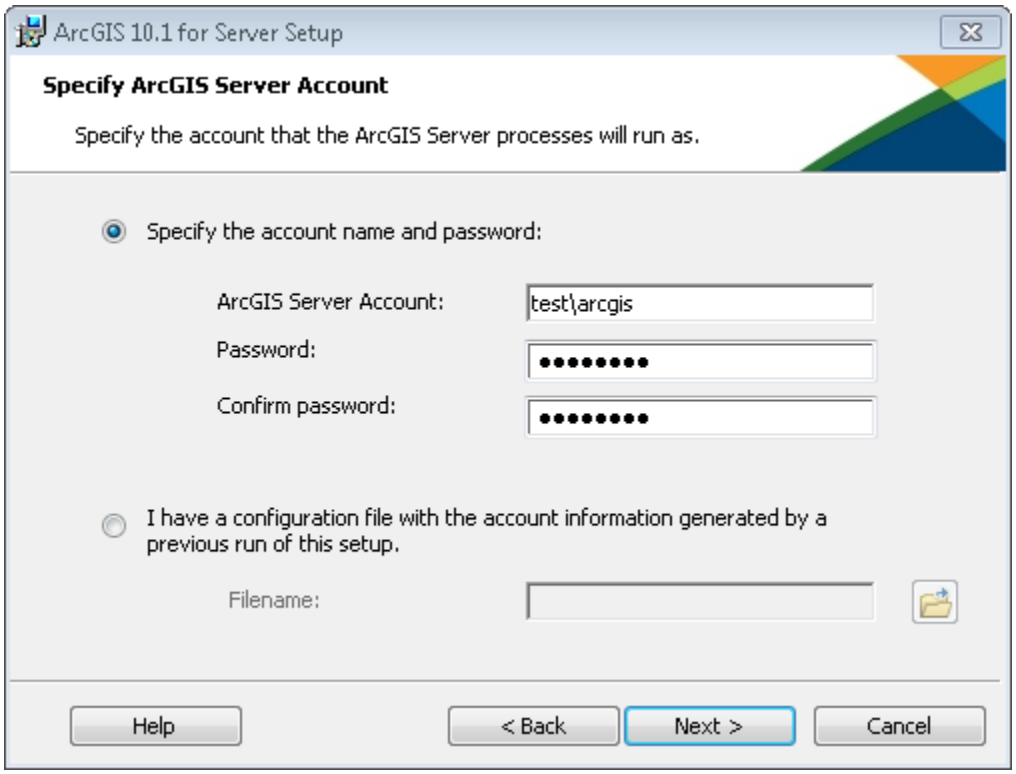

9. In the next screen, you will be given the option to export your configuration; this is recommended and will be helpful during upgrades. Click next and install.

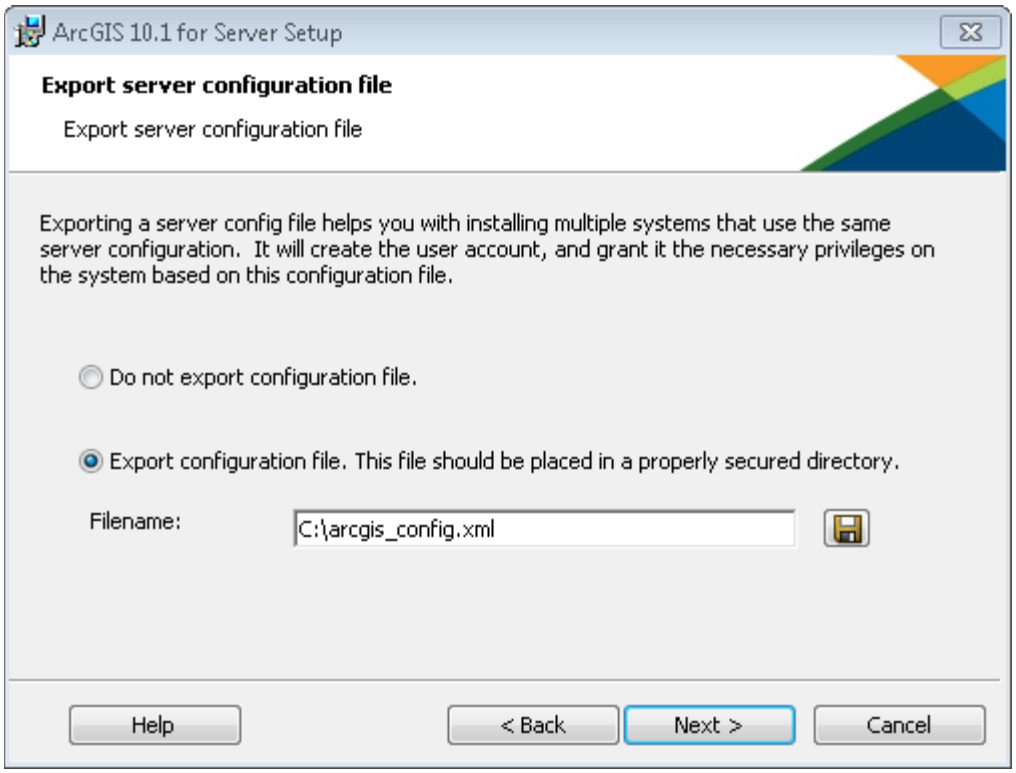

10. When finished, you will need to activate. Refer to Enabling ArcSDE for Workgroup Databases section for more detail about the activation process.
11. When finished activation, you will be prompted to create a new site.

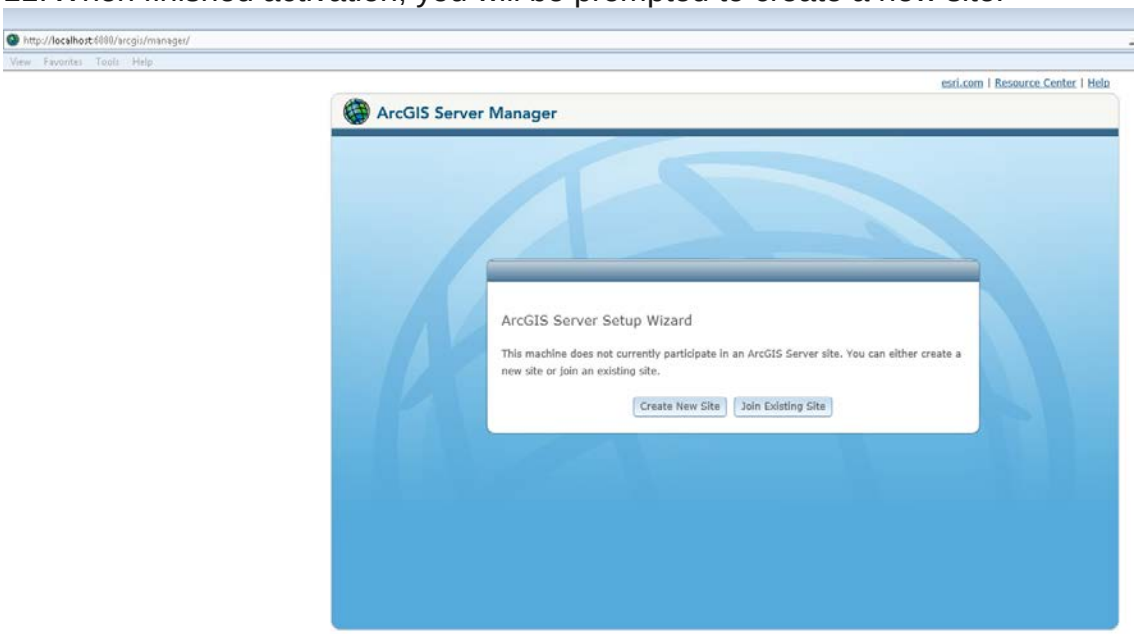

12. Click on create a new site and supply the username and password for the site administrator account. Click next.

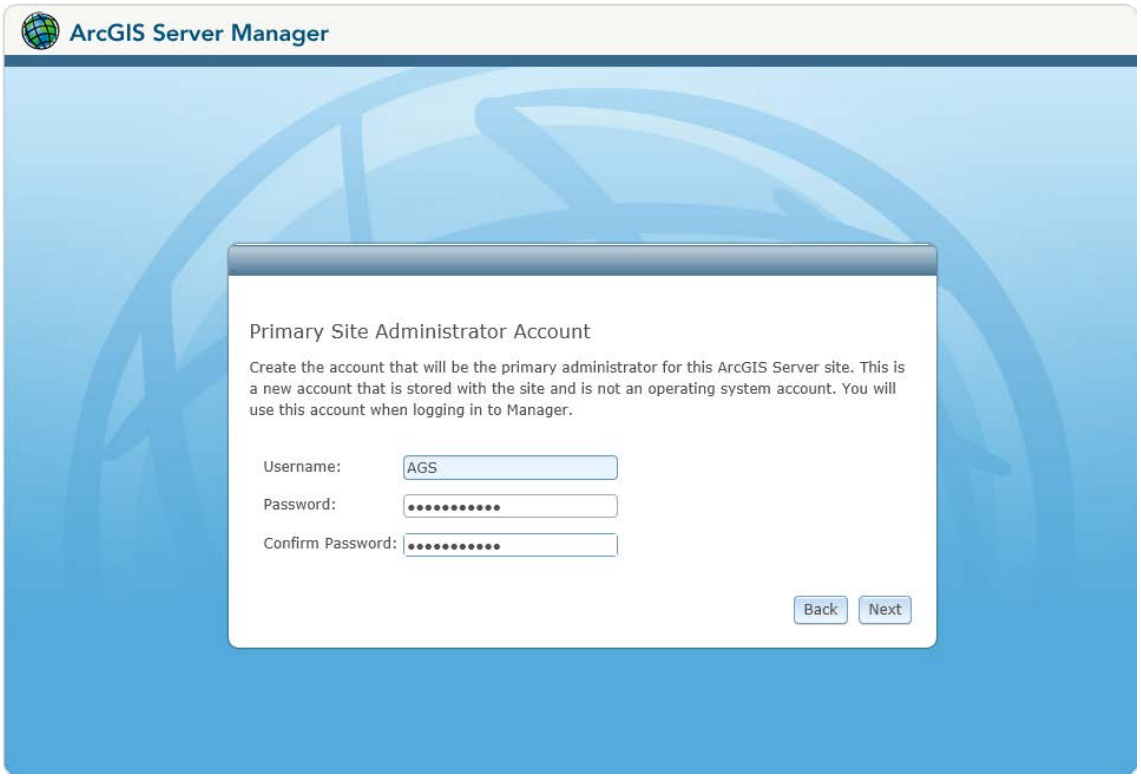

13. Specify the root server directory and configuration storage. In AGS v10.1, your server configuration is stored in the config-store directory and your working directories are stored in directories (arcgiscache, arcgisjobs, arcgisoutput, arcgissystem). In this example we will take the defaults.

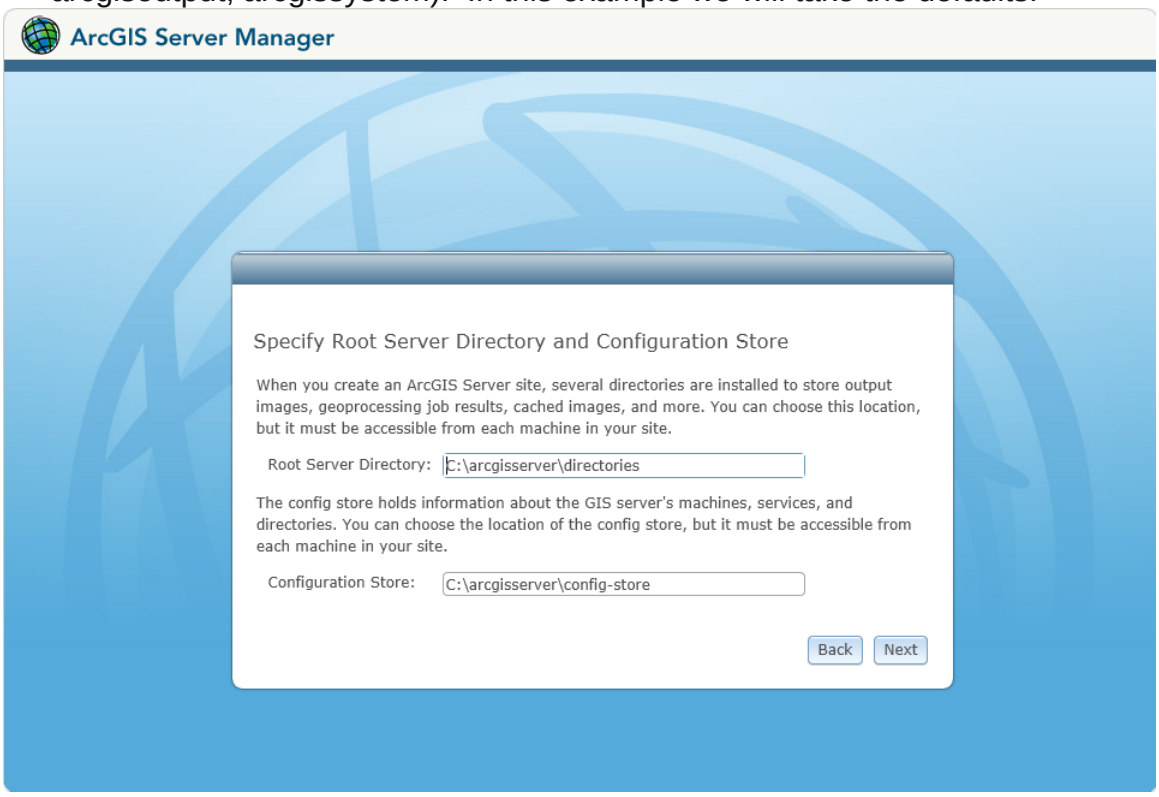

14. Click next and finish when done.

ArcGIS for Server does not need a web client to run. ESRI chose to use Tomcat as the servlet engine for ArcGIS for Server and can be a stand-alone system if used internally. However, if you want to serve webpages out through IIS, consider installing a Web Adapter. ArcGIS for Server web-handler will act as a proxy server that will pass requests to and from ArcGIS for Server. Web Adapter will also be able to create an alias URL location for your ArcGIS for Server. ArcGIS for Server by itself needs to have port 6080 as the rest service endpoint for the URL. For example: http://<internalservername>:6080/arcgis/rest/services is the URL without the Web Adapter. When using the Web Adapter, requests can be masked through an alias URL designation http://<URLname>/<virtualdirectory>/rest/services. The Web Adapter takes care of the requests to and from ArcGIS for Server using port 6080 and translating this to the user. The Web Adapter for ArcGIS for Server can be located on ArcGIS for Server machine running IIS or on a different server within the LAN or a DMZ. For internet and intranet based ArcGIS for Server services, it's recommended to use a DMZ with the firewall open from DMZ web server to the internal ArcGIS for Server on port 6443. The Web Adapter will be installed on the DMZ web server and using ArcGIS for Server installation media. Also, it is highly recommended that the Web Adapter uses HTTPS as the protocol within a secure design as well as running ArcGIS for Server as HTTPS. In ArcGIS for Server design section, we have added a diagram of what this actually looks like.

Connect to ArcGIS for Server from ArcCatalog

1. Open ArcCatalog and go to the TOC. Expand GIS Servers.

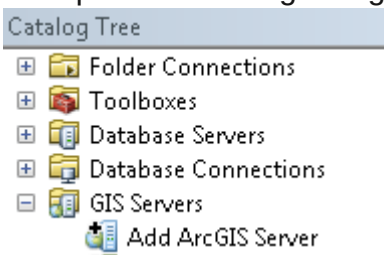

2. Click on Add ArcGIS for Server and choose Administer GIS Server. As a note, you can add users to be able to publish gis services without having administrative rights.

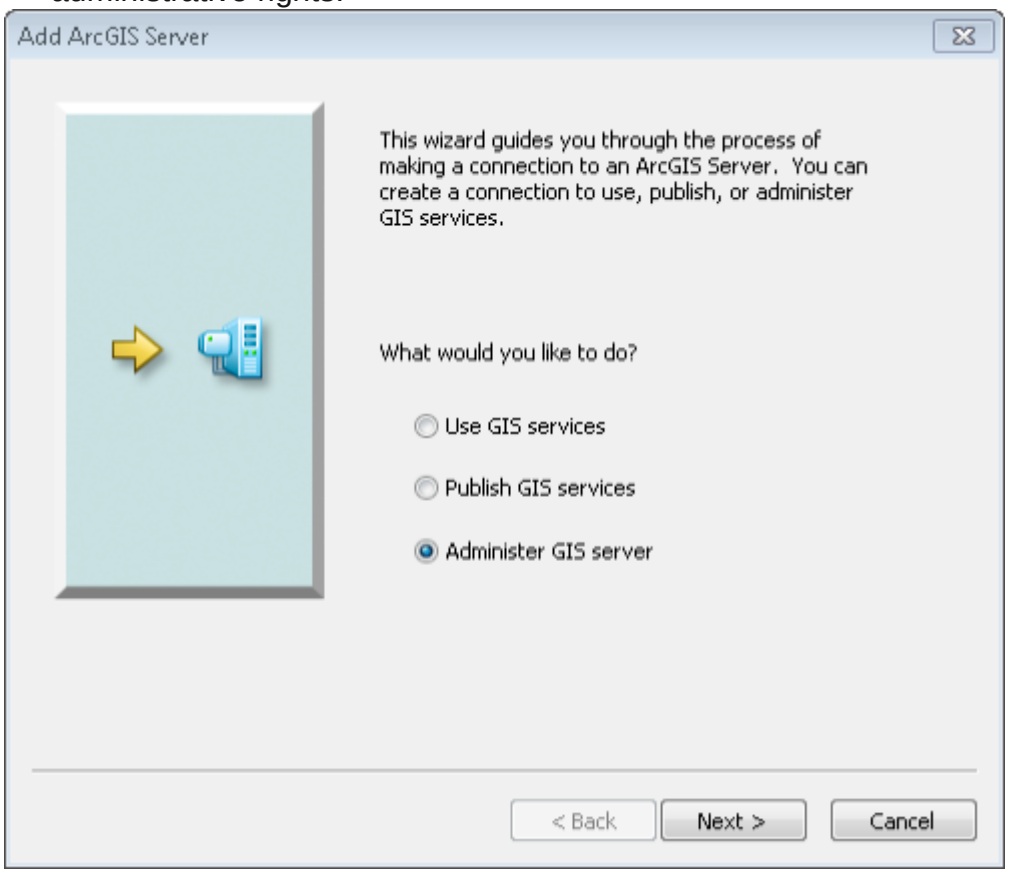

3. In the general dialog, you need to add the URL and the username and password for administrating the server. Make you use the exact case for your username and password as ArcGIS for Server is now case sensitive. If Web Adapter is installed, you can use the Web Adapter URL to connect to ArcGIS for Server if it is setup to accept management.

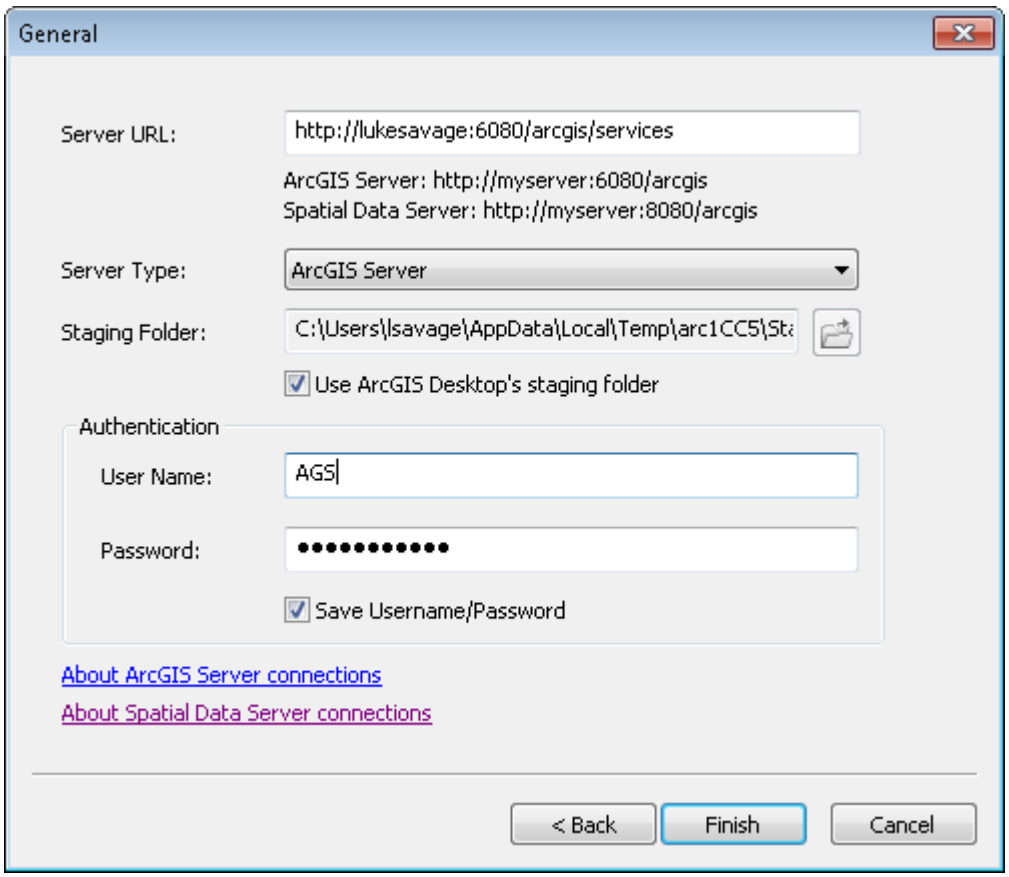

4. Click finish.

Installing and Configuring ArcGIS for Server Web Adapter

1. In the installation media of ArcGIS for Server, click on ArcGIS Web Adapter IIS. Click next and accept terms.

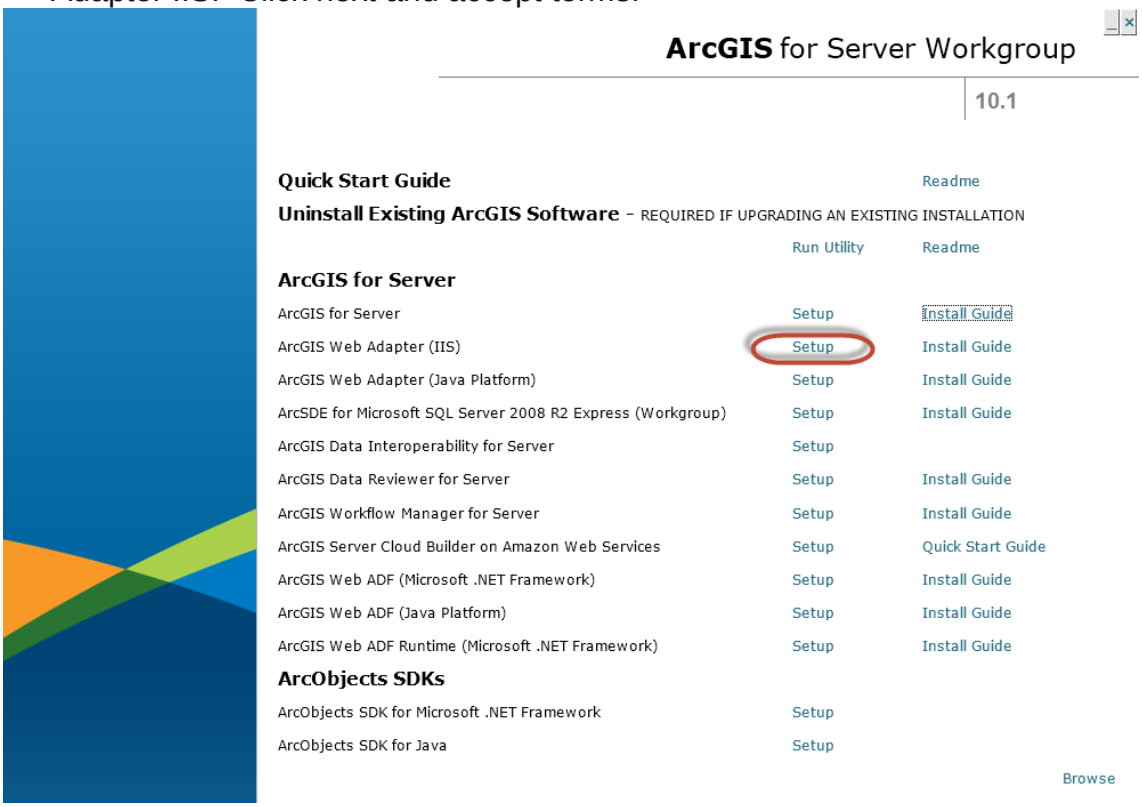

2. Enable the cross-domain Policy Files so that it will handle Silverlight and Flex clients. Click next.

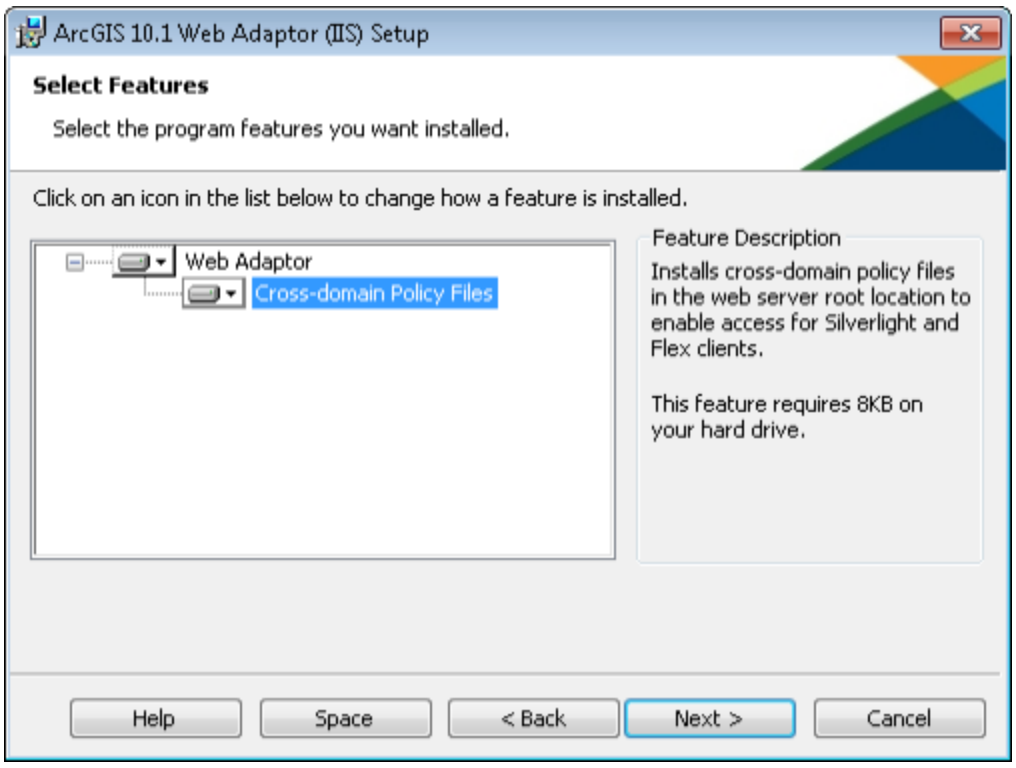

3. In the new Virtual Directory section, create your own virtual directory that will be used to alias ArcGIS for Server URL location. In this example, the URL change will be http://<URLname>/RockStar/rest/services. Click next.

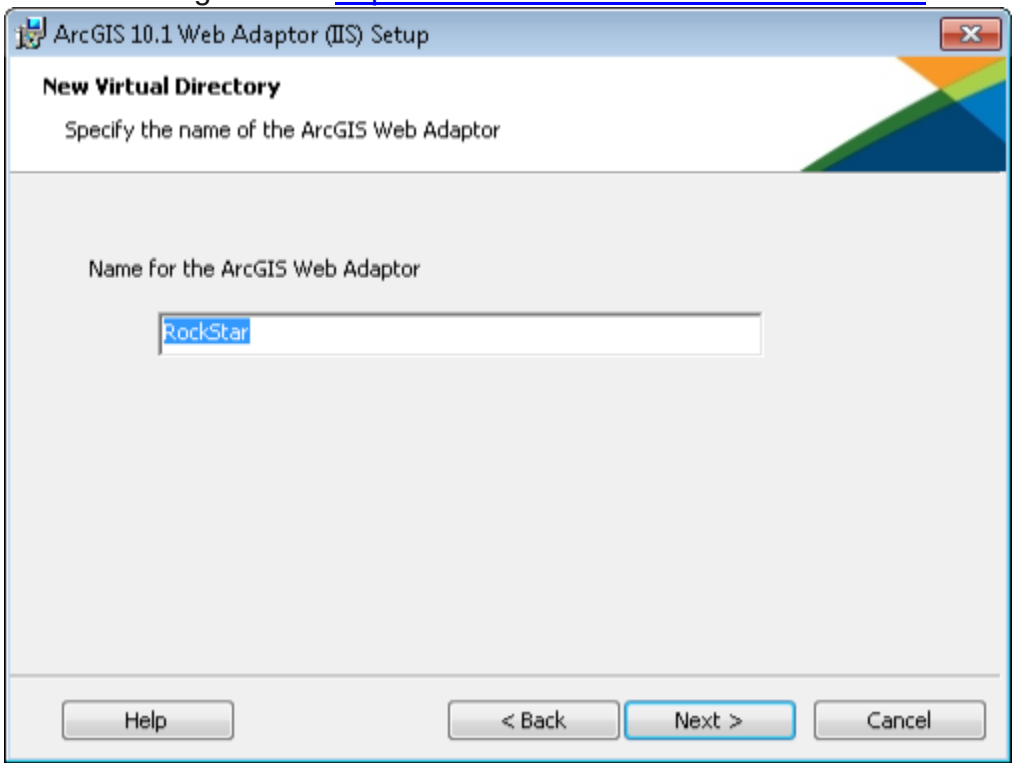

4. Click Install.

5. After finished, you will be rerouted to a site that will allow you to configure your ArcGIS Web Adapter. There are many options but the great thing about the Web Adapter is you can be security conscience. Make sure you figure out if you want management internal only or allow management from both the Web Adapter and ArcGIS for Server. Enabling administrative access to your site through the web adaptor is fine but discuss this option with your network engineer before enabling it to make sure the vulnerability is acceptable. Shared Key is optional and if you want to use it and need a token using a 128 bit encryption, please reference the below link. http://www.csgnetwork.com/wepgeneratorcalc.html (use the ASCII key as the token can only be 16 characters long). Click Configure.

Note: If you don't see the Rest/services URL and it's giving you a code 500, delete your configuration store under the arcgisserver folder and go to the http://<internalservername>:6080/arcgis/manager **and create a new site. This will reset your configuration of your site and will automatically add your web adapter in. Registration of sites sometimes gets corrupted when the web adapter doesn't complete correctly so this step is sometimes required after fixing the web adapter.**

6. If your web adaptor install didn't configure the IIS web pool properly, follow the link below to fix the problem. http://support.esri.com/en/knowledgebase/techarticles/detail/40406
7. Once configured, you should see a change in the green section of the web adaptor page.

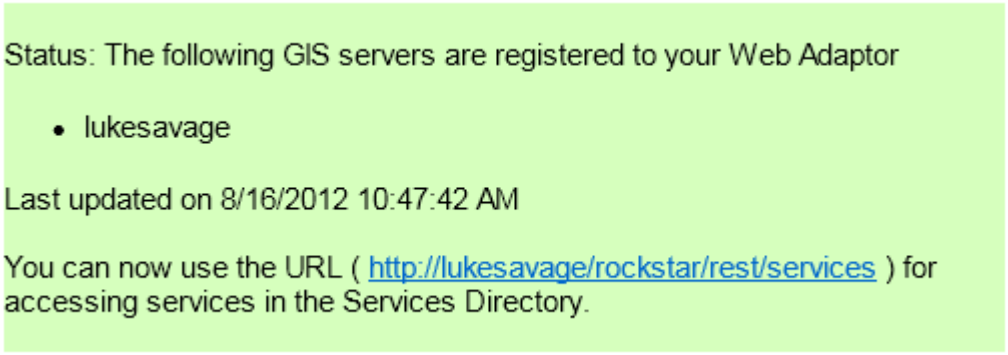

SQL Server Specifications

- Server should have 2-4 cores minimum and 4GB of RAM so that SQL Server can use at least 2GB dedicated for the instance (sql server express can only use up to 1GB RAM per instance). For a Lucity and an Enterprise Geodatabase SQL Instance configuration, I would suggest 4 cores minimum with at least 8GB of RAM for any new database server. Training, Development or Test servers can have less; 2 cores with at least 3GB of RAM which 1GB is dedicated to the SQL Server instance.
- RAID 10 configurations has become a standard RAID best practice. Please don't use onboard RAID controllers as they are not performance based.
- Don't load ArcGIS for Server on a database server. ArcGIS for Server is a memory hog and will fight for computer resources. If you are tight on budget, it would be better to build two mediocre systems than one big system that has both database and ArcGIS for Server on it (SAN and/or Virtual iSCSI configuration is an exception to the rule).
- Plan for data scalability 5 years out. If you have imagery that consumes 20GB and you know you will acquire a new image at least once within the 5 years, scale out four times the consumption rate. Space is cheap and there is no need to be caught in a low memory situation. Remember, digital imagery space is higher as they have a smaller focal length per image. If you acquired analog images in the past, digital images can be 2-3 times the memory size.
- If in a virtual environment, make sure you have enough RAM and storage space for your servers. In the new v10.1 server architecture, you could have as much as three to four servers minimum depending on your system design.

SQL Instance Setup

Install SQL Server for ArcSDE for Workgroup

1. You can install from ArcGIS for Server disk or you can download installation from Microsoft for SQL Server 2012 Express 64bit with Advanced Services.
2. Click on the SQL Server executable.

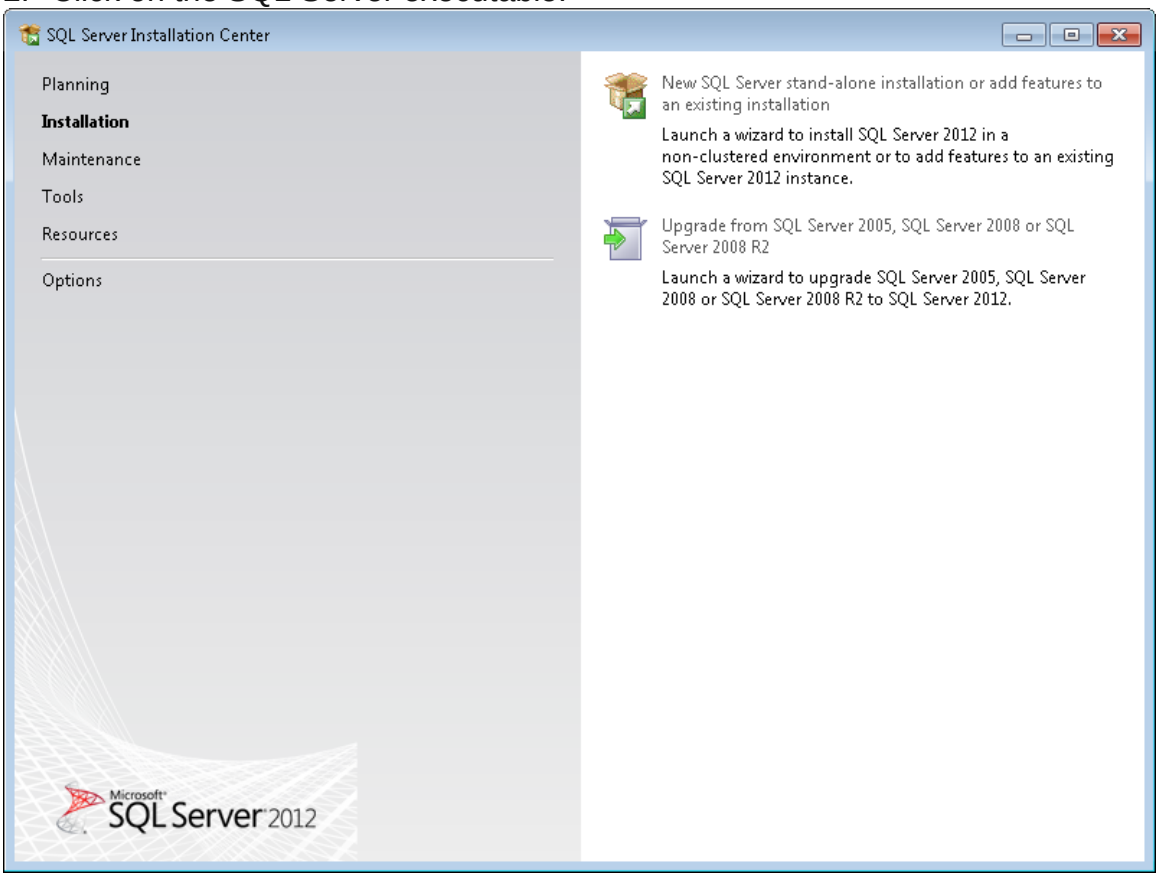

3. Click on new sql server stand-alone installation or add features to an existing installation. In this example, we will be using a fresh installation.
4. You can assign different locations for your shared feature directory, most of the time you'll take the defaults and click next. If you don't want reporting services, uncheck the 'Reporting Services –Native' checkbox. If you don't know and may in the future, you can continue with it checked.

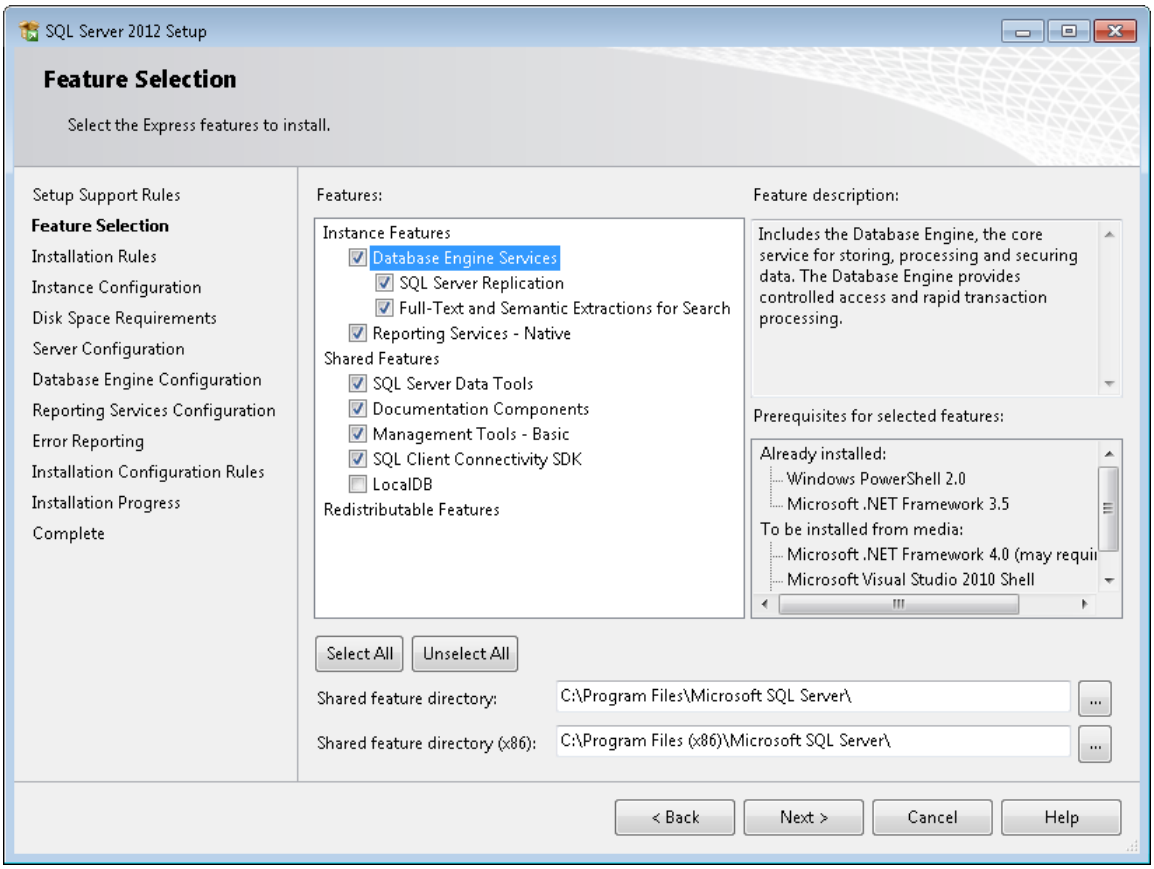

5. You can assign different instance names, install as a default instance, and create reporting services directory but in this particular example, we will use the defaults.

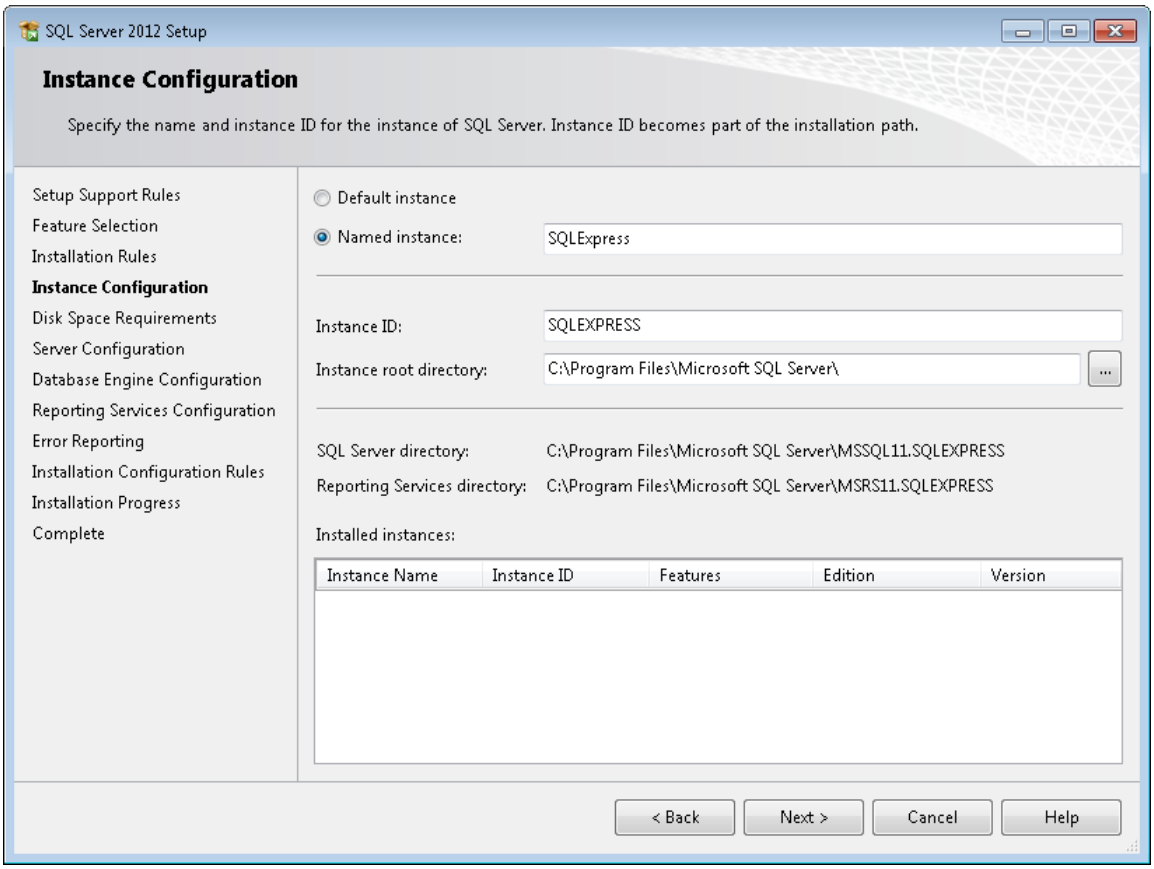

6. In this next dialog, turn on the sql browser to automatic startup

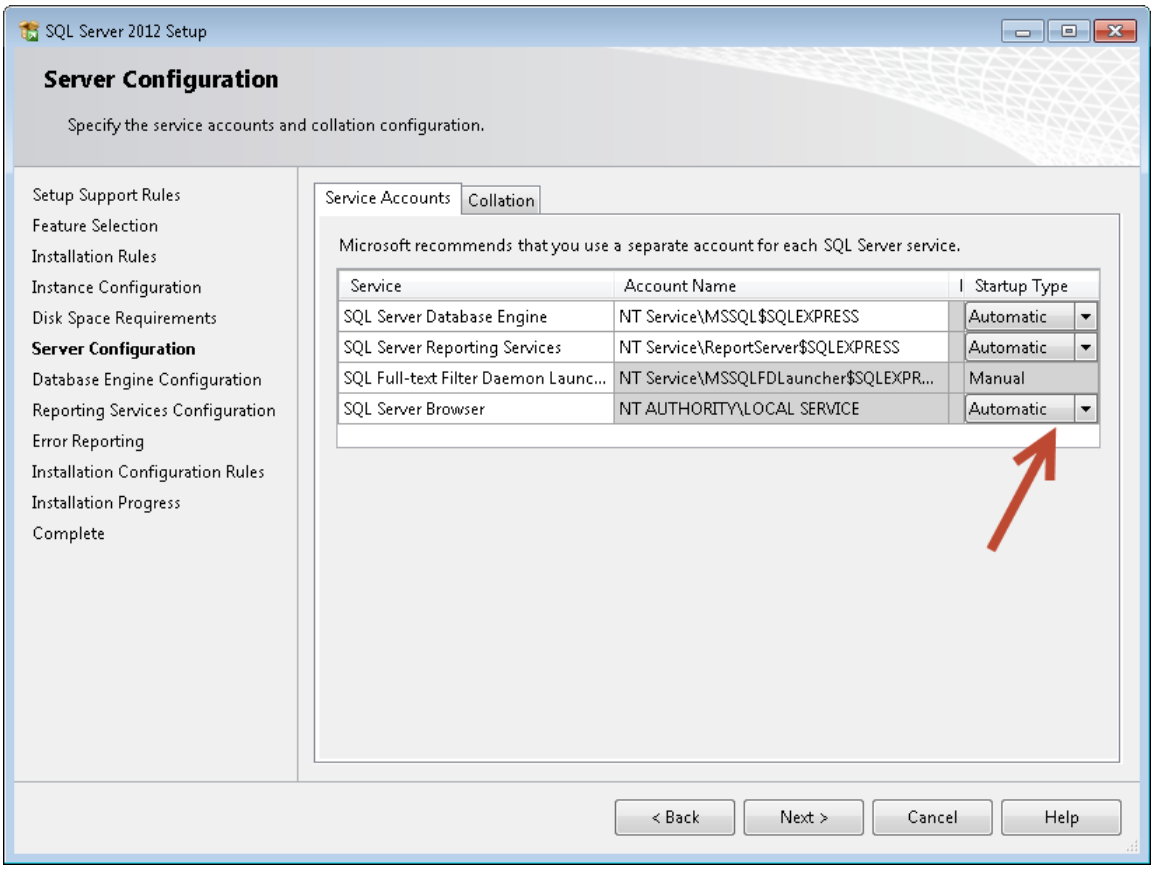

7. For Server Configuration, it is good to enabled mixed mode. This is the case where you have an administrator and you would like to have SA as an admin user that is constant. Windows authentication only is bad if the person that administrates ever leaves and you disable their active directory account by accident. It's not the end of the world, but this is an IT assurance issue. For me, it's better to be safe than sorry. Store your SA password in a safe location that can be accessed in case of administrative role changes. Again, vulnerability and stability assurance measures are the key to a healthy working server environment. Reduce the number of server admins so you have a single point of entry into the database as an administrator. For those using AD domain or LDAP, please change the account name to include a domain user so you can browse the local server when added in a domain. SQL Server Browser, SQL Server Agent and the database engine is necessary to change to be network aware if using backup maintenance plans (database maintenance plans are enabled for SQL Server Standard, Developer or Enterprise only).

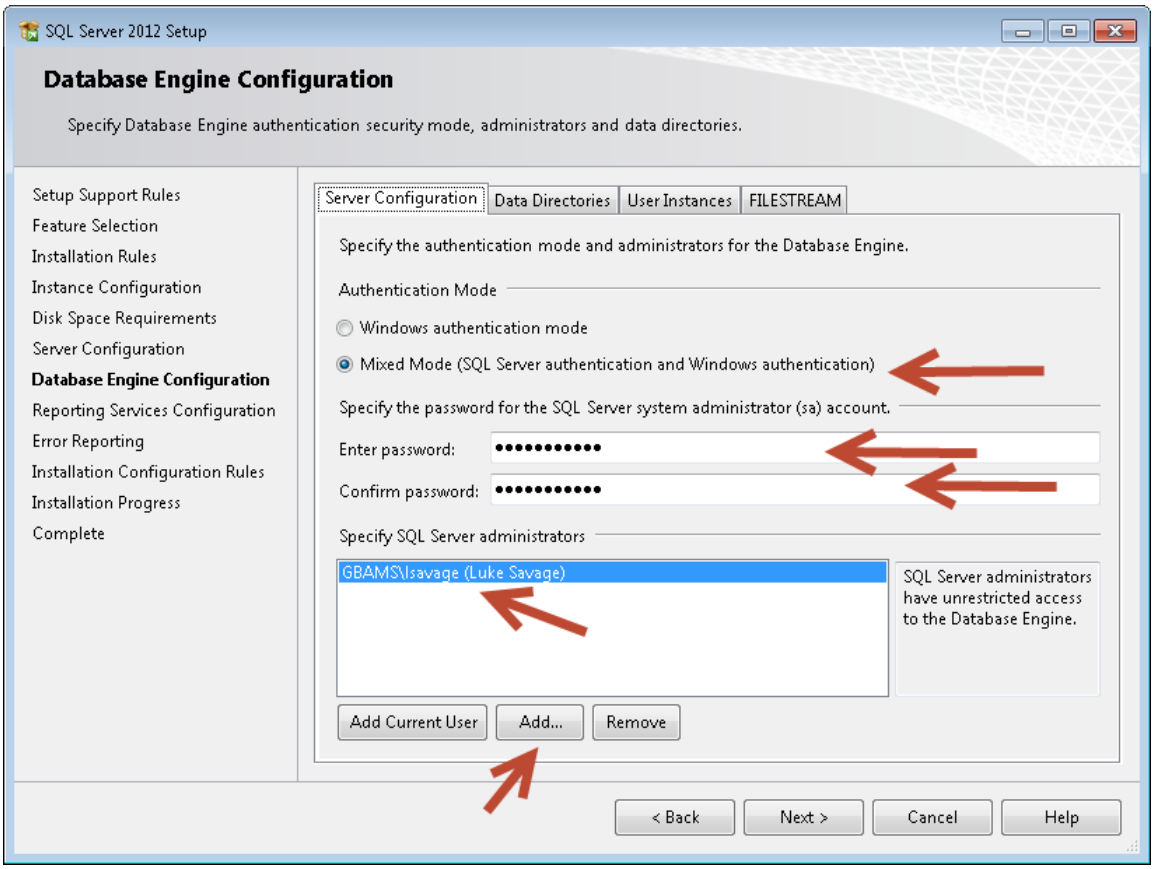

8. On the data directories tab, it is good to change the location for easy backup queries and administration. However, system tables will use the Microsoft default location within the root directory of the databases. I would not recommend enabling user instances on the user instance tab as well as filestreaming unless you plan on loading large documents outside of the Geodatabase as blobs. Click next when complete.

