IBM WebSphere Performance Pack: Web Content Management with IBM AFS Enterprise File System

Marco Pistoia, Tom Menner, Catherine Milligan, Bobby Gia Pham

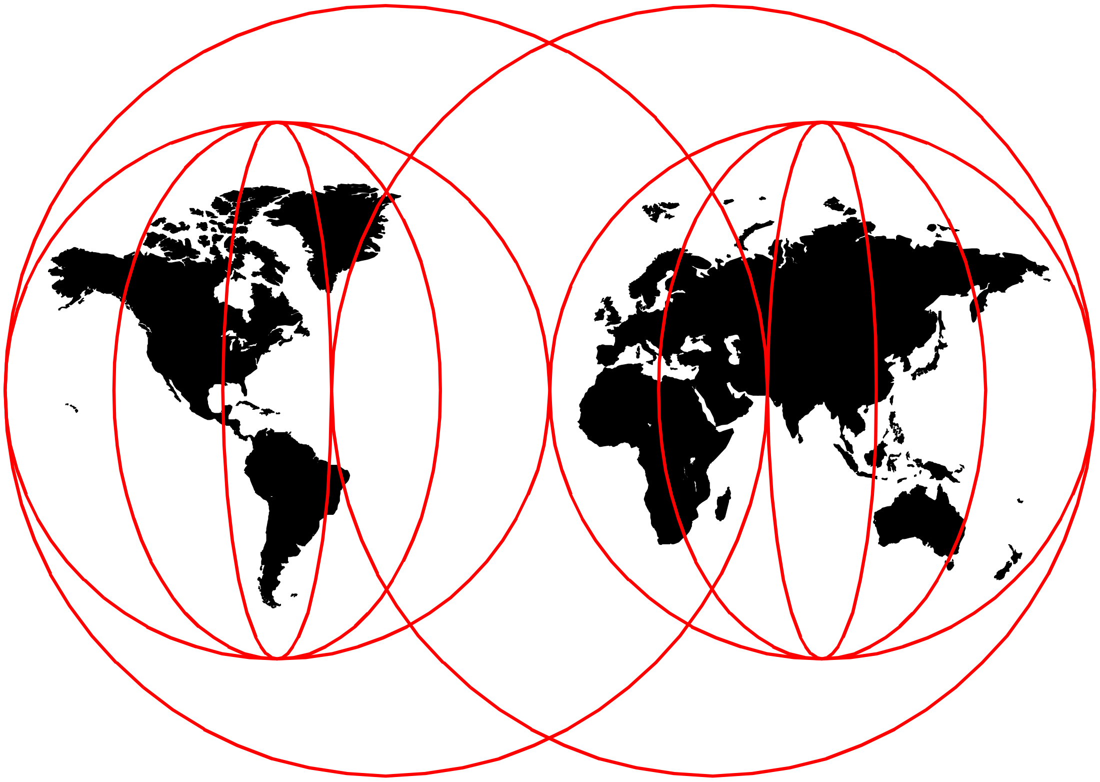

International Technical Support Organization

http://www.redbooks.ibm.com

SG24-5857-00

International Technical Support Organization

SG24-5857-00

IBM WebSphere Performance Pack: Web Content Management with IBM AFS Enterprise File System

October 1999

Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix A, “Special Notices” on page 427.

First Edition (October 1999)

This edition applies to IBM AFS Enterprise File System, the File Sharing component of WebSphere Performance Pack Version 2, for use with the AIX, Solaris, and Windows NT operating systems.

Comments may be addressed to:
IBM Corporation, International Technical Support Organization
Dept. HZ8 Building 678
P.O. Box 12195
Research Triangle Park, NC 27709-2195

When you send information to IBM, you grant IBM a non-exclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

Contents

Preface

IBM WebSphere Performance Pack is Web infrastructure software that addresses the scalability, reliability, and performance needs of e-business applications in both local and geographically distributed environments. Its functions incorporate leading-edge and robust caching, file management, and load balancing, which together compensate for the inherent weakness of the Internet to support critical business applications and expectations.

This redbook will give you a clear understanding of the features of IBM AFS Enterprise File System, the File Sharing component of IBM WebSphere Performance Pack. It shows how to plan for, install, configure, use, tune and troubleshoot this component and offers specific implementation examples. Moreover, it helps explain how to build complex scenarios that involve all the components of IBM WebSphere Performance Pack, to give you a better understanding of the technologies involved.

Note that this publication was written in conjunction with two other volumes about WebSphere: *IBM WebSphere Performance Pack: Load Balancing with IBM SecureWay Network Dispatcher*, SG24-5858, and *IBM WebSphere Performance Pack: Caching and Filtering with IBM Web Traffic Express*, SG24-5859. To realize the most benefit, all three volumes should be obtained.

The Team That Wrote This Redbook

This redbook was produced by a team of specialists from around the world working at the International Technical Support Organization, Raleigh Center.

The leader of this project was Marco Pistoia.

Marco Pistoia is an Advisory Software Engineer, working as a project leader at the International Technical Support Organization Raleigh Center. He writes extensively and teaches IBM classes worldwide on all areas of the e-business Application Framework, WebSphere, Java and Internet security. Marco holds a Master of Science degree with honors in Pure Mathematics from the University of Rome. Before joining the ITSO, he was a System Engineer in IBM Italy. He received an Outstanding Technical Achievement Award in 1996.

Tom Menner is a Technical Consultant, based in New York City, for the IBM Transarc Laboratory in Pittsburgh. He has worked for the IBM Transarc Laboratory since 1994, first as an AFS, UNIX and NT systems administrator and then as a DCE/DFS/AFS consultant and trainer. His areas of expertise include DCE, DFS and AFS across many different platforms.

Catherine Milligan is a Senior I/T Specialist at IBM Transarc Lab in Pittsburgh where she leads course development efforts for File System products, namely AFS and DFS, and teaches courses in Systems Administration and User Training. Her most recent project was developing the IBM WebSphere Performance Pack Server Administration course with hands-on labs. Catherine holds a B.S. degree in Applied Mathematics from Carnegie Mellon University and a Masters degree from the University of Pittsburgh School of Education. Prior to joining Transarc, Catherine worked as a Research Programmer, User Consultant, and Education and Training Coordinator at the Pittsburgh Supercomputing Center.

Bobby Gia Pham is a Technical Consultant in Raleigh, North Carolina. He has over 15 years of experience in large scale distributed systems. He has worked at IBM for 2 years. His areas of expertise include AFS, and DCE/DFS.

Thanks to the following people for their invaluable contributions to this project:

Poh Yee Tiong
IBM Singapore

Corinne Letilley
IBM Canada

Mary Ann DelBusso, Felicia Ferlin, Tony Mauro, Melissa Lyle
IBM Transarc Laboratory, Pittsburgh

Rick Schenck, Jerry Gschwind, Susan Hanis
IBM Research Triangle Park, North Carolina

Shawn Walsh, Jorge Ferrari, Tim Kearby, Margaret Ticknor, Pat Donleycott, Tate Renner, Linda Robinson, Gail Christensen
International Technical Support Organization, Raleigh Center

Vincenzo Iovine, Stefano Pischedda
IBM SEMEA Sud, Italy

Karen Gelveles
IBM Boca Raton

Comments Welcome

Your comments are important to us!

We want our redbooks to be as helpful as possible. Please send us your comments about this or other redbooks in one of the following ways:

- Fax the evaluation form found in “ITSO Redbook Evaluation” on page 439 to the fax number shown on the form.
- Use the online evaluation form found at `http://www.redbooks.ibm.com/`
- Send your comments in an Internet note to `redbook@us.ibm.com`

Part 1. AFS Enterprise File System Management and Administration

Chapter 1. IBM WebSphere Performance Pack Concepts

IBM WebSphere Performance Pack is Web infrastructure software that addresses the scalability, reliability and performance needs of e-business applications in both local and geographically distributed environments. Its functions incorporate leading-edge and robust caching, file management and load balancing, that together compensate for the inherent weakness of the Internet to support critical business applications and expectations.

IBM WebSphere Performance Pack has been developed using IBM's extensive experience with very demanding Web sites.

IBM WebSphere Performance Pack is composed of three main components, which permit you to reduce Web server congestion, increase content availability and improve Web server performance:

1. **File sharing**

   The file-sharing component, known as IBM AFS Enterprise File System (AFS), is an enterprise file system that enables cooperating hosts (clients and servers) to efficiently share file system resources across both local area networks (LANs) and wide area networks (WANs). It provides non-disruptive real-time replication of information across multiple servers, which guarantees data consistency, availability, global stability and administrative efficiency, necessary by large distributed Web sites or by Web sites with volatile content requiring considerable administrative effort to maintain content links and URLs to file I/O mapping.

2. **Caching and filtering**

   The caching and filtering component, known as IBM Web Traffic Express (WTE), is a caching proxy server that provides highly scalable caching and filtering functions associated with receiving requests and serving URLs. With tunable caching capable of supporting high cache hit rates, this component can reduce bandwidth costs and provide more consistent rapid customer response times.

3. **Load balancing**

   The load balancing component, known as IBM SecureWay Network Dispatcher (ND), is a server that is able to dynamically monitor and balance TCP servers and applications in real time. The main advantage of the load balancing component is that it allows heavily accessed Web sites to increase capacity, since multiple TCP servers can be dynamically linked in a single entity that appears in the network as a single logical server.

WebSphere Performance Pack V2 offers a fourth component named Common Configuration. This new feature allows for centralized configuration of the main components and can be run remotely on a separate machine connected to the network.

The WebSphere Performance Pack components, which were previously unavailable in a single Internet software offering, can increase the scalability, availability and reliability of your Web site while reducing infrastructure costs.

Installation procedures permit selection of which components to install, and specification of on which machine(s) the selected component(s) should be located. Subject to installation needs and operating platforms, components can coexist on a single machine or can be distributed over multiple machines.

IBM WebSphere Performance Pack Version 2 is supported on the following platforms:

- Any IBM RISC/6000-based machine running IBM AIX 4.2.1 or later and Java Runtime Environment (JRE)[1] 1.1.6 or later
- Any SPARC workstation running Sun Solaris 2.6 or later and JRE 1.1.6 or later – JRE 1.1.7B or later is required to use ND
- Any Intel x86 PC running Microsoft Windows NT[2] 4.0 and JRE 1.1.6 or later – JRE 1.1.7B is required to use ND

If you are planning to use the Common Configuration, then you need the Java Development Kit (JDK)[3] 1.1.6 or higher and a Java 1.1-enabled Web browser, such as:

- Netscape Communicator V4.08 or higher
- Microsoft Internet Explorer V4.01 with the fix pack, or higher
- Sun HotJava V1.1 or higher

1.1 AFS Enterprise File System in Distributed Computing Environments

AFS has provided scalable file administration and file sharing for large enterprises for many years, based upon its use of a virtual name space to make naming and logical directory structures of files independent of their physical location. AFS clients and AFS servers are used to establish this virtual name space capability. In typical local area network (LAN) file systems, this is achieved by installing AFS clients on user workstations, communicating with an AFS server that manages the I/O operations associated with the actual files. In a Web site, the AFS clients can be installed on HTTP servers to reduce the administrative effort associated with maintaining URL to file I/O mapping relationships. In addition, HTTP servers that are simultaneously AFS clients can significantly increase the connectivity capacity to Web server content and can provide local and geographically distributed access efficiency.

AFS is a central and scalable file system:

- It is *central* because AFS brings together all of the files within the file system into a single name space. Every AFS user shares this same name space, making all AFS files easily available from any AFS machines. With AFS, the name of a file is independent of both the file's and the user's physical location, contributing to ease of file sharing and resource management.
- It is *scalable* because AFS is able to manage a very large number of files, spread across many geographical locations. When remote files, residing on AFS servers, are accessed by remote AFS clients, they are cached on the client machines to improve performance. This makes remote working across global distances feasible, since it is possible to access your own files from sites many thousands of miles away as if they were local.

[1] You can download the JRE from `http://www.javasoft.com`.
[2] IBM WebSphere Performance Pack can be installed on either Windows NT Server or Workstation V4.0. However, Windows NT Server is required for PCs running the ND functions Interactive Session Support (ISS) Nameserver and Observers.
[3] You can download the JDK from `http://www.javasoft.com`.

Both small and large-scale distributed environments benefit from AFS mechanisms to reduce server and network load:

- AFS caches data on client machines to reduce subsequent data requests directed at file servers, substantially reducing network and server loads. Servers keep track of data given to clients through *callbacks*, guaranteeing cache consistency without constant queries to the server to see if the file has changed. It is important to underline that AFS also allows disk cache, not just memory cache. This is a key advantage of AFS over other shared file systems.
- The AFS remote procedure call (RPC) reads and writes data to an RPC stream, further improving the efficiency of data transfer across a local or wide area network.

Extended security is guaranteed through Kerberos authentication and access control lists (ACLs). AFS Kerberos-based authentication requires that users prove their identities before accessing network services. Once authenticated, AFS access control lists give individual users or groups of users varying levels of permission to perform operations on the files in a directory.

The AFS component also offers replication techniques for file system reliability. Multiple copies of frequently accessed (but infrequently changed) data are replicated on multiple file servers within a cell. When accessing this information, a client will choose among the available servers that house replicas. If one server is unavailable or unreachable, the client will go to another server. Replication also reduces the load on any particular server by placing frequently accessed information on multiple servers.

Moreover, management utilities are provided to ease the load of system administrators in growing environments. Backup, reconfiguration and routine maintenance are all done without any system down time. Files remain available to users during these operations. This is done by creating online clones of volumes.

AFS commands are RPC-based. Administrative commands can be issued by any authenticated administrator from any client workstation. System databases track data location information, authentication information and protection groups. These databases are replicated on multiple servers, and are dynamically updated as information changes. Server processes accomplish many tasks automatically, such as restarting servers, tracking file locations and updating file servers with new binaries and configuration files.

1.2 Caching and Filtering to Manage Internet Traffic and Bandwidth Demand

WTE, the caching and filtering component of IBM WebSphere Performance Pack, is both a caching proxy server and a content filter. The advanced caching of this component minimizes network bandwidth and ensures that end users spend less time when retrieving the same content multiple times.

This component acts as a gateway for multiple clients and performs basic Web server duties, such as receiving requests and serving URLs.

A traditional proxy server receives a request for a URL from a client and it forwards the request to the destination content server. WTE does something more; it can save or cache the Web documents it retrieves, and serve subsequent requests for those documents from its local cache. The client gets the requested information faster and network bandwidth is reduced.

This component of IBM WebSphere Performance Pack also offers other key features of advanced caching, such as:

- The ability to handle very large caches
- An option to automatically refresh the cache with the most frequently accessed pages
- The possibility to cache even those pages where the header information says to fetch them every time
- Configurable daily garbage collection, to improve server performance and ensure cache maintenance
- Remote Cache Access (RCA), a new function – available for the first time with IBM WebSphere Performance Pack V2 – that allows multiple WTE machines to share the same cache by using a distributed file system, such as AFS, therefore reducing redundancy of cached content

Moreover, WTE allows you to set content filtering at the proxy server level, rather than or in addition to the browser level, where content filtering could be easily compromised or over-ridden. This way, offensive contents will not be displayed on the client’s browser, depending on the parameters used in the configuration. Content filtering in WTE can use:

- Platform for Internet Content Selection (PICS) rules guiding use of rating labels - such as Recreational Software Advisory Council on the Internet (RSACi) criteria for inappropriate language, nudity or violence - placed in HTML or HTTP headers or third-party content rating label distributions
- Lists of URLs/sites for which access is to be blocked
- APIs for filtering applications

1.3 Load Balancing and Server Monitoring Capabilities

The Internet has grown so rapidly over the last few years, that you are probably looking for a way to handle your company's share of that traffic. If this growth is not properly handled, users get slow response or refused connections, creating an unsatisfactory user experience which may cause the user never to visit your site again. Internet sites can become unstable or even fail under critical load conditions. What is needed is a solution that balances the load effectively and protects the user from these bad experiences.

ND, the load balancing component of IBM WebSphere Performance Pack, has been developed to address these limitations and provide customers with advanced functions to meet their site’s scalability and availability needs. It consists of three functions: the Dispatcher, Interactive Session Support (ISS) and the new Content Based Routing (CBR) function, available for the first time with IBM WebSphere Performance Pack V2. These three functions can be deployed separately or together in various configurations to suit a wide variety of customer application requirements:

- You can use the Dispatcher function to balance the load on the server within a local area network or wide area network using a number of weights and measurements that are set dynamically.
- ISS is a DNS-based load monitoring component (daemon) that can be installed on each of your servers. This group of daemons is called an *ISS cell*. One of the members of the cell becomes a spokesman for the load monitoring service.

  You can use the ISS function to balance the load on servers within a local area network or wide area network using a domain name server round-robin approach or a more advanced user-specified approach. ISS periodically monitors the level of activity on a group of servers and detects which server is the least heavily loaded.

  ISS provides an *observer* interface to enable other applications to use the load monitoring service. Observers watch the cell and initiate actions based on the load. Application servers with the ISS load monitor daemon installed can pass periodic load reports to the Dispatcher using the Dispatcher observer. The results of these reports can be factored into the load-balancing performed by Dispatcher.
- You can use the CBR function (which must be installed and configured to work together with WTE) to load balance traffic based on the content of a client's URL request. CBR also offers a Cookie Affinity feature that allows requests from a particular HTTP client session to be load balanced to the same server for a specified time period. The client session will maintain affinity for a particular server without relying on the IP address of the client.

These three components offer a high availability feature:

- The Dispatcher high availability feature involves the use of a secondary machine that monitors the main, or primary, machine and stands by to take over the task of load balancing, should the primary machine fail at any time. This feature is available on all the platforms where IBM WebSphere Performance Pack is supported, without using High Availability Cluster Multi-Processing (HACMP).
- In ISS high availability, all the nodes in a site work together to eliminate any single point of failure.
- Multiple CBR machines can be load balanced in turn using ND, therefore granting CBR high availability.

The high availability feature provided by the Dispatcher function can be successfully used even in other configurations, for example to guarantee firewall high availability.

1.4 Building Record-Breaking Web Sites

IBM WebSphere Performance Pack allows you to design several architectures to enhance the performance of your Web site. Figure 1 on page 8 offers an idea of the multiple configurations that can be obtained combining the WebSphere Performance Pack components:

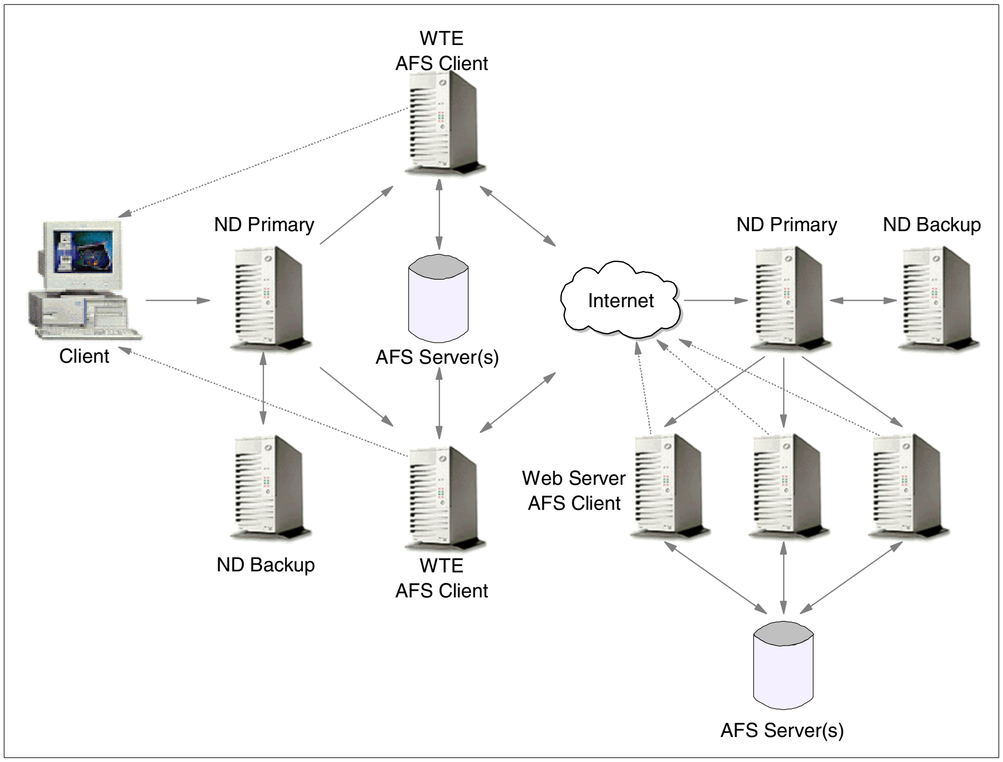

Figure 1. How to Implement a WebSphere Performance Pack Environment

- The WTE component minimizes response time and network bandwidth utilization by providing Web content caching. It also ensures reliable content filtering at the proxy server level. Multiple WTE servers can be load balanced.
- The ND component distributes the load between multiple clustered Web servers and WTE servers. As soon as a client request arrives, the load balancing machine uses sophisticated monitoring tools and then forwards the request to the least loaded server. High availability is provided by a backup Dispatcher machine, which monitors the state of the primary machine and takes over the primary machine if it fails.
- The clustered Web servers can share the same content by using AFS. This ensures scalability, high availability, reliable access to replicated data and an efficient security model for access and group management.

  AFS can also be integrated with RCA, so that clustered WTE servers can share the same Web cache therefore eliminating the need to fetch the same Web pages multiple times.
- The server selected by the ND machine can then respond directly to the client without any further involvement of the ND machine. Since there is no need for the server response to go back through the same physical path, a separate high-bandwidth connection can be used.

IBM used the WebSphere Performance Pack technology to create a scalable and reliable system that efficiently handled unprecedented traffic volumes. On February 17th, 1998, at 12:41 (Japan Standard Time), the official Web site of the Olympic Winter Games in Nagano made Internet history by logging a staggering 98,226 hits per minute. Less than a week later, a new all-time record was established with a peak load of more than 103,400 hits per minute, while still providing normal response time. The Internet site of the 1998 Nagano Olympic Winter Games is recognized by the Guinness Book of World Records.

In addition to the Winter Olympics site, IBM has built some of the largest Web sites in the world. For example, IBM hosted the Deep Blue chess match, the 1996 Olympic Games, the U.S. Open tennis tournament, Wimbledon, the Masters golf tournament, and the official Web site of the 1998 French Open tennis championship.

By using WebSphere Performance Pack on your Web site, you will have an efficient Web site, capable of providing fast responses. Your Web site will be able to handle very large amounts of simultaneous requests, without any major delay. Furthermore, the high availability features built into WebSphere Performance Pack will make the Web services available even when one or more of your server machines should unexpectedly fail.

The following figure shows an example of what your Web customers would not see if your Web site uses WebSphere Performance Pack:

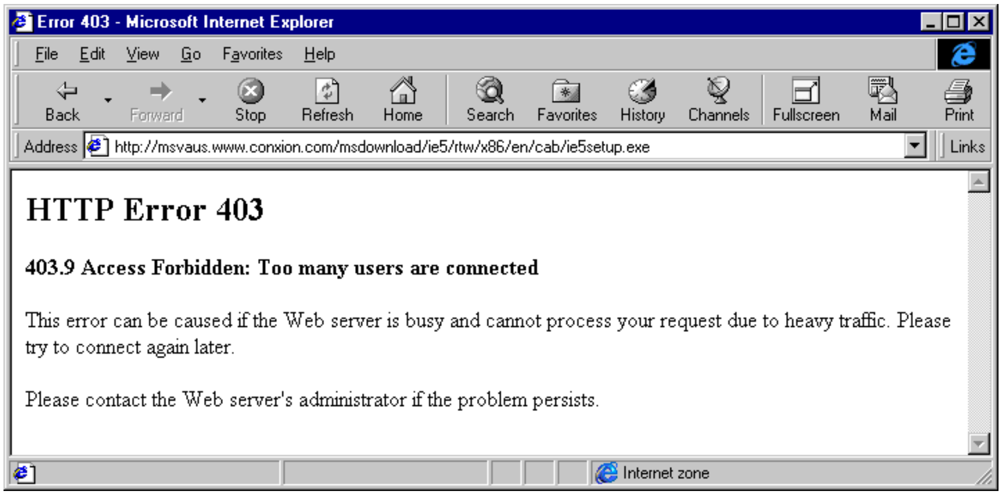
Error 403 - Microsoft Internet Explorer
File Edit View Go Favorites Help
Back
Forward
Stop
Refresh
Home
Search
Favorites
History
Channels
Fullscreen
Mail
Print
Address
http://msvaus.www.conxion.com/msdownload/ie5/rtw/x86/en/cab/ie5setup.exe
Links
Internet zone

HTTP Error 403

403.9 Access Forbidden: Too many users are connected

This error can be caused if the Web server is busy and cannot process your request due to heavy traffic. Please try to connect again later.

Please contact the Web server's administrator if the problem persists.

Figure 2. Don't Let This Happen to You!

1.5 What Is New in Version 2

In addition to the functions already available in Version 1.0, described in the IBM redbook *IBM WebSphere Performance Pack Usage and Administration*, SG24-5233, IBM WebSphere Performance Pack enhancements in Version 2 include:

- Quality of service enhancements
- New ND functions
- New WTE functions
- New AFS functions
- New functions available by combining components
- Integrated configuration assistance

1.5.1 Quality of Service Enhancements

WebSphere Performance Pack now provides support for differentiated quality of service, as discussed in the following sections.

1.5.1.1 User Classes

User classes give you the ability to change the level of service depending on the identity of the user. Preferred customers would expect to get better service. In Version 1.0, ND provided this capability by allowing rules-based load-balancing based on the client IP address, and WTE allowed PICS filtering based on the client IP address or the user and group. In Version 2, ND and WTE add the ability for rules-based load-balancing based on HTTP headers including `Cookie`, `Referer`, and `User-Agent`.

Differentiated service based on user identity lets you determine what level of service should be provided. For example, a frequent buyer can be routed to a higher capacity server or given access to additional content (for example, sale information) not made available to an unknown customer. The information used to identify the user comes from cookies or header information in the HTTP request.

1.5.1.2 Service Classes

Server classes give you the ability to change the level of service depending on the information or application requested by the client. Services involving a purchase would get better service than requests for information. In Version 1.0, WTE provided this capability through its reverse proxy function by allowing requests to be redirected to different servers based on the contents of the URL. In Version 2, ND and WTE add the ability for rules-based load-balancing based on the protocol, host, and path portion of the URL.

Differentiated service based on the service requested lets you give preferential service to some requests. For example, you may want to give preferential service to a *proceed to checkout* request over a *search* request.

1.5.2 New ND Functions

ND has been enhanced to include several new features. These are discussed in the following sections.

1.5.2.1 Server Directed Affinity API

In previous releases, administrators could define a particular port to be configured as *sticky*, meaning all requests to that port remain with a particular physical server for a short period of time. When a client connects to a sticky port, an entry for that client’s IP address is made in the affinity table, and a time stamp is set. If a new connection from the same client arrives, and the time stamp has not expired, then the new connection is sent to the same server as before. This function is still available in Version 2, but it is no longer the only option.

Customers can now write their own software using the *Server Directed Affinity* (SDA) Application Programming Interface (API) to implement an SDA agent, which communicates with a listener in the Dispatcher. This software can then manipulate the Dispatcher affinity tables to:

- Query the contents
- Insert new records

- Remove records

1.5.2.2 Binary Logging and Statistics

Dispatcher now provides a log in binary format. The command line and the GUI have been enhanced to provide access to the binary log information. A sample Java program is provided to allow customers to access and manipulate the log.

1.5.2.3 Remote Administration

In the previous release, the Dispatcher administration GUI could only be run on the same machine where Dispatcher was installed. Now remote administration is possible, since the Dispatcher administration GUI can be run on a separate machine.

1.5.2.4 Authenticated Administration

Authentication is provided to make remote administration more secure. The communication between the ND server and an ND administration client is authenticated using a key pair.[4] These keys are generated when the ND server is started for the first time.

1.5.2.5 Wildcard Cluster

Wildcard cluster is the capability to define a cluster which will receive traffic which is not destined for this particular machine, or which is destined for this machine, but for an address which is not defined as a cluster. The traffic will be intercepted and routed to a default cluster. This feature also makes it easier to configure multiple aliases on the same Dispatcher machine to use the same port and server configuration.

1.5.2.6 Wildcard Port

Wildcard port is the capability to define default actions when no port matches a particular request. You can use this to create a load balancing configuration for traffic to any port that has not been explicitly defined in your Dispatcher configuration, for example in a firewall load balancing environment, or to discard requests for ports that have not been configured.

1.5.2.7 ISS GUI

The GUI for ISS is essentially an `isscontrol` command generator much like the Dispatcher GUI. In other words, the user makes a change in the GUI, and an isscontrol command is generated and issued. Asynchronously, the issd daemon tells the GUI that the configuration has changed and the GUI refreshes itself.

1.5.3 New WTE Functions

The enhancements to the WTE component of WebSphere Performance Pack are described in the following sections.

1.5.3.1 Transparent Proxy

Transparent proxying means that the client software is totally unaware of the existence of the intermediate proxy server. Normally, if a client browser uses a proxy server, then the browser must be configured to specify the address and port of the proxy server. This is no longer necessary with transparent proxy, in fact the client is unaware that an intermediate proxy is in the network.

[4] A *key pair* is a matching pair of public and private keys, used for digital signatures and asymmetric encryption.

To use transparent proxy, the router, which may be a Dispatcher machine, is programmed to redirect requests to the WTE transparent proxy. WTE then intercepts all HTTP requests on port 80 that are targeted at some server out in the Internet. The request is parsed and processed, and may be satisfied from the transparent proxy's cache.

Note that in WebSphere Performance Pack V2, transparent proxy is only supported on the AIX platform and works for HTTP requests only.

