



Environmental Systems Research Institute, Inc., 380 New York St., Redlands, CA 92373-8100 USA • TEL 909-793-2853 • FAX 909-307-3014

ESRI Systems Integration Technical Brief

ArcSDE Tiered Hardware Configurations

Rev. 3 April 2, 2004

Copyright © 2004 ESRI
All rights reserved.
Printed in the United States of America.

The information contained in this document is the exclusive property of ESRI. This work is protected under United States copyright law and other international copyright treaties and conventions. No part of this work may be reproduced or transmitted 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 ESRI. All requests should be sent to Attention: Contracts Manager, ESRI, 380 New York Street, Redlands, CA 92373-8100, USA.

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

U.S. GOVERNMENT RESTRICTED/LIMITED RIGHTS

Any software, documentation, and/or data delivered hereunder is subject to the terms of the License Agreement. In no event shall the U.S. Government acquire greater than RESTRICTED/LIMITED RIGHTS. At a minimum, use, duplication, or disclosure by the U.S. Government is subject to restrictions as set forth in FAR §52.227-14 Alternates I, II, and III (JUN 1987); FAR §52.227-19 (JUN 1987) and/or FAR §12.211/12.212 (Commercial Technical Data/Computer Software); and DFARS §252.227-7015 (NOV 1995) (Technical Data) and/or DFARS §227.7202 (Computer Software), as applicable. Contractor/Manufacturer is ESRI, 380 New York Street, Redlands, CA 92373-8100, USA.

@esri.com, 3D Analyst, ADF, AML, ARC/INFO, ArcAtlas, ArcCAD, ArcCatalog, ArcCOGO, ArcData, ArcDoc, ArcEdit, ArcEditor, ArcEurope, ArcExplorer, ArcExpress, ArcFM, ArcGIS, ArcGrid, ArcIMS, ArcInfo Librarian, ArcInfo, ArcInfo—Professional GIS, ArcInfo—The World's GIS, ArcLogistics, ArcMap, ArcNetwork, *ArcNews*, ArcObjects, ArcOpen, ArcPad, ArcPlot, ArcPress, ArcQuest, ArcReader, ArcScan, ArcScene, ArcSchool, ArcSDE, ArcSdl, ArcStorm, ArcSurvey, ArcTIN, ArcToolbox, ArcTools, ArcUSA, *ArcUser*, ArcView, ArcVoyager, *ArcWatch*, ArcWeb, ArcWorld, Atlas GIS, AtlasWare, Avenue, *BusinessMAP*, Database Integrator, DBI Kit, ESRI, ESRI—Team GIS, ESRI—The GIS People, FormEdit, Geographic Design System, Geography Matters, Geography Network, GIS by ESRI, GIS Day, GIS for Everyone, GISData Server, *InsiteMAP*, MapBeans, MapCafé, MapObjects, ModelBuilder, MOLE, NetEngine, PC ARC/INFO, PC ARCPLOT, PC ARCSHELL, PC DATA CONVERSION, PC STARTER KIT, PC TABLES, PC ARCEDIT, PC NETWORK, PC OVERLAY, Rent-a-Tech, *RouteMAP*, SDE, SML, Spatial Database Engine, StreetEditor, StreetMap, TABLES, the ARC/INFO logo, the ArcAtlas logo, the ArcCAD logo, the ArcCAD WorkBench logo, the ArcCOGO logo, the ArcData logo, the ArcData Online logo, the ArcEdit logo, the ArcEurope logo, the ArcExplorer logo, the ArcExpress logo, the ArcFM logo, the ArcFM Viewer logo, the ArcGIS logo, the ArcGrid logo, the ArcIMS logo, the ArcInfo logo, the ArcLogistics Route logo, the ArcNetwork logo, the ArcPad logo, the ArcPlot logo, the ArcPress for ArcView logo, the ArcPress logo, the ArcScan logo, the ArcScene logo, the ArcSDE CAD Client logo, the ArcSDE logo, the ArcStorm logo, the ArcTIN logo, the ArcTools logo, the ArcUSA logo, the ArcView 3D Analyst logo, the ArcView Business Analyst logo, the ArcView Data Publisher logo, the ArcView GIS logo, the ArcView Image Analysis logo, the ArcView Internet Map Server logo, the ArcView logo, the ArcView Network Analyst logo, the ArcView Spatial Analyst logo, the ArcView StreetMap 2000 logo, the ArcView StreetMap logo, the ArcView Tracking Analyst logo, the ArcWorld logo, the Atlas GIS logo, the Avenue logo, the *BusinessMAP* logo, the Data Automation Kit logo, the Digital Chart of the World logo, the ESRI Data logo, the ESRI globe logo, the ESRI Press logo, the Geography Network logo, the MapCafé logo, the MapObjects Internet Map Server logo, the MapObjects logo, the MOLE logo, the NetEngine logo, the PC ARC/INFO logo, the Production Line Tool Set logo, the *RouteMAP* IMS logo, the *RouteMAP* logo, the SDE logo, The World's Leading Desktop GIS, *Water Writes*, www.esri.com, www.geographynetwork.com, www.gisday.com, and Your Personal Geographic Information System are trademarks, registered trademarks, or service marks of ESRI in the United States, the European Community, or certain other jurisdictions.