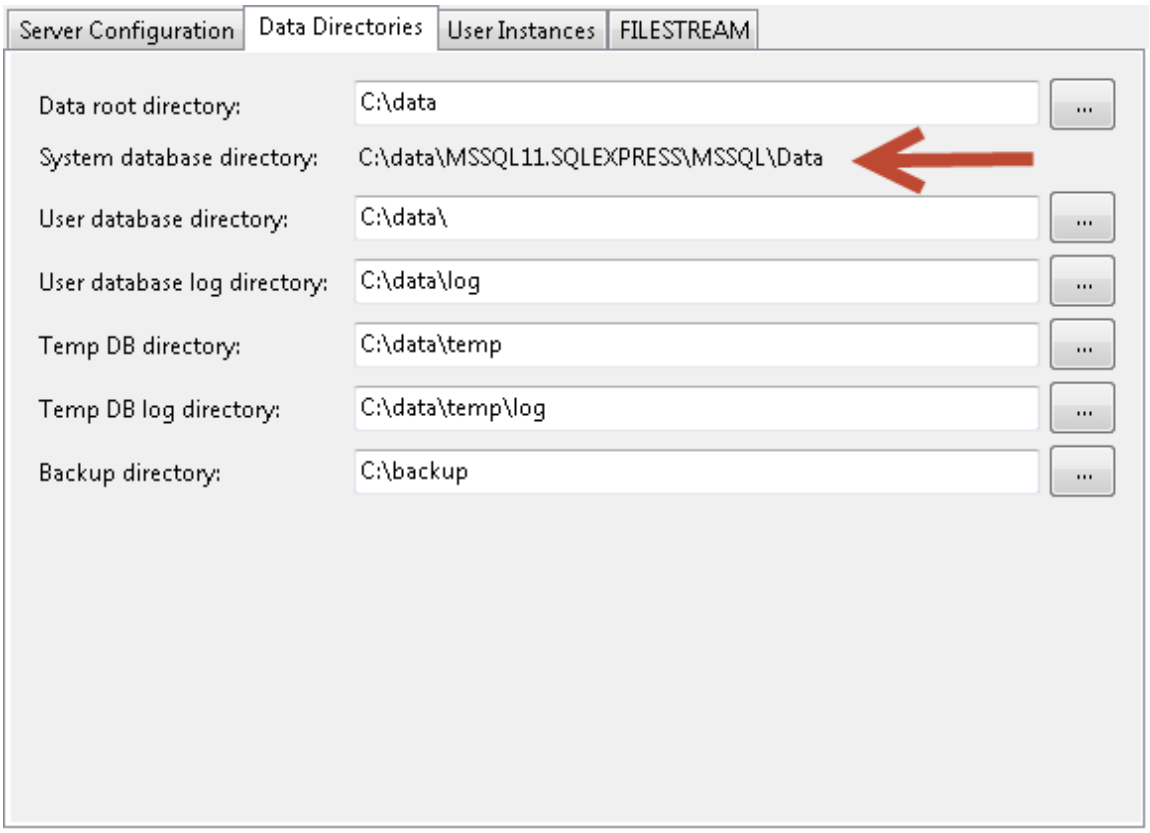

9. Optional: In the reporting services configuration, if you are using the reporting services, install and configure. If not, I would install and not configure. I chose to install but not configure right away. If you chose not to install at all during the setup process, you will not see this dialog. Click next.

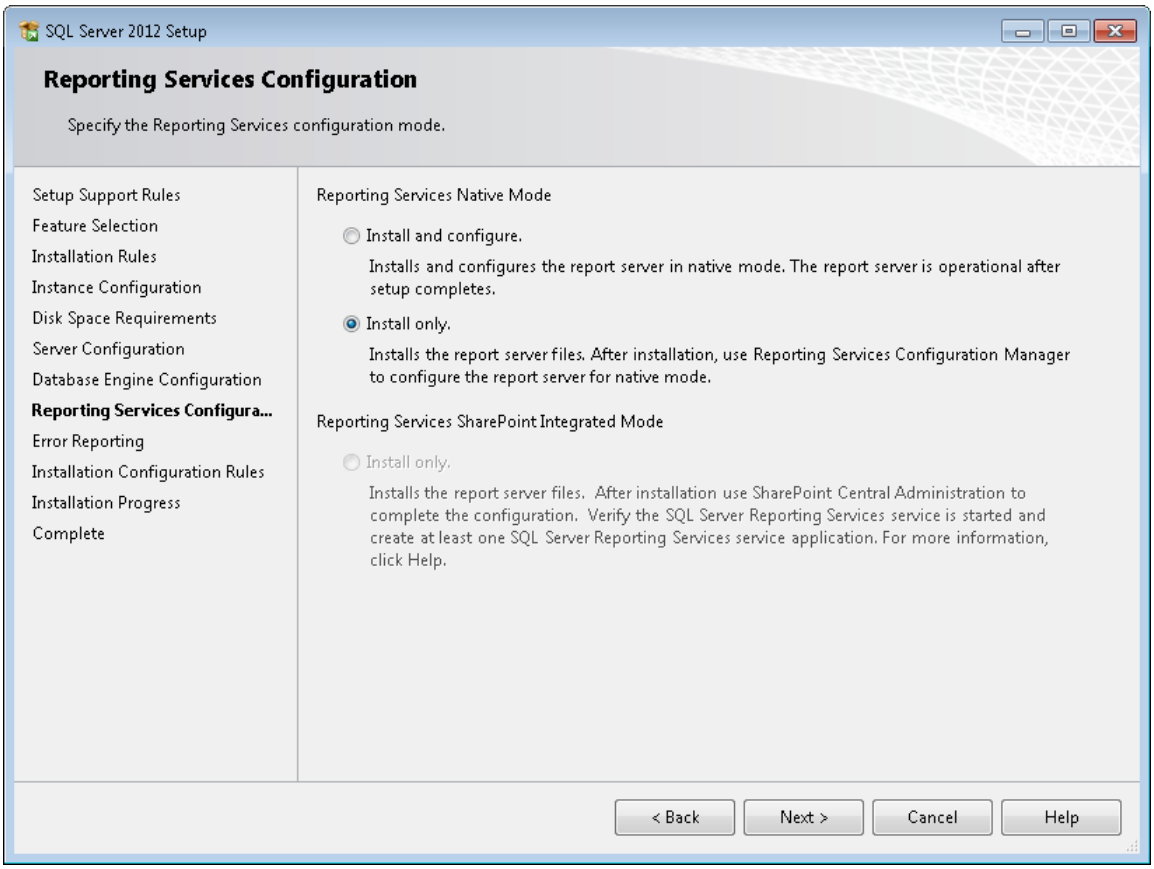

10. Your choice of sending Microsoft the error reports and usage automatically. I generally do this for a development or test environment. For a production environment, leave unchecked. Click Next.

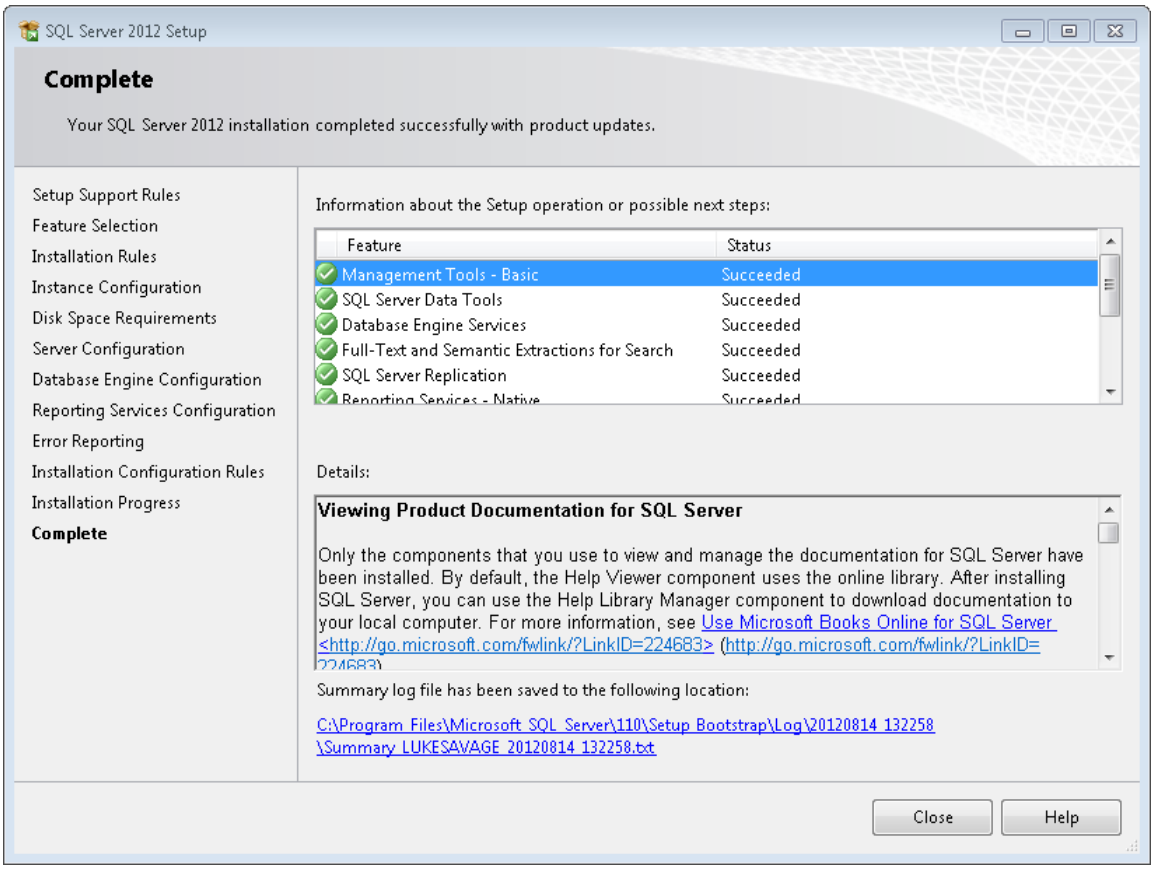

11. If complete, you have successfully created a sql instance. If not, go through the troubleshooting information listed in the error log.
12. Go to start/all programs/Microsoft sql server 2012/sql server management studio. If you had sql server 2008 before, you can import profile settings.

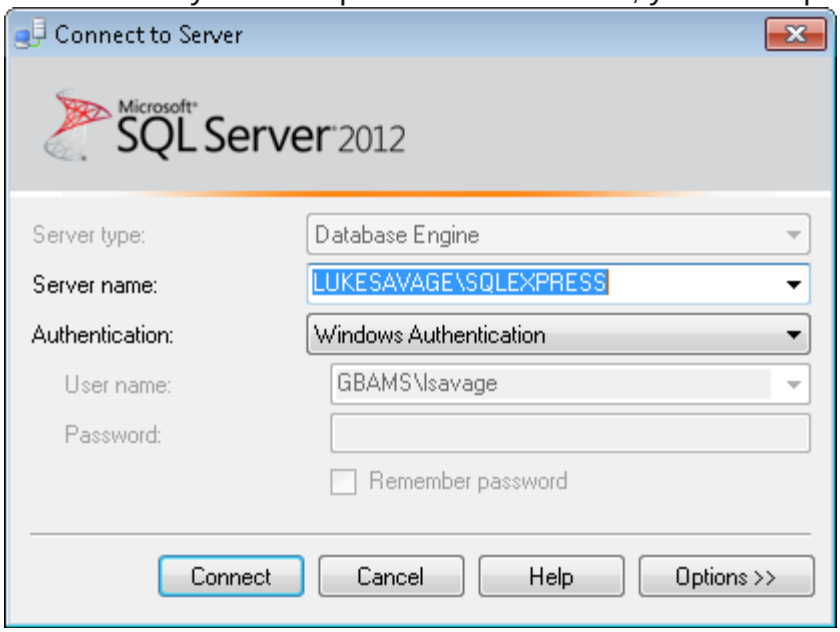

13. Click connect
14. Right click the instance you connected to and choose properties

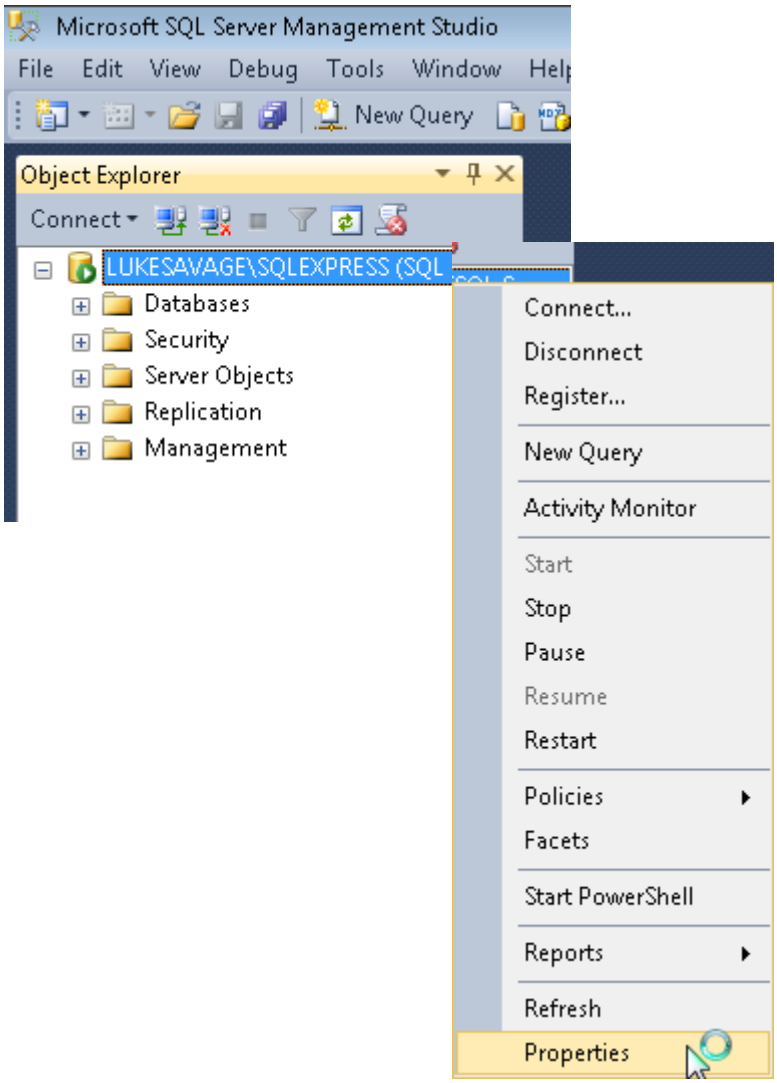

15. Go to the memory tab and click on maximum server memory. You will want this to be a realistic number as the default puts an ungodly amount of maximum ram that will far exceed your server limit. Because we are treating this as a production server, you will want to dedicate resources to SQL Server. For the operating system, at least give the operating system 3 GB of RAM and for SQL Server adjust appropriately. In this case, I'm giving SQL Server 2 GB of RAM but in can be more depending on your instance load and available memory. All other defaults are fine.

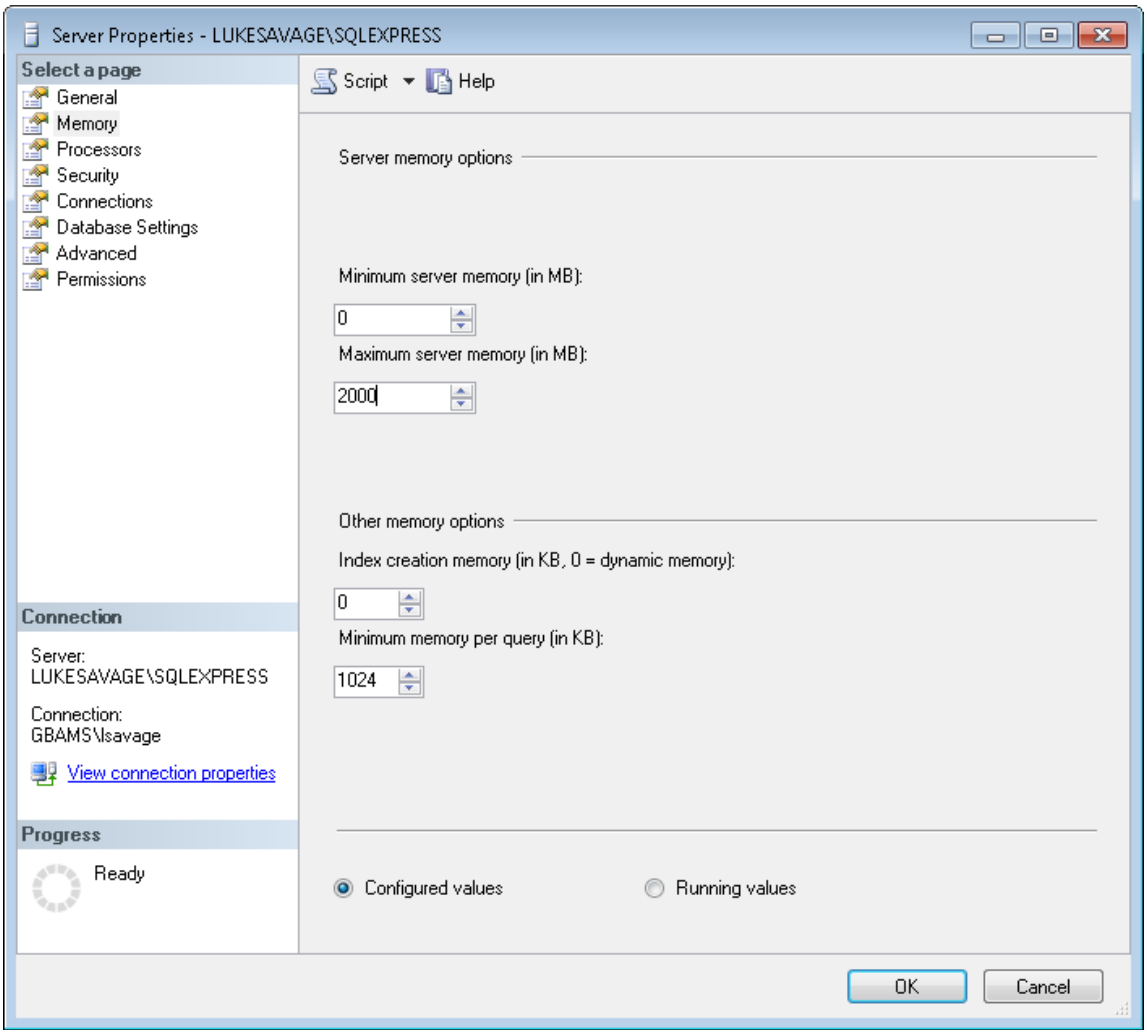

16. Tab down to Processors in the TOC, I would highly recommend enabling 'Boost SQL Server Priority' on the server. Especially if you are loading ArcGIS for Server on the same server, which I wouldn't recommend. However, if you are forced to use one server, ArcGIS for Server will try to push SQL Server out of the way so boosting SQL Server priority will help keep the re. This item will dedicate resources to SQL Server and keep the resources available for SQL Server. **DO NOT ENABLE WINDOWS FIBERS (lightweight pooling)!** This is for older technologies and ArcSDE from v10 and higher uses XML calls to the SQL Server database so keep windows fibers disabled.

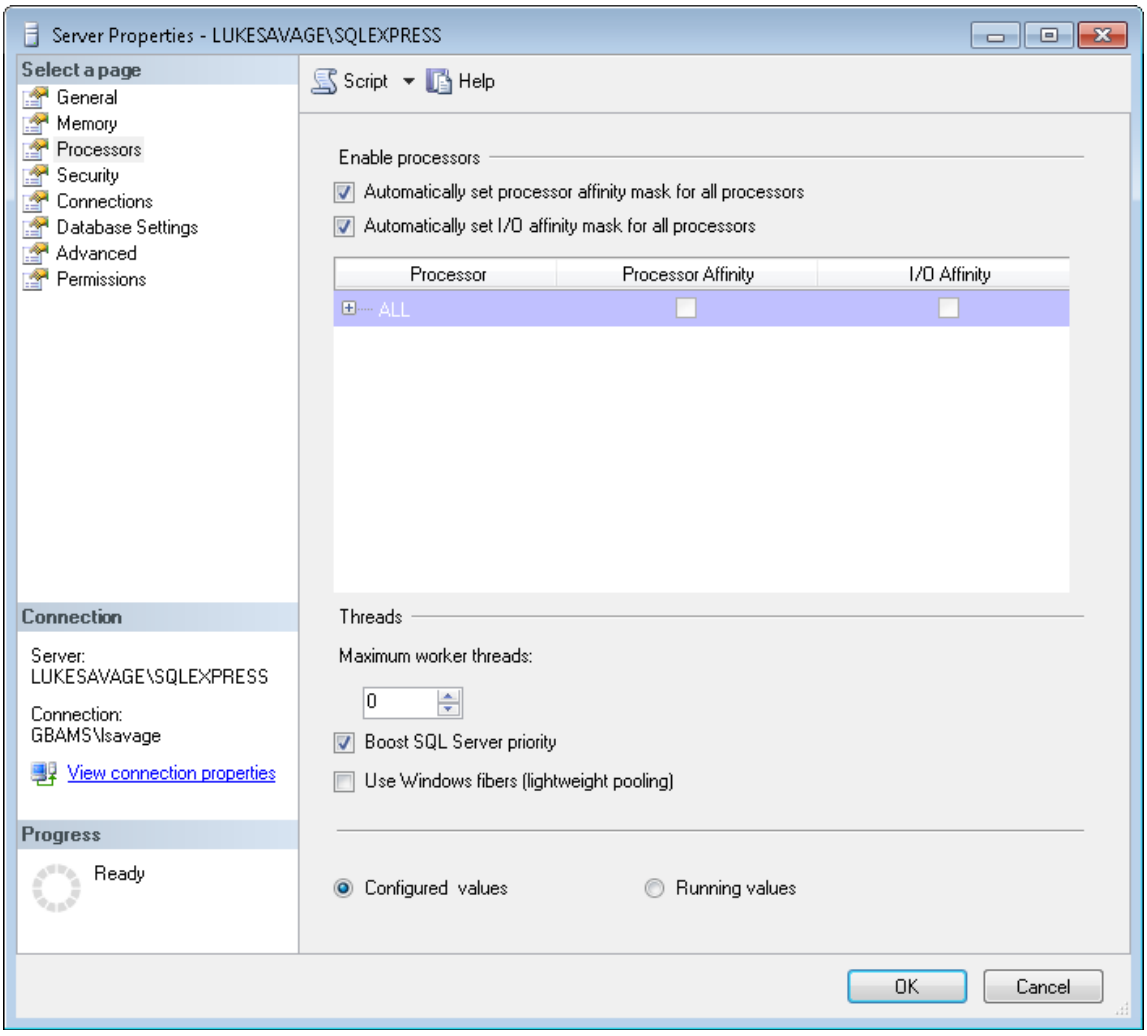

17. Double the network packet size for the Geodatabase connections in the advanced menu item (from 4096 to 8192). For network traffic and performance, the Geodatabase network packets need the additional size increase.

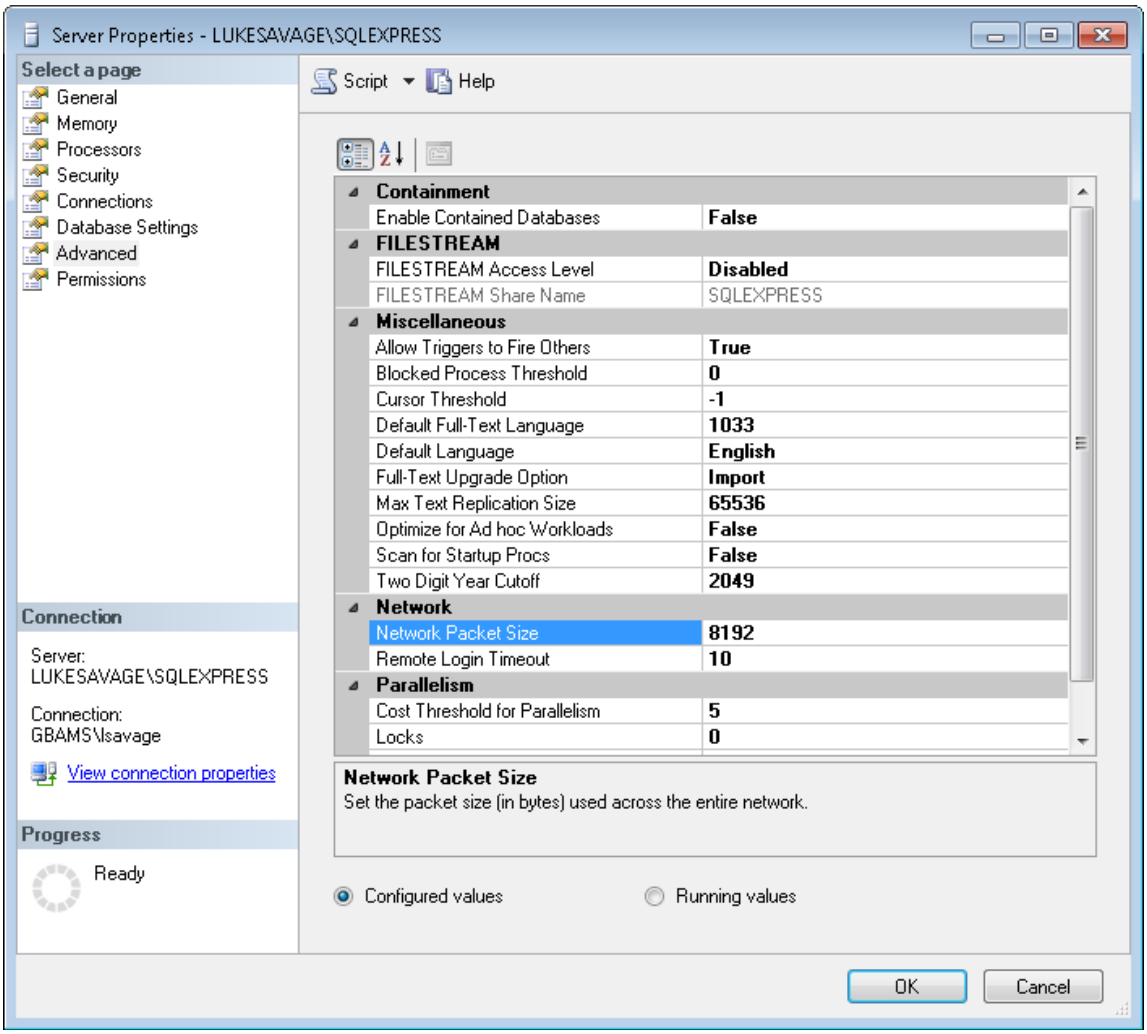

SQL Server Enterprise Setup

Very similar to the workgroup setup above only that I recommend using default instance and locating all databases for Lucity and ArcGIS in one instance for ease of administration and resource dedication.

Enabling ArcSDE for Workgroup Databases

1. Click on the setup for ArcSDE for Microsoft SQL Server 2008 R2.

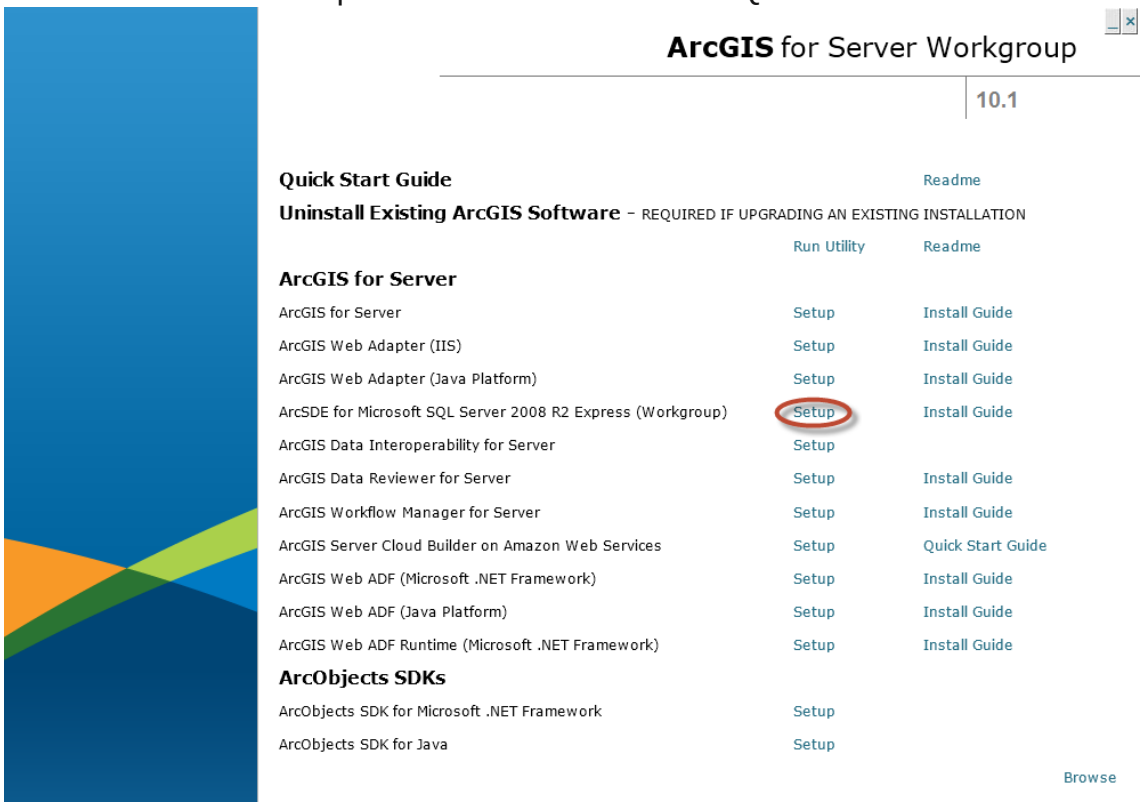

2. Click Next

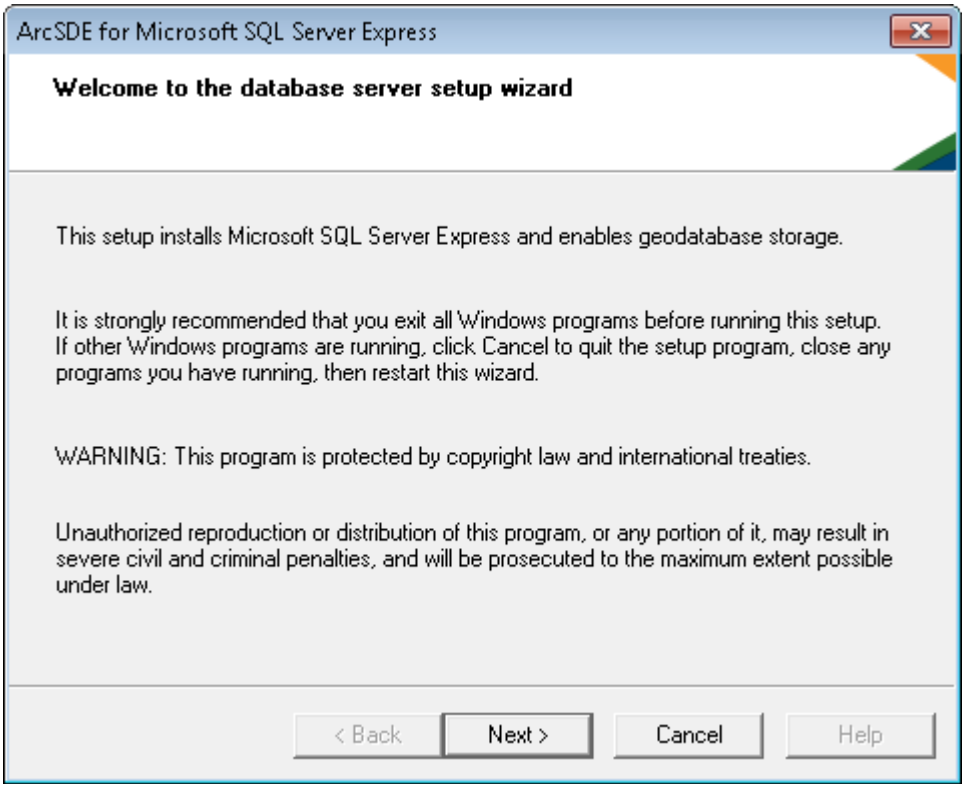

3. Uncheck the checkbox next to Microsoft SQL Server Express Edition

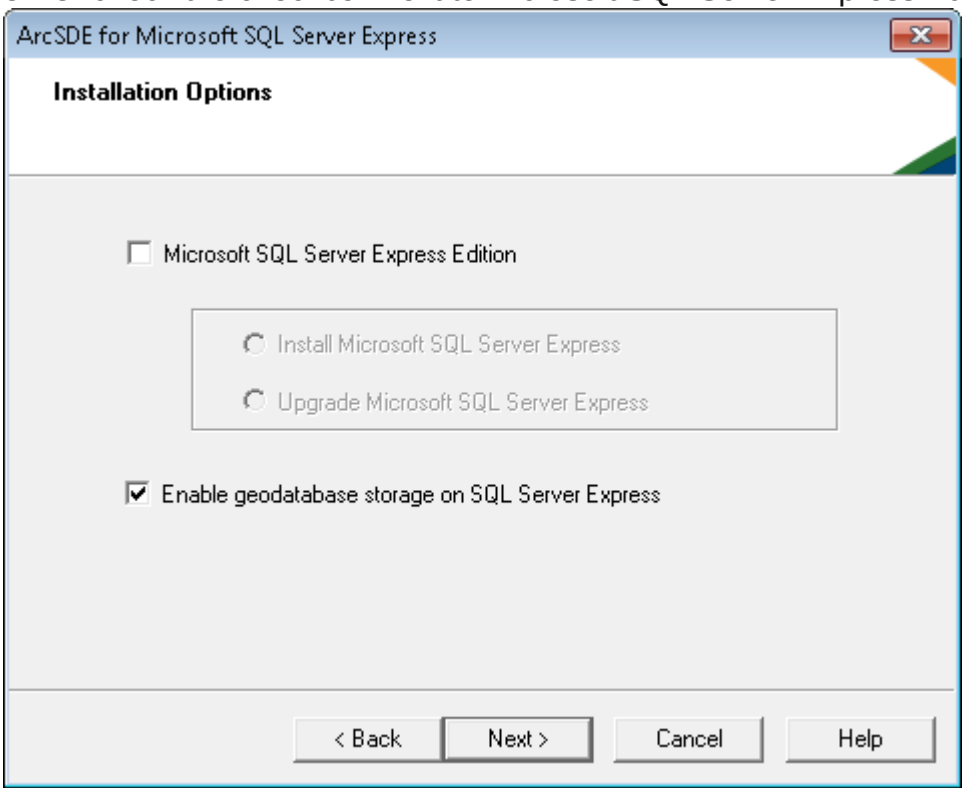

4. Click next
5. Supply your instance name and windows login

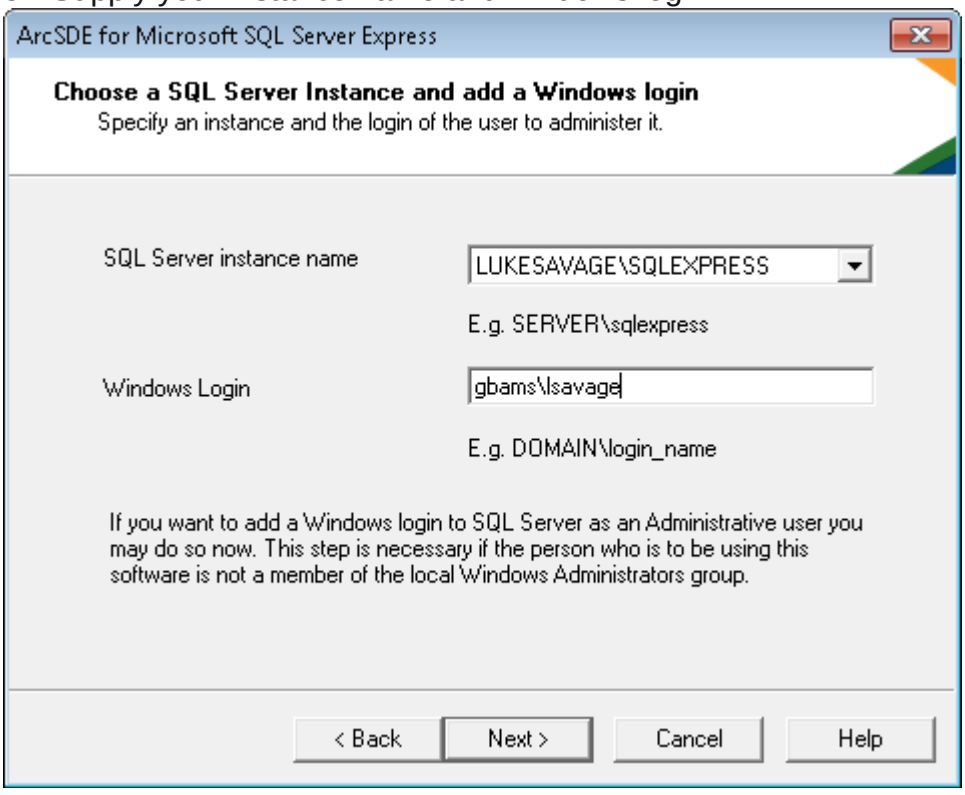

6. If successful, click ok and move to step 24. If not, fix the connection issues.

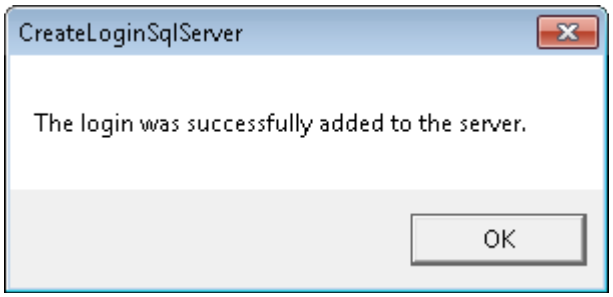

7. The next section will authorize the SQL Server Express instance. For SQL Server enterprise, this is different as the subsequent topics will show you how to authorize Enterprise Geodatabases for SQL Server. Click next

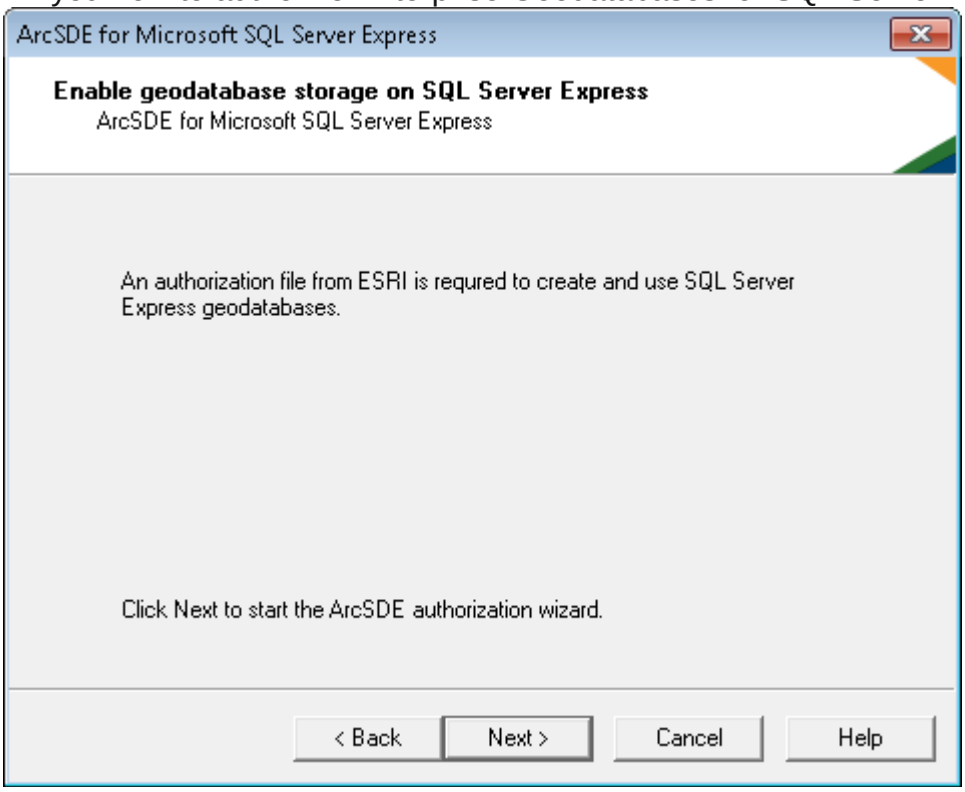

8. The authorize option allows you to point to an authorize file supplied from ESRI or to authorize the software installation. In this example, we will be choosing the later. Click next

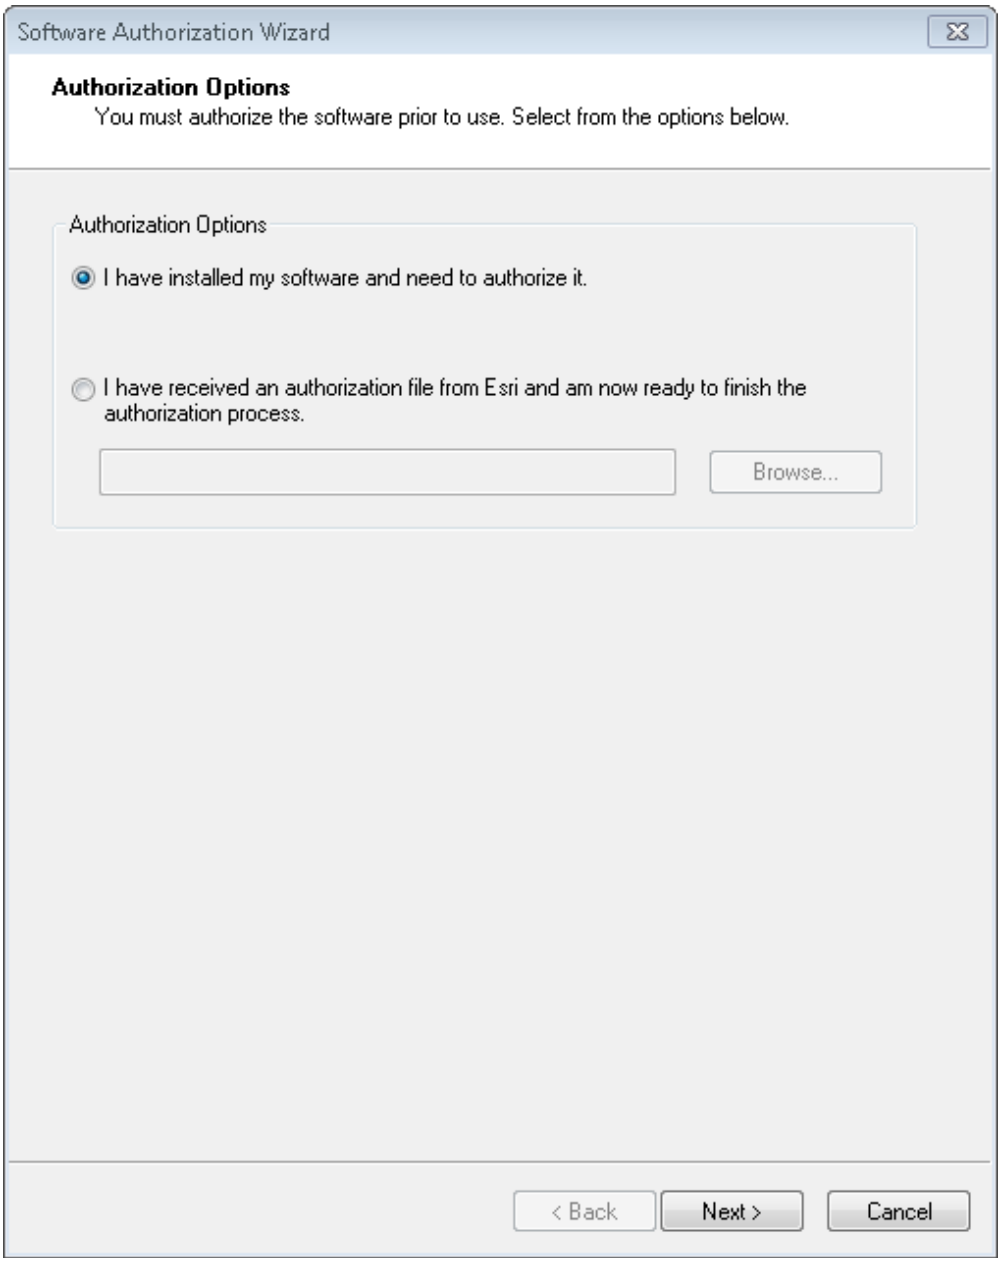

9. Again, you have the choice of email, website or through the internet. We will be performing the later. Click next.

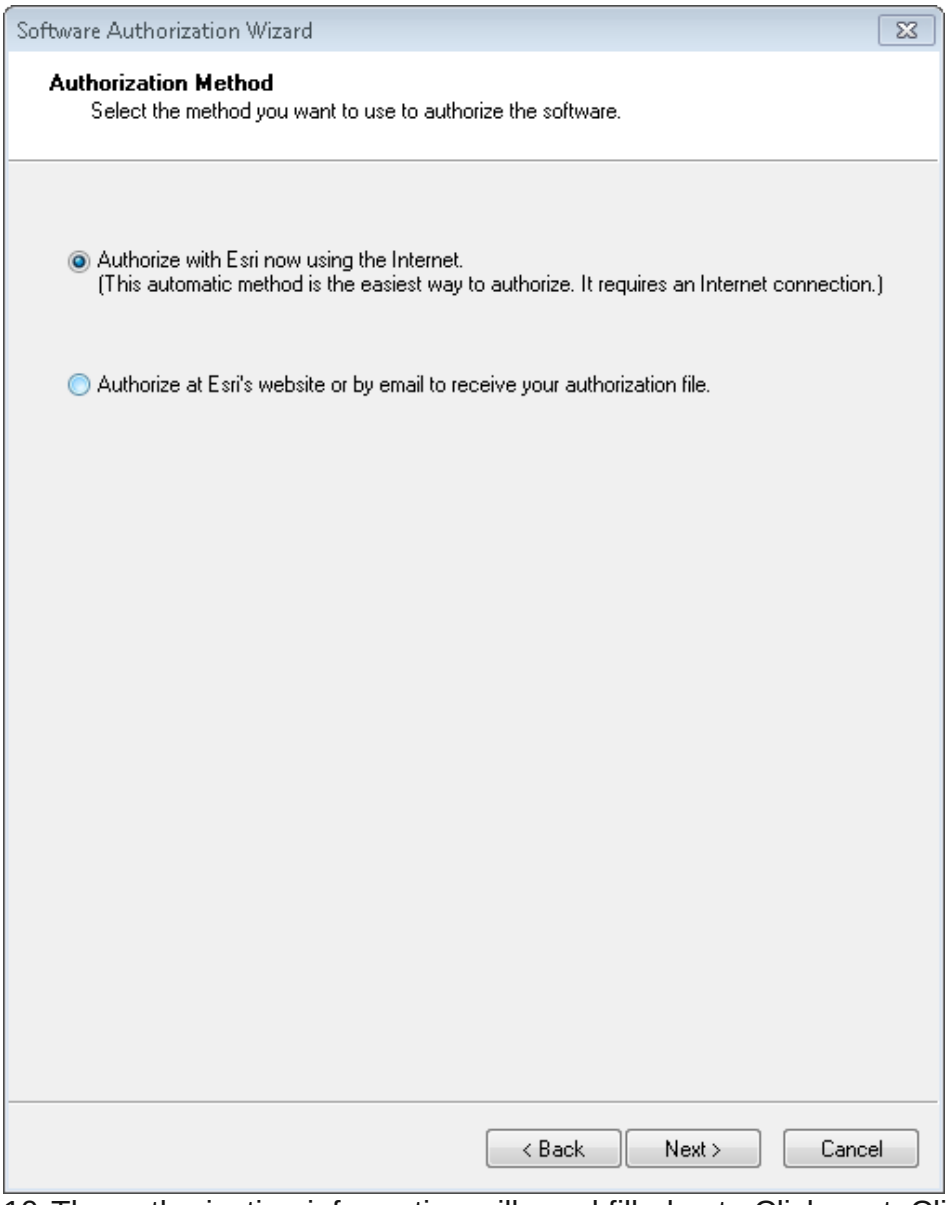

10. The authorization information will need filled out. Click next. Click next.
11. In software authorization number dialog, you will need to supply the authorization number that ESRI has supplied to you for ArcGIS for Server Workgroup. Click next.