1.5.3.2 Proxy Autoconfiguration Support

WTE now supports *automatic proxy configuration*, a feature of Netscape Navigator 2.0 (and later) and Microsoft Internet Explorer V 4.0 (and later). This feature provides a form of transparency, in that clients do not have to configure their browser to point to a specific proxy or SOCKS server, but to an automatic configuration file, as shown in the following figure:

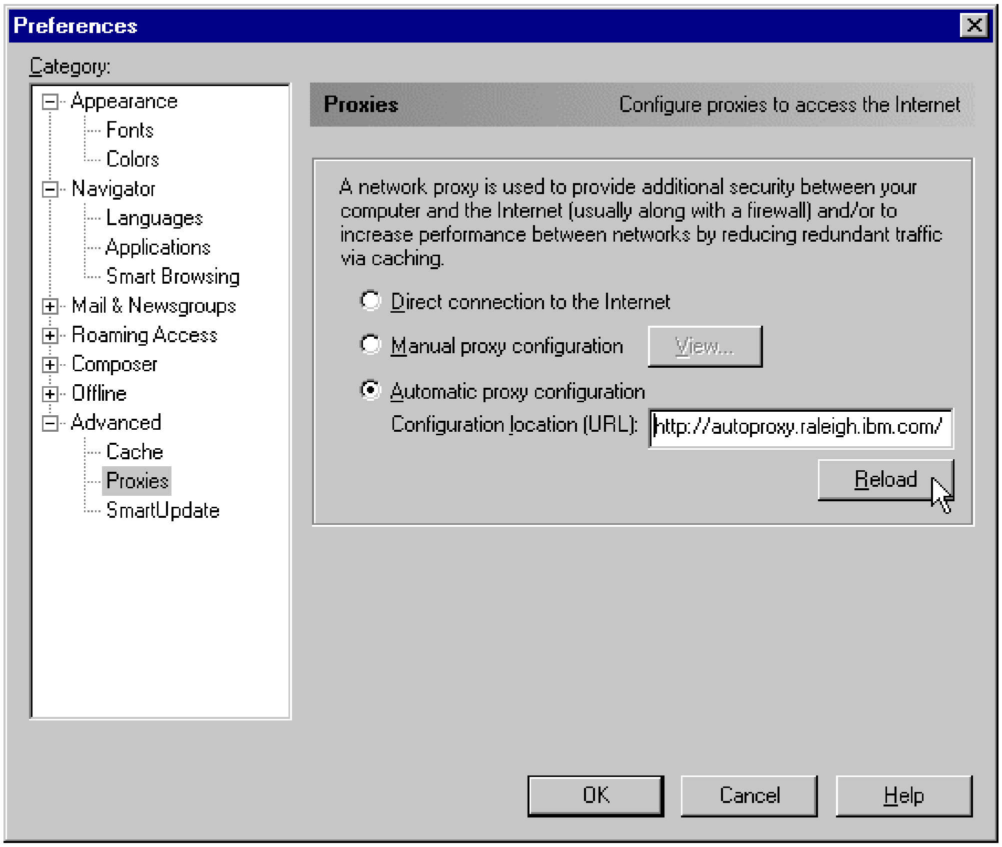

A network proxy is used to provide additional security between your computer and the Internet (usually along with a firewall) and/or to increase performance between networks by reducing redundant traffic via caching.

○ Direct connection to the Internet

○ Manual proxy configuration

● Automatic proxy configuration

Configuration location (URL): http://autoproxy.raleigh.ibm.com/

Figure 3. Automatic Proxy Configuration in Netscape Navigator

This lets the system administrator modify the configuration with little impact to the clients, who update their automatic configuration files and are automatically directed to the new configuration. Server administrators can use this to reroute requests when servers are down, to balance workload, to send requests for specific URLs to specific proxies, or other reasons specific to their installation.

1.5.3.3 FTP Proxy Enhancements

The WTE FTP proxy code now includes FTP PUT capability, improved authentication to prompt for the user ID and password (instead of requiring it in the URL), and a configuration directive to allow the user to specify whether FTP URLs will be treated as relative URLs or absolute URLs.

1.5.3.4 SNMP Subagent and MIB Support

WTE provides a Simple Network Management Protocol (SNMP) management information base (MIB) and SNMP subagent so you can use any SNMP-capable network management system, such as Tivoli NetView or Tivoli Distributed Monitoring to monitor your proxy server's health, throughput, and activity. The MIB data describes the proxy server being managed, reflects current and recent server status, and provides server statistics.

1.5.3.5 Performance Improvements

The cache architecture has been restructured to map URLs onto the cache file system more efficiently. This speeds retrieval of cached objects, uses disk space more efficiently within the cache, and speeds cache garbage collection. The cache also uses write-behind techniques for greater throughput.

Additionally, WTE now includes caching of Domain Name System (DNS) server lookup results, which can improve response time and reduce network load.

1.5.3.6 HTTP 1.1-Compliant Proxy Server

WTE is now an HTTP 1.1-compliant proxy server. WTE identifies itself as an HTTP 1.1 server and sends HTTP 1.1 in the outbound flows to origin servers. Persistent connections are supported from the client to the proxy, and from the proxy to the origin server. HTTP 1.1 cache control headers are processed and used to determine if a Web document is able to be cached. WTE receives and processes chunked data sent by HTTP 1.1 origin servers; WTE will unchunk data before giving control to a data filter or transmogrifier plug-in.

WTE provides new directives that allow the administrator to override certain HTTP 1.1 cache control headers. For example, query strings – URLs with a question mark (`?`) in them – are not considered cacheable by the HTTP 1.1 protocol, but WTE provides a directive that allows the administrator to specify which query strings should be cached.

WTE also provides an aggressive caching directive that allows the administrator to override the `Cache-Control: no-cache` header in Web documents.

1.5.3.7 Customization Exits

Several enhancements have been made to the WTE API to simplify writing an application plug-in. New request steps (*exit points*) have been added: *transmogrifier* and *garbage collection* (GC) advisor.

The transmogrifier gives the application write access to the outgoing data stream while the GC advisor allows the plug-in to influence garbage collection decisions:

- The transmogrifier step is intended to be used by applications that wish to perform transformations on the HTTP response data stream. Examples include converting Adobe PDF files to HTML, converting high resolution images to lower resolution quality, or translating pages from one language to another. The transmogrifier step is an extension of the data filter step.

  The WTE enhancements allow the application to specify multiple transmogrifiers, thus allowing the application to have multiple plug-ins, each of them performing different transformations on the data.

WTE introduces a correlator mechanism that eases state maintenance in the plug-in; WTE now automatically determines the content length of the response data, and the application no longer has to buffer the data to determine the content length. WTE also makes HTTP header processing easier; the response headers can now be extracted and set using API variables.

- The GC Advisor step is called for each file in the cache during the garbage collection process and allows the application to influence which files are kept and which are discarded.

1.5.3.8 Tivoli Ready

WTE is Tivoli Ready, which means it can be managed through either the Tivoli Enterprise Console (TEC), or through Tivoli Global Enterprise Manager (GEM). Supported Tivoli configurations include managing IBM applications that are installed on Tivoli-managed nodes, PC-managed nodes and endpoints in a distributed environment.

With the purchase of an IBM software product that carries the Tivoli Ready logo, you have the ability to manage your IBM software products through the Tivoli Enterprise management products, allowing you to automatically discover, monitor, and inventory one or more Tivoli Ready applications.

This Tivoli Ready instrumentation, when configured, provides you with the ability to:

- Graphically view the health of WTE through Tivoli GEM 2.2, TEC 3.1 or higher consoles
- Inventory WTE using Tivoli Inventory Version 3.2

1.5.3.9 Variant Caching

Variant caching extends the capabilities of the transmogrifier and allows WTE applications to request that WTE cache another version (*variant*) of the original document retrieved from the Web. This need arises when a plug-in performs a transformation on the original page retrieved from the Web. For example, if a plug-in translates a page from English to Italian, it would be advantageous to be able to cache not only the original document but the variant, as well. In fact, transformations of Web data are very CPU-intensive and degrade the performance of the proxy server.

By caching the variant, the number of transformations required is decreased, which improves overall proxy caching performance. The WTE API provides new predefined functions, HTTPD_variant_lookup() and HTTPD_variant_insert(), to allow the application to find a variant that has already been cached, and to insert a new variant into the cache.

1.5.3.10 Enhanced Log Maintenance

To effectively manage the space requirements for the four logs (error, access, proxy access, and cache access) generated by WTE, two improved functions have been implemented. They are compression and purging:

- Compression can be specified for the logs by date. The logs are compressed and stored with no regard for storage constraints.
- Purging is done by date and size of each file. The maximum size for each log is set during configuration of WTE. When the logs are purged by date, if the

maximum size for a log is still exceeded, daily logs are removed until the maximum size is no longer exceeded.

1.5.3.11 Error Message Personalization

The error messages sent to the client's browser are kept in HTML pages, which lets you change or alter the message. You can use this to give additional information, give additional instructions, or include things such as logos.

1.5.3.12 Secure Request Filtering

Filtering for Secure Sockets Layer (SSL) requests is done by using PICS filtering on the non-secure home page of the URL and applying the receive or block decision to the secure request.

1.5.3.13 Configuration Enhancements

Configuring the server using the browser, local or remote, is improved and more user friendly. The configuration screen is divided into three areas: navigation, workspace, and header. Navigation between different areas of the configuration process is always available in the navigation portion of the screen with a click of the mouse. The forms to be filled in are displayed in the workspace. The header portion contains a Help button and a Restart Server button. The server can be restarted anytime with the Restart Server button. When the Help button is clicked, another instance of the browser that does not interfere with the form being used displays the help information. The help information is divided into three types: field definition of the form being viewed, task oriented, and server general help.

1.5.4 New AFS Enterprise File System Function

The AFS component of WebSphere Performance Pack now provides some new features.

1.5.4.1 AFS Server for Windows NT

WebSphere Performance Pack Version 1.0 included AFS client and server for AIX and Solaris, but only the client was included for Windows NT. WebSphere Performance Pack V2 includes AFS Version 3.5 for Windows NT, which offers AFS server for Windows NT. With this component, you can store AFS files and directories, and run processes that provide servers on a Windows NT machine.

1.5.4.2 AFS Control Center

The AFS Control Center is a set of Windows NT-based tools for managing AFS cells. The AFS Control Center includes a User Administration graphical user interface (GUI) for account management and a Server Manager GUI for volume management. The Control Center helps simplify AFS server and account administration by letting administrators manage entire environments from a single Windows NT workstation, for example, monitoring file server utilization, transferring collections of files across servers, performing load balancing, or managing AFS accounts.

1.5.5 New Functions Available by Combining Components

As we said, WebSphere Performance Pack is composed of three main components: ND, WTE and AFS. Although these components can be installed separately, they are not different products; especially in this new version of WebSphere Performance Pack, new functions have been provided to integrate these components together and supply services that were not previously available.

1.5.5.1 Content Based Routing

Content Based Routing (CBR) is a new ND function that provides load balancing enhancements to the reverse proxy capabilities of WTE. CBR works in conjunction with WTE, which must be installed on the same machine.

The same rules used for Dispatcher can be used with CBR to load-balance requests over different sets of servers based on the client IP address, the entire URL, the protocol portion of the URL, the host portion of the URL, the path portion of the URL, the Referer HTTP header, or the User-Agent HTTP header.

1.5.5.2 Peak Load Management

Peak load management is the ability to detect and react to sudden increases in activity. In Version 1.0, ND allowed rules-based routing to alter the load-balancing algorithm based on the current connection rate, the number of active connections, and the time of day. In Version 2, an advisor for Dispatcher enhances load management on WTE by preventing Dispatcher from sending new requests to a WTE node that is engaged in garbage collection or cache refresh.

1.5.5.3 Remote Cache Access

RCA provides a way to share cache content among proxies. In many scenarios, multiple proxies are deployed *near* (in network terms) each other. Typically, they will have a load balancer in front of them for load balancing and high availability reasons. Each proxy has its own cache, and if the cached data cannot be shared, this results in cache space being wasted as multiple copies of the same document are stored. Also, since each cache is smaller than the sum of all the disks, the cache hit rate is lower (due to the smaller cache size).

RCA is a new, powerful feature implemented in WTE, and not available in the previous release of IBM WebSphere Performance Pack. RCA allows multiple proxy servers to cooperate to form cache arrays. Using RCA, multiple proxy servers can distribute the cache contents across their combined, logical cache to improve hit rates and reduce redundancy of cached content.

WTE uses new caching algorithms and information-sharing technologies to enable an Internet service provider (ISP) to manage its servers more efficiently by storing information where it is more likely to be needed and delivering it more efficiently to customers. These enhancements reduce transmission costs and eliminate the need for ISPs to replicate information in redundant proxy servers.

Although RCA is a new feature of WTE, it is best used in conjunction with the other components of WebSphere Performance Pack. In fact, RCA enables multiple peer WTE proxy servers, load balanced by one ND machine, to share the contents of their caches utilizing a shared file system, such as AFS, DFS, NFS, or Windows NT file sharing. We recommend the use of AFS, since it offers nondisruptive real-time replication of information across multiple servers, data consistency, availability, global stability and data consistency.

1.5.5.4 Integrated Configuration Assistance

To facilitate proxy caching, the install program automatically sets up caching with WTE. WebSphere Performance Pack also provides wizards to make configuration easier. Invoked from a browser, these wizards can configure key

WebSphere Performance Pack scenarios that involve all of the WebSphere Performance Pack components.

1.6 Who Can Benefit

IBM WebSphere Performance Pack allows you to design and use multiple architectures. The scenarios described in this section provide specific examples of how various ISP implementations can benefit from the use of IBM WebSphere Performance Pack.

1.6.1 Content Hosting Internet Service Providers

Content hosting ISPs can use WebSphere Performance Pack to more efficiently support and distribute the content from their own Web sites, and to provide more response- and cost-effective access to other sites.

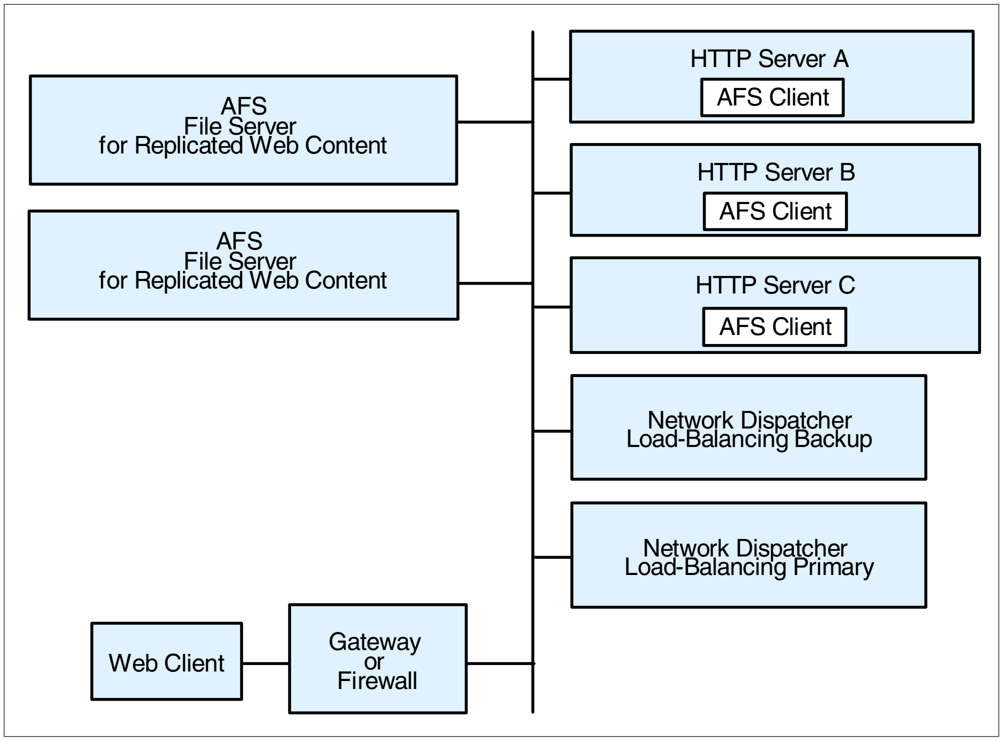

Figure 4. Content Hosting Internet Service Providers

Within a content hosting *server farm*, the WebSphere Performance Pack components can be configured to provide high availability and accessibility as follows:

- Content for hosted Web sites can be distributed over multiple volumes using the AFS virtual name space to simplify administration of page content. Because all of the Web servers need to have equal access to the Web content, a shared file system is an obvious choice for manageability of the Web content. AFS is a superior file system in this environment because of its replication capabilities, which provide improved availability and scalability. The disk caching capability of the AFS client ensures that network accesses to the file server are minimized.

- Scalability of access to the content can be achieved by adding multiple HTTP servers that are simultaneously AFS clients to access AFS server content, and by using efficient load balancing to dispatch requests to the HTTP server with the best capacity to handle the requests. ND, the load balancing component of IBM WebSphere Performance Pack guarantees improved availability and scalability by allowing a farm of Web servers to provide a single Web site image to clients.
- High availability can be maintained by using AFS file replication capabilities, and by configuring a hot standby ND component.
- Proxy caching can be used to provide more responsive access to content from other sites, as well as to optimize backbone network traffic capacity.
- Both availability and performance may be further enhanced by geographically distributing HTTP servers with AFS clients, together with proxy caching for other Internet content closer to user access points, such as points of presence (POPs) and network access points (NAPs).

Thus, content hosting service providers or corporate webmasters can benefit from the non-disruptive replication and distribution capabilities of the file sharing functions, local and wide area load balancing, and proxy caching. Flexible configuration of these components can ensure that requests are directed to the most appropriate local or remote location, and can enable location outages or routine maintenance schedules to be handled without disrupting customers.

The IBM WebSphere Performance Pack components can be used in conjunction with firewalls and authentication gateways to provide secure access where desired, and the load balancing function of WebSphere Performance Pack can also be used to scale these capabilities.

1.6.2 Corporate Web Sites and Content Aggregators

Use of IBM WebSphere Performance Pack by corporate Web sites and content aggregators is similar to that of content hosting ISPs.

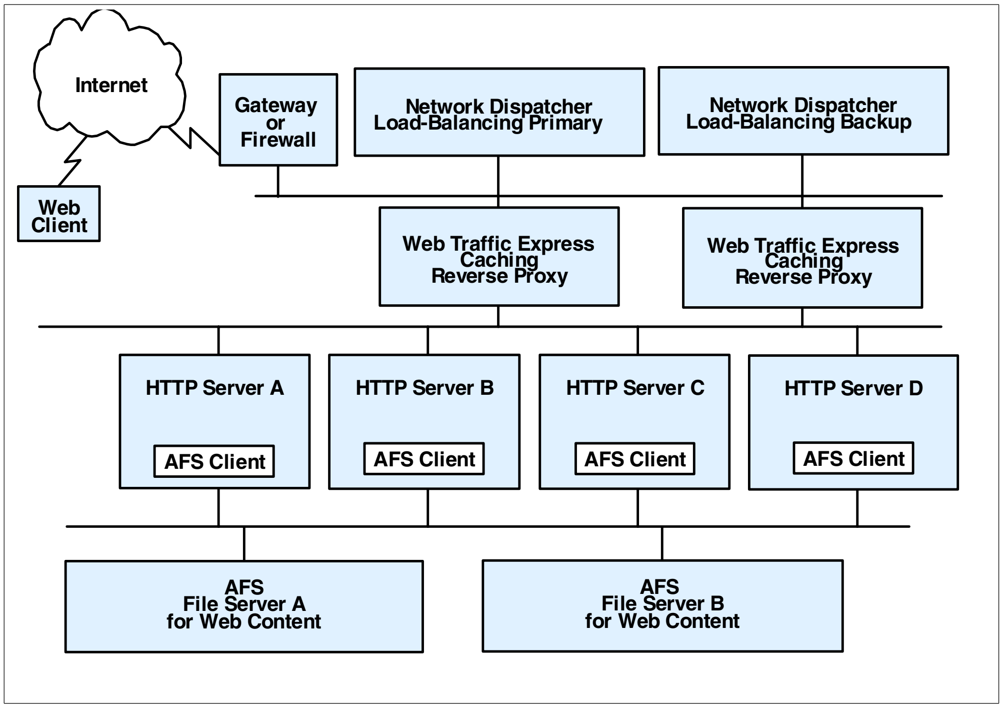
Internet
Web Client
Gateway or Firewall
Network Dispatcher Load-Balancing Primary
Network Dispatcher Load-Balancing Backup
Web Traffic Express Caching Reverse Proxy
Web Traffic Express Caching Reverse Proxy
HTTP Server A
AFS Client
HTTP Server B
AFS Client
HTTP Server C
AFS Client
HTTP Server D
AFS Client
AFS File Server A for Web Content
AFS File Server B for Web Content

Figure 5. Corporate Web Sites and Content Aggregators

In many cases, corporate Web sites and content aggregators need to maintain a demilitarized zone (DMZ) to ensure that access to Web content by employees, business partners, or customers does not expose internal computer resources to unauthorized users or hackers. However, firewalls cannot be deployed between AFS clients and their servers. In such instances, AFS replication can be used to establish read-only AFS servers within the DMZ. Here multiple AFS clients and the load balancing ND function can be used to provide the degree of scalability necessary to satisfy users. In addition, WTE caching and filtering proxy servers may be deployed on the same machines or on separate systems to filter or optimize access from within the corporation to external Web sites.

Many corporate Web sites are located at head office or regional locations, while branch offices or business partners may need frequent access to the content. Most offices of this type have relatively low line speed (somewhat less than 1.5 Mbps) network access to the regional or corporate sites. Relatively few users can, with concurrent usage, use all the available bandwidth with resulting erratic response times. At these locations HTTP servers with AFS clients combined with general purpose caching can provide more consistent user response time.

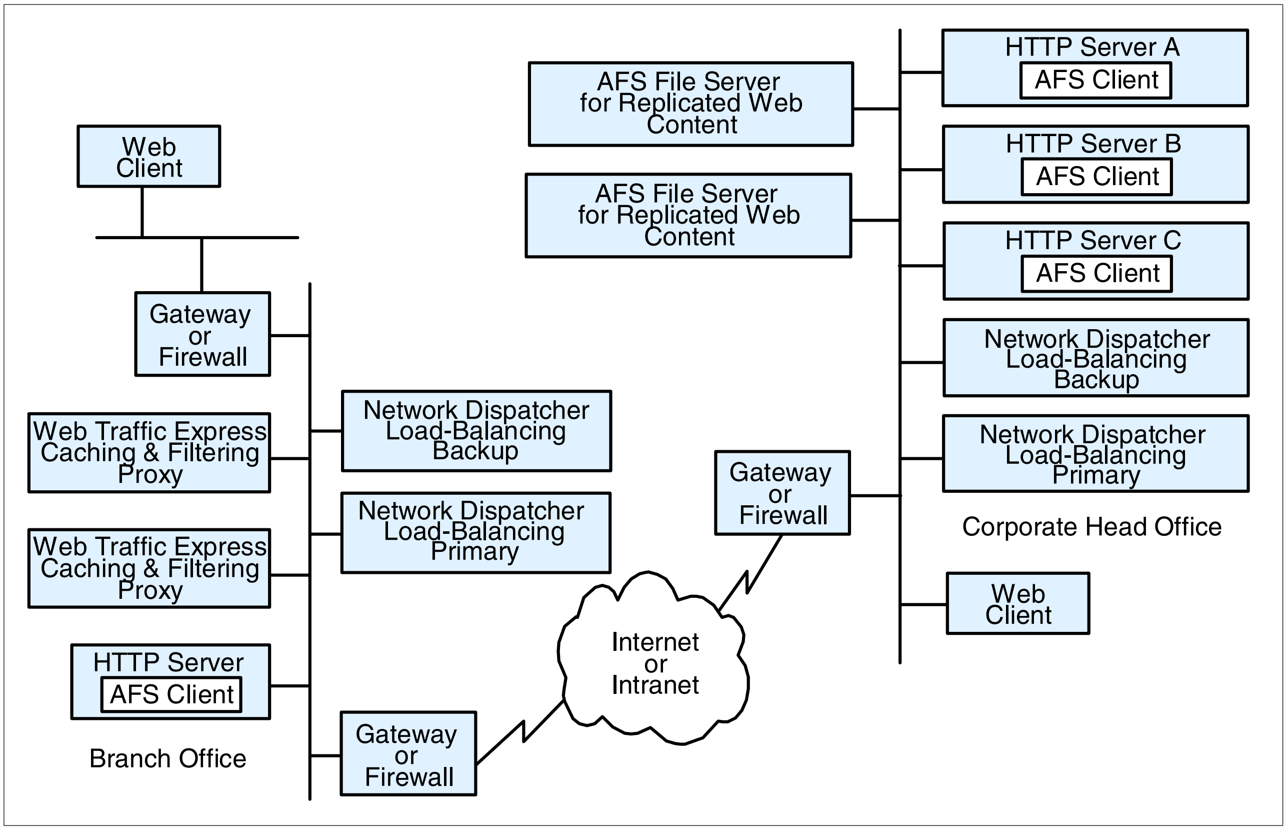
HTTP Server A
AFS Client
AFS File Server for Replicated Web Content
HTTP Server B
AFS Client
Web Client
AFS File Server for Replicated Web Content
HTTP Server C
AFS Client
Gateway or Firewall
Network Dispatcher Load-Balancing Backup
Web Traffic Express Caching & Filtering Proxy
Network Dispatcher Load-Balancing Backup
Network Dispatcher Load-Balancing Primary
Gateway or Firewall
Network Dispatcher Load-Balancing Primary
Corporate Head Office
Web Traffic Express Caching & Filtering Proxy
Web Client
Internet or Intranet
HTTP Server
AFS Client
Gateway or Firewall
Branch Office

Figure 6. Head Office - Branch Office

Some industries also have small branch offices in addition to larger branch or regional offices. A simple deployment of WTE caching and filtering on a single platform, eventually combined with a firewall, is probably sufficient to provide more consistent response time for employees in the small branch offices.[5] In the larger branches or regional offices, it may be desirable to have more than one WTE proxy server and to use the load balancing ND functionality to provide better resource management.

Two approaches can be used to reduce redundant page storage and to address the effectiveness of caching in these locations:

1. Hierarchical caching

   Optimize a primary (initial) cache for higher page hit rate by favoring more files of a smaller size, and a hierarchical cache for higher hit byte rate by favoring fewer files of a larger size.

2. RCA

   Configure RCA together with the load balancing function of ND and with AFS shared file storage to reduce the index size of each proxy and to increase the scalability of the caching storage.

[5] IBM has recently released a new product in the WebSphere family, called WebSphere Performance Pack WebSphere Cache Manager. This product is a caching and filtering proxy server obtained by the same code base as WTE. Cache Manager is available on Windows NT and Linux. Only a subset of the functions available in WTE are implemented in Cache Manager. For this reason, Cache Manager is particularly indicated for small companies or for small branch offices of large industries.

1.6.3 Corporate Headquarters Buildings or Large Campuses

On large campuses or in corporate headquarter buildings, the size of the campus or number of personnel frequently lead to the creation of smaller local area networks (LANs) interconnected by routers and a backbone LAN. Busy LAN servers combined with increasing use of Web server applications can result in congestion on the backbone LAN segments. This can be reduced by installing AFS client-enabled HTTP servers and general proxy caching on user LAN segments. The AFS clients can provide caching for corporate Web content, while the WTE proxy server can provide caching and filtering for external Internet access.

Design considerations for the smaller LAN segments are similar to those for small branch offices discussed earlier.

1.6.4 Backbone Internet Service Providers

Backbone ISPs typically provide co-location and/or peering services for other ISPs in addition to content hosting for large national or international corporations. In many cases they may provide *virtual ISP services* for other service providers such as content hosting ISPs. Backbone ISP customers are increasingly demanding both high availability and differentiated service levels, and backbone ISPs are responding by enhancing their Internet infrastructures. Features demanded by backbone ISPs include:

- Load balancing for a variety of traffic (for example, mail and FTP in addition to Web traffic)

  In addition, requests are made for traffic balancing management and high availability for authentication servers and management systems.

- Proxy caching

  To minimize the effects of *hot potato* routing and Web traffic *surges*, backbone ISPs are typically installing highly scalable caches at major peering points and network interconnections points such as network access points (NAPs).

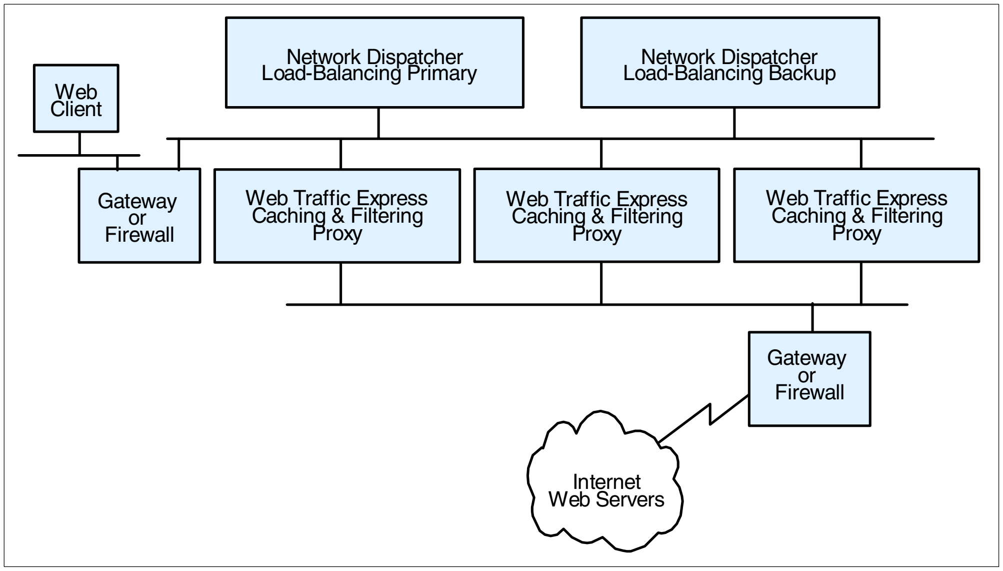
Network Dispatcher Load-Balancing Primary
Network Dispatcher Load-Balancing Backup
Web Client
Gateway or Firewall
Web Traffic Express Caching & Filtering Proxy
Web Traffic Express Caching & Filtering Proxy
Web Traffic Express Caching & Filtering Proxy
Gateway or Firewall
Internet Web Servers

Figure 7. Hierarchical Caching

International ISPs in particular need caching to reduce the costs associated with trans-oceanic links. Such installations can make very effective use of the RCA feature of WTE, a variation of the traditional caching function that when used together with load balancing and shared file storage can dramatically reduce redundancy of page storage.

