ArcSDE for Microsoft SQL Server Administration

Rob Stauder, ESRI

Rstauder@esri.com

Agenda

- Setting up the DBMS and Server
- The Default Configuration
- Using the dbtune table
- Maintaining Performance

Setting up the Server

- Supported DBMS/OS revisions
- SDE and non-SDE databases
- Filegroups and Files
- Logins, Users and Roles
- SQL Server and OS Settings

DBMS/OS Support

- **ArcSDE 8.0.2 for MS SQL Server:**
  - **MS SQL Server 7.0 sp1**
    - **MSDE, Desktop, Standard, Enterprise Edition**
  - **Windows NT 4.0 Workstation, sp4**
  - **Windows NT 4.0 Enterprise Edition, sp4**
  - **Windows 2000 Professional**
  - **MDAC 2.5**

DBMS/OS Support

- **ArcSDE 8.1 for MS SQL Server**
  - **MS SQL Server 7.0**
    - **MSDE, Desktop, Standard, Enterprise Edition**
  - **MS SQL Server 2000**
  - **Windows NT 4.0 sp6a**
    - **Workstation, Server, Enterprise Edition**
  - **Windows 2000**
    - **Professional, Server**

SDE and Non-SDE Databases

- **SDE Database**
  - Required
  - Container for SDE and GDB system tables and SDE system table stored procs
  - Holds all user logfiles
  - Default for ArcInfo
- **Non-SDE Databases**
  - Not required
  - Spatial data only
  - Data stored procedures
  - Must explicitly connect
  - Sde login must be database user

SDE and Non-SDE Databases

- **Multiple Databases:**
  - **ArcSDE tables are fully qualified**
    - **Db.owner.table**
  - **DML across databases (select,insert,update,delete)**
  - **NO DDL across databases! (create table, drop table…)**

SDE and Non-SDE Databases

- Database
  - Autogrow files in large increments
  - Stripe filegroups across multiple disks
  - Allows for more parallel disk access
- Logfile
  - Autogrow in large increments
  - Separate from os paging file
  - Separate from data files
  - Should have own disk.

Files and Filegroups

- **Filegroups are administrative groupings of data files.**
- **Filegroups allow you to assign tables to specific files (location on disks).**
  - `Create table contour (shape integer) on landbase`

Files and Filegroups

- **Files:**
  - **Physical allocation of space within a filegroup**
  - **A unit of parallelism**
  - **Smallest unit of recovery**

Files and Filegroups

- Files are filled using proportional fill strategy.
- Employ to create simple striping.
- Improve throughput by allowing parallel scans of data and reduces disk queuing.
- Autogrow in large increments

