



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

ArcIMS Configuration Performance Factors

November 12, 2003

Copyright © 2003 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

Several important design and configuration issues must be considered when building an ArcIMS configuration. For example, the servers must be sized correctly to handle expected peak transaction rates, enough network capacity must be in place to handle peak download rates, and the site must be tuned to perform efficiently. This technical brief will explore several key ArcIMS configuration factors including the impact those factors have on performance.

Background

There are several ArcIMS configuration areas that can be addressed that have an impact on performance. Covering all of the areas would go beyond the scope of this Technical Brief, which is to cover several factors that can have a significant impact on performance. The ArcIMS configuration areas that will be addressed are listed below and relate to system architecture design, service configuration, and site configuration including the virtual servers associated with creating maps. Other available virtual servers, such as the Metadata Server and Query Server, will not be addressed. Also, areas involving application design will not be addressed.

- Map Server Capacity Planning
- Virtual Server Type
- Spatial Servers and Threads
- Data Source Type
- Output Image Type
- Output Image Resolution
- JPEG Compression Quality
- Using ArcMap as a Map Service Client

Map Server Capacity Planning

ArcIMS map server capacity planning, including tiered configuration options, high-availability options, firewall options, and server sizing are addressed in Sections 6.0 and 7.0 of the *System Design Strategies* white paper.

<http://www.esri.com/library/whitepapers/pdfs/sysdesig.pdf>

Virtual Server Type

There are three main ArcIMS Virtual Servers involved with creating maps: Image Server, ArcMap Server, and Feature Server.

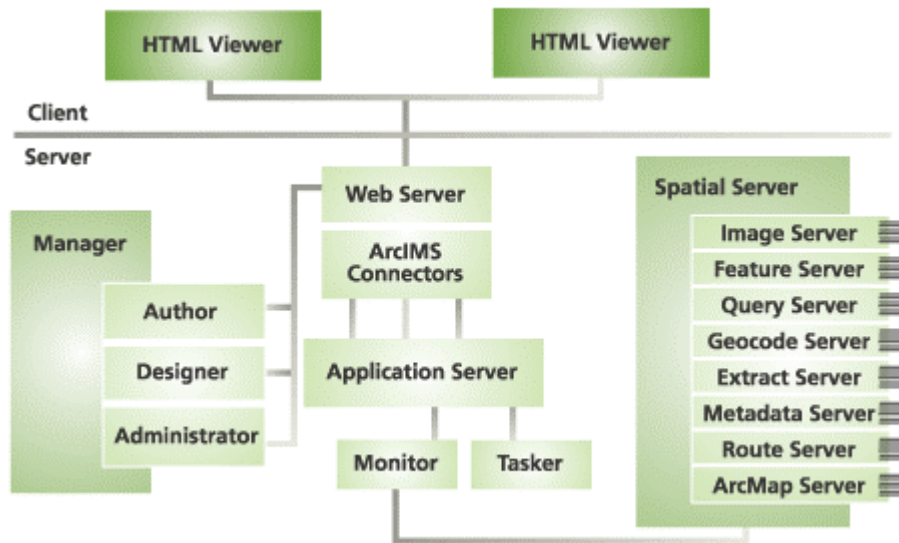
Image Server – A streamlined set of mapping components that provide effective mapping functionality in the form of an output map image that is downloaded to the end client device. The client application can either be a Web browser or a GIS application such as ArcMap, ArcView, or ArcExplorer.

ArcMap Server – Based on the ArcGIS COM architecture, ArcMap Server takes advantage of all the advanced data access and cartographic capabilities included with the ArcGIS authoring environment. As with Image Server, an output map image is created and downloaded to the end client device. The client application can either be a Web browser or a GIS application such as ArcMap, ArcView, or ArcExplorer.

Feature Server – Provides compressed feature streaming to the end client device for display and analysis. The end client must process the features directly, and therefore must use a GIS application such as ArcMap, ArcView, or ArcExplorer. Data are cached at the client to reduce network data transfer.

These three Virtual Servers are part of the Spatial Server container and execute on the defined map server as defined by the ArcIMS architecture (see Figure 1).