Other companies and products mentioned herein are trademarks or registered trademarks of their respective trademark owners.

Introduction

ArcSDE is part of a flexible GIS data management architecture and can be configured in either a two-tier or three-tier hardware configuration with multiple GIS client connection options. Choosing an appropriate hardware and connection configuration depends on several factors including available server capacity, high-availability requirements, licensing, and corporate configuration standard practices. Several of these factors will be discussed in the following sections as the different configuration options are presented.

ArcSDE Connections Options

Prior to discussing the various ArcSDE hardware options, the ArcSDE connection options require explanation. There are two client-to-server connection options available with ArcSDE: an Application Server Connect (ASC) option with ArcSDE middleware installed on the Database Server, and a Direct Connect (DC) option, which uses Database Management System (DBMS) client software to connect directly to the database without the use of ArcSDE middleware.

The ASC option is commonly referred to as a three-tiered *software* configuration and is not to be confused here with a three-tiered *hardware* configuration. This document will only refer to tiers in the context of hardware configurations.

Application Server Connect

Application Server Connect requires an ArcSDE service that is typically installed on the Database Server as part of the ArcSDE installation process. There are two primary processes that comprise an ArcSDE service. First is the GIOMGR process, which is the ArcSDE process that listens for GIS client connection requests. In response to a GIS client connection request, the GIOMGR spawns a separate GSRVR process for each client connection, which in turn connects to the DBMS and becomes the broker for the GIS client to database connection. The GSRVR process uses the existing DBMS server libraries on the Database Server to connect with the DBMS, therefore a DBMS client install is not required on the ArcSDE/Database Server. The traffic from the GIS client to the GSRVR process is comprised of ESRI application specific instructional commands requesting an operation. The traffic from each GSRVR process to the database is comprised of actual SQL statements required to fulfill the request.

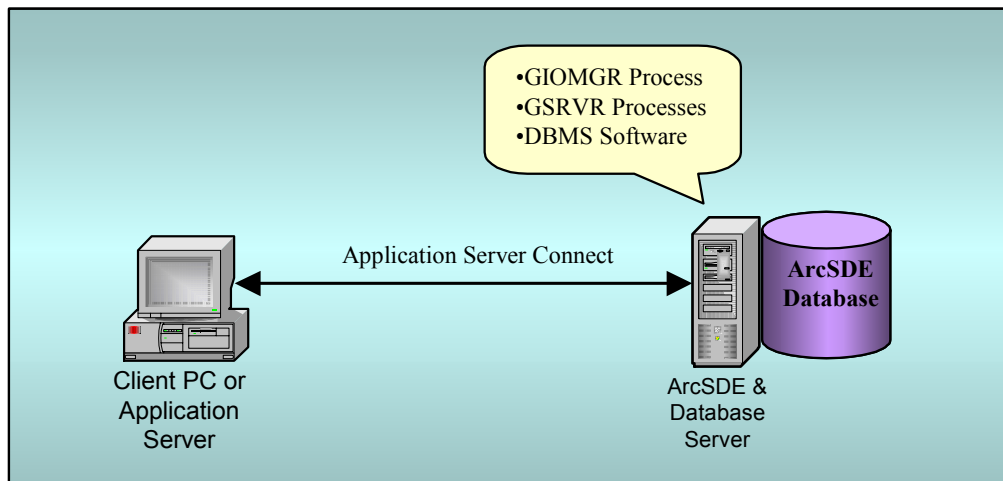
Direct Connect

Direct Connect does not require an ArcSDE service, though it does require that the same ArcSDE repository/schema be installed within the DBMS. With DC, the GSRVR functionality is integrated into the GIS client process in the form of a set of DLLs. The GSRVR-enabled client application process uses DBMS client software to communicate over the network directly with the DBMS. Therefore, the appropriate DBMS client software must be installed and configured on the GIS client. With DC, SQL statements are communicated directly between the DBMS client and the DBMS on the server.