Files and Filegroups

Filegroup Example

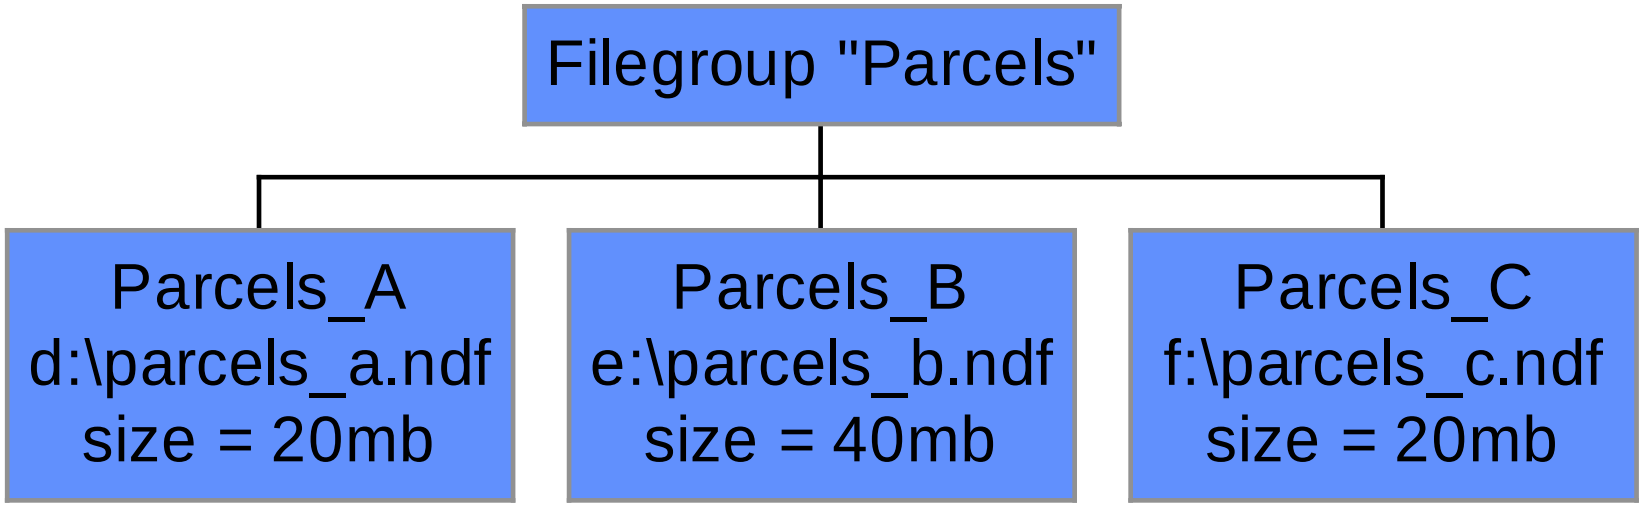

Files and Filegroups

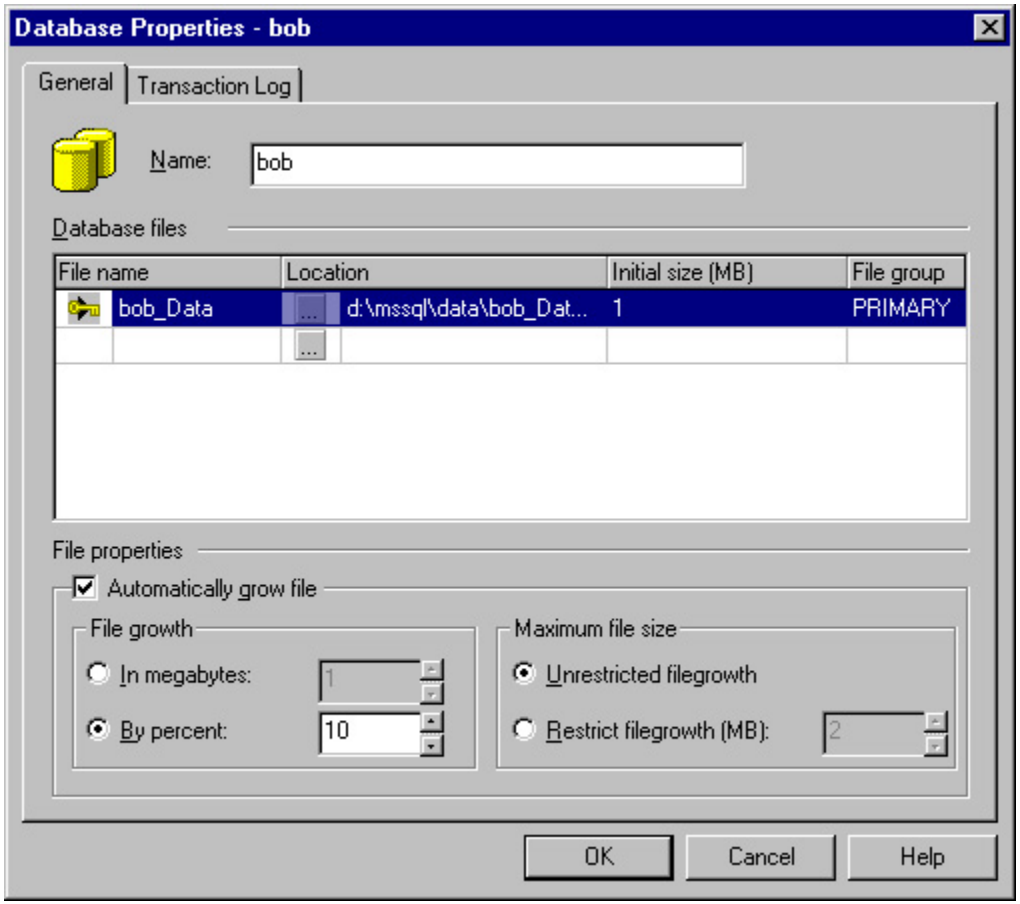

Files and Filegroups

- **Considerations:**
  - **For random I/O (majority of queries), create filegroups that span multiple disks.**
    - **More smaller disks are better than larger fewer disks.**
    - **High probability that pages accessed will be found on one or more disks.**
  - **For sequential I/O, allocate a single disk to a filegroup. Do not mix random and sequential data on the same drive.**