Figure 1
ArcIMS Architecture



Since map generation using Image Server is generally not as complex as compared to ArcMap Server, the map service time with Image Server is typically less resulting in maximum peak transaction rates. Testing has shown that peak transaction rates for an equivalent service published under both types of virtual servers are roughly 10% higher for Image Server. Image Server should be used for mapping applications unless specific

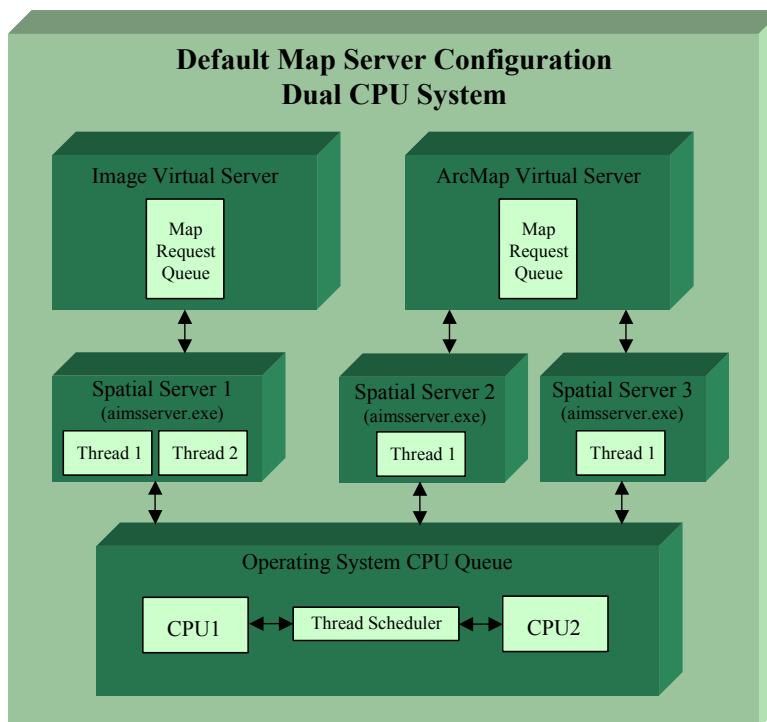
application functional needs require the use of the alternate ArcMap Server virtual server. Feature Server is for situations where raw features are required to be sent to the end client application. Since features are downloaded as opposed to an image of the features, network traffic with Feature Server will be increased substantially when compared to an equivalent image download from Image Server or ArcMap Server.

Recommendation: Use Image Server for basic maps and ArcMap Server for complex maps requiring additional cartographic elements, etc. Use Feature Server only when a requirement exists to download raw features to the client application. Feature Server should only be used on networks where bandwidth is not an issue and when there is a specific need to stream data to support client analysis, etc.

Spatial Servers and Threads

ArcIMS provides several configuration options for spatial servers (defined as a aimsserver.exe process) and their associated instances, or threads. For the Image Server virtual server, each spatial server has 1-n number of threads associated with it for parallel processing. For ArcMap Server, only one thread per spatial server is allowed and parallel processing is achieved by the use of multiple spatial servers. By default, the Image Server virtual server is configured with one spatial server with two threads while the ArcMap Server virtual server is configured with two spatial servers with one thread each. Figure 2 depicts a default map server configuration for Image and ArcMap virtual servers on a dual-CPU system.

