



ArcSDE Admin Tools

Alan Jackson
Jim Jarvis

ArcSDE Admin Tools

What we will cover:

**Overview of tools for administering
ArcSDE databases**

Demos of the database tools

Brief Introduction

What is SDE?

- Take a standard **RDBMS** Table and add a **GEOMETRY** column (with X,Y's inside).
- This is old news (pre SDE).
- Storage is simple.
- Attribute query & retrieval easy.

SDE solves these problems

- SDE provides a super fast **INTEGER** based topology engine (Why? Because of overlay processing.)
- SDE provides a very efficient method of **SPATIAL INDEXING** (utilizing grids)

Normal SQL Queries

A typical SQL query:

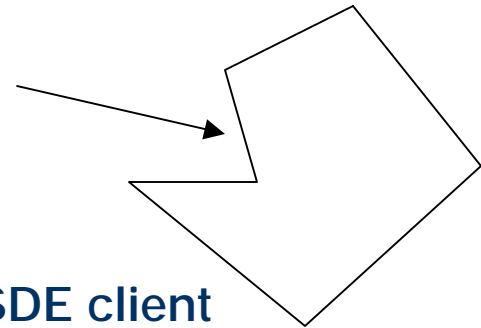
```
SELECT    <COLUMN(S)>  
FROM      <TABLE(S)>  
WHERE     <COLUMN> <RELATION> <VALUE>
```

```
SQL> select pop96, area from counties where name =  
'SAN DIEGO';
```

SDE Spatial Queries

SDE will allow you to make these kinds of queries by providing geometric column types and topology relational operators.

`SELECT pop96, geometry from counties
where area < 25000 and
geometry [is inside or touching] this`



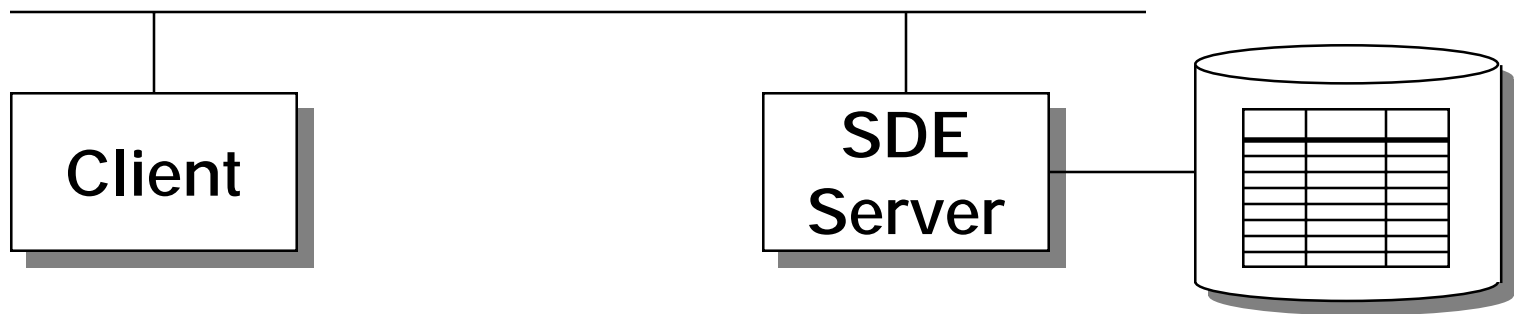
Note: This is conceptual. The actual implementation is programmed into the SDE client

Spatial Types and Functions


```
SELECT count(*)  
FROM buildings, bridges  
WHERE Contains(Buffer(bridges.position,  
15.0), buildings.footprint);
```

Note: This is the actual SQL specified in OGC test T51 - Spatial Types and Functions.

Client Server Queries




Connection on startup




Client connects to server

Query is initiated




Client submits query to server

SDE/RDBMS do their work




Server finds features satisfying constraints

Results shipped back



Server sends features to client over network


Client Rendering (Drawing)



Client processes one-at-a-time in a loop


S Table and Grids

Query a Land Parcels Layer



Spatial Query from Client


"retrieve all land parcels that **overlap** a ruptured storage tank's contamination plume"



Client submits spatially-constrained query

S Table Queried First


"retrieve all land parcels that overlap a ruptured storage tank's contamination plume"



Server determines which parcels share an index grid with the plume

Simple Envelope Test


"retrieve all land parcels that overlap a ruptured storage tank's contamination plume"



Server finds shapes with overlapping envelopes

Topology Engine Compares the Rest

"retrieve all land parcels that overlap a ruptured storage tank's contamination plume"






- Topology Engine will accept or reject features base on complex topology comparisons
- The process is: Grid->Envelope->Feature

Business and Feature Tables

WILDERNESS AREAS

Name	Boundary
Big Flat	8984
Upper Valley	2170
East Ridge	3642

F9

FID	Points
8984	
2170	
3642	

Overview of SDE RDBMS Tables and Their Uses

LAYERS

LAYER_ID	TABLE_NAME	...
2	ROADS	

ROADS

SHAPE	...
6582	
6583	
6584	

F2

FID	...
6582	
6583	
6584	

S2

SP_FID	...
6582	
6583	
6584	

How to install SDE

- Follow the SDE install instructions in the order specified
- Use the wizards in the SDE install when available
- Use RDBMS GUI tools to create users and databases. (or scripts if you must)

How to install SDE

- ArcSDE 8 will allow to put the SDE software under a directory of any name you choose. For example:

`/sdeadmin/sdeoracle804`

`/sdeadmin/sdeoracle8i`

Are valid SDE directory names.

How to install SDE

- **Exceptions:**
- **tcp address are automatically set by most 3.0.x installs. Duplicate tcp addresses will not work.**
- **To upgrade to ArcSDE 8 from SDE 3.x you do NOT change your data.**

How to install SDE

- ODBC
- Know your ODBC drivers:
- DB2 allows SDE database to be aliased to a different name.
- Use in-memory option for Informix IUS. Performance will increase up to 15%.