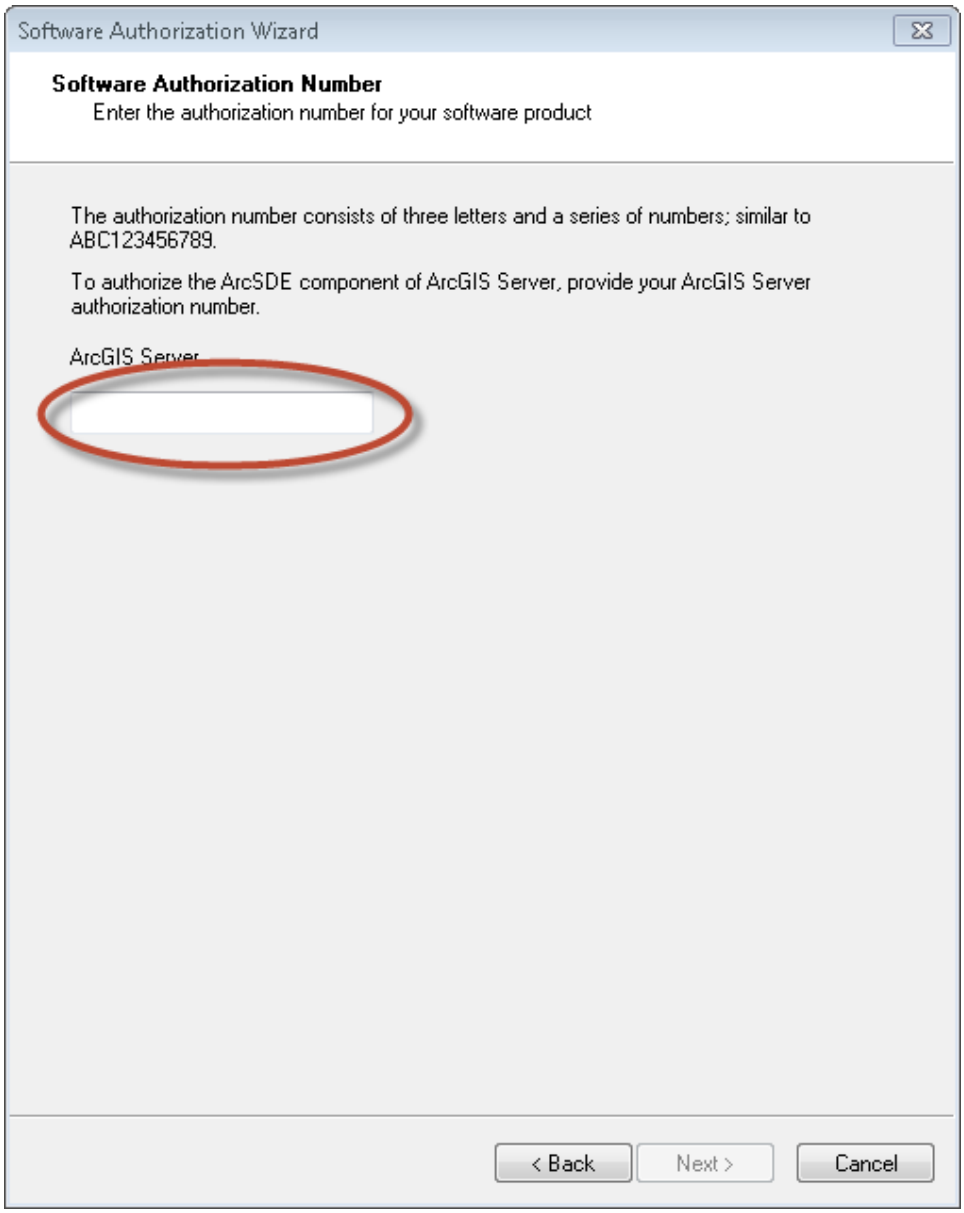

12. If you do not receive a congratulations message as it may be an internet interruption. If you are not successful, contact your ESRI representative for an authorization file. Click Finish.

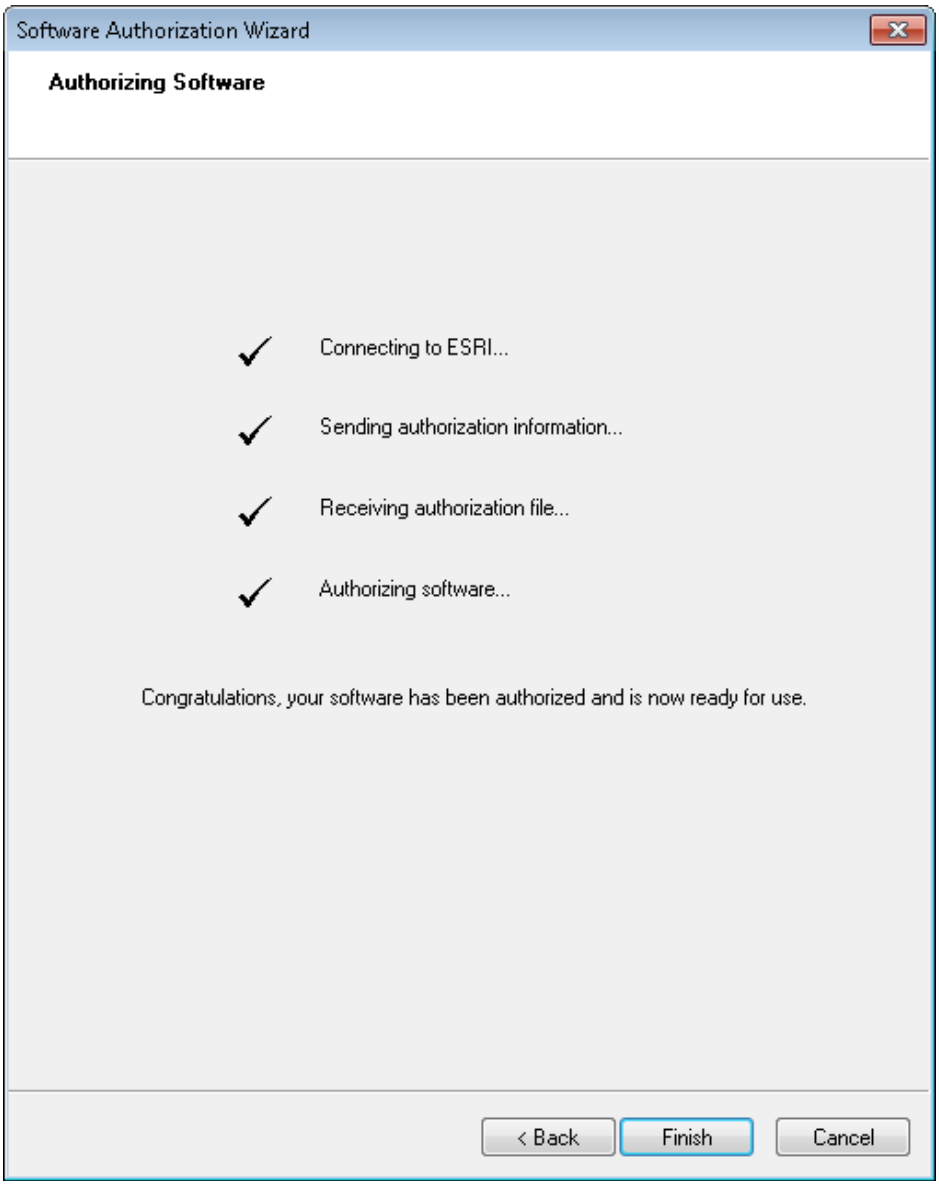
Software Authorization Wizard
Authorizing Software
Connecting to ESRI...
Sending authorization information...
Receiving authorization file...
Authorizing software...
Congratulations, your software has been authorized and is now ready for use.
< Back
Finish
Cancel

Install SQL Server Native Client

You will need to install the exact version of SQL Server Native Client on each workstation that will connect using ArcMap or ArcCatalog to the Geodatabase.

13. When using SQL Server 2012 with ArcSDE v10.1, each client (ArcGIS for Desktop, Server) needs the SQL Server Native Client 64 bit. Also, these versions need to match. For example, SQL Server 2012 original install media is a different version than the latest updated SQL Server 2012 from Microsoft updates. If the native client doesn't match the patched version of SQL Server, it won't work. ESRI is requiring anyone connecting to ArcSDE using an installed sql native 64 bit client with identical 2012 versions or SQL Server 2008 R2 native client. Latest SQL Server 2012 Native 64 bit Client download link: http://go.microsoft.com/fwlink/?LinkID=239648&clcid=0x409

Connect and Create Geodatabase in Workgroup

14. Open ArcCatalog or ArcMap -> Embedded ArcCatalog and expand the database server's item in the TOC. Click on Add Database Server.

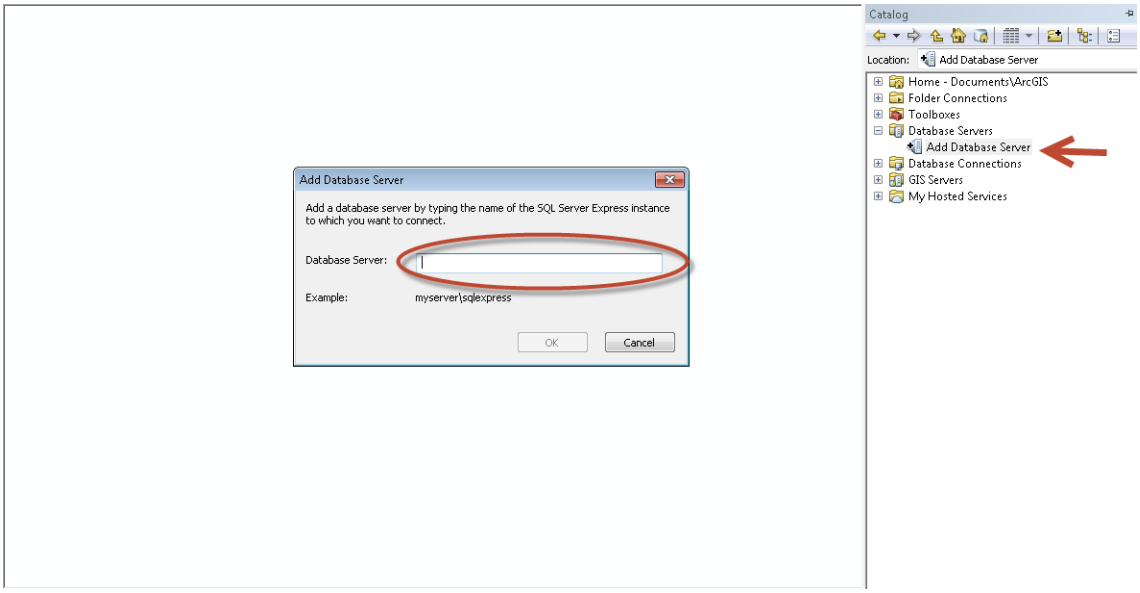

15. After the connection has been established, you have the choice of creating a Geodatabase by right clicking on the instance and choosing new Geodatabase.

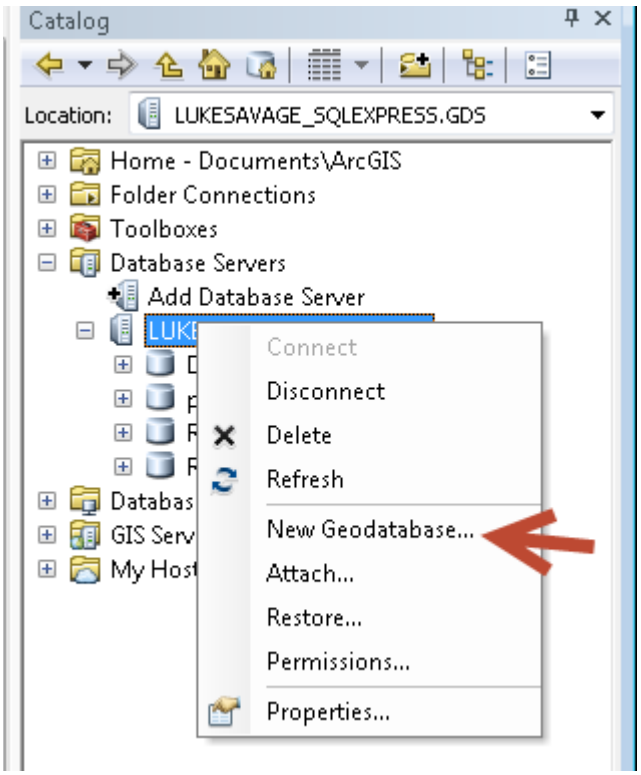

16. Create new Geodatabase and change the location if needed of the database file. Also, you have the option of changing the initial size but the autogrow method will be established. Click OK.

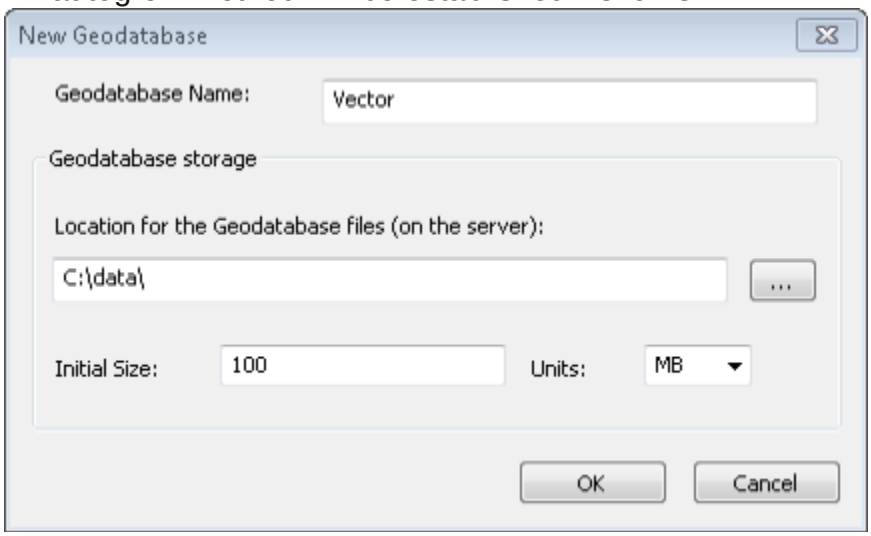

17. Once complete you will see a new database in your instance. Afterwards, you can add users to the instance by right clicking on the instance in the TOC and choosing permissions.

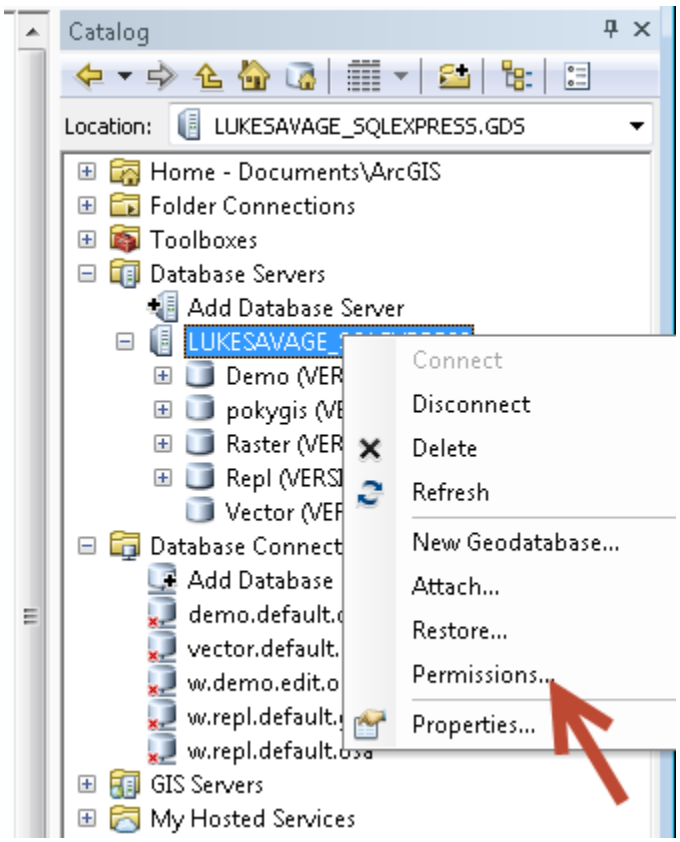

18. This is where you can add users and groups within your organization from Active Directory.

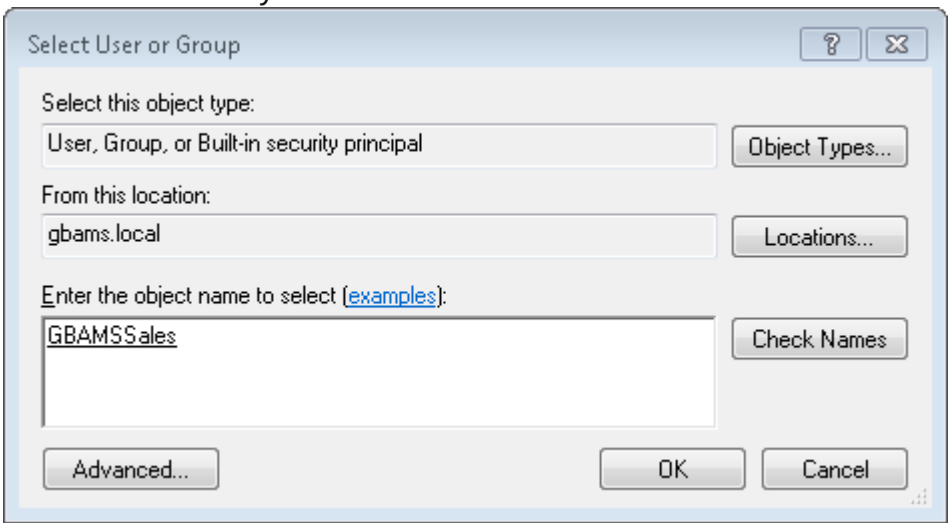

19. Once you click ok, you then need to make a decision if they are going to be server administrators or not. In this example, they are not server administrators. Click ok or apply.

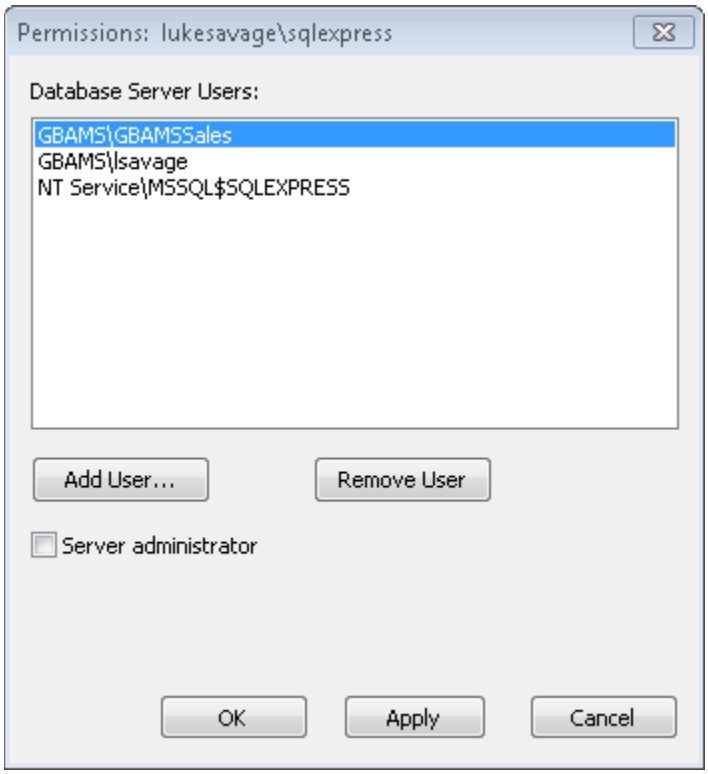

20. Next, you need to assign the group or user per Geodatabase if you have more than one. Because I separate Raster and Vector Geodatabases, the permissions change for users and groups depending on the intent for use. Right click on the new Geodatabase and select administration permissions in the drop down menu.

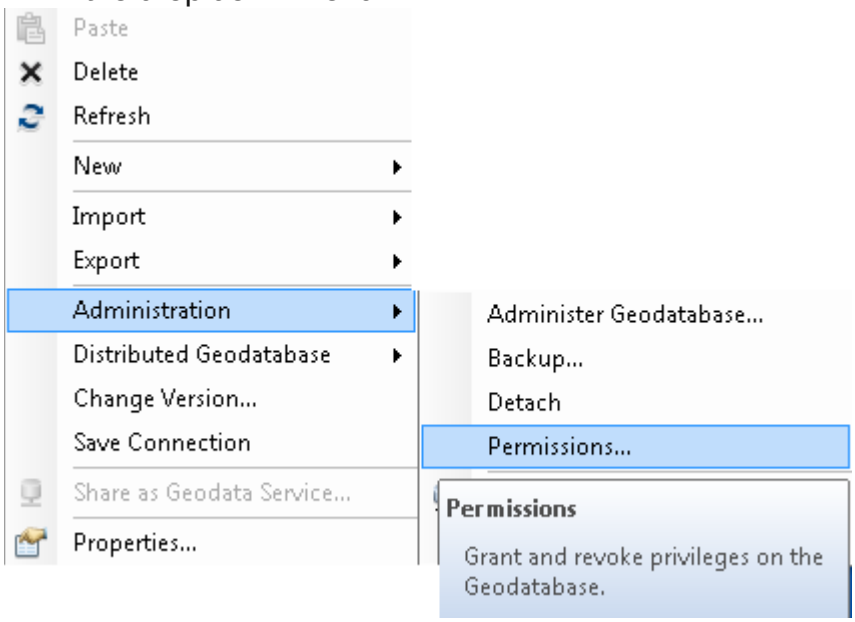

21. In the permissions: <database> dialog, select the group or user you would like to add permissions to and select the appropriate permission. In this example, I will be giving the GBAMSSales group read/write permissions. Click ok or apply. Notice the other options available for Workgroup for

example, backup, Administer Geodatabase, detach, save connection, compress, Geodatabase maintenance, and change versions. We will explore these tasks in subsequent sections.

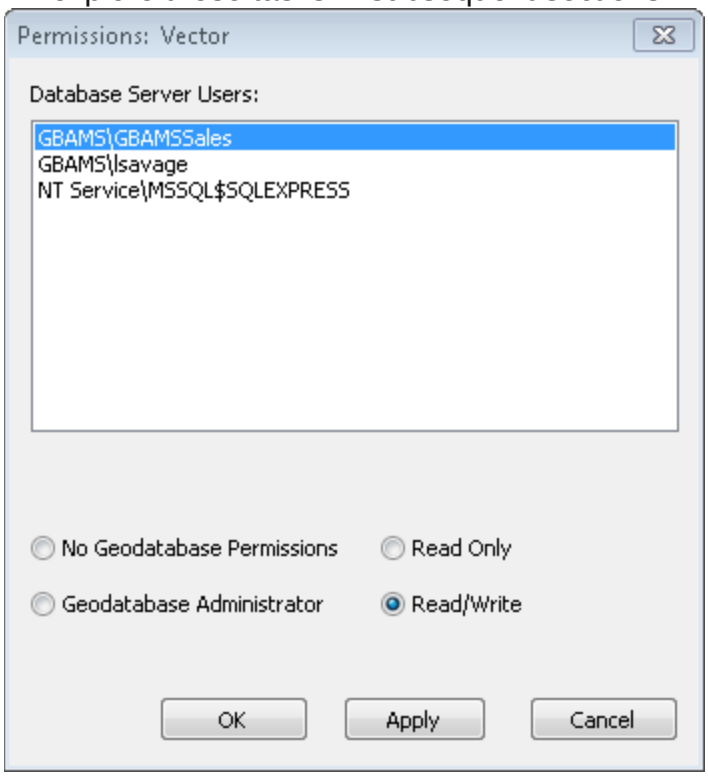

22. For your users who edit and publish maps, you will want them to create a connection to the database. You can add database connections using the database connection utility as well. This is a primary connection for anyone running full versions of SQL Server. Go to the TOC of ArcCatalog again and expand Database Connections.

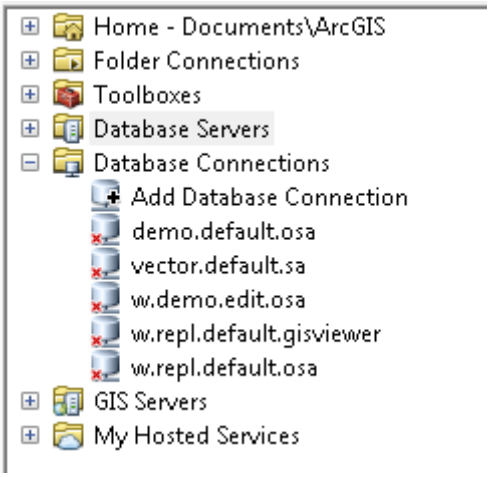

23. Click on 'Add Database Connection'. Add instance and the database to the connection strings. You have the choice of authenticating using the operating system or SQL Server user accounts. In this example, we are going to use the OSA method. Click ok.

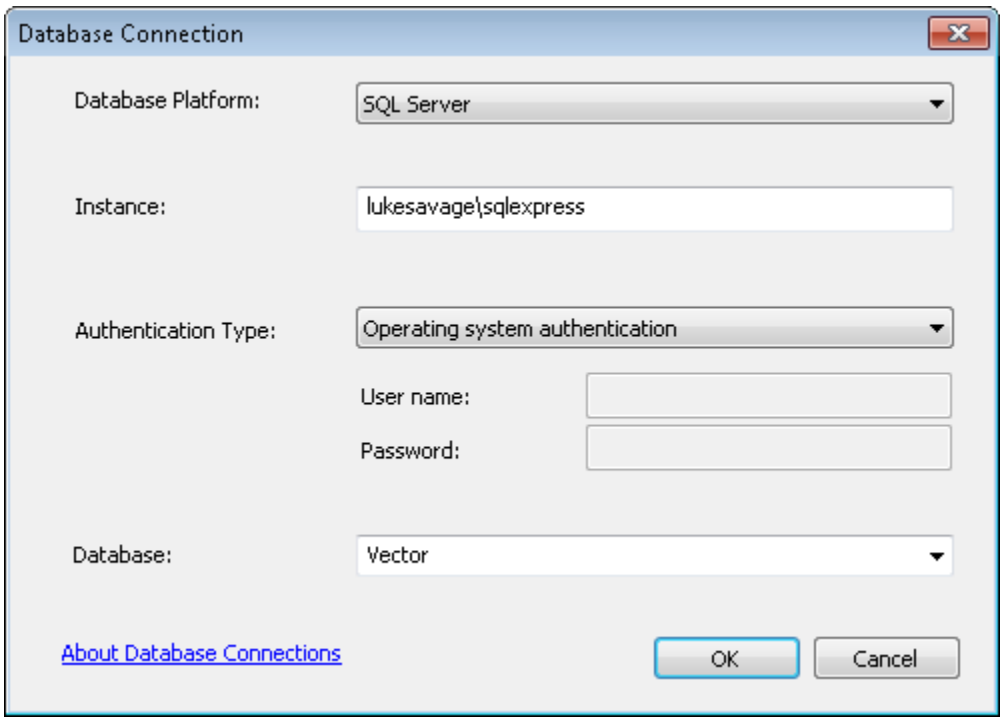

24. The connection string for the name of the connection is always arbitrary. Recommend using the following naming convention.
    - **<instance if more than one instance; if not skip>.<database>.<version>.<OSA or SQL user>**
    - For this example, we will be renaming the connection to vector.default.osa.
    - Right click on the new connection and select rename. Type in the new name.

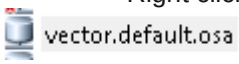

Registering the Geodatabase in ArcGIS for Server

25. In order for ArcGIS for Server to connect to your database, you need to add ArcGIS for Server user as read or read/write depending on your uses.
26. You can either grant permissions to ArcGIS for Server service user through the Geodatabase Administration toolset in ArcCatalog or you can add the user in SQL Server and then grant permissions of the data to the user. In this exercise, we'll go through and add ArcGIS for Server user to SQL Server. This exercise is beneficial for GIS workflow scenarios. We will be assigning ArcGIS for Server service user that has read/write access to the database with no DBO or other special privileges.
27. Go to SQL Server Management Studio and add the new user to your Geodatabase or Geodatabases. Open Security within the SSMS TOC and right click on logins. Select new login. Add ArcGIS for Server Service user to the Geodatabase. Recommend assign a default database to general groups or users to be other than master.

28. Go to User Mapping and select on the appropriate databases you want ArcGIS for Server service user to connect to. Make sure it has db_datareader or db_datawriter enabled depending on the functionality you want.

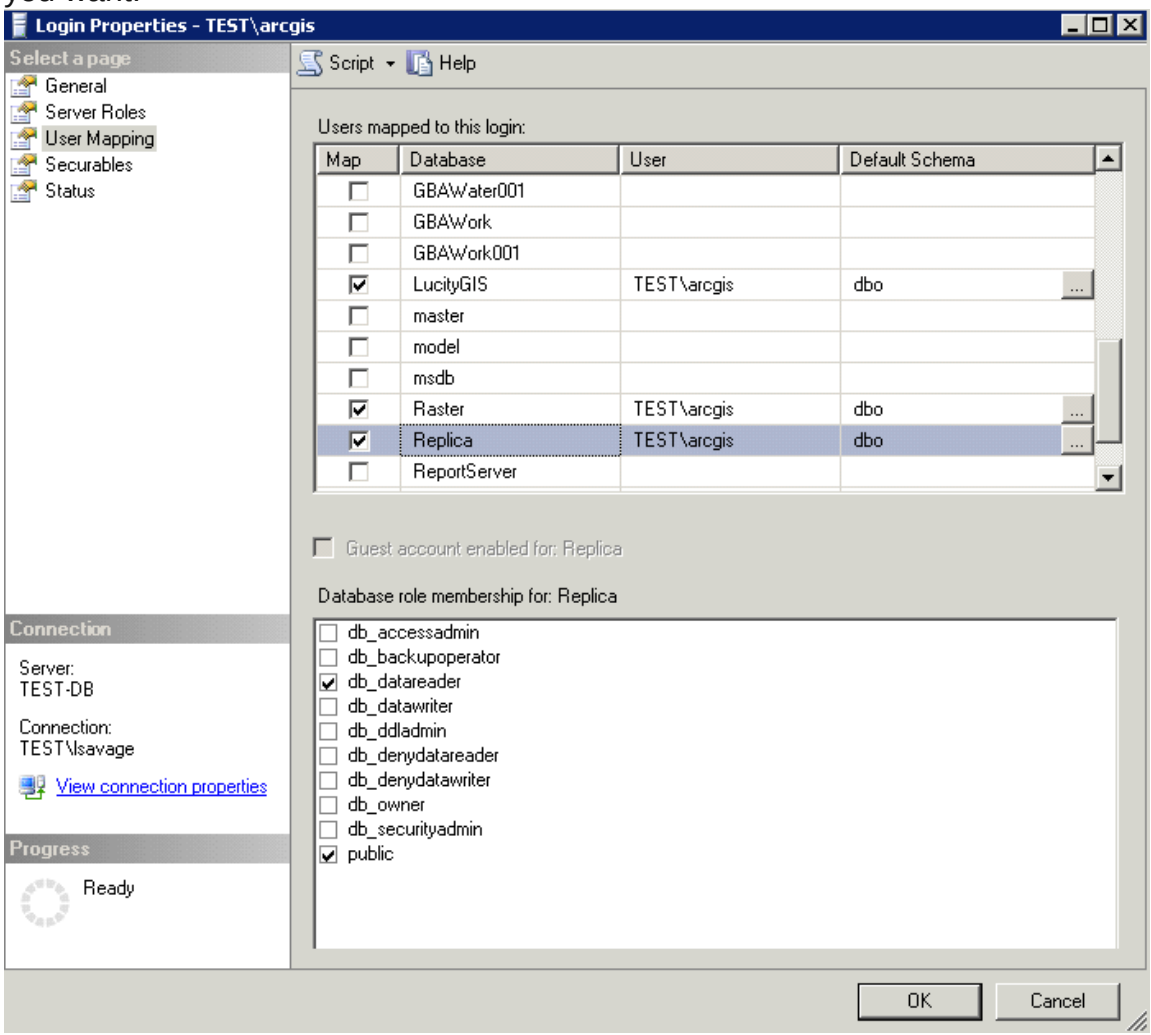

29. Click ok and go to your workstation and open ArcCatalog. Go to the GIS Servers section within the TOC of ArcCatalog. Connect to the administration connection by double clicking the connection or right click and choose connect.

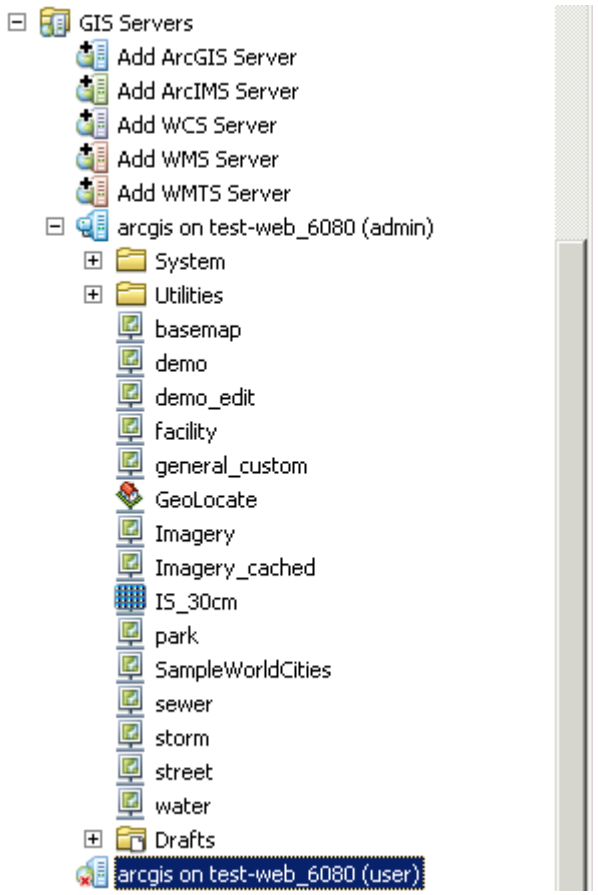

30. Once connected, right click on the connection (in this case arcgis on test-web_6080 (admin)) and choose server properties.

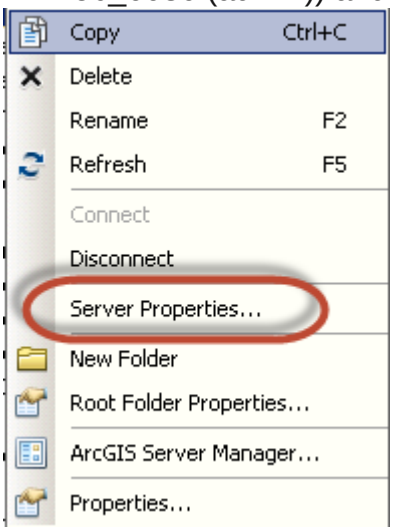

31. Go to the data store tab and click the plus button.

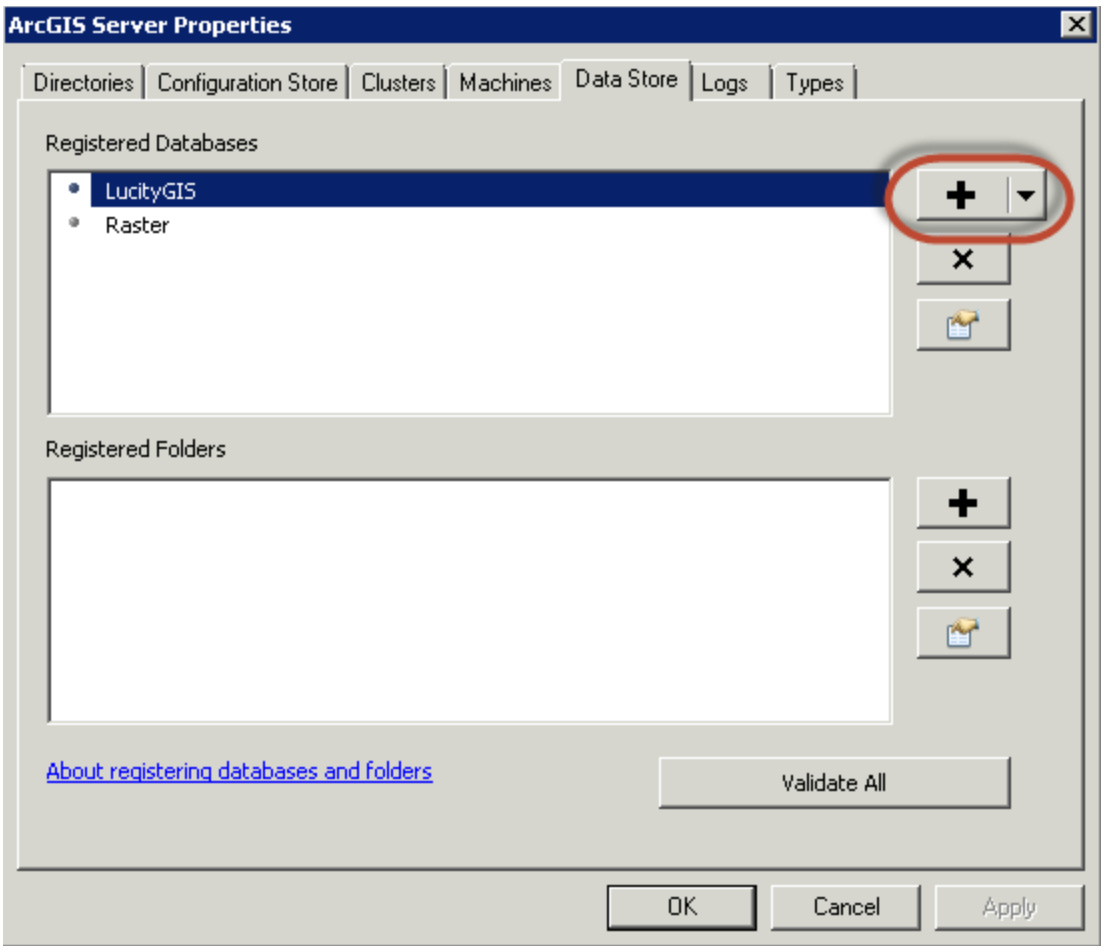

32. Notice I already have LucityGIS and Raster inside of my data store and now I want to connect to the Replica. Add the replica name and click add for publisher database connection.

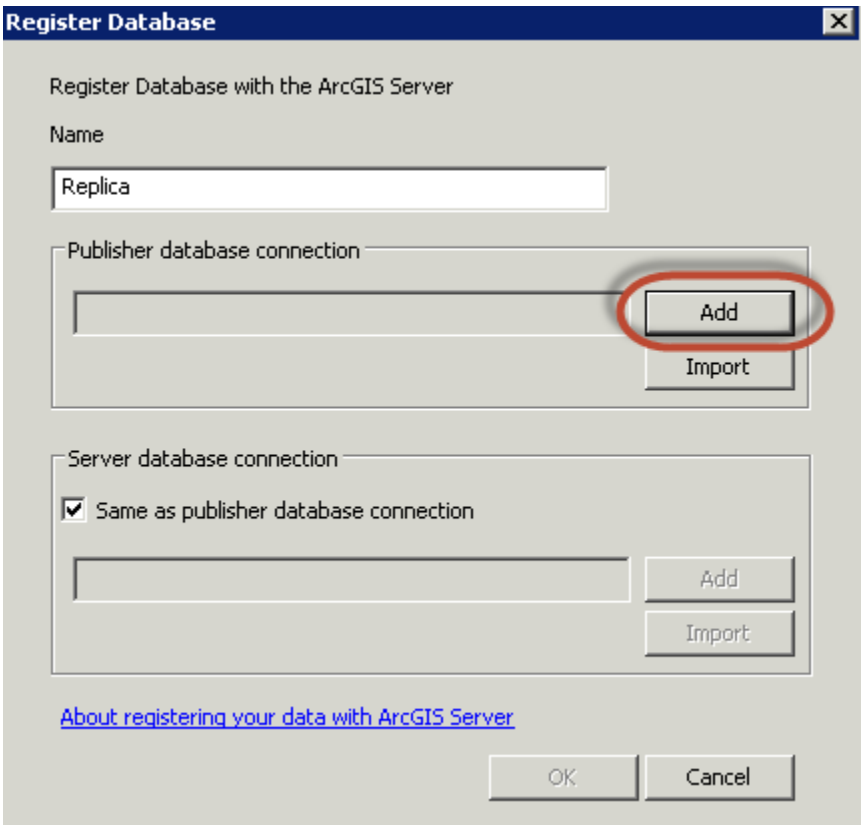

33. Add the appropriate information in instance and database. Click ok and ok.

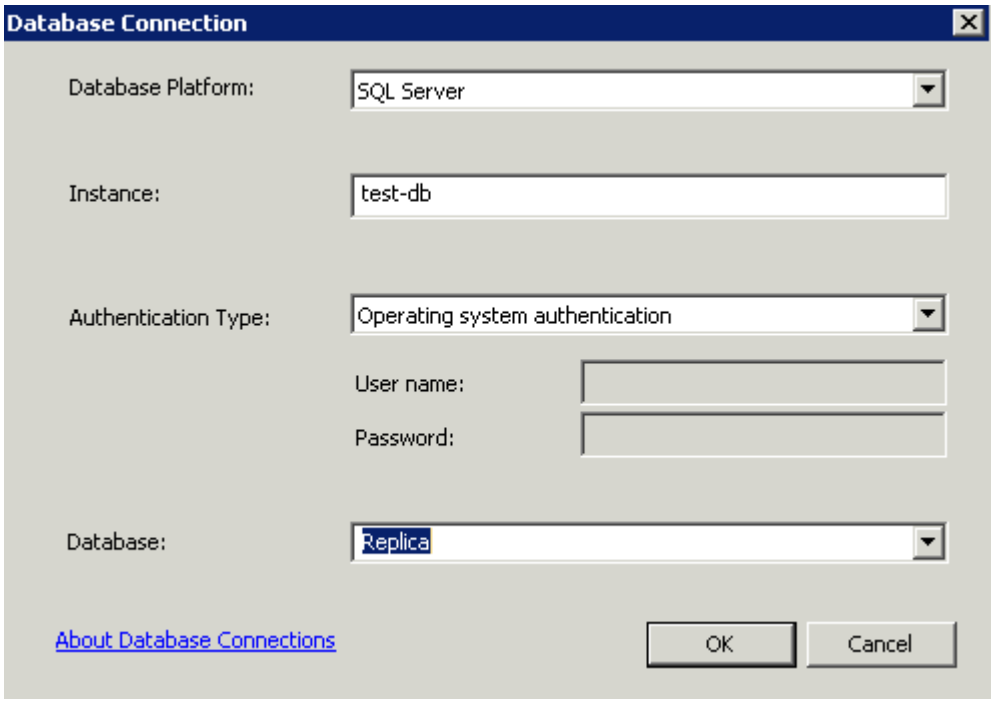

34. Your database should have a green check next to it to signify completed successfully.

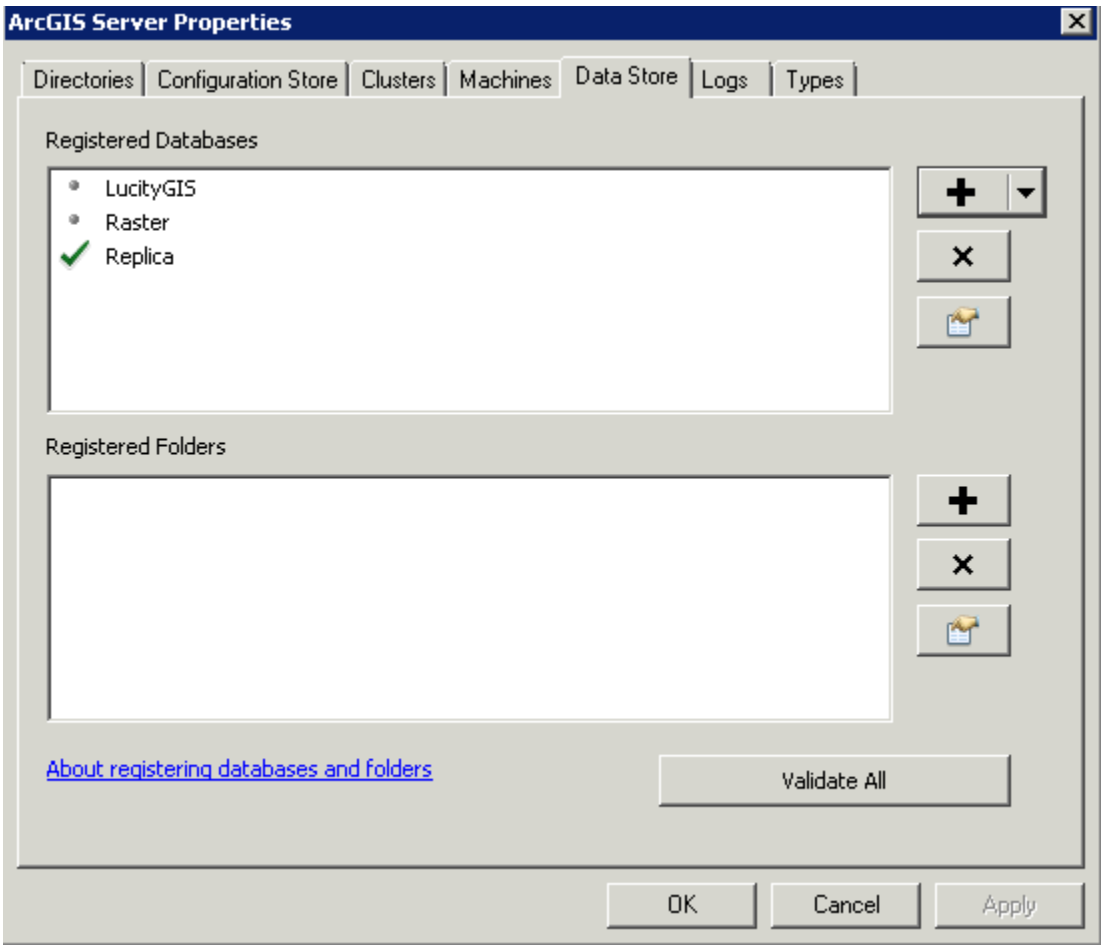

35. Now you are ready to publish map services from connections to the SQL Server Enterprise Geodatabase.

Create Enterprise Geodatabase (formerly ArcSDE)

1. Choose the database platform you will be connecting to. Since Lucity is SQL Server friendly, this will demonstrate the sql server platform type.
2. Within SQL Server, an instance is the container or the repository of the databases.
3. Database would be a name given to the database. In SQL Server's case, you will be specifying a name of the database. In this example, we will be creating a database called "vector".
4. Since I have mixed mode turned on for sql authentication, I chose to use the sa account. As the sysadmin, I could have created a database using my OSA account; however, for organizations that have changes, I would recommend creating the database using the SA as this will tag the database to be owned by sa. For most cases, a mixed mode sql instance is the best practice. This gives your organization the flexibility of using both the sa account and OSA when managing databases. **If a DBA is let go and they have created databases using their credentials and the SQL Server instance is set to only use OSA credentials, you will not like the outcome.**
5. Uncheck sde owned schema as this was the old way ESRI used to use for the ArcSDE setup and configuration. DBO schemas are much preferred as they have the most flexibility of use within a windows environment.
6. Authorization file is the file created after authenticating ArcGIS for Server (i.e. c:\program files\ESRI\ License10.1\sysgen\keycodes)

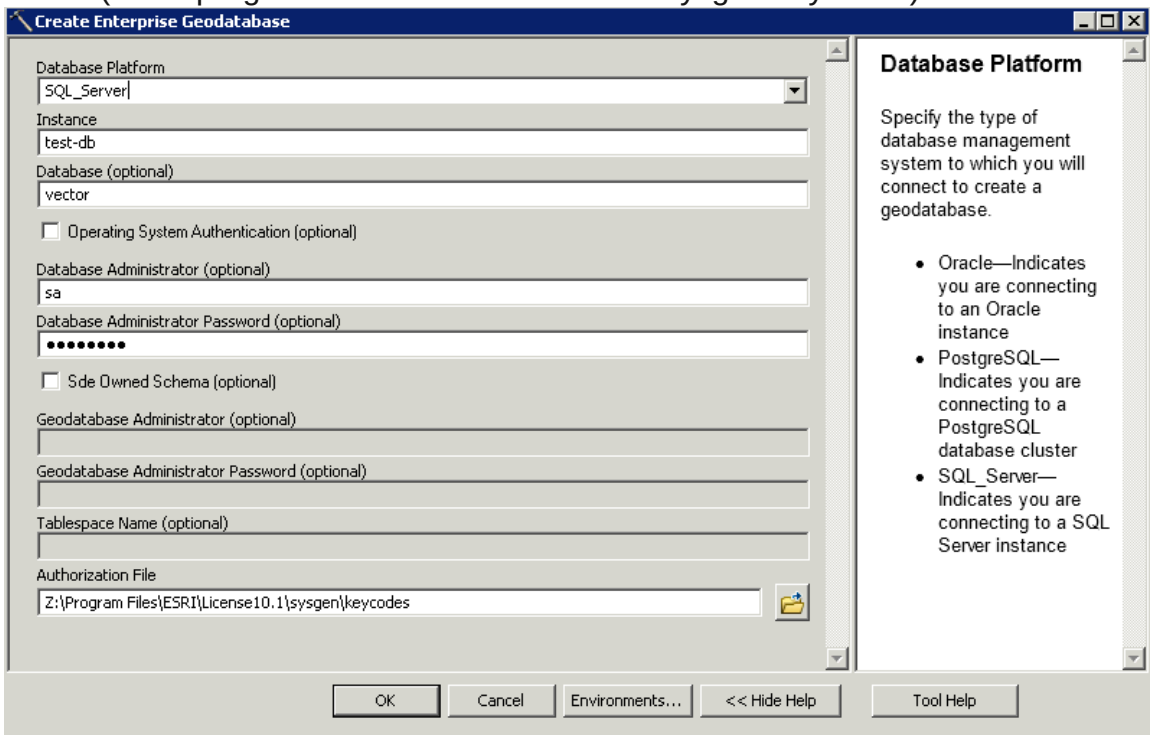

Publishing and Configuring Map Services

1. Open ArcMap and locate your .mxd to publish to ArcGIS Server
    - Recommend: If you have an image, recommend separating the image from the vector map services for performance reasons.
    - Recommend: Cached Image Service
    - Recommend: Converting dynamic labels to Annotation and scale.
    - Recommend: Make your map services scale dependent so not everything shows up in a small scale extent.
    - Recommend: Use ArcGIS Server Web Adapter if using a DMZ. Please consult your IT Administrator for setting up a DMZ.
2. When cartographic and scaling operations are complete, go to the file menu and click on 'Share As/Service'.