Figure 2
Default Map Server Configuration

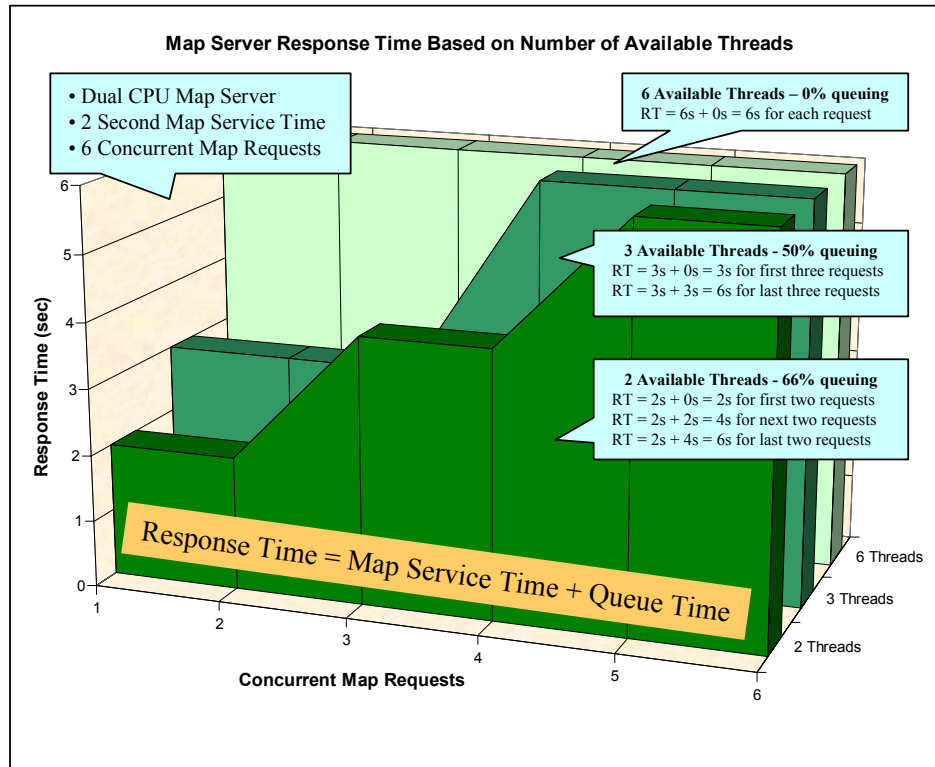


With this example, a total of four threads from three different spatial servers are available for the two CPUs. However, depending upon the application, more threads may be required to take advantage of all the CPU resources. Also, as the number of CPUs increases on a given system, additional threads will need to be configured to take full advantage of the CPU resources. For example, if a map server is configured with a system having four CPUs, it may be necessary to add four additional threads or more to take full advantage of the CPU resources. Keep in mind that map services would have to be published using both virtual server types for this configuration to be effective. That is, at least one service that used the Image Virtual Server and another that uses the ArcMap Server Virtual Server. Testing and tuning should be performed for each implementation to determine the appropriate number of threads required.

Not only does the total number of threads running impact performance but also the manner in which they are configured. Testing on Windows platforms has shown that using more than four threads per spatial server results in reduced performance. For example, using twelve threads on a single spatial server on a four-CPU system under full load resulted in only 80% CPU utilization, but using three spatial servers with four threads each resulted in 97% CPU utilization. Assuming the map server is running efficiently, higher CPU utilization generally results in higher transaction rates and therefore efficient use of hardware resources. The trade off with creating and using additional spatial servers as opposed to just adding more threads is additional memory consumption.

The performance goal with configuring and tuning ArcIMS threads is to maximize the transaction rate capability of a given map server. This is accomplished by fully utilizing the map server CPU resources (90+% utilization) while minimizing the number of threads so that map server response times do not grow to unacceptable levels. For example, if you configure six Image Server threads on a dual-CPU system and that system receives six concurrent requests, queue time for each request will be zero but the map service time for each request will be roughly three times longer than if only two threads were configured. This would result in slower response times for all the requests instead for just those requests that are forced to wait in the queue under a thread-limited configuration. For a two-thread configuration, instead of simultaneously processing all six requests, the CPUs would process the first two requests while the other four waited in the queue, resulting in a response time for the first two requests that is equal to the map service time (response time = service time + queue time). For the second set of requests, response time would be equal to the map service time plus the queue time while waiting for the first two requests to process, and so on. Figure 3 illustrates this discussion and shows how the number of configured threads affects response time performance out of the map server. It assumes a single request consumes all resources from one CPU for the duration of the map service time. The bottom line is that it is best to configure adequate threads to take advantage of the map server hardware, which maximizes transaction rates, while leaving additional requests in the virtual server queue for future processing once a thread becomes available.

