

GIS by ESRI

ESRI® White Paper Series
March 1996



Spatial Database Engine™ (SDE™)

Copyright © 1996 Environmental Systems Research Institute, Inc.
All rights reserved.
Printed in the United States of America.

The information contained in this document is the exclusive property of Environmental Systems Research Institute, Inc. This work is protected under United States copyright law and other international copyright treaties and conventions. ESRI grants the recipient of the ESRI information contained herein the right to freely reproduce, redistribute, rebroadcast, and/or retransmit this information for personal, noncommercial purposes, including teaching, classroom use, scholarship, and/or research, subject to the fair use rights enumerated in Sections 107 and 108 of the Copyright Act (Title 17 of the United States Code). No part of this work may be reproduced or transmitted for commercial purposes in any form or by any means, electronic or mechanical, including photocopying and recording, or by any information storage or retrieval system, except as expressly permitted in writing by Environmental Systems Research Institute, Inc. All requests should be sent to Environmental Systems Research Institute, Inc., 380 New York Street, Redlands, CA 92373-8100 USA, Attention: Contracts Manager.

The information contained in this document is subject to change without notice.

RESTRICTED RIGHTS LEGEND

Use, duplication, and disclosure by the government are subject to restrictions as set forth in FAR §52.227-14 Alternate III (g)(3) (JUN 1987), FAR §52.227-19 (JUN 1987), or DFARS §252.227-7013 (c)(1)(ii) (OCT 1988), as applicable. Contractor/Manufacturer is Environmental Systems Research Institute, Inc., 380 New York Street, Redlands, CA 92373-8100 USA.

ESRI, ARC/INFO, PC ARC/INFO, ArcView, and ArcCAD are registered trademarks; ARC COGO, ARC NETWORK, ARC TIN, ARC GRID, ARC/INFO LIBRARIAN, ARCPLOT, ARCEDIT, TABLES, Application Development Framework (ADF), ARC Macro Language (AML), Data Automation Kit, Avenue, FormEdit, ArcSdl, ArcBrowser, ArcDoc, *ARCLine*, ARCSHELL, IMAGE INTEGRATOR, DATABASE INTEGRATOR, DBI Kit, ARC/INFO, ArcTools, ArcStorm, ArcScan, ArcExpress, ArcPress, Maplets, Spatial Database Engine, SDE, PC ARCEDIT, PC ARCPLOT, PC DATA CONVERSION, PC NETWORK, PC OVERLAY, PC STARTER KIT, PC ARCSHELL, Simple Macro Language (SML), ArcUSA, ArcWorld, ArcScene, ArcCensus, ArcCity, *BusinessMAP*, the ESRI corporate logo, the ESRI globe logo, the ARC/INFO logo, the ARC COGO logo, the ARC NETWORK logo, the ARC TIN logo, the ARC GRID logo, the ARCPLOT logo, the ARCEDIT logo, the Avenue logo, the ArcTools logo, the ArcStorm logo, the ArcScan logo, the ArcExpress logo, the PC ARC/INFO logo, the ArcView logo, the ArcCAD logo, the ArcData logo, the SDE logo, the ArcPress Logo, the ESRI GISTOOLS logo, ARCware, *ARC News*, ArcSchool, ESRI—Team GIS, ESRI—The GIS People, GIS by ESRI, ARC/INFO—The World's GIS, Geographic User Interface (GUI), Geographic User System (GUS), Geographic Table of Contents (GTC), Your Personal Geographic Information System, GIS for Everyone, and The World's Leading Desktop GIS are trademarks; and the ArcData Publishing Program, ARCMail, ArcQuest, ArcWeb, ArcOpen, www.esri.com, @esri.com, and Rent-a-Tech are service marks of Environmental Systems Research Institute, Inc.

The names of other companies and products herein are trademarks or registered trademarks of their respective trademark owners.