Two-Tier Hardware: Application Server Connect (ASC) Configuration

An ArcSDE two-tier ASC hardware configuration consists of a GIS client system running the client application and a Database Server that is running the ArcSDE service and housing the GIS database (Figure 1).

Figure 1
Two-Tier Application Server Connect ArcSDE Configuration



Following are pros and cons for the two-tier ASC configuration:

Pros:

- Simplest to implement (requires ArcSDE installation on the Database Server).
- No additional client configuration required.
- Assuming equivalent CPUs on the client/server systems, ASC decreases GIS client CPU load by 5-10% relative to DC, which reduces client application server scalability (ArcIMS, ArcGIS Server, Terminal Server, etc.).
- Prior to ArcSDE 9.0, ASC network traffic was roughly half that of DC. At ArcSDE 9.0, depending on the DBMS and data type, much less network traffic is utilized and can be nearly equal to, or in some cases, less than ASC.
- Client and Server application backward compatibility is supported. GIS applications can be upgraded without upgrading ArcSDE and the database schema. This is typically true but may depend on the particular software release.

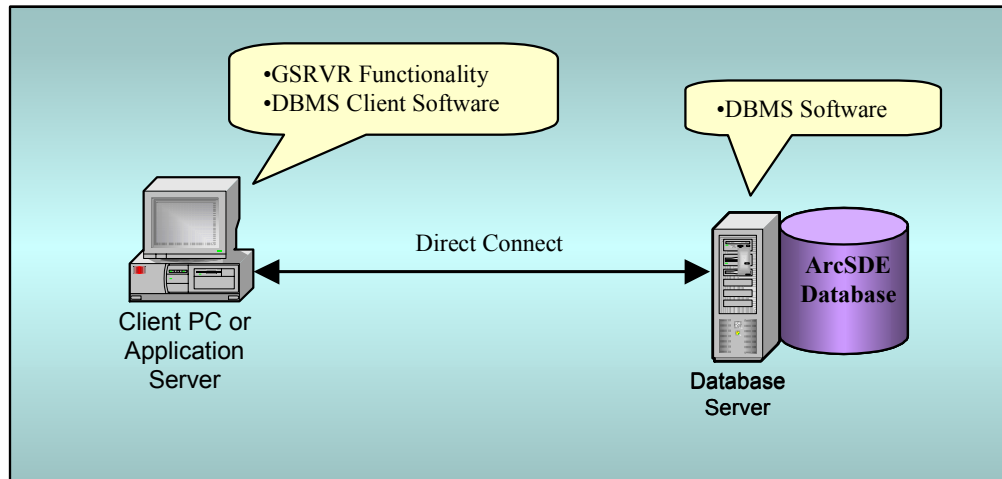
Cons:

- Assuming equivalent CPUs on the client/server systems, ASC increases Database Server CPU load by roughly 20-40% relative to DC, which reduces Database Server scalability.
- Increases memory requirements on the Database Server due to the GIOMGR and GSRVR processes.
- Does not support Oracle's Real Application Cluster (RAC) with Transparent Application Failover (TAF). However, it does support non-transparent failover clustering such as Microsoft's Cluster Services.

Two-Tier Hardware: Direct Connect (DC) Configuration

An ArcSDE two-tier DC hardware configuration consists of a client system running the client GIS application and the associated DBMS client software. The Database Server does not require an ArcSDE service (Figure 2).

Figure 2
Two-Tier Direct Connect ArcSDE Configuration



Following are pros and cons for the two-tier DC configuration:

Pros:

- Assuming equivalent CPUs on the client/server systems, DC reduces Database Server CPU load by 20-40% relative to ASC due to off-loading of GSRVR processes, which provides an increase in Database Server scalability.
- Reduces memory requirements on the Database Server by off-loading the GSRVR process to the client GIS application.
- Eliminates the requirement to run application middleware (ArcSDE) on the Database Server. ArcSDE install is still required to create repository, etc.
- Supports Oracle RAC/TAF.