Files and Filegroups

- When to Use
  - Multiple disk controllers.
  - Multiple processors
  - Hardware striping (Raid)
  - Separate data from system tables
- When not to use
  - Single controller
  - Single processor
  - Ease of admin most important
  - Single user database

Logins, Users and Roles

- ‘sde’ login
  - ArcSDE 8.0.2 - must have *SQL Server Authenticated* ‘sde’ login.
  - ArcSDE 8.1 – ‘sde’ login not necessary

    BUT
  - ArcInfo Desktop 8.1 requires ‘sde’ login.

Logins, Users and Roles

General login and user information:

- **Logins have access to the sql server.**
- **Logins are granted access to databases and become database ‘users.’**
- **Permissions are granted at the database level.**

Logins, Users and Roles

ArcSDE 8.0.2 login to user rule:

- **All logins must have create table permission in the sde database.**
  - **To create their logfiles**
  - **Because logfiles are a dbms table, you should use different accounts for all users to reduce contention.**
- **All logins must have create table permission in any other database if they will own data.**

Logins, Users and Roles

ArcSDE 8.0.2	Sde db	Other db
Sde login	Add as a user Create table and procedure	Add as a user Create table
Other login	Add as a user Create table	Add as a user Create table

Logins, Users and Roles

- **ArcSDE 8.1 Login to User Rule:**
  - **If a user will own data in a database, that user must have CREATE TABLE and CREATE PROCEDURE granted in that database.**
  - **User must be able to create ‘I<reg_id>’ procedures.**

Logins, Users and Roles

ArcSDE 8.1	Sde db	Other db
Sde login	Add as a user Create table and procedure	Add as a user
Other login	Add as a user Create table Create procedure (if login will own data)	Add as a user Create table Create procedure (if login will own data).

Logins, Users and Roles

Roles:

- **Fixed Server, Database, User-Defined**
  - **Fixed Server roles span server**
  - **Database roles are specific to a particular database**
    - **User-Defined roles are a subset of Database roles**
  - **Recommendation: Assign users to roles with caution. Be very careful using Fixed Server roles.**

Logins, Users and Roles

‘sde’ as ‘dbo’

- The sde login will become ‘dbo’ if:
  - The sde login is added to the sysadmin fixed server role.
  - The sde login is added to the dbcreator fixed server role and creates the sde database (sde owns sde).
- SDE can have dbo privileges but remain ‘sde’ if:
  - Sde is added to the dbowner database role if the database is created by another account.

Logins, Users and Roles

‘sde’ as ‘dbo’ –

- Sde in sysadmin or dbcreator role (and creates the database) – System tables will be owned by ‘dbo.’
  - Not supported at 8.0.2
  - Supported at 8.1 NOTE: Arcinfo 8 does not support this configuration.
- Sde in dbowner role – System tables owned by ‘sde’.

Logins, Users and Roles

- ArcSDE 8.1 does not require an sde login.
- NOTE: ArcInfo 8 desktop does not support this feature.
- Create the sde service with –p “”

```
sdeservice -o create -p "" -l ham -d SQLSERVER
```

- The service is created as your nt/win2k login. Your login’s group will be mapped to an equivalent role in the database. An nt admin group member becomes a sql server sysadmin member.

Server and DBMS Configurations

- **SQL Server Settings**
  - sp_configure ‘show advanced options’,1
  - Many settings require ‘reconfigure with override’
  - Enterprise Manager
  - Some settings require server stop and restart
  - Use dbcc freeproccache and dbcc dropcleanbuffers to clear procedure and buffer cache when testing settings.