Spatial Database Engine[™] (SDE[™])

An ESRI White Paper

Contents	Page
Need for Better Data Management	1
Need for Providing Better Data Access	1
The New Technology	1
Key Benefits of SDE	1
Key Features of SDE	2
Supports Large Databases	3
Efficient Spatial Object Model	3
Spatial Object Types	3
Attribute Data	4
Spatial Analysis	5
SDE Architecture	6
Performance	8

Spatial Database EngineTM (SDETM)

Over the past twenty-five years, enormous volumes of digital spatial data have been collected through various techniques such as automated mapping, geographic information systems, and image processing. Millions of files contain geographic information in a variety of formats, often tailored to specific projects. These data support a wide variety of applications in many disciplines including government, land management, resource management, transportation, health care, and business.

Need for Better Data Management

As awareness grows that geographic information can be used for more general purposes, demand for access increases. However, data access may be constrained by the need for a geographic information system (GIS) or desktop mapping package, constraints that often preclude broader use of geographic information.

Need for Providing Better Data Access

Organizations naturally want to leverage their significant time and financial investment in geographic data by providing access for applications ranging from general purpose to mission critical. While easily shared among a small group, the distribution of data across an organization has been more problematic. This has often resulted in multiple copies of geographic data being distributed among various groups. Without highly structured procedures to manage this replication, the consistency and integrity of the data are at risk, ultimately putting decisions based on that data at risk.

The New Technology

Taking advantage of advances in client/server computing and database management technologies, Environmental Systems Research Institute, Inc. (ESRI), has created a technology that allows you to manage your geographic data, disseminate it based on efficient spatial analysis queries, deliver it anywhere on your network, and embed the query and analysis functionality into any application. This new technology is the Spatial Database EngineTM (SDETM).

Key Benefits of SDE

Application developers, database administrators, and end users alike benefit from SDE.

The developer can

- Build high-speed applications.
- Tailor applications to specific user needs.
- Embed sophisticated geoprocessing into applications.
- Choose the development environment.
- Deliver applications on a variety of platforms.

The database administrator can

- Increase availability of geographic data.

- Provide multi-user access to a single database.
- Tune the database for performance.
- Manage millions of geographic features.
- Maintain geographic data integrity.
- Integrate geographic data into the corporate data management structure.
- Manage geographic data using facilities of a relational database management system (RDBMS).

The End User End users are the ultimate beneficiaries. They get fast applications, built to their specifications on popular computing platforms, delivered to their work site.

Key Features of SDE Implemented in a standard relational database environment, SDE represents geographic features (e.g., land records, roadways, sales territories) using a robust geometric data model, complemented with the full set of relational database data types (e.g., integer, date, Binary Large Object [BLOB]) for representing attribute information. In addition to the data management facilities inherited from the relational database environment, SDE implements a client/server, cooperative processing architecture, ensuring maximum performance for hundreds of users. The client application programming interface, delivered for a number of different computing platforms, is easily integrated into any application development environment.

Using these open systems constructs, SDE acknowledges the overall popularity and availability of geographic data, and addresses the pressure to capitalize on the large investments in building spatial data sets, the requirements to provide fast, multiple-user access to them, and the need to combine the management and use of geographic data with other traditional and multimedia data in one seamless computing environment.

With the enormous breakthrough in combining very large spatial databases with complex spatial query tools that efficiently deliver geographic features from a relational database, SDE represents a significant evolution in spatial data access.

The benefits provided by SDE are supported by a number of design features.

- Geospatial**
- Very large spatial data sets
 - Exceptionally fast spatial data retrieval
 - Flexible topologically constrained searches
 - Dynamic spatial overlay direct on the database
 - Logically continuous nontiled spatial data sets

- Open Systems**
- Cooperative processing client/server architecture
 - TCP/IP network access to data via tunable intelligent asynchronous buffering
 - Two-phase transaction control
 - Shared multi-user environment

Application Development

- Well-defined, open application program interface (API)
- Selectable graphic user interface (GUI) client software
- Selectable, industry-standard RDBMS spatial data storage
- Flexible security scheme for shared environment
- Support for multiple computing platforms
- Client support from ESRI's ArcView Version 3 and MapObjects software

Supports Large Databases

SDE uses a unified data model to maintain spatial and attribute data in an RDBMS. It can manage millions of spatial features that might be encountered in databases containing all of the roads for the United States, all of the land parcel records for one state, all of the customer locations of a credit card company, all of the wells of an oil company, or all of the stands for a forestry company.