The SDE install crashed

Use:

- `sde.lic`- contains ESRI license information
- `sdeerr.log`-contains lots of stuff
- `giomgr.log`-rarely useful for debugging installs
- **Don't use** the C developers guide

The SDE install crashed

- Check the README.wri file:
- Contains a variety of debugging info including error codes and suggestions about the causes of each error code.

The SDE install crashed

- Check the SDE service:
- Use `sdeservice -o list`
- Make sure the service you tried to start is the right one
- Make sure the service is in the right path

The SDE install crashed

- Check the Event viewer (NT):
- Contains a description of the error.
- Some databases (SQL Server and Oracle) raise error messages which often appear in the system or application event viewer.

The SDE install crashed

- Check the license manager:
- Start by checking sde.lic
- Use Imutil or use the license manager GUI (NT) to find the status of the license manager, features licensed.

The SDE install crashed

- Check the ODBC connection:
- Try to login to the database with the SDE account and create a table
- Use ODBC test. Available for free at:
<http://www.microsoft.com/data/download.htm>

The SDE install crashed

- Check the database log. Examples:
- SQL Server: mssql7/log/errorlog
- Oracle:ORCALRT for alerts
- DB2: db2diag.log
- Informix: ol_<database name>.log

The SDE install crashed

Check the registry (NT):

.... **And be very careful**....

Loading data

- **Allocating space for tables**
- **Option 1: Allocate by files**
- **Recommended for transaction based systems**
- **Option 2: Allocate on demand**
- **Recommended for read-only or fixed data sized systems (data warehouses)**

Loading data

- Using a Redundant Array of Independent Disks (RAID)
- RAID 5: Best, highest cost solution
- Access many drives simultaneously
- Best throughput.
- Backup required

Loading data

- Using a Redundant Array of Independent Disks (RAID)
- RAID 10: most cost effective solution
- Uses disk mirroring and striping
- Each disk has a copy of the data
- Data is spread across many disks
- Backup required

Loading data

- **Tip:**
- **Use RAID 10 and put transaction logs on a separate disk not managed by RAID for best, cost effective performance solution.**
- **If you use RAID 5 put the transaction logs on a disk not managed under RAID 5. Nothing works better than RAID 5.**

DBTune file

- Used by SDE to allocate space for data

- Oracle:

```
##DCWDSET_TABLESPACE DCW
TEMP_TABLESPACE TEMP
INDEX_TABLESPACE DCW_IXF_TBLSP
F_INIT          40960
F_NEXT          40960
F_IX1_INIT      40960
F_IX1_NEXT      40960
A_TBLSP
DCWA_INIT       40960
A_IX1_INIT      40960
S_TBLSP         DCW_IX
```

DBTune file

•SQL Server

```
##DEFAULTS
AUTOCOMMIT_MODE          1
PACKET_SIZE              8192
F_STORE                  big_fas
F_IX1_STORE              big_idx
F_CLUSTER                1
A_STORE                  big_fas
S_IX1_STORE              big_idx
S_IX1_FILL               80
S_CLUSTER                1
S_IX2_STORE              big_idx
S_IX2_FILL               80
P_STORE                  big_p
P_IX1_STORE              big_idx
P_IX1_FILL               80
END
```

DBTune file

- Used by SDE to allocate space for data.

- Informix 9.2 : Datablade

```
##DEFAULTS
```

```
A_TBLSP WORLD
```

```
A_IX_TBLSP WORLD
```

```
A_SBLOB_OBS SBLOBDBS
```

```
END
```

Estimating Table Size

Table size is 75% of shapefile:

Shapefile	Table size
-----------	------------

100 MB	75 MB
--------	-------

DBF	
-----	--

100 MB	75 MB
--------	-------

Raster (BSQ)	
--------------	--

100 MB	100 MB
--------	--------

(ArcSDE 8.1 will support user defined compression of rasters.)

give away: layer_get_size

- Analyzes tables to calculate exact size.
- Outputs a DBTUNE.SDE entry

Usage: `layer_get_size <owner/passwd>`
`<layer_name> <layer_number> <keyword>`
`<multiplier> <next_scale>`

```
unix> layer_get_size sde/sde counties 2 COUNTIES 1.0 0.1
```

Loading Data

Creating Grids:

- Single most important aspect of tuning spatial data
- Implemented as S tables in SDE
- Often determined by users running the application.
- These tips will get the **DBA** started with optimizing performance. Tuning parameters for **applications** may differ.

Inside the S Table

- S Table Columns:
 - SP_FID, GX, GY, EMINX, EMINY, EMAXX, EMAXY
 - SP_FID is the join key.
- GX and GY are grid coordinates used in the S table search.
- The Envelope values are for trivial rejection of disjoint features. This avoids the fetch from the F table.

S Table Indexing

SDE 3.0

S<n>_IX1 on (GX, GY)

S<n>_IX2 on (SP_FID)

SDE 3.0.1

S<n>_IX1 on (GX, GY, and all of the rest)

S<n>_IX2 **NOT MAINTAINED!**

SDE 3.0.2/ArcSDE 8

S<n>_IX1 on (GX, GY, and all of the rest)

S<n>_IX2 on (SP_FID)

How is the S table used?

- To put the layer into normal io mode
- To add and delete records
- SQL is used to create the spatial index:
 - All the columns in the S table are indexed (concatenated). Example:

