

Microsoft TerraServer

SQL Server 7.0

Tom Barclay, Robert Eberl, Jim Gray,
John Nordlinger, Guru Raghavendran,
Don Slutz, Greg Smith, Phil Smoot
Microsoft Research and Development

John Hoffman, Natt Robb III
Aerial Images

Hedy Rossmeyssl, Beth Duff, George Lee,
Theresa Mathesmier, Randall Sunne
United States Geological Survey

Lee Ann Stivers, Ken Goodman
Digital Equipment Corporation

June 1998

Summary: The Microsoft® TerraServer stores aerial and satellite images of the earth in a Microsoft SQL Server™ database served to the public through the Internet. It is the world's largest atlas, combining five terabytes of image data from the United States Geodetic Survey, SOVINFORMSPUTNIK, and Encarta® Virtual Globe. Internet browsers provide intuitive spatial and gazetteer interfaces to the data. The TerraServer demonstrates the scalability of Microsoft Windows NT® Server, Enterprise Edition version 4.0 and SQL Server, Enterprise Edition running on Digital hardware including the AlphaServer 8400 and StorageWorks™ storage system. The TerraServer is also an E-Commerce application. Users can buy the right to use the imagery using Microsoft Site Servers managed by the USGS and Aerial Images. This paper describes the TerraServer's design and implementation.

Contents

[The Microsoft TerraServer](#)

[Database Design](#)

[Microsoft TerraServer Hardware](#)

[Assessment](#)

[Summary](#)

[Acknowledgments](#)

[For More Information](#)

The Microsoft TerraServer



Figure 1. The TerraServer hardware

The TerraServer has five terabytes of uncompressed satellite and aerial images of urban areas, compressed to one terabyte of database data. It serves these images onto the Internet with a graphical and intuitive user interface. The application demonstrates several things:

- **Information at your fingertips.** This is the most comprehensive world atlas anywhere—and it is available to anyone with access to the Internet.
- **Windows NT Server, Enterprise Edition and SQL Server, Enterprise Edition version 7.0 scale.** The TerraServer fills eight large cabinets: one for the Digital Alpha 8400 processors, and seven cabinets for the 324 disks—almost three terabytes (TB) of raw disk storage and 2.3 TB of RAID5 storage.
- **Windows NT and SQL Server, Enterprise Edition are excellent for serving multimedia and spatial data onto the Internet.**
- **Microsoft Site Server Commerce Edition can help sell images over the Internet.**

TerraServer is a multimedia database that stores both classical text and numeric data, as well as multimedia image data. In the future, most huge databases will be comprised primarily of document and image data. The relational metadata is a relatively small part of the total database size. TerraServer is a good example of this new breed of multimedia databases.

The Application

An Interesting Internet Server. TerraServer is designed to be a compelling Internet application. It tries to be interesting to almost everyone, everywhere, to be offensive to no one, and to be relatively inexpensive to build and operate. It is hard to find data like that—especially a terabyte of such data. A terabyte is nearly a billion pages of text—four million books. A terabyte holds 250 full-length movies. It is a lot of data.

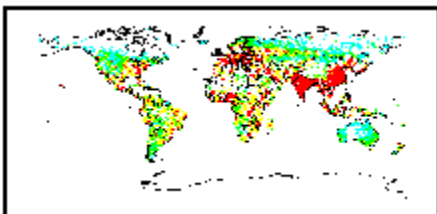


Figure 2. World population density

Satellite Images of the Urban World. Pictures have a universal appeal, so it was natural to pick a graphical application. Aerial images of the urban world seemed to be a good application. The earth's surface is about 500 square tera-meters. 75 percent is water, 20 percent of the rest is above 70° latitude. This leaves about 100 square tera-meters. Most of that is desert, mountains, or farmland. Less than 4 percent of the land is urban. The TerraServer primarily stores images of urban areas. Right now, it has nearly five square tera-meters—and it grows as more data becomes available.

Cooperating with the United States Geological Survey (USGS): The USGS has published aerial imagery of many parts of the United States. These images are approximately one-meter resolution (each pixel covers one square meter). We have a Cooperative Research Agreement (CRADA) with the USGS to make this data available to the public. We have loaded all the published USGS data (3 TB raw, 0.6 TB compressed). This is 30 percent of the United States. As additional data becomes available, it will be loaded into the TerraServer. This data is unencumbered and can be freely distributed to anyone. It is a wonderful resource for researchers, urban planners, and students. The picture at left shows a baseball game in progress near San Francisco. You can see the cars, but one-meter resolution is too coarse to show people.