It does this in a single continuous database with multi-user access.

Efficient Spatial Object Model

Geographic features, such as restaurant locations, scenic routes, or vacation regions, are stored using a spatial object type and a set of attributes linked by a unique feature identifier. Efficient geographic feature representation requires more than just a point object; it requires a robust set of spatial object types with distinct characteristics and behaviors.

Addressing this required flexibility, SDE can represent linear features as either a line that is not allowed to intersect itself, a line that can only intersect at endpoints, one where self-intersection is not a constraint, or lines that connect at endpoints to form a network. Similarly, area features can be represented as a single polygon, as a polygon with one or more holes, or as many polygons.

An SDE data set, called a layer, is a set of geographic features of uniform object type having the same attribute set. Features are indexed by layer using a nontiled, space-continuous scheme, promoting fast retrieval operations, even with a very large number of features.

For example, a land parcel, under a topological or cartographic model, is not stored as a complete object. Instead, it is stored as a number of node-to-node paths. Additionally, attribute data are not indexed directly to the parcel object, but rather to a place somewhere in the interior of the parcel. This can be a disadvantage because every access of a parcel polygon requires a reconstruction involving many disk accesses. The object entity model employed by SDE stores each parcel as a complete polygon, and therefore can retrieve the entire object with a single disk access.

In SDE terms, the spatial object types are point, point cluster, spaghetti, line string, ring, polygon, donut polygon, and network.

Spatial Object Types

SDE currently supports three object classes containing seven object types. Spatial objects may be either two or three dimensional. The SDE administrator declares one or more allowable object types for each layer. For instance, land parcels must be of area feature class only. Additionally, SDE supports the concept that a feature type is also a network. A network feature type may only contain LINESTRING and RING objects.

Point Class Object Types	<ul style="list-style-type: none">■ POINT Object—A single coordinate point. This object type is used to represent point locations such as water wells, hydrants, or control points.■ POINTCLUSTER Object—An arbitrary grouping of nonconnected coordinate points. The advantage of this object is speed—the ability to retrieve many points with a single disk access. However, attribute data cannot be associated with individual points easily (although it can be done). This object type is useful for efficient storage and retrieval of elevation point data.
Linear Class Object Types	<ul style="list-style-type: none">■ SPAGHETTI Object—A linear path consisting of (n) connected points. No integrity checking is made. This object type is useful for importing linear data from other systems where the integrity of the data is suspect and automatic SDE feature integrity checking is not desired.■ LINESTRING Object—A linear path consisting of (n) connected points where nonconnected lines may not cross or touch.■ RING Object—A linear path consisting of (n) connected points where the starting and ending points are identical. Logically the feature has no area although visually it is identical to a polygon. Nonconnected lines may not cross or touch. This object type is used to represent such features as ring roads.
Area Class Object Types	<ul style="list-style-type: none">■ POLYGON Object—A closed path consisting of (n) connected points where the starting and ending points are identical. This object represents an area feature. Nonconnected lines may touch at the vertices to form an inverted polygon. These objects logically represent the area contained by the polygon, inclusive of the edges. This object type is used to represent such features as land parcels.■ DONUTPOLYGON Object—A multiple-shell polygon consisting of a single outer shell containing one or more interior shells. The interior shells represent null areas or "holes" Individual shells may not touch. The outer shell may include inversions (like a donut that touches at one point). Interior "null" shells may include exversions (two donut shells that touch at a single point are logically a single connected shell). These objects logically represent the area contained by the outer polygon shell, inclusive of the edges, except for the interior shells which represent null areas. This object type is used to represent such features as lakes which have islands.
Network Class Feature Types	<p>If a feature type (e.g., road centerlines) is declared to be a network, the only allowable object types are LINESTRING and RING. For each feature type declared a network, additional node-to-node tables are maintained by SDE. Network processing functionality is provided for those feature types declared a network.</p>
Attribute Data	<p>SDE provides a transparent link between spatial objects and user-defined, nonspatial attribute data. Spatial objects can be stored or retrieved with selected attribute data in a single request using dynamic attribute masks. In addition to the standard relational data types, BLOBs may be stored as attribute data.</p>
SQL Support	<p>SDE supports the use of standard ANSI SQL 89 constraint clauses for all spatial object retrievals including spatial searches. The SQL clauses may constrain both aspatial</p>

attributes and spatial attributes such as area, length, or feature type.