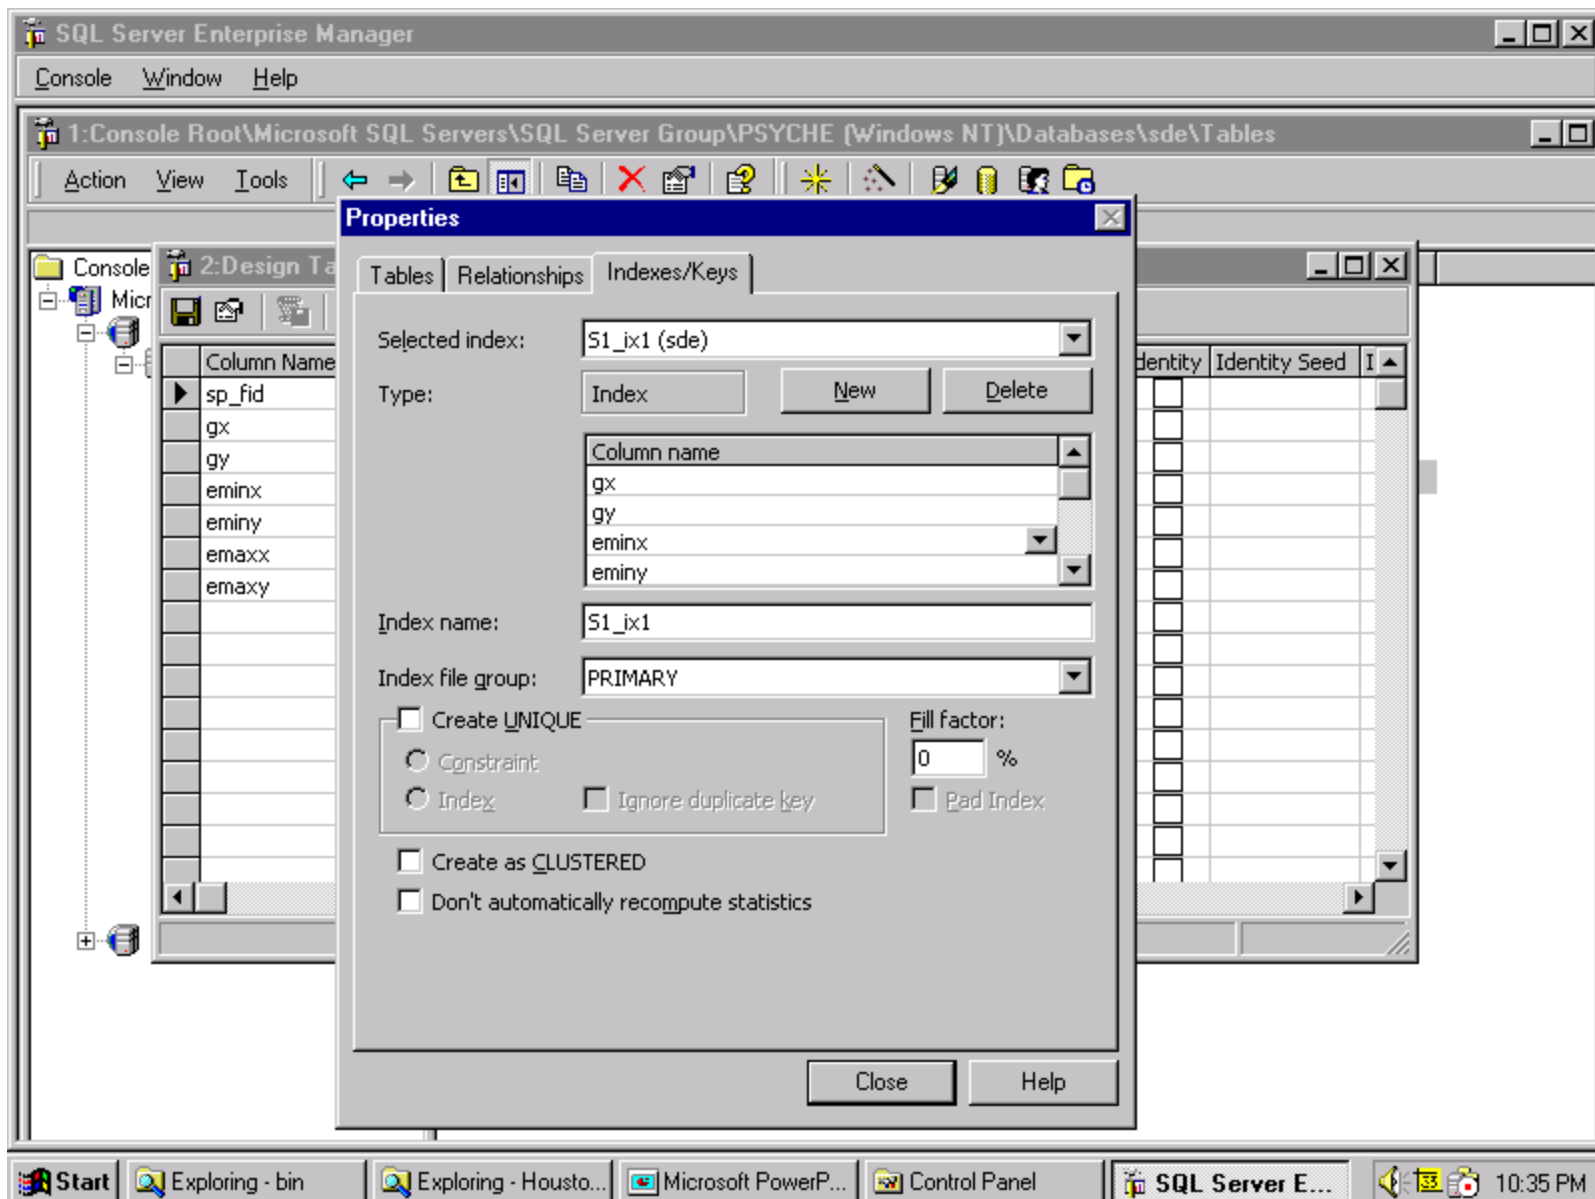
```
create unique index county on db2se.county_db2_shp(se_row_id)
```

How is the S table used?

- **Exceptions:**
- **Informix Datablade:** Uses Rtree to create a non-spatial index
- **DB2 Datajoiner:** Creates it's own spatial grid
- **Data converters do not use grid specifications in these 2 databases**

Demo: Enterprise Manager

Using The SQL Enterprise
Manager to examine indexes



Learn from this idea

- Indexing both the “**where**” columns **AND** the “**select**” columns helps Business Table Queries in the same way that it helped the S table queries.
- Use this to help speed your attribute queries.
- Index **BOTH** spatial column and your favorite column.
- Remember: Indexes are **ORDER DEPENDENT!!!**

Grid Recommendations

First Level Grid should be set to the **LARGER** of:

- Your layers average **feature size**.
- or
- Your average **Querying Area**.
- or
- Use **sdestats** utility
- First order grids are fine 99% of the time.