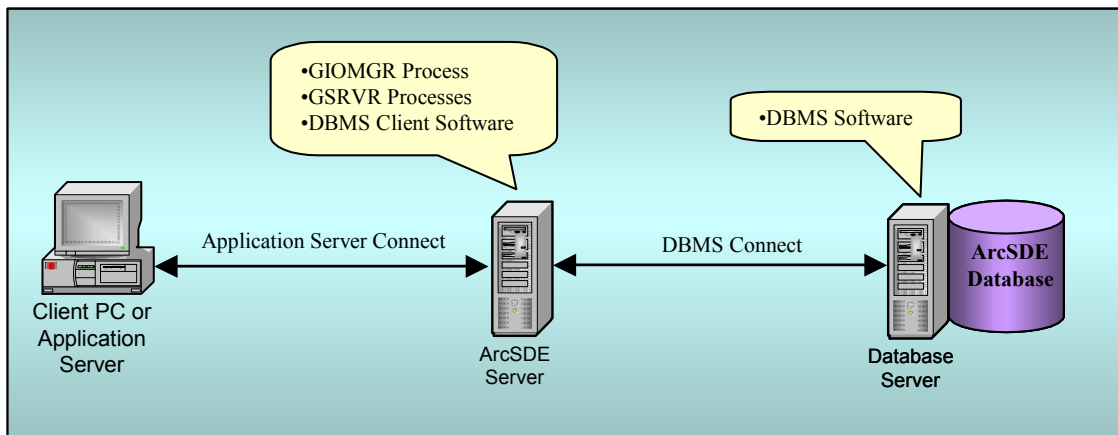
Cons:

- Requires DBMS client software installation and configuration on each GIS client.
- Assuming equivalent CPUs on the client/server systems, DC increases GIS client CPU load by 5-10% relative to ASC, which reduces client application server scalability (ArcIMS, ArcGIS Server, Terminal Server, etc.).
- Prior to ArcSDE 9.0, DC network traffic was roughly twice that of ASC. At ArcSDE 9.0, depending on the DBMS and data type, much less network traffic is utilized and can be nearly equal to, or in some cases, better than ASC.
- Client and Server application backward compatibility is not supported. GIS client software must be upgraded concurrently with a database schema upgrade.

Three-Tier Hardware: Application Server Connect (ASC) Configuration

For an ArcSDE three-tier hardware configuration (Figure 3), ArcSDE middleware is installed on one or more middle-tier ArcSDE Application Servers. This effectively creates an ASC from the GIS client to the ArcSDE Server and a DBMS connect from the ArcSDE Server to the Database Server. An example of when this configuration would be used is where it is not desirable to install application middleware on the Database Server and where DC is not a feasible option. The appropriate DBMS client software must be installed and configured on the middle-tier ArcSDE Server to support the DBMS connection to the Database Server.

Figure 3
Three-Tier Application Server Connect ArcSDE Configuration



Following are pros and cons for the three-tier ASC configuration:

Pros:

- Reduces Client/Server CPU requirements by off-loading the GSRVR processes to a middle-tier server, which increases Database Server scalability.
- Reduces memory requirements on the Database Server by off-loading the GSRVR process to the client GIS application.

Cons:

- Most complex configuration to implement and requires additional middle-tier server hardware that must be sized to handle GSRVR CPU processing and memory requirements. According to current system design models, since 20% of the server-side load is GSRVR processing, as many as 100 users per CPU would be supported at the middle tier. Memory requirements at the middle tier would vary depending upon the application but a good starting point is 1-2 GB per CPU.
- Increases network traffic due to additional ArcSDE Server DBMS connection to the Database Server. A network bottleneck can be avoided by creating a dedicated network connection between the ArcSDE and Database Server.

- High-availability requirements would drive the need for redundant middle-tier servers and a method to load balance client load. If occasional middleware and database application mix is allowed, both systems could be configured in the same cluster where each server acts as a backup for the other.
- Though this configuration supports Oracle's RAC at the Database Server tier, TAF would not be available if a middle-tier ArcSDE node failed since all the GSRVR connections would be lost, even if the middle-tier ArcSDE servers were participating in an active/passive cluster.
- Overall performance may vary depending on the number of CPUs and their relative performance at the various tiers. A performance bottleneck in the middle or back-end tier will hamper overall system performance.

Conclusion

ArcSDE supports several multi-tiered software and hardware configurations and different client connection types. Choosing the appropriate configuration depends on various technical and business goals and constraints. When choosing a configuration, it is important to consider network, memory and CPU loads to avoid potential performance bottlenecks. Support and installation issues should be considered as well.

Support

Enterprise GIS system design is addressed in the *System Design Strategies* white paper at <http://www.esri.com/library/whitepapers/pdfs/sysdesig.pdf>. For answers to additional GIS capacity planning and solution questions, contact ESRI Systems Integration at sihelp@esri.com. For technical support, contact ESRI Technical Support at <http://support.esri.com>.