Typical configurations for backbone ISPs would include load balancing for authentication gateways (such as authentication servers and subscriber management application servers) as well as for WTE caching proxy servers. Because of the amount of traffic and the desire for high availability, such caching servers would likely use the RCA feature and would thus also be configured as AFS clients. Therefore, there would also be an AFS server with the file content.

1.6.5 Access Internet Service Providers

Access ISPs need to provide both more consistent response time to their customers and to conserve their backbone network link and access charges. Thus caching at POPs would address these needs. These configurations would be similar to those for backbone ISP solutions.

Because access service providers target many of the small- to medium-size business customers, there is also an opportunity for them to create revenue producing services, by deploying smaller caching devices on customer premises but configuring and managing them as an ISP service. For this scenario, the caching would look much as it does for corporate branch offices.

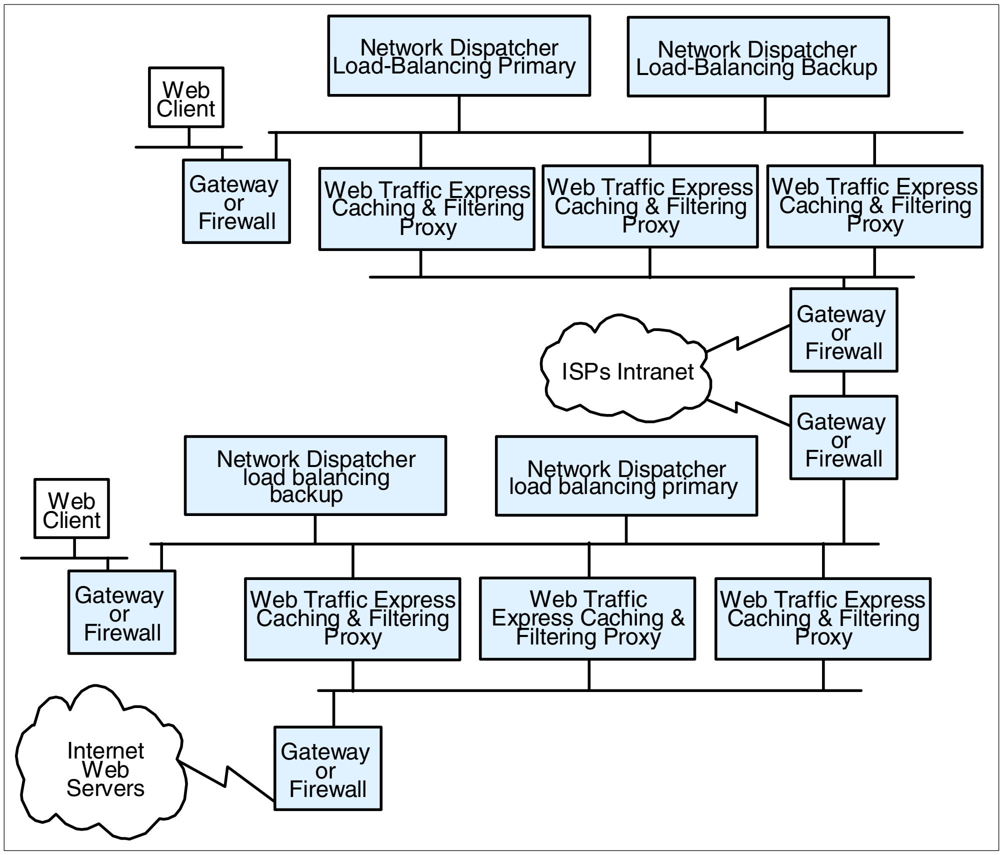
Network Dispatcher Load-Balancing Primary
Network Dispatcher Load-Balancing Backup
Web Client
Gateway or Firewall
Web Traffic Express Caching & Filtering Proxy
Web Traffic Express Caching & Filtering Proxy
Web Traffic Express Caching & Filtering Proxy
Gateway or Firewall
ISPs Intranet
Gateway or Firewall
Network Dispatcher load balancing backup
Network Dispatcher load balancing primary
Web Client
Gateway or Firewall
Web Traffic Express Caching & Filtering Proxy
Web Traffic Express Caching & Filtering Proxy
Web Traffic Express Caching & Filtering Proxy
Internet Web Servers
Gateway or Firewall

Figure 8. Access Internet Service Providers

1.6.6 Access ISPs with Subscriber Home Page Hosting

Access ISPs providing subscriber home page hosting typically do so on multi-homed servers, allowing subscribers to have their own domains. Such servers may have frequent changes, requiring the ISP to dedicate considerable time to administering the servers from a directory and backup perspective. For example, if the ISP is successful in recruiting new users, therefore having to expand the server capacity, it may involve reorganizing the various servers to spread the load and to allow room for individual sites to expand. The file management services provided by IBM WebSphere Performance Pack enable the ISP to reduce this maintenance workload, and to provide high availability replication for subscribers willing to pay for continuous availability. They also allow the ISP to scale up the site capacity with additional AFS clients, without the need to reorganize the underlying Web servers. ND load balancing can ensure that despite the growth users get rapid response time. In addition, when combined with subscriber management routings, to intelligently assign DHCP IP addresses, it becomes possible to offer differentiated classes of services for customers paying a premium.

1.7 Other IBM WebSphere Offerings

This section describes the other IBM WebSphere products that can interact with IBM WebSphere Performance Pack in an e-business Application Framework environment: IBM WebSphere Application Server and IBM WebSphere Studio.

Notice that IBM WebSphere Performance Pack can be integrated with a number of other IBM e-business products, such as IBM eNetwork Firewall, IBM HTTP Server and Lotus Domino Go Webserver. In this redbook, there are several examples and scenarios where all these e-business products are used together to create a powerful and secure Web site.

1.7.1 IBM WebSphere Application Server

IBM WebSphere Application Server lets you achieve your *Write Once, Run Anywhere* goal for Java servlet development. The product consists of a Java-based servlet engine that is independent of both your Web server and its underlying operating system.

WebSphere Application Server offers a choice of server plug-ins that are compatible with the most popular server APIs. The supported Web servers are:

- IBM HTTP Server
- Apache Server
- Domino
- Lotus Domino Go Webserver
- Netscape Enterprise Server
- Netscape FastTrack Server
- Microsoft Internet Information Server

IBM WebSphere Application Server V2.0 is available in Standard Edition and Advanced Edition.[6] In addition to the servlet engine and plug-ins, WebSphere Application Server Standard Edition provides:

- Implementation of the JavaSoft Java Servlet API, plus extensions of and additions to the API

[6] See `http://www.software.ibm.com/webservers/appserv/`.

- Sample applications demonstrating the basic classes and the extensions
- The IBM WebSphere Application Server Manager, a graphical interface making it easy to:
  - Set options for loading local and remote servlets
  - Set initialization parameters
  - Manage servlets
  - Specify servlet aliases
  - Create servlet chains and filters
  - Administer and monitor Enterprise Java Services (EJS) components
  - Enable Lightweight Directory Access Protocol (LDAP) directory support
  - Log servlet messages
  - Enable JVM debugging
  - Monitor resources used by Application Server
  - Monitor loaded servlets, active servlet sessions, and JDBC connections
  - Monitor errors, events, exceptions, and log output
  - Create dumps and data snapshots
  - Dynamically enable and disable tracing
- A connection management feature that caches and reuses connections to your JDBC-compliant databases

  When a servlet needs a database connection, it can get one from the pool of available connections, eliminating the overhead required to open a new connection for each request.
- Additional Java classes, coded to the JavaBeans specification, that allow programmers to access JDBC-compliant databases

  These data access beans provide enhanced function while hiding the complexity of using relational databases. They can be used in a visual manner in an integrated development environment.
- Support for dynamic page content called JavaServer Pages (JSP)

  JSP technology lets you produce dynamic Web pages with server-side scripting. The result is to separate your presentation logic (for example, the HTML code that defines your Web site structure and appearance) from your business logic (for example, the Java code that accesses a database for information to display on the Web site). For flexibility, JSP files can include any combination of inline Java, `<SERVLET>` tags, National Center for Supercomputing Applications (NCSA) tags, and JavaBeans.
- Enablement for LDAP supported directory services
- Modules and a command line interface for integrating Application Server and Apache Server into the Tivoli environment for distributed monitoring and operations.
- eXtensible Markup Language (XML) Document Structure Services

WebSphere Application Server Advanced Edition provides all the features of Standard Edition, plus:

- Enterprise Java Services

  This function is provided to run and manage applications coded to Sun's Enterprise JavaBeans (EJB) specification.

- Common Object Request Broker Architecture (CORBA) support, enhanced to provide both bean-managed and container-managed persistence.

IBM WebSphere Application Server will also be available in an Enterprise Edition, which will include the same features as the Advanced Edition, plus:

- TXSeries support, IBM's world-class transactional application environment
- Component Broker (CB) support, with its fully distributed object and business process integration capabilities

1.7.2 IBM WebSphere Studio

IBM WebSphere Studio is a suite of tools that can be used by people involved in creating and maintaining Web sites. It allows your team to:

- Easily create Java beans, database queries, and Java servlets using the Studio wizards
- Group your Web site files into projects and folders
- Edit and update the files with your preferred tools
- Publish all or part of the Web site on any of your WebSphere Application Server systems
- Maintain the files locally on individual workstations, or in a central location using a source control system

IBM WebSphere Studio is the tool complement of IBM WebSphere Application Server. This suite of tools makes it easier to design, develop and maintain dynamic, interactive Web sites.

The tools provided by IBM WebSphere Studio are:

- Studio workbench

  The workbench helps you manage and maintain your Web site applications and files. It uses projects and folders to group and organize the files and lets you perform all other functions from this central location. The Studio supports all file types. You can create files, open them, edit them, copy them, move them around, delete them, publish them. And, you can do all this right from the workbench using the workbench menu functions, the wizards, and your preferred Web development software.

- Studio wizards

  The wizards are the fastest way to add dynamic content to your Web pages. They help you retrieve information from common databases, use server-side JavaBeans, capture information about your customers, and register Web visitors. You don't have to be expert at SQL syntax or Java programming. The wizards walk you through step-by-step and then generate sophisticated servlet code for you.

- Companion products

For added convenience the Studio comes with an integrated set of companion products. Everything is included for you to build, manage, and publish complete Web sites:

- NetObjects ScriptBuilder

  Use this text-based editor for files that contain common markup, scripting, and programming languages, such as HTML, DHTML, and JavaServer Page (JSP) extensions, JavaScript, JScript, and Java. Its features make the scripting process easy. You can quickly preview scripted Web pages, reference and add language elements, and navigate to embedded functions and objects.

- NetObjects Fusion

  Build your Web sites in a visual manner, dragging and dropping pages to create an interconnected hierarchy. You can view individual pages or the entire site structure. This powerful tool lets you include images, multi-media, and dynamic HTML, and control the visual appearance of an entire site from one central location.

- NetObjects BeanBuilder

  BeanBuilder allows you to put an applet in your Web site in a very short time. This visual authoring tool lets you quickly combine Java beans into new applets.

- VisualAge for Java, Professional Edition

  If you are familiar with Java programming, you can use this robust, full-function environment to create and customize Java components. This award-winning tool includes advanced functions, such as incremental compilation and the ability to invoke methods while debugging. You can use it to build sophisticated Java beans that you can use in the Studio basic servlet wizard. And, in turn, you can also use it to modify the Java servlets and beans generated by the Studio wizards.

- WebSphere Application Server

  What good is server-side Java logic without a server that can handle it? For your convenience, the best place to publish your Web sites is right in the package. The WebSphere Studio comes with a complimentary copy and a developer's license for the WebSphere Application Server.

- Base HTTP servers

  The WebSphere Application Server runs on several HTTP servers and more than one operating system. But, just in case you don't have one handy, the Apache Web server and a base HTTP server are included for good measure.

Chapter 2. IBM AFS Enterprise File Systems Concepts

This chapter introduces the file sharing component of IBM WebSphere Performance Pack Version 2. This component is based on a distributed file system from the Andrew research project at Carnegie Mellon University, that was called the *Andrew File System*. The principals involved in this research project formed Transarc Corporation and this file system developed into a product. In the process, the *Andrew* name was discontinued and AFS became a registered trademark of Transarc Corporation. As Transarc changes into IBM as the IBM Transarc Lab, all Transarc trademarks and logos are transferring to IBM. Hence, the file sharing component of Performance Pack is also referred to as IBM AFS Enterprise File System.

2.1 What Is New in Version 2?

The AFS component of WebSphere Performance Pack V2.0 provides AFS Version 3.5 server and client processes and the AFS Control Center for the Windows NT operating system. The previous version of WebSphere Performance Pack, on the Windows NT platform, included only the AFS client component. WebSphere Performance Pack V2.0 continues to include AFS Version 3.4a server and client processes for IBM AIX and Sun Solaris operating systems.

2.1.1 AFS 3.5 Server for Windows NT

Performance Pack V2.0 offers AFS Version 3.5 server and client processes on Windows NT. An AFS server machine, running Windows NT Server or Workstation V4.0, can perform one or more of the following functions:

- AFS File Server

  A machine running AFS File Server processes stores files. It also delivers files to AFS clients and receives files from AFS clients.

- AFS Database Server

  A machine running AFS Database Server processes maintains databases of information about AFS security, users, protection groups, and volumes.

- AFS Backup Server

  A machine running the AFS Backup Server process maintains a database containing information used by the AFS Backup System to dump data to tape (for example, volumes that will be dumped together, the schedules that will be followed, and contents of tapes).

- AFS System Control Machine (SCM)

  A machine running AFS system control server processes stores master copies of AFS system configuration files and distributes these files to other AFS server machines that request them.

AFS 3.5 for Windows NT and Service Pack Version

At the time of the writing of this redbook, AFS 3.5 is only supported on Windows NT Server or Workstation V4.0 (Service Pack 3).

WebSphere Performance Pack V2.0 also provides the AFS Server Configuration wizard to facilitate the configuration of AFS server machines.

2.1.2 AFS 3.5 Control Center for Windows NT

WebSphere Performance Pack V2.0 includes the AFS Control Center on Windows NT. This component consists of two Windows-based graphical user interface (GUI) tools:

- The AFS Server Manager tool performs administrative tasks relating to AFS volumes and AFS services on AFS server machines, running Windows NT, AIX, or Solaris.
- The AFS Account Manager tool assists with user and group account administration tasks.

Both of these tools only run on Windows NT machines.

2.1.3 AFS Management with the Common Configuration Assistance

WebSphere Performance Pack V2.0 integrates a Common Configuration framework to assist in AFS management, through the use of wizards, forms, and documentation. By using a browser, these wizards will assist you in the management of Web contents on your AFS servers through tasks such as setting up, adding, replicating, moving, and deleting AFS Web volumes.

2.2 Features of AFS

AFS is a *distributed file system* that allows users to share and access all of the files stored in a network of computers as easily as they access the files stored on their local machines. The file system is called distributed because files may reside on many different machines, but are available to users on every machine.

In fact, AFS stores files on a subset of the machines in a network, called *File Server* machines. File Server machines provide file storage and delivery service, along with other specialized services, to the other subset of machines in the network, the client machines. These machines are called *clients* because they make use of the servers' services while doing their own work. In a standard AFS configuration, clients provide computational power, access to the AFS files and other general purpose tools to the users seated at their consoles. There are generally many more client workstations than File Server machines.

AFS servers run a number of *server processes*, so-called because each provides a distinct specialized service: one handles file requests, another tracks physical location of data, a third manages security, and so on. To avoid confusion, AFS documentation always refers to *server machines* and *server processes*, not simply to *servers*.

2.2.1 The Concept of an AFS Cell

An AFS *cell* is a collection of client and server machines that form an administrative unit. A cell includes the machines, the AFS data on the File Server machines, the AFS databases on the Database Server machines and the accounts created to access the collection of machines and their resources.

The cell name is the second element in the fully qualified path of AFS data. For example, data in the transarc.com cell would reside under /afs/transarc.com in the directory tree.

Each AFS user account and each AFS machine belongs to exactly one cell. This is referred to as the user’s or machine’s *local* cell. A user may be able to access other cells, which are referred to as *foreign* cells.

2.2.2 Single File Name Space

Although your AFS cell is administratively independent, you will probably want to organize your local collection of files (your *file space* or *directory tree*) so that users from other cells can also access this information. AFS allows cells to combine their local file spaces into a global file space, and does so in such a way that file access is transparent. Users do not need to know anything about a file’s physical location in order to access the file. All they need to know is the pathname of the file, which includes the cell name. Thus every user at every machine sees the collection of files in the same way, meaning that AFS provides a uniform *name space* to its users.

2.2.3 AFS Volumes

AFS groups files into volumes, making it possible to distribute files across many machines and yet maintain a uniform name space. A volume functions like a container for a set of related files, keeping them all together on one partition. In fact a *volume* is a collection of related files and directories and is part of the directory tree. This enables server administrators to break the directory tree down into smaller units that can then be stored on File Servers. Volumes can vary in size, but are by definition smaller than a partition, since partitions hold one or more volumes. We recommend a maximum size of 2 GB. You can read or write to AFS volumes that are larger than 2 GB, but you cannot perform typical AFS volume operations, such as dumping, restoring, moving or replicating the volume.

Volumes are important to server administrators and users for several reasons. Their small size makes them easy to move from one partition to another, or even between machines. The server administrator can maintain maximum efficiency by moving volumes to keep the load balanced evenly. In addition, volumes correspond to directories in the file space; most cells store the contents of each user home directory in a separate volume. Thus the complete contents of the directory move together when the volume moves, making it easy for AFS to keep track of where a file is at a certain time. Volume moves are recorded automatically, so users do not have to keep track of file locations.

2.2.4 Replication and Caching

AFS incorporates special features on server machines and client machines that help make it efficient and reliable:

- On server machines, AFS allows administrators to replicate *frequently accessed but infrequently changed volumes*, such as those containing Web pages with product information. *Replication* means putting an identical ReadOnly copy (sometimes called a *replica*) of a volume on more than one File Server machine. The crash of one File Server machine housing the volume does not interrupt users’ work, because the volume’s contents are still available from other machines. Replication also means that one machine does not become overburdened with requests for files from a popular volume.

- On client machines, AFS uses *caching* to improve efficiency. When a user on a client workstation requests a file, a process named *Cache Manager*, running on the client, sends a request for the data to the proper File Server machine. The user does not need to know which machine this is; the Cache Manager determines file location automatically. The Cache Manager receives the file from the file server and puts it into the cache, an area of the client machine's local disk or memory dedicated to temporary file storage.

  Caching improves efficiency because the client does not need to send a request across the network every time the user wants the same file. Network traffic is minimized, and subsequent access to the file is especially fast because the file is stored locally.

2.2.4.1 Cached Data Validation

AFS has a way of ensuring that the cached file stays up-to-date, called a *callback*. When a user or process on an AFS client machine makes an update to a file, that file will be updated in the Cache Manager for an AFS client. When the file is changed, and the update is finished, the file is written back to the AFS server, in the appropriate ReadWrite volume. When a file server's copy of the file is updated, that server will notify each of the AFS clients that have a copy of that file. Those clients will be informed that a new copy of that file is available and they will need to retrieve that copy the next time a user or a process requests that specific file.

Notice that in order to copy modified data to each of the replicas, the system administrator must issue a particular command:

```
vos release
```

The volserver server process handles this replication task, but it is not automatic. The administrator decides when to make this modification available and runs the command above to do this.

2.2.5 Content Backup

Besides the volume that the user works with (the user's *home volume*), a snapshot of the user's directory is usually made every night. Thanks to the design of AFS, this snapshot is a volume, called a *backup volume*, which can also be attached to the file name space to allow easy retrieval of files mistakenly deleted. It is common for backup volumes to be created for all volumes in the system, not only user volumes, so the archive process can use this stable and consistent image when writing the data to tape. The server administrator can set up the system, so users themselves can correct careless file deletions with no system intervention.

Notice that the creation of a backup volume every night is *not* automatic. The server administrator sets this up if desired with a bos cron job. This is similar to a UNIX cron job in that it runs at a specified time, but under control of the BOS Server. For example, each server could have the following process running:

```
Instance backup, (type is cron) currently running normally.
Auxiliary status is: run next at Thu Nov  5 13:00:00 1998.
Process last started at Wed Nov  4 13:00:03 1998 (4 proc starts)
Last exit at Wed Nov  4 13:00:09 1998
Command 1 is '/usr/afs/bin/vos backupsys user. server1 -localauth'
Command 2 is '03:00'
```

This process, which is run on server1, will create a backup copy of each volume whose name begins with `user` on the server called server1 at 3:00 a.m. each morning.

You can obtain a similar output by entering the command:

```
bos status server1 -long
```

2.2.6 AFS Security

Even in a cell where file sharing is especially frequent and widespread, it is not desirable that every user have equal access to every file. One way AFS provides adequate security is by requiring that servers and clients prove their identities to one another before they exchange information. This procedure, called *mutual authentication*, requires that both server and client demonstrate knowledge of shared secret information (such as a password) known only to the two of them. Mutual authentication guarantees that servers provide information only to authorized clients and that clients receive information only from legitimate servers.

AFS utilizes algorithms and other procedures based on *Kerberos*. This technology was originally developed by the Massachusetts Institute of Technology's Project Athena.

Users themselves control another aspect of AFS security, by determining who has access to the directories they own. Notice that every AFS directory has an *access control list* (ACL). For any directory a user owns, the user can build an ACL that grants or denies access to the contents of the directory. An ACL pairs specific users (or groups of users) with specific types of access rights. There are seven separate access rights and up to 20 different users or groups of people may appear on an ACL. We will discuss more details of this in 2.6.6, "Protecting AFS Files and Directories" on page 38.

2.3 Why Do I Need IBM AFS Enterprise File System?

AFS, the File Sharing component of IBM WebSphere Performance Pack, enables the easy sharing of information, and also builds a management infrastructure that supports ever-increasing file sharing demands without a proportional increase in cost. It is not just a way to share files, but it is a way for an organization to control its distributed file services.

From what we have seen, AFS has many benefits, such as:

- High scalability because of architectural protocol efficiencies
- Ease of sharing via its single, shared name space
- Reliable access to replicated data

- Single point of administration based on distributed services
- An efficient security model for access and group management
- Superior support for clients and desktops of different architectures

2.4 AFS and NFS

Network File System (NFS) from Sun Microsystems is another distributed file system option. NFS lacks many of the scalability and performance features available with AFS, and does not implement the common name space feature of AFS. Other major differences between AFS and NFS are in the areas of performance, availability, management and security.

Interoperability between AFS and NFS is possible with Transarc's NFS/AFS Translator. This tool is particularly effective for NFS installations that are transitioning to AFS. Throughout a transition period, machines that have not been converted to AFS can access files stored in the AFS name space. The Translator also provides an AFS solution for those system types that are not directly supported with a full AFS port.

2.5 Machine Responsibilities

AFS is a distributed file system that uses the client/server model of computing. AFS servers are physically secure machines that store AFS files and directories, databases of AFS information, and AFS configuration files. AFS client machines are used by server administrators and users to access information from AFS servers and issue commands to interact with AFS servers (for example, create protection groups or manage server processes).

AFS Users in WebSphere Performance Pack Scenarios

When we refer to AFS users in WebSphere Performance Pack scenarios, we are referring to Web authors responsible for creating Web pages that are stored in AFS file space or Web servers configured to access Web pages from AFS file space.

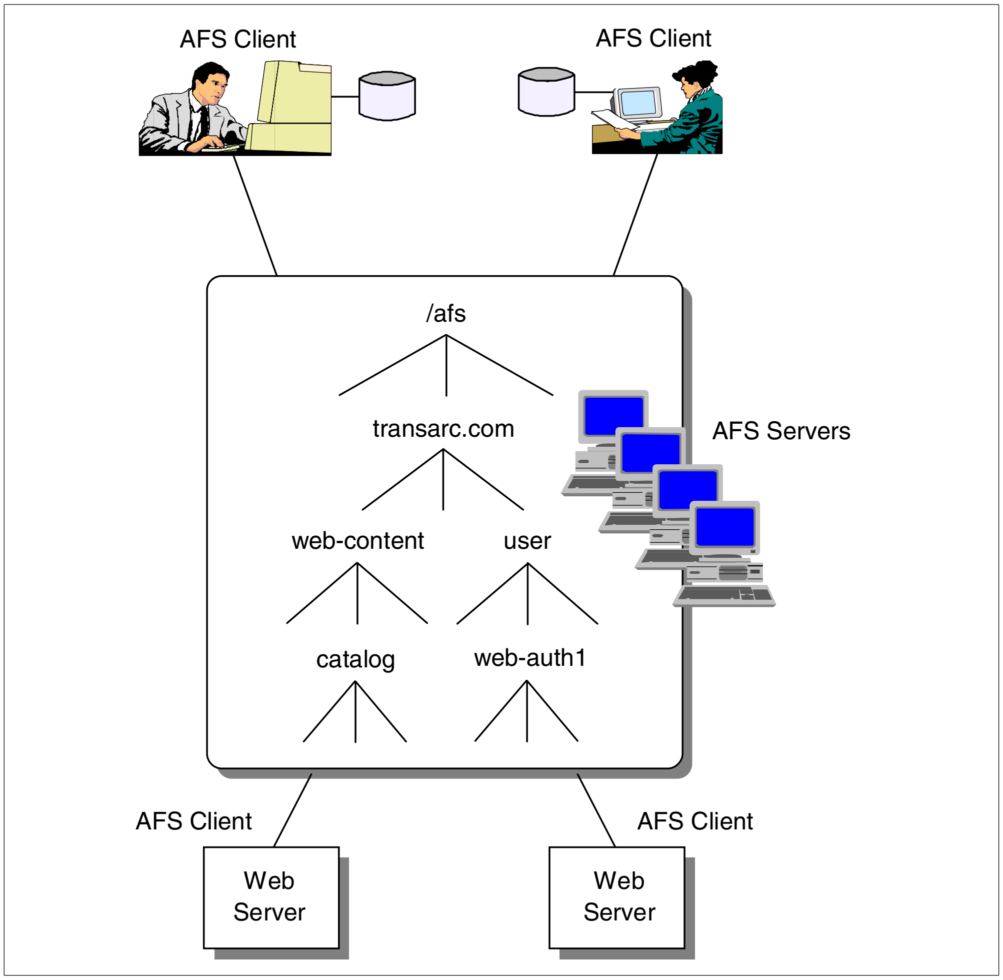

Figure 9. Machine Responsibilities

2.5.1 AFS Servers

An AFS server's responsibilities are determined by the processes it runs. The responsibilities of AFS servers include:

- **File Server**

  A machine running AFS File Server processes stores files. It also delivers files to AFS clients and receives files from AFS clients.

- **AFS Database Server**

  A machine running AFS Database Server processes stores and maintains AFS Databases:

  - The *Authentication Database* stores an encrypted password for each account, an encrypted server key, and account security information.
  - The *Volume Location Database* (VLDB) stores information about volumes and their physical location.
  - The *Protection Database* contains information about AFS users and the corresponding user IDs, AFS groups and their corresponding group IDs, and group membership.

- The *Backup Database* contains information used by the AFS Backup System to dump data to tape (for example, data that will be dumped together, the schedules that will be followed, and contents of tapes).

  The AFS server on Windows NT makes a distinction between a Database Server and a Backup Server. The Database Server on Windows NT houses the Authentication, Protection, and Volume Location Database. The Backup Server on Windows NT only houses the Backup Database.

- **AFS System Control Machine (SCM)**

  A machine functioning as an AFS SCM stores master copies of AFS system configuration files and distributes these files to other AFS server machines that request them. It is important that the configuration files are the same on each AFS server in your environment. They contain information about the machine's local cell, the names of Database Servers in the cell, server encryption keys, and a list of server administrators. You should have one SCM per cell.

- **AFS Binary Distribution Machine**

  A machine functioning as an AFS Binary Distribution Machine (BDM) stores master copies of AFS binaries for a particular operating system and distributes these files to other AFS server machines that request them. You may choose to set up one BDM for each operating system type. This allows you to install AFS binaries on one machine and have all AFS servers of the same operating system type pull the binaries from that machine. It is not necessary to have a BDM, although it is recommended if you have multiple servers of the same operating system type.

2.5.2 AFS Clients

An AFS Client machine runs the Cache Manager, a set of kernel modifications responsible for communicating with processes on AFS servers and accessing AFS files and directories (see 2.2.4, "Replication and Caching" on page 29). The Cache Manager is responsible for communicating with processes on AFS server machines and accessing AFS files and directories.

2.6 The AFS File System

AFS files and directories are organized into a directory tree. The root of the directory tree is /afs. All AFS client machines see the same directory structure under /afs. This is known as the *single file name space*.

To access an AFS file or directory, the user just specifies the path. The user does not need to know the physical location of the file or directory. In fact, the physical location could change and the user would continue to use the same path. This is known as *location independence*.