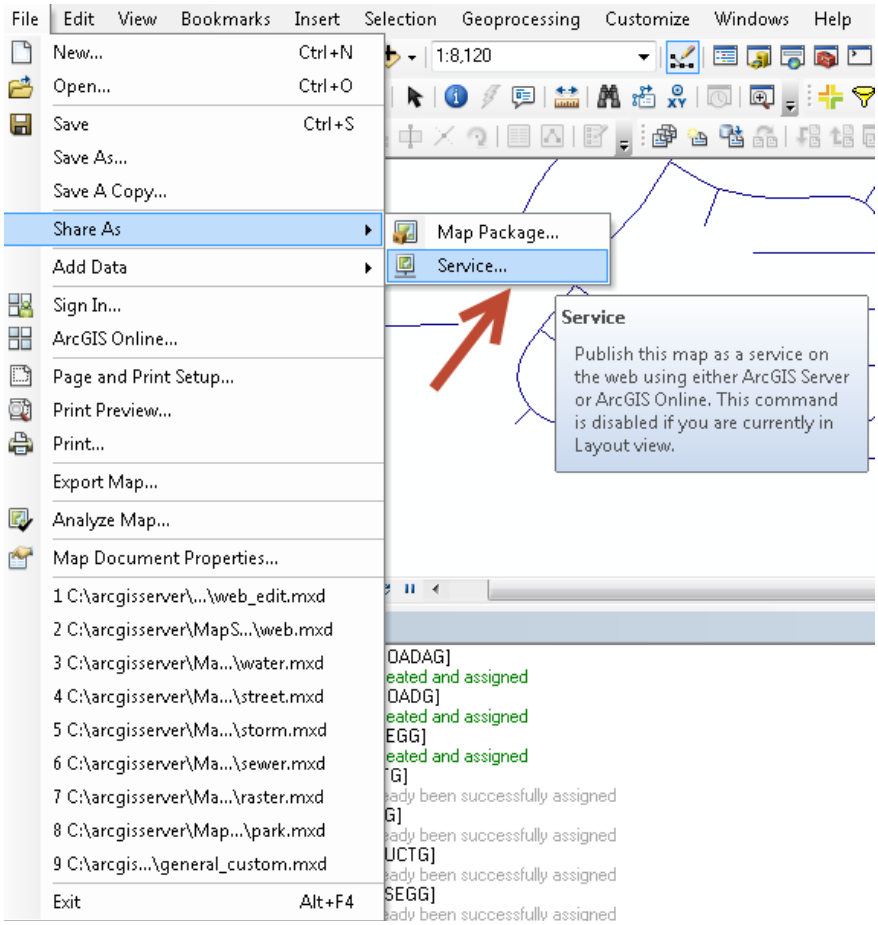

3. Share as Service dialog prompts. There are three choices. Choose Publish a service and click next.

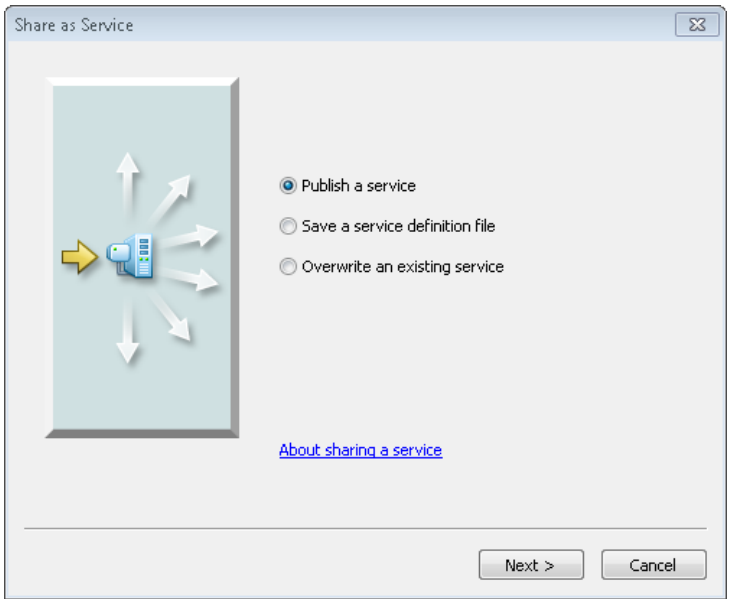

4. Publish a Service dialog appears and give it a name. Click Next.

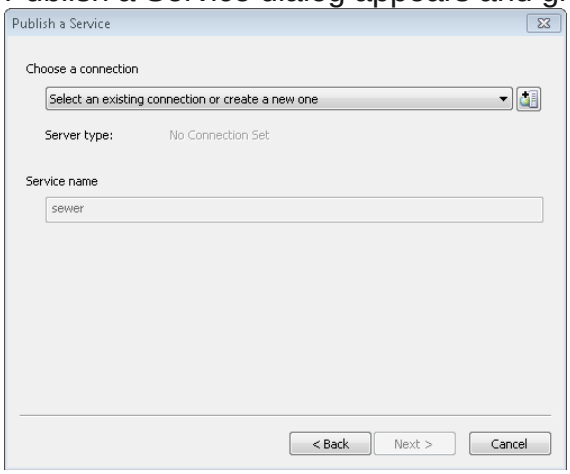

5. Service Editor dialog will appear and this is where you add your settings for your map service.

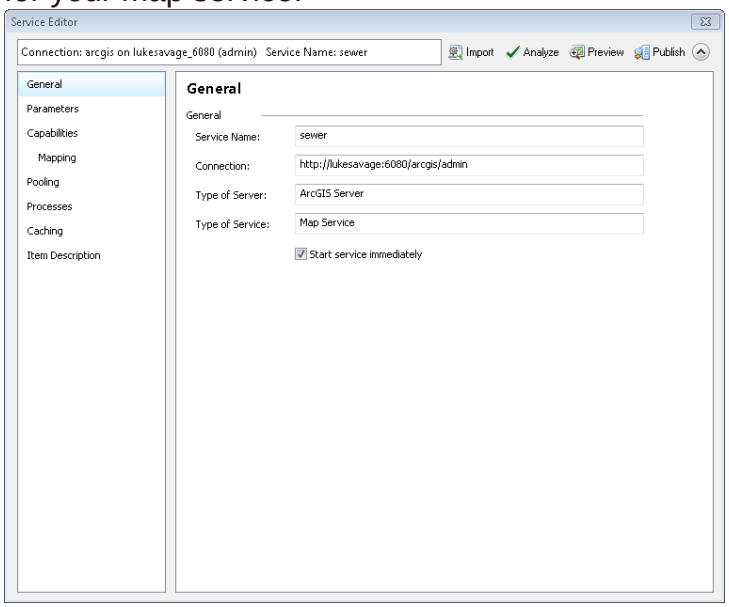

6. Click on Analyze at the top of the dialog. After the analyze tool has completed, it will list errors, warnings and messages letting you know if there are any issues with the performance of the data being rendered to ArcGIS Server.

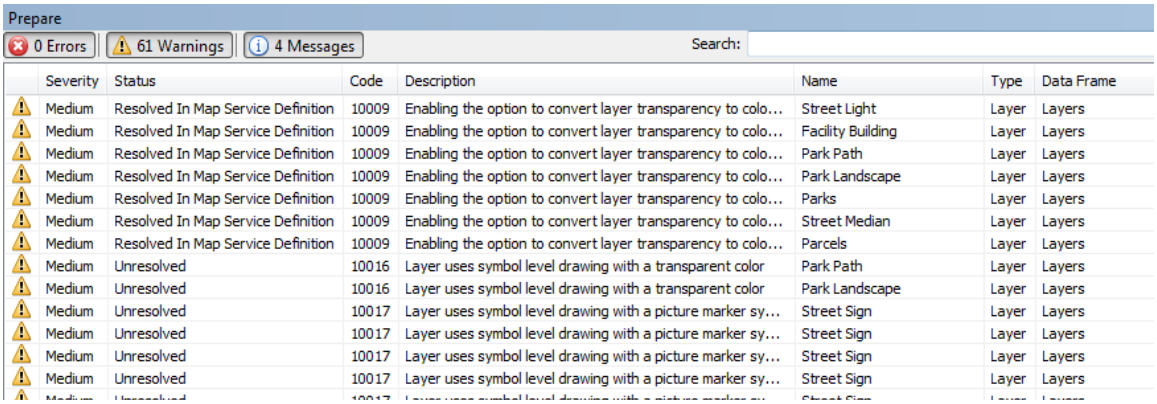

	Severity	Status	Code	Description	Name	Type	Data Frame
⚠	Medium	Resolved In Map Service Definition	10009	Enabling the option to convert layer transparency to colo…	Street Light	Layer	Layers
⚠	Medium	Resolved In Map Service Definition	10009	Enabling the option to convert layer transparency to colo…	Facility Building	Layer	Layers
⚠	Medium	Resolved In Map Service Definition	10009	Enabling the option to convert layer transparency to colo…	Park Path	Layer	Layers
⚠	Medium	Resolved In Map Service Definition	10009	Enabling the option to convert layer transparency to colo…	Park Landscape	Layer	Layers
⚠	Medium	Resolved In Map Service Definition	10009	Enabling the option to convert layer transparency to colo…	Parks	Layer	Layers
⚠	Medium	Resolved In Map Service Definition	10009	Enabling the option to convert layer transparency to colo…	Street Median	Layer	Layers
⚠	Medium	Resolved In Map Service Definition	10009	Enabling the option to convert layer transparency to colo…	Parcels	Layer	Layers
⚠	Medium	Unresolved	10016	Layer uses symbol level drawing with a transparent color	Park Path	Layer	Layers
⚠	Medium	Unresolved	10016	Layer uses symbol level drawing with a transparent color	Park Landscape	Layer	Layers
⚠	Medium	Unresolved	10017	Layer uses symbol level drawing with a picture marker sy…	Street Sign	Layer	Layers
⚠	Medium	Unresolved	10017	Layer uses symbol level drawing with a picture marker sy…	Street Sign	Layer	Layers
⚠	Medium	Unresolved	10017	Layer uses symbol level drawing with a picture marker sy…	Street Sign	Layer	Layers
⚠	Medium	Unresolved	10017	Layer uses symbol level drawing with a picture marker sy…	Street Sign	Layer	Layers

- If errors, you need to resolve before proceeding.
- If warnings, you may or may not address these issues but ESRI highly recommends fixing these issues.
- Messages are helpful hints on low priority performance issues. Most of the time, these do not need to be addressed as issues as they are intended.

7. Go to parameters in the TOC for the Service Editor dialog and change the anti-aliasing to Normal. Not all map services will show artifacts but if so normal is a good place to start.

   Note: *Anti-Aliasing: A graphics technique that blends foreground and background pixels near edges of objects to trick your eye into seeing smoother borders. You can use this option if unwanted artifacts appear in your map displays, for example, jagged lines, wavy lines or bands, and moiré patterns.*

Use the following options to get the graphics effect and performance that will meet your map service needs:

- *None: No antialiasing is performed.*
- *Fastest: Minimal antialiasing is performed, optimized for speed.*
- *Fast: Some antialiasing is performed, optimized for speed with better quality than can be achieved with Fastest.*
- *Normal: A good balance of speed and quality.*
- *Best: The best quality antialiasing. This option takes the longest to render.*

- Force text anti-aliasing should be enabled by default.

8. Change the maximum number of records returned by the server. This will increase the records returned by one query. 3000 is a good number to return unless your users need more.

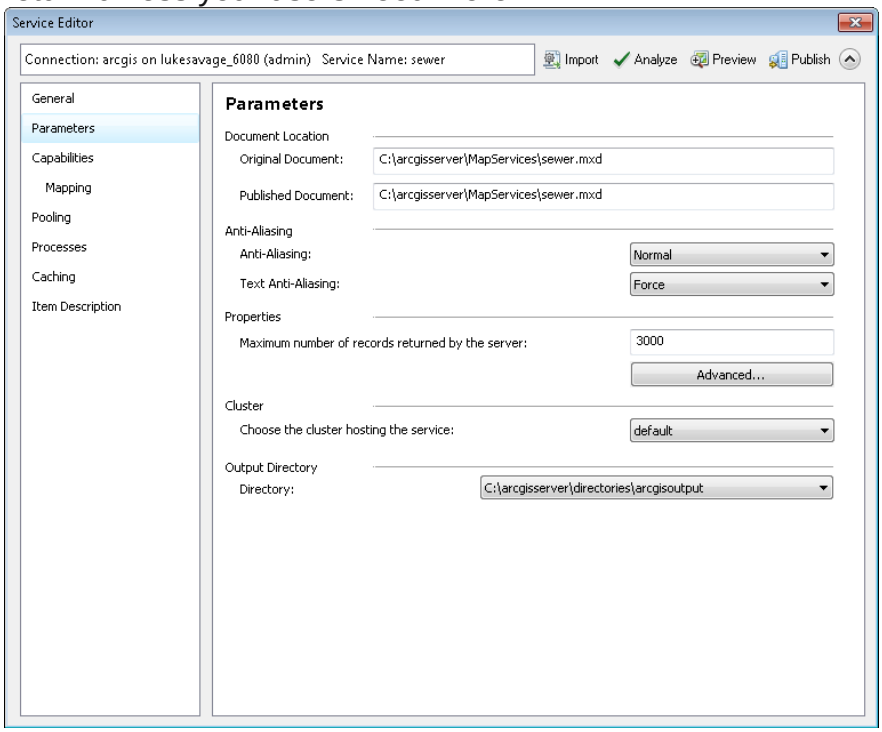

Unlock Schema

7. Fixed in ArcGIS 10.1 SP1 from bug number NIM082716, setting SchemaLockingEnabled to false in Service Properties of ArcGIS for Server 10.1 SP1 disables schema locking. For those who wish to serve their production data as a map service, this is extremely helpful if you need to make any kind of schema changes in your database.
8. Within the service editor dialog when publishing a map service, go to Parameters within the menu tabs on the left and click on advanced.

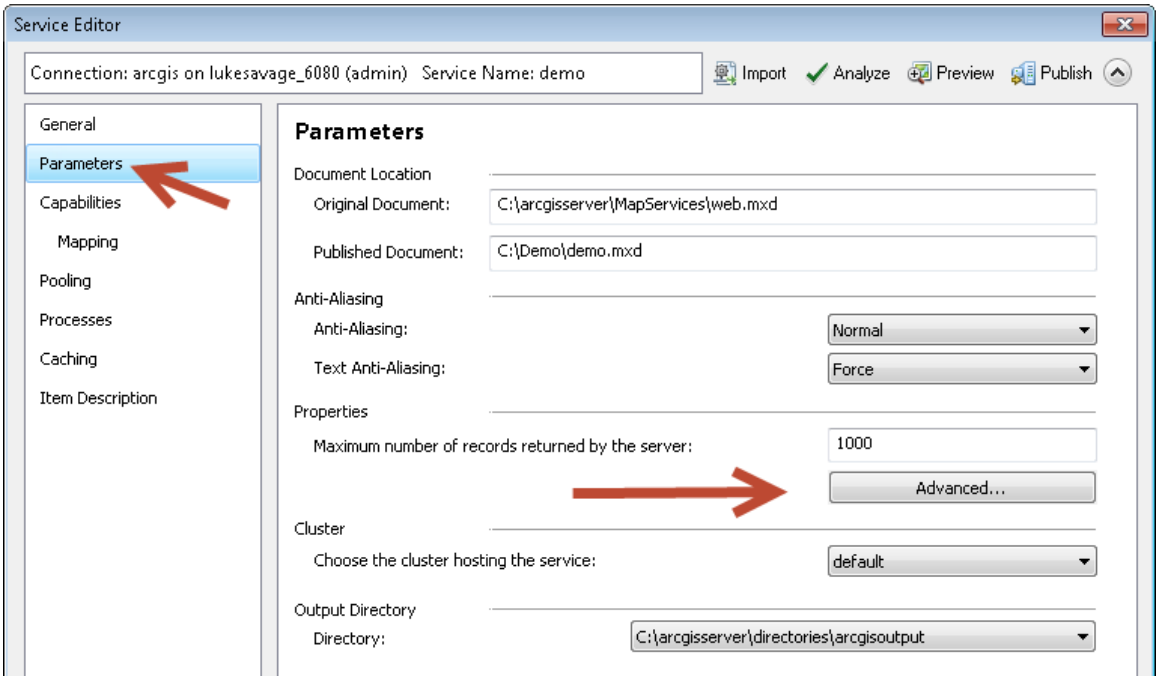

9. Click on the ‘Advanced’ button and change the values for schemaLockingEnabled to false if you don’t want the map service to lock your database. Click ok when finished.

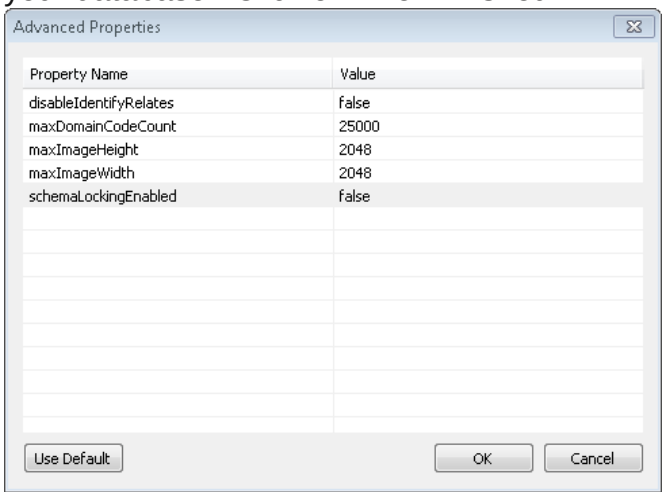

10. Go to the Capabilities tab and make sure the Mapping service is the only box checked.

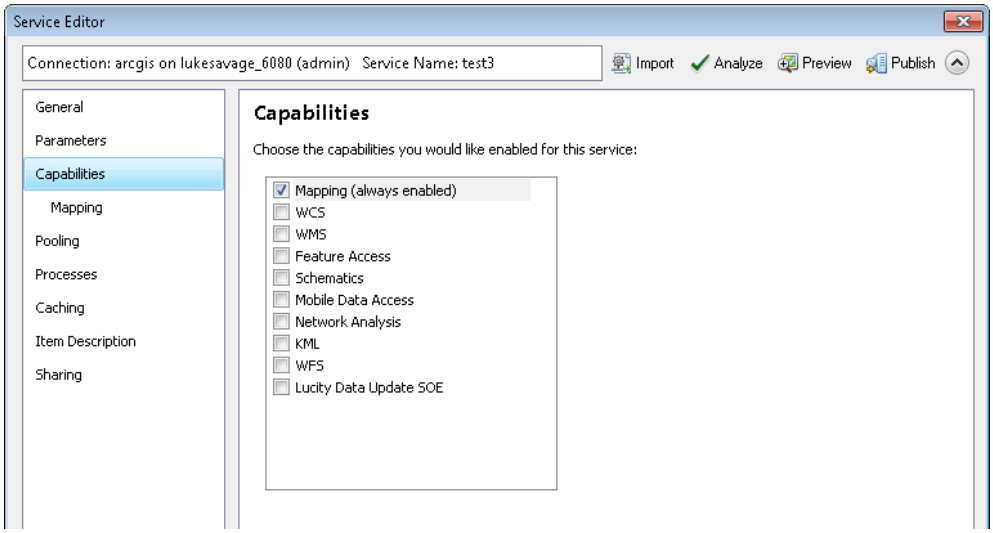

11. In pooling, usually 2 minimum and 5 maximum will suffice for most organizations. For large organizations, increase based on user load.

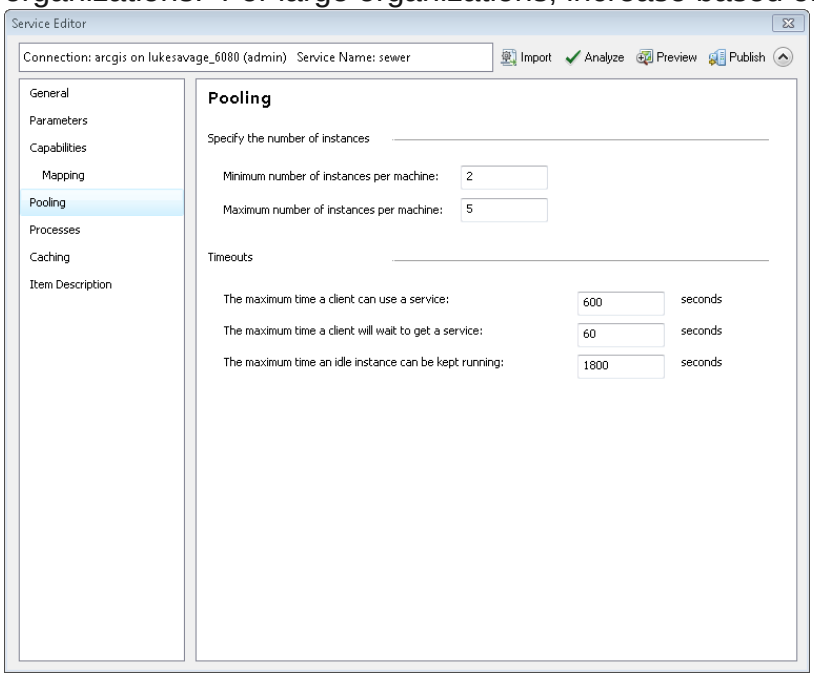

12. In Processes, go to recycle this configuration every: <number> and change to a desirable value. Usually, 12 hours is a good recycled time frame.
13. Change the starting at value to 6:00 am or whenever you would like it to start

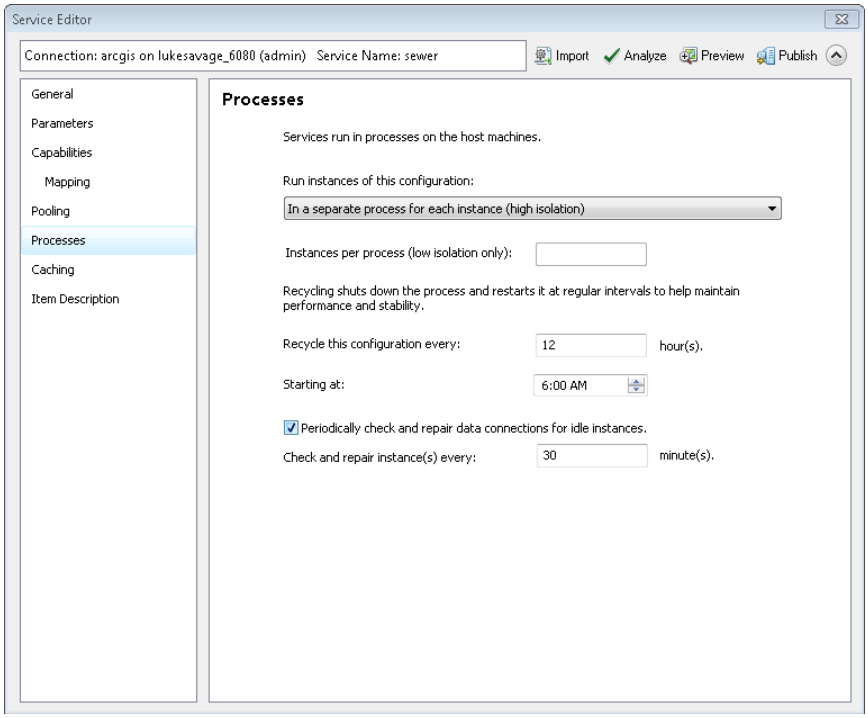

14. If you want to cache, click on caching and change the values to a suitable number. For basemaps and imagery, recommend caching. For interactive map services, leave caching turned off which is by default.
15. If you want to see the speed of your map service, click on the preview icon in the menu toolbar. Otherwise, click on publish.

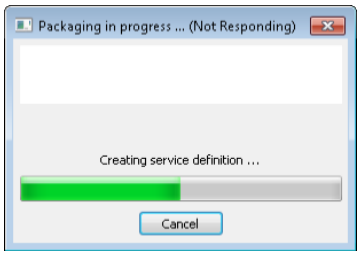

16. Below are Lucity map service deployment scenarios.
    - You may have multiple map services for each group if needed. It really depends on your setup with replication, versioning and how complex this may be. Each organization may be different on how they want to handle distributing maps to the users.
    - Create Basemaps and Imagery services (tiled) for distribution to your users. Consider your users and what content they would like to see in Basemaps.
    - Make sure you're basemap services are using the same coordinate systems (prefer wkid) in the publishing map document. This is important when going to and from basemap services so that it can honor the zoomed extent. For ArcGIS Online services and mashing up with local services, consider Mercator.
    - Make sure operational layers are using wkid for the coordinate system in your map document.

Lucity Data Update SOE Configuration

This server object extension to ArcGIS for Server pushes edits from the Lucity Web and Desktop to the Geodatabase when users edit attributes that are linked.

1. Open ArcMap
2. Right click on Layers in the table of contents and make sure the coordinate system is using the same spatial reference as the operational layers in your map services. In this example, we are using WKID 2868.

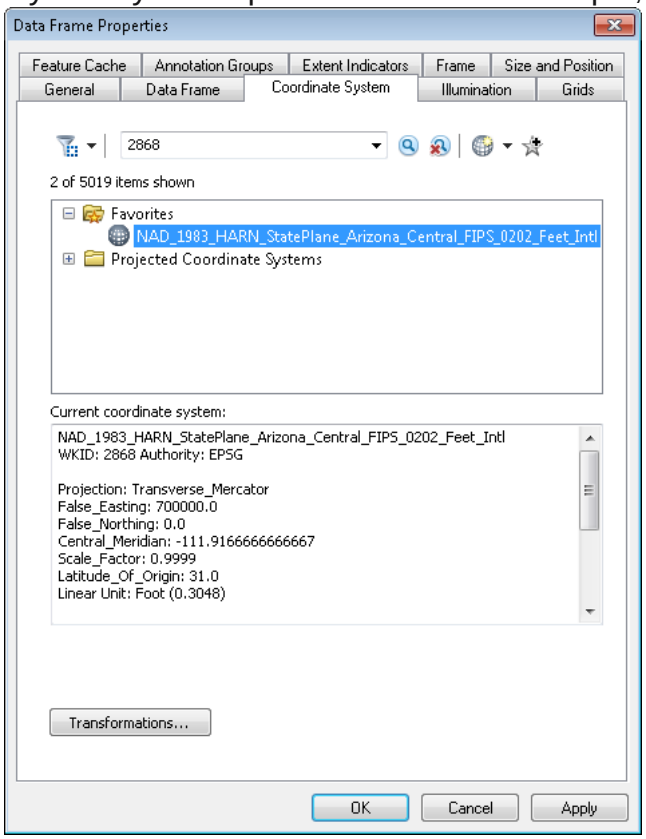

3. Add feature class that is mapped to Lucity. If multiple GIS connections that are synchronizing to Lucity, create a separate map service for each connection with only one feature class per database. The connection to the database must use a user with editing rights.

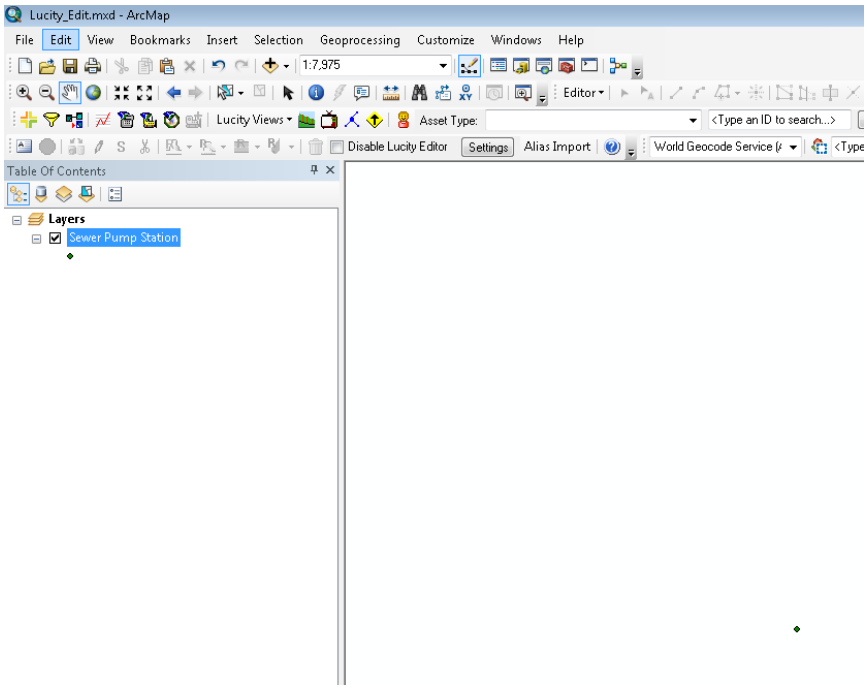

4. Go to the File/Share As/Service menu.
5. Publish a service

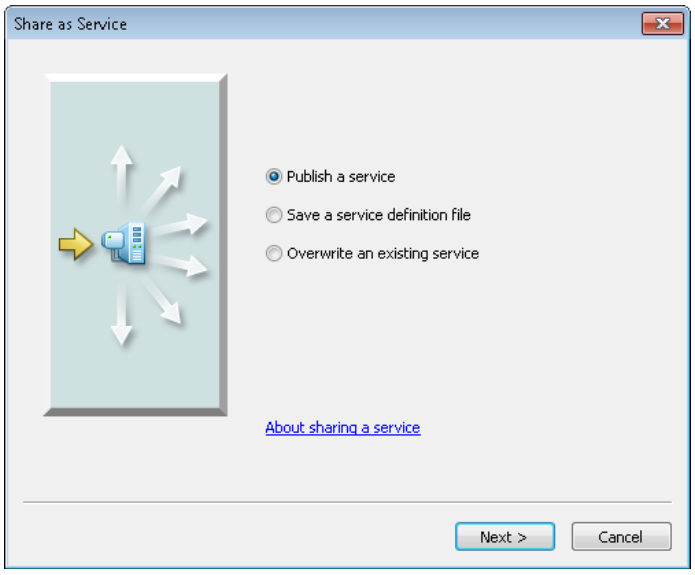

6. Click Next
7. Create a service name. Remember, the name used is case sensitive in the rest service URL

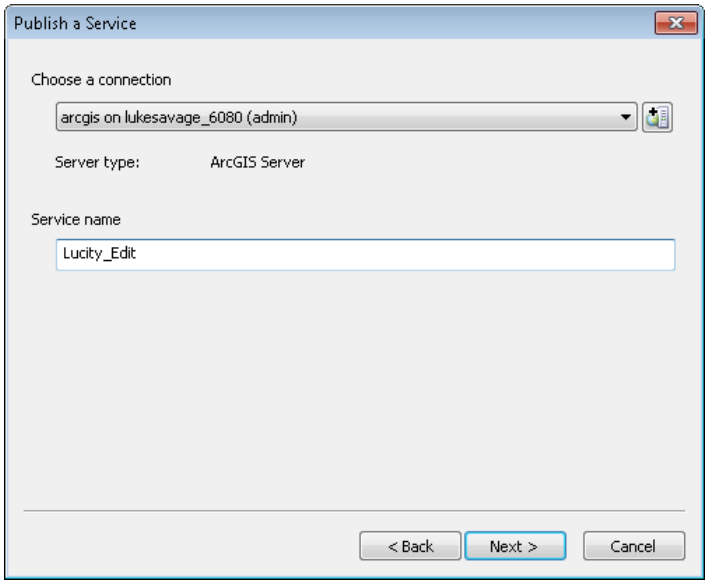

8. Click Next
9. Publish a Service in the root folder or a designated folder.

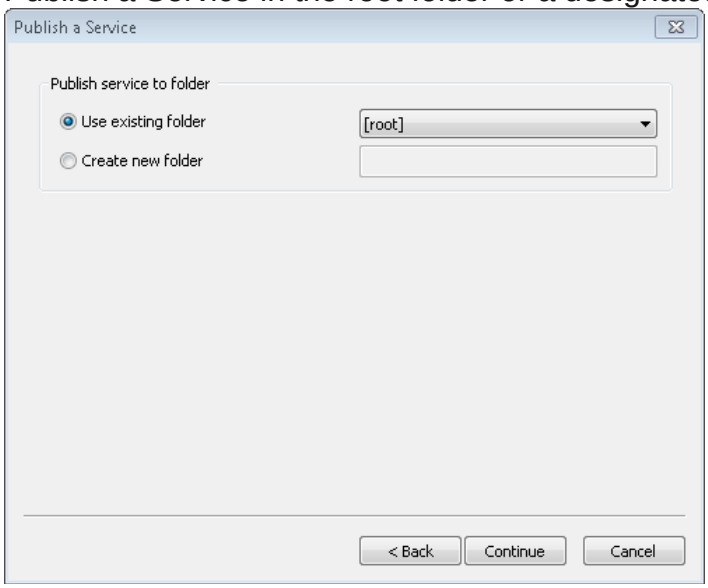

10. Click Continue
11. In the Service Editor tabbed menu/Table of Contents, click on parameters.
12. Under properties, click on the advanced button

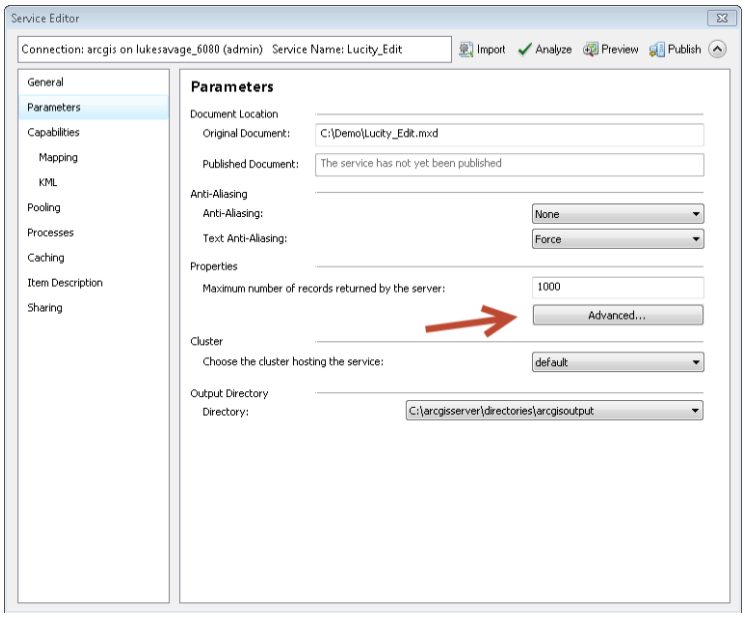

13. Change the values for schemaLockingEnabled to ‘false’ if you don’t want the map service to lock your database. Click ok when finished.

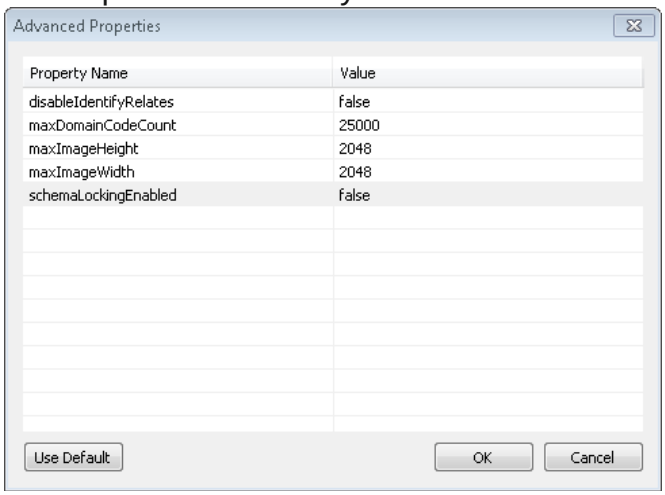

14. Go to the Capabilities tab and make sure Mapping and the Lucity Data Update SOE capability is checked.

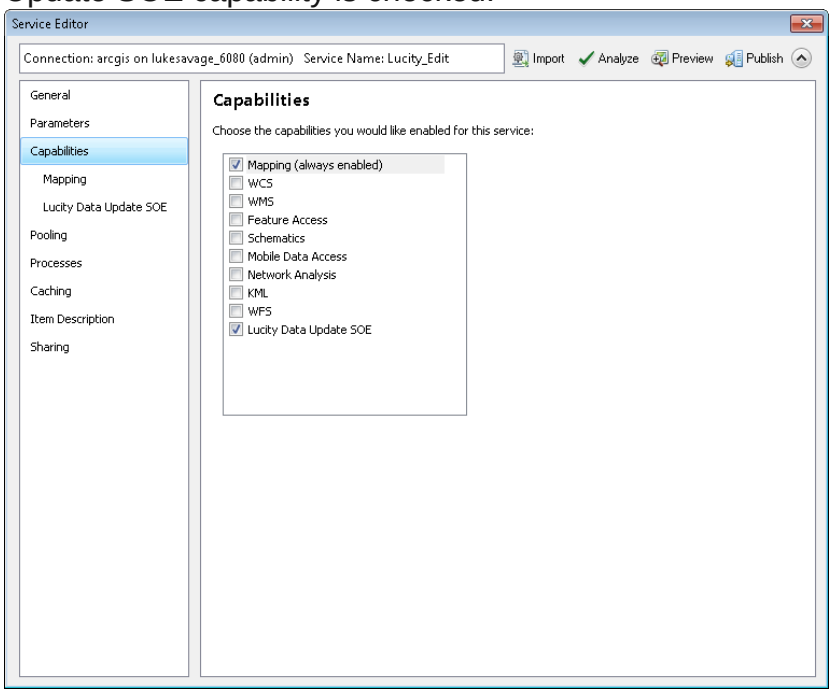

15. In pooling, usually 1 minimum and 3 maximum will suffice for most organizations. For large organizations, increase based on user load but be careful of the number of pools per CPU core.

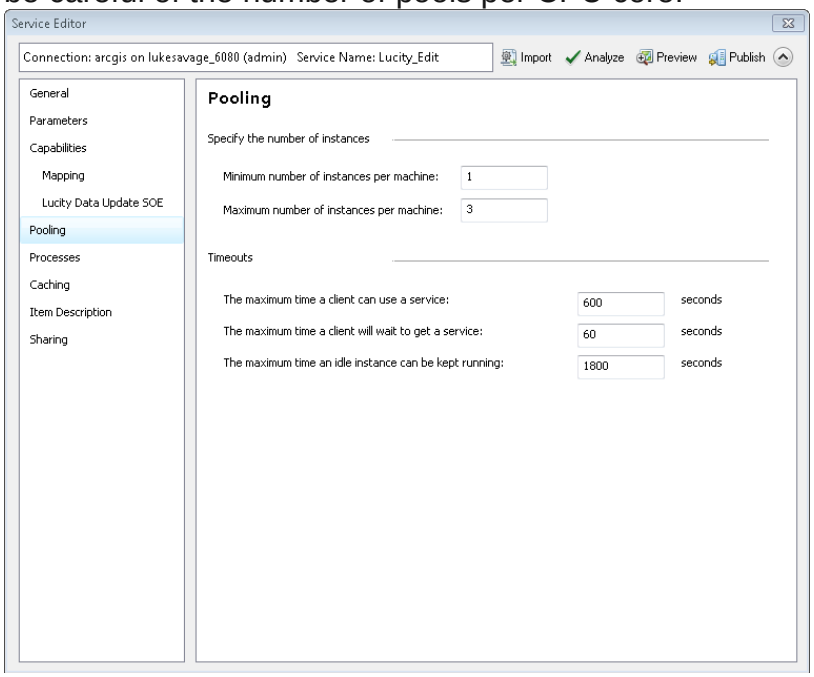

16. In Processes, go to recycle this configuration every: <number> and change to a desirable value. Usually, 12 hours is a good recycled time frame.

17. Change the starting at value to 6:00 am or whenever you would like it to start

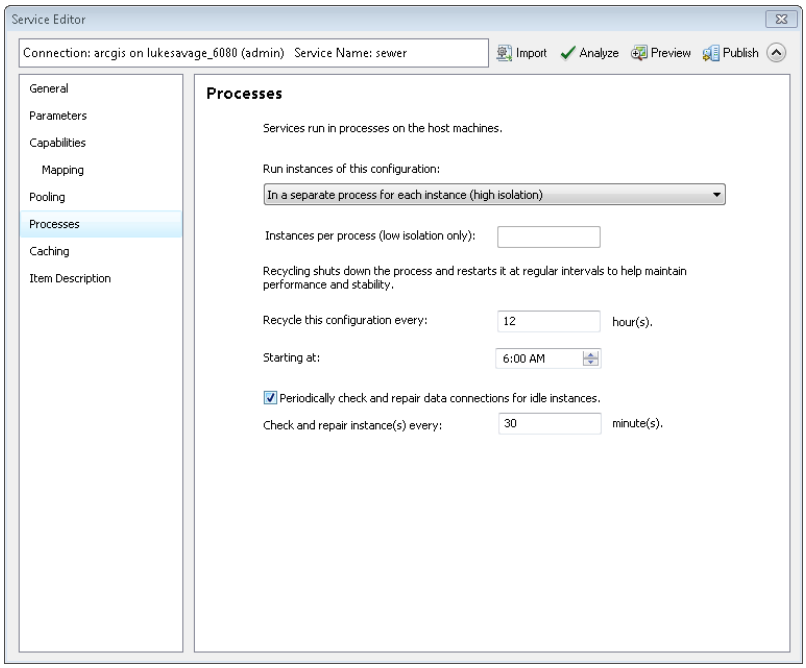

18. Click on publish.

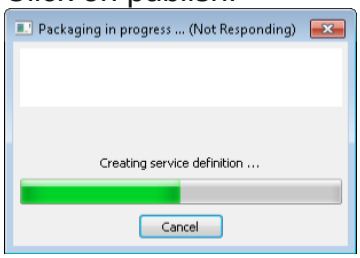

19. Below are definitions for the server object extension for Lucity and deployment scenarios.
    - **Lucity Data Update SOE**. This server object extension is used to push edits from the Lucity Web and Desktop applications to the Geodatabase when users edit relevant attributes in the web and desktop. You will only need one map service using the Lucity to Geodatabase update SOE capability per Geodatabase.
    - For production environments, we suggest creating a dedicated service for the Lucity to Geodatabase Update SOE so it is not shared with services used for map display. Only use one feature (if multiple databases, use one map service with one feature per Geodatabase; there will be two map services created, one for each geodatabase; each map service will have one feature class for performance reasons.

Lucity Data Update SOE Logic

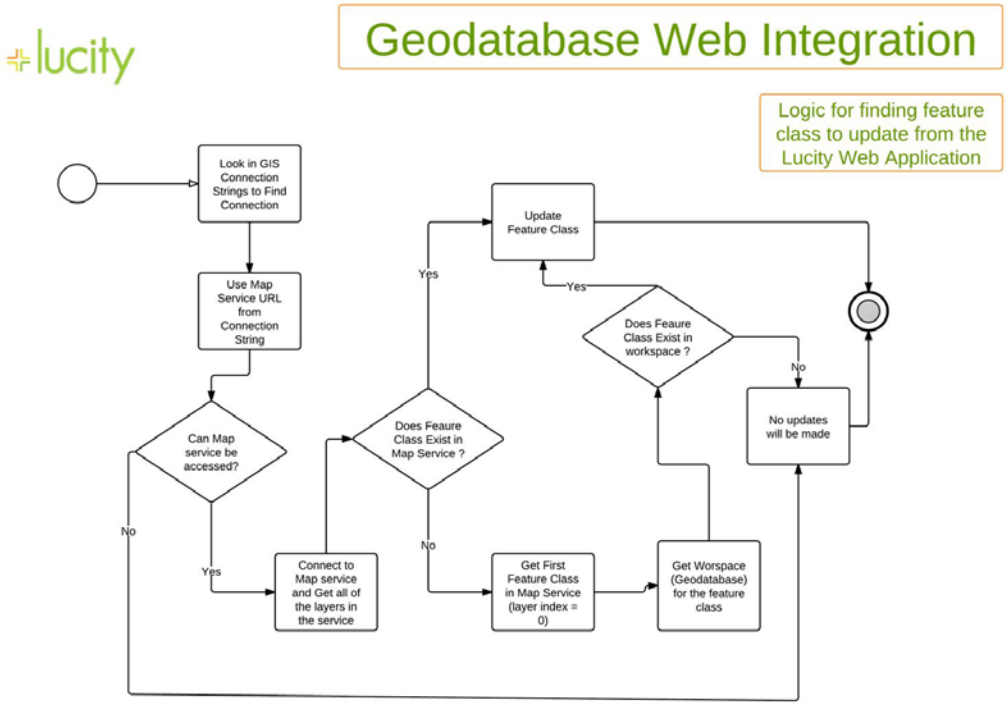

How to Create a Feature Service

Preparing Map for a Feature Service

New in Lucity 7.60, we have created a new redlining tool in Lucity WebMap. To configure the Geodatabase and ArcGIS for Server to enable this new capability, we must create a feature class or feature classes depending on the type of redlining capability you desire. We support points, lines and polygons. Below we will demonstrate the creation of a feature dataset, three feature classes and published them to a feature service.

1. Open ArcMap
2. Click on the Catalog button if it is not already open.

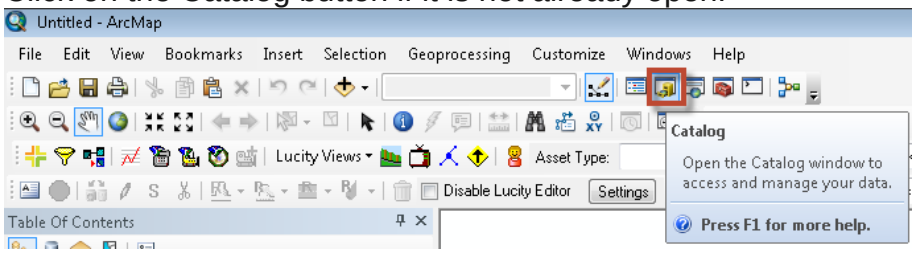

3. In the Catalog Tree, go to database connections and open your production GIS database.

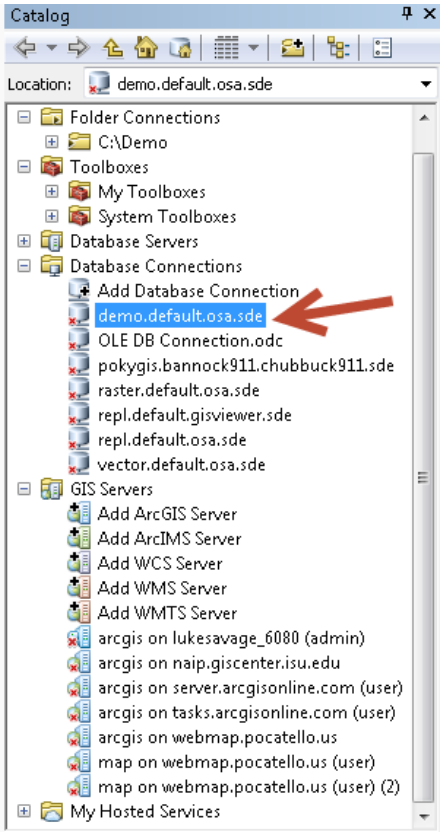

4. Right click on the connection and choose new and a submenu will appear.
5. Click on Feature Dataset

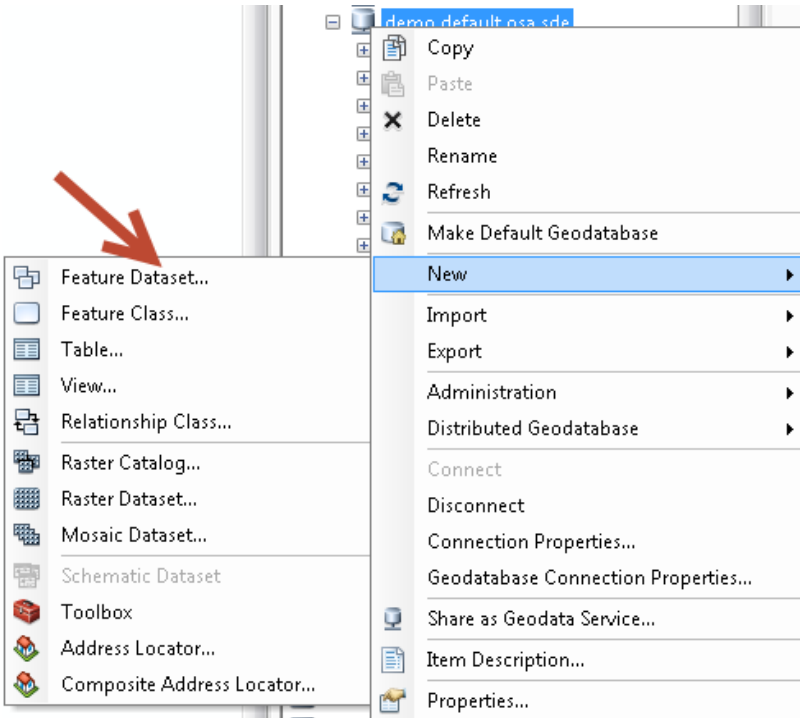

6. Give the Feature Dataset a name. In this example we will call it Redlining

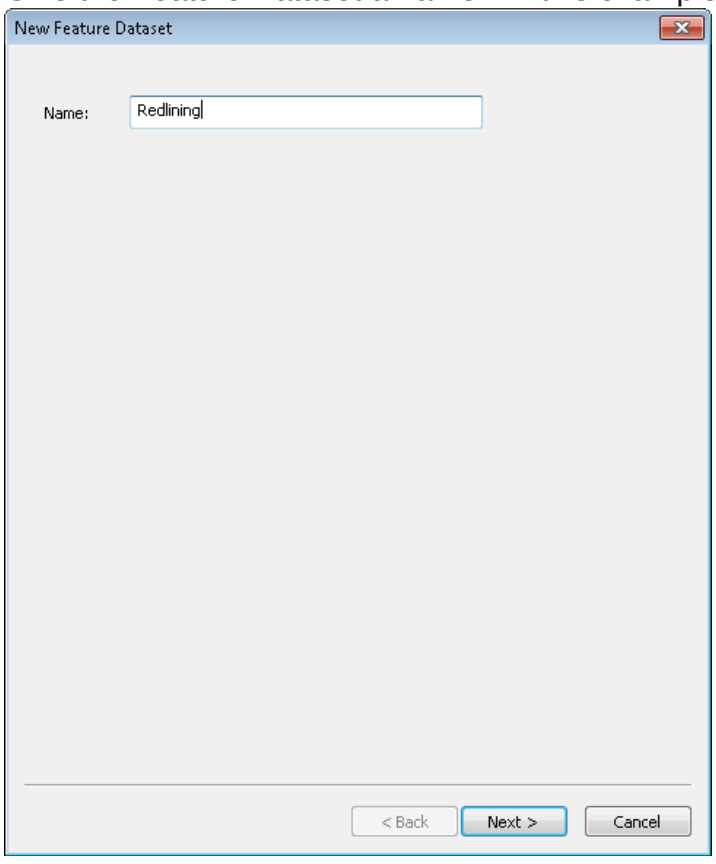

7. Click Next
8. Set your coordinate system for the feature dataset

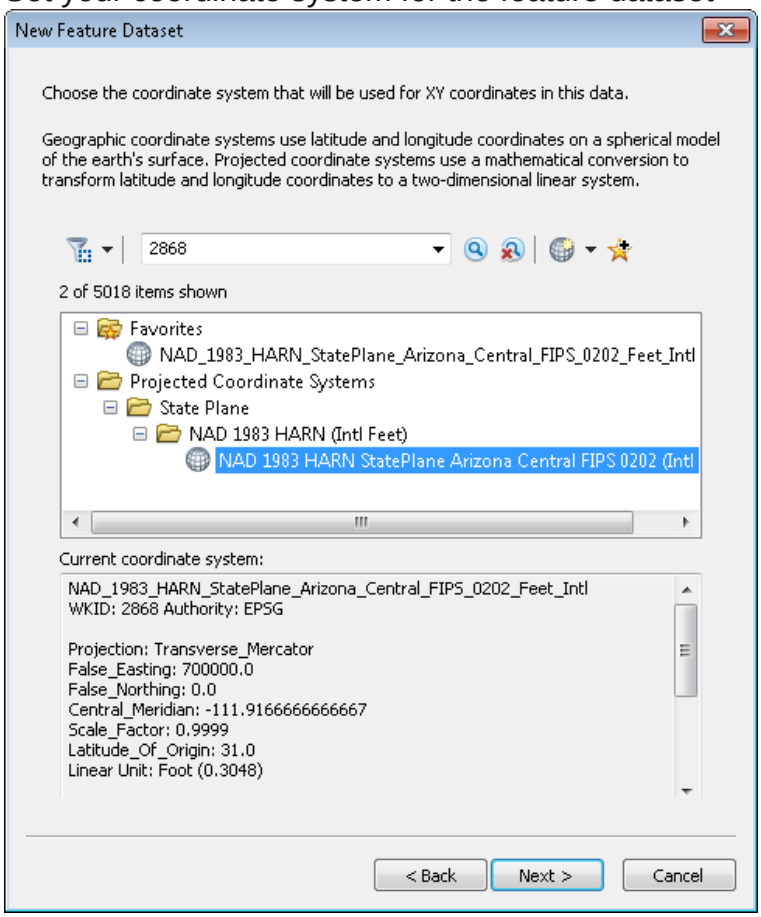

9. Click Next
10. Click Next unless you need vertical coordinate systems applied to the data
11. Set your Z, M and XY tolerance. In this example, we will take the defaults.