Figure 3. A USGS 1-meter resolution image of Candlestick Park near San Francisco

Working with SOVINFORMSPUTNIK (the Russian Space Agency) and Aerial Images. To be interesting to everyone everywhere, TerraServer must have worldwide coverage. The USGS data covers much of the continental United States. There is considerable imagery of the planet, but much of it either has poor quality (10 meter to 1-km resolution), has not been digitized, or is encumbered. SOVINFORMSPUTNIK and their representative, Aerial Images, have some of the best data and were eager to cooperate. The Russians and Aerial Images contributed two square tera-meters of imagery (1.56-meter resolution). This data is trademarked SPIN-2, meaning satellite-2-meter imagery. They intend to deliver an additional 2.4 square tera-meters over the next year.



Figure 4. A SPIN-2 1.6-meter image of Atlanta's Olympic stadium

TerraServer is the largest world atlas. The SOVINFORMSPUTNIK SPIN-2 imagery covers Rome, Athens, Hong Kong, New York, Chicago, Seattle, and many other cities. TerraServer has more data in it than *all* the HTML pages on the

Internet. If printed in a paper atlas, with 500 pages per volume, the information would fill a collection of 2,000 volumes. It grows by 10,000 pages per month. Clearly, this atlas must be stored online. The USGS data (the three square tera-meters) is seven times larger. This data is a world asset that will likely change the way geography is taught in schools, the way maps are published, and the way we think about our planet.

TerraServer as a business. Slicing, dicing, and loading the SPIN-2 and USGS data is a continuing process. Today, the TerraServer stores a terabyte. Aerial Images, Digital, and Microsoft are operating the TerraServer on the Internet (<http://terraserver.microsoft.com/>). Microsoft views TerraServer as a demonstration of the scalability of Windows NT Server and Microsoft SQL Server, Enterprise Edition. Digital views it as a demonstration of their Alpha and StorageWorks servers. The USGS is participating as an experiment to present USGS data to a wider audience through the Internet. They operate an online store that allows anyone to download copies of the USGS images. SOVINFORMSPUTNIK and Aerial Images view TerraServer as a try-and-buy distributor for their intellectual property. They make coarse-resolution (8-meter, 16-meter, and 32-meter) imagery freely available. The fine-resolution data is viewable in small quantities, but customers must buy the right to use the good imagery. All the SPIN-2 images are watermarked, and the high-resolution images are lightly encrypted.

Site Server Commerce Edition, a new business model for the Internet. Aerial Images' business model is likely to become a textbook case of Internet commerce. Because they use the Internet to sample and distribute their images, Aerial Images has very low distribution costs. This allows them to sell imagery in small quantities and large volumes. Microsoft helped USGS and Aerial Images set up Microsoft Site Servers that accept credit-card payments for the imagery. A **Download** button on the image page takes the user to these Site Servers (Microsoft has no financial interest in these transactions). You can buy a detailed image of your neighborhood for a few dollars.