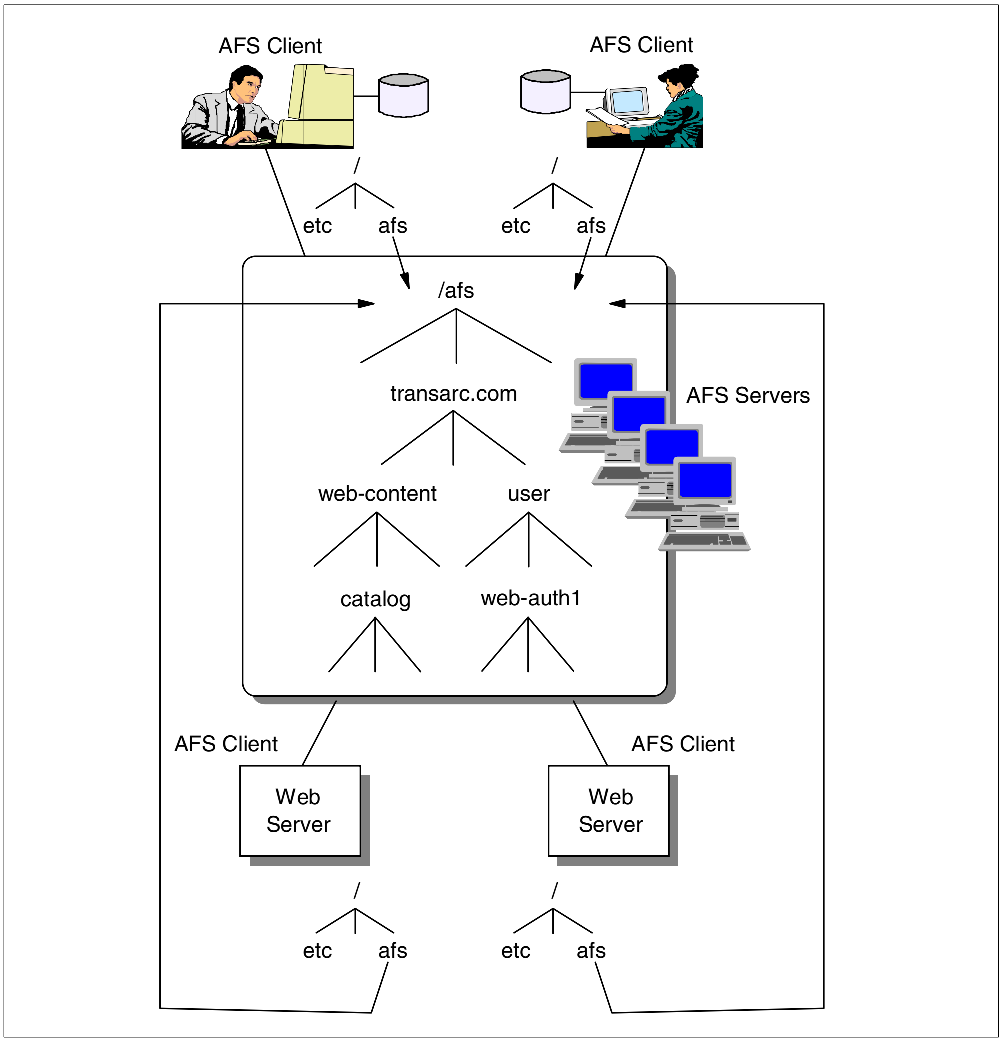

Figure 10. AFS File System

2.6.1 Volumes

The AFS directory tree is broken down into smaller subtrees. Each subtree is a collection of related files and directories called a *volume*. AFS volumes provide a unit of administration for replicating frequently accessed files and directories, moving data from one physical location to another. AFS volume size should be restricted to 2 GB. As we mentioned in 2.2.3, “AFS Volumes” on page 29, it is possible to have a volume larger than 2 GB, but then it will not be possible to perform certain administrative tasks, such as dumping, restoring, moving or replicating the volume.

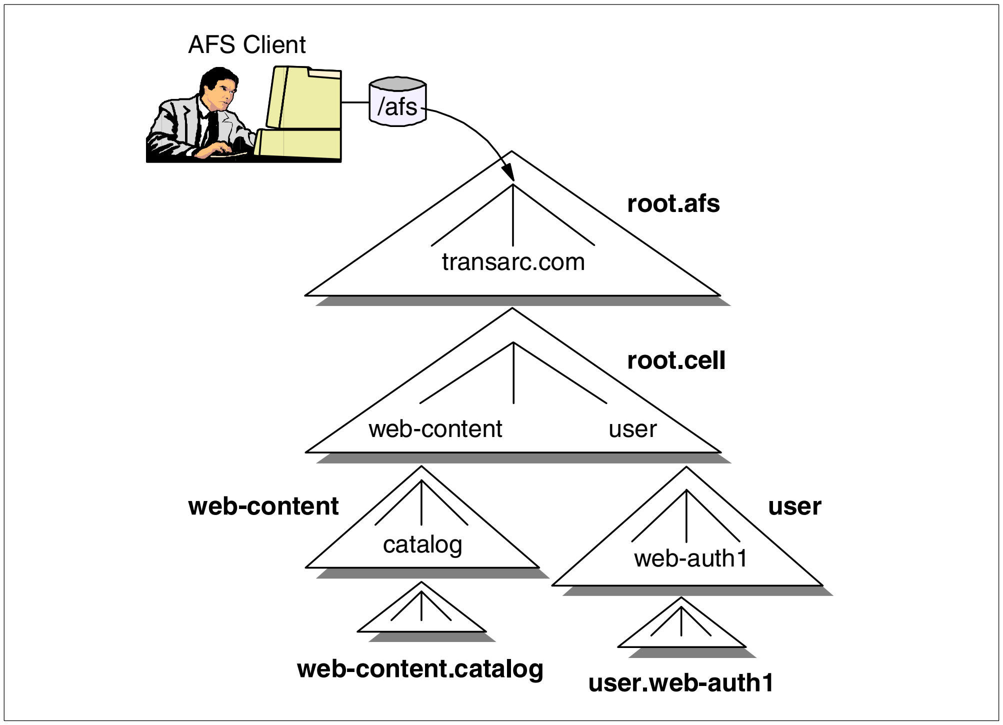

Figure 11. AFS Volumes

Each AFS volume is identified by a name. The root volume is called root.afs by convention. The next volume is typically the root.cell volume.

AFS volumes are stored on the disk partitions of File Server machines. AFS partitions are often referred to as *vice partitions* because they have names in the form /vicep*x*, where *x* is a letter in the range from a to z and from aa to iv. AFS partition names must be unique on a per server basis. In other words, it is possible to have a partition called /vicepa on the server wspp1 and also have a partition called /vicepa on the AFS server wspp2. The AFS partition size is constrained by the underlying operating system.

2.6.2 Mount Points

The point at which a volume is attached to the AFS directory tree is called a *mount point*. To users, mount points look like subdirectories. A mount point contains the name of the volume being attached to the directory tree. It does not contain information about the volume's physical location.

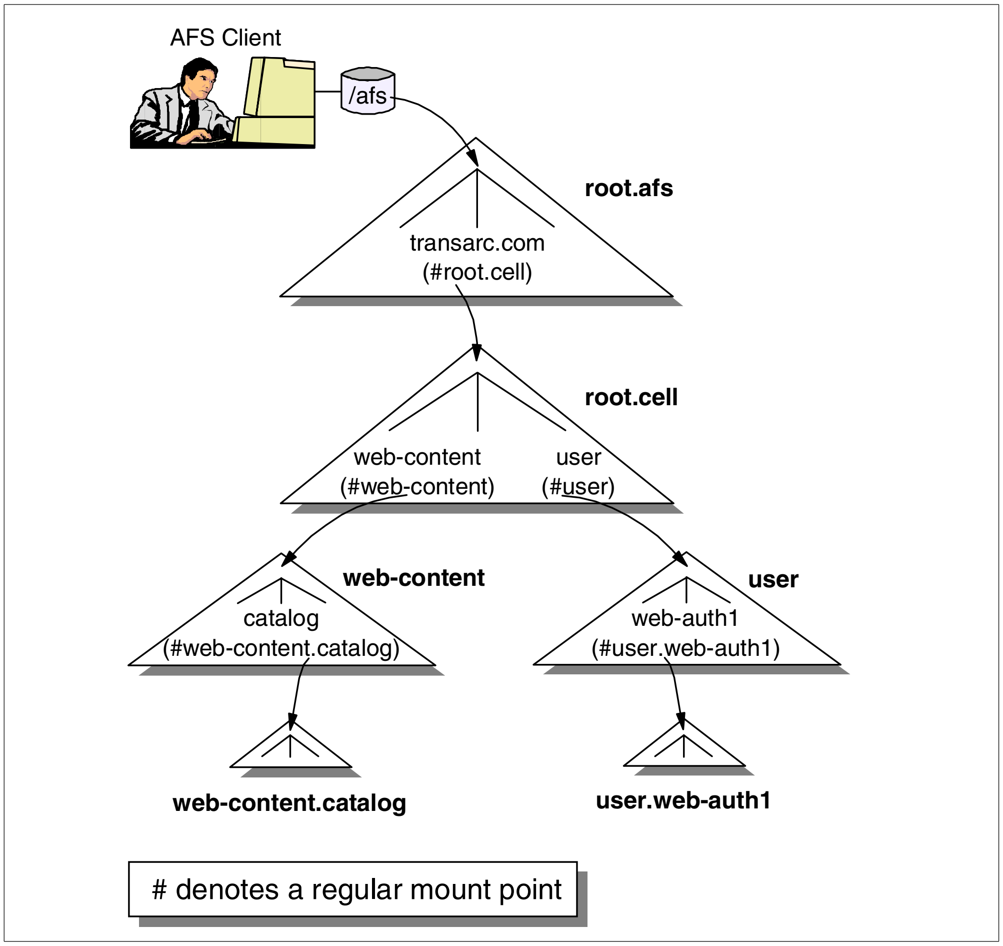

Figure 12. AFS Mount Points

2.6.3 Volume Location Database (VLDB)

When a user specifies a path for an AFS file, the AFS client machine's Cache Manager interprets the path. The Cache Manager can extract the volume name from the mount point. To access the volume, the Cache Manager must determine where the volume resides, so it contacts the VLDB. The VLDB notifies the Cache Manager of the volume's location(s). The Cache Manager then contacts a File Server on which the volume resides and requests the file.

2.6.4 Caching

When the Cache Manager receives the requested file, it caches the file on the AFS client machine's local disk. Subsequent requests for the file may use the cached copy instead of fetching another copy from the File Server.

2.6.5 Callback Mechanism

When the File Server sends the file to the Cache Manager, it also sends a *callback*. The File Server will notify the Cache Manager, by breaking the callback, if the file changes. When a user requests a file that is already cached, the Cache Manager checks to see if the file has a valid callback. If so, the request is satisfied from the cache. If the callback has been broken, the Cache Manager retrieves the file from the File Server.

2.6.6 Protecting AFS Files and Directories

ACLs are used to protect AFS data. Every AFS directory has an ACL that controls access to the directory and its contents. The ACL is applied at the directory level, not the file level as in UNIX.

An ACL consists of *entries*. Each entry specifies who has permission to access the directory and files within the directory. The entry also specifies what access rights are granted to that user or group of users.

There are seven access rights, or permissions, that you can assign on an ACL:

- lookup (l)

  This permission allows you to list the contents of the directory and examine the directory's ACL. It is necessary in conjunction with other permissions.

- read (r)

  This permission allows you to read the contents of files in the directory.

- insert (i)

  This permission allows you to add new files or subdirectories to the directory. You can do this by creating the files with an editor or by copying the files from another location.

- write (w)

  This permission allows you to modify the contents of files and to use the UNIX `chmod` command on files in the directory.

- delete (d)

  This permission allows you to remove or move files from the directory.

- lock (k)

  This permission allows you to run programs that set an advisory lock on files in the directory.

- administer (a)

  This permission allows you to change the ACL on the directory.

Server administrators and users can define AFS protection groups that contain a list of AFS users or machines. This protection group can be added to ACLs to give all members of the group the same access rights.

2.7 AFS Commands

There are two types of AFS commands: *simple commands* and *compound commands*. Simple commands consist of one word. Syntax is available for each of these simple commands by issuing the command with the `help` switch. Compound commands consist of two components: a command suite and a subcommand.

AFS compound commands share a common command syntax:

```
suite subcommand -switch instance[+] -flag
```

A list of the AFS command suites is summarized in the following table:

Table 1. AFS Command Suites

Command Suite	Purpose	Used by
`fs`	To interact with a File Server and to configure the Cache Manager	Users and server administrators
`pts`	To create, delete, and modify information in the Protection Database	Users and server administrators
`vos`	To create, delete, and modify a volume structure and/or a VLDB entry	Primarily server administrators
`bos`	To administer server processes and configuration files	Primarily server administrators
`backup`	To use the AFS Backup System	Server administrators
`kas`	To create, delete, and modify information in the Authentication Database	Primarily server administrators
`uss`	To create user accounts	Server administrators

A command suite has multiple subcommands. When issuing a compound command, specify the command suite followed by the subcommand. To list the subcommands available for a command suite, specify the command suite followed by `help`. For example:

```
fs help
```

Once you know a command suite and subcommand, you can obtain help in one of two ways:

- For a description, the alias (if it exists), and syntax, specify:

  suite `help` *subcommand*

  For example:

```
# fs help setacl
fs setacl: set access control list
aliases: sa
Usage: fs setacl -dir <directory>+ -acl <access list entries>+ [-clear ]
[-negative ] [-id ] [-if ] [-help ]
```

- For only the command syntax, specify:

  suite subcommand `-help`

  For example:

```
# fs setacl -help
Usage: fs setacl -dir <directory>+ -acl <access list entries>+ [-clear ]
[-negative ] [-id ] [-if ] [-help ]
```

2.8 AFS Authentication

When you authenticate, you prove your identity to AFS. Before you authenticate, you have access to public AFS data, data that the system:anyuser group has permission to access via ACLs, and data on the client machine’s local disk that you have permission to access via the operating system's permissions. After you authenticate, you have additional access to protected AFS data, data that you or an AFS protection group you belong to has permission to access via the ACLs.

You can authenticate, or prove your identity to AFS, by logging in via an AFS modified login or by issuing the `klog` command. You receive a token as evidence that you have been authenticated in a cell.

AFS includes a modified version of login for most AFS client machine supported system types. On AIX, this can be implemented by using the AFS modified login binary in place of the vendor-provided login binary. AFS integrated logon for Solaris involves installation and configuration of a Pluggable Authentication Module (PAM); this integrates authentication mechanisms, such as login, to provide security infrastructure for all authenticated access to and from the machine. On Windows NT, you implement integrated login by selecting a check box for integrated login in the AFS Client Configuration window. This modified login looks like the local operating system login, but the user is authenticated to AFS at login time. At login time, the user specifies a user name and password. If the password matches the user's password in the Authentication Database, the user is logged in to the local operating system and is granted an AFS token. If the password does not match the user's password in the Authentication Database, the user does not receive a token, but the user is logged in to the operating system if the password matches that of the local operating system.

To successfully authenticate via AFS modified login, AFS user names must match those of the local operating system. On Windows NT, passwords must also match.

A token is evidence that the user is authenticated in a cell. Tokens have a default lifetime of 25 hours, although this can be changed by server administrators. The maximum token lifetime is 720 hours (30 days). When a token's lifetime is exceeded, the token expires and the user must reauthenticate. A user can reauthenticate before a token expires to refresh the token's lifetime.

2.9 AFS Server Concepts

AFS servers are physically secure machines that store AFS files and directories, databases of AFS information, and AFS configuration files.

2.9.1 Types of AFS Servers

In 2.5.1, "AFS Servers" on page 33, we saw that there are four types of AFS servers:

- AFS File Server
- AFS Database Server
- AFS System Control Machine
- AFS Binary Distribution Machine

An AFS server can, and usually does, have multiple responsibilities. For example, it is common to configure an AFS server to function as a File Server and a Database Server. You do this by running the processes that give the machine these responsibilities.

2.9.2 AFS Server Processes

All AFS servers run the Basic OverSeer (BOS) Server process, bosserver. This process is responsible for starting and monitoring other AFS server processes on that machine. It knows which processes to start based on configuration files the server administrator creates on the local machine.

A machine functioning as an AFS File Server runs the fs *trio of processes*:

- The *fileserver* process runs continuously during normal operation to handle requests at the file and directory level. This is the process you interact with when you access data from AFS file space.
- The *volserver* process also runs continuously during normal operation and handles operations at the volume level. This is the process you interact with when you create, delete, or move volumes.
- The *salvager* process checks the system for internal consistency, and repairs errors it finds. It is analogous to UNIX fsck, but unlike fsck, it understands AFS volume structures and only operates on AFS data. The salvager runs automatically at startup time if the bosserver determines that the server was not properly shut down. It also runs automatically if the fileserver process fails. A server administrator can invoke the salvager automatically via the `bos` command suite, via the `salvager` command, or via the AFS Control Center's Server Manager.

While these are three separate processes, the BOS Server treats them as a unit and understands dependencies between the processes.

A machine functioning as an AFS Database Server runs separate processes for each of the four AFS databases:

- The *kaserver* process is responsible for maintaining the Authentication Database. It also verifies user identity at login by requiring a password, grants the user a token as proof to AFS server processes that the user has authenticated, and provides a way for server and client processes to prove their identities to each other. You can use the `kas` command suite to interact with the kaserver process.
- The *ptserver* process is responsible for maintaining the Protection Database and providing information when the fileserver process needs to retrieve information about an AFS user or protection group. You can use the `pts` command suite to interact with the ptserver process.
- The *vlserver* process, not to be confused with the volserver process that runs on File Servers, is responsible for maintaining the VLDB and for providing the Cache Manager with information about volumes and volume location. You can use the `vos` command suite to interact with the vlserver process.
- The *buserver* process is responsible for maintaining the Backup Database and for providing an interface to the AFS Backup System. You can use the `backup` command suite to interact with the buserver process.

The Update Server processes are responsible for transferring information from SCMs and BDMs to other AFS servers. These processes include upserver and upclient. Although the name may suggest otherwise, these processes run on AFS servers, not AFS clients:

- The *upserver* process runs on SCMs and BDMs and is responsible for distributing the contents of a specified directory, such as /usr/afs/etc on an SCM or /usr/afs/bin on a BDM, to other servers that request them.
- The *upclient* process is responsible for checking a specified directory on a particular machine and retrieving any files that have changed.

In addition to these processes, AFS servers may also run the runntp process. This process synchronizes an AFS server's clock with the clock on another machine. It is important for the clocks on all AFS servers in a cell to be synchronized for AFS services to function properly. If there is another time-keeping mechanism already running on the server, it is not necessary to run this process.

2.9.3 Local Directories

AFS creates and relies upon information in a number of directories on a server's local disk:

- On UNIX systems, the /usr/afs/etc directory contains cell-wide server configuration files that must be the same on every AFS server in your cell. On Windows NT, the directory where those files are located by default is C:\WSPP\AFS\Server\usr\afs\etc. You should use commands from the `bos` command suite to update the files. The files include:
  - CellServDB

    The CellServDB file is an ASCII file that contains a list of the local cell's Database Servers.
  - UserList

    The UserList is an ASCII file that contains the names of system administrators in the local cell.
  - KeyFile

    The KeyFile contains server encryption keys. A *server encryption key* is a string of octal numbers that is used to encrypt/decrypt packets of information. This file contains one key that matches the key in the Authentication Database for the AFS entry.
  - ThisCell

    The ThisCell file is an ASCII file that specifies the cell to which the server machine belongs.
- On the UNIX platform, the /usr/afs/local directory contains machine-specific configuration files. On Windows NT, these files are by default located in the C:\WSPP\AFS\Server\usr\afs\local directory. These files differ from server machine to server machine. You should use commands from the `bos` command suite to update the files. The files include:
  - BosConfig

    The BosConfig file contains information about AFS processes that the bosserver is responsible for. It defines which processes should be running on the server and the state they should be in. This file also stores information that specifies when the bosserver checks for new binaries in its /usr/afs/bin directory and when it restarts all of the machine's AFS server processes.

  - SALVAGE.fs

    The SALVAGE.fs file controls the behavior of the Salvager. If this file exists when the bosserver starts, the bosserver will run the Salvager before starting other File Server processes. This file is created when the bosserver starts the fileserver process and is removed when the bosserver properly shuts down the fileserver process.

  - salvage.lock

    The salvage.lock file is a zero-length file that exists when the Salvager is running. The purpose of this file is to allow only one server administrator to run the Salvager on the File Server at a time.

  - NoAuth

    The NoAuth file is a file that should not be present under normal circumstances. If present, AFS server processes are not required to check authorization.

  - sysid

    The sysid file contains up to the maximum number of IP addresses supported for each server machine; this number in AFS 3.4a and 3.5 is 15. The sysid file also contains a unique host identifier which distinguishes the server from all others in the cell. The file is automatically created at startup time. Since this file contains server-specific information, server administrators should not copy this file from one File Server to another.

- On UNIX systems, the /usr/afs/bin directory contains AFS server and command suite binaries. By default, on Windows NT these binaries are located in the C:\WSPP\AFS\Server\usr\afs\bin directory. Since the binaries will be the same on every AFS server of the same operating system version, you can maintain master copies on the BDM and use upclient/upserver processes to distribute the binaries to other AFS servers with the same operating system version.

- If a UNIX machine is functioning as an AFS Database Server, the /usr/afs/db directory contains the four AFS database files, which are:

  - kaserver.DB0, the Authentication Database
  - prdb.DB0, the Protection Database
  - vldb.DB0, the Volume Location Database
  - bdb.DB0, the Backup Database

  On Windows NT, the directory where the files above are located by default is C:\WSPP\AFS\Server\usr\afs\db.

- The /usr/afs/logs directory on UNIX systems contains log files associated with various AFS server processes. It may also contain old versions of the log files if they exist and core files generated by the AFS server processes. The log files found in this directory include:

  - BosLog, generated by the bosserver process
  - FileLog, generated by the fileserver process
  - VolserLog, generated by the volserver process
  - SalvageLog, generated by the salvager
  - AuthLog, generated by the kaserver process
  - VLLog, generated by the vlserver process
  - BackupLog, generated by the buserver process

On WIndows NT, these log files are by default located in the C:\WSPP\AFS\usr\afs\logs directory.

2.10 AFS Client Concepts

AFS client machines are used by server administrators and users to access information from AFS servers and issue commands to interact with AFS servers (for example, create protection groups or manage server processes).

2.10.1 Client Responsibilities

Recall that AFS clients run the Cache Manager, a set of kernel modifications responsible for communicating with processes on AFS servers and accessing AFS files and directories. When a user requests an AFS file, the Cache Manager interprets the path to extract the volume name from the mount point, requests volume location information from the VLDB, fetches the file from the File Server, and caches the file locally.

2.10.2 Processes

The afsd process runs on AFS client machines. This process is started from the client machine's initialization file. The afsd process initializes the Cache Manager by transferring information into kernel memory and starting several AFS daemons.

When the afsd process starts, some of its responsibilities include:

- Setting a field in kernel memory that defines the cell to which the machine belongs
- Putting names and IP addresses of Database Servers from local and foreign cells into kernel memory
- Defining the directory under which the AFS directory tree resides, typically /afs
- Determining whether to use disk or memory cache
- Defining the name of the local directory to be used for caching (on UNIX, typically /usr/vice/cache)
- Setting cache size and cache parameters
- Randomly selecting a File Server in the local cell to get time from, unless the `-nosettime` switch is used with the afsd binary

2.10.3 Local Directories

On the UNIX platform, AFS relies upon information in two directories on a client machine's local disk:

- The files necessary to configure the client machine are by default located in the /usr/vice/etc directory. They include:
  - afsd binary file

    As described in 2.10.2, "Processes" on page 44, the afsd binary initializes the Cache Manager.
  - cacheinfo file

The cacheinfo file stores information about the cache. It specifies the directory where the Cache Manager mounts AFS, the name of the cache directory, and the cache size.

  - ThisCell file

    The ThisCell file contains the name of the client machine's local cell.

  - CellServDB file

    The CellServDB file contains the names and IP addresses of Database Servers in the local cell and in foreign cells.

  - dkload (AIX) and modload (Solaris) subdirectory

    The dkload directory on AIX contains AFS libraries that need to be loaded into the client machine's kernel, system-specific binary for dynamic kernel loading, and files to help you call afsd with arguments appropriate for the client machine's cache size, number of users, and RAM.

    The modload directory on Solaris contains files necessary to build a kernel that incorporates AFS modifications if they exist for the specific operating system.

- The /usr/vice/cache directory contains files that the Cache Manager creates or maintains for its own use. The contents include:

  - V- files

    V- files are used to store chunks of data in the cache. The file names are in the form V*digits*, such as V3367. The number of V- files in the directory depends upon the size of the cache.

  - CacheItems

    The CacheItems file is a cache overhead file that stores information about each V- file in the cache.

  - VolumeItems

    The VolumeItems is another cache overhead file that records information about volumes that the Cache Manager has accessed. This includes volume-to-mount point relationships and volume-location information.

On Windows NT, the Cache Manager is a process, obtained by running the afsd_service.exe executable file, by default located in the C:\WSPP\AFS\Client\Program directory. The files necessary to configure the Cache Manager on Windows NT are located in the operating system directory, which is by default C:\WINNT.

2.10.4 The @sys Variable

AFS uses the @sys variable to simplify administration of operating system-specific files. The @sys variable expands to the current system's CPU and operating system type, for example: sun4x_56 for Solaris 2.6, rs_aix42 for AIX 4.2.1, and i386_nt40 for Windows NT 4.0. When the Cache Manager encounters the @sys variable in a path, @sys is replaced with the machine's system type. The following figure offers a graphical representation of the usage of the @sys variable on the Solaris and AIX platforms:

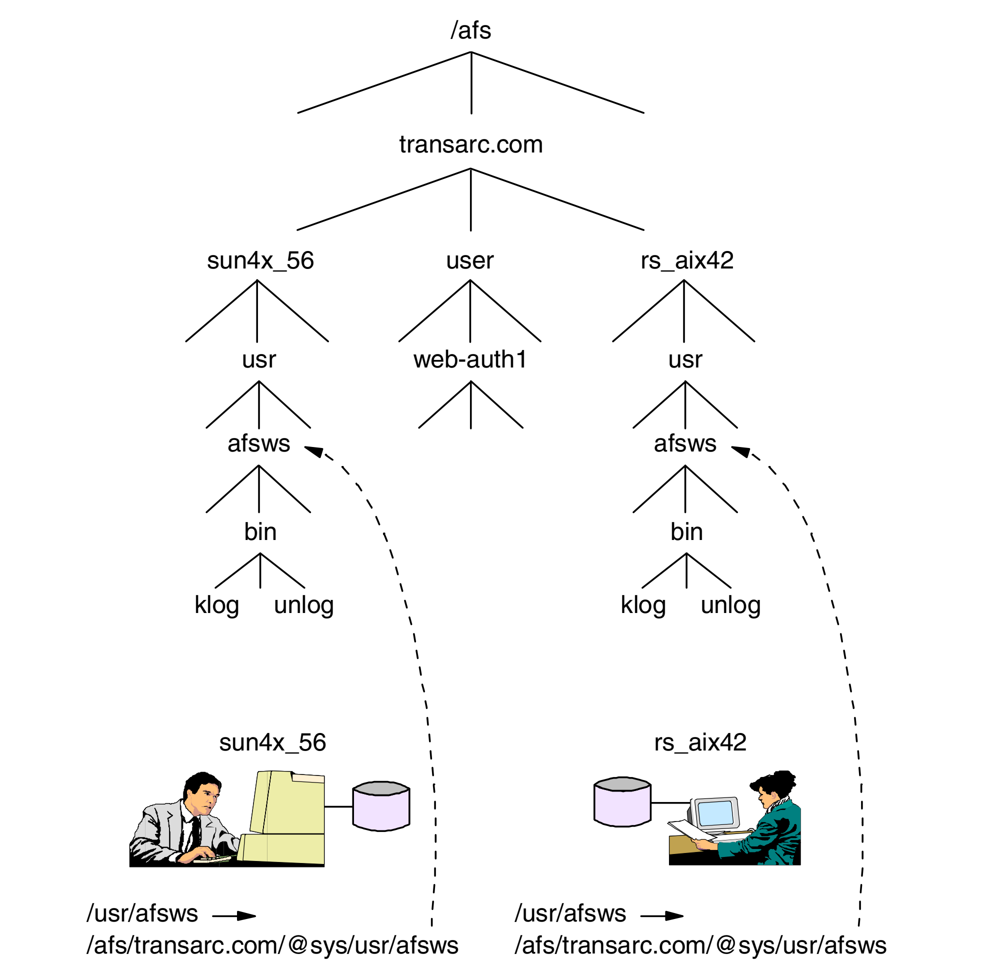

Figure 13. @sys Variable Usage

2.11 AFS Volume Concepts

The AFS directory tree is broken down into smaller subtrees. Each subtree is a collection of related files and directories called a *volume*. AFS volumes provide a unit of administration for replicating frequently accessed files and directories, moving data from one physical location to another.

2.11.1 Volume Headers

Volume headers contain pointers to where the volume's files and directories are located on disk. Although an AFS partition is formatted as the vendor-supplied file system, data is stored in a different manner. If a server administrator logs in to an AFS File Server and lists the contents of a vice partition, the contents would include volume headers, as shown in the following diagram:

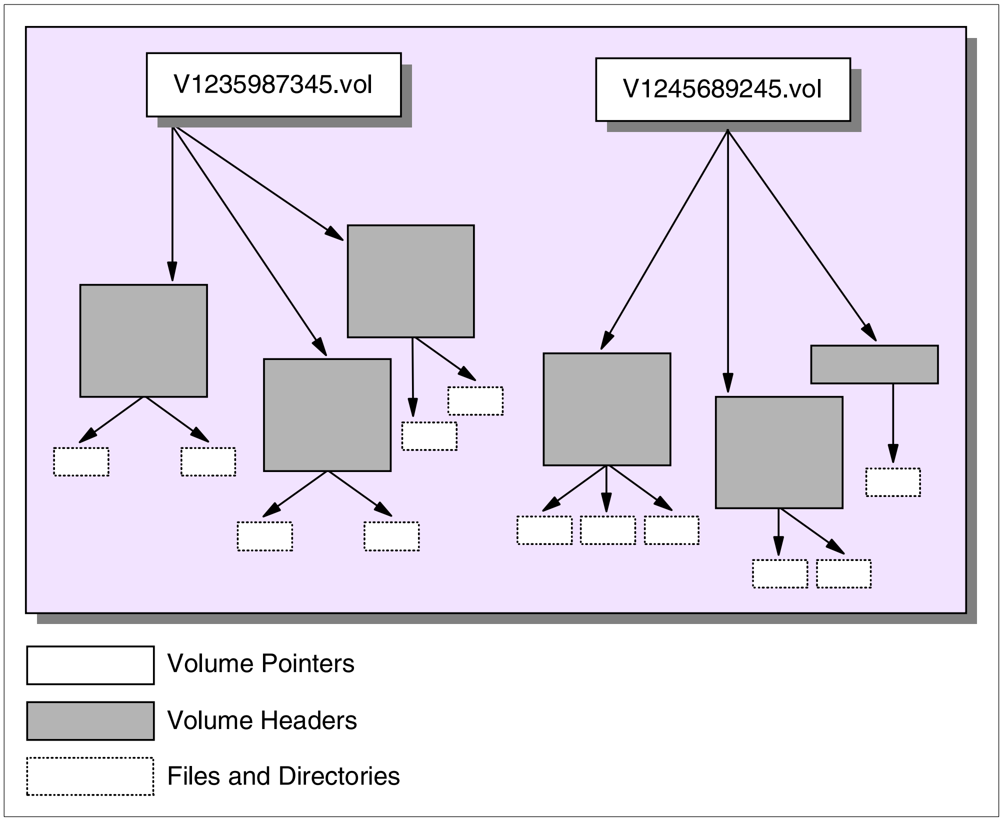

Figure 14. Relationship between Volume Pointers, Headers, Files and Directories

2.11.2 Types of Volumes

AFS has three types of volumes:

- A *ReadWrite* volume contains modifiable copies of AFS files and directories. Each volume has a ReadWrite version.
- A *ReadOnly* volume contains a non-modifiable copy, *replica*, of the contents of a ReadWrite volume. Each ReadWrite volume can have up to eleven ReadOnly replicas. The name of the ReadOnly volume is the name of the ReadWrite volume with a .readonly extension (for example, root.afs.readonly).
- A *Backup* volume contains a non-modifiable snapshot of the contents of the ReadWrite volume. Each ReadWrite volume can have at most one Backup version. The name of the Backup volume is the name of the ReadWrite volume with a .backup extension (for example, root.afs.backup).