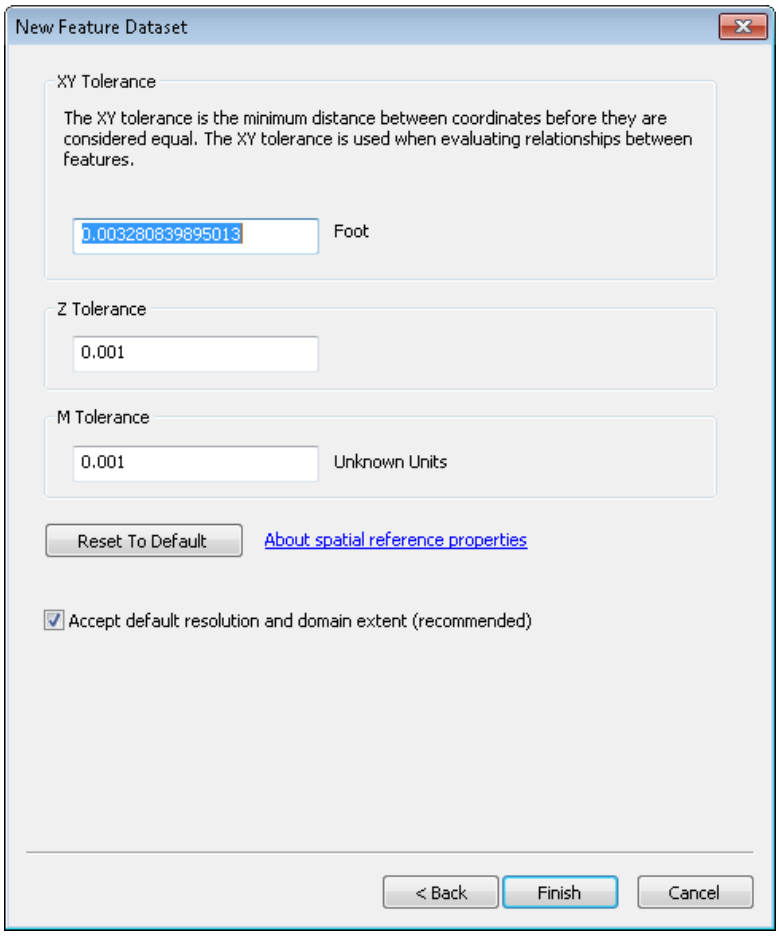

12. Click Finish
13. In the Database Connection, right click on the newly created Feature Dataset and choose new and submenu Feature Class

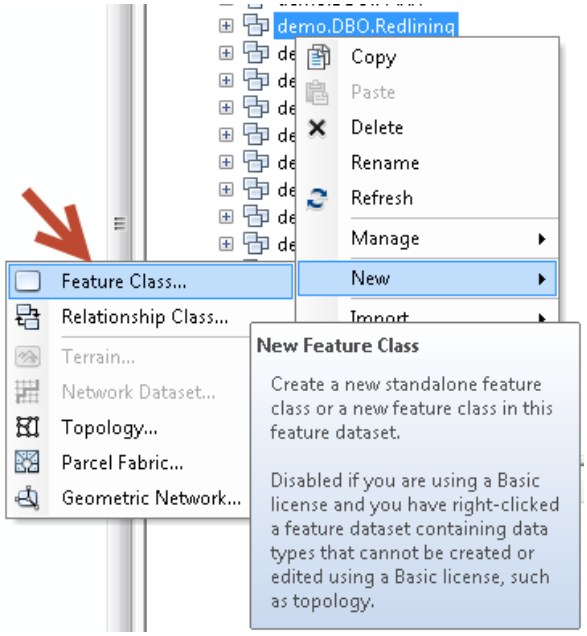

14. In this example we will create a point Feature Class and call in RPoint

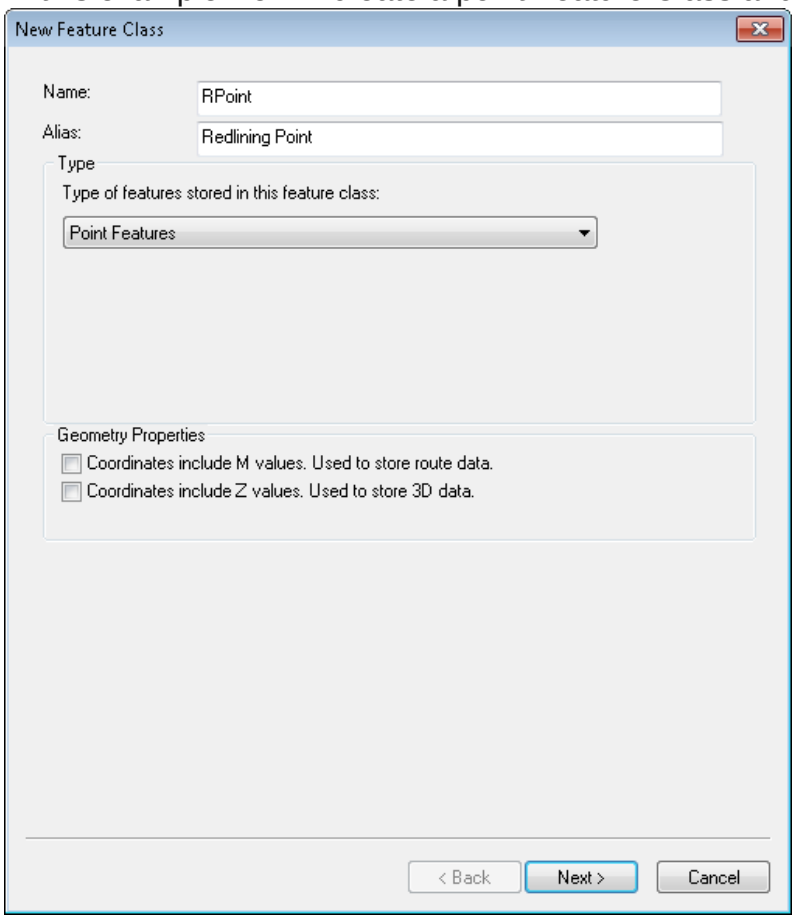

15. Click Next
16. Except defaults unless you are using configuration keywords
17. Click Next

18. Add desired fields to your redlining point feature class. If you plan on using editing user, date and last editing fields, make sure you turn on editor tracking after creating the feature class.

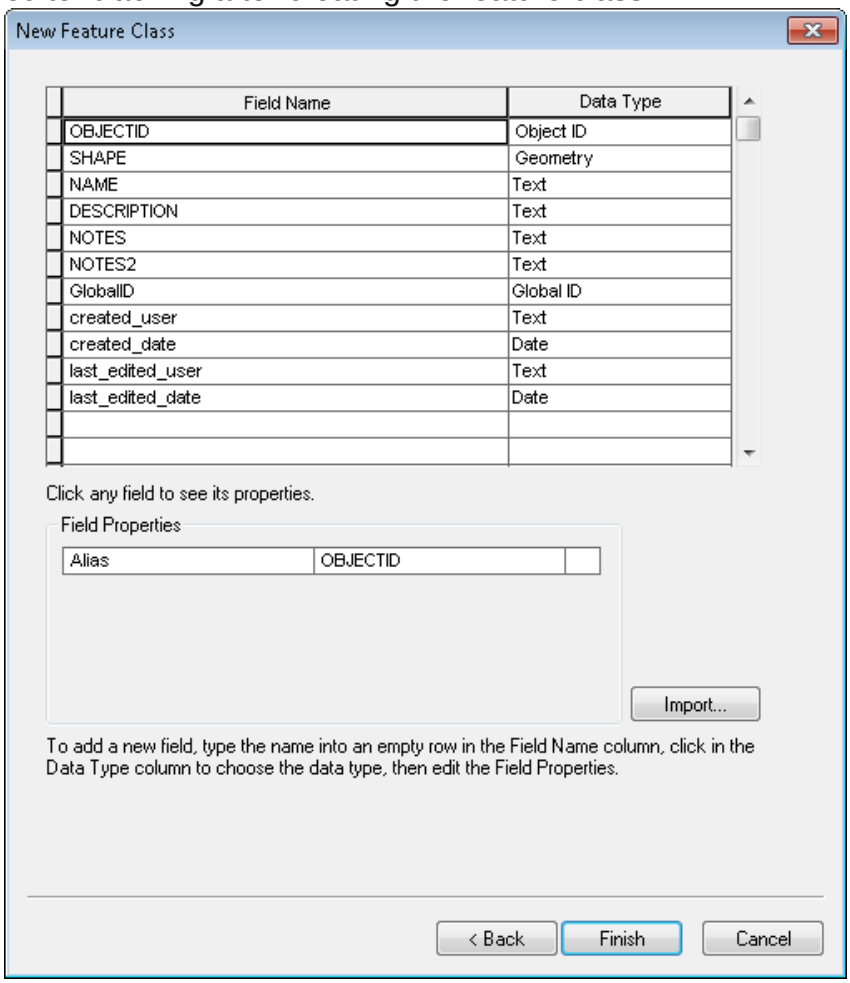

19. Click Finish
20. Repeat steps 13 through 18 to create polygon and line Feature Classes if desired.
21. Make sure to assign credentials for access to the newly create feature dataset and optional versioning before publishing.
22. Optional: Right click on each feature class that you would like to enable edit tracking and choose properties.
23. Go to the tab Editor Tracking and click on Enable Editor Tracking
24. Map the fields to the appropriate Feature Class fields
25. Recommend using UTC

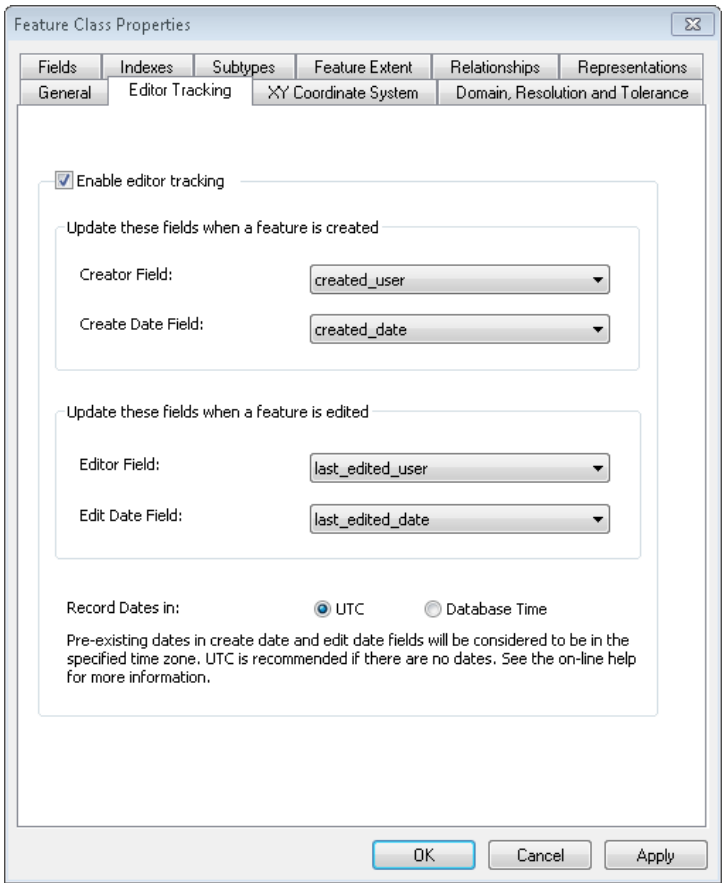

26. Click Apply or ok.

Publishing a Feature Service

27. Go to the File/Share As/Service menu option to create a service

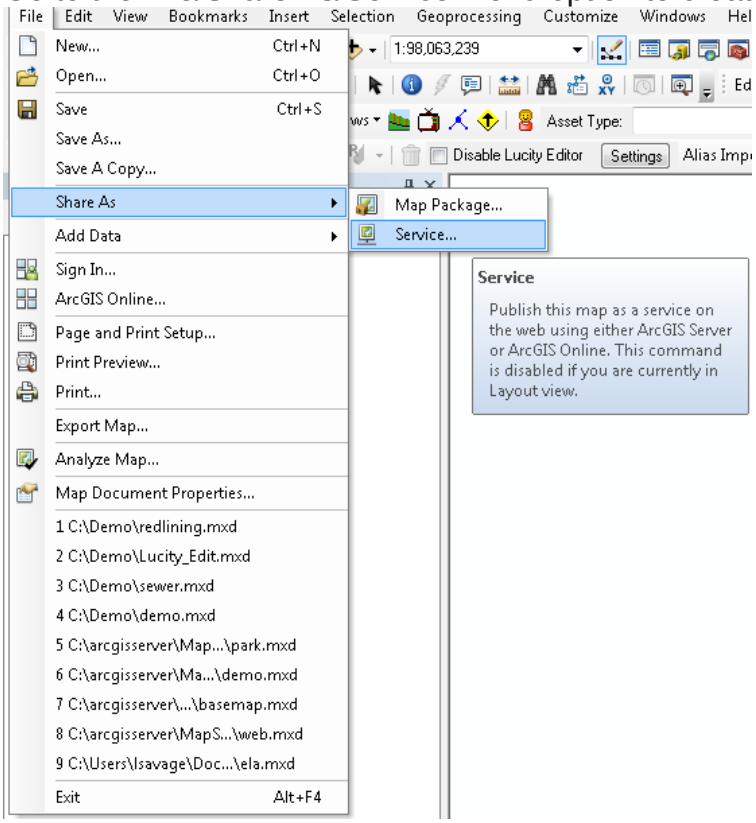

28. Select Publish a Service

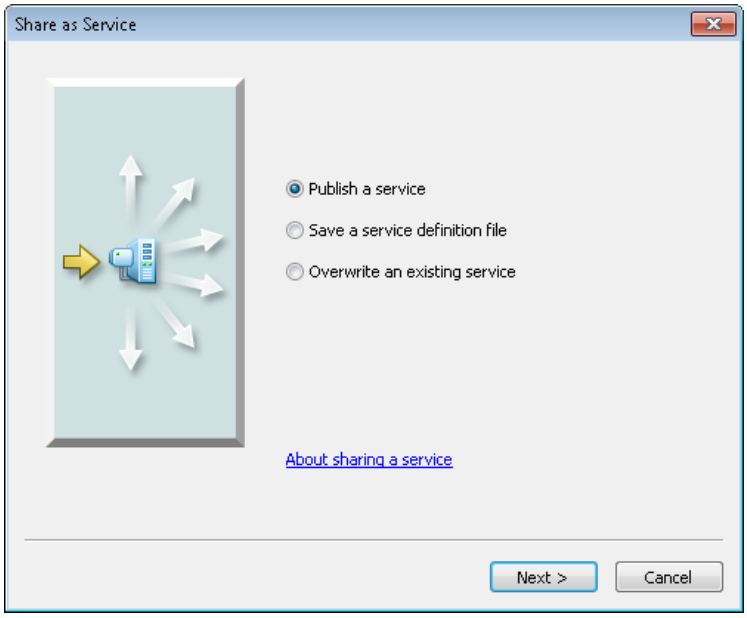

29. Click Next

30. Choose server connection and create service name

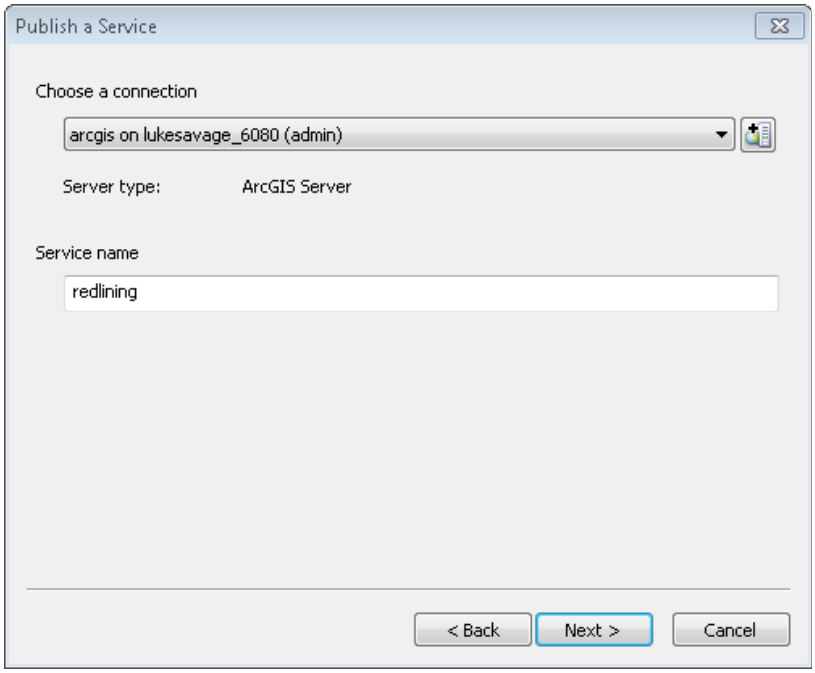

31. Click Next
32. Use in folder or in the root of ArcGIS service directory
33. Click Continue
34. Go to Parameters and click on Advanced

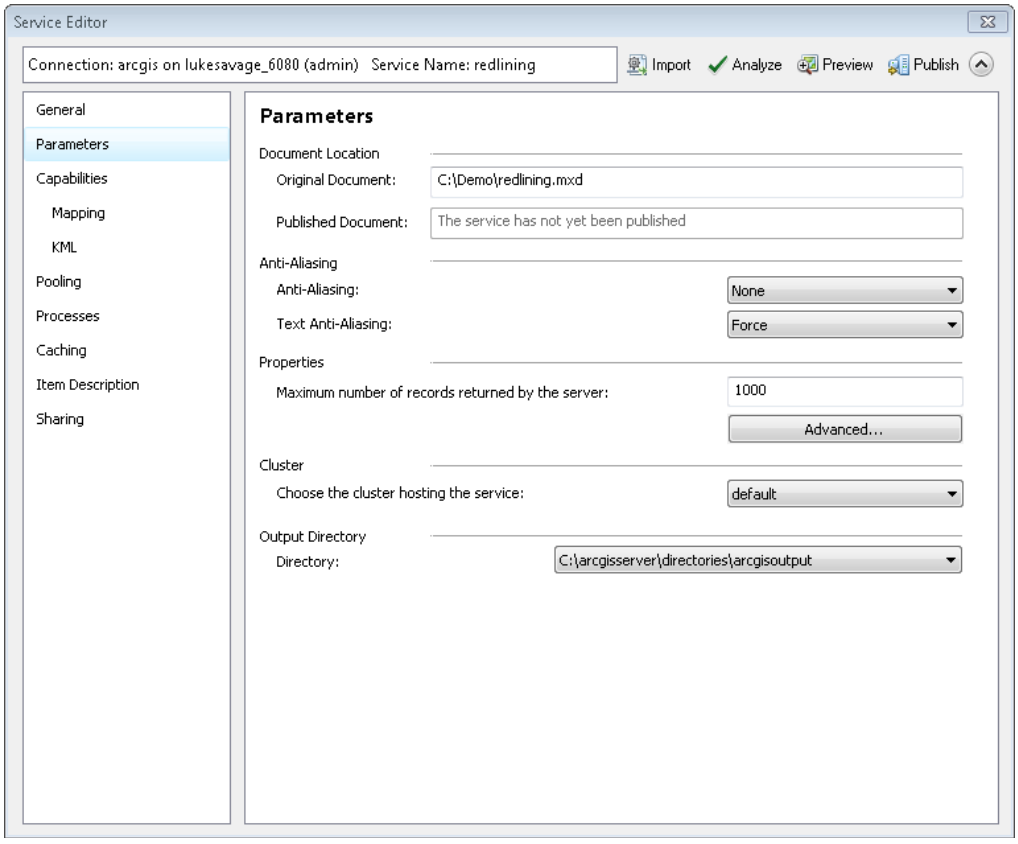

35. Set SchemaLockingEnabled value to ‘false’

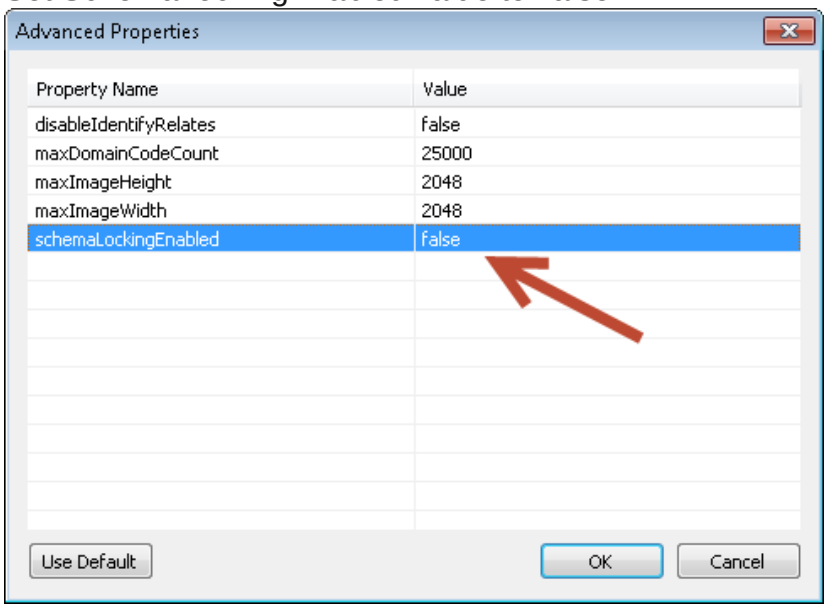

36. Click Ok
37. Go to the Capabilities tab and make sure Mapping and Feature Access is selected.

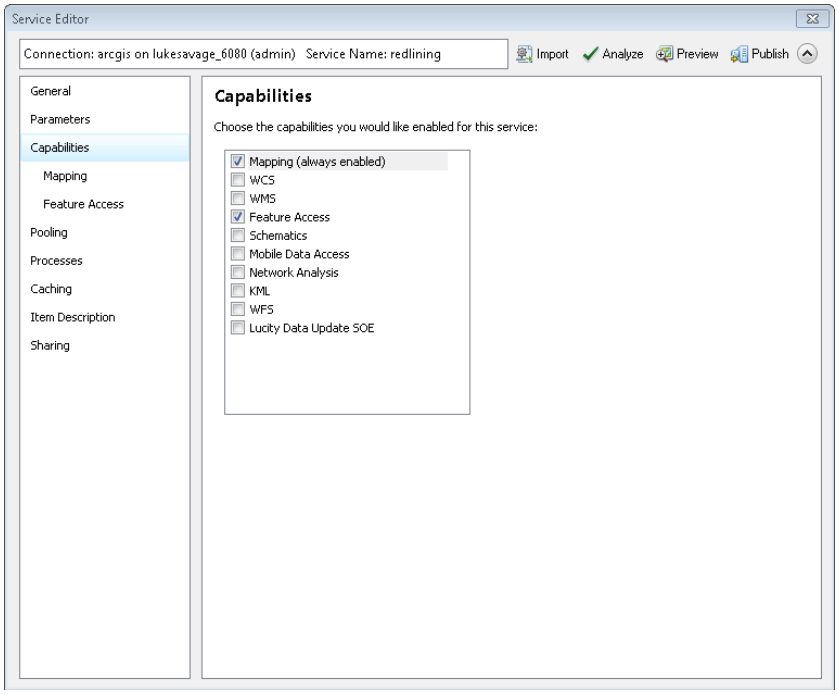

38. Click on the submenu tab ‘Feature Access’ and add the selection ‘Enable ownership-based access control on features’. This option will only allow the created user features to be updated by the owner of the features but only allow the other users to query the created data.

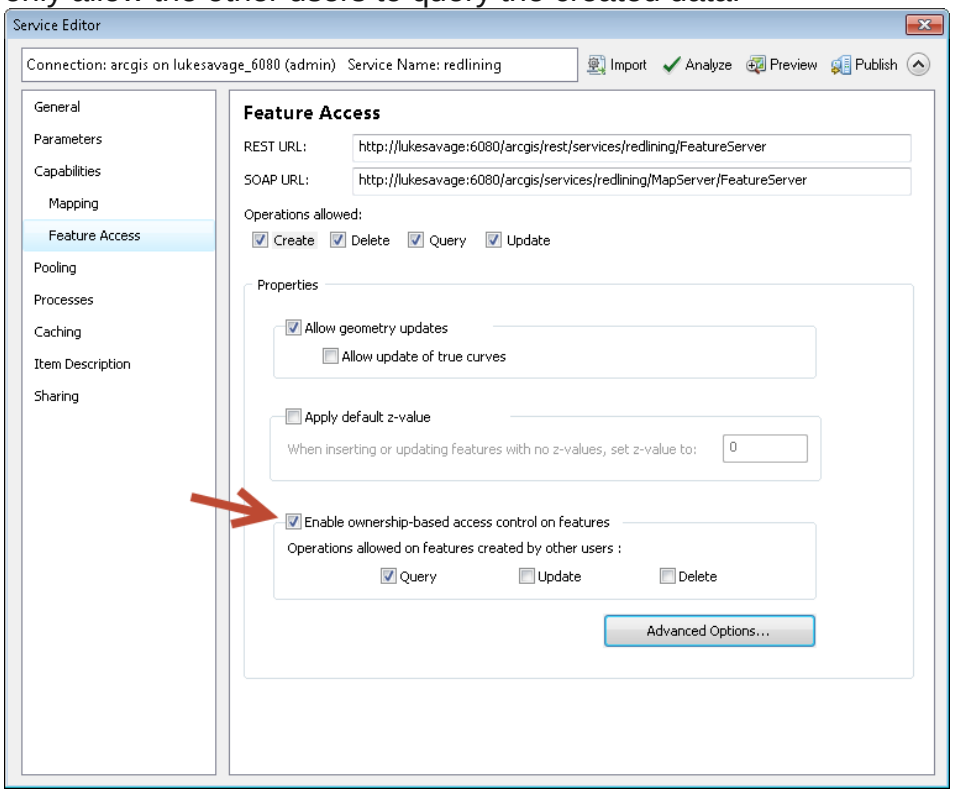

39. Analyze to make sure no errors are present.
40. Click Publish
41. Optional: If you wish to secure the feature service, go to the below section of ‘Security a Map Service’ and add each user account to have access to the newly created map/feature service.

Enabling the Geometry Service

We need to enable the geometry service for Lucity GIS Web Map application to allow for buffering, projecting, calculating lengths and areas.

1. Log into ArcGIS Server Manager.

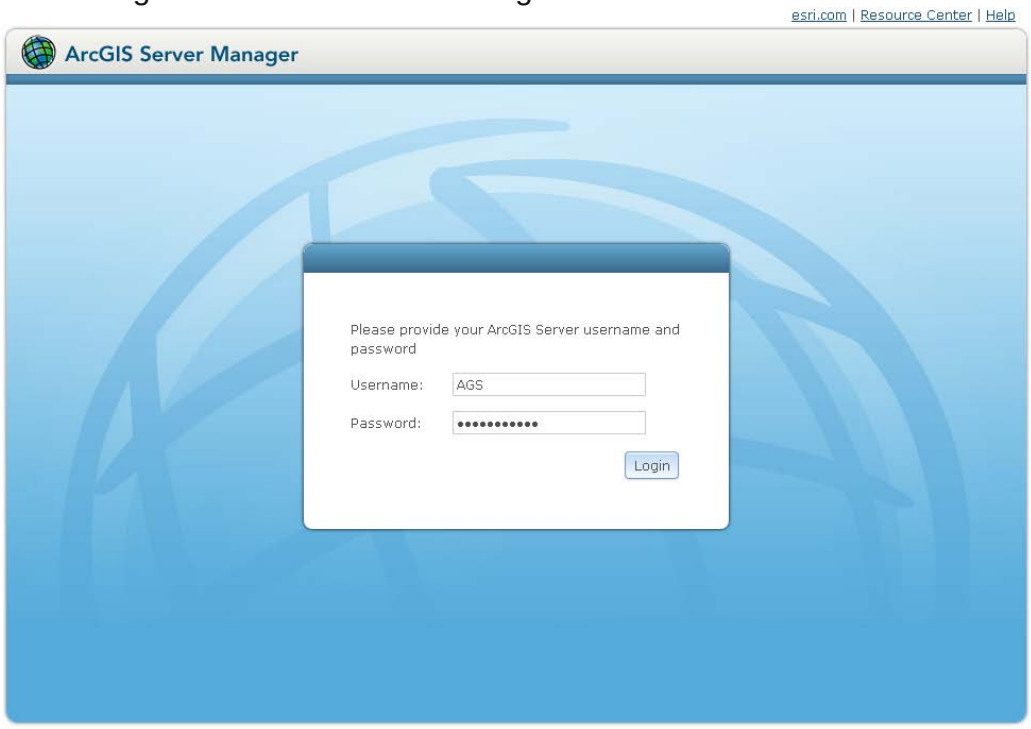

2. Once logged in, go to the site root tab on the left side of the browser and click on the folder called ‘Utilities’

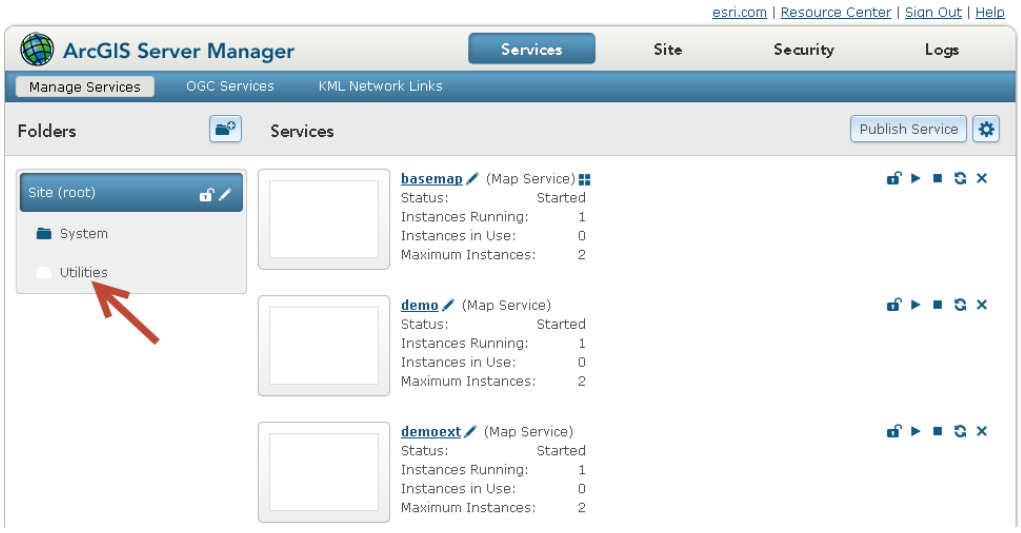

3. Go to the Geometry Service and click on the play icon ▶ to start if the service is stopped
4. Once complete, the service is now available.

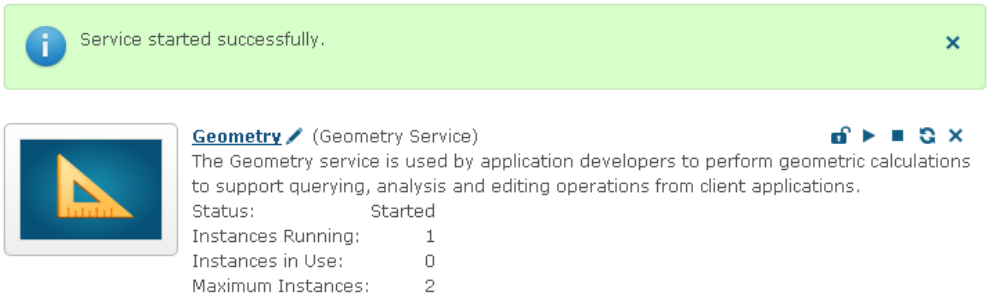

Creating a Geocoding Service

Geocoding services allow Lucity GIS Web Map to find and display addresses on a map and see how they relate to surrounding features. We do support composite geocoding services.

1. Open ArcCatalog
2. Create a geolocator in your publishing database
3. For dual ranges, make sure to set the 'Match with no zones' value to Yes

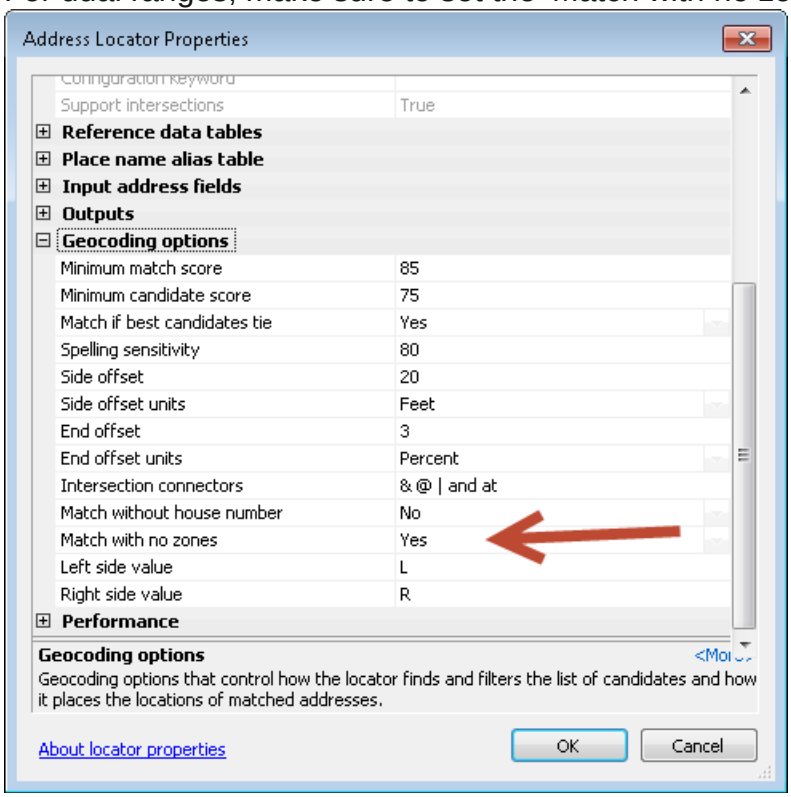

4. Right click on the geolocator in the Catalog Tree and click on Share As/Geocode service

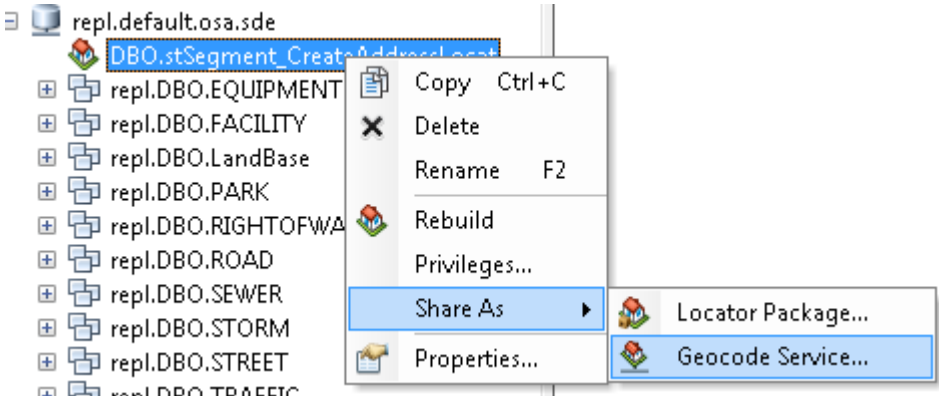

5. Publish a service
6. Click Next

7. Make the service name simple so it's easier to remember.

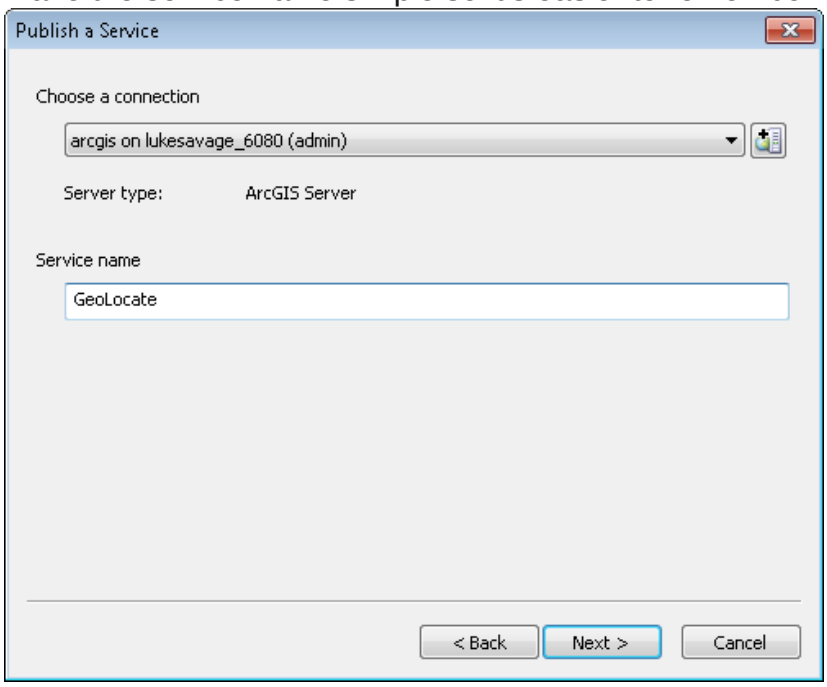

8. Click Next
9. Choose a folder or keep default location as root
10. Click Continue
11. Select Parameters and change the interactive Find if you want more than 500 records returned by the server.

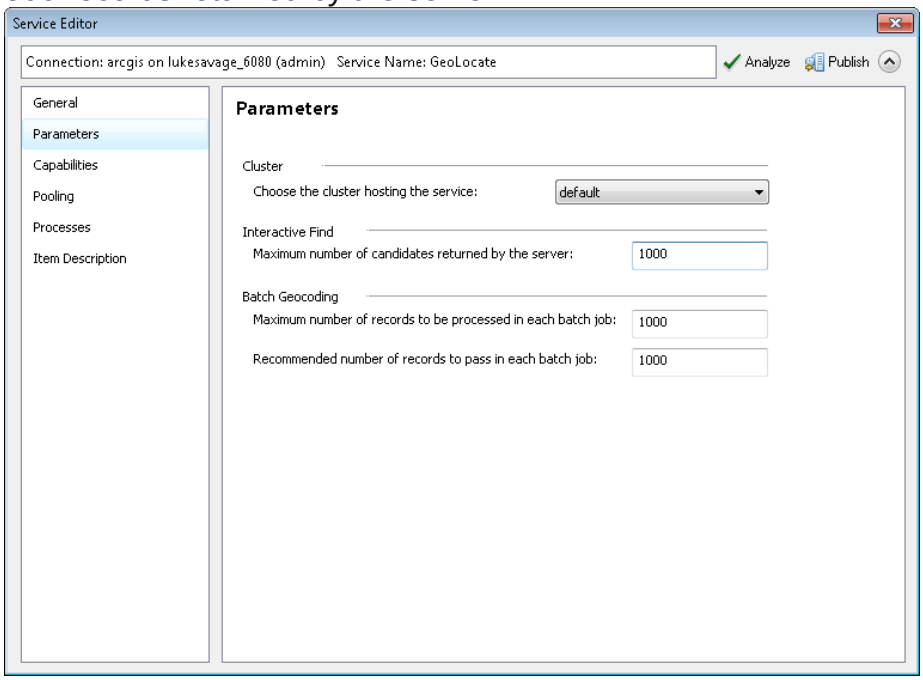

12. Analyze and fix any errors. If no errors, publish

Securing a Map Service

Starting at Lucity 7.60, we are supporting ArcGIS for Server secure map services. In this section, we will describe the deployment of secure map services as a built-in user store.

1. Log into ArcGIS Server Manager.

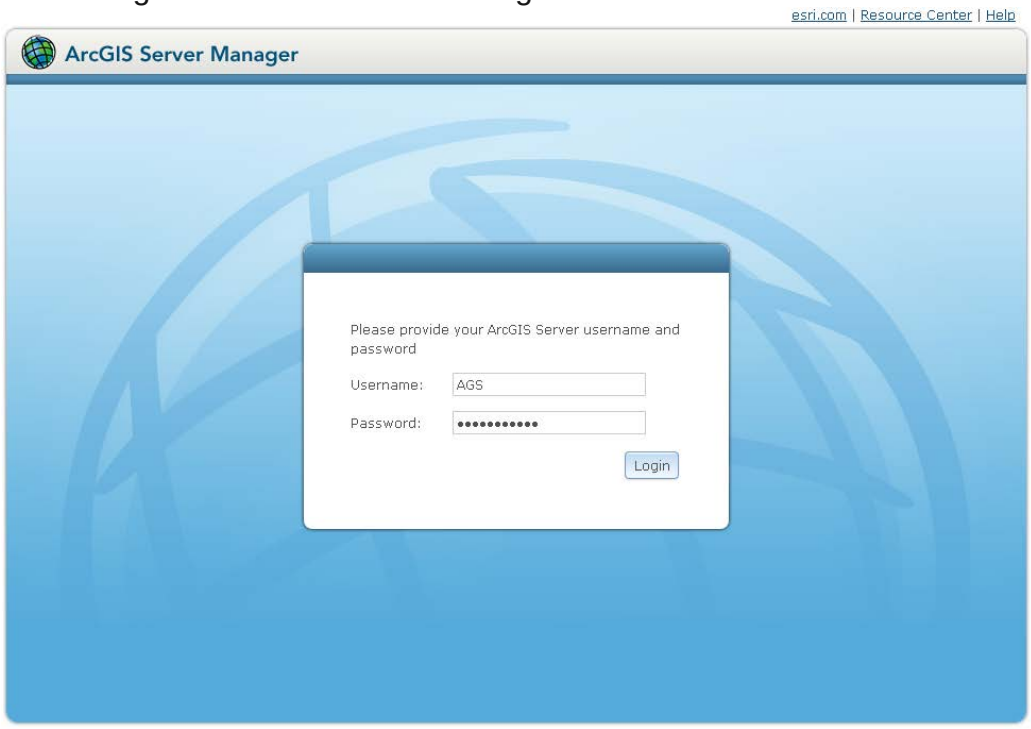

2. Go to the Security menu to configure the security features.
3. Click on configuration settings pencil to configure the type of data store.

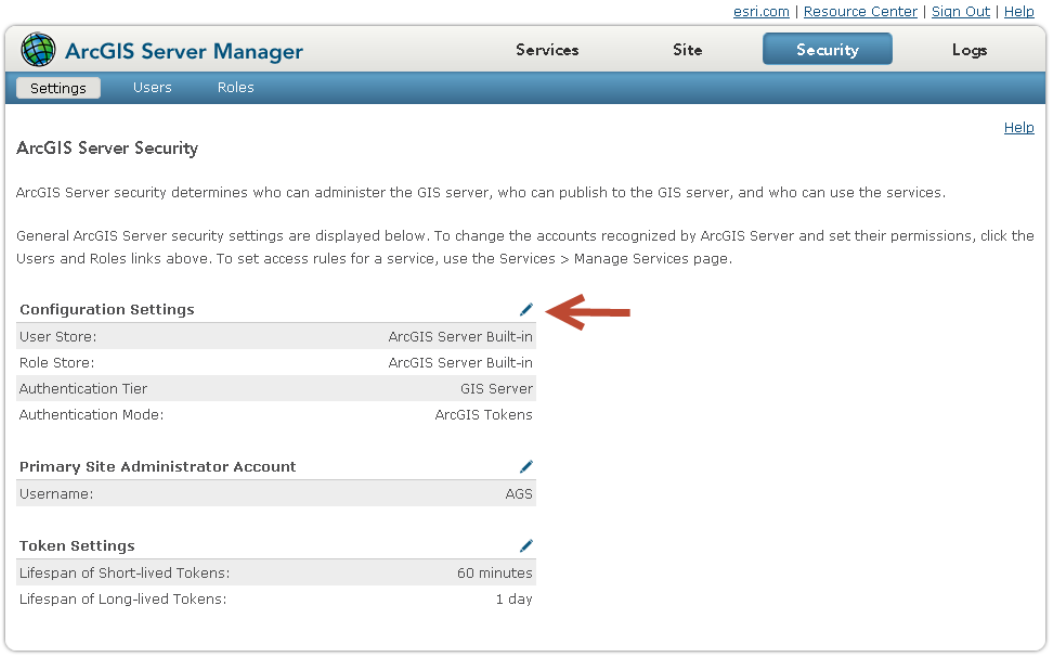

4. Again, we will be configuring a built-in store at the ArcGIS for Server level not the web store or the windows domain/LDAP configuration in this example.

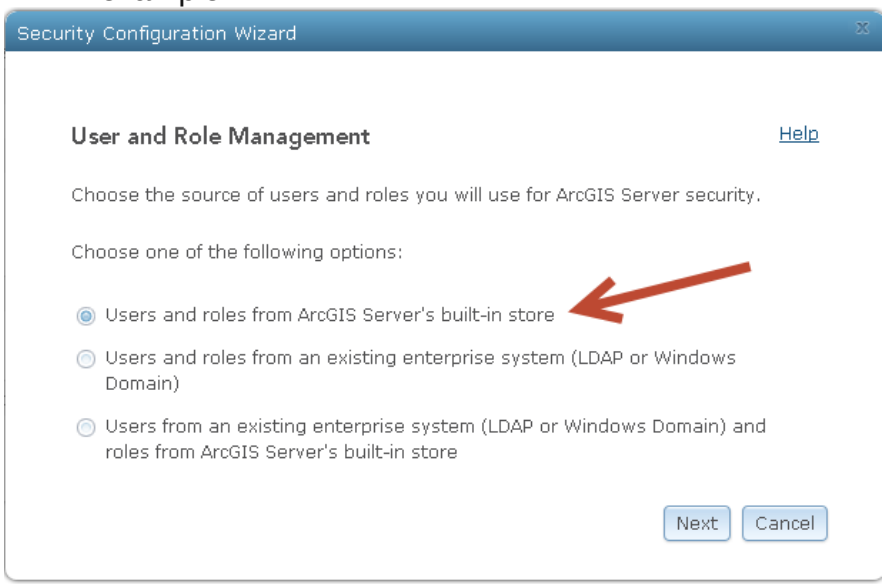

5. Click Next
6. Finish after reviewing configuration
7. Optional (if using tokens), configure your Long-Lived tokens to the appropriate IT policy. In this example, the lifespan for Long-lived tokens will be 1 day.

