



White Paper

Performance and Recommended Setup for Autodesk MapGuide

When you plan and set up your Autodesk MapGuide® application, you will want to maximize performance on your site. There are several factors that affect performance. This white paper discusses those factors and helps you make choices about your setup.

How Much Traffic Will Your Website Get?

The most important factor affecting performance is traffic at your website. In particular, you need to determine how many requests per second you will need to support during peak periods on your site. To get an idea of how many requests you will have, ask yourself these questions:

- **How many users will the site support?** For example, will it support the public and potentially thousands of users, or will there just be a handful of users?
- **How often will users visit the site?** Do field technicians check in at the site just once in the morning, or will users be accessing the site throughout the day?
- **How many requests will a user require during each visit?** For example, will users be able to look up information in a few steps, or will they need to pan, zoom, generate reports, and perform other tasks that send requests to the server?
- **When will users visit the site?** For example, will there be peak hours when everyone visits, or will requests be spread more evenly throughout the day? For example, if technicians tend to visit the site during a three-hour period in the morning, but most of them access the site during the same 30-minute period, those 30 minutes would be considered your peak period.

After answering these questions, you should be able to determine the maximum number of requests per second you'll need to support. Multiply the number of users you estimate will hit during your peak period (in hours) by the number of times they will visit during those hours, then multiply that number by the number of requests you estimate they'll make to the server during each visit. Lastly, divide that number by the number of seconds in the peak period (one hour equals 3600 seconds). The equation looks like this

$$\frac{\text{users} \times \text{visits per user} \times \text{requests per visit}}{\text{peak hours} \times 3600} = \text{requests per second}$$

For example, assume that your site supports 100 users, 80 of whom visit the site once a three-hour period, and that each one's tasks result in about 10 requests per visit. The calculation would be as follows:

$$\frac{80 \times 1 \times 10}{3600} = 0.22 \text{ requests per second}$$

3 x 3600

. This is a rather small number, so you can assume that so far, you should be fine with the minimal hardware and software requirements listed in the *Autodesk MapGuide User's Guide*. However, this methodology is very inexact, so you should round up, particularly if the peak hours also correspond with peak hours on your intranet or the Internet, in which case performance will be slowed. Also, as you will see in the next section, the number of requests per second is only one factor. The type of data you are serving is also a major factor.

What Type of Data Do You Have?

In addition to simultaneous requests, you need to take into consideration the type of data you are serving. Are your maps primarily vector data, which loads quickly, or do you have lots of raster images (such as satellite photos), which are much larger and therefore take more time to process? Map authors can improve performance by including only pertinent data on each layer, but if you have a lot of raster data to serve, you will need more server power than if you had only vector data.

You also need to determine whether your data set is simple, medium, or complex. This can be a difficult generalization to make, but it will help you in determining the server configuration that will maximize performance—the next step of the process. Use these guidelines to help you determine your data's complexity:

- **Simple:** Data includes small vector data files, raster images, and local data sources. Maps use simple themes and have no more than 5 layers visible simultaneously, with less than 100 features visible simultaneously per layer.
- **Medium:** Data includes larger vector data files, several megabytes of raster images, small to midsize client-server or local data sources, and small Autodesk® GIS Design Server or Oracle® data sets. Maps use more complex themes and have 5 to 10 layers visible simultaneously, with 100 to 500 features visible simultaneously per layer.
- **Complex:** Data includes very large vector data files, very large raster images and catalogs, large client-server data sources, and large Autodesk GIS Design Server or Oracle data sets. Maps use complex themes and have more than 10 layers visible simultaneously, with more than 500 features visible simultaneously per layer.

What Configuration Will You Need?

Now that you have an idea of how many requests per second you will have to process and are aware of your data's complexity, you can begin to plan your server configuration.

In the following table, find your data complexity in the first column, and then find your requests per second in that row. This will put you into one of three categories: small, medium, or power.

	Small	Medium	Power
Simple map/data	1 request/second	2 requests/second	4+ requests/second
Medium map/data	.5 request/second	1 request/second	2+ requests/second
Complex map/data	.25 request/second	.5 request/second	1+ requests/second

The following are general guidelines for minimum recommended system configurations for small, medium, and power configurations:

- **Small:** Single server with 200MHz+ Pentium II processor and 128MB of RAM.
- **Medium:** One or two servers with 400MHz+ Pentium III processors (consider using dual processors) and 128–256MB of RAM each. If you use more than one server, set them up for load balancing.
- **Power:** Two or more servers with dual 600MHz+ Pentium III processors set up for load balancing, 256MB of RAM each, with raster data on one server and vector data and databases on another. If data is of medium or high complexity, store file-based databases on the same server as the vector data and store server-based databases on a separate server altogether.

Load balancing means setting up multiple servers so that requests are distributed equally among them. If you have problems with outages, or if it is absolutely critical that your site not go down, set up a *mirrored architecture*. For example, if you have two servers, you could add a second server that exactly duplicates each server's structure. Then if one of the servers goes down, a backup server can take its place. You can also combine a mirrored architecture with load balancing. For more information on performance enhancements, including using multiple servers, refer to the "Architecture and Performance" section in the *Autodesk MapGuide User's Guide*.