2.11.3 Volume Traversal Rules

Recall that the VLDB contains volume location information. A Cache Manager contacts the VLDB to learn information about a particular volume. The VLDB tells the Cache Manager where the volume and its replicas reside. If replicated, server preferences determine which server to contact for a replica. The decision whether to access the ReadWrite version or a ReadOnly replica is based on *traversal rules*.

Traversal rules define how a client chooses between the ReadWrite version and ReadOnly replica if the volume is replicated. The rules are as follows:

- If root.afs is replicated, use a ReadOnly volume.
- If root.afs is not replicated, use the ReadWrite volume.
- As the path is traversed
  - If you are currently in a ReadOnly volume and the next volume is
    - Replicated, you must use a ReadOnly volume
    - Not replicated, you must use the ReadWrite volume
  - If you are currently in a ReadWrite volume, you must access a ReadWrite volume

Traversal rules assume that you did not force a mount point to contain the name of a ReadOnly volume.

The following figure offers a graphical representation of traversal rules:

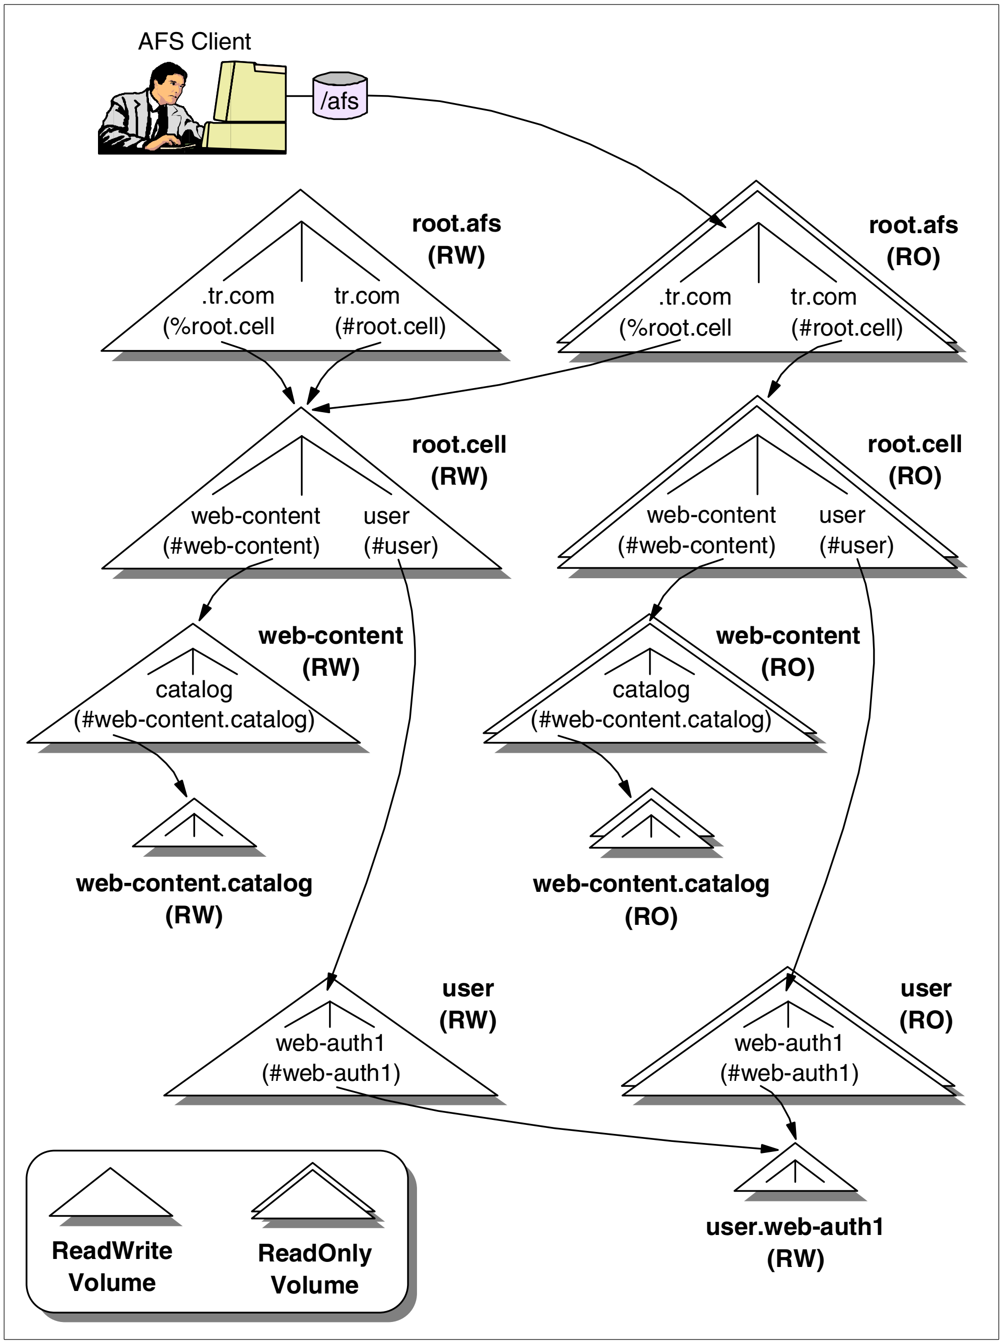

Figure 15. How Traversal Rules Work

Notice in the root.afs volume, there is a ReadWrite mount point, .tr.com, that contains the name of the root.cell volume. We know this is a ReadWrite mount point because it contains a percent symbol (%) instead of a hash symbol (#).

The Cache Manager does not support the previously stated traversal rules when it crosses or interprets a ReadWrite mount point. Instead, the fact that a mount point is ReadWrite indicates that the Cache Manager should access a ReadWrite version of the specified volume.

By convention, if the Regular mount point that contains the name of the root.cell volume is called *cellname*, and the ReadWrite mount point that contains the version of the root.cell volume is called *.cellname*.

Since it is best to let the Cache Manager follow traversal rules, and ReadWrite mount points do not follow traversal rules, there should only be one ReadWrite mount point in your directory tree. In fact, when the Cache Manager is on a ReadWrite volume, it always accesses the ReadWrite version of the next volume. Therefore, you only need one ReadWrite mount point high up in your namespace to put the Cache Manager on the ReadWrite path, and guarantee that the Cache Manager can access the ReadWrite version of all volumes.

When a Web author wants to access a file in the AFS directory tree, he specifies the path. Suppose he wanted to access /afs/tr.com/web-content/filex. This would take the Web author to the file in the ReadOnly version of the web-content volume. If the Web author wanted to access a modifiable version of the file, he must specify the path /afs/.tr.com/web-content/filex because it leads to the ReadWrite version of the web-content volume.

When a server administrator is setting up the AFS directory structure, it is important to create the ReadWrite mount point to the root.cell volume before replicating the root.afs volume. If the server administrator forgets to do this, it may not be possible to access the ReadWrite volumes in the directory tree and create mount points to attach volumes into the directory tree.

In particular, if a server administrator forgets to replicate the root.afs volume, replicated copies of all volumes mounted beneath root.afs will never be accessed. Suppose a server administrator forgot to replicate the root.afs volume but did replicate the root.cell volume in the tr.com cell. The replicas of the root.cell volume would not be used. There could be up to 11 replicas of the root.cell volume that are taking up space on the File Servers but are not being accessed. A user would specify path /afs/tr.com to access the root.cell volume. In this case, the Cache Manager would use traversal rules to determine which version of the root.afs volume to access. Since root.afs is not replicated, the Cache Manager must access the ReadWrite version. The Cache Manager extracts the name of the root.cell volume from the tr.com mount point and contacts the VLDB to determine where the volume resides. Even though there are replicas of this volume, the Cache Manager uses traversal rules. The rules indicate that if you are in a ReadWrite volume, you must access a ReadWrite volume.

A similar situation occurs if you forget to replicate any volume in the directory tree. If you forget to replicate a volume, replicas of all volumes mounted beneath that volume will never be accessed.

2.11.4 Naming Suggestions

When you create a volume, you create its ReadWrite version. The name you select should be no more than 22 characters. The limit is really 31 characters, but includes the .readonly or .backup extensions. When you create a ReadOnly replica or a Backup volume, the .readonly and .backup extensions are automatically appended to the volume name.

The name you select for a volume should reflect the volume contents. Server administrators should be able to examine volume names from the VLDB or from volume headers and have an idea of what the volume contains.

The volume name should also reflect where the volume is attached in the directory tree. This is useful if you have to delete a volume. When you delete a volume, you should also delete the mount point that attaches the volume to the directory tree. If your volume name and mount point do not correspond, you will have to find another way to determine where the volume was attached to the directory tree so you can delete its mount point.

2.12 How AFS Works

The name space concept starts with a typical desktop client computer. All AFS client machines run a Cache Manager process. The Cache Manager maintains information about the identities of the users logged into the machine, finds and requests data on their behalf, and keeps retrieved files on local disk. The local caching reduces network traffic and server load.

The central AFS file name space is key. While all the file and directory names in the name space are related to physical data on the servers, the names are visible in a single, enterprise-wide, and consistent directory structure. The servers perform all their work in concert to provide a seamless file name space visible by all desktops and all users. Any changes to the name space, any new files or any changed data are publicized to clients immediately.

Unlike past systems, which made use of the /etc/ file system on a client to map between a local directory name and a remote file system through a mounting operation, AFS does its mapping (file name to location) at the server. This process has the advantage of making the served file space location independent. Location independence means that a user does not need to know which File Server holds the file. The user only needs to know the path name of a file.

AFS, the File Sharing component of IBM WebSphere Performance Pack, has important applications in HTTP environments.

Typically, an HTTP daemon sends to Web browsers files available from the machine on which it executes, but the process does not care whether the files that it serves are stored on a local disk or on an AFS file server, since access to AFS files is indistinguishable from access to local files on an AFS client machine. Many HTTP servers will run with files on the local server, so when you see a path, it is mapping to a file that is stored locally. From the HTTP server's perspective, there is nothing different from pulling a file out of the AFS directory structure. The Cache Manager and all the other elements that go into providing this unified name space make it possible for the HTTP Server to make use of the AFS client to access a file from the AFS directory tree.

Chapter 3. AFS Installation and Configuration on AIX

This chapter summarizes the installation prerequisites for the AFS component of IBM WebSphere Performance Pack on AIX systems, and describes the installation and configuration procedures necessary to set up the first AFS machine, an additional AFS File Server, and an AFS Client:

- The first AFS machine in the cell will be designated as a Database Server, a File Server, a System Control Machine (SCM), a Binary Distribution Machine (BDM), and a client.
- An AFS File Server is also referred to as a File Sharing Server in WebSphere Performance Pack. Although there is no technical requirement for a machine to be configured as both AFS File Server and AFS Client, the AFS configuration script of WebSphere Performance Pack configures both server and client services for an AFS File Server.
- An AFS Client is also referred to as a File Sharing Client in Performance Pack. For an AFS Client, the AFS configuration script of WebSphere Performance Pack configures only AFS Client processes.

To demonstrate the installation and configuration of AFS on AIX systems, we define an AFS cell called sushi.itso.ral.ibm.com and use the following IBM RS/6000 machines:

Table 2. List of Machines in the AFS Cell, sushi.itso.ral.ibm.com

Machine Role	Host Name	Operating System
First AFS Machine	rs60002.itso.ral.ibm.com	AIX 4.3.2
AFS File Server	rs600030.itso.ral.ibm.com	AIX 4.3.0
AFS Client	rs600023.itso.ral.ibm.com	AIX 4.3.1
AFS Client	rs600022.itso.ral.ibm.com	AIX 4.3.1

3.1 AIX System Requirements

To prepare your AIX system for the installation of the AFS component of WebSphere Performance Pack, you need:

- Any IBM RISC System/6000 or IBM Power Series Family machine with one or more processors
- AIX Version 4.2, 4.2.1, 4.3, 4.3.1, or 4.3.2
- A minimum of 128 MB of RAM
- A minimum of 256 MB of swap space

  For performance reasons, we recommend that you allocate twice the amount of RAM for swap space.
- An operational machine on the network
- A minimum of 20 MB of available disk space on the /usr file system to install the AFS files
- A minimum of 40 MB of available disk space for the AFS disk cache

Although AFS memory caching is available, memory caching does not provide better performance than disk caching. We recommend that you use disk caching for an AFS client.

- At least one AIX journaled file system (JFS) dedicated to storing AFS data for each AFS File Server.
- Java Runtime Environment (JRE) 1.1.6
- Administrative privileges

  You need to have root access to the AIX machine in order to install, configure, and start AFS server and client processes.

3.1.1 Paging Space Allocation

You need to allocate adequate paging space on your AFS machines. On AIX systems, the command:

```
bootinfo -r
```

displays the amount of real memory in KB, and the command:

```
lsps -a
```

displays the amount of paging space. For performance reasons, the amount of paging space must be at least twice the amount of memory on your machine.

The following session demonstrates the usage of these two commands:

```
# bootinfo -r
131072

# lsps -a
Page Space  Physical Volume   Volume Group    Size   %Used  Active  Auto  Type
hd6         hdisk0            rootvg         256MB      33     yes   yes    lv
```

Figure 16. Displaying the Amount of Memory and Paging Space on an AIX System

To change the size of the paging space, use the `chps` command.

3.1.2 Name Resolution Service

We recommend that you use a name service in your environment, such as the Domain Name System (DNS), to resolve host names of machines. It is difficult to maintain accurately host names of machines with local host tables.

On AIX systems, the `host` command resolves a host name into an Internet address or an Internet address into a host name. By using the `host` command, you can verify that both the forward and reverse name resolution service is functioning correctly.

The following session illustrates the usage of the `host` command for a machine with the host name of rs600024, the IP address of 9.24.104.161, and the DNS domain name of itso.ral.ibm.com.

```
# hostname
rs600024

# host rs600024
rs600024.itso.ral.ibm.com is 9.24.104.161

# host 9.24.104.161
rs600024.itso.ral.ibm.com is 9.24.104.161
```

Figure 17. Verifying Forward and Reverse Host Name Resolution on an AIX System

3.1.2.1 Network Connection

The AIX machine has to be operational on the network. For each machine where you are planning to install and configure AFS, verify network configuration and integrity.

On UNIX systems, such as IBM AIX and Sun Solaris, the following commands can assist you in determining your network status:

- To show the network routing table:

  `netstat -nr`

- To send a network request to a host or gateway:

  `ping` *`host`*

- To display the statistics of the network interface:

  `netstat -ni`

An example is shown in the following figure:

```
# netstat -nr
Routing tables
Destination      Gateway           Flags   Refs     Use  If   PMTU  Exp  Groups

Route Tree for Protocol Family 2 (Internet):
default          9.24.104.1        UG        9    68765  tr0     -   -
9.24.104/24      9.24.104.161      U        19    51050  tr0     -   -
127/8            127.0.0.1         U         2      175  lo0     -   -

Route Tree for Protocol Family 24 (Internet v6):
::1              ::1               UH        0        0  lo0 16896   -

# ping -c 5 -s 1024 9.24.104.1
PING 9.24.104.1: (9.24.104.1): 1024 data bytes
1032 bytes from 9.24.104.1: icmp_seq=0 ttl=64 time=4 ms
1032 bytes from 9.24.104.1: icmp_seq=1 ttl=64 time=3 ms
1032 bytes from 9.24.104.1: icmp_seq=2 ttl=64 time=3 ms
1032 bytes from 9.24.104.1: icmp_seq=3 ttl=64 time=3 ms
1032 bytes from 9.24.104.1: icmp_seq=4 ttl=64 time=3 ms

----9.24.104.1 PING Statistics----
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max = 3/3/4 ms

# netstat -ni
Name  Mtu   Network     Address            Ipkts Ierrs    Opkts Oerrs  Coll
lo0   16896 link#1                          37800     0    37822     0     0
lo0   16896 127         127.0.0.1           37800     0    37822     0     0
lo0   16896 ::1                             37800     0    37822     0     0
tr0   1492  link#2      0.4.ac.96.93.87    502988     0    82637     0     0
tr0   1492  9.24.104    9.24.104.161       502988     0    82637     0     0
```

Figure 18. Verifying Network Connection on an AIX System

Figure 18 on page 56 demonstrates the usage of these commands from an AIX machine with the IP address 9.24.104.161:

- The command:

  `netstat -nr`

  shows the network routing table, verifying that the machine IP address is 9.24.104.161 and the default gateway IP address is 9.24.104.1.

- The command:

  `ping 9.24.104.1`

  verifies the network connection between the local machine, 9.24.104.161, and the default gateway, 9.24.104.1, by sending an Internet Control Message Protocol (ICMP) echo request to the default gateway. The PING statistics confirm network connectivity between the two machines.

- The command:

  `netstat -ni`

  displays the statistics of the network interface, indicating that there are no packet errors or collisions.

3.1.3 Time Synchronization

Before configuring additional AFS server or client machines, manually verify that the clock on the new machine is relatively synchronized with the clock on the first AFS machine. If the clock is off on the new machine, correct it before continuing. After you configure additional AFS server or client machines, AFS uses the Network Time Protocol (NTP) to synchronize time amongst a cell’s server machines in the AFS cell. An AFS client machine synchronizes its clock with the clock on a server machine chosen at random.

3.1.4 JRE Version

The WebSphere Performance Pack installation process uses the Java InstallShield’s setup class. You need to ensure that the Java Runtime Environment (JRE) V1.1 (or higher) is installed. At the time of this writing, the recommended JRE version for AIX is 1.1.6.

To quickly determine the Java version, enter:

```
java -version
```

To ensure that the JRE has been correctly installed on your AIX system, enter:

```
lslpp -l "Java.rte.*"
```

The following session displays the state of the JRE on an AIX machine:

```
# lslpp -l "Java.rte.*"
  Fileset                      Level  State      Description
  ----------------------------------------------------------------------------
Path: /usr/lib/objrepos
  Java.rte.bin               1.1.6.3  COMMITTED  Java Runtime Environment
                                                 Executables
  Java.rte.classes           1.1.6.3  COMMITTED  Java Runtime Environment
                                                 Classes
  Java.rte.lib               1.1.6.3  COMMITTED  Java Runtime Environment
                                                 Libraries
```

If the output of the `lslpp` command does not indicate level 1.1.6 or higher, or the State of this software product is not `COMMITTED`, you will have to install JRE V1.1.6 on this machine. The CD distribution of WebSphere Performance Pack contains the latest Java Development Kit (JDK) in the directory aix/jdk. Actually, you would not need the full JDK, but only its subset known as JRE, which contains just the Java Virtual Machine (JVM), the Java platform core classes, and supporting files. In other words, the JRE is the smallest set of executables and files that constitute the standard Java platform and it contains only the run-time part of the JDK: no compiler, no debugger, no tools.

To install the JDK on AIX, refer to the IBM Redbook, *Network Computing Framework Component Guide*, SG24-2119.

3.1.5 Disk Space Requirements on UNIX Systems

Since the disk requirements are very similar for AIX and Solaris, in this section we unify the discussion for both these platforms. In 4.1.7, “Disk Space Requirements on Solaris Systems” on page 101, we will simply point to this section for more details.

WebSphere Performance Pack installs the AFS component in the following directories: /usr/lpp/WSPP (AIX) and /opt/WSPP (Solaris), /usr/afs, and /usr/vice. You must allocate adequate disk space for these directories:

- /usr/lpp/WSPP (AIX) or /opt/WSPP (Solaris)

  This directory contains WebSphere Performance Pack setup files, configuration scripts, and installation logs. It only requires less than 1 MB of disk space.

- /usr/afs

  This directory contains the AFS server binaries; system administrator command and user command binaries; database, configuration, and log files. You will need approximately 20 MB for this directory. Depending on the size of the AFS databases and the growth of the AFS log files, you may need to allocate more disk space.

- /usr/vice

  This directory contains the AFS client binaries, AFS modules, and the default AFS disk cache directory, /usr/vice/cache. The size of an AFS disk cache depends on the number of users on this machine, the size of the files and the number of files that will be accessed. On an AIX system, an AFS disk cache size should not exceed 85% of the size of its disk partition as reported by the `df` command.

  AFS Disk Cache File System

  The AFS disk cache directory can be configured from any valid AIX Journaled File System (JFS). It is desirable to have a dedicated file system for the AFS disk cache directory to guarantee that no other application will be competing for that disk space.

  We select a local disk directory on the existing /usr file system for the AFS cache directory, /usr/vice/cache, and the initial cache size of 40 MB.

Since all the above mentioned directories reside in /usr, we recommend that you have approximately 80 MB of free disk space in the /usr file system.

On AIX systems, if you need to increase the size of the /usr file system, enter the following command:

```
smitty chjfs
```

Select **/usr** from the list of defined file systems on your machine. Press Enter to continue to the Change / Show Characteristics of a Journaled File System window as shown in the following session:

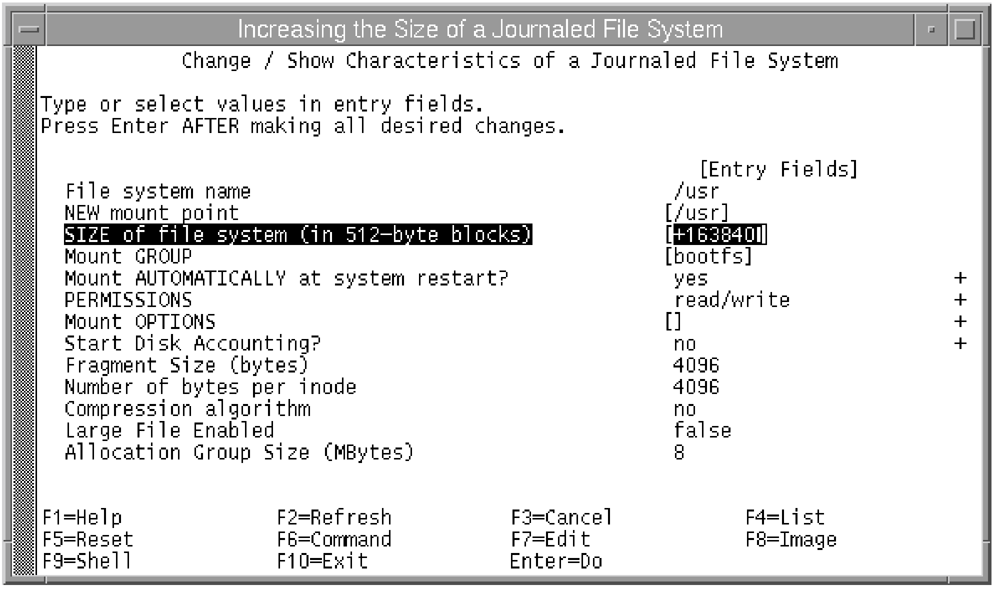

```
              Change / Show Characteristics of a Journaled File System

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                                          [Entry Fields]
  File system name                                      /usr
  NEW mount point                                      [/usr]
  SIZE of file system (in 512-byte blocks)             [+163840]
  Mount GROUP                                          [bootfs]
  Mount AUTOMATICALLY at system restart?                yes                    +
  PERMISSIONS                                           read/write             +
  Mount OPTIONS                                        []                      +
  Start Disk Accounting?                                no                     +
  Fragment Size (bytes)                                 4096
  Number of bytes per inode                             4096
  Compression algorithm                                 no
  Large File Enabled                                    false
  Allocation Group Size (MBytes)                        8

F1=Help             F2=Refresh          F3=Cancel           F4=List
F5=Reset            F6=Command          F7=Edit             F8=Image
F9=Shell            F10=Exit            Enter=Do
```

Figure 19. Increasing the Size of the /usr File System

In the `SIZE of file system (in 512-byte blocks)` field, enter the desired size of the /usr file system. The value can be an absolute value or an increment specified by the plus sign (+). In this example, 80 MB is equivalent to 163,840 512-byte blocks as illustrated in Figure 19 on page 59. Press Enter to start increasing the size of the /usr file system. Upon successful completion, press F10 to exit.

3.2 Configuring AFS Partitions on AIX Systems

Every AFS File Server must have a disk partition dedicated to be an AFS partition to store AFS volumes. By convention, the first AFS partition created on an AFS file server machine is called /vicepa, the second partition /vicepb, etc. On an AIX system, an AFS partition has to be created on a standard or compressed JFS.

AFS Partitions and Large File Enabled JFSs

On AIX systems, the AFS File Server processes do not recognize AFS partitions implemented on large file enabled JFSs. If you attempt to use a large file enabled JFS as an AFS partition, you will receive the following message in the /usr/afs/logs/FileLog file:

```
date /vicepx is a big files filesystem, ignoring it
```

To create a JFS to be used as an AFS partition, you need to perform the following operations:

- Create an AIX logical volume
- Add a standard or compressed JFS to this AIX logical volume
- Mount this file system

In the example shown in the following session screen, the AIX machine has five file systems:

```
# df -k
Filesystem    1024-blocks      Free %Used    Iused %Iused Mounted on
/dev/hd4            65536     40428   39%     2475     8% /
/dev/hd2           512000     79216   85%    20412    16% /usr
/dev/hd9var         16384     14136   14%      485    12% /var
/dev/hd3            16384     15744    4%       38     1% /tmp
/dev/hd1            16384     14556   12%      134     4% /home
```

The following session demonstrates the steps to create a 64 MB AFS partition:

1. To create an AIX logical volume, enter:

   `smitty mklv`

   The Add a Logical Volume window appears as shown in Figure 20 on page 60.

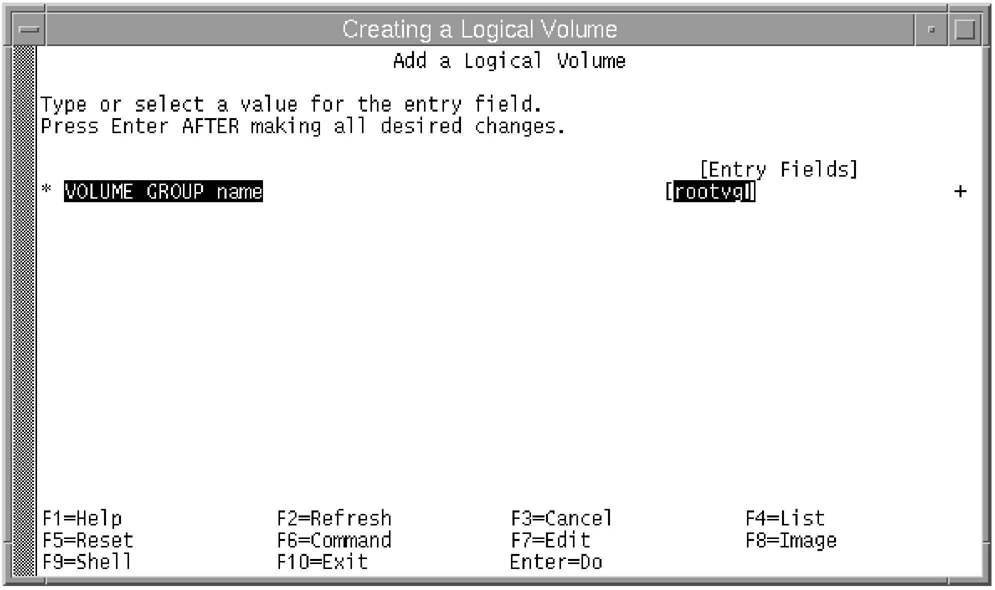

```
                          Add a Logical Volume

Type or select a value for the entry field.
Press Enter AFTER making all desired changes.

                                                   [Entry Fields]
* VOLUME GROUP name                               [rootvg]          +
```

Creating a Logical Volume
F1=Help F2=Refresh F3=Cancel F4=List
F5=Reset F6=Command F7=Edit F8=Image
F9=Shell F10=Exit Enter=Do

Figure 20. Selecting the Volume Group

Press F4 to select the volume group in which a new logical volume will be created. Press Enter to continue.