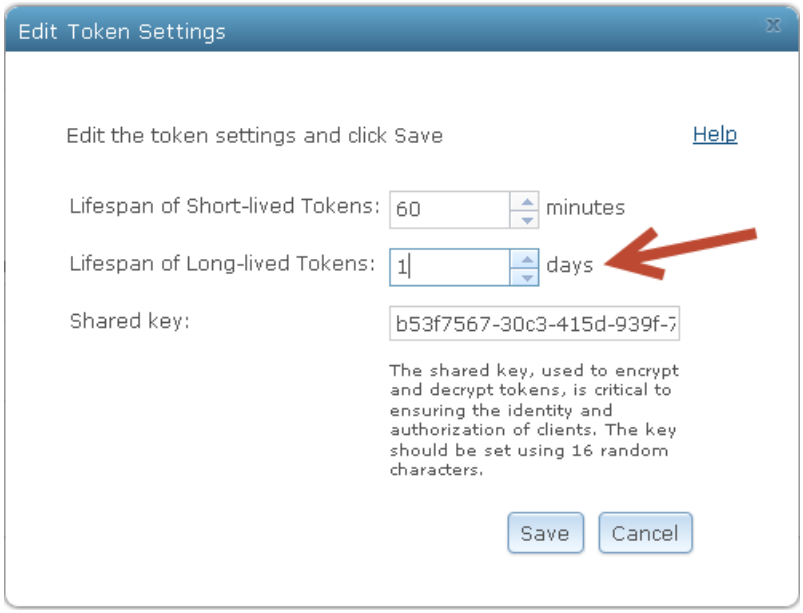

8. Click Save
9. Go to the roles submenu of security

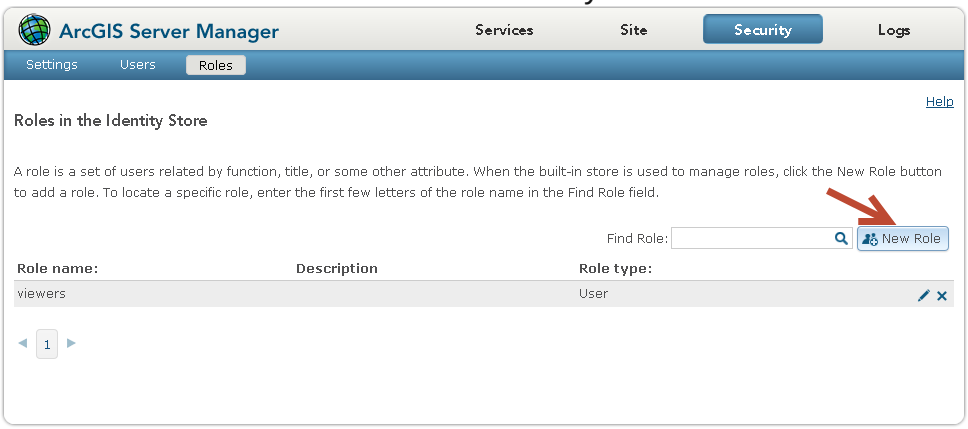

10. Click new role
11. Add a role called ‘viewer’. In this example, we will create a role used for the end user. You can similarly create roles for administrators and publishers in the same way. If you already had a user create, you can add them to the role by clicking the person with the plus sign icon

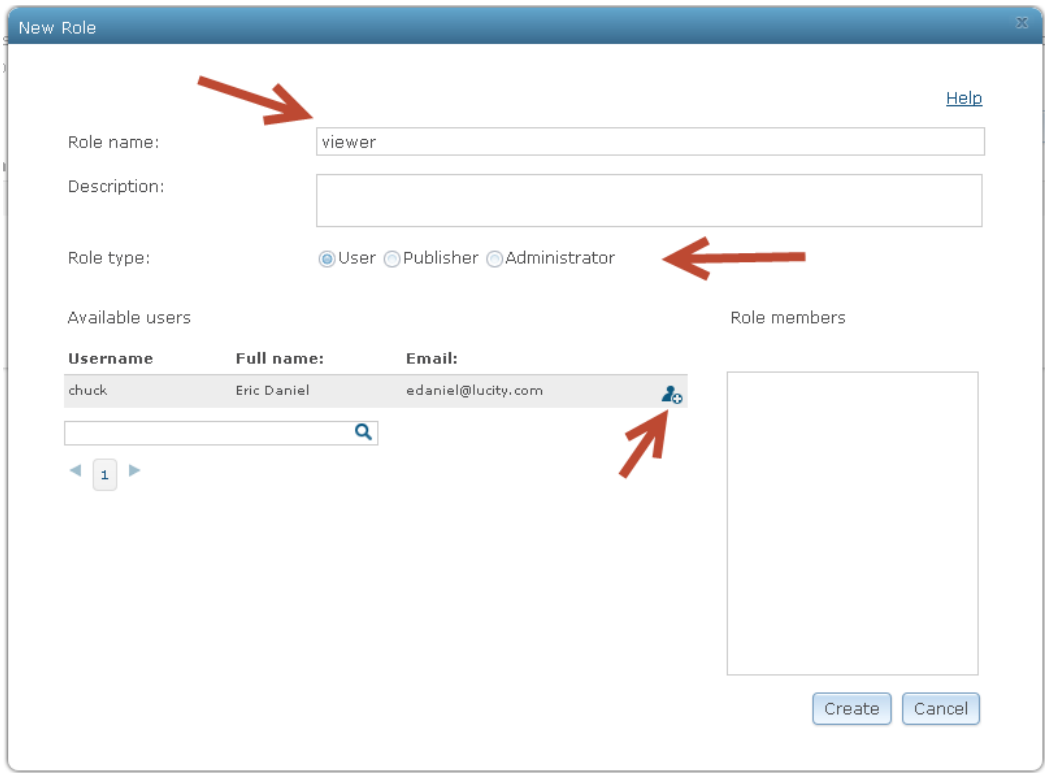

12. Once finished, click on create.
13. In this example, we will create a user. Go to the submenu item called ‘users’.
14. Click on new user.

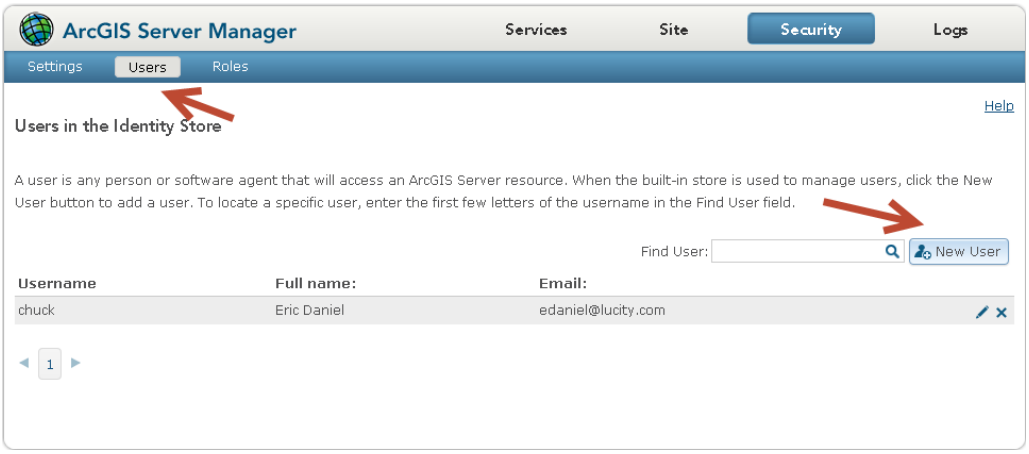

15. Once the ‘new user’ dialog appears, add the user information. In this example, we will be adding the newly create role to the user during the creation of the user account.
16. Add the information in the ‘new user’ dialog
17. Click on the person with the plus sign icon

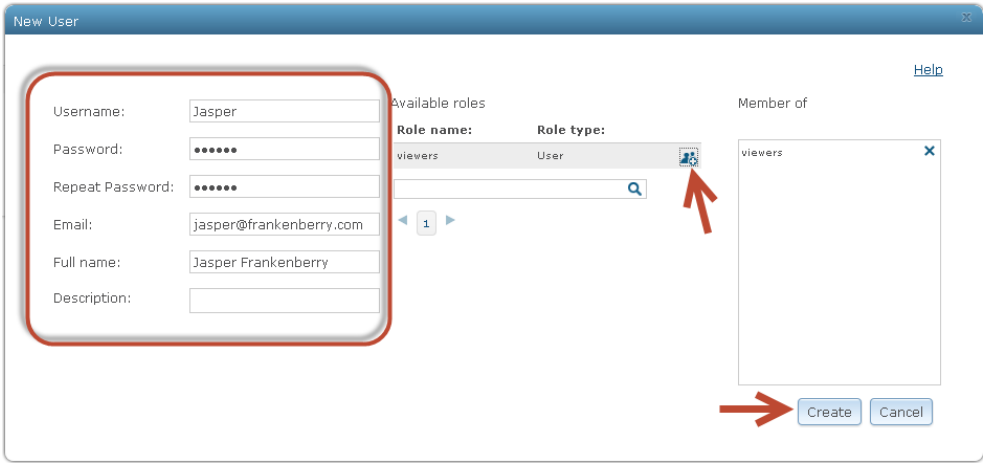

18. Click create
19. We now have the new user added to the users list.

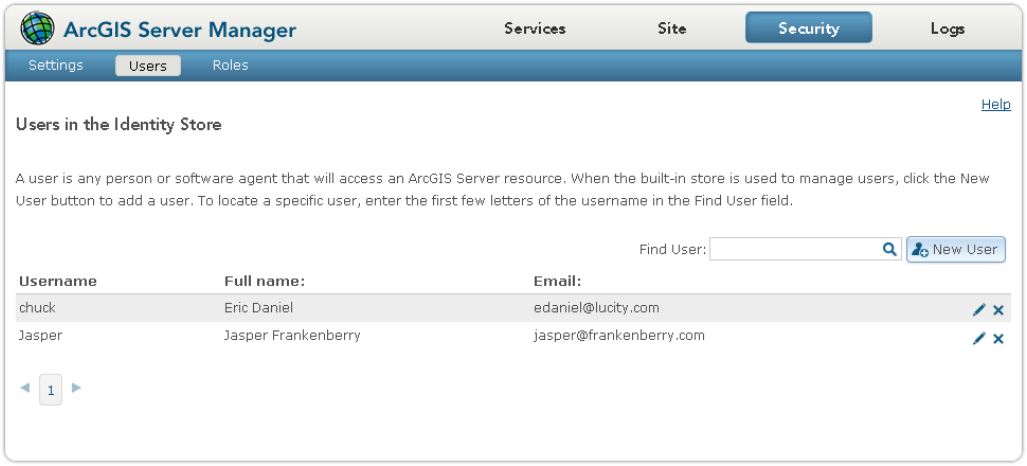

20. Assuming you published a map service for the intention to secure, go to the main menu and click on ‘services’

21. Click on the lock icon
22. Click on the radio button ‘Private, available only to selected users’ to enable security. You have the choice to allow access to all users who are logged in which gives access to all users in the ‘built-in’ users security store but in this example we only want to use the ‘viewers’ role.
23. A message will appear ‘no roles have been selected’. This is a friendly reminder that there have been no roles selected.

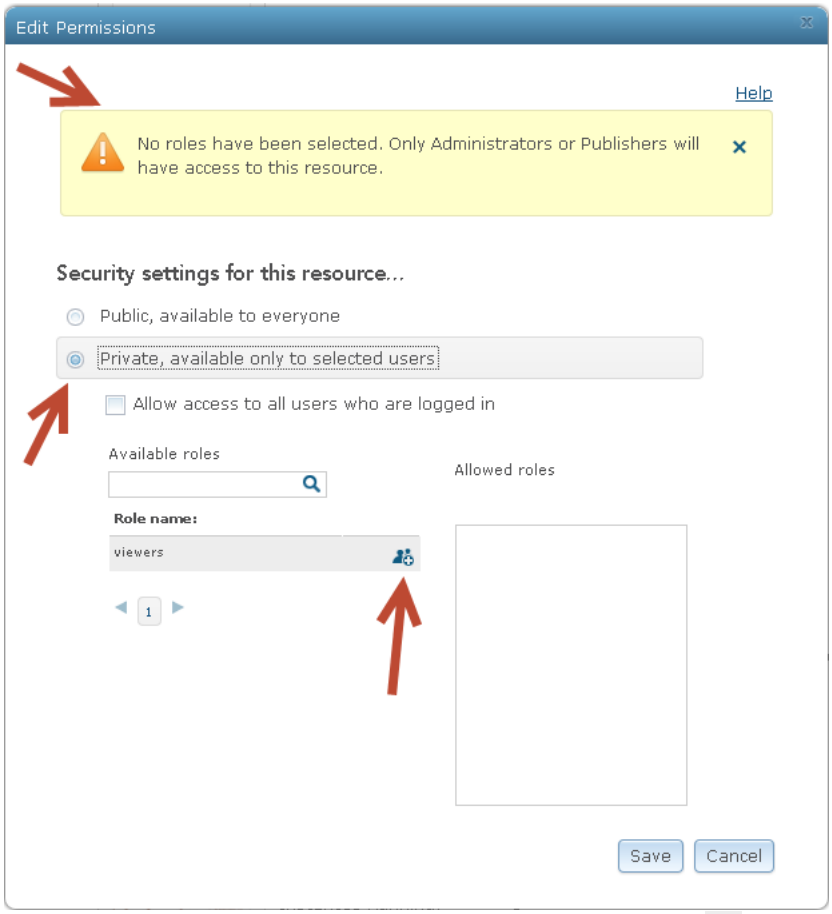

24. Click on the person with the plus sign icon under roles. This will add the 'viewers' role into the 'Allowed roles' dialog.

25. Click Save
26. For the map service, you should now see the lock icon in the locked position.

27. Test the map service security. You can do this by going to the rest endpoint of the map service.
    Example:
    http://<servername>/<alias>/rest/services/<nameofmapservice>/MapServer
28. Supply the newly created user credentials and click on login

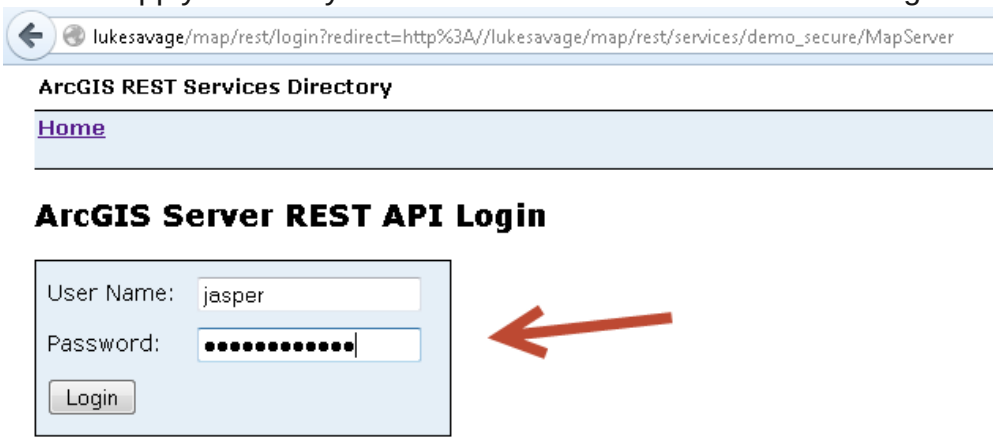

29. If you see the rest service directory open, you have successfully secured your map service.

Maintenance

Geodatabase Administration Tools

Starting in ArcGIS v10.1, the interface has changed so an admin user can administrate the Enterprise Geodatabase.

1. Go to ArcCatalog and connect to a Geodatabase as an administrator with Dbo privileges. Once connected, right click on the connection and choose administration.

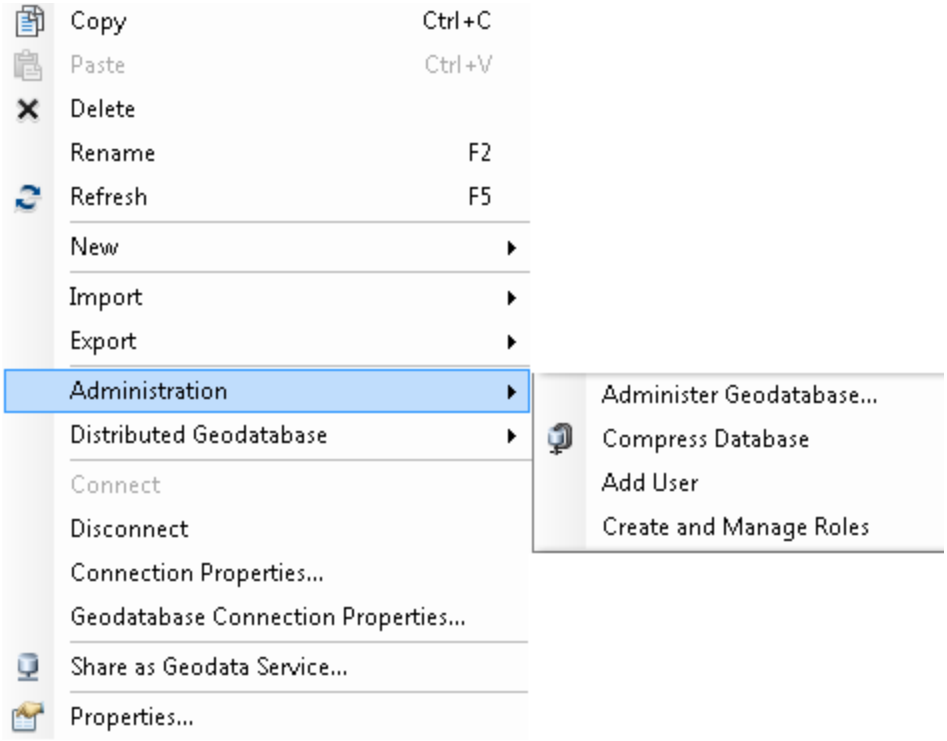

2. In the administration menu, you will see four choices. Compress is now integrated within the administration menu. Also, you can add users and create/manage roles in your enterprise Geodatabase without touching SQL Server. Roles will allow you to give permissions to the role without having to give permissions to each and every user. Think of a role as a group.
3. Click on Add User. You have two choices, operating system authentication or sql server authentication which gives an arbitrary login and password not related to windows but related to sql server only. You have the choice of adding the user to a role.

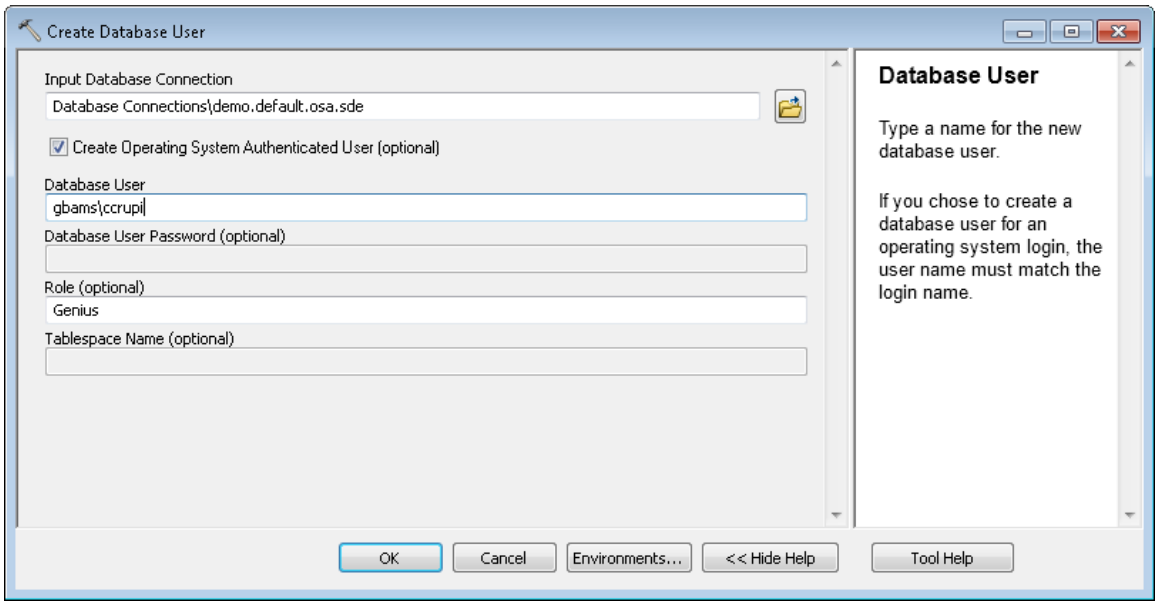

4. Go to Administer Geodatabase on the administration menu, the Geodatabase Administration dialog will appear. You can create versions, check connections and locks.

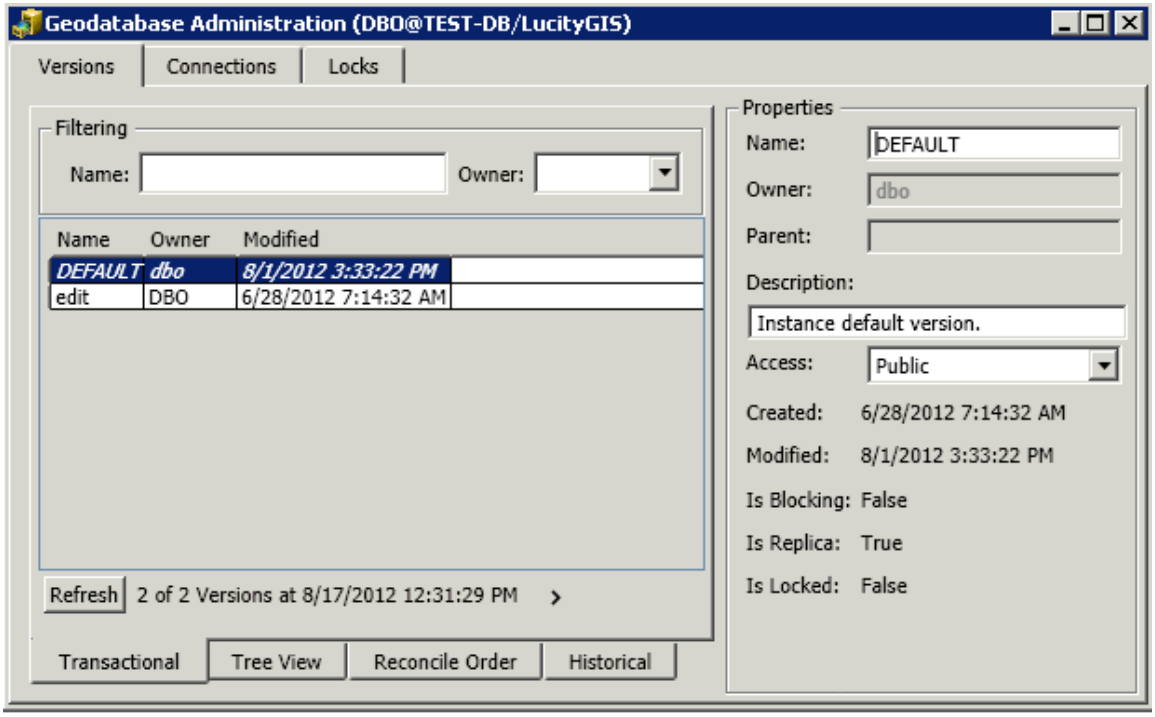

Create Role

There are two different ways you can create roles. One way is to create an SDE role and assign users to the roles. The other way is to create Active Directory groups, add active directory groups to SQL Server allowing for Active Directory groups to act like roles. The great thing is that Active Directory is then managed by your IT user administration system without needing to manage all the users through SQL Server. We will show both ways below.

Create Role within ArcCatalog

1. Right click on the database connection and go to Create and Manage Roles. You can access this from ArcToolbox Data Management Tools as well.

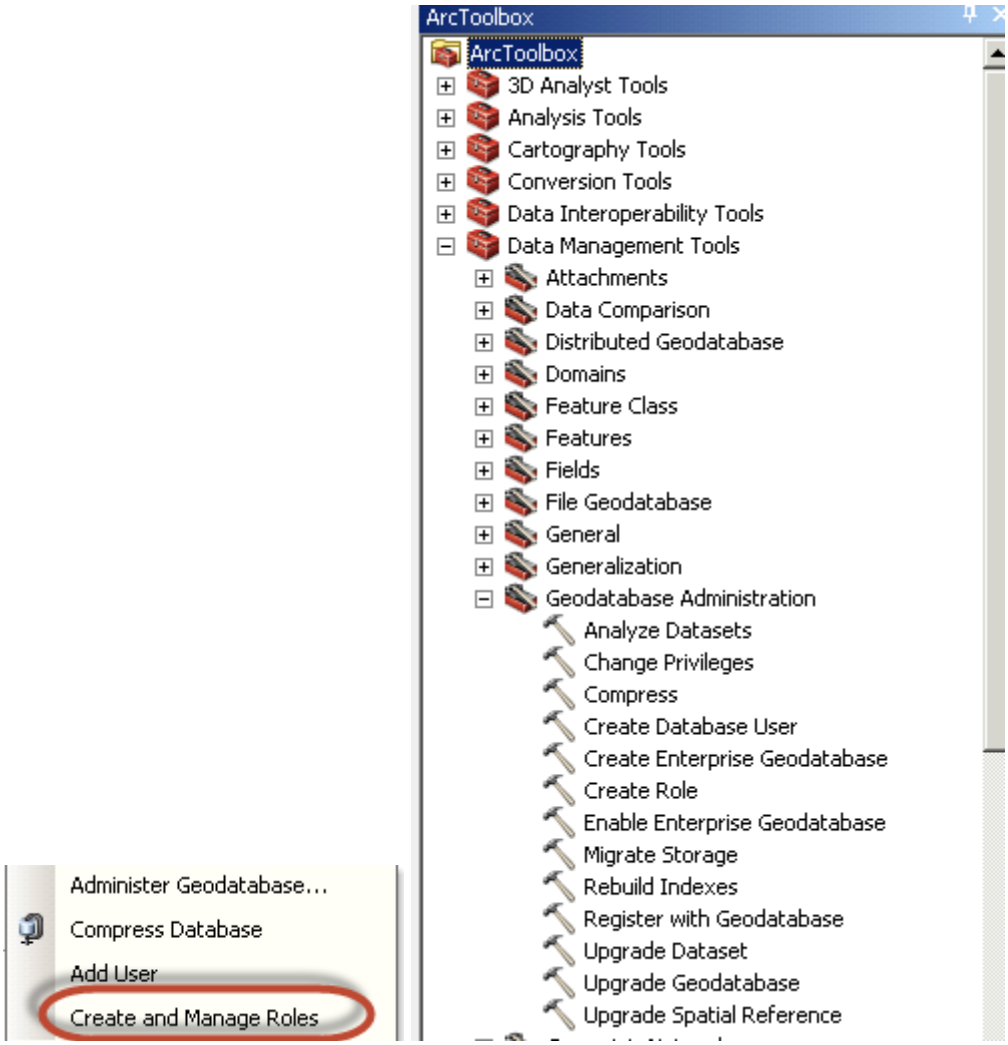

2. The create role dialog will appear. You can grant or revoke users from roles from this interface.

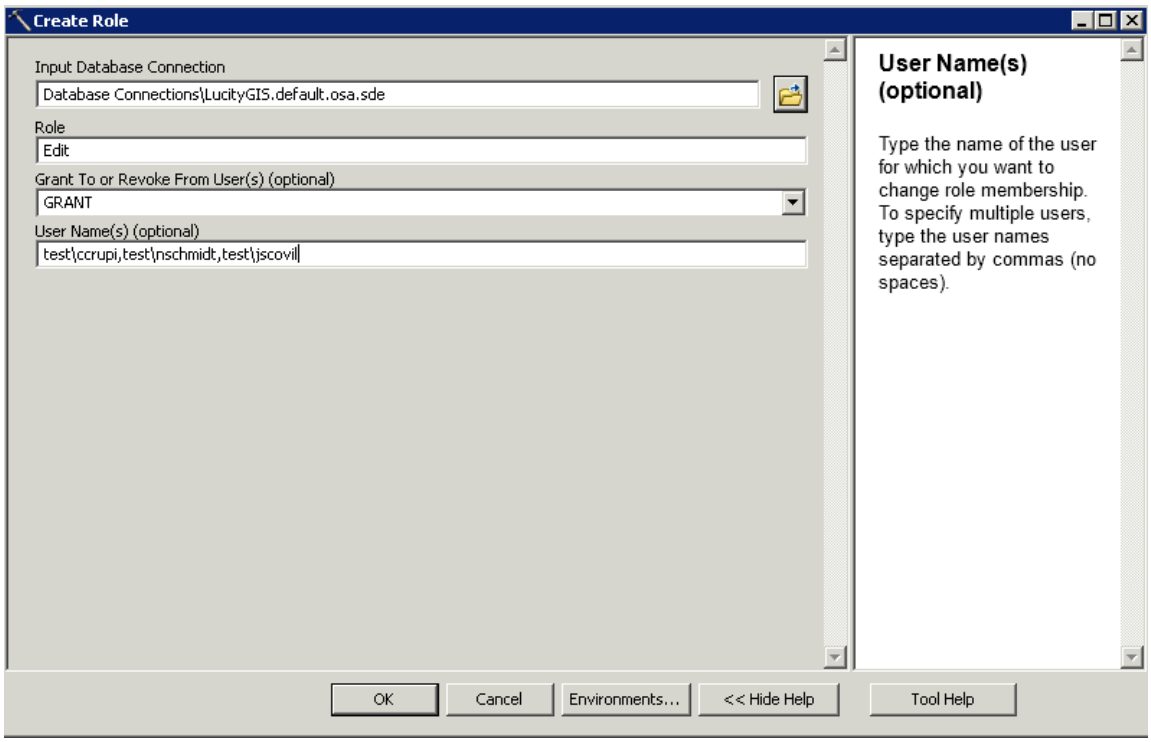

3. Create Active Directory Groups for the Enterprise Geodatabase
4. Open the Active Directory Users and Computers dialog from your domain controller and add a group by right clicking on users and select new group. Add a name such as GISView. The default settings are fine. Click ok.

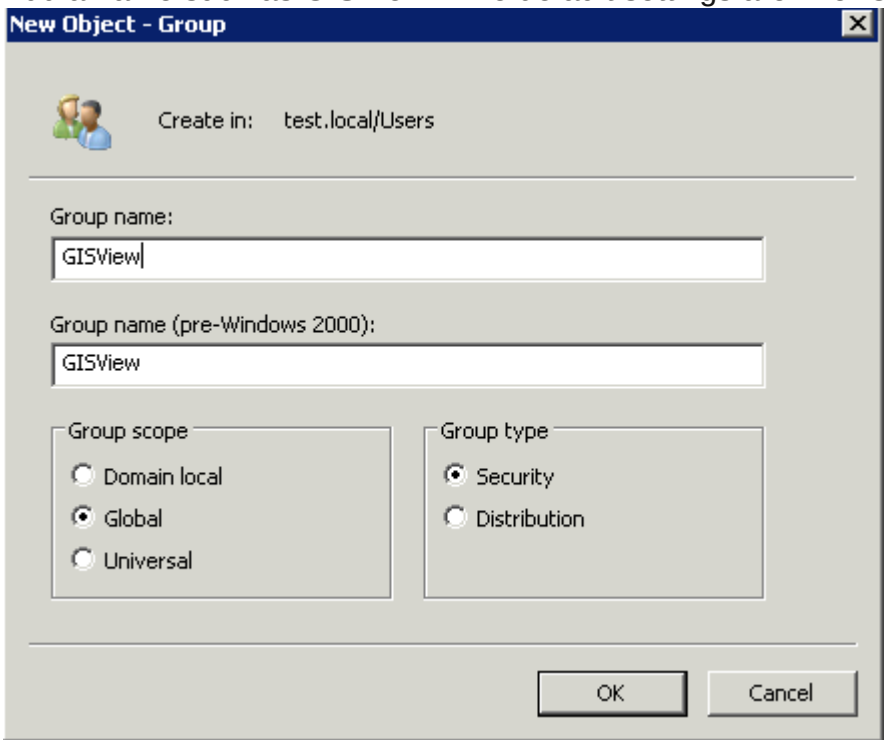

5. Double click on the new group and add members within the members tab.

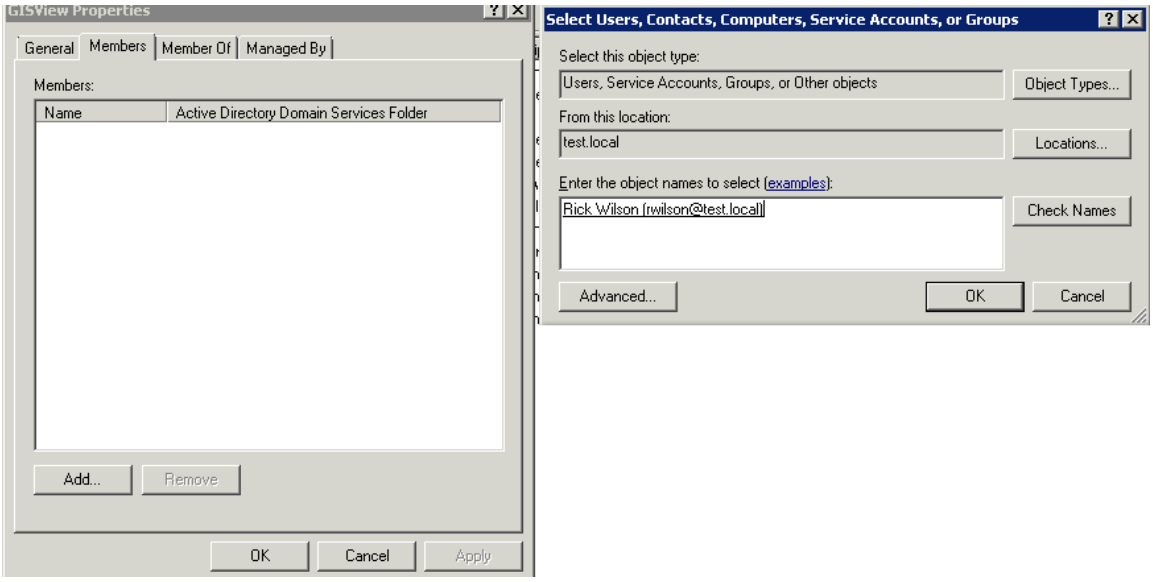

6. Add the member and click ok.
7. Open SQL Server Management Studio and add new group to SQL Server by going to the TOC of the instance connection and expand security. Right click on the Login folder and select new login.

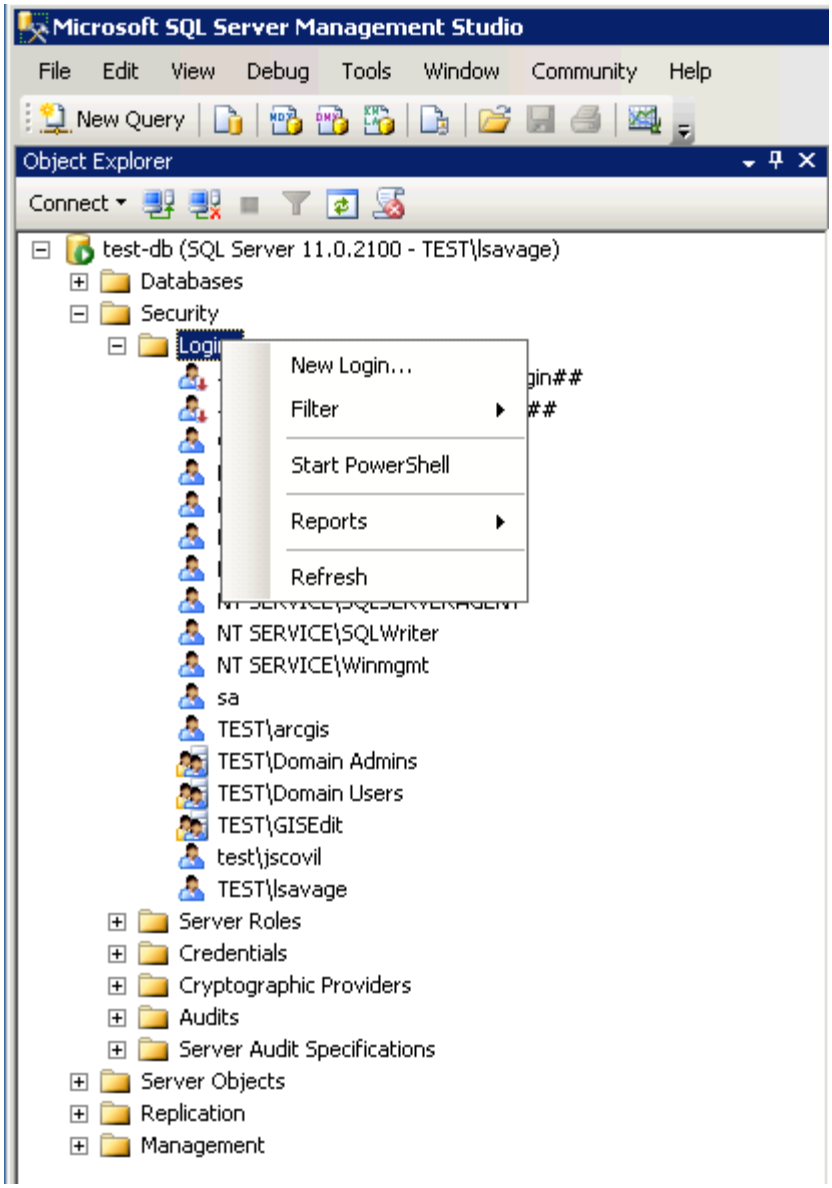

8. Click on Search and make sure you are connecting to your domain as the location and add groups as a selectable item from Object Types. Click ok when you've found the group. In this example, we're using test\GISView. Make sure the default database is selected to something other than master. Since these users are going to a view only group, we're assigning the group to the replica as the default database.

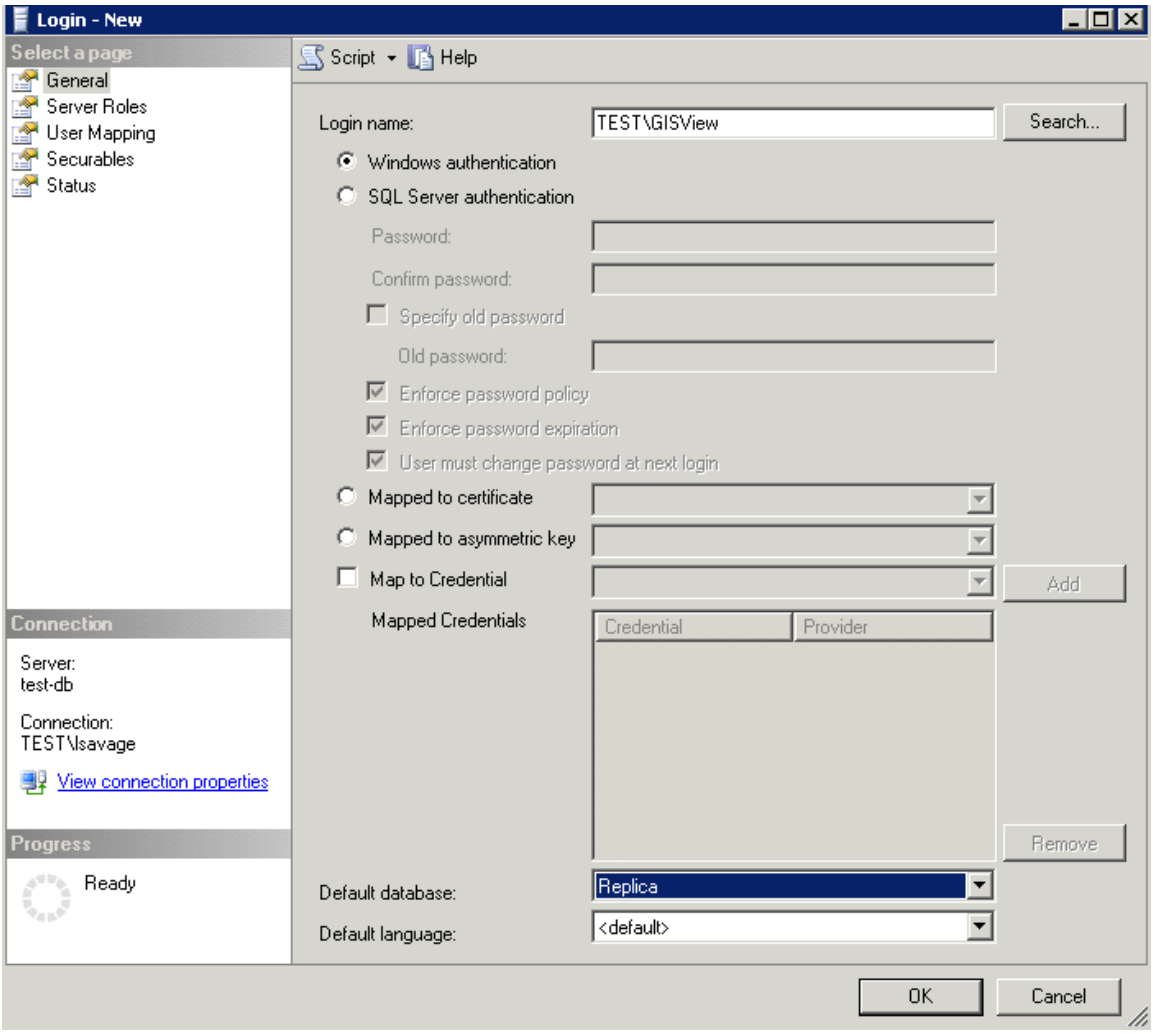

9. In the TOC of the Login – New dialog, click on user mapping. Check the boxes next to the databases you want the group to be assigned to. Just keep database role membership for <database> set to public. We'll let the Enterprise Geodatabase control privileges. Click ok.

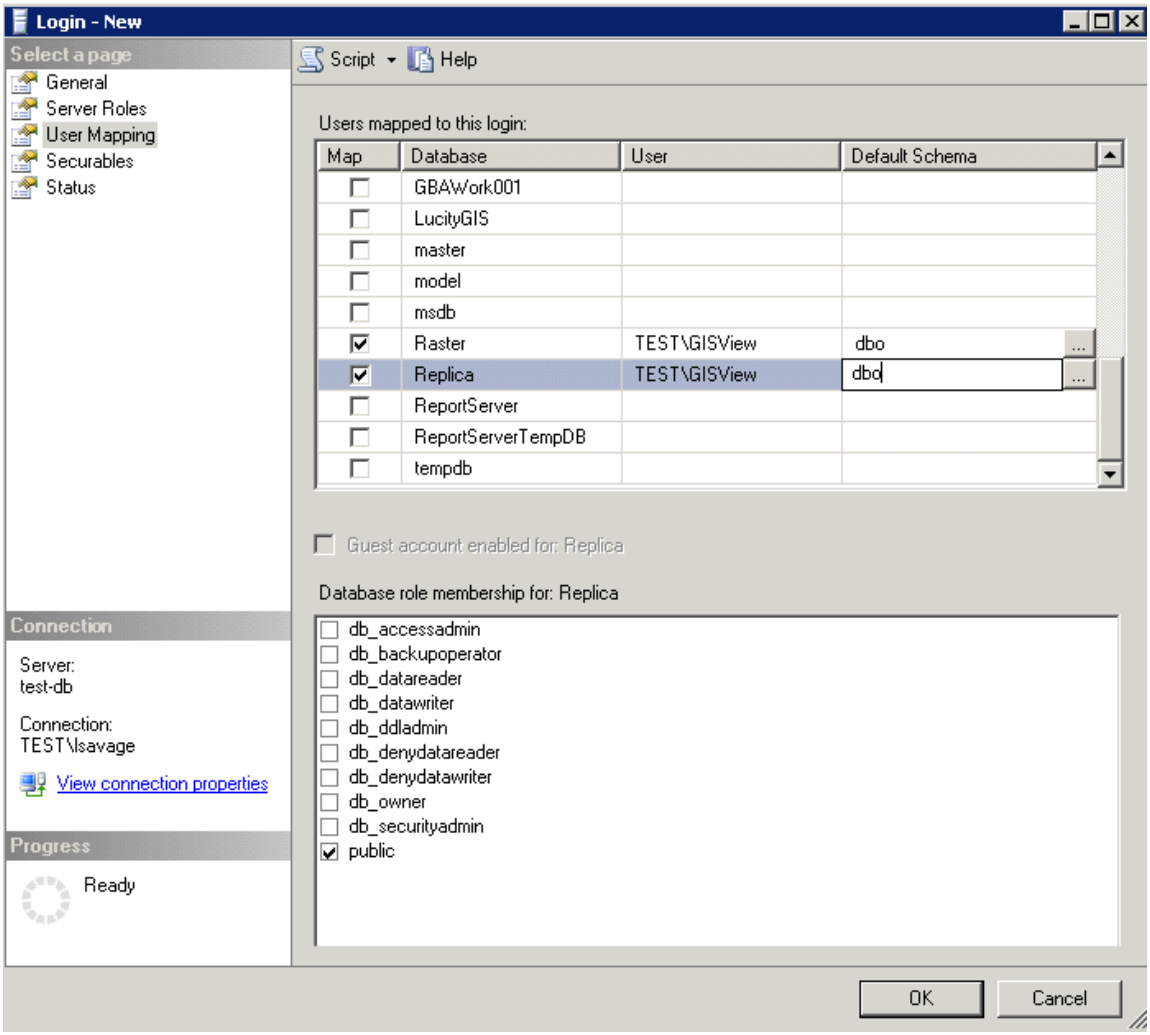
Login - New
Select a page
General
Server Roles
User Mapping
Securables
Status
Script
Help
Users mapped to this login:
Map
Database
User
Default Schema
GBAWork001
LucityGIS
master
model
msdb
Raster
TEST\GISView
dbo
Replica
TEST\GISView
dbo
ReportServer
ReportServerTempDB
tempdb
Guest account enabled for: Replica
Database role membership for: Replica
db_accessadmin
db_backupoperator
db_datareader
db_datawriter
db_ddladmin
db_denydatareader
db_denydatawriter
db_owner
db_securityadmin
public
Connection
Server:
test-db
Connection:
TEST\lsavage
View connection properties
Progress
Ready
OK
Cancel

Change/Add Privileges

In the previous exercise, we created two types of role scenarios. One is using the enterprise Geodatabase tools to create a role with multiple users which can be revoked per user. The other is allowing Active Directory Groups to be assigned like roles. Now, we are going to assign privileges to these two types of roles.

1. Go to ArcCatalog database connection as administrator and right click on the data you want to add privileges to. Remember, you can multiply select one object at a time (Feature Datasets, Feature Classes, and Tables). In this example, we'll be adding the active group we created. Select Privileges.

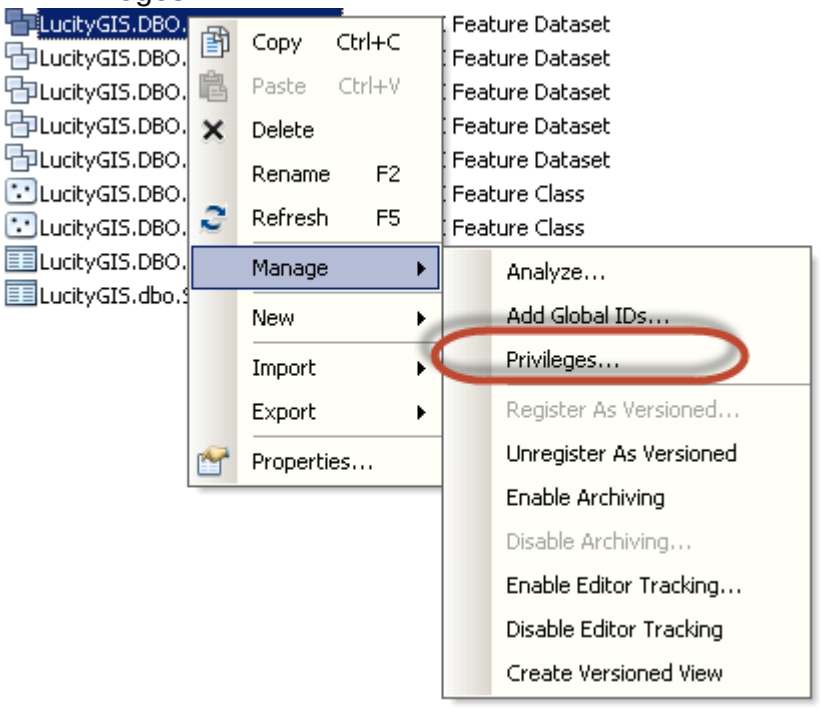

2. Click on the user/role you would like to add. Click ok.

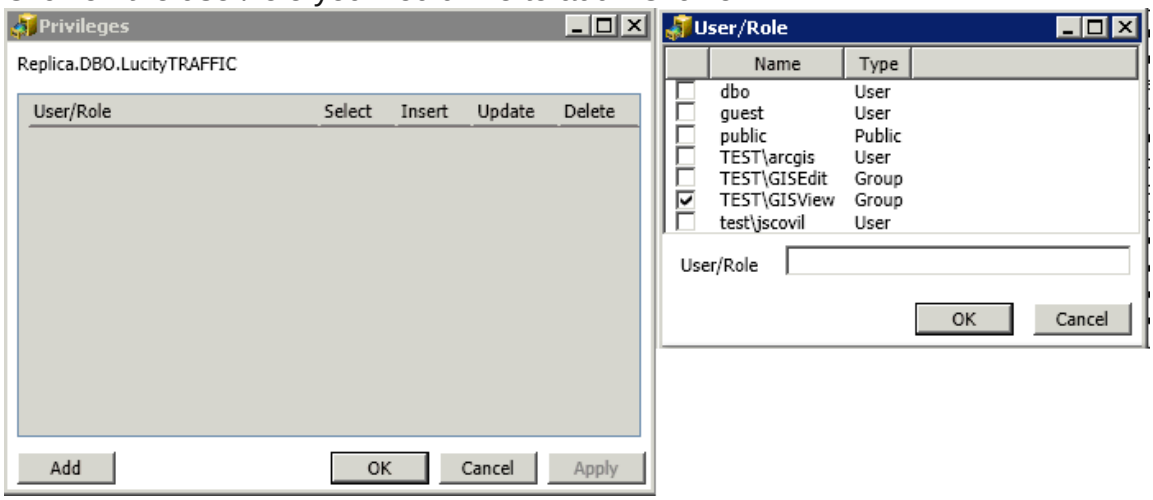

3. Since the GIS View group is read-only, only keep the select box checked. Click ok.

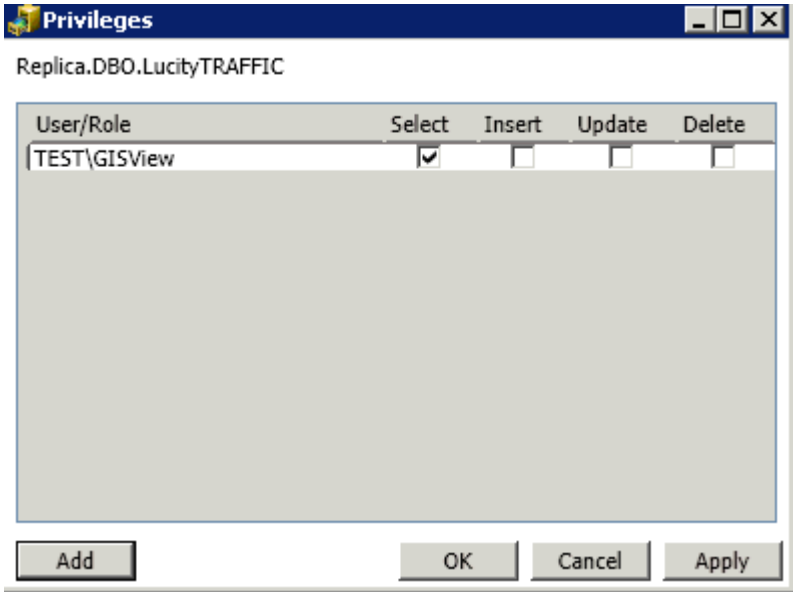

Analyze and Index

Analyze and index is something that you want to do on a regular basis if you creating and adding features to your database. This will allow for better faster access to your database. Make sure ArcGIS for Server services are stopped before rebuilding indexes.