Figure 3
Map Server Thread Configuration Performance

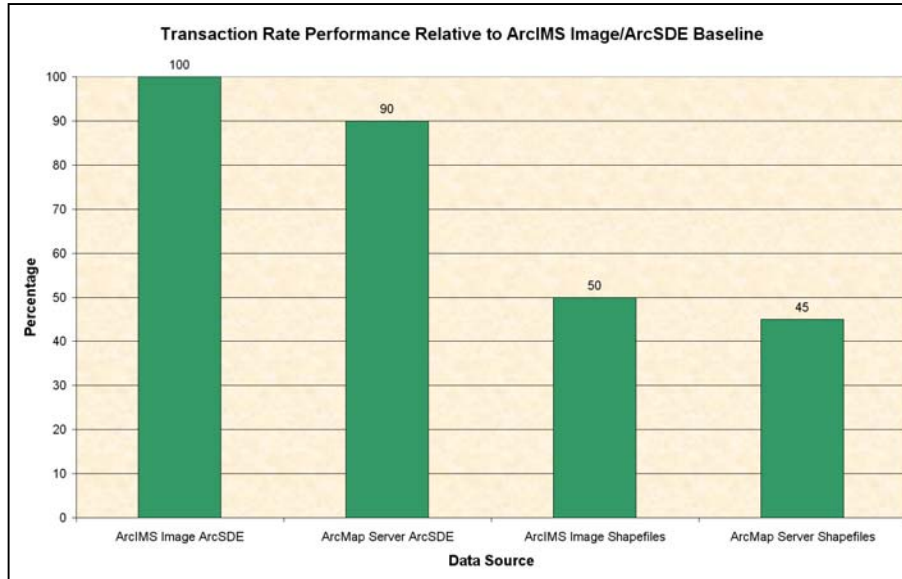


Recommendation: Determine the optimum number of threads by testing the application to see how many threads are needed to consume the CPU capacity of the map server(s) while maintaining reasonable response times. Remember, the performance goal is to maximize transaction rates by maximizing the use of the hardware without elongating map server response time too much by having too many threads configured. Using two and possibly three threads per CPU will typically provide the load required to effectively utilize the map server resources. Also, avoid configuring more than four threads per Image Server spatial server instance since performance will likely decrease. Instead, create additional spatial servers with fewer threads each.

Data Source Type

The data source chosen to support published ArcIMS services, in conjunction with the selected virtual server type, will affect transaction rate performance. Several variables must be considered when comparing performance such as the number of layers, the size of the shapefiles, the number of features retrieved, use of rasters, etc. Figure 4 shows the relative transaction rate performance based on different data sources, in terms of percentages, using an ArcIMS Image Service (9 vector layers plus 1 raster layer) with an ArcSDE data source as the baseline.

Figure 4
Transaction Rate Performance Based on Data Source



Recommendation: If the site is medium to large in size and supporting many transactions (several thousand per hour and higher), then the data should be accessed from ArcSDE to ensure optimum performance. Even for small sites, if the shapefiles accessed are large (100 MB+), ArcSDE is a better option to ensure acceptable map service response times.

Output Image Type

ArcIMS supports three output image types: JPEG, GIF, and PNG (8 and 24 bit). ArcMap Server also supports .pdf, .ai, .eps, and .emf output formats with a GET_LAYOUT ArcXML request. Following is a brief description of the three main output image types and when they are appropriate for use.

JPEG (Joint Photographic Experts Group)

- Uses “lossy” compression.
- Supports 16.7 million colors.
- Best suited for photographs and complex graphics.
- A compromise must be made between output image size and the quality of the image, i.e., higher quality settings result in less compression, which results in larger output image file sizes.
- No transparency support.

GIF (Graphics Interchange Format)

- Uses “loss-less” compression.
- Supports 256 colors.
- Best suited for simple images with few distinct colors.
- Requires GIF license to be purchased.
- Image sizes can be smaller or larger than comparable JPEG depending upon the number of distinct colors used, etc.
- Supports transparency.