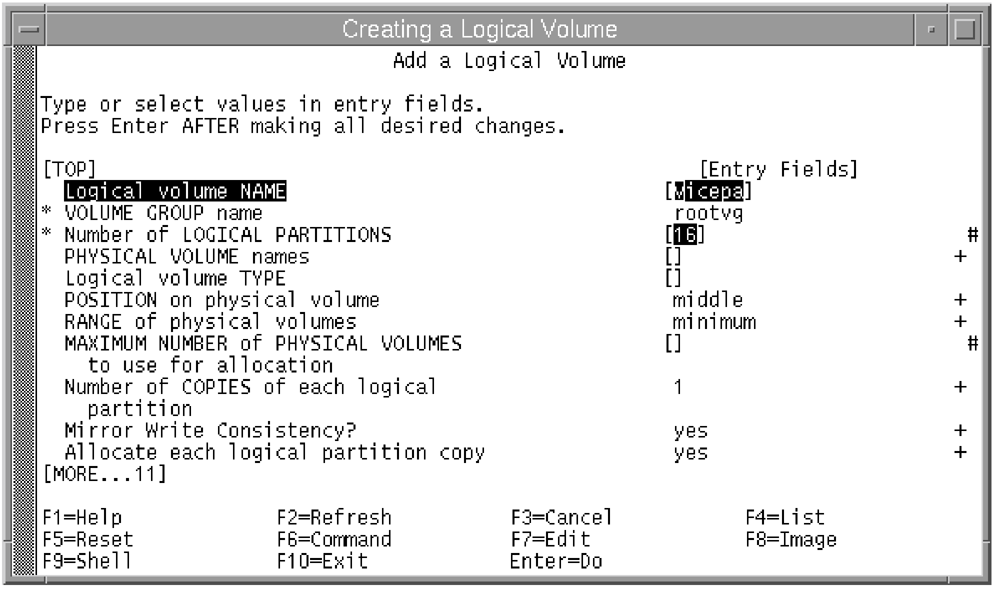

```
                              Add a Logical Volume

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

[TOP]                                                   [Entry Fields]
  Logical volume NAME                                  [vicepa]
* VOLUME GROUP name                                     rootvg
* Number of LOGICAL PARTITIONS                         [16]                    #
  PHYSICAL VOLUME names                                []                     +
  Logical volume TYPE                                  []
  POSITION on physical volume                           middle                +
  RANGE of physical volumes                             minimum               +
  MAXIMUM NUMBER of PHYSICAL VOLUMES                   []                      #
    to use for allocation
  Number of COPIES of each logical                      1                     +
    partition
  Mirror Write Consistency?                             yes                   +
  Allocate each logical partition copy                  yes                   +
[MORE...11]

F1=Help             F2=Refresh          F3=Cancel           F4=List
F5=Reset            F6=Command          F7=Edit             F8=Image
F9=Shell            F10=Exit            Enter=Do
```

Figure 21. Creating an AIX Logical Volume

From the Add a Logical Volume window (see Figure 21 on page 61), perform the following steps:

- In the `Logical volume NAME` field, type the AIX logical volume name.

  In this example, vicepa is the logical volume name.

- In the `Number of LOGICAL PARTITIONS` field, type the number of AIX logical partitions.

  In this example, we allocate 16 AIX logical partitions (4 MB per partition) for this AIX logical volume.

- Press Enter to start creating a new AIX logical volume.

2. To add a JFS to this AIX logical volume, enter:

   ```
   smitty crjfslvstd
   ```

   From the Add a Standard Journaled File System window (see Figure 22 on page 62), perform the following steps:

   - Press F4 in the `LOGICAL VOLUME name` field to select the name of the logical volume, **vicepa**
   - Type `/vicepa` in the `MOUNT POINT` field
   - Press F4 to select **yes** in the `Mount AUTOMATICALLY at system restart?` field
   - Press Enter to start adding a new AIX JFS

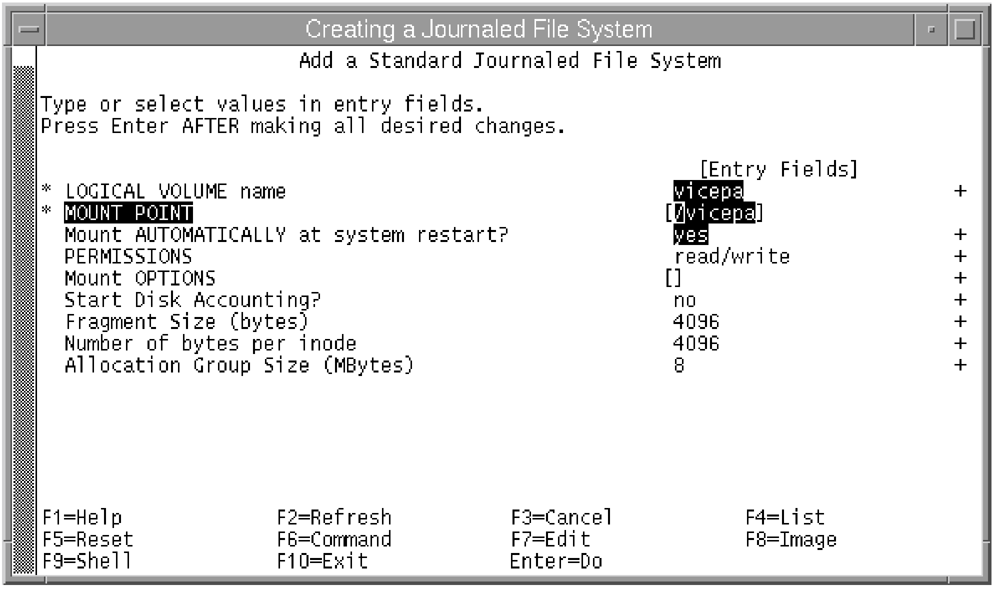

```
                    Creating a Journaled File System
                  Add a Standard Journaled File System

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                                      [Entry Fields]
* LOGICAL VOLUME name                                 vicepa              +
* MOUNT POINT                                        [/vicepa]
  Mount AUTOMATICALLY at system restart?              yes                 +
  PERMISSIONS                                         read/write          +
  Mount OPTIONS                                      []                   +
  Start Disk Accounting?                              no                  +
  Fragment Size (bytes)                               4096                +
  Number of bytes per inode                           4096                +
  Allocation Group Size (MBytes)                      8                   +




F1=Help              F2=Refresh           F3=Cancel            F4=List
F5=Reset             F6=Command           F7=Edit              F8=Image
F9=Shell             F10=Exit             Enter=Do
```

Figure 22. Creating a Journaled File System for an AFS Partition

3. To mount the previously created JFS, enter the command:

   ```
   mount /vicepa
   ```

   To confirm the available disk space on your file systems, enter the command:

   ```
   df -k
   ```

   The following session shows the new AFS partition and other file systems on an AIX machine:

```
# df -k
Filesystem    1024-blocks      Free %Used    Iused %Iused Mounted on
/dev/hd4            65536     40428   39%     2475     8% /
/dev/hd2           512000     79216   85%    20412    16% /usr
/dev/hd9var         16384     14136   14%      485    12% /var
/dev/hd3            16384     15744    4%       38     1% /tmp
/dev/hd1            16384     14556   12%      134     4% /home
/dev/vicepa         65536     63344    4%       32     1% /vicepa
```

AIX File System Helper

During the installation process, the setup program replaces the AIX file system helper program with an AFS-aware helper. But it does not correctly save the original AIX copy. You can manually save the original copy of the AIX helper program by entering the following command:

```
cp -p /sbin/helpers/v3fshelper /sbin/helpers/v3fshelper.orig
```

3.3 Installing AFS Files on AIX Systems

This section describes the procedure to install the necessary AFS files to configure an AIX machine to be an AFS machine, including the first AFS machine, an additional AFS File Server, and an AFS client.

To install IBM AFS Enterprise File System, the File Sharing component of WebSphere Performance Pack, on AIX systems, perform the following tasks:

1. Login as root.
2. Insert the IBM Websphere Performance Pack CD-ROM in the CD-ROM drive.
3. From a command prompt, enter the following commands:

```
mkdir /CD-ROM
mount -rv cdrfs /dev/cd0 /CD-ROM
cd /CD-ROM/aix
```

4. Start the Java InstallShield installation program by entering:

```
java setup
```

   The setup program displays the Welcome to the IBM WebSphere Performance Pack installation program window (see Figure 23 on page 63).

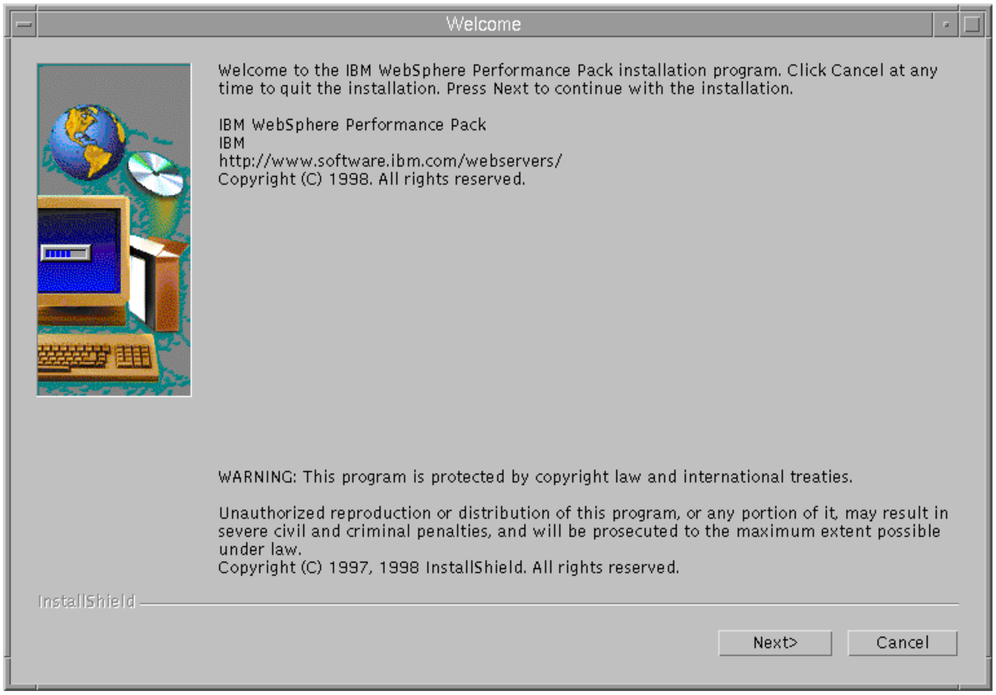

Welcome to the IBM WebSphere Performance Pack installation program. Click Cancel at any time to quit the installation. Press Next to continue with the installation.

IBM WebSphere Performance Pack
IBM
http://www.software.ibm.com/webservers/
Copyright (C) 1998. All rights reserved.

WARNING: This program is protected by copyright law and international treaties.

Unauthorized reproduction or distribution of this program, or any portion of it, may result in severe civil and criminal penalties, and will be prosecuted to the maximum extent possible under law.
Copyright (C) 1997, 1998 InstallShield. All rights reserved.

Figure 23. The IBM WebSphere Performance Pack Welcome Window

5. Click **Next** to continue. The Choose Destination Location window appears (see Figure 24 on page 64).

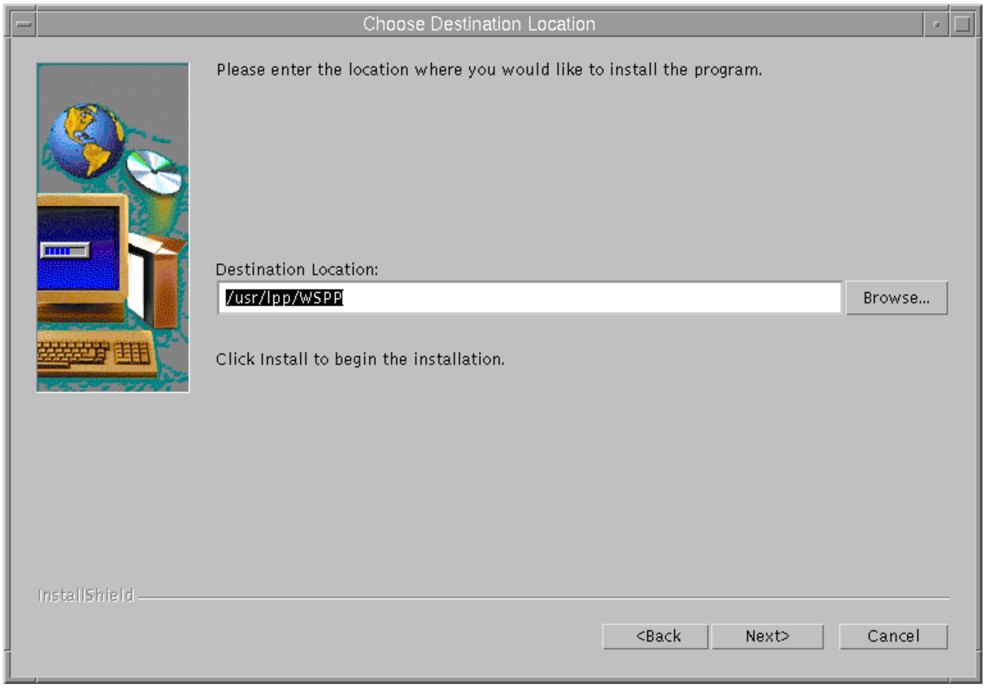

Choose Destination Location

Please enter the location where you would like to install the program.

Destination Location: /usr/lpp/WSPP

Click Install to begin the installation.

Figure 24. The Choose Destination Location Window

6. You need to specify the path to the location where you would like to install the WebSphere Performance Pack setup files, configuration scripts, and installation logs. You can click **Browse...** to select your own path. We chose the default path, /usr/lpp/WSPP. After making your selection, click **Next** to continue.
7. If the specified destination location does not exist, a dialog box appears to inquire whether you would like to have it created. If necessary, click **Yes** to create this directory.
8. You will be prompted to select the WebSphere Performance Pack components you want to install on your system, as shown in the following figure:

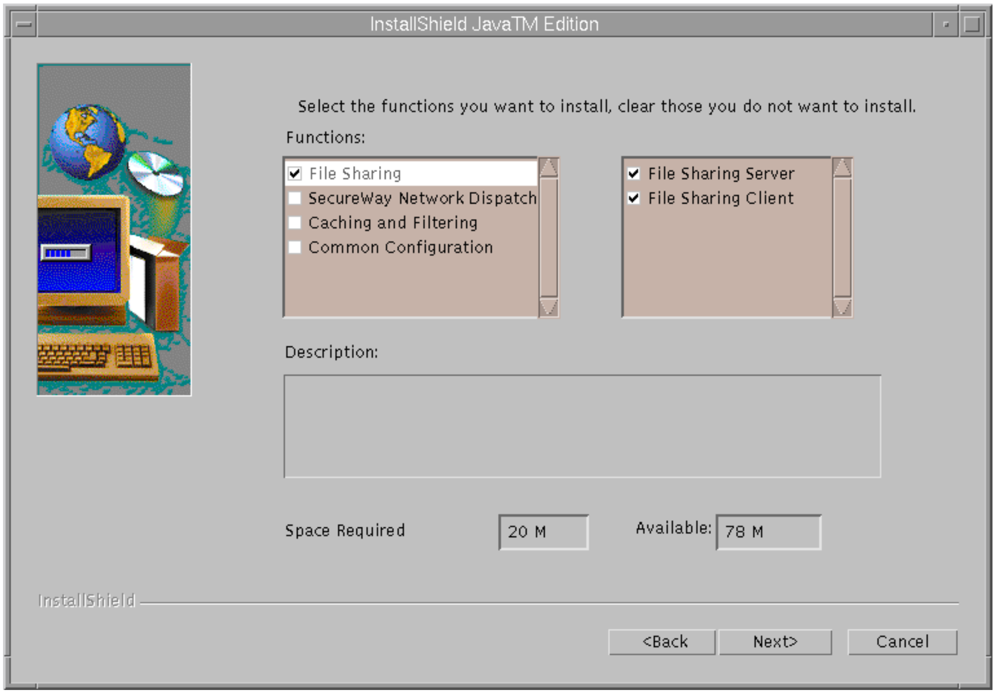

Select the functions you want to install, clear those you do not want to install.

Functions:

- ☑ File Sharing
- □ SecureWay Network Dispatch
- □ Caching and Filtering
- □ Common Configuration

- ☑ File Sharing Server
- ☑ File Sharing Client

Description:

Space Required: 20 M

Available: 78 M

Figure 25. The File Sharing Selection Window

Select the **File Sharing** check box and the setup program displays a list of File Sharing components in the window.

In WebSphere Performance Pack, the File Sharing Server, or AFS Server, must be installed with the File Sharing Client, or AFS Client. In this way, a File Sharing Server machine runs both AFS Server processes and AFS Client processes.

- If you are installing the first AFS machine or another AFS Server machine, you must select both the File Sharing Server and File Sharing Client. Select the **File Sharing Server** check box and the **File Sharing Client** check box is automatically set, as shown in Figure 25 on page 65.
- If you are installing an AFS client machine, you only have to select the **File Sharing Client** check box.

You need to choose one of the above selections and click **Next** to continue.

9. On the next pop-up window, you have an option to replace the current version of any WebSphere Performance Pack product that has already been installed.

   - If you are installing the first AFS machine and there are no other WebSphere Performance Pack products residing on this machine, check **Yes** to replace any copies of the File Sharing components eventually installed in the selected destination location.
   - If you are not sure, check **No** to keep the current copies.

   Then, click **Install** to start the installation of the chosen AFS components from the CD-ROM distribution to local disk.

   An Installing window appears showing the progress of the installation.

10. The Installation Complete window appears upon successful completion of the copying of files from the CD-ROM distribution, as shown in Figure 26 on page 66.

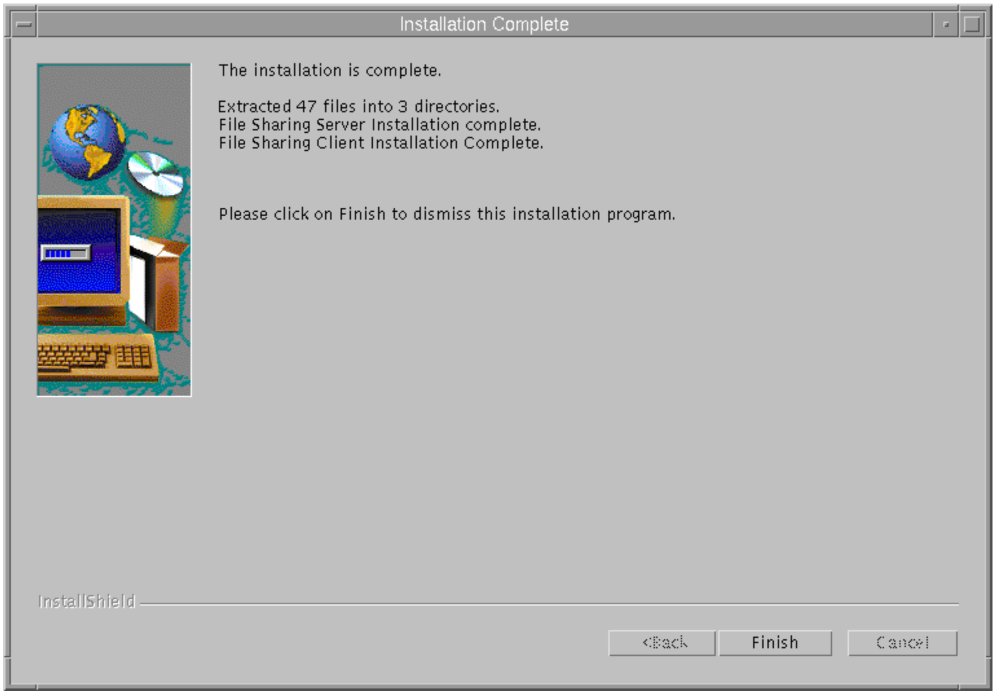

The installation is complete.

Extracted 47 files into 3 directories.
File Sharing Server Installation complete.
File Sharing Client Installation Complete.

Please click on Finish to dismiss this installation program.

Figure 26. The Installation Complete Window

Click **Finish** to complete the installation process.

3.3.1 WebSphere AFS Directories

The installation program creates the following directories:

- /usr/lpp/WSPP (AIX) or /opt/WSPP (Solaris)

  This directory contains WebSphere Performance Pack configuration files. The afs subdirectory has the AFS configuration shell script, afsConfigure.sh.

- /usr/afs/bin

  This directory contains AFS binaries, including server executables, system administrator and user commands.

- /usr/afs/etc

  This directory contains common server configuration files, such as a list of database server machines, the local cell name, the server encryption keys, and a list of system administrators in the cell. Initially, this directory is empty.

- /usr/afs/local

  This directory contains local server configuration files specific for each AFS server machine. Initially, this directory is empty.

- /usr/afs/db

  This directory contains database files. Initially, this directory is empty.

- /usr/afs/backup

This directory contains backup configuration files. Initially, this directory is empty.

- /usr/afs/logs

  This directory contains server log files. Initially, this directory is empty.

- /usr/vice/etc

  This directory contains client configuration files and kernel extension files. Initially, this directory is empty.

As we mentioned in “AIX File System Helper” on page 62, on AIX systems the installation program also replaces the AIX file system helper program, v3fshelper, with an AFS-aware helper. The file system checker program, fsck, uses the helper program to perform file system integrity checks. The AIX fsck program discards files from directories that are not part of the AIX JFS, including the /vicep*x* directories. These directories are non-standard AIX directories, containing information that make up AFS volumes. On the other hand, the AFS-aware v3fshelper correctly handles both AFS and AIX file systems.

> **AIX fsck and AFS Partitions**
>
> Never run the standard AIX file system checker program on an AFS File Server. It discards AFS volumes.

The installation program correctly replaces the AIX helper program with the Transarc one. But it does not save the original AIX executable. In fact, by comparing the listing of the /sbin/helpers directory and the /usr/afs/bin directory, you can see that the AFS-aware v3fshelper executable has the same identical size as the v3fshelper.noafs file, which should be the non-AFS original version:

```
# cd /sbin/helpers
# ls -l v3fshelper*
-r-xr-xr-x   1 bin      bin       168634 Mar 06 02:14 v3fshelper
-r-xr-xr-x   1 root     system    168634 Mar 06 02:14 v3fshelper.noafs

# cd /usr/afs/bin
# ls -l v3fshelper*
-rwxr-xr-x   1 root     system    168634 Mar 06 02:14 v3fshelper
```

Another confirmation that these files are indeed identical comes by running the `cmp` AIX comparing command. For this reason, we recommend that prior to performing the installation, you save a copy of the original version of the v3fshelper executable file (see again “AIX File System Helper” on page 62).

3.3.2 WebSphere AFS Configuration Script

The WebSphere Performance Pack configuration process uses a configuration script called afsConfigure.sh. During the installation of the File Sharing component of Performance Pack, the setup program creates this script in the directory tree of the specified destination directory. The default path is /usr/lpp/WSPP (AIX) or /opt/WSPP (Solaris). If you are using the default path, the setup program installs this script in the /usr/lpp/WSPP/afs (AIX) or /opt/WSPP/afs (Solaris) directory. This script assists you in the following tasks:

- Configuring the first AFS machine

- Configuring additional AFS File Server machines
- Configuring AFS Client machines

3.4 Configuring the First AFS Machine on AIX Systems

Before configuring an AIX machine to be the first AFS machine, you must complete the following tasks:

- Install the necessary AFS files from the WebSphere Performance Pack CD-ROM distribution to the machine's local disk, as shown in 3.3, "Installing AFS Files on AIX Systems" on page 62.
- Create at least one AFS partition to store AFS volumes, as illustrated in 3.2, "Configuring AFS Partitions on AIX Systems" on page 59.

Then, you can continue with what is indicated in 3.4.1, "Using the WebSphere AFS Configuration Script" on page 68 and 3.4.2, "Verifying the First AFS Machine Configuration" on page 75.

3.4.1 Using the WebSphere AFS Configuration Script

To start the configuration of the first AFS machine, perform the following tasks:

1. Login as root
2. From a command prompt, enter the following commands:

   ```
   cd /usr/lpp/WSPP
   cd afs
   ```

3. Start the AFS configuration script by entering:

   ```
   ./afsConfigure.sh
   ```

   While the script is running, since authorization checking is disabled during the configuration process, you may safely ignore error messages similar to:

   ```
   bos: no such entry (getting tickets)
   bos: running unauthenticated
   libprot: no such entry Could not get afs tokens, running unauthenticated
   vsu_ClientInit: Could not get afs tokens, running unauthenticated
   ```

4. After the script displays an option menu, enter `1` to instruct the AFS configuration script to configure this machine as the SCM, as shown in the following session:

   ```
   # cd /usr/lpp/WSPP
   # cd afs
   # ./afsConfigure.sh
   Please choose from the following options
   1. Configure this machine as your SCM.
   2. Configure this machine as a File Server.
   3. Configure this machine as an AFS Client.
   4. Quit the program.
   Enter choice: 1
   ```

   In WebSphere Performance Pack, the role of SCM is played by the first AFS machine, which is also an AFS Database Server, an AFS File Server, a BDM for its AFS system type, and an AFS Client.

The AFS configuration script, afsConfigure.sh, proceeds to configure the first AFS machine. You will be prompted to interact with the system to provide the information necessary to complete the configuration. The configuration steps are listed in the next steps.

5. You must enter the name of an AFS partition on this first AFS machine. You should specify the full path name of this partition. In our environment, we configure /vicepa as the first AFS partition on the first AFS machine, as shown in the following session:

```
Enter the name of your partition: /vicepa
```

6. The script starts the Basic OverSeer (BOS) Server process with authorization checking disabled, since the Authentication Database does not exist yet. The following command is automatically executed by the script:

```
/usr/afs/bin/bosserver -noauth
```

7. You should enter the name of your cell, by using standard Internet domain name format. You want to select a unique cell name for the AFS global name space. The cell name can have up to 64 characters, including only lowercase letters, numbers, underscores, dashes, and periods. In our environment, our AFS cell name is sushi.itso.ral.ibm.com, as shown in the following session:

```
Enter the name you want your cell to be: sushi.itso.ral.ibm.com
```

   You need to select your cell name wisely, because the process of changing the cell name is very complicated. This procedure updates the following server common configuration files: /usr/afs/etc/ThisCell and /usr/afs/etc/CellServDB.

8. The script starts the Authentication Server process by executing the following command:

```
/usr/afs/bin/kaserver
```

   In the BOS configuration file, /usr/afs/local/BosConfig, the script creates the following entries:

```
bnode simple kaserver 1
parm /usr/afs/bin/kaserver
end
```

   In the AFS database directory, /usr/afs/db, the script initializes the Authentication Database files kaserver.DB0 and kaserver.DBSYS1.

9. You should enter a password for the afs server key and the admin administrative account, as shown in the following session:

```
Enter your AFS and admin user password: XXXXXXXX
```

   The script creates these two entries, `afs` for the afs server key and `admin` for the admin administrative account respectively, in the Authentication Database. It designates the admin user as a privileged account.

   It places the admin account in the server administrative list, /usr/afs/etc/UserList, and the afs server encryption key in the key file,

/usr/afs/etc/KeyFile. The script uses the same password for both entries. Please, remember this password!

AFS Password and afsConfigure.sh

With the AFS configuration script, afsConfigure.sh, the entered password string echoes visibly on the screen!

10.The script starts the Protection Server process by executing the following command:

```
/usr/afs/bin/ptserver
```

In the /usr/afs/local/BosConfig file, the script creates the following entries:

```
bnode simple ptserver 1
parm /usr/afs/bin/ptserver
end
```

In the /usr/afs/db directory, the script initializes the protection database files prdb.DB0 and prdb.DBSYS1. It also creates an `admin` entry in the protection database, and places it in the system administrator group, system:administrators.

11.The script starts the Volume Location Server process by executing the command:

```
/usr/afs/bin/vlserver
```

In the /usr/afs/local/BosConfig file, the script creates the following entries:

```
bnode simple vlserver 1
parm /usr/afs/bin/vlserver
end
```

In the /usr/afs/db directory, the script initializes the volume location database files vldb.DB0 and vldb.DBSYS1.

12.The script starts the Backup Server process by executing the command:

```
/usr/afs/bin/buserver
```

In the /usr/afs/local/BosConfig file, the script creates the following entries:

```
bnode simple buserver 1
parm /usr/afs/bin/buserver
end
```

In the /usr/afs/db directory, the script initializes the volume location database files bdb.DB0 and bdb.DBSYS1.

13.The script starts the fs trio of processes, which includes the File Server, the Volume Server, and the Salvager (see 2.9.2, “AFS Server Processes” on page 41). The following commands are automatically executed by the script:

```
/usr/afs/bin/fileserver
/usr/afs/bin/volserver
/usr/afs/bin/salvager
```

In the /usr/afs/local/BosConfig file, the script creates the following entries:

```
bnode fs fs 1
parm /usr/afs/bin/fileserver
parm /usr/afs/bin/volserver
parm /usr/afs/bin/salvager
```

```
end
```

14. The script creates the root volume for AFS, root.afs, on the specified AFS partition, /vicepa, on this first AFS machine.

15. The script starts the server portion of the Update Server by executing the following command:

    ```
    /usr/afs/bin/upserver -clear /usr/afs/bin
    ```

    Since WebSphere Performance Pack uses the international edition of AFS without encryption, the Update Server distributes only the AFS binary files residing in /usr/afs/bin. Due to the sensitive nature of the server common configuration files, /usr/afs/etc, we recommend that you use the `bos` command to manually update these files.

    In the /usr/afs/local/BosConfig file, the script creates the following entries:

    ```
    bnode simple upserver 1
    parm /usr/afs/bin/upserver -clear /usr/afs/bin
    end
    ```

16. The script starts the controller process for the NTP daemon by executing the following command:

    ```
    /usr/afs/bin/runntp -localclock
    ```

    The script assumes that this first AFS machine does not have an external time source to synchronize its clock and that it uses its internal clock to synchronize time amongst machines in the cell.

    In the /usr/afs/local/BosConfig file, the script creates the following entries:

    ```
    bnode simple runntp 1
    parm /usr/afs/bin/runntp -localclock
    end
    ```

At this point in the configuration process, Step 5 on page 69 through Step 16 on page 71, your first AFS machine is running the database server processes, the File Server processes, and the Update Server process for its AFS system type. To complete the configuration of this first AFS machine, the script proceeds to configure the AFS Client process.

17. The script defines the cell information for AFS client programs in the following files: /usr/vice/etc/ThisCell and /usr/vice/etc/CellServDB.