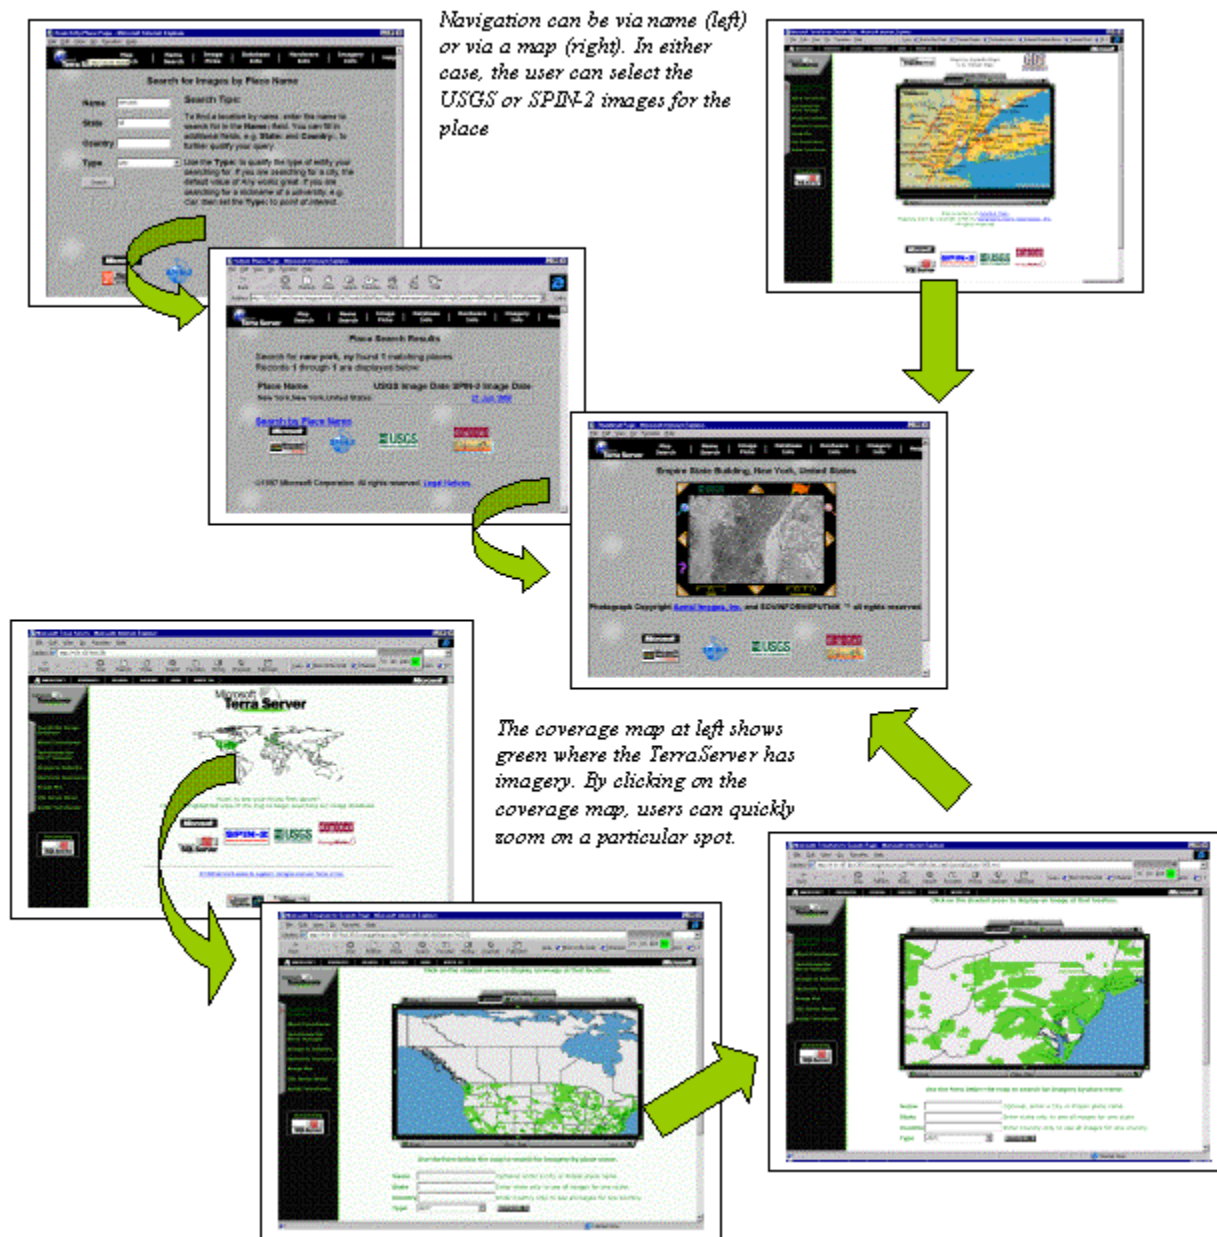


Figure 5. Navigation through TerraServer

User Interface to the Microsoft TerraServer

Navigation through database searches. The TerraServer can be accessed from any Web browser (for example, Internet Explorer, Netscape Navigator). Full resolution SPIN-2 imagery requires the Web browser to support Java applets. Any Web browser that supports HTML tables and display of Jpeg data can host the TerraServer user interface. Navigation can be spatial through a point-and-click map control based on Microsoft's Encarta World Atlas. Clients only knowing the place name can navigate textually by presenting a name to the Encarta Virtual Globe Gazetteer. The gazetteer knows the names and locations of 1.1 million places in the world. For example, "Moscow" finds 28 cities, while "North Pole" finds 5 cities, a mining district, a lake, and a point-of-interest. There are 378 listings for San Francisco in the gazetteer. The user can select the appropriate member from the list. The map control displays the 40-km map of that area. The user can then pan and zoom with this map application and can select the USGS and SPIN-2 images for the displayed area.

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.