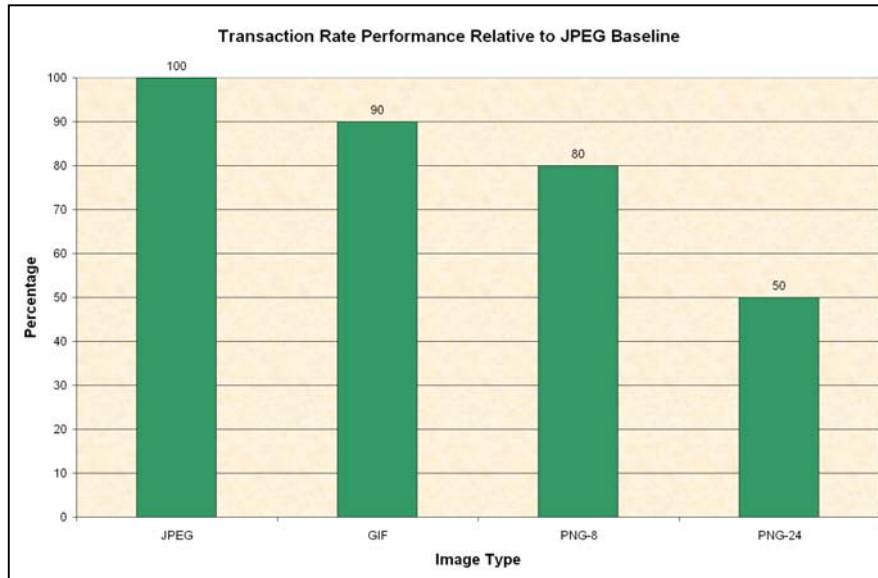
PNG (Portable Network Graphics)

- Uses “loss-less” compression.
- PNG-8 (8 bit) supports 256 colors.
- PNG-24 (24 bit) supports 16.7 million colors.
- Created as a free license alternative to GIF.
- PNG-8 image sizes are typically 5-25% smaller than equivalent GIF.
- PNG-24 image sizes can be several hundred kilobytes in size and are often not suitable for the Internet, but will work fine on an Intranet.
- Supports transparency.

For Internet deployments, one of the performance design goals should be to provide functional maps while minimizing the output image size. Maps destined for the Internet should not exceed 100 KB with a design goal range of 25-75 KB. If the output images are too large, download time for users with minimal bandwidth will be too large resulting in loss of interest by the end user. Also, large output images can result in an Internet bandwidth bottleneck that could effectively reduce peak transaction rates, even though the map servers have remaining capacity. Section 4.0 of the *System Design Strategies* white paper (<http://www.esri.com/library/whitepapers/pdfs/sysdesig.pdf>) covers this topic in detail.

The selected output image type also affects ArcIMS transaction rate performance. This is because certain image types require more CPU resources than others, resulting in longer map service times for more complex image types. For example, a PNG-24 requires more resources than an equivalent GIF image because of the number of colors supported, etc. Also, images with more colors naturally result in larger output image file sizes. An appropriate output file type should be selected based on application needs, bandwidth availability, and scalability requirements. Figure 5 shows the relative transaction rate performance of the different image types, in terms of percentages, using an ArcIMS Image Service (9 vector layers plus 1 raster layer) with JPEG as the baseline.

Figure 5
Transaction Rate Performance Based on Image Type



Recommendation: Use JPEG for the output image type when possible for optimal performance and output image size. If transparency is required, use GIF or PNG. PNG-24 can be used for high-quality output images, though using it results in large output image sizes and is, therefore, most appropriate for Intranet deployments.