18. You are prompted to enter the size of the AFS disk cache in 1 KB blocks. The default AFS cache directory is /usr/vice/cache. In this example, we chose an AFS disk cache size of 40,000 1 KB blocks (about 40 MB):

    ```
    Enter the number of 1 KB blocks to devote to the cache: 40000
    ```

    As we said in 3.1.5, “Disk Space Requirements on UNIX Systems” on page 57, an AFS disk cache size should not exceed 85% of the size of its disk partition as reported by the `df` command. Make sure that there is adequate disk space to accommodate the desired cache size.

19. The script creates the AFS cache information file, /usr/vice/etc/cacheinfo, with the following entry:

    ```
    /afs:/usr/vice/cache:40000
    ```

20. The script creates the root of the AFS name space, /afs.

21. The script starts the AFS cache manager process by executing the following command:

```
/usr/vice/etc/afsd -stat 2000 -dcache 800 -daemons 3 -volumes 70 -nosettime
```

22. The script creates and mounts top level volumes of the AFS file tree of your cell.

In summary, the AFS configuration script prompts you for the name of the AFS partition, the name of the AFS cell, the afs and admin password, and the size of the AFS disk cache in 1 KB blocks.

In our environment, the name of our first AFS partition is /vicepa, the name of our AFS cell is sushi.itso.ral.ibm.com, and the size of our AFS disk cache is 40,000 1 KB blocks.

The next several text figures contain the transcript of our configuration session of the first AFS machine on an AIX system called rs60002:

```
# cd /usr/lpp/WSPP
# cd afs
# ./afsConfigure.sh
Please choose from the following options
1. Configure this machine as your SCM.
2. Configure this machine as a File Server.
3. Configure this machine as an AFS Client.
4. Quit the program.
Enter choice: 1
This machine is your SCM


Setting up first AFS Server
nfsd was found on this machine

Enter the name of your partition: /vicepa
Creating license file

Starting the bosserver

bosserver started successfully

Enter the name you want your cell to be: sushi.itso.ral.ibm.com

Setting the cell name to sushi.itso.ral.ibm.com
bos: no such entry (getting tickets)
bos: running unauthenticated

Starting the Authentication Server
bos: no such entry (getting tickets)
bos: running unauthenticated

Enter your AFS and admin user password: XXXXXXXX
AFSpassword is XXXXXXXX

Creating AFS and ADMIN users...
bos: no such entry (getting tickets)
bos: running unauthenticated
bos: no such entry (getting tickets)
bos: running unauthenticated
bos: no such entry (getting tickets)
bos: running unauthenticated

Starting the Protection Server
bos: no such entry (getting tickets)
bos: running unauthenticated

Adding admin to system:administrators group
libprot: no such entry Could not get afs tokens, running unauthenticated.
User admin has id 1
libprot: no such entry Could not get afs tokens, running unauthenticated.
```

Figure 27. (Part 1 of 3). Configuration of the First AFS Machine on an AIX System

```
Starting the Volume Location Server
bos: no such entry (getting tickets)
bos: running unauthenticated

Starting the Backup Server
bos: no such entry (getting tickets)
bos: running unauthenticated

Starting the File Server, Volume Server and Salvager
bos: no such entry (getting tickets)
bos: running unauthenticated

Creating the cell's root volume root.afs
vsu_ClientInit: Could not get afs tokens, running unauthenticated.
Volume 536870912 created on partition /vicepa of rs60002

Starting the Server Portion of the Update Server
bos: no such entry (getting tickets)
bos: running unauthenticated

Starting the Controller for NTPD
bos: no such entry (getting tickets)
bos: running unauthenticated

Enter the number of 1 KB blocks to devote to the cache: 40000
 U.S. Government Restricted Rights.
====================================
If you are licensing the Software on behalf of the U.S.
Government ("Government"), the following provisions apply
to you. If the Software is supplied to the Department of
Defense ("DoD"), it is classified as "Commercial Computer
Software" under paragraph 252.227-7014 of the DoD Supplement
to the Federal Acquisition Regulations ("DFARS") (or
any successor regulations) and the Government is acquiring
only the license rights granted herein (the license rights
customarily provided to non-Government users).  If the
Software is supplied to any unit or agency of the Government
other than DoD, it is classified as "Restricted Computer
Software" and the Government's rights in the Software are
defined in paragraph 52.227-19 of the Federal Acquisition
Regulations ("FAR") (or any successor regulations) or, in
the case of NASA, in paragraph 18.52.227-86 of the NASA
Supplement in the FAR (or any successor regulations).

afsd: My home cell is 'sushi.itso.ral.ibm.com'
afsd: Creating '/usr/vice/etc/AFSLog'
CreateCacheFile: Creating cache file '/usr/vice/etc/AFSLog'
afsd: Forking rx callback listener.
afsd: Forking AFS daemon.
afsd: Forking Check Server Daemon.
afsd: Forking 3 background daemons.
```

Figure 28. (Part 2 of 3). Configuration of the First AFS Machine on an AIX System

```
afsd: Calling AFSOP_CACHEINIT: 2000 stat cache entries, 4000 optimum cache
files, 40000 blocks in the cache, flags = 0x0, dcache entries 800
afsd: Sweeping workstation's AFS cache directory.
SweepAFSCache: Creating '/usr/vice/cache/CacheItems'
CreateCacheFile: Creating cache file '/usr/vice/cache/CacheItems'
SweepAFSCache: Creating '/usr/vice/cache/VolumeItems'
CreateCacheFile: Creating cache file '/usr/vice/cache/VolumeItems'
SweepAFSCache: Creating '/usr/vice/cache/V0'
CreateCacheFile: Creating cache file '/usr/vice/cache/V0'
SweepAFSCache: Creating '/usr/vice/cache/V1'
CreateCacheFile: Creating cache file '/usr/vice/cache/V1'
SweepAFSCache: Creating '/usr/vice/cache/V2'
CreateCacheFile: Creating cache file '/usr/vice/cache/V2'
...
SweepAFSCache: Creating '/usr/vice/cache/V3997'
CreateCacheFile: Creating cache file '/usr/vice/cache/V3997'
SweepAFSCache: Creating '/usr/vice/cache/V3998'
CreateCacheFile: Creating cache file '/usr/vice/cache/V3998'
SweepAFSCache: Creating '/usr/vice/cache/V3999'
CreateCacheFile: Creating cache file '/usr/vice/cache/V3999'
afsd: 0 out of 4000 data cache files found in sweep 1.
afsd: 4000 out of 4000 data cache files found in sweep 2.
afsd: All AFS daemons started.
afsd: Forking trunc-cache daemon.
afsd: Mounting the AFS root on '/afs', flags: 0.

Tokens held by the Cache Manager:

User's (AFS ID 1) tokens for afs@sushi.itso.ral.ibm.com [Expires Apr  1
21:53]
   --End of list--
All volumeID/name mappings checked.
Volume 536870915 created on partition /vicepa of rs60002
Added replication site rs60002 /vicepa for volume root.afs
Added replication site rs60002 /vicepa for volume root.cell
Released volume root.afs successfully
Released volume root.cell successfully
All volumeID/name mappings checked.
Volume status for vid = 536870913 named root.afs.readonly
Current disk quota is 5000
Current blocks used are 4
The partition has 126894 blocks available out of 131072

Volume status for vid = 536870916 named root.cell.readonly
Current disk quota is 5000
Current blocks used are 2
The partition has 126894 blocks available out of 131072
```

Figure 29. (Part 3 of 3). Configuration of the First AFS Machine on an AIX System

3.4.2 Verifying the First AFS Machine Configuration

To verify that the configuration of the first AFS machine completed successfully, perform the following steps:

1. Display the status of AFS processes on this first AFS machine by entering the following command:

   ```
   /usr/afs/bin/bos status hostname -long
   ```

   where *hostname* is the host name of this first AFS machine. Figure 30 on page 76 shows a sample output from the above command.

```
# /usr/afs/bin/bos status rs60002 -long
Bosserver reports inappropriate access on server directories
Instance kaserver, (type is simple) currently running normally.
    Process last started at Wed Mar 31 20:17:34 1999 (1 proc starts)
    Command 1 is '/usr/afs/bin/kaserver'

Instance ptserver, (type is simple) currently running normally.
    Process last started at Wed Mar 31 20:18:40 1999 (1 proc starts)
    Command 1 is '/usr/afs/bin/ptserver'

Instance vlserver, (type is simple) currently running normally.
    Process last started at Wed Mar 31 20:19:41 1999 (1 proc starts)
    Command 1 is '/usr/afs/bin/vlserver'

Instance buserver, (type is simple) currently running normally.
    Process last started at Wed Mar 31 20:20:39 1999 (1 proc starts)
    Command 1 is '/usr/afs/bin/buserver'

Instance fs, (type is fs) currently running normally.
    Auxiliary status is: file server running.
    Process last started at Wed Mar 31 20:21:39 1999 (2 proc starts)
    Command 1 is '/usr/afs/bin/fileserver'
    Command 2 is '/usr/afs/bin/volserver'
    Command 3 is '/usr/afs/bin/salvager'

Instance upserver, (type is simple) currently running normally.
    Process last started at Wed Mar 31 20:22:39 1999 (1 proc starts)
    Command 1 is '/usr/afs/bin/upserver -clear /usr/afs/bin'

Instance runntp, (type is simple) currently running normally.
    Process last started at Wed Mar 31 20:22:41 1999 (1 proc starts)
    Command 1 is '/usr/afs/bin/runntp -localclock'
```

Figure 30. Displaying Status of AFS Server Processes on the First AFS Machine

   The AFS command:

   ```
   bos status
   ```

   displays the status of AFS Server processes listed in the BOS configuration file, /usr/afs/local/BosConfig, as shown in Figure 31 on page 77. We recommend that you do not edit this file directly. You should use the AFS `bos` command suite to modify this file.

```
# cat /usr/afs/local/BosConfig
restarttime 11 0 4 0 0
checkbintime 3 0 5 0 0
bnode simple kaserver 1
parm /usr/afs/bin/kaserver
end
bnode simple ptserver 1
parm /usr/afs/bin/ptserver
end
bnode simple vlserver 1
parm /usr/afs/bin/vlserver
end
bnode simple buserver 1
parm /usr/afs/bin/buserver
end
bnode fs fs 1
parm /usr/afs/bin/fileserver
parm /usr/afs/bin/volserver
parm /usr/afs/bin/salvager
end
bnode simple upserver 1
parm /usr/afs/bin/upserver -clear /usr/afs/bin
end
bnode simple runntp 1
parm /usr/afs/bin/runntp -localclock
end
```

Figure 31. Content of the /usr/afs/local/BosConfig File for the First AFS Machine

2. List the running AFS processes on this first AFS machine by entering the following command:

   `ps -ef | grep afs`

   Figure 32 on page 77 shows a sample output of the above command.

```
# ps -ef | grep afs
root 12944     1   0 20:13:44      -  0:00 /usr/afs/bin/bosserver -noauth
root 13110 12944   0 20:17:34      -  0:00 /usr/afs/bin/kaserver
root  3216 12944   0 20:18:40      -  0:00 /usr/afs/bin/ptserver
root 14190 12944   0 20:19:41      -  0:00 /usr/afs/bin/vlserver
root 14448 12944   0 20:20:41      -  0:00 /usr/afs/bin/buserver
root 14706 12944   0 20:21:41      -  0:00 /usr/afs/bin/fileserver
root 14964 12944   0 20:21:41      -  0:00 /usr/afs/bin/volserver
root 12672 12944   0 20:22:41      -  0:00 /usr/afs/bin/upserver -clear /usr/afs/bin
root 13534 12944   0 20:22:41      -  0:00 /usr/afs/bin/runntp -localclock
root 15222     1   0 20:24:46     -  0:00 /usr/vice/etc/afsd -stat 2000 -dcache 800 -daemons 3 -volumes 70 -nosettime -verbose
root 15480     1   0 20:24:46     -  0:00 /usr/vice/etc/afsd -stat 2000 -dcache 800 -daemons 3 -volumes 70 -nosettime -verbose
root 15738     1   0 20:24:46     -  0:00 /usr/vice/etc/afsd -stat 2000 -dcache 800 -daemons 3 -volumes 70 -nosettime -verbose
root 15996     1   0 20:24:46     -  0:00 /usr/vice/etc/afsd -stat 2000 -dcache 800 -daemons 3 -volumes 70 -nosettime -verbose
root 16254     1   0 20:24:46     -  0:00 /usr/vice/etc/afsd -stat 2000 -dcache 800 -daemons 3 -volumes 70 -nosettime -verbose
root 16512     1   0 20:24:46     -  0:00 /usr/vice/etc/afsd -stat 2000 -dcache 800 -daemons 3 -volumes 70 -nosettime -verbose
root 16770     1   0 20:24:46     -  0:00 /usr/vice/etc/afsd -stat 2000 -dcache 800 -daemons 3 -volumes 70 -nosettime -verbose
root 17028     1   0 20:24:46     -  0:00 /usr/vice/etc/afsd -stat 2000 -dcache 800 -daemons 3 -volumes 70 -nosettime -verbose
root 17286     1   0 20:24:46     -  0:00 /usr/vice/etc/afsd -stat 2000 -dcache 800 -daemons 3 -volumes 70 -nosettime -verbose
root 17544     1   0 20:24:46     -  0:00 /usr/vice/etc/afsd -stat 2000 -dcache 800 -daemons 3 -volumes 70 -nosettime -verbose
root 17802     1   0 20:24:46     -  0:00 /usr/vice/etc/afsd -stat 2000 -dcache 800 -daemons 3 -volumes 70 -nosettime -verbose
root 18060     1   0 20:24:46     -  0:00 /usr/vice/etc/afsd -stat 2000 -dcache 800 -daemons 3 -volumes 70 -nosettime -verbose
root 18318     1   0 20:25:26      -  0:00 /usr/vice/etc/afsd -stat 2000 -dcache 800 -daemons 3 -volumes 70 -nosettime -verbose
```

Figure 32. Listing of AFS Processes on the First AFS Machine

3. The AFS command:

   `kas examine`

displays information about an entry in the Authentication Database. A user with administrative privileges has the `ADMIN` flag in his or her Authentication Database entry. This user can issue any AFS `kas` command. To verify that the admin user has this flag, enter the following command:

```
/usr/afs/bin/kas examine admin
```

The following session shows that the admin user has indeed the `ADMIN` flag. Note that the octal presentation of the key is only visible when authorization checking is disabled. You can ignore any error message since authorization checking is still disabled during the configuration process.

```
# /usr/afs/bin/kas examine admin
Password for root: (Press Space and <Enter>)
kas:examine: Auth. as root to AuthServer failed: user doesn't exist
Proceeding w/o authentication


User data for admin (ADMIN)
  key (0):p7\032[\346\310u\002, last cpw: Wed Mar 31 20:18:40 1999
  password will never expire.
  An unlimited number of unsuccessful authentications is permitted.
  entry never expires.  Max ticket lifetime 25.00 hours.
  last mod on Wed Mar 31 20:18:40 1999 by <none>
  permit password reuse
```

Figure 33. Examining the AFS admin User

4. The AFS command:

   ```
   bos listusers
   ```

   lists the users from the /usr/afs/etc/UserList file. These users can issue any privileged AFS `bos` and `vos` command. To confirm that the `admin` user has an entry in the /usr/afs/etc/UserList file, enter the following command:

   /usr/afs/bin/bos listusers *hostname*

   where *hostname* is the host name of this first AFS machine. The following session shows that the `admin` user appears indeed in the /usr/afs/etc/UserList file:

```
# /usr/afs/bin/bos listusers rs60002
SUsers are: admin
```

Figure 34. Listing of the Privileged Users

5. The AFS command:

   ```
   pts membership
   ```

   lists the groups to which a user or machine belongs, or lists the users and machines that are members of each group. A user who belongs to the system:administrators group can issue any AFS `pts` and `fs` command. This user also has implicit administer (`a`) rights on every ACL in the AFS file system of the local cell (see 2.6.6, “Protecting AFS Files and Directories” on page 38). To verify that the admin user is a member of the system:administrators group, enter the following command:

```
/usr/afs/bin/pts membership admin
```

The following session shows that the admin user is indeed a member of the system:administrators group:

```
# /usr/afs/bin/pts membership admin
Groups admin (id: 1) is a member of:
  system:administrators
```

Figure 35. Displaying AFS Group Membership of the admin User

To verify that the AFS client is functioning properly, you need to perform the tasks in 3.6.2, “Verifying the AFS Client Configuration” on page 92.

To configure the AFS initialization file to start AFS at system reboot, you need to follow the procedures in 3.7, “Starting AFS at System Reboot on AIX Systems” on page 94. After completing the AFS initialization file configuration, reboot your machine and verify that all AFS processes are running normally.

3.5 Configuring an AFS File Server Machine on AIX Systems

Before configuring an AIX machine to be an AFS File Server in an existing AFS cell, you must complete the following tasks, which are prerequisites to the instructions in 3.5.1, “Using the WebSphere AFS Configuration Script” on page 79 and 3.5.2, “Verifying the AFS File Server Configuration” on page 87:

- Install the necessary AFS files from the WebSphere Performance Pack CD-ROM distribution to the machine’s local disk, as shown in 3.3, “Installing AFS Files on AIX Systems” on page 62.
- Create at least one AFS partition to store AFS volumes, as illustrated in 3.2, “Configuring AFS Partitions on AIX Systems” on page 59.

3.5.1 Using the WebSphere AFS Configuration Script

The WebSphere AFS configuration script assumes that any AFS File Server is of the same operating system type as the first AFS machine. If you are configuring an AFS File Server of a new system type, you must create a server instance of the Update Server and not a client instance of the Update Server.

At the time of the writing of this book, the WebSphere AFS configuration script did not correctly configure the Update Server for a heterogeneous cell. For example, if both the first AFS machine and the AFS File Server machine are AIX systems, the configuration script correctly configures the client instance of the Update Server on the AFS File Server; however, if the first AFS machine is an AIX system and the AFS File Server machine is a Solaris system, the configuration script incorrectly configures the client instance of the Update Server on the AFS File Server, resulting in the Solaris AFS File Server accepting AIX binary updates from the first AIX AFS machine.

To correctly configure AFS Servers in a heterogeneous cell with various AFS system types, you need to modify the WebSphere AFS configuration script as shown in Figure 36 on page 80, Figure 37 on page 81, and Figure 38 on page 82:

1. Remove the highlighted lines in the afsConfigure.sh file, as shown in Figure 36 on page 80.

```
...
setupFS()
{
   echo "Setting up File Server"
...
#  Create an instance of the Update Server
#  So first query the cell name and machine name of primary machine
   while :
   do
   {
     echo "Enter the name of your cell: \c"
     read CELLNAME
     if [ "${CELLNAME}" = "" ]
     then
     {
       continue
     }
     fi
     break
   }
   done

   while :
   do
   {
     echo "Enter the hostname of your first AFS Server: \c"
     read SCM
     if [ "${SCM}" = "" ]
     then
     {
       continue
     }
     fi
     break
   }
   done
   /usr/afs/bin/bos create $MACHINE_NAME upclientbin simple "/usr/afs/bin/upclient $SCM -clear
/usr/afs/bin" -cell $CELLNAME

#  Start the runntp process
   /usr/afs/bin/bos create $MACHINE_NAME runntp simple /usr/afs/bin/runntp -cell $CELLNAME

#  Start the fs process, which binds the file server, volume server and salvager
   /usr/afs/bin/bos create $MACHINE_NAME fs fs /usr/afs/bin/fileserver /usr/afs/bin/volserver
/usr/afs/bin/salvager -cell $CELLNAME
   sleep $SLEEPTIME

#  Start the Client.
   setUpAFSClient
} # setup FS
```

Figure 36. Listing of the setupFS() Function in the WebSphere AFS Configuration File afsConfigure.sh

2. Add the highlighted lines in the afsConfigure.sh file, as shown in Figure 37 on page 81 and Figure 38 on page 82.

```
...
getBDM()
{
    while :
    do
    {
        echo
        echo "$MSG \c"
        read BDM
        if [ "${BDM}" = "" ]
        then
        {
            continue
        }
        fi
        break
    }
    done
}

setupFS()
{
   echo "Setting up File Server"
...

#  Create an instance of the Update Server
#  So first query the cell name and machine name of primary machine
   while :
   do
   {
     echo "Enter the name of your cell: \c"
     read CELLNAME
     if [ "${CELLNAME}" = "" ]
     then
     {
       continue
     }
     fi
     break
   }
   done
```

Figure 37. (Part 1 of 2). Modifications to the WebSphere AFS Configuration Script afsConfigure.sh

```
 while :
   do
     echo
     echo "The AFS system type of this AFS File Server is"
     echo "    1 - The same as the first AFS machine"
     echo "    2 - The same as an existing AFS Server"
     echo "    3 - A new system type"
     echo ""
     echo "Please enter the AFS system type for this AFS File Server: \c"
     read TYPE

     case $TYPE in
     1 ) MSG="Enter the host name of the first AFS machine:"
         getBDM
         echo
         echo "Starting the client portion of the Update Server"
         /usr/afs/bin/bos create $MACHINE_NAME upclientbin simple "/usr/afs/bin/upclient $BDM
-clear /usr/afs/bin" -cell $CELLNAME
         ;;

     2 ) MSG="Enter the host name of the AFS Server with the same system type:"
         getBDM
         echo
         echo "Starting the client portion of the Update Server"
         /usr/afs/bin/bos create $MACHINE_NAME upclientbin simple "/usr/afs/bin/upclient $BDM
-clear /usr/afs/bin" -cell $CELLNAME
         ;;
     3 ) echo
         echo "Starting the server portion of the Update Server"
         /usr/afs/bin/bos create $MACHINE_NAME upserver simple "/usr/afs/bin/upserver -clear
/usr/afs/bin" -cell $CELLNAME
         ;;

     * ) echo
         echo "Unknown AFS system type"
         continue
         ;;
     esac

     break
   done

#  Start the runntp process
   /usr/afs/bin/bos create $MACHINE_NAME runntp simple /usr/afs/bin/runntp -cell $CELLNAME

#  Start the fs process, which binds the file server, volume server and salvager
   /usr/afs/bin/bos create $MACHINE_NAME fs fs /usr/afs/bin/fileserver /usr/afs/bin/volserver
/usr/afs/bin/salvager -cell $CELLNAME
   sleep $SLEEPTIME

#  Start the Client.
   setUpAFSClient
} # setup FS
```

Figure 38. (Part 2 of 2). Modifications to the WebSphere AFS Configuration Script afsConfigure.sh

To start the configuration of an AFS File Server machine, perform the following tasks:

1. Log in as root.
2. Copy or get through FTP the contents of the AFS Server common configuration directory, /usr/afs/etc, from the first AFS machine to the directory with the same name on this machine. The files in the /usr/afs/etc directory include:
   - CellServDB
   - KeyFile
   - License
   - ThisCell
   - UserList
3. From a command prompt, enter the following command:

   ```
   cd /usr/lpp/WSPP/afs
   ```

   While the script is running, since authorization checking is disabled during the configuration process, you may safely ignore error messages similar to:

   ```
   bos: no such entry (getting tickets)
   bos: running unauthenticated
   libprot: no such entry Could not get afs tokens, running unauthenticated
   vsu_ClientInit: Could not get afs tokens, running unauthenticated
   ```

4. Start the AFS configuration script by entering:

   ```
   ./afsConfigure.sh
   ```

5. After the script displays an option menu, enter **2** to instruct the AFS configuration script to configure this machine as an AFS File Server, as shown in the following session screen:

   ```
   # cd /usr/lpp/WSPP/afs
   # ./afsConfigure.sh
   Please choose from the following options
   1. Configure this machine as your SCM.
   2. Configure this machine as a File Server.
   3. Configure this machine as an AFS Client.
   4. Quit the program.
   Enter choice: 2
   ```

Notice that in WebSphere Performance Pack, the AFS File Server is automatically also an AFS Client.

The AFS configuration script, afsConfigure.sh, proceeds to configure an AFS File Server. You will be prompted to interact with the system to provide the information necessary to complete the configuration. The configuration steps are listed below:

1. You are prompted to enter the name of an AFS partition on this AFS File Server. You must specify the full path name of this partition. In our environment, we configure /vicepa as the first AFS partition on this AFS File Server, as shown in the following session:

   ```
   Enter the name of your partition: /vicepa
   ```

2. The script starts the Basic OverSeer (BOS) Server with authorization checking disabled. To do this, the script automatically executes the following command:

   ```
   /usr/afs/bin/bosserver -noauth
   ```

3. You are prompted to enter the name of your cell, by using standard Internet domain name format. You need to know the name of an existing AFS cell. In our environment, the existing AFS cell name is sushi.itso.ral.ibm.com, so this is the cell name we enter, as shown in the following session:

   ```
   Enter the name of your cell: sushi.itso.ral.ibm.com
   ```

4. Enter the host name of the first AFS machine for the selected AFS cell. If you have DNS in your environment, we recommend that you use the fully qualified DNS host name to refer to your first AFS machine. In our environment, the fully qualified DNS host name of the first AFS machine in the AFS cell sushi.itso.ral.ibm.com is rs60002.itso.ral.ibm.com, so this is the value we enter, as shown in the following session:

   ```
   Enter the hostname of your first AFS Server: rs60002.itso.ral.ibm.com
   ```

5. The script creates an instance of the client portion of the Update Server by executing the following command:

   ```
   /usr/afs/bin/upclient BDM -clear /usr/afs/bin
   ```

   where *BDM* is the host name of the BDM of the same AFS system type as the AFS File Server that you are configuring. In WebSphere Performance Pack, the BDM is usually the SCM or the first AFS machine. This host name is obtained in Step 4 on page 84.

   In the /usr/afs/local/BosConfig file, the script creates the following entries:

   ```
   bnode simple upclientbin 1
   parm /usr/afs/bin/upclient rs60002.itso.ral.ibm.com -clear /usr/afs/bin
   end
   ```

   > **Server or Client Instance of the Update Server**
   >
   > The WebSphere AFS configuration script assumes that any AFS File Server is of the same system type as the first AFS machine. If you are configuring an AFS File Server of a new system type, you must create a server instance of the Update Server and not a client instance of the Update Server.

6. The script starts the controller process for the NTP daemon by executing the following command:

   ```
   /usr/afs/bin/runntp
   ```

   In the /usr/afs/local/BosConfig file, the script creates the following entries:

   ```
   bnode simple runntp 1
   parm /usr/afs/bin/runntp
   end
   ```

7. The script starts the fs trio of processes, which includes the File Server, the Volume Server, and the Salvager (see 2.9.2, “AFS Server Processes” on page 41). The following commands are automatically executed by the script:

```
/usr/afs/bin/fileserver
/usr/afs/bin/volserver
/usr/afs/bin/salvager
```

In the /usr/afs/local/BosConfig file, the script creates the following entries:

```
bnode fs fs 1
parm /usr/afs/bin/fileserver
parm /usr/afs/bin/volserver
parm /usr/afs/bin/salvager
end
```

At this point in the configuration process, Step 1 on page 83 through Step 7 on page 84, your AFS File Server is running the File Server processes, and the Update Client process of its AFS system type. To complete the configuration of this AFS File Server, the script proceeds to configure the AFS Client process.

8. The script defines the cell information for AFS client programs in the following files: /usr/vice/etc/ThisCell and /usr/vice/etc/CellServDB.
9. You are prompted to enter the size of the AFS disk cache in 1 KB blocks. The default AFS cache directory is /usr/vice/cache. In this example, we choose an AFS disk cache size of 40,000 1 KB blocks (about 40 MB):

   ```
   Enter the number of 1 KB blocks to devote to the cache: 40000
   ```

   Make sure that there is adequate disk space to accommodate the desired cache size.
10. The script creates the AFS cache information file, /usr/vice/etc/cacheinfo, with the following entry:

    ```
    /afs:/usr/vice/cache:40000
    ```
11. The script creates the root of the AFS name space, /afs.
12. The script starts the AFS cache manager process by executing the following command:

    ```
    /usr/vice/etc/afsd -stat 2000 -dcache 800 -daemons 3 -volumes 70 -nosettime
    ```

In summary, the configuration script prompts you for the name of the AFS partition, the name of an existing AFS cell, the host name of your first AFS machine, and the size of the AFS disk cache in 1 KB blocks.

In our environment, the name of the AFS partition on this AFS File Server is /vicepa, the name of our existing AFS cell is sushi.itso.ral.ibm.com, the host name of our first AFS machine is rs60002.itso.ral.ibm.com, and the size of our AFS disk cache is 40,000 1 KB blocks.