Spatial Analysis

SDE provides a robust set of geometric and spatial analysis functions to be used again and again in applications such as real estate inquiry, buffering around environmentally protected areas, checking for environmental hazards, and so on. These functions determine the relationship between geographic features including whether they cross, share a common point or boundary, are equal, share a common area, or whether one is contained in another. In addition, SDE supports clipping, dissolve, buffer generation, distance calculation, polygon overlay, and network processing.

With these, and over 100 other functions, the possibilities for performing spatial analysis seem endless. You could query the database to locate all the land parcels that intersect the location of a water main; or locate the parcels that are partially or completely within 500 feet of a high voltage power line; or identify the areas in a county that are greater than five acres in size, are at least a half mile away from a major roadway, and a quarter mile away from any rivers or streams. These spatial queries can be further constrained by using a standard structured query language (SQL) "WHERE" clause against the attributes of the geographic features.

Output from spatial queries can be used in mapping applications, or in other applications that require geometric analysis without the mapping component. This means you can embed spatial analysis in an application without having to invoke traditional GIS technologies.

Rapid Spatial Searches

SDE supports a wide variety of spatial search methods to determine almost any kind of relationship between spatial entities. Spatial relationships are determined dynamically—*there is no need to prestore relationships*. The design of the database and spatial object model together with a unique search algorithm provide unprecedented performance in retrieval of spatial objects from the database. The performance achieved by SDE with very large databases would not be possible using conventional spatial retrieval methodologies and models.

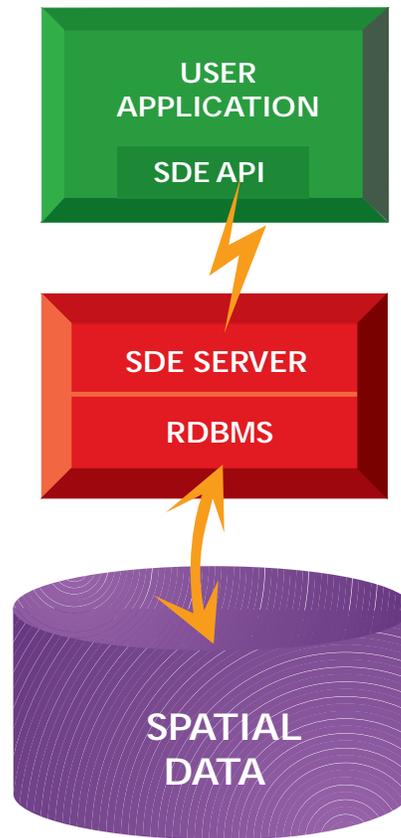
In addition to the standard SQL attribute constraint, a number of spatial relationships can be tested between two geographic features, a primary and a secondary. These relationships include the following:

- **Envelope test.** This search method only requires that the coordinate envelopes of the primary and secondary feature overlap or touch. The coordinate envelope of a feature consists of the minimum and maximum ranges of its coordinate values. This is the fastest search method because only an envelope test is performed. Typical usage is for screen displays and approximate searches.
- **Common point.** This search method specifies that the primary and secondary features must share at least one identical common point. A typical usage would be to find all parcels of land adjacent to a given parcel, assuming that topological integrity exists between parcels.