Never make grids **smaller** than your features

Demo: sdestats

A free grid size utility from ESRI

sdestats

Sdestats -o size -l counties,shape -F 0,1,10 -D sde -u sde -p go -r
output.txt

Value	Count	Pctile
0	30	0.96
0.1	24	1.72
0.2	98	4.84
0.3	610	24.27
0.4	1010	56.43
0.5	575	74.75
0.6	287	83.89
0.7	113	87.48
0.8	98	90.61

Grids: Final Remarks

- Even though Grids are tuned to an average query, there is almost always a “sweet spot” setting that works well for **ANY** query.
- Ballpark settings typically get you close enough.
- Grid tuning is a Black Art. Perform benchmarks.

Demo: Loading data into ArcSDE

Using the Feature Manipulation
Engine (FME) from Safe
Software to load data

The screenshot displays the FME for Smallworld (Evaluation License) application window. The main interface includes a menu bar (File, Edit, Mapping File, Preferences, Help) and a toolbar with various icons. Two dialog boxes are open:

- Set Translation Parameters:**
 - Source: Format: **ESRI Shape**, Dataset: **D:\ESRI\ESRIDATA\USA\Counties.shp**, Coordinate System: **Unknown**
 - Destination: Format: **ESRI Spatial Database Engine V3.x (SDE3)**, Dataset: **sde**, Coordinate System: **Same as source**
 - Semantic Mapping Files: **Automated Shape File Translation**
 - Feature Processing Functions: (empty field)
 - Clip Source Data To Feature:
 - Show all formats
- SDE Output Settings:**
 - Database Connection: Server: **psyche**, User ID: **sde**, Password: *******, Instance Name: **esri_sql7**
 - Table Parameters: Configuration: **DEFAULTS**, Dimension: **2**, Minimum X: **-180**, Minimum Y: **-90**, Scale: **1000**, Grid 0 Size: **1000**, Grid 1 Size: **1000**, Grid 2 Size: **1000**

The Windows taskbar at the bottom shows the Start button, Control Panel, Explorer windows, Windows Explorer, and the FME application. The system clock indicates 10:02 PM.

A word on ANALYZING TABLES

SQL> analyze table parcels calculate statistics;

Temp, Log and rollback space

After your layer loads, will it go into NORMAL_IO mode? If not, here are a few suggestions:

- Be sure your **TEMP** and **RBS's** are big enough.
- Check that your **TEMP** table space default storage is OK.
- Leave the layer in **NORMAL_IO BEFORE** you load it.
- Use a **COMMIT** frequency value.

Tuning through the GIO manager

The gio manager has two type of parameters:

Data transport and array sizing.

MAXBUFSIZE: Size of transport buffer

MINOBJECTSIZE: Minimum number of objects or rows per transport buffer

MINBUFSIZE: Minimum number of bytes per transport buffer

RASTERCOLUMN: Size of image tile (ArcSDE 8.1)

Tuning through the GIO manager

Admin/DBA should increase the size of
these parameters

For example:

Increasing the transport buffer to
300K accounts for 15M of memory usage for
50 users for transport memory only

Increase the MINBUFSIZE to 1/2 the MAXBUFSIZE
Setting the MINBUFSIZE too high will increase the wait
time to fill the buffer .

Tuning through the GIO manager

Array buffer:

MAXBLOBSIZE	1000000	# Maximum BLOB size allowed for storage
BLOBMEM	500000	# Maximum BLOB size for in-memory storage
SHAPEBUFSIZE	4000	# Shape array buffer size
SHAPEPTSBUFSIZE	400000	# Shape POINTS array buffer size
ATTRBUFSIZE	50000	# Attribute array buffer size
SPINDEXBUFSIZE	28000	# Spatial index array buffer size
BLOBBUFSIZE	30000	# BLOB(attribute) array buffer size
MAXARRAYSIZE	100	# Max. array fetch size
MAXARRAYBYTES	550000	# Max. array bytes allocated per stream
AUTOCOMMIT	500	# Auto-commit frequency within user transaction

Tuning through the GIO manager

Tuning SHAPEPTSBUFFSIZE to the optimal setting is critical to performance.

The SDE server estimates the average size of all features based on the array size (MAXARRAYSIZE) and the size the points buffer (SHAPEPTSBUFFSIZE).

If a feature exceeds this size and cannot fit, it is flagged as truncated and fetched separately.

Tuning through the GIO manager

For example:

FType = (X,Y,Z,Measures) 4-bytes per type.
X-Y only is an FType of '8'

**SHAPEPTSBUFSIZE = (AVG(numofpts) * FType) *
MAXARRAYSIZE**

Tuning through the GIO manager

DB2 Datajoiner and Informix Datablade users:

Proper setting of the Blob parameters (MAXBLOBSIZE and BLOBMEM) is essential for performance.

These databases store features as Geometry datatypes.

We use the word BLOB because of a naming convention that started in SDE 2.x. The word BLOB is used for compatibility only.

Tuning the SGA (Oracle)

Start your SGA with about 50-60% of memory then use `cache.sql` and `latch.sql` to measure performance.

Let the instance run for awhile before you run them

Cache and latch recommend memory settings for the instance

They are free and are included with SDE

They are located in the tools directory

Load Testing the Tuned SDE

To test your settings use **SDETIME**

SDEtime is designed for load testing of SDE databases.

SDEtime executes random spatial queries on an SDE database.

It is free.

It is written in MapObjects 2.0 using Visual Basic 6.0

Load Testing the Tuned SDE

SDETIME Demo

SDE Testing

Buffer Size: 0.5 Interval:(sec) 2 Length: (sec) 10

Procedure	Start	End	Lapse(s)	nRecs	nPoints	Size	ReadLapse(s)
Get Data:	22:19:10	22:19:12	2.1830	2	139	3.11k	0.01
Get Data:	22:19:14	22:19:15	0.6210	7	317	8.17k	0.03
Get Data:	22:19:28	22:19:28	0.4410	4	144	4.08k	0.01
Get Data:	22:19:30	22:19:31	0.4300	6	46	3.39k	0.02
Get Data:	22:19:33	22:19:33	0.4410	5	160	4.78k	0.02

See result file: D:\HoustonSDE\sdetime.log (-106.76, 40.56; -105.76, 41.56)

Tuning Myths

If you do the right stuff:

- De-fragment your data.
- Spread out your I/O
- Tune Oracle
- Tune the SDE grids
- Tune using giomgr.def
- Index Properly

You keep making SDE **faster** a little at a time.