Server and DBMS Configurations

SQL Server Settings

- Affinity Mask
  - Specify processors used
  - Set with sp_configure or enterprise manager, restart server
  - Default = all
- Autoshrink Database
  - Disable with sp_dboption or with enterprise manager's db properties
- Autogrow Files
  - Let files grow in large increments – set with sql or enterprise manager
- Index Create Memory
  - Allocate memory to index building, set with sp_configure

Server and DBMS Configurations

SQL Server Settings

- Lightweight Pooling – reduce context switching with ‘fibers’
  - Set to 1 with sp_configure or Enterprise Manager
- Max Async IO – Max # of outstanding asynchronous I/O requests to a data file.
  - Max setting = 255, default = 32, boost with high performance RAID; Can over saturate I/O subsys with requests
- Max Worker Threads – Max # of threads (or fibers) possibly available. Default = 255
- Network Packet Size – Default = 4096, set with sp_configure or dbtune table

Server and DBMS Configurations

SQL Server Settings

- **Priority Boost –** If set, sql server runs at a higher priority than other processes. Set in Enterprise Manager
- **Recovery Interval** – Postpones automatic checkpoints.
- **Set Working Set Size –** If set to 1, sql server will not get paged out, even when idle. Use only when min and max server memory are set. Set with sp_configure.
- **Tempdb –** Make it large so it doesn't have to autogrow. Should not be on a fault protected disk.

Server and DBMS Configurations

Windows NT Settings

- Task management - Set performance boost to foreground application to none
  - Set under control panel-system-performance
- Resource Allocation - Maximize Throughput for Network Applications
  - Set under control panel-network-select server, then properties
- Pagefile.sys - away from transaction log
- Network Protocols - removed unused

Server and DBMS Configurations

- **Get MDAC 2.5!**
  - **www.microsoft.com/data**
  - **Check mdac version using component checker (see www.microsoft.com/data**
  - **check mdac version by checking version of msado15.dll in C:\Program Files\Common Files\System\ado. Version should be 2.50.4403.9**

Default Configuration

- ArcSDE 8 “out of the box” optimizations:
  - Firehose cursors
  - Clustered indexes
- Dbtune.sde configuration file now SDE_dbtune database table
- SDE_dbtune used to control data placement or “tune down” server.

Default Configuration

Firehose Cursors

- generate multiple connections made to a server – up to two connections per use of SE_stream_create.
  - Increase giomgr.defs' `MAXSTREAMS` parameter in multi-user environment
- Forego use of tempdb to populate work tables
- Eliminate need for temporary stored procedures
- Allow execution of stored procedures through SE_stream_prepare sql, etc.
- Permit use of TABLOCKX and UPDLOCK hints in select statements.

Default Configuration

Clustered Indexes at ArcSDE 8.1

- f<layer_id>_uk1 - Feature Table fid field
- d<reg_id>_idx2 - Deletes Table DELETED_AT
- a<layer_id>_ix1_a - Adds Table Shape field
- s<layer_id>_ix1 - Spatial Index table's covering index
- a<layer_id>_ix1 - Business table's shape field
- sde_logfiles_data_idx2 – SDE_logfiles_data table's sde_row_id column
- sde_logfiles_uk – SDE_logfiles table's logfile_name column

- One clustered index per table – The SQL Server query optimizer favors clustered indexes because these indexes organize data around the index key.

Default Configuration

- **Implications**
  - **Use SDE_dbtune table to disable default clustering**
  - **Page splits will occur on highly dynamic data and cause fragmentation of your tables.**
  - **Use index fill factor to delay page splits but don't set too low. FILL_FACTOR=%fill of index pages.**
  - **Rebuild indexes to reorganize your data when extent fragmentation occurs.**
  - **A clustered index controls the location of table data. Data resides at the index leaf level.**

Using the SDE_dbtune Table

Control data placement with –k switch or loading parameters with ArcInfo desktop tools.

- **Reference filegroup name in config_string field of SDE_dbtune table. Value becomes part of the "… on…" statement.**
- **Referenced Filegroup must exist in the connected database.**
- **If no filegroup is referenced, primary filegroup in database is used.**
- **A clustered index defines physical ordering of rows for a table's data pages!**

Using the SDE_dbtune Table

keyword	parameter_name	config_string
HYDRO	F_INDEX_1	WITH FILLFACTOR = 75 ON HYDRO

In this excerpt from an SDE_dbtune table, the keyword hydro points to configuration parameter F_INDEX_1. This parameter will create an index on the hydro filegroup by appending its config_string to a create index statement:

```
Create index F44_uk1 on f44.fid with fillfactor = 75 on HYDRO
```