Output Image Resolution

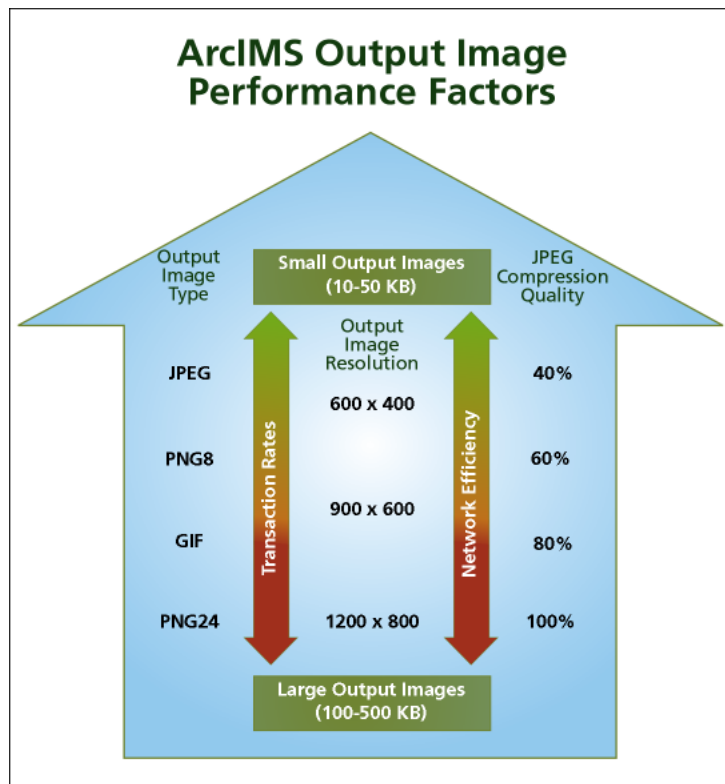
Map image resolution should always be considered, and limited in many cases, when designing ArcIMS image service-based applications. Chosen image resolution, which is based on the actual x and y image dimensions, affects transaction rate performance and output image file size. The larger the resolution, the longer the processing time required to generate the output image which translates into lower transaction rates. Also, larger resolution images result in larger output image files, which impacts client download time and network bandwidth utilization. For example, testing has shown that an image service using 600 x 400 resolution results in twice the transaction rates and one-half to one-third the output image size as compared to the same image service using 1200 x 800 resolution.

Recommendation: Use smaller resolution images when possible. An output image size of 600 x 400 is a good starting point for a typical map display. Web programming techniques can be used to force the output image size to remain at a pre-determined size regardless of the client request.

JPEG Compression Quality

JPEG images are created using compression techniques that are represented by the use of a quality value. The lower the quality value, the higher the amount of compression. This results in a smaller output image file size with the trade-off being reduced image quality. The default quality value for both Image Server and ArcMap Server is 85, which provides very good quality output images. For Internet applications, an appropriate quality value should be selected based on the acceptable image quality. This ensures that the output image file size is kept as low as possible, which conserves bandwidth, and at the same time maximizes transaction rates for bandwidth-limited configurations. Also, setting the quality value above 85 can cause the output image file size to grow dramatically with little or no noticeable improvement in image quality. The JPEG image quality value can be changed for the Image Server by modifying `<PARAMETER name="jpegquality" value="85" />` in the aimsms.cfg file. This is a global setting that will affect all Image Server services that use the JPEG format. Further, ArcIMS must be restarted after making the change. The JPEG quality value for ArcMap Server is fixed at 85 and cannot be changed. Figure 6 summarizes the last three output image topics discussed in terms of their impact on ArcIMS performance. Choosing options nearer to the top of Figure 6 will provide improved transaction rate and network performance. This chart is based on results from a specific ArcIMS Image Service application using 9 vector layers and 1 raster layer. Results may vary depending upon the application.

Figure 6
Output Image Performance Factors



Recommendation: To minimize the output image size for Internet or bandwidth-limited networks, increase the amount of JPEG compression by lowering the JPEG quality value while maintaining the desired image quality.

Using ArcMap as a Map Service Client

ArcIMS supports several clients including Web browsers and GIS applications such as ArcMap, ArcView, and ArcExplorer. Typically, the output image type is pre-determined by the service and the client receives that image type following a request. However, the configured output image type can be ignored and replaced with an image type requested by the client. When ArcMap is used as a client, it does just this and always requests a PNG-24 output image type. If the service is intended for an Intranet, use of PNG-24 may be appropriate. However, if the service is to be accessed by ArcMap over the Internet or on a bandwidth-limited network, the PNG-24 image type should not be used, unless the application functionality absolutely demands it. The configuration setting `spatialServer.AllowOutputTypeChange` in the `Esrimap_prop` file can be used to override the image type request by the client. Setting this value to false will cause ArcIMS to ignore the request for a different output image type by the application. This will result in the use of the output image type defined by the map service.

Recommendation: If ArcMap is to be used as an Image Server or ArcMap Server client, set the `spatialServer.AllowOutputTypeChange` parameter in the `Esrimap_prop` file to false for Internet or bandwidth-limited networks. If high quality images are needed for internal users that have plenty of available bandwidth, a second identical service can be created that uses the PNG-24 output image type.

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>.