Reality

- SDE throughput is a **multi-step** process.
- You're only as fast as your slowest step.
- Speeding up a faster step **WON'T** help.
- Focus on the things that matter.

Tune what Matters

- Things that make a **BIG** difference
 - Proper Attribute Indexing
 - Grid Sizing
 - Application Logic
- Things that may not make a big difference
 - Fragmentation
 - Spreading out your I/O
 - Oracle fine tuning
 - Set database to Read-only

Tune what Matters

- **ALWAYS** calculate the extent of the layer

`sdelayer -o alter -E calc`

Views and Synonyms

- Views between tables in different databases in increasingly popular.
- Synonyms are usually a carryover of SDE 2.
- Avoid changing column names in SDE tables:
"This is walking into a place of worship and feeling the need to rearrange the furniture."
Unidentified SDE team member.

Utility locations

Free SDE utilities at:

<http://www.esri.com/devsupport/devconn/sde>


and:

ftp://ftp.esri.com/pub/staff/vangelo/sde/se_toolkit-1.1.0.tar.gz

SDE Admin Tools

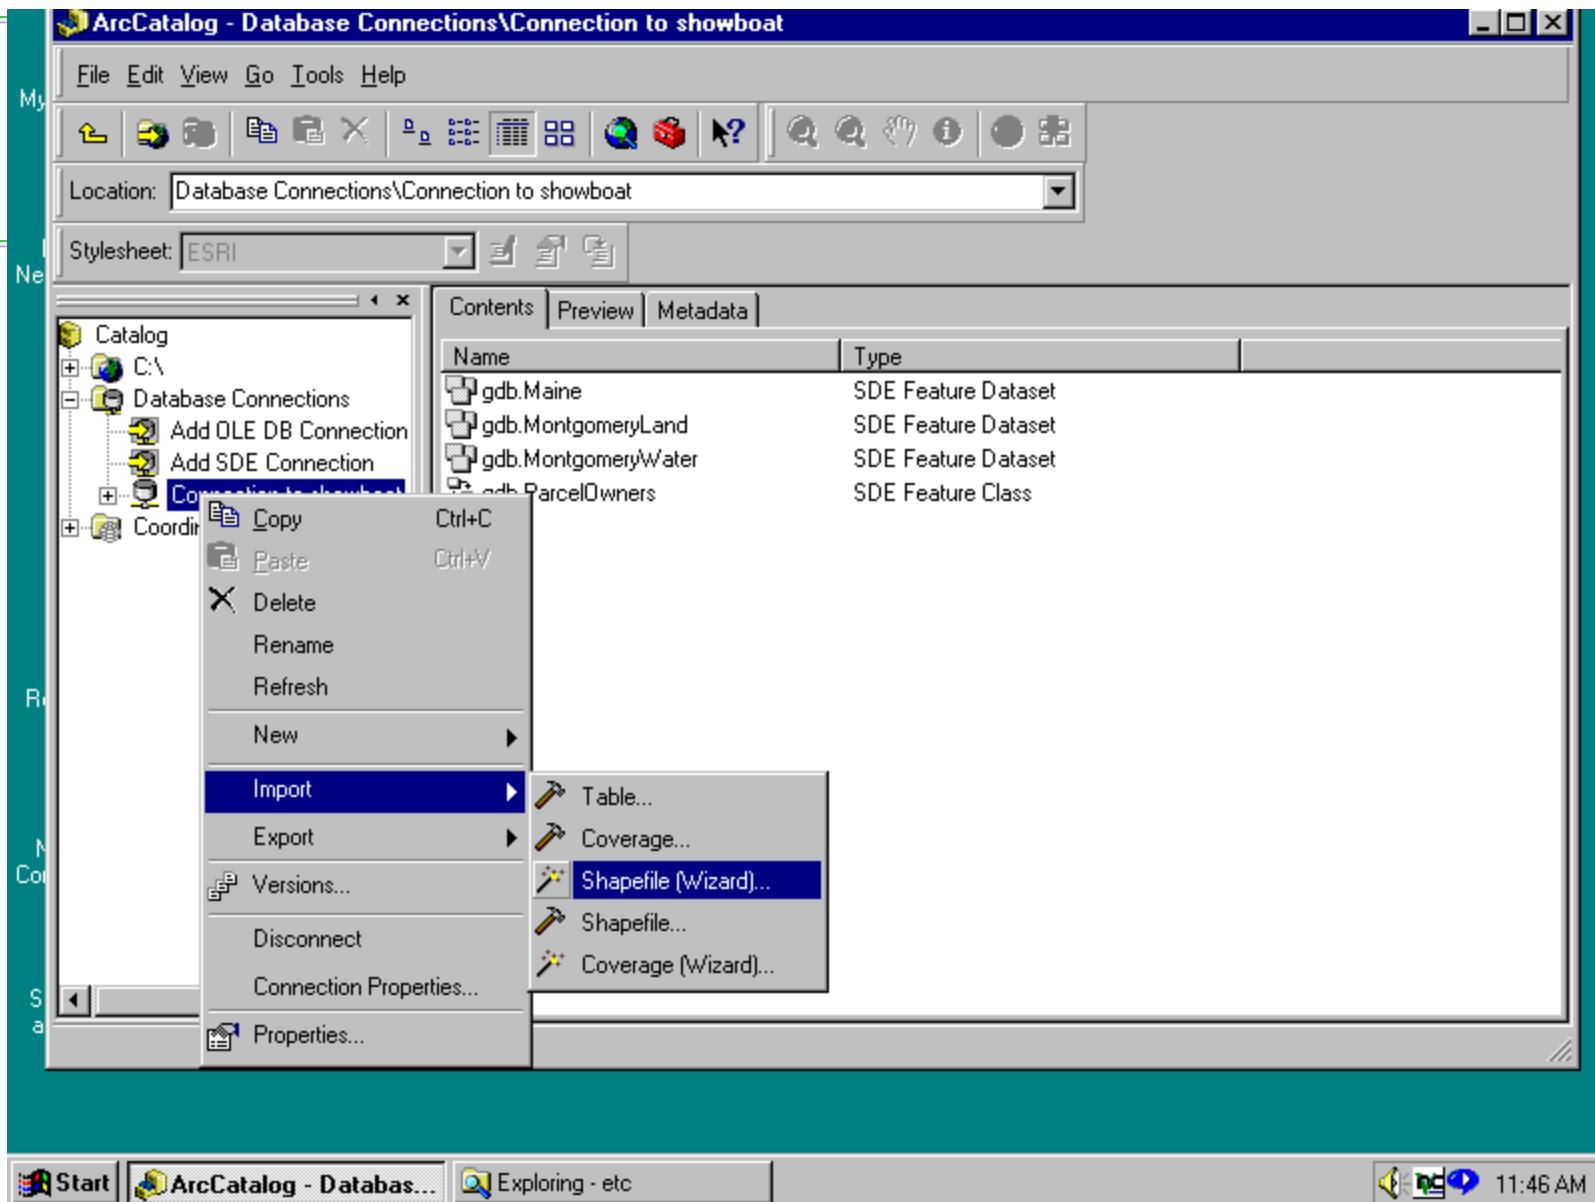
- SDA - ESRI Sweden
- ArcCatalog
- SDX - On Line Evolution
- Java SDE tool
- SDE HTML
- SDE Loader
- SDE-SQL Server Admin tools