Using the SDE_dbtune Table

- **What will happen here?**
- **Which Filegroup will contain the feature table?**

keyword	parameter_name	config_string
DEFAULTS	F_STORAGE	ON feats
DEFAULTS	F_IX1_CLUSTER	1
DEFAULTS	F_INDEX_1	WITH FILLFACTOR = 75 ON HYDRO

Using the SDE_dbtune Table

Data Loading

- 195+ SQL Server reserved keywords
  - Pipe, national,dummy,percent,precision,etc
  - May cause data loading to fail
- Column types:
  - 802 doesn't support bit or guid datatype
  - 8.1 supports bit and guid datatypes

Using the SDE_dbtune Table

To Tune down the server

- NUM_DEFAULT_CURSORS
  - 0 = server cursors
  - -1 = all firehose (default)
  - number > 0 to employ that many firehose cursors on concurrent SDE fetching streams.

Using the SDE_dbtune Table

To disable clustering:

- **Set the config_string = 0 for *_Cluster parameter_name value**

keyword	parameter_name	config_string
DEFAULTS	B_IX1_CLUSTER	0
DEFAULTS	F_IX1_CLUSTER	0
DEFAULTS	S_IX1_CLUSTER	0
DEFAULTS	S_IX2_CLUSTER	0

Maintaining Performance

Performance and tuning…

- **Monitor system using Performance Monitor**
- **Review and adjust Windows NT/2k settings**
- **Review and adjust SQL Server configuration settings**
- **Make only one change at a time and measure its effect – tune from a baseline/test**
- **Do periodic database maintainence**

Maintaining Performance

Page Splits and Extent Fragmentation

- Page splits occur when a full data page with rows ordered by clustered keys incur an insert or an *delete-insert* update. The page *overflows* causing the storage engine to allocate a new page and move approximately half the page's contents to it.
- 8 8kb pages comprise a table extent. A uniform extent holds contiguous data. A mixed extent does not.
- New pages derived from page splits may not be allocated from uniform extents. This tends to fragment your tables causing *extent fragmentation.*

Maintaining Performance

- Highly dynamic data can cause page splits.
- Delay page splitting with FILL_FACTOR
- Monitor table fragmentation with dbcc showcontig:

```
declare @id integer
select @id = OBJECT_ID('ogis.f1')
dbcc showcontig(@id)
```

- Fix extent fragmentation by rebuilding clustered index.
- Dbcc dbreindex, create index with drop existing, sdelayer –o load_only_io (unversioned data only)

Maintaining Performance

DBCC Showcontig (table_id,index_id):

```
DBCC SHOWCONTIG scanning 's45' table...
Table: 's45' (1293247662); index ID: 1, database ID: 7
TABLE level scan performed.
- Pages Scanned................................: 3
- Extents Scanned..............................: 3
- Extent Switches..............................: 2
- Avg. Pages per Extent........................: 1.0
- Scan Density [Best Count:Actual Count].......: 33.33% [1:3]
- Logical Scan Fragmentation ..................: 66.67%
- Extent Scan Fragmentation ...................: 33.33%
- Avg. Bytes Free per Page.....................: 2901.7
- Avg. Page Density (full).....................: 64.15%
```

Maintaining Performance

Post dbcc dbreindex('thad.s45')

DBCC SHOWCONTIG scanning 's45' table...
Table: 's45' (1293247662); index ID: 1, database ID: 7
TABLE level scan performed.
- Pages Scanned.....: 2
- Extents Scanned.....: 1
- Extent Switches.....: 0
- Avg. Pages per Extent.....: 2.0
- Scan Density [Best Count:Actual Count].....: 100.00% [1:1]
- Logical Scan Fragmentation .....: 0.00%
- Extent Scan Fragmentation .....: 0.00%
- Avg. Bytes Free per Page.....: 305.5
- Avg. Page Density (full).....: 96.23%

Maintaining Performance

- **To Rebuild Clustered Indexes**
  - **Sdelayer -o load_only_io - sdelayer -o normal_io (See sql script in faq at http://www.esri.com/devsupport/arcsde/samples/arcsde_online/sdehelp.htm) NOTE: This is only valid for unversioned data**
  - **dbcc dbreindex(db.owner.table)**

- **Will force data reorganization, refill data pages to FILL_FACTOR setting, prevent or delay page splits.**