- **Line cross.** Primary and secondary features that cross the point of intersection cannot be a vertex in both features. A typical usage would be to find errors in topological integrity.
- **Common line.** This search method specifies that the primary and secondary features must share at least one identical common line.
- **Area intersect.** At least one of the features must be an area feature, and the other feature must be wholly or partially inside it. Mere adjacency is not considered to be an area intersect. Typical usage would be to determine the zoning of a parcel.
- **Edge touch or area intersect.** This search method specifies that the primary and secondary features must touch or cross; or one feature must be wholly or partially inside the other feature.
- **Area intersection without edge touch.** One feature is wholly inside the other (area) feature and their boundaries do not touch in any way.
- **Wholly contained.** If one feature is an area feature, the other feature must be wholly inside it, inclusive of the first feature's boundary. If both features are linear features, the one feature's path must be wholly included in the other's. If one feature is a point feature, it must be one of the other feature's vertices.
- **Wholly contained without edge touch.** One feature must be an area feature; the other feature must be wholly inside it, and their boundaries may not intersect or touch.
- **Point inside polygon.** Regardless of the primary feature type, this test performs a point-inside-polygon test using only the first point of the primary feature. The secondary must be an area feature and the point used for the test must be inside it. This method is typically used to identify area features by pointing.
- **Identical.** Both features must be identical in terms of feature type and coordinate description. This method is typically used to find duplicate data.

SDE Architecture

The architecture of the ESRI SDE reflects that of a database. SDE serves as the repository of a spatial object model and, as such, provides spatial services to client tasks. The services provided by SDE have a peer-to-peer relationship with the services provided by conventional relational databases. The client application is the coordinating agent among the various databases. SDE uses an underlying RDBMS (Oracle 7.2) that is invisible to the end user and application developer. The underlying relational database actually contains all of the spatial and aspatial attribute data stored under the domain of SDE. SDE resides above the RDBMS in a layered approach. The SDE API interface is the access point to all SDE functionality. SDE is tightly coupled to the underlying RDBMS with a database interface module. This module is the interface to all geographic and attribute data stored within the underlying relational database. The SDE interface module has been designed to keep a vendor-neutral relationship with the RDBMS to the greatest degree possible.



GIS Integrator's Toolset

SDE is provided without a graphical front end. This approach allows the graphical development tools to be separated from basic spatial functionality. As an integrator's toolset, SDE is, therefore, compatible with any and all of the commercially available graphic development tools such as the Extensible Virtual Toolkit (XVT).

Processing Model

SDE employs a true cooperative processing client/server model. Each client task is assigned a dedicated server task. This multithreaded design means that SDE is scalable; that is, it will automatically take advantage of servers which have multiple processors. The SDE server operates under the UNIX operating system on a variety of platforms. All server tasks are executed on the platform where the server is installed. The SDE system monitors all client and server connections and maintains system integrity.

Support for Multiple Computing Platforms

A product that is to be used in a large shared database environment must be able to support multiple computing platforms. Often client tasks reside on a different kind of computer than that of the server task. For data transfer between different computing platforms with compatible machine architecture (e.g., Data General and Sun), SDE uses standard TCP/IP protocol. The transfer of data between heterogeneous hardware platforms employs XDR (eXternal Data Representation) conversions to provide data compatibility between machines of different architecture.

For example, a server task running a UNIX server creates an XDR data representation for all data to be transferred to a PC client task. When the data are received by the PC client,

they are transparently converted to machine-specific format using a parallel XDR library function.

Client tasks may run in a variety of environments and on a variety of platforms:

- Client applications developed to run under Windows NT can run in native Windows NT mode or under DOS/Windows 3.1 by utilizing the WIN32s interface.
- Client applications may also run on a variety of UNIX platforms. Client tasks may also be run on the same platform as the SDE server.
- The developer can choose among a variety of environments for SDE applications. ArcView Version 3 offers GUI support for SDE as well as complete access to the SDE Client API through the Avenue object-oriented scripting language. MapObjects, an object linking and embedding (OLE) compliant custom control, offers SDE support. MapObjects can be used with Visual Basic, Delphi, Powerbuilder, and other popular Windows-based development environments. And, of course, the C programmer can use the Client API directly.