Default SDE Admin Tool



SDE Admin Tools

- **ArcCatalog:** The new data loading interface for ArcSDE.
- **Available only on the NT.**
- **For data import/export**
- **Uses wizards.**



SDA - ESRI Sweden

- Built on the C API using C++
- Can create and run scripts for layer management/manipulation
- Can spatially enable DBMS tables
- Connection monitoring
- The best tool for administrating SDE 3.x databases.
- For more information see:
<http://www.esri-sweden.com/products/sda.shtml>

Spatial Database Administrator - [Connection1 - [psyche, esri_ius, sde, sde]]

File Edit View Layer Table Window Help

Layers Tables

ID	Table	Column	T.	Access	Created
8	SDE.SOILS	FEATURE	a	sel ins upd del	3/12/99 2:32:48 PM
9	SDE.ZIPCODES	FEATURE	a	sel ins upd del	3/12/99 2:50:23 PM

Create layer

Description: roads

Table*: SDE.ZIPCODES

Spatial column*: line

Minimum ID: 1

Creation keyword: DEFAULTS

Initial number of shapes: 0

Average points per shape: 0

Array size: 0

Shape types*

- Nil shapes (n)
- Points (p)
- Lines (s)
- Simple lines (l)
- Areas (a)
- Multipart shapes (+)

Dimension: 2-D

Grid sizes

Grid size 1*: 1000

Extra data

- Annotations (A)
- CAD data (c)

User-defined extent

Min X: Max X:

Min Y: Max Y:

Clear Calculate Index

* = Required fields when creating layer.

Help Cancel OK

Coordinate system

Projected: WGS_1972_UTM_Zone_7S

Geographic:

None

Description:

```
PROJCS["WGS_1972_UTM_Zone_7S",GEOGCS["GCS_WGS_1972",DATUM["D_WGS_1972",SPHEROID["WGS_1972",6378135,298.26]],PRIMEM["Greenwich",0],UNIT["Degree",0.0174532925199433]],PROJECTION["Transverse
```