Maintaining Performance

- **Schedule reindexing with the Database Maintenance Wizard.**
- **Table Statistics - leave autostatistics enabled to keep table statistics up to date.**
  - **Optimizer will detect when they are stale and update them.**

Maintaining Performance

Application Tuning

- Don’t display huge feature classes at full extent
  - Use sdegroup to reduce the number of records in a table.
  - Sdegroup creates a new layer of multi-part shapes.
  - Based upon tile size – tile size = size of commonly used extent.
  - Use scale ranges to avoid lengthy redraws
- Use sde views to create stored queries of your data (read-only)
  - Create with sdetable –o create_view
  - Remove with sdetable –o delete
- Use ArcInfo edit caches
- Use the client and server’s task manager watch cpu usage

Diagnosing Bottlenecks

Did performance degrade?

- Table fragmentation
- Disk fragmentation
- More users added?

Database layout

- Use SDE_dbtune
- Spatial Index
- Size of delta tables

Diagnosing Bottlenecks

Monitoring the Processor

- **Average utilization should not exceed 90%, peaks above ok. If avg > 80% you may have a bottleneck**
- **Perfmon Processor Object:**
  - **%Processor time counter – use per processor**
  - **%User Time – %time spent on application processing, not os or system functions.**
  - **Task manager performance tab**
- **Perfmon System Object:**
  - **Processor Queue Length-# of threads waiting to run. If > 2/processor, bottleneck could be due to processors, workload, change/rebuild indexes.**

Diagnosing Bottlenecks

- **Perfmon system object**
  - **Context switches/second: if >10,000 enable lightweight pooling. If still > 10,000/second add disks.**

Monitoring Disk I/O

- **Perfmon PhysicalDisk object**
  - **Use PhysicalDisk counters for RAID, LogicalDisk otherwise**
  - **Issue diskperf –Y[E] and reboot**
  - **Disk Reads/sec, Disk Writes/sec, Disk Transfers/sec:**
    - **Transfers is an aggregate of reads and writes**
  - **Avg Disk Sec/Read, Avg Disk Sec/Write – Avg time for a read/write from a disk**
  - **Avg Disk Queue Length: avg # of read/writes queued on disk or disk array.**
    - **if > 2, add more drives and spread data across more disks.**

Diagnosing Bottlenecks

Monitoring Memory

- **Perfmon Memory Object:**
  - **Available Bytes: amount of free memory**
    - **If low, you may have to add memory**
  - **Pages/Sec: # pages paged into or out of memory**
- **Perfmon Buffer Manager Object:**
  - **Buffer Cache Hit Ratio: % requests found in cache**
  - **Free Buffers: # free buffers available to SQL Server**
- **Perfmon Process Object:**
  - **Page Faults/sec: Use with Pages/sec to id process causing page faults.**

Diagnosing Bottlenecks

Object: Processor					
		0	1	2	3
	% Processor Time	31.491	22.874	19.756	32.542
	% User Time	30.248	21.449	19.212	31.118
Object: PhysicalDisk					
		_Total	0	1	
	Avg. Disk Queue Length	0.731	0.729	0.002	
	Avg. Disk sec/Read	0.006	0.006	0	
	Avg. Disk sec/Write	0.025	0.026	0.007	
	Disk Reads/sec	86.11	86.11	0	
	Disk Transfers/sec	95.549	95.265	0.284	
	Disk Writes/sec	9.44	9.155	0.284	
	Avg. Disk sec/Transfer	0.008	0.008	0.007	
Object: System					
	Context Switches/sec	1922.719			
	Processor Queue Length	0			
Object: Memory					
	Pages/sec	0.818			
	Page Faults/sec	431.518			
Object: SQLServer:Buffer Manager					
	Buffer Cache Hit Ratio	99.383			

ArcSDE 8.1 New Features

- Replication
  - Transactional
  - Snapshot
- Direct Connect Driver
  - ArcSDE Without giomgr
- Standby Server Support
- Performance Improvements

Additional Resources

- ArcSDE FAQ online:
  - http://www.esri.com/devsupport/arcsde/samples/arcsde_online/sdehelp.htm
- ESRI Developer Support:
  - http://www.esri.com/devsupport/index.html
- SQL Server Website:
  - http://www.microsoft.com/sql
- ESRI Systems Integration
- ESRI Educational Services


geography
gis
esri
our global network
ESRI