Other Server, Network, and Database Factors

You now have an idea of the minimum server configuration you will need to maximize performance. However, several other factors also affect performance. This section provides a list of these factors, which you should consider when planning your setup.

Server

- **CPU**
- **RAM**
- **Disk performance**
- **Virtual memory**
- **Other applications running on the server:** Does the server host your mail server, Oracle Server, SQL Server, or other server? If so, consider moving Autodesk MapGuide Server to another computer.

Network/Internet

- **LAN/Internet connection speed:** How quickly can Autodesk MapGuide talk to other servers and clients?
- **Network traffic/load:** How many other simultaneous requests are the web server, Internet, and network handling? What is peak load and when does it occur?
- **Proxy server/firewall settings**
- **Web server:** Is it set up efficiently?
- **MapAgent:** Are you using the fastest MapAgent for your web server? Use the ISAPI or NSAPI MapAgent where applicable.
- **File storage:** Are you storing many of the files Autodesk MapGuide Server needs (SDFs, raster files, database files) locally on the server? Or are they being retrieved from a shared network location?

Databases

- **OLE DB providers versus ODBC drivers:** With Autodesk MapGuide 6, you can still use your existing ODBC drivers via the OLE DB provider for ODBC. However, ODBC is generally much slower than OLE DB, and adding the extra layer of the OLE DB provider for ODBC will further decrease performance. Whenever possible, you should eliminate ODBC connections and use a supported OLE DB provider to connect Autodesk MapGuide Server to your data source.
- **Connection speed:** How fast is your connection to the database server (Oracle Server, SQL Server, etc.)?
- **Database server speed:** How quickly can your database server respond? Is it under constant heavy use? Can you break out your data to another database server that is less heavily loaded?
- **Number of servers:** Is the database server on the same computer as Autodesk MapGuide Server? The rule of thumb is to avoid running your client-server database server, such as Oracle Server or SQL Server, on the same computer as Autodesk MapGuide Server. Running them on the same computer decreases performance, because Autodesk MapGuide Server and the database server are competing for the same memory and CPU resources to service requests. Conversely, if you are using smaller, file-based desktop data sources—such as Microsoft® Access, Microsoft Excel, dBASE®, or text/CSV—try to keep the data files local to the server computer to optimize file I/O speed.
- **Database design:** Are you using indexes, keys, and database design optimizations to your advantage?
- **SQL complexity and optimization:** Is your SQL optimized for your database? Techniques will vary depending on your database vendor and your data; however, optimized SQL can greatly improve database response times.
- **Database efficiency:** Are you using an efficient database for your data and data access needs? Ensure that the database used suits the purpose, size, and complexity of your data. For example, if you expect your data to grow extensively and be shared with multiple applications or users, you will want to use an enterprise-class client-server database rather than a local desktop database. On the other hand, if you have a relatively small quantity of data and Autodesk MapGuide is the only or one of a very few applications using the data, it will be less costly and more efficient to use a local desktop database to service your data requests.

Other Performance Improvement Measures

In addition to optimizing your hardware setup, you can take several other steps with maps and data to improve performance.

1. Optimize Raster Layers

Raster image files can be quite large. Thus, map authors should work efficiently when creating raster map layers. There are several ways to optimize the performance of raster layers:

- **Limit visibility by scale.** In most cases, you should author your maps so that raster layers are visible only at large scales (when zoomed in). Limiting the scale at which a raster layer is visible will reduce the amount of data that the Autodesk MapGuide Server must transmit. In addition to setting limits on the layers, you may want to set scale limits on the raster image files themselves, to avoid inadvertently creating layers that

display raster data when zoomed out. You can set scale limits on raster image files and other resources using Autodesk MapGuide Server Admin.

- **Subsample large images.** If you need to display large raster data sets at small scales (when zoomed out), you should consider subsampling the raster images. Subsampling reduces the amount of detail and data in the file. For best results, you can use the subsampled image when the map is zoomed out and the original image when zoomed in. For more information, see step 4, “Generalize and Subsample Data.”
- **Create tiled TIFFs.** You can improve performance by converting large raster images to the tiled TIFF format. Tiling images in TIFF format causes Autodesk MapGuide Server to use much less memory and allows it to randomly access different parts of the image better. To convert images to tiled TIFF format, use Autodesk MapGuide Raster Workshop or Autodesk® CAD Overlay®; for more information, contact your nearest Autodesk reseller.

