

# The Microsoft TerraServer™

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

Lee Ann Stivers, Ken Goodman  
Digital Equipment Corporation

June 1998

Technical Report  
MSR-TR-98-17

Microsoft Research  
Microsoft Corporation  
One Microsoft Way  
Redmond, WA 98052

cs.DB/9809005 2 Sep 1998

# Microsoft TerraServer™

---

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



Lee Ann Stivers, Ken Goodman  
Compaq Corporation



June 1998

## Abstract

The Microsoft TerraServer stores aerial and satellite images of the earth in a SQL Server Database served to the public via the Internet. It is the world's largest atlas, combining five terabytes of image data from the United States Geodetic Survey, Sovinform Sputnik, and Encarta Virtual Globe™. Internet browsers provide intuitive spatial and gazetteer interfaces to the data. The TerraServer demonstrates the scalability of Microsoft's Windows NT Server and SQL Server running on Compaq AlphaServer 8400 and StorageWorks™ hardware. 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.

## Table of Contents

<b><i>The Microsoft TerraServer™</i></b> .....	<b>1</b>
The Application.....	1
User Interface to the Microsoft TerraServer.....	4
Server Design .....	6
<b><i>Database Design</i></b> .....	<b>9</b>
Coordinate Systems .....	9
Image Format .....	11
Database Themes.....	13
Logical Database Design for Image Data .....	15
Database Design for Gazetteer .....	17
Physical database design .....	20
Data Loading .....	21
Backup & Restore .....	23
<b><i>Microsoft TerraServer Hardware</i></b> .....	<b>25</b>
<b><i>Assessment</i></b> .....	<b>27</b>
<b><i>Other Scalability Projects</i></b> .....	<b>28</b>
<b><i>Summary</i></b> .....	<b>29</b>
<b><i>Acknowledgments</i></b> .....	<b>30</b>
<b><i>References</i></b> .....	<b>32</b>



## The Microsoft TerraServer™

---

The TerraServer has five terabytes of 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 and SQL Server can scale up to huge nodes:** The TerraServer fills eight large cabinets: one for the Compaq Alpha 8400 processors, and seven cabinets for the 324 disks -- almost three terabytes of raw disk storage and 2.3 TB of RAID5 storage.
- **Windows NT and SQL Server are excellent for serving multi-media and spatial data onto the Internet.**
- **Microsoft Site Server can help sell images over the Internet.**



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

### The Application

**An Interesting Internet Server:** TerraServer is designed to be compelling Internet application. It tries to be interesting to almost everyone everywhere, be offensive to no one, and 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 — 4 million books. A terabyte holds 250 full-length movies. It is a lot of data.



**Satellite Images of the Urban World:** Pictures are big and 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% is water, 20% 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% 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.

# 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.