1. Go to your database connection and select on the objects you want to analyze and rebuilt index. Right click on the select items and go to manage/analyze. There is no feel good complete dialog given so let the wheel turn until complete.

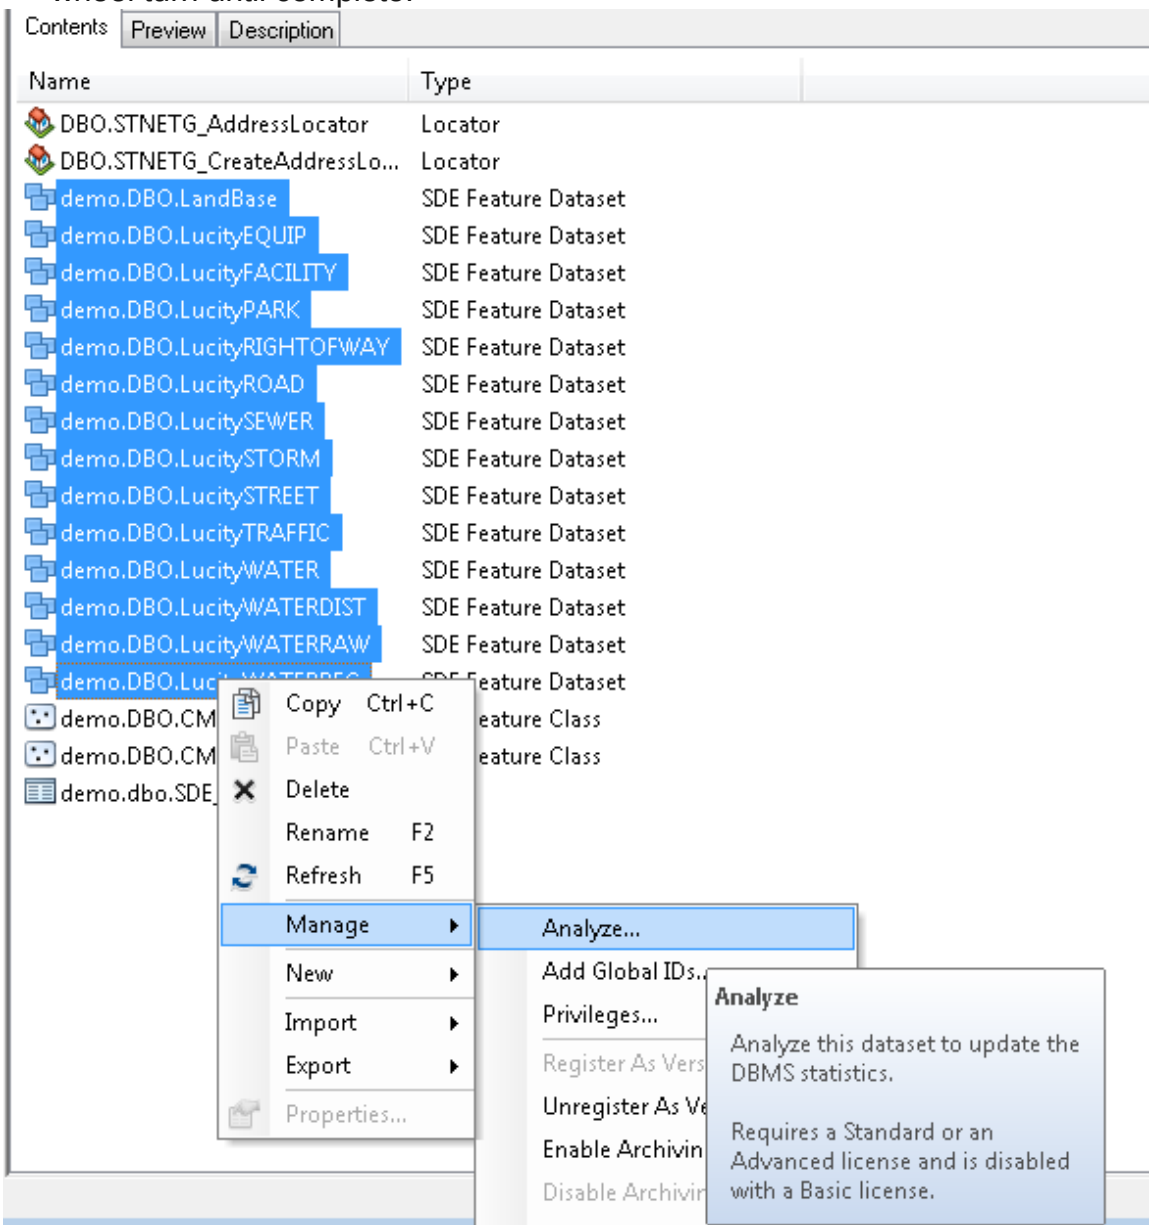

Compression

What is Compression?

Is a process of moving rows of data from the delta (add and delete) tables to the base tables of the feature classes. This is a versioning performance process.

- If your data is in default and you have the data versioned, you should compress.
- Performance and Data Integrity Tip: Analyze statistics and rebuilt indexes before and after compression
- Data Integrity Tip: Compress at least once a week
- Data Integrity Tip: Compress after large loads of data or major changes in data

There are several ways to compress your database. Below are the different ways to compress your database.

1. Using Python
    a. Open ArcToolbox and go to data management tools/Geodatabase Administration/Compress
    b. Open the Python dialog by going to the Geoprocessing Menu and left click on Python
    c. A dialog window will open for Python. Left click on the Compress tool in ArcToolbox and drag the tool within the Python window
    d. Notice that the Python dialog has been updated with the new function

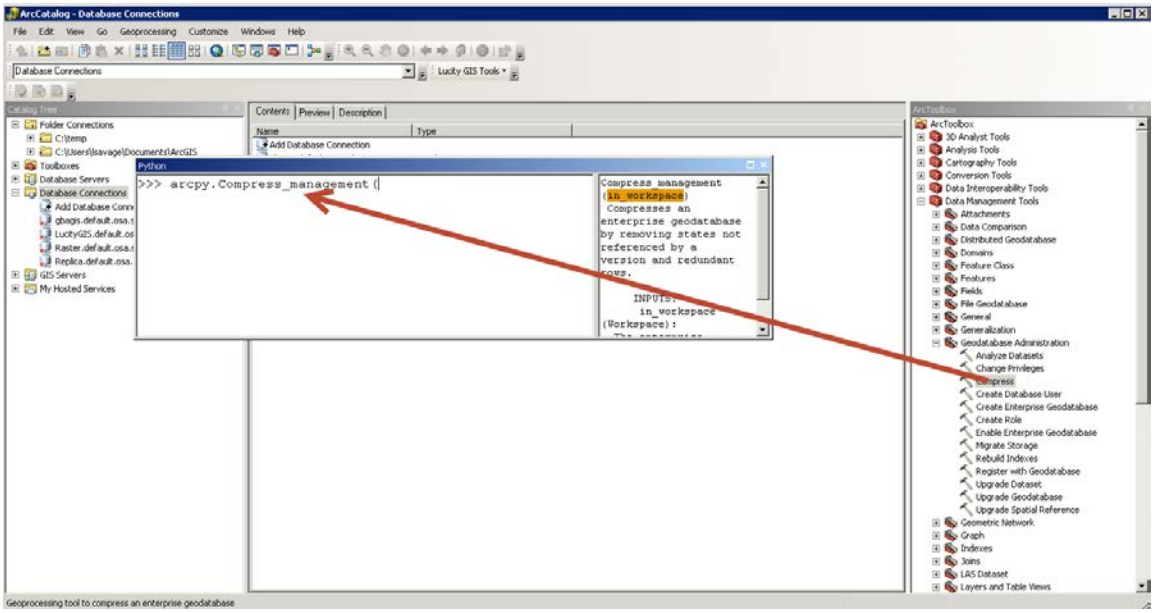

    e. Now open the database connection you would like to compress. Make sure the connection has rights to perform a compress. Drag and drop the connection into the Python dialog

f. Hit the Enter button on the keyboard as this will execute the command. On the right, the dialog will show the start time and how long it took to complete the operation.

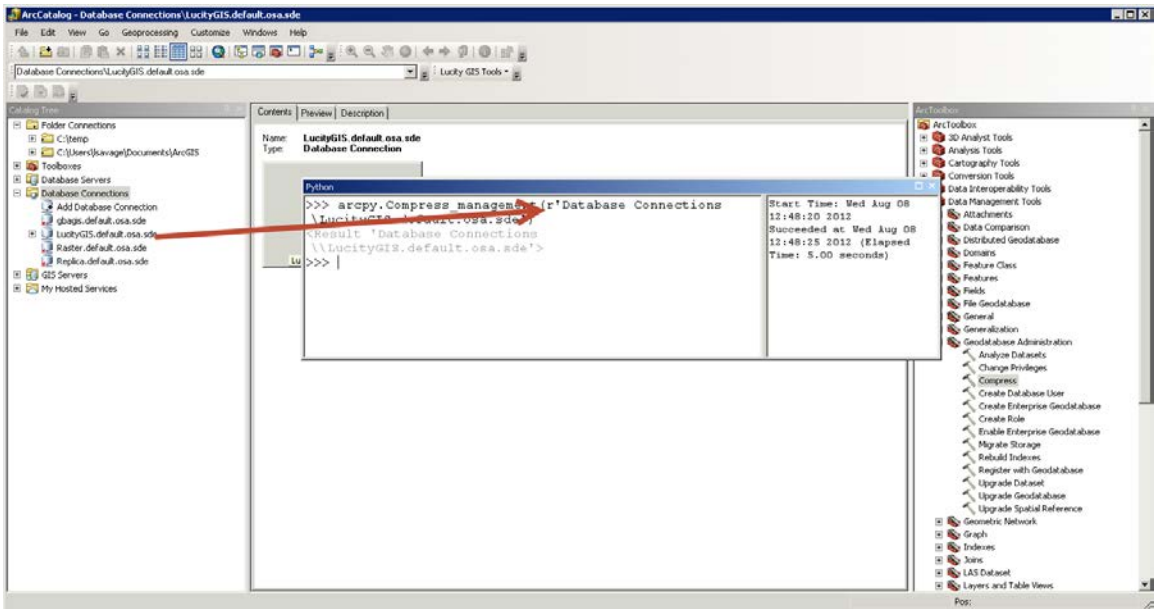

2. Using the ArcCatalog Tool

a. Right Click on a database connection in the ArcCatalog TOC
b. Choose Administration/Compress Database

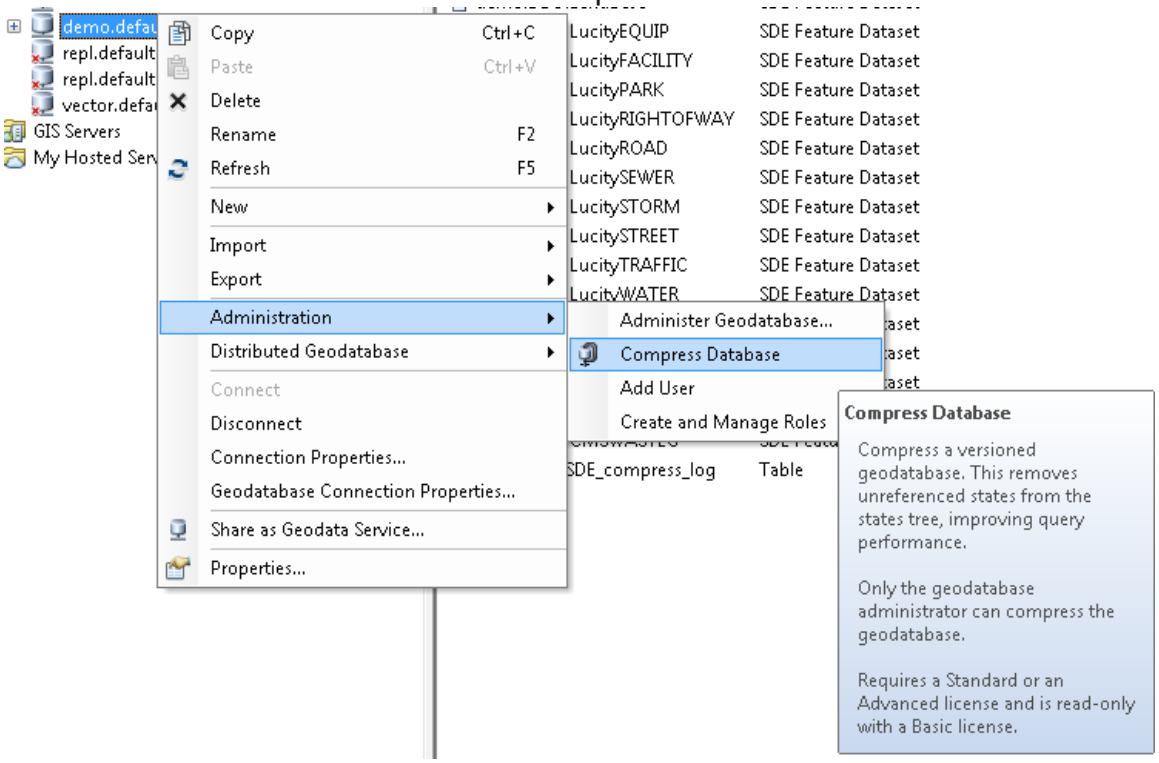

3. Command Line Compress; make sure whatever system connected has the command line installation of ArcSDE installed.
    a. Syntax: sdeversion -o compress -u <sa or admin> -p <password> -i sde:sqlserver:<instancename> -s <server> -D <database> -N

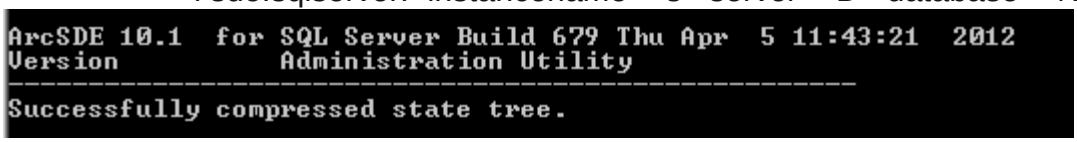

Note: Compress process will only compress data that is in read-only mode. If a feature class is being edited, the compress process will bypass the feature class table and only compress non-edited data. If you are having problems with compression, please reconcile, post all versions. Afterwards, delete the versions and make sure all state locks are terminated as well as all users are disconnected. After all users are disconnected, perform the compression routine. This is better known as "Zero State Compression". Afterwards, recreate your versions and you'll have clean delta tables.

Spatial and Attribute Indexes

In ArcGIS v10.1, spatial indexes are managed for you in a sql table. You can delete and recreate the spatial index from the feature class properties. Spatial indexes allow your spatial data to be queried faster within a grid concept. Each grid has a size based on the density of your spatial data. In the past, this was a calculation that needed to be performed by the GIS DBA or GIS Analyst but is now automated for you. Kudos ESRI!

Spatial Indexes

1. If you are recalculating your feature extent and feel that you want to recreate your spatial index, right click on a feature class and choose properties.
2. When the feature class properties dialog appears, go to the indexes tab.
3. Click Delete and Create under Spatial Index at the bottom of the dialog. Click Apply.

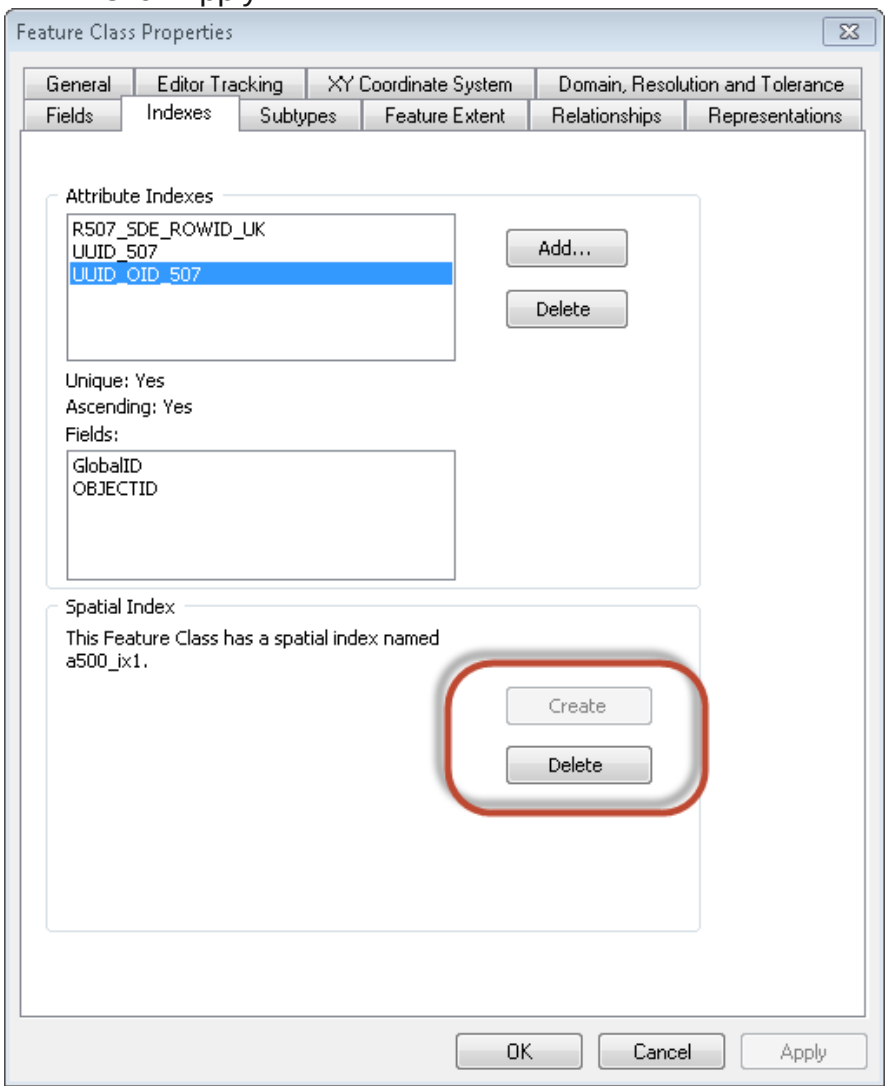

4. This process will recreate the feature class spatial index table for the feature class.

Attribute Indexes

Attribute indexes are used to speed the performance of data queries used on a consistent basis. As an example, we are going to setup an attribute index for Lucity GIS data. Both the common ID and the Lucity AutoID will speed up your query performance capability for each feature class that is linked to Lucity when assigning an attribute index.

5. In the same dialog as the previous exercise (Feature Class Properties), please go to the indexes tab.
6. Click on the Add button next to Attribute Indexes.

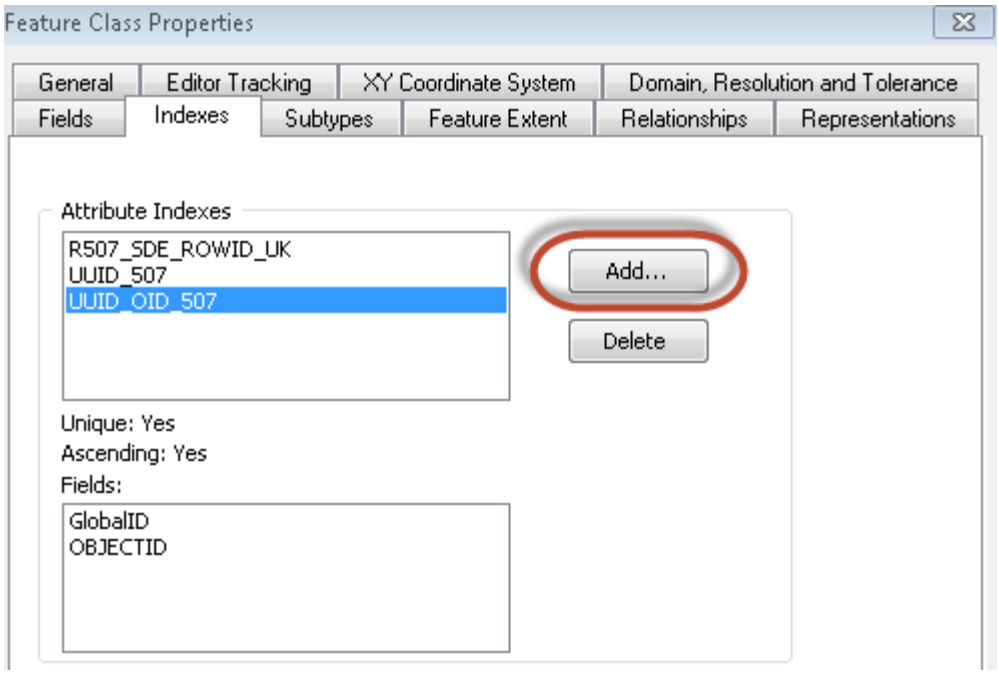

7. Create LucityIDX for the Name and add the NTG_ID which is linked to the Lucity AutoID.

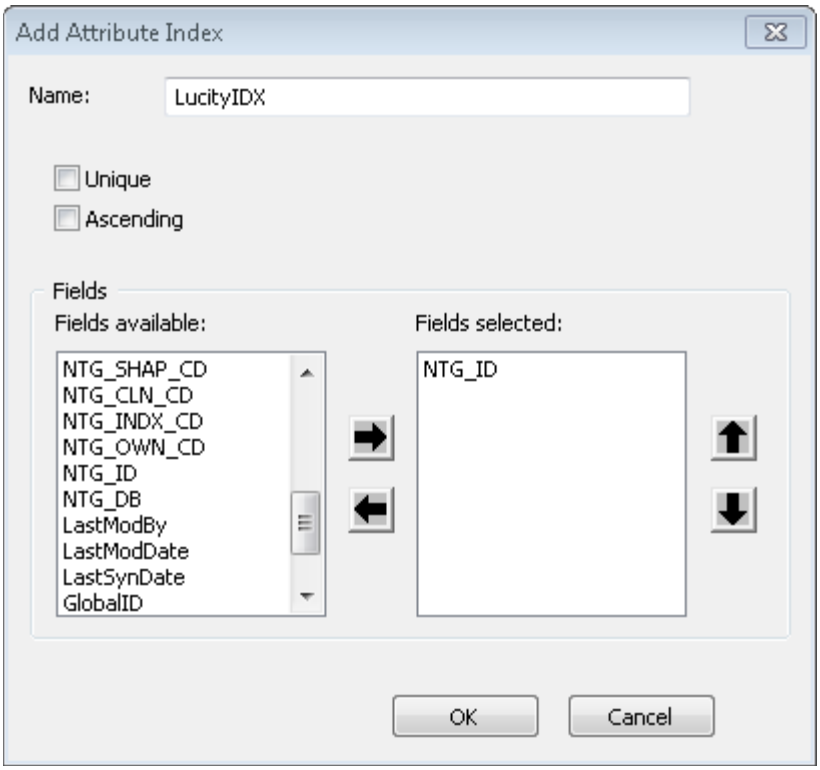

8. Click ok and Apply.
9. Complete the same operation again by adding LucityCMN index and associate it with the field NTG_NUMBER which is linked to the Lucity AltID.

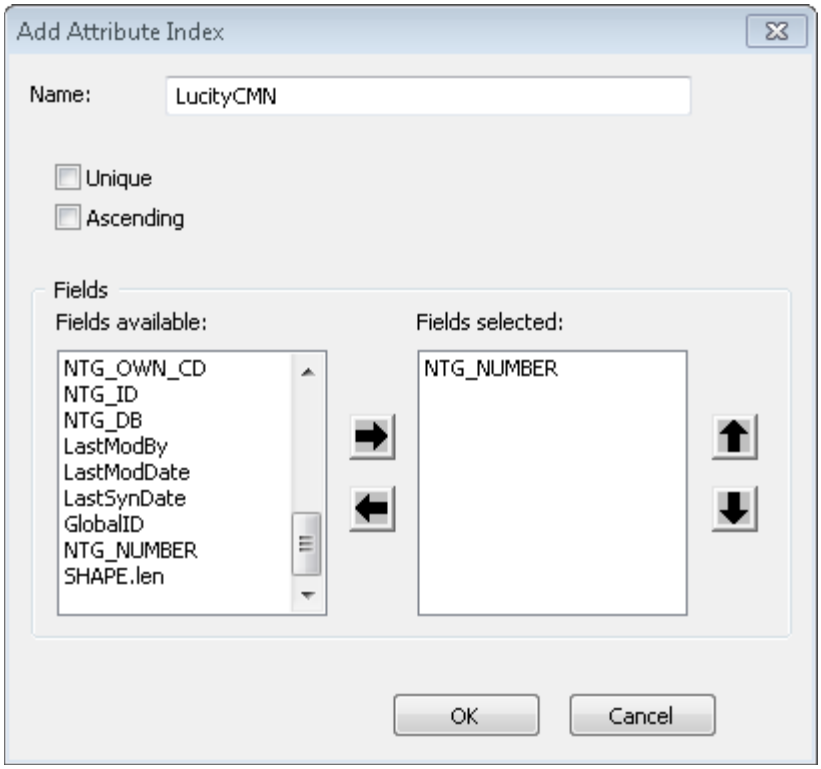

10. Click Ok and Apply.

Tune SQL Server space for Geodatabases

In terms of SQL Server space for Geodatabases, the functionality of the Geodatabases is key to how much space you will need. Recommend splitting up the vector and the raster data into two separate Geodatabases. Raster should be in its own database. This will allow for easy to manage backups and system configuration.

1. When loading Raster data calculate the decompressed size of the image and set this as the file size for the initial load. The logfile can be somewhat small as ESRI has changed the transaction of the loads from one long transaction to iterative raster load transactions. A good rule of thumb, create logfiles for Raster 1/3 of the size of the Raster database. For the data file size, find out the decompression size of the complete mosaic raster data or single image size as a collection. You can choose to use the new mosaic option for raster as this will create a virtual mosaic dataset like the terrain datasets but that's totally up to you as the administrator.

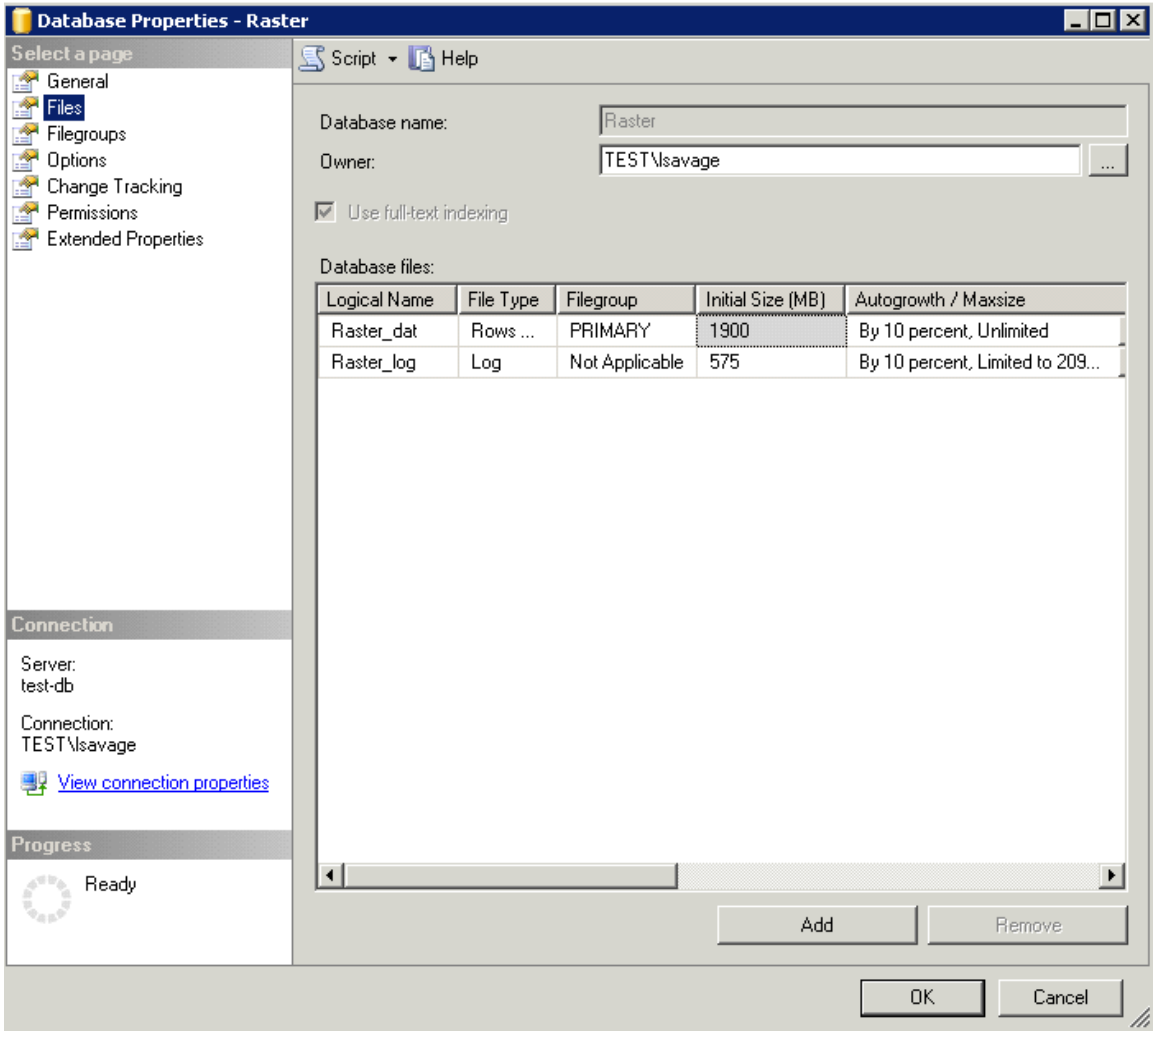

2. For logfiles, don't let the log files use unlimited file growth as a looped transaction could bring your server down.

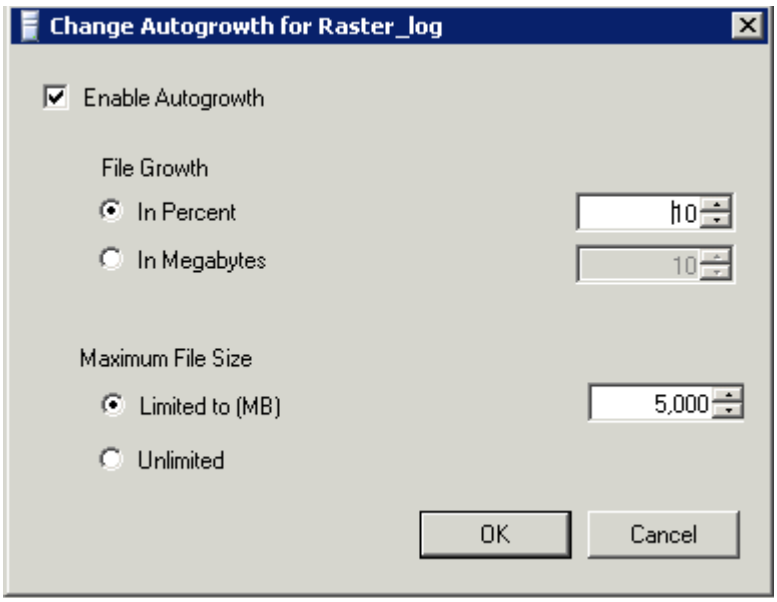

3. Do the same for the data filegroup.

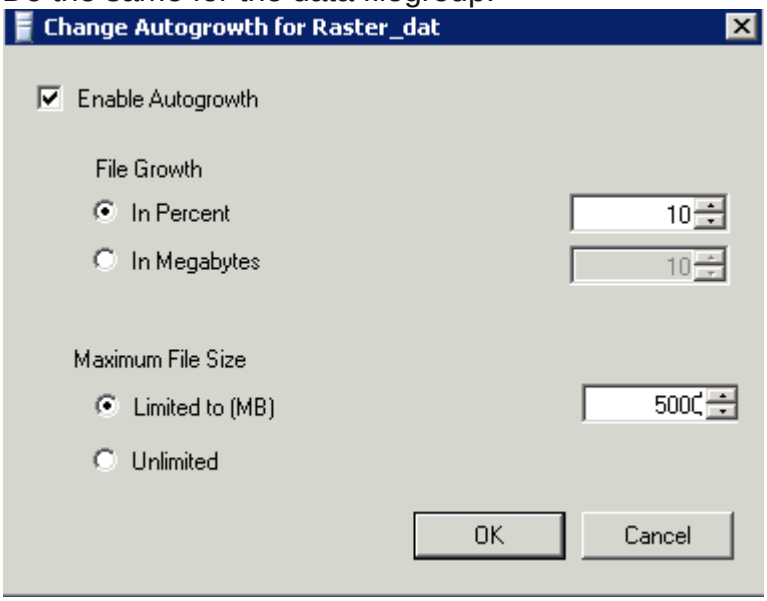

4. Most county Raster datasets are 30 GB to 40 GB once loaded in the Enterprise Geodatabase. You may want to compensate for the increase. It's faster to pre-allocate space for Raster before you load.
5. After loading, analyze through ArcCatalog and assign permissions. Make sure you backup, truncate logfile and shrink the database before making it available to the public.
6. If large loads for vector data or a Geodatabase that is used as an editing Geodatabase, perform similar operations.
7. Set your growth to a manageable amount within your data storage. For Vector, you can use percentage. For Raster, you will be adding Raster data per load and sometime every year to three years at a time.

Backups

If you want to create backups using the SQL Server maintenance wizard, this is an easy step by step process. However, in our example below, we will be touching on creating backups from SQL Script, using SQL Commands and windows task manager. Command line seems a bit quicker when running through each operation. Below is the information about each batch script. In order to run the complete procedure below, you will need to have each database backup recovery model setting turned to 'Full'. Simple model is good for a development environment but only is as good as the last backup. For cleanup, please use a file copy and delete batch routine to reset the backups as backups will not overwrite in sql script. You can use a simple copy and delete command in command line.

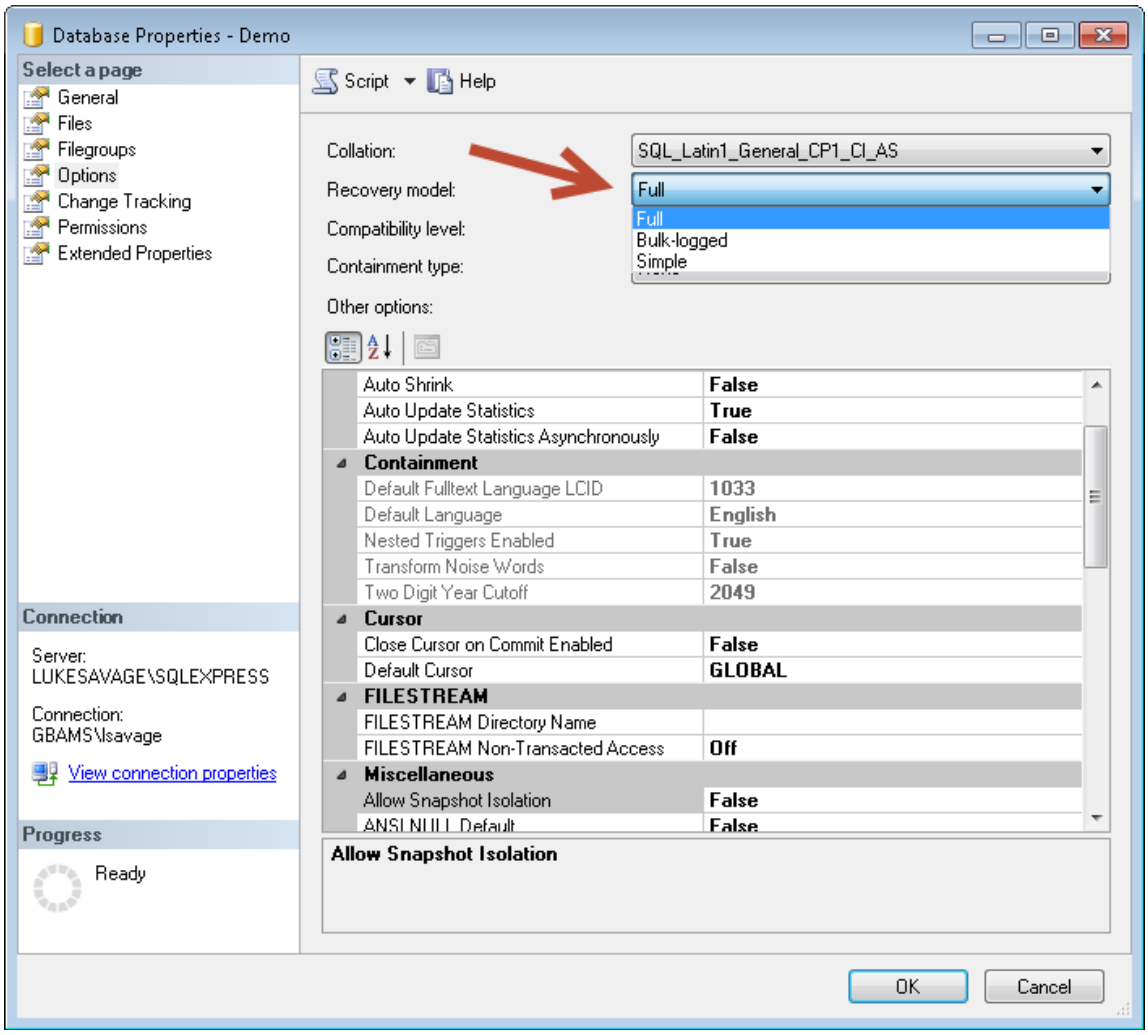

Full Backup: This will backup the entire database.
Intermediate Backups: Captures changes to the database since the last full backup (also known as differential backups.
Truncate Logs: Truncates logs to free up space
Shrink: Shrinks data and log files.
Analyze and Index: Analyze/update statistics, and updates indexes

1. Create a full backup script. For every database, create a separate action.

```
BACKUP DATABASE [DEMO] TO  DISK =
N'C:\backup\SQLEXPRESS\DEMO.bak' WITH NOFORMAT, NOINIT,
RETAINDAYS = 13,  NAME = N'DEMO-Full Database Backup', SKIP,
NOREWIND, NOUNLOAD,  STATS = 10
GO
```

2. Create an intermediate or differential script.

```
BACKUP DATABASE [DEMO] TO  DISK =
N'C:\backup\SQLEXPRESS\DEMO.bak' WITH DIFFERENTIAL, NOFORMAT,
NOINIT, RETAINDAYS = 13,  NAME = N'DEMO-Full Database Backup', SKIP,
NOREWIND, NOUNLOAD,  STATS = 10
GO
```

3. Create a backup Log script.

```
BACKUP LOG [DEMO] TO  DISK = N'C:\backup\SQLEXPRESS\DEMO.bak'
WITH NOFORMAT, NOINIT, RETAINDAYS = 13,  NAME = N'DEMO-
Transaction Log  Backup', SKIP, NOREWIND, NOUNLOAD,  STATS = 10
GO
```

4. Create a shrink script.

```
USE [DEMO]
GO
DBCC SHRINKFILE (N'DEMO_Log' , 0, TRUNCATEONLY)
GO
```

5. Create an update stats and index script.

```
EXEC sp_updatestats
go
EXEC sp_MSforeachtable @command1='print "?"', @command2='set
QUOTED_IDENTIFIER ON;ALTER INDEX ALL ON ? REBUILD WITH
(FILLFACTOR=90,ONLINE=OFF)'
```

We will be calculating an index of 90% capacity for fillfactor (which gives 10% growth for data indexes). Also, we will retain the backups for 13 days and starting over on the 14th day. For each script we will create a batch file to call the .sql files.

6. Create batch routines under root of c:\.
    - Full Backup batch file (example: full_backup.bat)

```
cd\
sqlcmd -S lukesavage\sqlexpress -i c:\backup\SQLEXPRESS\full_backup.sql -o
c:\backup\SQLEXPRESS\full_backup_log.txt
exit
```

    - Intermediate or differential backup (example: intermediate_backup.bat)

cd
sqlcmd -S lukesavage\sqlexpress -i c:\backup\SQLEXPRESS\intermediate_backup.sql -o c:\backup\SQLEXPRESS\inter_backup_log.txt
exit

- Backup Logs and Truncate (example: backup_logs.bat)

cd
sqlcmd -S lukesavage\sqlexpress -i c:\backup\SQLEXPRESS\truncate_logs.sql -o c:\backup\SQLEXPRESS\truncate_log.txt
exit

- Shrink Data and Log files (example: shrink.bat)

cd
sqlcmd -S lukesavage\sqlexpress -i c:\backup\SQLEXPRESS\shrink.sql -o c:\backup\SQLEXPRESS\shrink_log.txt
exit

- Update Stats and Index batch file (example: update_stats_index.bat)

cd
sqlcmd -S lukesavage\sqlexpress -d Demo -i c:\backup\SQLEXPRESS\update_stats_index.sql -o c:\backup\SQLEXPRESS\VECTOR_update_stats_index_log.txt
exit

7. Create Windows Scheduled tasks for each backup file.
    - Full Backup: running once a week on a Saturday Early Morning
    - Differential or Intermediate backup: Running once a day except Saturday at the same time as the Full Backup
    - Backup Logs: Run once a day or twice a day depending on your workload and editing. Make sure you try to backup logs during lower workload hours.
    - Shrink Data and Logs: Run right after the backup log routines
    - Analyze, update statistics and indexes: Run once a day and/or before and after backups. Most will run this once a day after backups occur.
    - Create backup copy and cleanup procedures using command line scripts or delete utilities that can be scheduled

Automated Maintenance Scripts

With the release of ArcGIS v10.1, you have the capability to create python scripts to analyze, update statistics, and update index files. However, in this example we will show how to use ArcSDE and SQL Server command line.

SQL Server and ArcSDE Scripts

1. Install ArcSDE with Command line tools (**ArcGIS 10.2 is the last version to include command line tools. Post 10.2, plan to write automation scripts in python**)
2. Go to SQL Server Studio Management and open designer.
3. Create query. In this query, it will analyze and update statistics and rebuild indexes with a fill factor of 90%, which means you have 10% growth for indexes.

EXEC sp_updatestats
go
EXEC sp_MSforeachtable @command1='print "?"', @command2='set QUOTED_IDENTIFIER ON;ALTER INDEX ALL ON ? REBUILD WITH (FILLFACTOR=90,ONLINE=OFF)'

4. Save query as update_stats_index.sql in the root of c:\
5. Create a Batch file that contains the following script. Notice the output log file created so you can keep track of your operations and if they fail troubleshoot the issue effectively. Items in Red signify zero state compression where all data is pushed to base tables. If zero state compression is not needed, delete the lines in Red below.
6. Save as <name>_state_compression.bat in the root of c:\.

Note: In a versioned environment, recommend zero state compression at least once every 6 months for lightweight editing. Heavy data editing environment zero state compression should be accomplished every quarter. Normal compression should be taking place once a week. This is to ensure data integrity long term.

sdemon -o kill -t all -p <sa password> -i sde:sqlserver:<instance> -s <server name> -D <geodatabase1> -N
sdemon -o kill -t all -p <sa password> -i sde:sqlserver:<instance> -s <server name> -D <geodatabase2> -N
sqlcmd -S localhost -U sa -P <sa password> -D <geodatabase1> -i c:\update_stats_index.sql -o c:\update_stats_preindex_log.txt
sqlcmd -S localhost -U sa -P <sa password> -D <geodatabase2> -i c:\update_stats_index.sql -o c:\update_stats_preindex_log.txt
sdeversion -o delete -i sde:sqlserver:<instance> -V dbo.edit -u sa -p <sa password> -s <server name> -D <geodatabase1> -N
sdeversion -o compress -u sa -p <sa password> -i sde:sqlserver:<instance> -s <server name> -D <geodatabase1> -N
sdeversion -o compress -u sa -p <sa password> -i sde:sqlserver:<instance> -s <server name> -D <geodatabase2> -N
sqlcmd -S localhost -U sa -P <sa password> -D <geodatabase1> -i c:\update_stats_index.sql -o c:\update_stats_index_postlog.txt
sqlcmd -S localhost -U sa -P <sa password> -D <geodatabase2> -i c:\update_stats_index.sql -o c:\update_stats_index_postlog.txt

sdeversion -o create -i sde:sqlserver:<instance> -V edit -P dbo.default -A public -u sa -p <sa password> -s <server name> -D <geodatabase1>
sqlcmd -S localhost -U sa -P <sa password> -D <geodatabase1> -i c:\update_stats_index.sql -o c:\update_stats_index_log.txt

7. Add a windows task in windows task scheduler on the server where ArcSDE binaries and command line tools reside.

Automate Replica Synchronization using Python

8. Go to ArcCatalog and open model builder. This is the Garfield way of creating python scripts for the non-savvy coder.

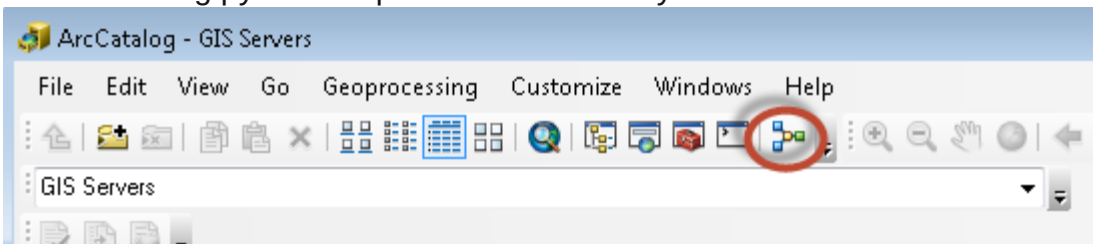

9. In Model Builder, open arctoolbox if not already open and go to Distributed Geodatabases/Synchronize Changes

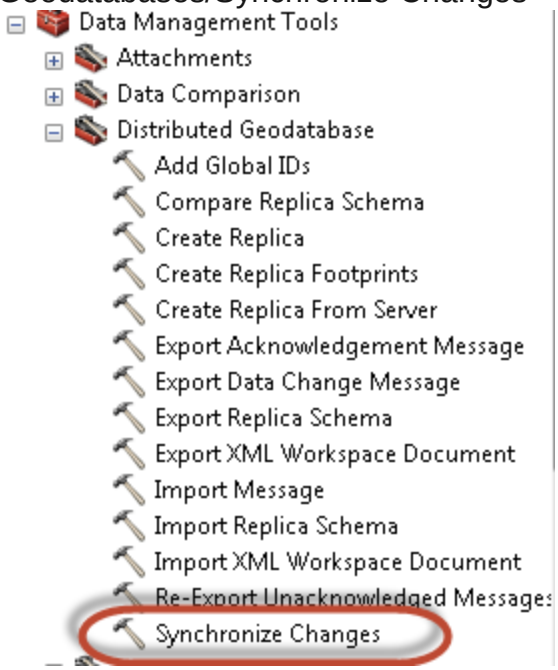

10. Drag and drop the tool into Model Builder.

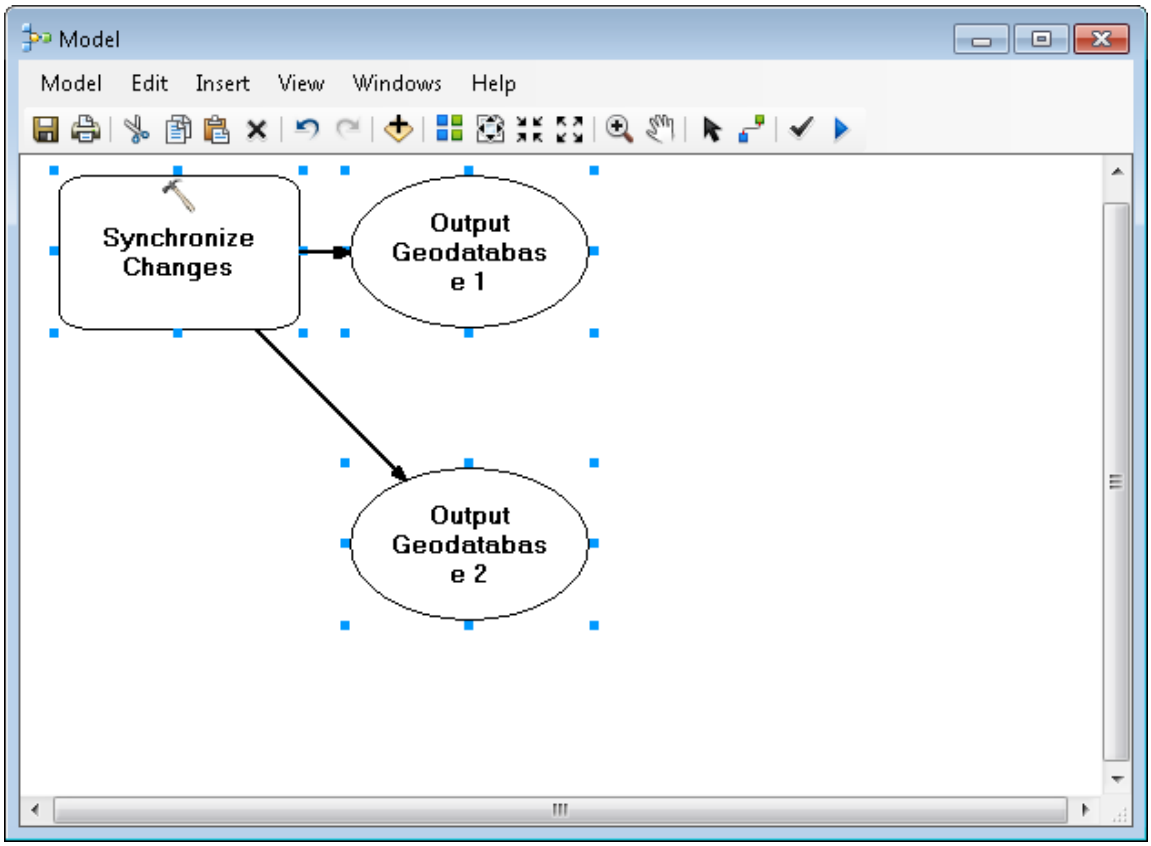

11. Once tool is in Model Builder, double click on Synchronize Changes box to open the dialog.
12. Fill in the information appropriately. For reference, please visit the Replica section.