2. Limit Visibility by Scale

You can limit the visibility of data in map layers in three ways:

- **Make dynamic map layers visible by scale.** To reduce the amount of data transferred, use Autodesk MapGuide Author to set up map layers to display at a higher zoom level (zoom in closer). For example, when a map layer’s visibility is set to the scale of 1:25,000, it downloads only 25 percent of the data that is downloaded when visibility is set to 1:50,000.
- **Make dynamic map layers available but not visible.** In Autodesk MapGuide Author, you can define layers that do not initially appear on the map but that the user can turn on and view as needed. These layers appear in the Autodesk MapGuide Viewer legend without a check mark. Autodesk MapGuide Author and Autodesk MapGuide Viewer request the data from Autodesk MapGuide Server and display the layer only when the layer name is checked.
- **Limit resources by scale.** In Autodesk MapGuide Server Admin, you can specify a scale limit for a resource, so that all map layers that reference the resource can display its data only when the map is at the specified scale or larger (zoomed in).

To implement this suggestion, in Autodesk MapGuide Server Admin use the Resources tab of the Properties dialog box to edit an existing resource or add a new one. Then enter the scale limit in the Scale Limit text box.

3. Use Dynamic and Static Map Layers

Dynamic layers store minimal information in the map window file (in MWF format) and download map layer data only as required. To improve performance, use dynamic map layers for large data sets and reserve static map layers for small, frequently accessed data sets, such as initial views and overview maps. Static layers store their data in the MWF file, so the data is downloaded from the server only once, when the author saves the MWF file. After that, data is read locally from the MWF file, nearly eliminating delay when the user returns to the static layers.

Note: Data stored in the MWF files is not as secure as that stored in dynamic layers, so use dynamic layers whenever the data is highly sensitive. When using dynamic layers, be sure to limit visibility as described in step 2, “Limit Visibility by Scale.”

4. Generalize and Subsample Data

Generalizing and subsampling data both speed access by reducing the size of the files that Autodesk MapGuide Server must transfer. You generalize data in structured display files

(SDFs) and subsample data in raster images. Generalizing SDF data eliminates line detail because the number of vertices is reduced. Subsampling raster images decreases the resolution. The following are some tips for generalizing or subsampling and working with your data:

- Use the Autodesk MapGuide SDF Loader to generalize the data in SDFs. For more information, refer to the Autodesk MapGuide SDF Loader help.
- Use the Autodesk MapGuide Raster Workshop to add your raster images to a raster image catalog file and subsample them. For more information, refer to the Autodesk MapGuide Raster Workshop help.
- Set up maps to use a generalized or subsampled version of data when zoomed out and the original version when zoomed in. To do this, create two map layers for the same data. On the first layer, use the generalized or subsampled data. Then set the display range so that the layer displays when zoomed out but becomes unavailable when zoomed in to a specified point. On the second layer, use the original data, and set the display range so that the layer becomes available when zoomed in to the point where the first layer stops.

For example, you might set the first layer's display range at 20,000,000–40,000,000, which would cause the generalized or subsampled data to display when the user is zoomed out on the map within that range. You would then set the second layer's display range from 1–19,999,999, which would cause the original version of the same data to display when the user is zoomed in within this second range.

For an idea of how much this can improve performance with raster images, consider a 50MB color GeoSpot image measuring 4000 x 4000 pixels and covering 144 square kilometers. You could subsample this image to 1000 x 1000, which would reduce the file size to 4MB. Then set up two map layers: one that references the original image when the map is zoomed in, and one that references the subsampled image when the map is zoomed out. This would prevent the server from having to send an additional 46MB of unnecessary data when the map is zoomed out.

5. Limit Data in Structured Display Files

When creating SDFs, make sure to use only one type of feature in each SDF, such as points in one and polylines in another. Layers can use only one type of feature data, so including multiple types of features in a single SDF simply slows access to the file. The best approach is to have a one-to-one relationship between SDFs and layers, where each file contains only the data needed for its corresponding layer. However, if two layers use some of the same data, it is better to have all the data for both layers in a single SDF rather than duplicating data in two files, to avoid problems when updating the data.

6 Data Access Methods

In general, OLE DB providers offer the best Autodesk MapGuide Server data access performance when compared to ODBC or DWG. The SDF and SHP providers, in particular, are top performers and are very close in their speed. No performance benefit is gained by converting the SHP data to SDF or vice versa.

Autodesk MapGuide Provider for Oracle is slightly slower performance-wise than either the SDF or SHP provider. Using an ODBC driver is the slowest data access method. Whenever possible, use the associated OLE DB provider instead of the ODBC driver to maximize performance.

Conclusion

We hope that the guidelines in this white paper help you to determine your performance needs and give you ideas for ways to improve the performance of your Autodesk MapGuide implementation. To discuss questions or your own experiences with fellow Autodesk MapGuide users, visit the discussion group at www.autodesk.com/mapguide-discussion.

autodesk

Autodesk, Inc.
111 McInnis Parkway
San Rafael, CA 94903
USA

Autodesk, Autodesk MapGuide, and CAD Overlay are either registered trademarks or trademarks of Autodesk, Inc., in the USA and/or other countries. All other brand names, product names, or trademarks belong to their respective holders.

© Copyright 2001 Autodesk, Inc. All rights reserved.