Performance

An important consideration of the SDE design team has been performance and transactional integrity. Since real-time information systems demand high performance, every effort has been made to maximize the speed of system response by taking advantage of current technology. SDE utilizes a client/server model which emphasizes asynchronous cooperative processing between the client and server tasks. This means that the server task actually predicts the demands of the client task in advance of the request. The SDE multithreaded architecture automatically takes advantage of multiprocessor hardware, and intelligent data buffering maximizes network throughput.

The cooperative processing client/server model employed by SDE is designed to provide maximum performance in the client/server environment:

- Operations that are particularly CPU intensive, such as buffer calculation and polygon overlay, are executed by the SDE client task to take advantage of workstation processing power. This avoids placing unnecessary burden on the SDE server and improves processing performance. The choice of processing location for any given function is automatic and transparent to the application developer.
- The SDE server task minimizes network congestion by performing all spatial searches and retrievals locally and transmitting to the client task only the features and data required by the application.
- Intelligent data buffering allows the client and server tasks to perform asynchronously. In effect, the server task predicts subsequent client data requirements, retrieves and buffers the required data, and is prepared to ship resultant data buffers across the network immediately upon request from the client task. This intelligent data buffering is transparent to the logic of the client task application. The client application, if desired, can dynamically "tune" the buffering characteristics during execution to improve performance based on characteristics of the spatial data.

ESRI GIS AND MAPPING SOLUTIONS

For more than 25 years ESRI has been helping people manage and analyze geographic information. Our solutions range from desktop mapping to enterprisewide GIS.

ArcView®—Easy-to-use GIS for Windows™, Macintosh, and a variety of UNIX platforms. ArcView software integrates spatial, tabular, and image data and is customizable and extensible.

ARC/INFO®—The World's GIS: automate, modify, manage, analyze, and display geographic information.

Spatial Database Engine™ (SDE™)—High-performance spatial database: true client/server architecture that performs fast, efficient, spatial operations and manages large, shared geographic data sets.

MapObjects™—An OLE Control and a collection of OLE automation objects that allow application developers to add mapping and GIS capabilities to their applications.

ArcCAD®—Brings the unmatched functionality of ARC/INFO software to AutoCAD®, the world's leading CAD software.

PC ARC/INFO®—Full-featured GIS for Windows and DOS-based PCs: comprehensive set of geoprocessing and cartography tools.

Data Automation Kit™—Geodata creation software: topological data creation and editing, data conversion, and map projection capabilities.

BusinessMAP™—Database mapping software that converts database records with ZIP Code fields to symbols on a United States map.

For more information, call
1-800-447-9778 (1-800-GIS-XPRT)
Send E-mail inquiries to info@esri.com
Visit ESRI on the World Wide Web at www.esri.com



CORPORATE OFFICE

ESRI
380 New York Street
Redlands, California
92373-8100 USA
Telephone: 909-793-2853
Fax: 909-793-5953

U.S. OFFICES

Alaska
Telephone: 907-344-6613
Boston
Telephone: 508-777-4543
California
Telephone: 909-793-2853, ext. 1906
Charlotte
Telephone: 704-541-9810
Denver
Telephone: 303-449-7779
Minneapolis
Telephone: 612-454-0600
Olympia
Telephone: 360-754-4727
Philadelphia
Telephone: 610-725-0901
San Antonio
Telephone: 210-340-5762
St. Louis
Telephone: 314-949-6620
Washington, D.C.
Telephone: 703-506-9515

INTERNATIONAL OFFICES

ESRI-Australia
Telephone: 61-9-242-1005
ESRI-Canada
Telephone: 416-441-6035
ESRI-France
Telephone: 33-1-46-23-6060
ESRI-Germany
Telephone: 49-8166-380
ESRI-Italia
Telephone: 39-6-406-961
ESRI-South Asia
Telephone: 65-735-8755
ESRI-Spain
Telephone: 34-1-559-4375
ESRI-Sweden
Telephone: 46-23-84094
ESRI-Thailand
Telephone: 66-2-678-0707
ESRI-(UK)
Telephone: 44-1-923-210-450

For the location of an international distributor, call 909-793-2853, ext. 1375