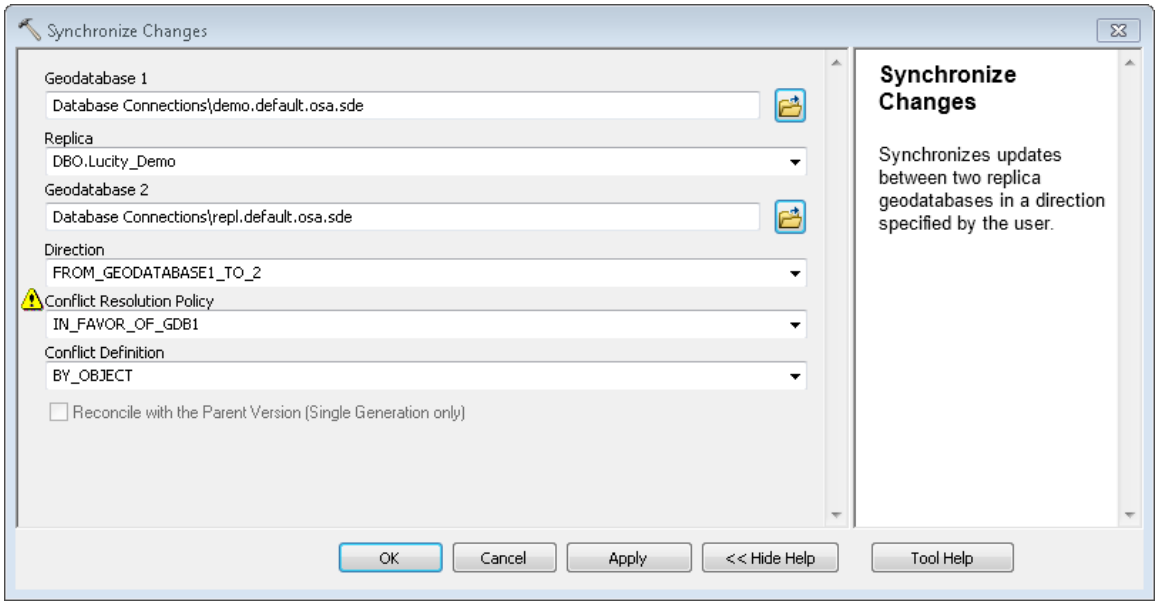

13. Click ok. Notice colors change in the dialog. Run the model to test by clicking on the run button.

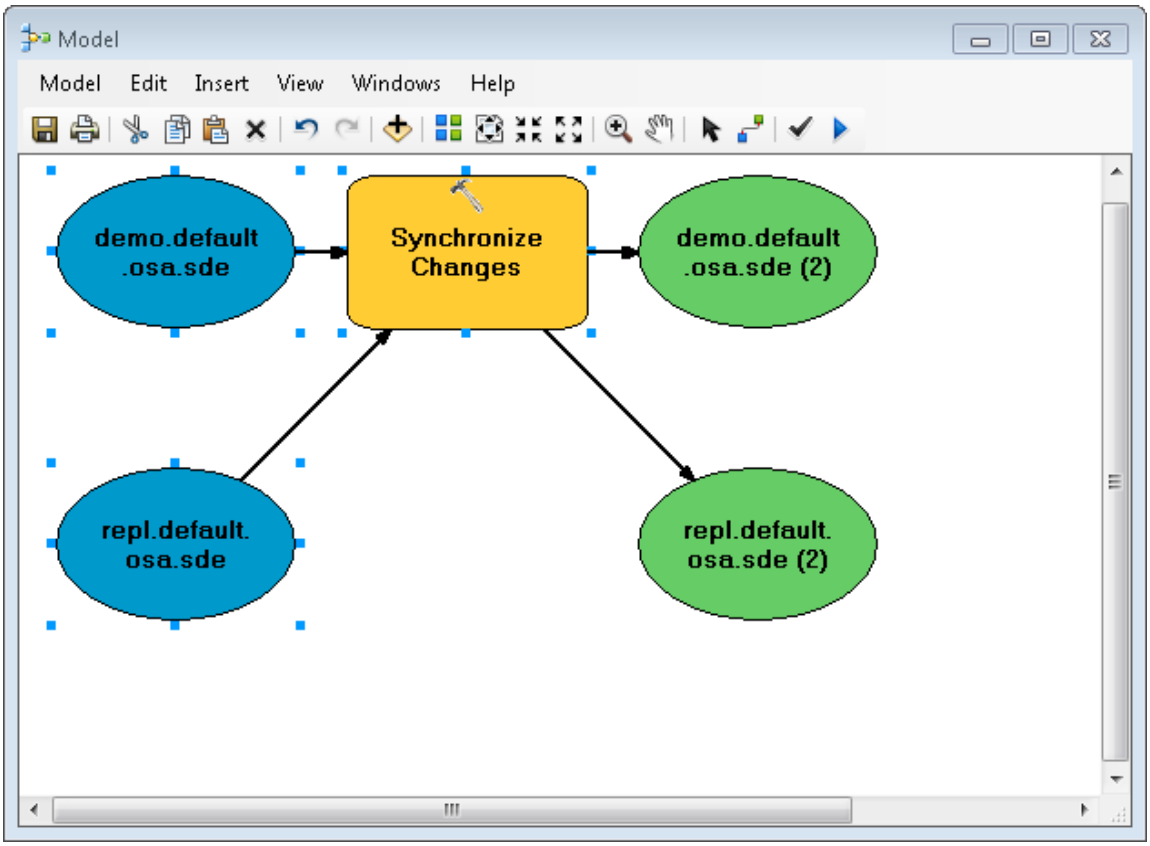

14. After successfully testing the model, go to the model menu and click on export/To Python Script.

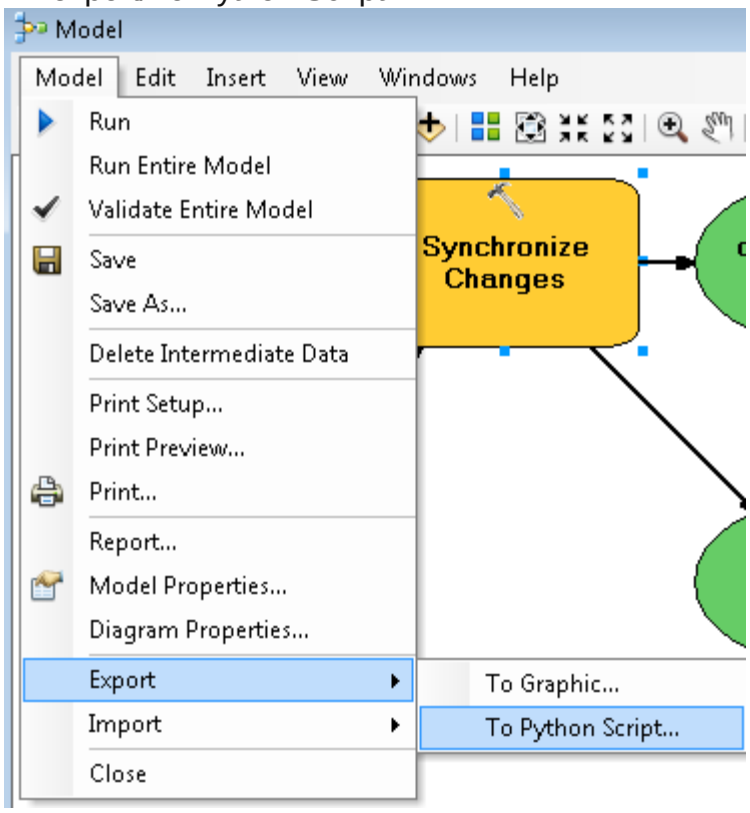

15. Save in location that has Python and ArcCatalog installed. For example, auto_synch.py under the root of c:\.
16. Edit the Python script so that it references arcinfo rather than arceditor. This can produce errors in your script. Also, add the RED content below to enable logging when running this script.

```
# Set the necessary product code
import arcinfo
# Import arcpy module
import arcpy
# Record Logs
arcpy.SetLogHistory(True)
# Local variables:
demo_default_osa_sde = "Database Connections\\demo.default.osa.sde"
repl_default_osa_sde = "Database Connections\\repl.default.osa.sde"
# Process: Synchronize Changes
arcpy.SynchronizeChanges_management(demo_default_osa_sde, "DBO.Lucity_Demo",
repl_default_osa_sde, "FROM_GEODATABASE1_TO_2", "IN_FAVOR_OF_GDB1",
"BY_OBJECT", "DO_NOT_RECONCILE")
```

17. Go to Windows Task Scheduler and create a basic task.

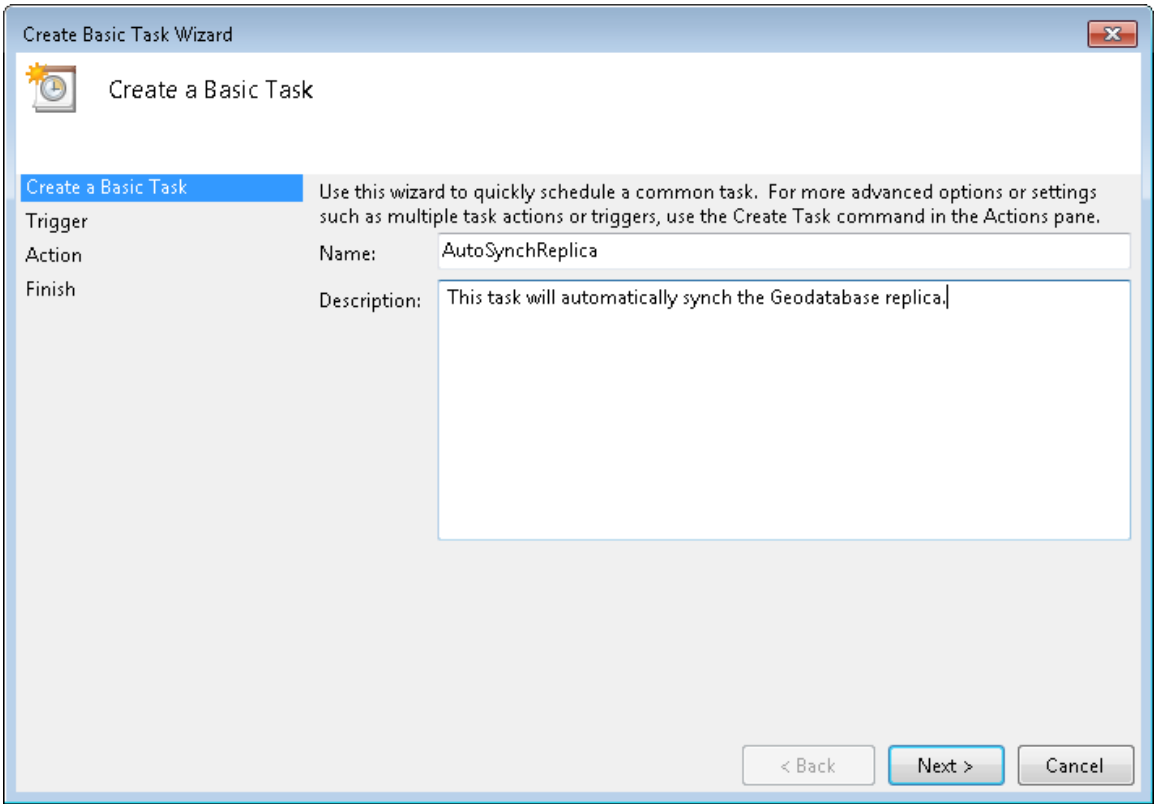

18. Click next and schedule the synchronization of the replica to whenever you want the synchronization to take place. Usually each day or once a week will fit most organizations. In this example we are going to setup a weekly schedule. Click next.

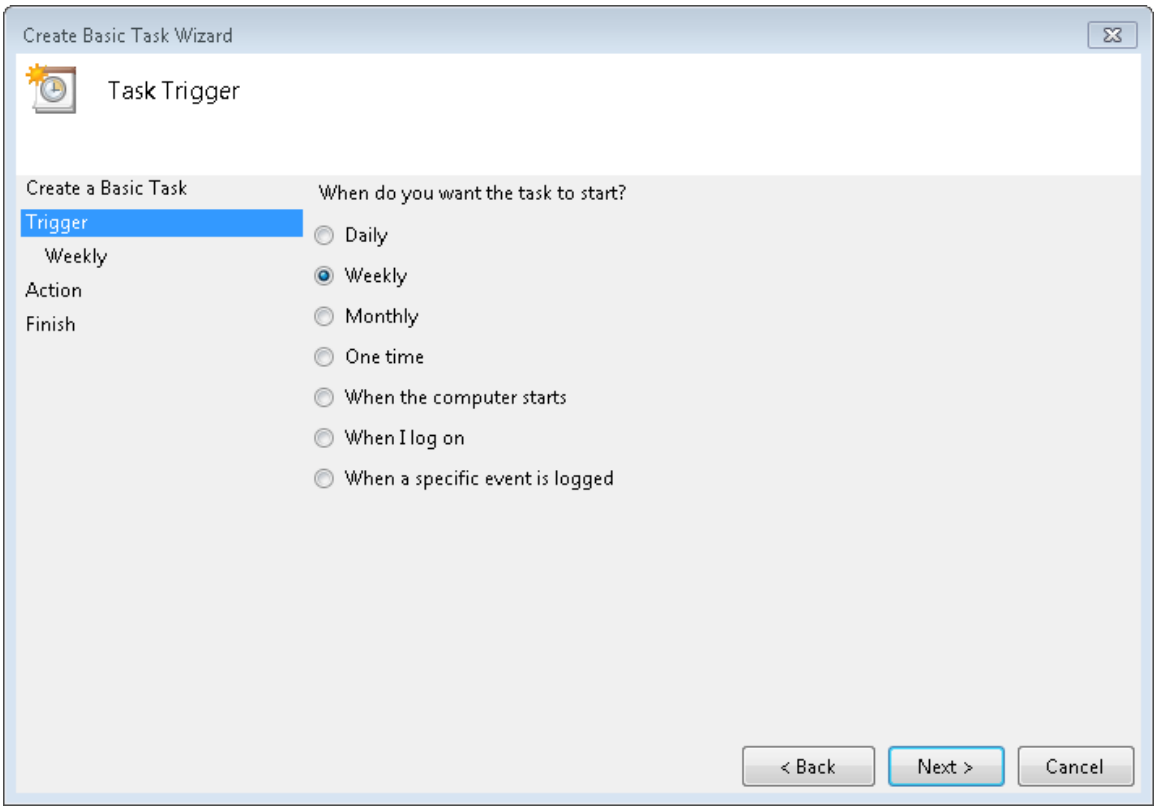

19. In the timeliness of this schedule, I want this to fire off before the backups have occurred so that my databases stay in tune. In my backup routines, I may have a performance maintenance script that analyzes, updates statistics and indexes as well as backups the transaction logs and shrinks the database. Therefore, in this example we'll set the synchronization to occur after I leave for the day on Friday. Click next.

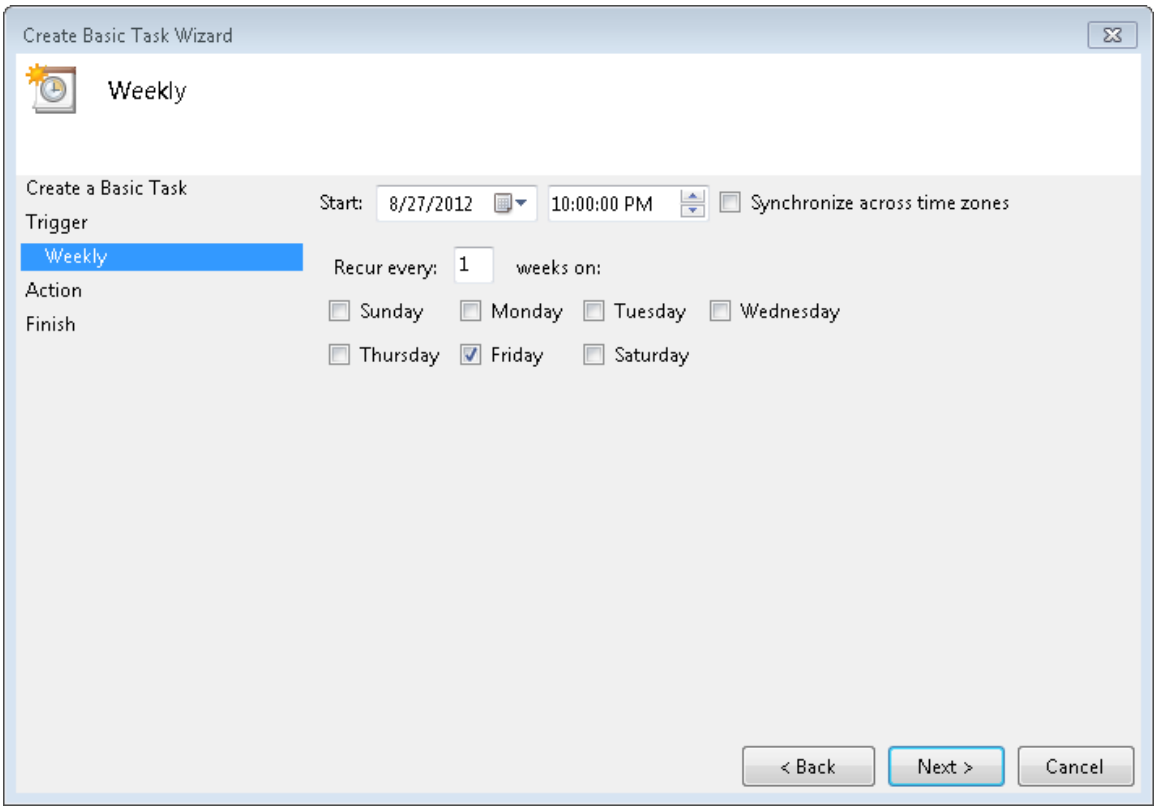

20. Choose start a program and click next.

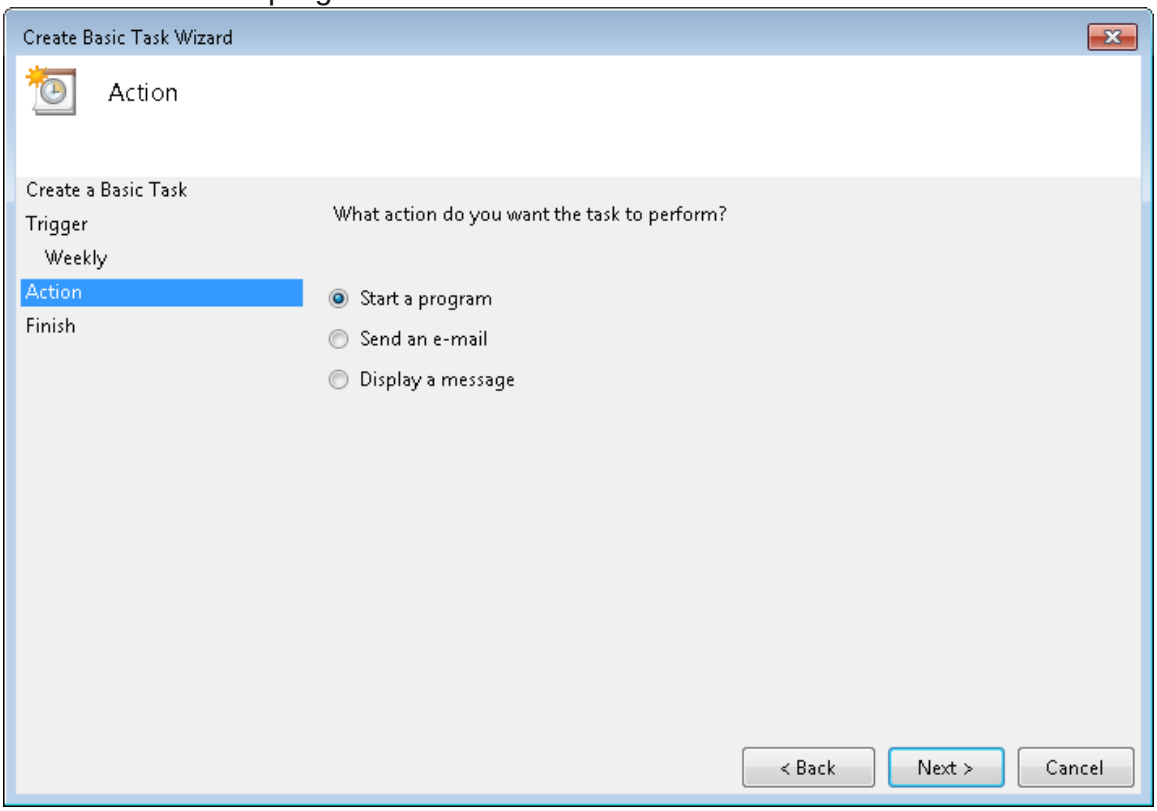

21. In Windows Vista and above, we need to add the python program as the script so that it knows what executable to use. In Add arguments

(optional):, you need to add the location of your saved python script. Also, in 'Start in (optional), you need to add the location folder of the python executable. Once complete, click next.

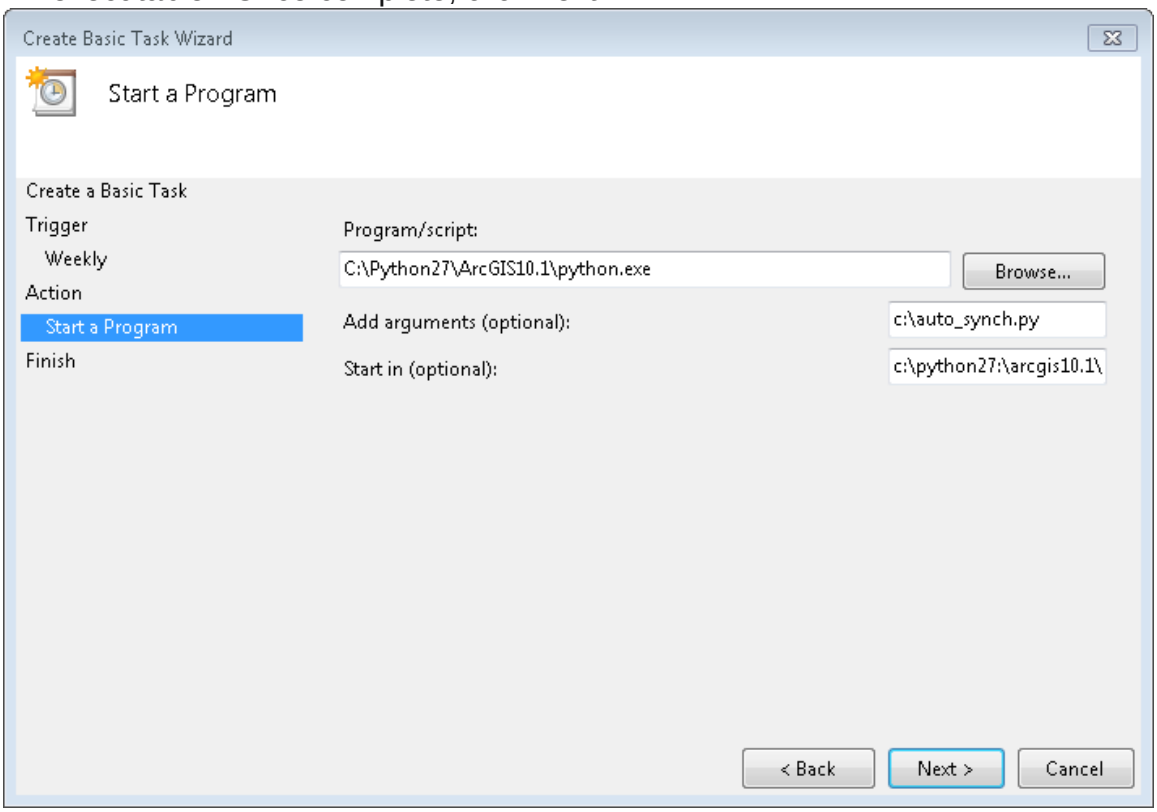

22. Click the 'Open the Properties dialog for this task when I click Finish' checkbox to verify administrative user to run this task and its credentials. Click Finish.

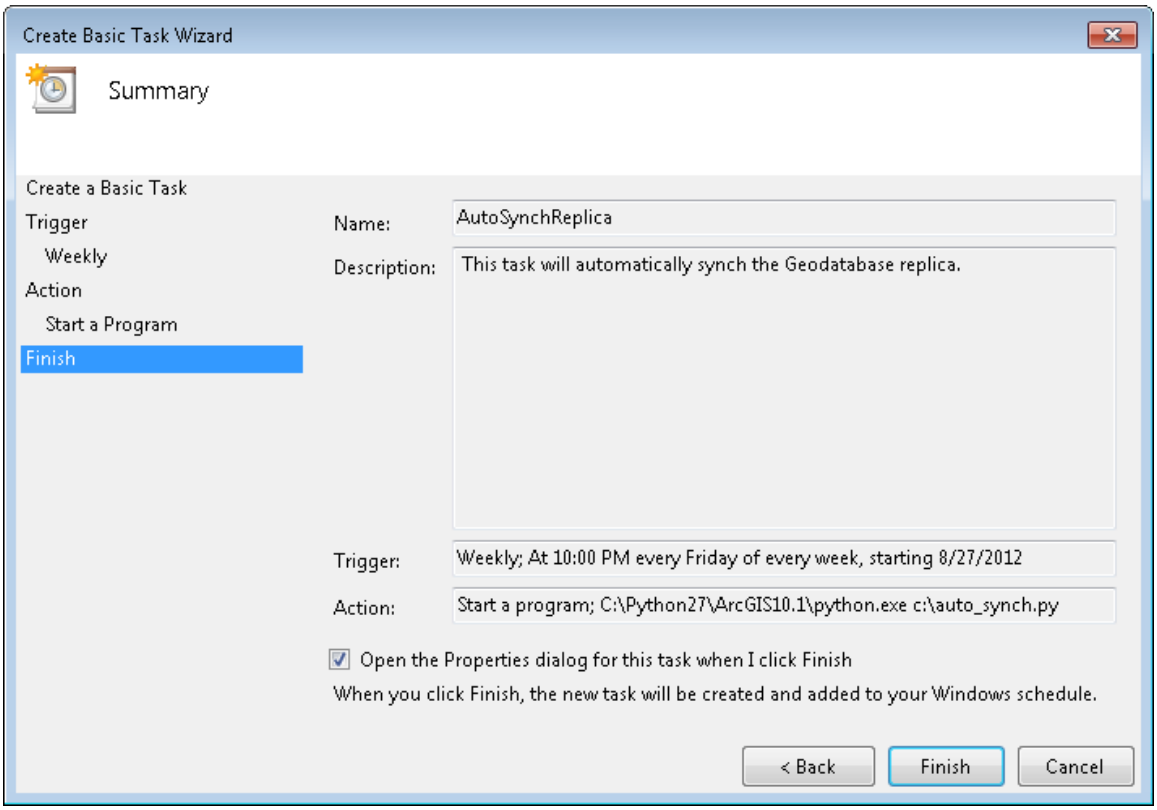

23. Check the ‘Run whether user is logged on or not’ radio button and check the checkbox next to ‘Run with highest privileges’.

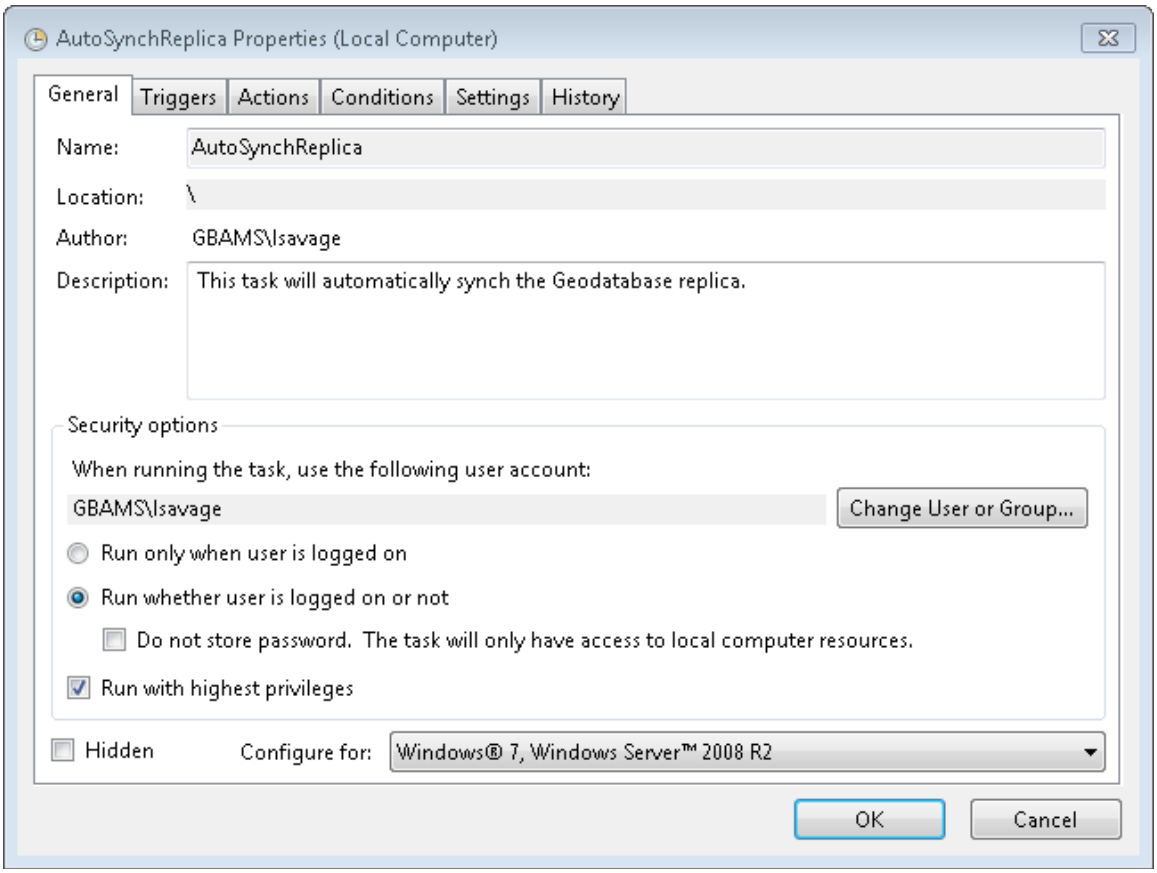

24. Test the task by right clicking on the new task and select run.
25. Check the ArcToolBox results history location by going to the following location.

C:\Users\<username>\AppDate\Roaming\ESRI\Desktop10.1\ArcToolbox\History

Options

Versioning

Is an alternative state of the database where you can make edits and changes that will not affect the base tables. When complete and edits are ready, the parent table will be reconciled and posted with the child version. Within the Geodatabase, there are two tables that store changes to the base data. These are called delta tables which are known as Add and Delete tables.

What are A and D tables?

- A tables are add tables. Anytime you add a record or change a record, the changes are added to the A table.
- D tables are delete tables. Anytime you delete a record, the delete rows are added to the D table.
- A and D tables are numbered based on the registration_id in the SDE_Table_Registry table.

Example of finding an Add Table.
If you have made an edit to a feature class, you need to get the registration_id for that feature class and go to the appropriate A table. This is stored in the sde_table_registry in SQL Server. For child versions, they are subsequent IDs such as default version would be the actual registration_id (A144) and the child version would be (A145). Same applies to the D tables.

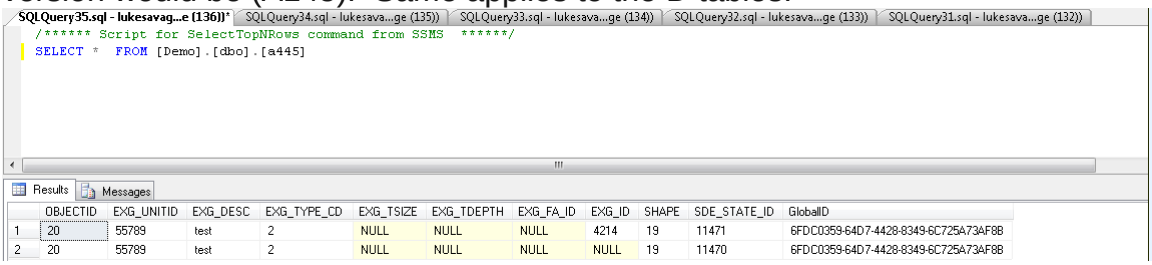

When Does Versioning Make Sense?

The option to moving edits to base tables is a scary thing for GIS people. Why?

- You can edit simple data only—points, lines, polygons, annotation, and relationships. You cannot edit a feature class in a topology, geometric network, or terrain.
- You cannot archive changes for the dataset.
- You cannot replicate the dataset.
- When you edit the DEFAULT version or post a version to the DEFAULT, you do not have the ability to resolve conflicts, so it is possible to overwrite another user's edits.

How to Create a Version

1. Right click on a Database Connection link in ArcCatalog TOC or right click in the white space in the contents tab in ArcCatalog

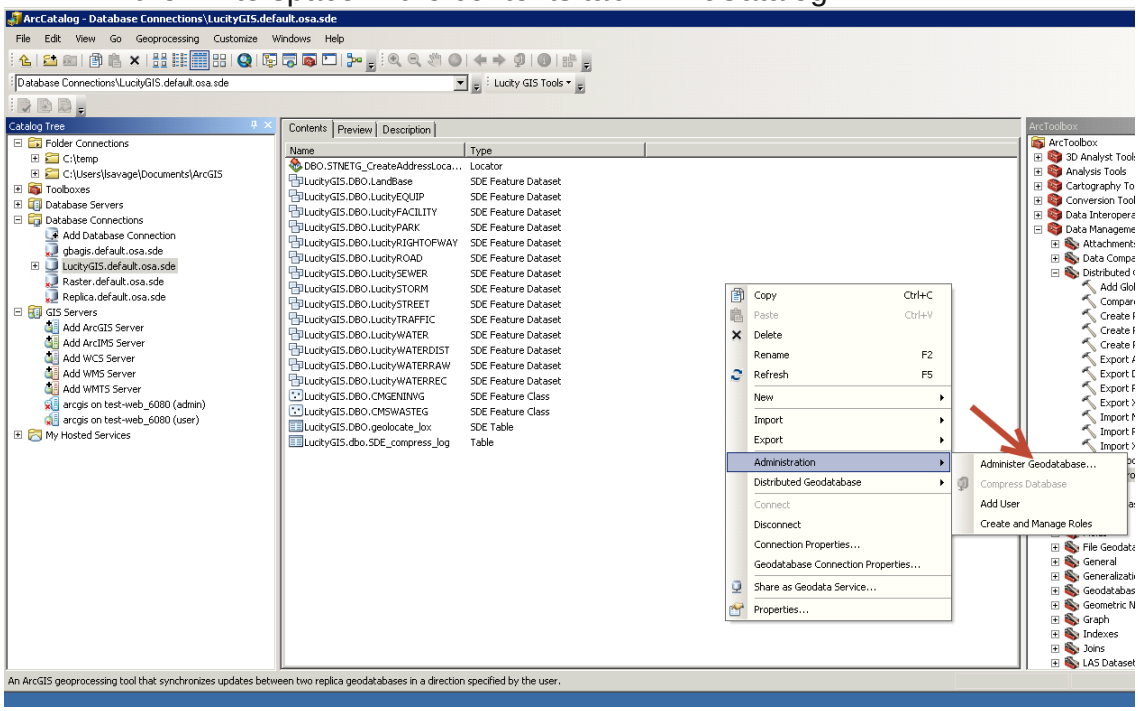

2. Click on Administration/Administrator Geodatabase.
3. Right click on the Default version and select new version

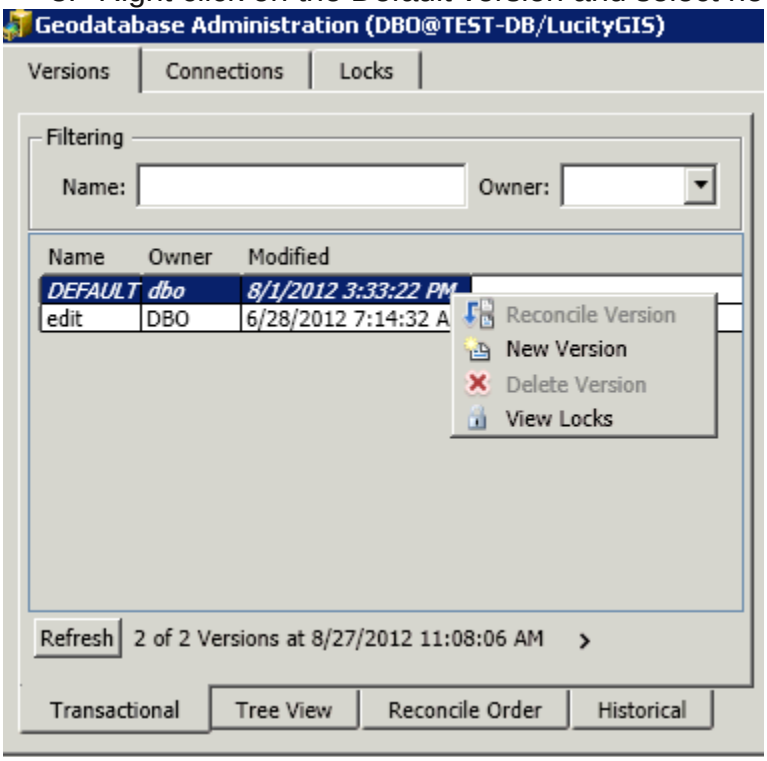

4. Create a Name for the new version and select public. Click Ok.

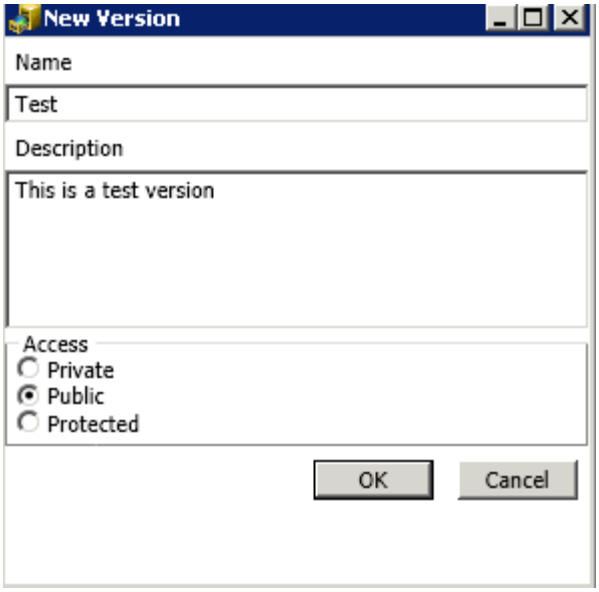

5. What's the difference between Access types
    - Private: Created user eyes only; no one has access unless owner of the version
    - Public: Everyone can see the version and create versions from. If you have edit privileges, you can edit the version.
    - Protected: Everyone can see the version but cannot edit the version unless owner. Everyone can create versions from.
6. This is now the child of Default and is listed in the Geodatabase Administration dialog.

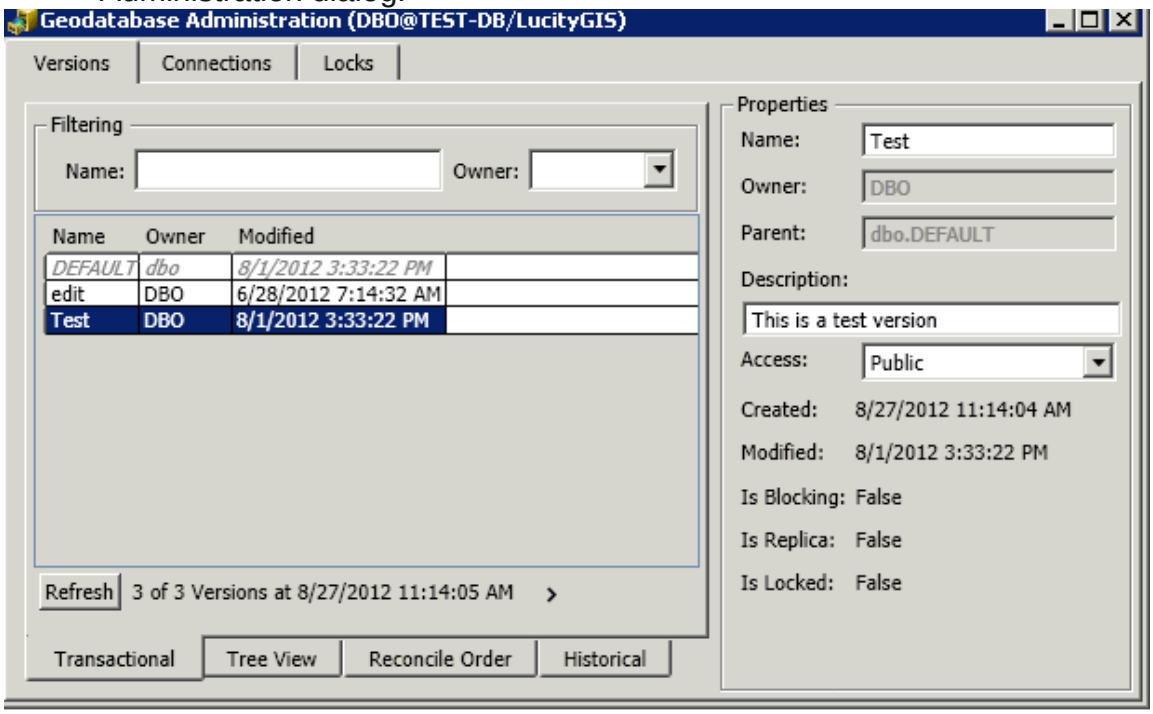

Replica

A replica is a copy of a database using GUID attributes that synchronize from the parent to the child databases or vice versa. There are several ways to create a replica but we will be creating the most likely replica used which is a one-way replica. One way replicas are easy to setup and are handy for the administrators to offload view only users off of the production database while performing synchronizations from time to time. Keeping the database in synch when production changes have been approved is a nice way to distribute your data to users while retaining data integrity, security and flexibility.

1. Prepare your data for a replica. Go to your database and right click on Feature Datasets, and go to manage/add global ids. Global Ids are necessary for the Geodatabase to keep track of changes in each database that is a part of a replica.

Note: Data must be registered as versioned

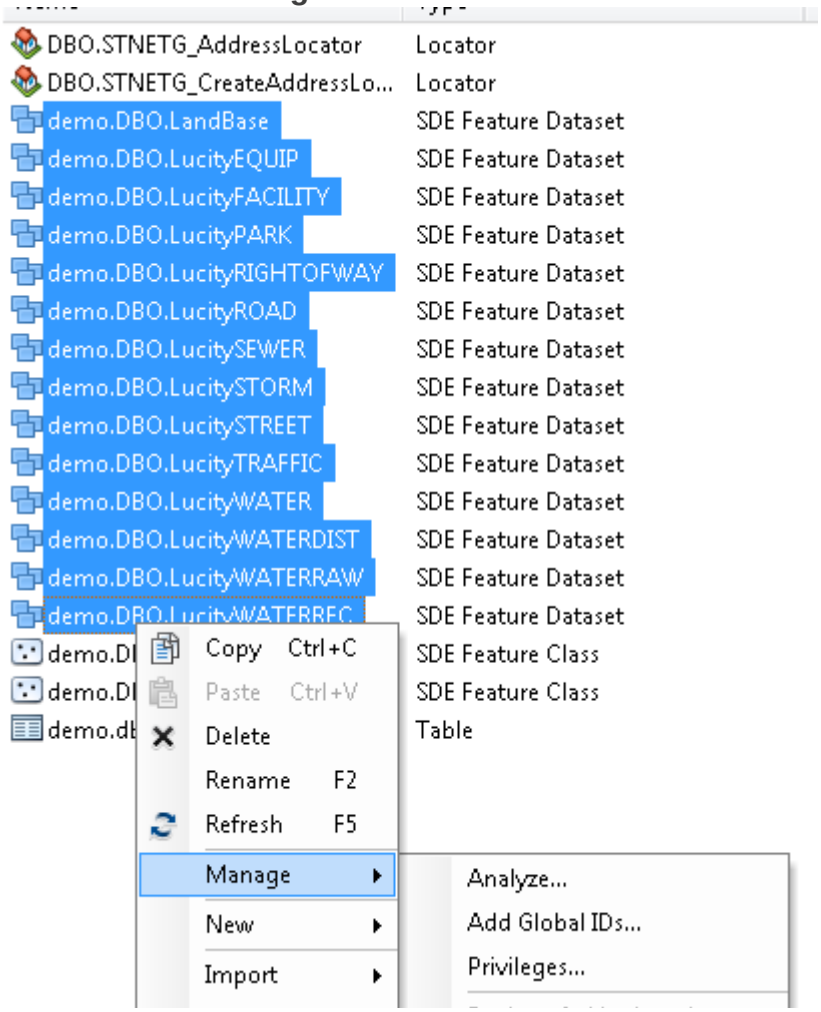

2. Do this for stand-alone Feature Classes as well.
3. To create a replica, go to ArcToolbox/Data Management Tools/Distributed Geodatabase and expand.
4. Click on Create Replica

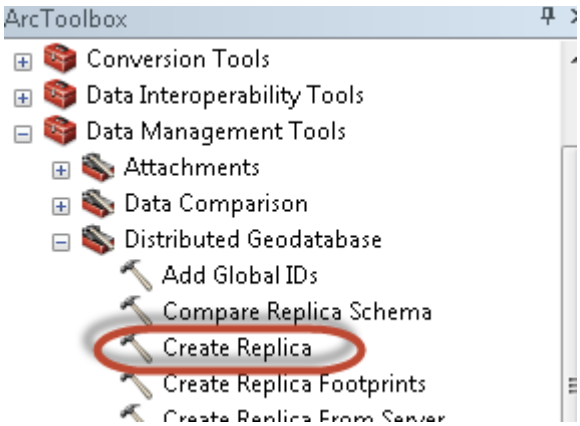

5. Browse or drag and drop data you want to replicate from into the Create Replica dialog window. Change the Replica Type to one-way replica. Add the replicated database connection to replicate the data to. Give the Replica a name.

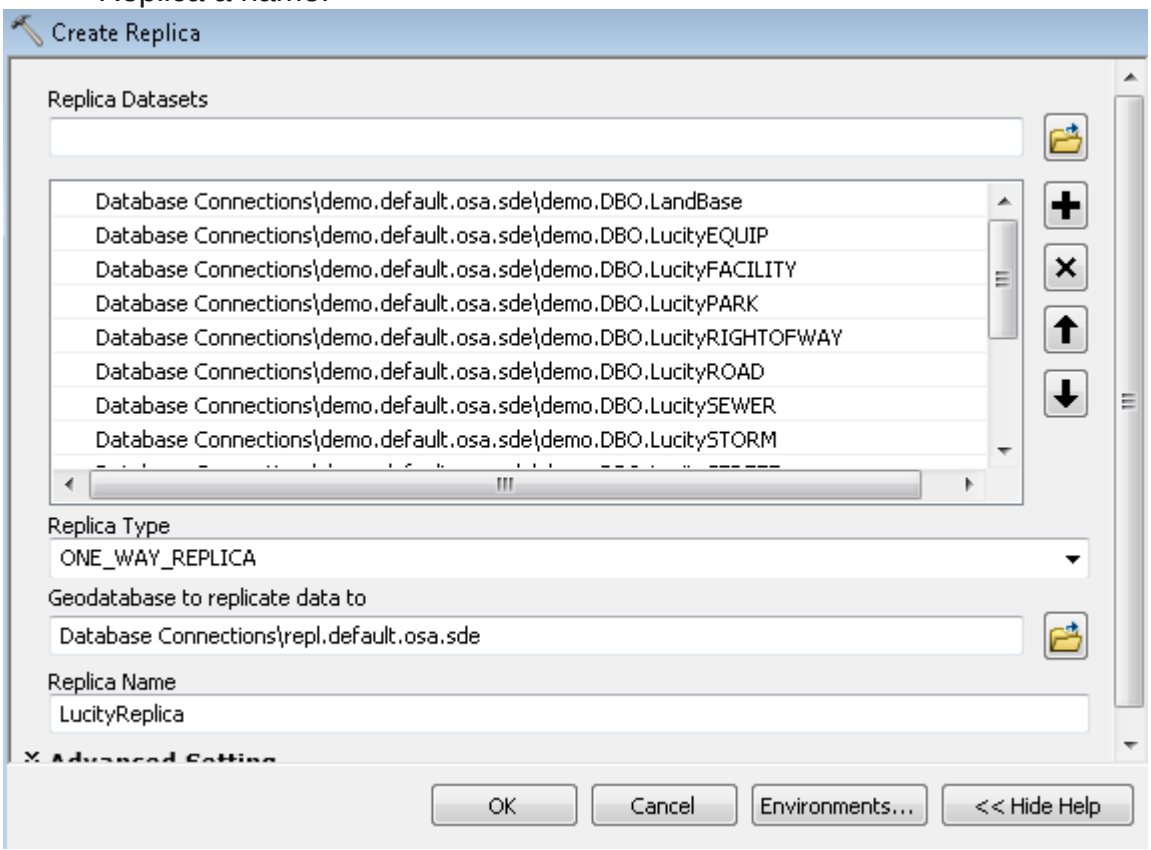

6. The advanced settings allow you to change the behavior of your replica. Usually, you would except the defaults and click ok.
7. Once replica is in place, you can assign privileges.