OK Cancel Help

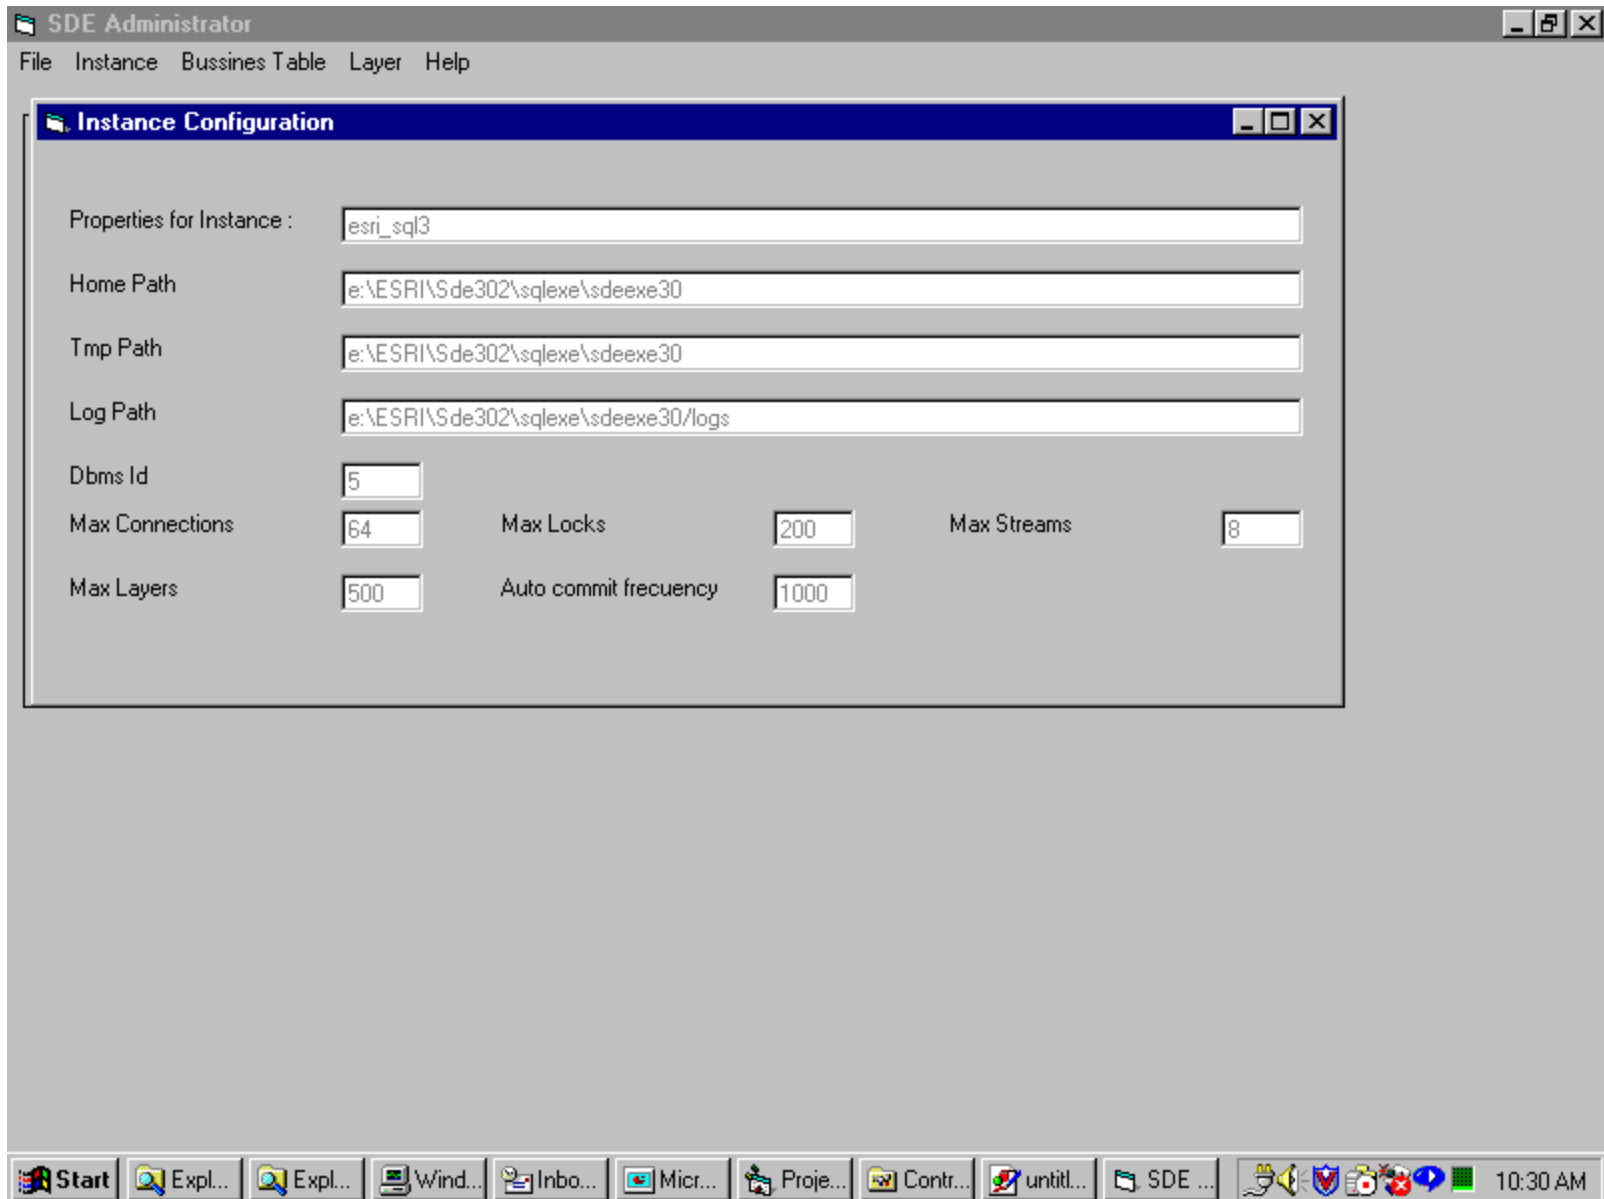
Ready

SDA 1.3

Demo

SDX (On Line Evolution Inc.)

- **ActiveX component for SDE administration**
- **Administration of instances and users of an instance**
- **Full support for layer management.**
- **Spatially enables existing DBMS tables.**
- **For more information see:**
<http://www.geocities.com/SiliconValley/Byte/1574/>



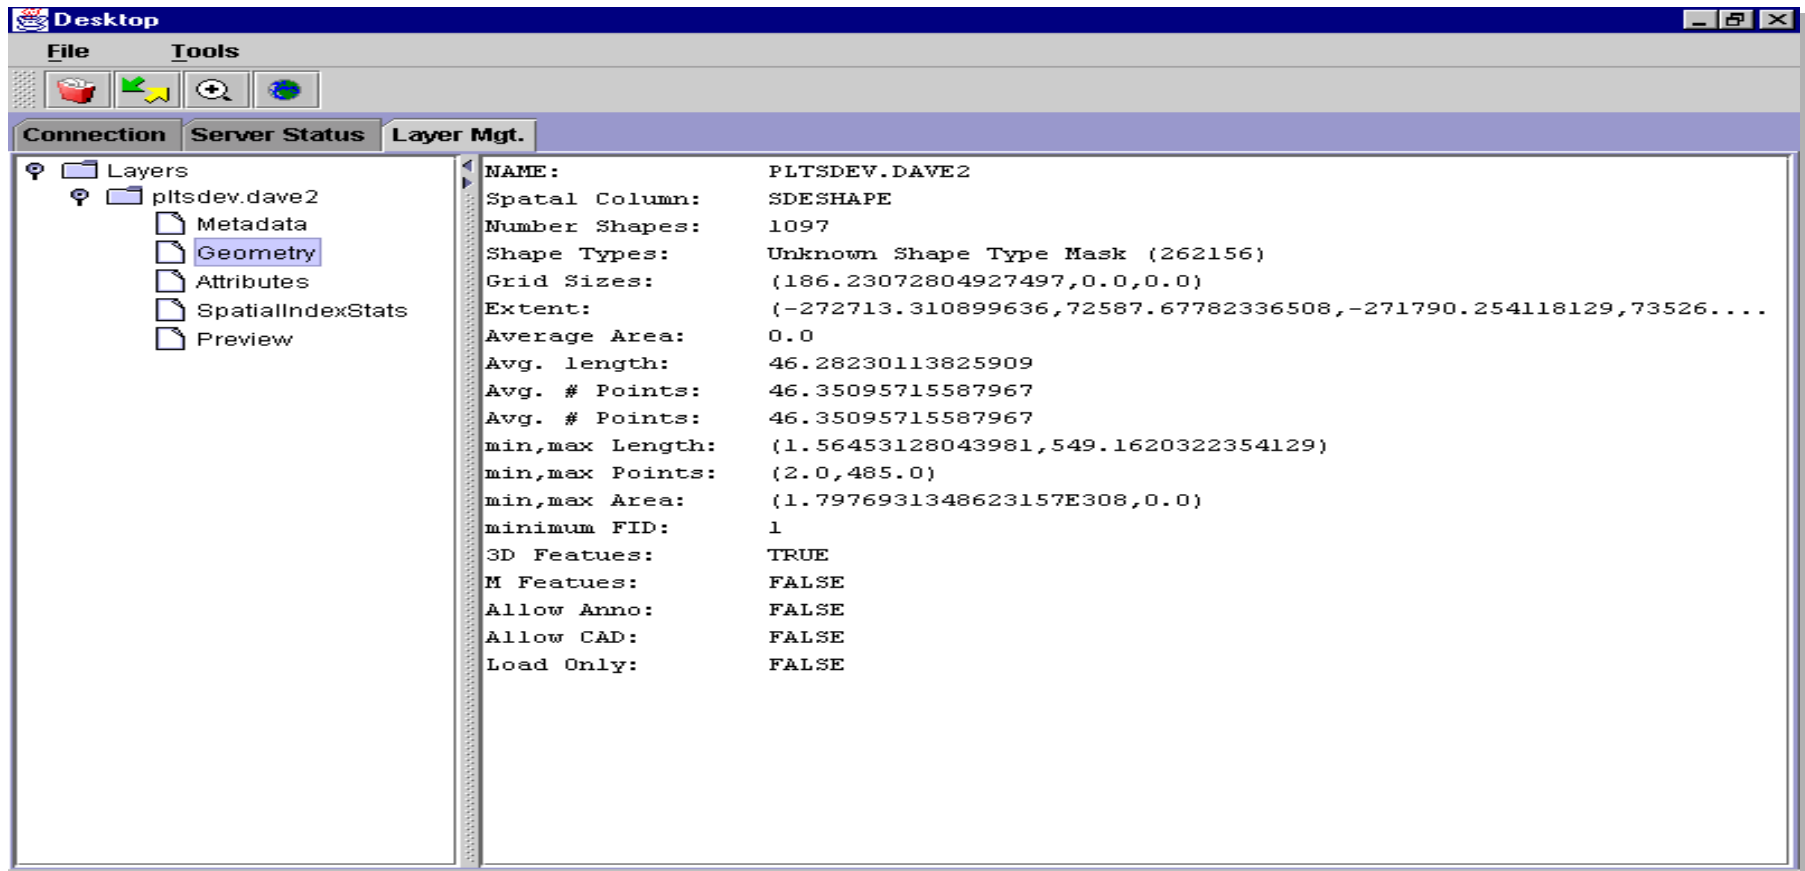
SDX Version 1.0 (beta)

Demo

Java Admin Tool

- **Freeware from ESRI**
- **Supports display of layers and layer attributes**
- **Runs under Windows NT and Unix**
- **Supports SDE 3.x and ArcSDE 8 beta 1**
- **Is being ported to ArcSDE 8 final**

Java Admin Tool




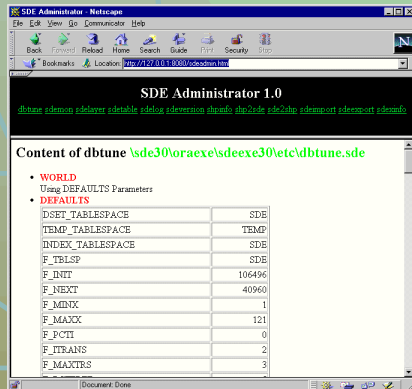
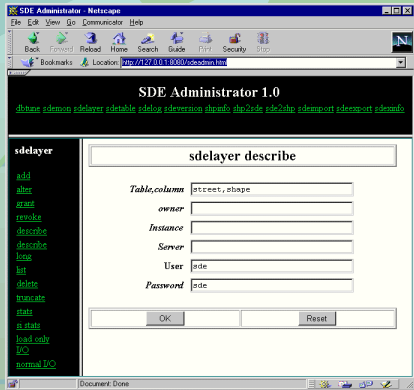
Java SDE Admin Tool

Demo

SDE HTML Requirements


- **FREEware from ESRI**
- **JavaScript Enabled Web Browser.**
- **Servlet Enabled Web Server.**
 - CGI version also available
- **SDE 3.x Version.**
- **Downloaded *.html, *.class files.**

SDE HTML Admin



```

##WORLD
END
##DEFAULTS
DSET_TABLESPACE SDE
TEMP_TABLESPACE TEMP
INDEX_TABLESPACE SDE
F_TBLSP SDE
F_INIT 106496
F_NEXT 40960
F_MINX 1
F_MAXX 121
F_PCTI 0
F_ITRANS 2
    
```




















SDE Loader

- **Freeware from ESRI**
- **Displays layers and attributes**
- **Import/Export of data**
- **Written in MapObjects 1.2 and Visual Basic.**

SQL Server Admin Tools

- **Freeware from ESRI**
- **Runs as stored procedures or wizards**
- **Performs layer validation and database creation**

ArcSDE Admin Tools: Summary

	Platform		ArcSDE version	
	NT	Unix	3.x	4.x
ArcCatalog				
SDA				
SDE Java				
Webtools				
SDELoader				
Command Line				

Geographic Knowledge Everywhere

Uc
99c

ESRI *Nineteenth Annual User Conference*