The next several text figures contain the transcript of our configuration session of the AFS File Server on an AIX system called rs600030:

```
# cd /usr/afs/etc
(Copy or FTP the contents of /usr/afs/etc from your first AFS machine)
# cd /usr/lpp/WSPP
# cd afs
# ./afsConfigure.sh
Please choose from the following options
1. Configure this machine as your SCM.
2. Configure this machine as a File Server.
3. Configure this machine as an AFS Client.
4. Quit the program.
Enter choice: 2
This machine is a File server

Setting up File Server
nfsd was not found on this system

Enter the name of your partition: /vicepa

Starting the bosserver

bosserver started successfully
Enter the name of your cell: sushi.itso.ral.ibm.com
Enter the hostname of your first AFS Server: rs60002.itso.ral.ibm.com
bos: no such entry (getting tickets)
bos: running unauthenticated
bos: no such entry (getting tickets)
bos: running unauthenticated
bos: no such entry (getting tickets)
bos: running unauthenticated

Enter the number of 1 KB blocks to devote to the cache: 40000
 U.S. Government Restricted Rights.
====================================
If you are licensing the Software on behalf of the U.S.
Government ("Government"), the following provisions apply
to you. If the Software is supplied to the Department of
Defense ("DoD"), it is classified as "Commercial Computer
Software" under paragraph 252.227-7014 of the DoD Supplement
to the Federal Acquisition Regulations ("DFARS") (or
any successor regulations) and the Government is acquiring
only the license rights granted herein (the license rights
customarily provided to non-Government users).  If the
Software is supplied to any unit or agency of the Government
other than DoD, it is classified as "Restricted Computer
Software" and the Government's rights in the Software are
defined in paragraph 52.227-19 of the Federal Acquisition
Regulations ("FAR") (or any successor regulations) or, in
the case of NASA, in paragraph 18.52.227-86 of the NASA
Supplement in the FAR (or any successor regulations).
```

Figure 39. (Part 1 of 2). Configuration of an AFS File Server on an AIX System

```
afsd: My home cell is 'sushi.itso.ral.ibm.com'
afsd: Creating '/usr/vice/etc/AFSLog'
CreateCacheFile: Creating cache file '/usr/vice/etc/AFSLog'
afsd: Forking rx callback listener.
afsd: Forking AFS daemon.
afsd: Forking Check Server Daemon.
afsd: Forking 3 background daemons.
afsd: Calling AFSOP_CACHEINIT: 2000 stat cache entries, 4000 optimum cache
files, 40000 blocks in the cache, flags = 0x0, dcache entries 800
afsd: Sweeping workstation's AFS cache directory.
SweepAFSCache: Creating '/usr/vice/cache/CacheItems'
CreateCacheFile: Creating cache file '/usr/vice/cache/CacheItems'
SweepAFSCache: Creating '/usr/vice/cache/VolumeItems'
CreateCacheFile: Creating cache file '/usr/vice/cache/VolumeItems'
SweepAFSCache: Creating '/usr/vice/cache/V0'
CreateCacheFile: Creating cache file '/usr/vice/cache/V0'
SweepAFSCache: Creating '/usr/vice/cache/V1'
CreateCacheFile: Creating cache file '/usr/vice/cache/V1'
SweepAFSCache: Creating '/usr/vice/cache/V2'
CreateCacheFile: Creating cache file '/usr/vice/cache/V2'
...
SweepAFSCache: Creating '/usr/vice/cache/V3997'
CreateCacheFile: Creating cache file '/usr/vice/cache/V3997'
SweepAFSCache: Creating '/usr/vice/cache/V3998'
CreateCacheFile: Creating cache file '/usr/vice/cache/V3998'
SweepAFSCache: Creating '/usr/vice/cache/V3999'
CreateCacheFile: Creating cache file '/usr/vice/cache/V3999'
afsd: 0 out of 4000 data cache files found in sweep 1.
afsd: 4000 out of 4000 data cache files found in sweep 2.
afsd: All AFS daemons started.
afsd: Forking trunc-cache daemon.
afsd: Mounting the AFS root on '/afs', flags: 0.
```

Figure 40. (Part 2 of 2). Configuration of an AFS File Server on an AIX System

3.5.2 Verifying the AFS File Server Configuration

To verify that the configuration of the AFS File Server completed successfully, perform the following steps:

1. To display the status of AFS processes on an AFS File Server, enter:

   `/usr/afs/bin/bos status` *`hostname`* `-long`

   where *`hostname`* is the host name of this AFS File Server. Figure 41 on page 88 shows a sample output from the command above.

```
# /usr/afs/bin/bos status rs600030 -long
bos: no such entry (getting tickets)
bos: running unauthenticated
Bosserver reports inappropriate access on server directories
Instance upclientbin, (type is simple) currently running normally.
    Process last started at Thu Apr  1 09:05:21 1999 (1 proc starts)
    Command 1 is '/usr/afs/bin/upclient rs60002.itso.ral.ibm.com -clear
                  /usr/afs/bin'

Instance runntp, (type is simple) currently running normally.
    Process last started at Thu Apr  1 09:05:51 1999 (1 proc starts)
    Command 1 is '/usr/afs/bin/runntp'

Instance fs, (type is fs) currently running normally.
    Auxiliary status is: file server running.
    Process last started at Thu Apr  1 09:05:51 1999 (2 proc starts)
    Command 1 is '/usr/afs/bin/fileserver'
    Command 2 is '/usr/afs/bin/volserver'
    Command 3 is '/usr/afs/bin/salvager'
```

Figure 41. Displaying Status of AFS Server Processes on an AFS File Server

The AFS command:

`bos status`

displays the status of AFS Server processes listed in the BOS configuration file, /usr/afs/local/BosConfig, as shown in Figure 42 on page 88. We recommend that you do not edit this file directly. You must use the AFS `bos` command suite to modify this file.

```
# cat /usr/afs/local/BosConfig
restarttime 11 0 4 0 0
checkbintime 3 0 5 0 0
bnode simple upclientbin 1
parm /usr/afs/bin/upclient rs60002.itso.ral.ibm.com -clear /usr/afs/bin
end
bnode simple runntp 1
parm /usr/afs/bin/runntp
end
bnode fs fs 1
parm /usr/afs/bin/fileserver
parm /usr/afs/bin/volserver
parm /usr/afs/bin/salvager
end
```

Figure 42. Content of the /usr/afs/local/BosConfig File for an AFS File Server

2. List the running AFS processes on this AFS File Server by entering the following command:

   `ps -ef | grep afs`

   Figure 43 on page 89 shows a sample output of the command above.

```
# ps -ef | grep afs
root 12700     1   0 09:04:19      -  0:00 /usr/afs/bin/bosserver -noauth
root 13416 12700   0 09:05:51      -  0:01 /usr/afs/bin/upclient rs60002.itso.ral.ibm.com -clear /usr/afs/bin
root 13674 12700   0 09:05:51      -  0:00 /usr/afs/bin/runntp
root 14190 12700   0 09:05:51      -  0:00 /usr/afs/bin/fileserver
root 14448 12700   0 09:05:51      -  0:00 /usr/afs/bin/volserver
root 11660     1   0 09:06:59     -  0:00 /usr/vice/etc/afsd -stat 2000 -dcache 800 -daemons 3 -volumes 70 -nosettime -verbose
root 12156     1   0 09:06:59     -  0:00 /usr/vice/etc/afsd -stat 2000 -dcache 800 -daemons 3 -volumes 70 -nosettime -verbose
root 13160     1   0 09:06:59     -  0:00 /usr/vice/etc/afsd -stat 2000 -dcache 800 -daemons 3 -volumes 70 -nosettime -verbose
root 14706     1   0 09:06:59     -  0:00 /usr/vice/etc/afsd -stat 2000 -dcache 800 -daemons 3 -volumes 70 -nosettime -verbose
root 14964     1   0 09:06:59     -  0:00 /usr/vice/etc/afsd -stat 2000 -dcache 800 -daemons 3 -volumes 70 -nosettime -verbose
root 15222     1   0 09:06:59     -  0:00 /usr/vice/etc/afsd -stat 2000 -dcache 800 -daemons 3 -volumes 70 -nosettime -verbose
root 15480     1   0 09:06:59     -  0:00 /usr/vice/etc/afsd -stat 2000 -dcache 800 -daemons 3 -volumes 70 -nosettime -verbose
root 15738     1   0 09:06:59     -  0:00 /usr/vice/etc/afsd -stat 2000 -dcache 800 -daemons 3 -volumes 70 -nosettime -verbose
root 15996     1   0 09:06:59     -  0:00 /usr/vice/etc/afsd -stat 2000 -dcache 800 -daemons 3 -volumes 70 -nosettime -verbose
root 16254     1   0 09:06:59     -  0:00 /usr/vice/etc/afsd -stat 2000 -dcache 800 -daemons 3 -volumes 70 -nosettime -verbose
root 16512     1   0 09:06:59     -  0:00 /usr/vice/etc/afsd -stat 2000 -dcache 800 -daemons 3 -volumes 70 -nosettime -verbose
root 16770     1   0 09:06:59     -  0:00 /usr/vice/etc/afsd -stat 2000 -dcache 800 -daemons 3 -volumes 70 -nosettime -verbose
root 17028     1   0 09:07:41     -  0:00 /usr/vice/etc/afsd -stat 2000 -dcache 800 -daemons 3 -volumes 70 -nosettime -verbose
```

Figure 43. Listing of AFS Processes on an AFS File Server

To verify that the AFS client is functioning properly, you need to perform the tasks in 3.6.2, “Verifying the AFS Client Configuration” on page 92.

To configure the AFS initialization file to start AFS at system reboot, you need to follow the procedures in 3.7, “Starting AFS at System Reboot on AIX Systems” on page 94. After completing the AFS initialization file configuration, reboot your machine and verify that all AFS processes are running normally.

3.6 Configuring an AFS Client Machine on AIX Systems

To configure an AIX machine to be an AFS client to an existing AFS cell, you must first install the necessary AFS files from the WebSphere Performance Pack CD-ROM distribution to the machine’s local disk, as shown in 3.3, “Installing AFS Files on AIX Systems” on page 62. After that, you can proceed with the instructions indicated in 3.6.1, “Using the WebSphere AFS Configuration Script” on page 89 and 3.6.2, “Verifying the AFS Client Configuration” on page 92.

3.6.1 Using the WebSphere AFS Configuration Script

To start the configuration of an AFS client machine, perform the following tasks:

1. Log in as root.
2. Copy or get through FTP the ThisCell and CellServDB files of the AFS client common configuration directory, /usr/vice/etc, from the first AFS machine to the directory with the same name on this machine.
3. From a command prompt, enter the following command:

   ```
   cd /usr/lpp/WSPP/afs
   ```

4. Start the AFS configuration script by entering:

   ```
   ./afsConfigure.sh
   ```

5. After the script displays an option menu, enter **3** to instruct the AFS configuration script to configure this machine as an AFS client, as shown in the following session:

```
# cd /usr/lpp/WSPP/afs
# ./afsConfigure.sh
Please choose from the following options
1. Configure this machine as your SCM.
2. Configure this machine as a File Server.
3. Configure this machine as an AFS Client.
4. Quit the program.
Enter choice: 3
```

The AFS configuration script, afsConfigure.sh, proceeds to configure an AFS client. During this process, you will be prompted to enter the information necessary to complete the configuration. The configuration steps are as follows:

1. Enter the name of your cell, by using standard Internet domain name format. You need to know the name of an existing AFS cell. In this example, the existing AFS cell name is sushi.itso.ral.ibm.com, as shown in the following session:

```
Enter your cell name: sushi.itso.ral.ibm.com
```

2. Enter the admin password:

```
Enter your admin password: XXXXXXXX
```

   Notice that the password shows up in clear text at the command prompt, although we have masked it here with a sequence of `x`s.

3. Enter the size of the AFS disk cache in 1 KB blocks. The default AFS cache directory is /usr/vice/cache. In this example, we chose an AFS disk cache size of 40,000 1 KB blocks (about 40 MB):

```
Enter the number of 1 KB blocks to devote to the cache: 40000
```

   Make sure that there is adequate disk space to accommodate the desired cache size.

4. The script creates the AFS cache information file, /usr/vice/etc/cacheinfo, with the following entry:

```
/afs:/usr/vice/cache:40000
```

5. The script creates the root of the AFS name space, /afs.
6. The script starts the AFS Cache Manager process by entering the following command:

```
/usr/vice/etc/afsd -stat 2000 -dcache 800 -daemons 3 -volumes 70
```

In summary, the configuration script prompts you for the name of an existing AFS cell, the AFS admin password, and the size of the AFS disk cache in 1 KB blocks.

In our environment, the name of our existing AFS cell is sushi.itso.ral.ibm.com, the AFS admin password is `xxxxxxxx`, and the size of our AFS disk cache is 40,000 1 KB blocks.

The next several text figures contain the transcript of the configuration session of an AFS client on an AIX system called rs600023:

```
# cd /usr/vice/etc
(Copy or FTP the ThisCell and CellServDB file from your first AFS machine)
# cd /usr/lpp/WSPP
# cd afs
# ./afsConfigure.sh
Please choose from the following options
1. Configure this machine as your SCM.
2. Configure this machine as a File Server.
3. Configure this machine as an AFS Client.
4. Quit the program.
Enter choice: 3
This machine will be an AFS Client

Setting up AFS Client
Enter your cell name: sushi.itso.ral.ibm.com
Enter you admin password: XXXXXXXX
nfsd was found on this machine

Enter the number of 1 KB blocks to devote to the cache: 40000
 U.S. Government Restricted Rights.
====================================
If you are licensing the Software on behalf of the U.S.
Government ("Government"), the following provisions apply
to you. If the Software is supplied to the Department of
Defense ("DoD"), it is classified as "Commercial Computer
Software" under paragraph 252.227-7014 of the DoD Supplement
to the Federal Acquisition Regulations ("DFARS") (or
any successor regulations) and the Government is acquiring
only the license rights granted herein (the license rights
customarily provided to non-Government users).  If the
Software is supplied to any unit or agency of the Government
other than DoD, it is classified as "Restricted Computer
Software" and the Government's rights in the Software are
defined in paragraph 52.227-19 of the Federal Acquisition
Regulations ("FAR") (or any successor regulations) or, in
the case of NASA, in paragraph 18.52.227-86 of the NASA
Supplement in the FAR (or any successor regulations).

afsd: My home cell is 'sushi.itso.ral.ibm.com'
afsd: Creating '/usr/vice/etc/AFSLog'
CreateCacheFile: Creating cache file '/usr/vice/etc/AFSLog'
afsd: Forking rx callback listener.
afsd: Forking AFS daemon.
afsd: Forking Check Server Daemon.
afsd: Forking 3 background daemons.
```

Figure 44. (Part 1 of 2). Configuration of an AFS Client on an AIX System

```
afsd: Calling AFSOP_CACHEINIT: 2000 stat cache entries, 4000 optimum cache
files, 40000 blocks in the cache, flags = 0x0, dcache entries 800
afsd: Sweeping workstation's AFS cache directory.
SweepAFSCache: Creating '/usr/vice/cache/CacheItems'
CreateCacheFile: Creating cache file '/usr/vice/cache/CacheItems'
SweepAFSCache: Creating '/usr/vice/cache/VolumeItems'
CreateCacheFile: Creating cache file '/usr/vice/cache/VolumeItems'
SweepAFSCache: Creating '/usr/vice/cache/V0'
CreateCacheFile: Creating cache file '/usr/vice/cache/V0'
SweepAFSCache: Creating '/usr/vice/cache/V1'
CreateCacheFile: Creating cache file '/usr/vice/cache/V1'
SweepAFSCache: Creating '/usr/vice/cache/V2'
CreateCacheFile: Creating cache file '/usr/vice/cache/V2'
...
SweepAFSCache: Creating '/usr/vice/cache/V3997'
CreateCacheFile: Creating cache file '/usr/vice/cache/V3997'
SweepAFSCache: Creating '/usr/vice/cache/V3998'
CreateCacheFile: Creating cache file '/usr/vice/cache/V3998'
SweepAFSCache: Creating '/usr/vice/cache/V3999'
CreateCacheFile: Creating cache file '/usr/vice/cache/V3999'
afsd: 0 out of 4000 data cache files found in sweep 1.
afsd: 4000 out of 4000 data cache files found in sweep 2.
afsd: All AFS daemons started.
afsd: Forking trunc-cache daemon.
afsd: Mounting the AFS root on '/afs', flags: 0.

Tokens held by the Cache Manager:

User's (AFS ID 1) tokens for afs@sushi.itso.ral.ibm.com [Expires Apr  9
20:41]
   --End of list--
All volumeID/name mappings checked.
```

Figure 45. (Part 2 of 2). Configuration of an AFS Client on an AIX System

Modifying the /etc/vfs File

On an AIX system, you have to add the following line to the /etc/vfs file:

```
afs 4 none none
```

Without this entry in the /etc/vfs file, the AIX `mount` command produces error messages, although the command is functioning correctly. This line also enables the AIX operating system to unmount AFS (the local /afs directory) correctly during shutdown.

3.6.2 Verifying the AFS Client Configuration

To verify that the configuration of the AFS client completed successfully, perform the following steps:

1. List the running AFS processes on an AFS client machine by entering the following command:

   ```
   ps -ef | grep afs
   ```

Figure 46 on page 93 shows a sample output of the command above:

```
# ps -ef | grep afs
root  3628     1   0 19:12:14      -  0:00 /usr/vice/etc/afsd -stat 2000 -dcache 800 -daemons 3 -volumes 70 -verbose
root 10992     1   0 19:12:14      -  0:00 /usr/vice/etc/afsd -stat 2000 -dcache 800 -daemons 3 -volumes 70 -verbose
root 11162     1   0 19:12:14      -  0:00 /usr/vice/etc/afsd -stat 2000 -dcache 800 -daemons 3 -volumes 70 -verbose
root 12494     1   0 19:12:14      -  0:00 /usr/vice/etc/afsd -stat 2000 -dcache 800 -daemons 3 -volumes 70 -verbose
root 12654     1   0 19:12:14      -  0:00 /usr/vice/etc/afsd -stat 2000 -dcache 800 -daemons 3 -volumes 70 -verbose
root 12922     1   0 19:12:14      -  0:00 /usr/vice/etc/afsd -stat 2000 -dcache 800 -daemons 3 -volumes 70 -verbose
root 13214     1   0 19:12:14      -  0:00 /usr/vice/etc/afsd -stat 2000 -dcache 800 -daemons 3 -volumes 70 -verbose
root 13418     1   0 19:12:14      -  0:00 /usr/vice/etc/afsd -stat 2000 -dcache 800 -daemons 3 -volumes 70 -verbose
root 13674     1   0 19:12:14      -  0:00 /usr/vice/etc/afsd -stat 2000 -dcache 800 -daemons 3 -volumes 70 -verbose
root 13932     1   0 19:12:14      -  0:00 /usr/vice/etc/afsd -stat 2000 -dcache 800 -daemons 3 -volumes 70 -verbose
root 14190     1   0 19:12:14      -  0:00 /usr/vice/etc/afsd -stat 2000 -dcache 800 -daemons 3 -volumes 70 -verbose
root 14448     1   0 19:12:14      -  0:00 /usr/vice/etc/afsd -stat 2000 -dcache 800 -daemons 3 -volumes 70 -verbose
root 14706     1   0 19:12:56      -  0:00 /usr/vice/etc/afsd -stat 2000 -dcache 800 -daemons 3 -volumes 70 -verbose
```

Figure 46. Listing of AFS Processes on an AFS Client Machine

2. Verify that the /afs tree is available with the command:

   `df -k`

   This is shown in the following session:

```
# df -k
Filesystem    1024-blocks      Free %Used    Iused %Iused Mounted on
/dev/hd4            32768     22032   33%     1525    10% /
/dev/hd2           577536     75124   87%    27909    20% /usr
/dev/hd9var         16384     13516   18%      494    13% /var
/dev/hd3            16384     12696   23%       37     1% /tmp
/dev/hd1            16384     15268    7%       66     2% /home
AFS              72000000  72000000    0%        0     0% /afs
```

Figure 47. Displaying the /afs File System

3. Examine the top directory level of the /afs tree as shown in the following session:

```
# cd /afs
# ls -la
total 20
drwxrwxrwx   2 root     system      2048 Mar 31 20:27 .
drwxr-xr-x  27 bin      bin         1024 Apr 03 01:29 ..
drwxrwxrwx   2 root     system      2048 Mar 31 20:27 .sushi.itso.ral.ibm.com
drwxrwxrwx   2 root     system      2048 Mar 31 20:27 sushi.itso.ral.ibm.com
```

Figure 48. Examining the /afs Tree

To configure the AFS initialization file to start AFS at system reboot, you need to follow the procedures in 3.7, “Starting AFS at System Reboot on AIX Systems” on page 94. After completing the AFS initialization file configuration, reboot your machine and verify that all AFS processes are running normally.

3.7 Starting AFS at System Reboot on AIX Systems

At system reboot, the system starts AFS by executing the AFS initialization file called /etc/rc.afs. This file is a script containing commands to:

1. Load the AFS extensions into the operating system kernel
2. Start the BOS Server process for an AFS Server
3. Start the AFS Cache Manager for an AFS Client

This AFS initialization file is invoked from the /etc/inittab file. In the /etc/inittab file, you need to create the following AFS entry after the NFS entry:

```
rcafs:2:wait:/etc/rc.afs > /dev/console 2>&1 # Start AFS Daemons
```

Figure 49 on page 94 illustrates the modification to the /etc/inittab file.

```
...
rcnfs:2:wait:/etc/rc.nfs > /dev/console 2>&1 # Start NFS Daemons
rcafs:2:wait:/etc/rc.afs > /dev/console 2>&1 # Start AFS Daemons
...
```

Figure 49. Creating the AFS Entry in the /etc/inittab File

You want to make the /etc/rc.afs file executable by entering the command:

```
chmod +x /etc/rc.afs
```

This AFS initialization file can be used on any AIX machine running AFS server and/or client processes.

WebSphere Performance Pack provides a similar AFS initialization file, rc.dkload.afs, in the directory /usr/lpp/WSPP/afs. This file contains control characters that make the file itself unexecutable in a UNIX environment. In fact, for the UNIX installation, this file comes in DOS text format, and should be converted to UNIX text format. To do this, you should enter the following commands:

```
tr -d '\015' < rc.dkload.afs > rc.dkload.afs2
mv rc.dkload.afs2 rc.dkload.afs
```

The next several output screens contain the listing of the AFS initialization file for AIX systems, /etc/rc.afs:

```
#!/bin/sh
#
# rc.afs: rc script for AFS on AIX platforms
#
# Install this script as /etc/rc.afs
# then make an entry in /etc/inittab after the NFS entry:
# rcnfs:2:wait:/etc/rc.nfs > /dev/console 2>&1 # Start NFS Daemons
# rcafs:2:wait:/etc/rc.afs > /dev/console 2>&1 # Start AFS Daemons
#

# Choose one depending on how much usage this client gets
SMALL="-stat 300 -dcache 100 -daemons 2 -volumes 50"
MEDIUM="-stat 2000 -dcache 800 -daemons 3 -volumes 70"
LARGE="-stat 2800 -dcache 2400 -daemons 5 -volumes 128"
OPTIONS=$MEDIUM

# Choose one depending on how you want NFS requests handled by AFS
# Use "none" if this machine won't be an AFS/NFS Translator.
# To make the machine an AFS/NFS Translator, use iauth for:
#   AIX 4.1.x  where x >= 5
#   AIX 4.2.x  where x >= 1
#   AIX 4.3.x  for any x
# otherwise use "nfs"
NFS_NONE="none"
NFS_NFS="nfs"
NFS_IAUTH="iauth"
NFS=$NFS_NONE

if [ "$NFS" = "nfs" ]; then
  ExportExt=""
  AFSExt=".trans"
  RMTSYS="-rmtsys"
elif [ "$NFS" = "iauth" ]; then
  ExportExt=".nonfs"
  AFSExt=".iauth"
  RMTSYS="-rmtsys"
else
  ExportExt=".nonfs"
  AFSExt=""
  RMTSYS=""
fi

# Load AFS into the kernel
cd /usr/vice/etc/dkload
/usr/vice/etc/dkload/cfgexport -a
/usr/vice/etc/dkload/export.ext${ExportExt}
/usr/vice/etc/dkload/cfgafs -a /usr/vice/etc/dkload/afs.ext${AFSExt}
if [ $? -ne 0 ]; then
  /bin/echo "Unable to load AFS extensions into kernel. Not starting
client."
  exit 1
fi
```

Figure 50. (Part 1 of 2). AFS Initialization File (/etc/rc) for AIX System

```
# Start bosserver
if [ -x /usr/afs/bin/bosserver -a -f /usr/afs/etc/ThisCell \
  -a -f /usr/afs/etc/CellServDB -a -f /usr/afs/etc/KeyFile \
  -a -f /usr/afs/etc/UserList -a -f /usr/afs/etc/License ]; then
  echo 'Starting bosserver' > /dev/console
  /usr/afs/bin/bosserver &
  OPTIONS="$OPTIONS -nosettime"
fi


#
# Check that all of the client configuration files exist
#
for file in afsd cacheinfo ThisCell CellServDB; do
  if [ ! -f /usr/vice/etc/${file} ]; then
    /bin/echo "/usr/vice/etc/${file} does not exist. Not starting AFS
client."
    exit 2
  fi
done


#
# Check that the root directory for AFS (/afs)
# and the cache directory (/usr/vice/cache) both exist
#
for dir in `awk -F: '{print $1, $2}' /usr/vice/etc/cacheinfo`; do
  if [ ! -d ${dir} ]; then
    echo "${dir} does not exist. Not starting AFS client."
    exit 3
  fi
done


# Start afsd
/usr/vice/etc/afsd $OPTIONS $RMTSYS


# Start AFS inetd services
if [ -x /etc/inetd.afs -a -f /etc/inetd.conf.afs ]; then
    /etc/inetd.afs /etc/inetd.conf.afs
fi
```

Figure 51. (Part 2 of 2). AFS Initialization File (/etc/rc) for AIX System

Chapter 4. AFS Installation and Configuration on Solaris

This section summarizes the installation prerequisites for the AFS component of WebSphere Performance Pack on Solaris systems. It also describes the installation and configuration procedures, including the first AFS machine, an additional AFS File Server, and an AFS Client. Since these procedures are very similar to the ones for AIX systems, we will refer you to the AIX configuration section as necessary.

4.1 Solaris System Requirements

To prepare your Solaris system for the installation of the File Sharing (AFS) component of WebSphere Performance Pack, you need:

- Any SPARC-compatible server or workstation machine
- Sun Solaris Version 2.6
- A minimum of 64 MB of RAM
- A minimum of 128 MB of swap space

  For performance reasons, we recommend that you allocate twice the amount of RAM for swap space.
- An operational machine on the network
- A minimum of 20 MB of available disk space on the /usr file system to install the AFS files
- A minimum of 40 MB of available disk space for the AFS disk cache

  Although AFS memory cache is available, memory caching does not provide better performance than disk caching. We recommend that you use disk caching for an AFS client.
- At least one UNIX file system (UFS) dedicated to storing AFS data for each AFS File Server.
- Java Runtime Environment (JRE) 1.1.7 or higher
- Administrative privileges

  You need to have root access to the Solaris machine in order to install, configure, and start the AFS server and client.

4.1.1 Paging Space Allocation

You need to allocate adequate paging (swap) space on your AFS machines. On Solaris systems, the command:

```
prtconf | grep Memory
```

displays the amount of real memory in MB, and the command:

```
swap -s
```

displays the amount of paging space. For performance reasons, the amount of paging space must be at least twice the amount of memory on your machine.

The following session demonstrates the usage of these two commands:

```
# prtconf | grep Memory
Memory size: 64 Megabytes

# swap -s
total: 33848k bytes allocated + 9936k reserved = 43784k used, 195116k available
```

Figure 52. Displaying the Amount of Memory and Paging Space on a Solaris System

To create additional paging space, use the commands `mkfile` and

```
swap -a
```

4.1.2 Name Resolution Service

We recommend that you use a name service in your environment, such as the Domain Name System (DNS), to resolve the host names of machines. It is difficult to maintain accurately host name resolution of machines with local host tables.

On Solaris systems, the `nslookup` command resolves a host name into an Internet address or an Internet address into a host name. By using the `nslookup` command, you can verify that both the forward and reverse name resolution services are functioning correctly.

The following session illustrates the usage of the `nslookup` command for a machine with the host name of sun, the IP address of 9.24.105.31, and the DNS domain name of itso.ral.ibm.com.

```
# hostname
sun
# nslookup sun
Server:  rsserver.itso.ral.ibm.com
Address:  9.24.104.108

Name:    sun.itso.ral.ibm.com
Address:  9.24.105.31

# nslookup 9.24.105.31
Server:  rsserver.itso.ral.ibm.com
Address:  9.24.104.108

Name:    sun.itso.ral.ibm.com
Address:  9.24.105.31
```

Figure 53. Verifying Forward and Reverse Host Name Resolution on a Solaris System

4.1.3 Network Connection

The Solaris machine has to be operational on the network. For each machine where you are planning to install and configure AFS, verify network configuration and integrity.

On UNIX systems, such as IBM AIX and Sun Solaris, the following commands can assist you in determining your network status:

- The command:

```
netstat -nr
```

shows the network routing table.

- The command:

  ```
  ping 9.24.105.1
  ```

  verifies the network connection between the local machine and the default gateway, 9.24.105.1, by sending an Internet Control Message Protocol (ICMP) echo request to the default gateway machine network interface.

- The command:

  ```
  netstat -ni
  ```

  displays the statistics of the network interface, and should indicate that there are no packet errors or collisions.

Figure 54 on page 99 demonstrates the usage of these commands from a Solaris machine with the IP address 9.24.105.10. Since these commands are similar to the ones used on AIX systems, for a detailed description of the session, refer to 3.1.2.1, “Network Connection” on page 55 for AIX systems.

```
# netstat -nr

Routing Table:
  Destination           Gateway           Flags  Ref   Use   Interface
-------------------- -------------------- ----- ----- ------ ---------
9.24.105.0           9.24.105.10           U        2     15  le0
default              9.24.105.1            UG       0     60
127.0.0.1            127.0.0.1             UH       0     22  lo0

# ping -snr 9.24.105.1 1024 5
1032 bytes from 9.24.105.1: icmp_seq=0. time=5. ms
1032 bytes from 9.24.105.1: icmp_seq=1. time=3. ms
1032 bytes from 9.24.105.1: icmp_seq=2. time=4. ms
1032 bytes from 9.24.105.1: icmp_seq=3. time=3. ms
1032 bytes from 9.24.105.1: icmp_seq=4. time=4. ms

----9.24.105.1 PING Statistics----
5 packets transmitted, 5 packets received, 0% packet loss
round-trip (ms)  min/avg/max = 3/3/5

# netstat -ni
Name  Mtu  Net/Dest      Address        Ipkts  Ierrs Opkts  Oerrs Collis Queue
lo0   8232 127.0.0.0     127.0.0.1      187    0     187    0     0      0
le0   1500 9.24.105.0    9.24.105.10    13922  384   3622   2     0      0
```

Figure 54. Verifying Network Connection on a Solaris System

4.1.4 Time Synchronization

Before configuring additional AFS server or client machines, manually verify that the clock on the new machine is relatively synchronized with the clock on the first AFS machine. If the clock is off on the new machine, correct it before continuing. After you configure additional AFS server or client machines, AFS uses the Network Time Protocol (NTP) to synchronize time amongst a cell’s server machines. An AFS client machine synchronizes its clock with the clock on a server machine chosen at random.

4.1.5 JRE Version

The WebSphere Performance Pack installation process uses the Java InstallShield's setup class. You need to ensure that the JRE is installed. At the time of this writing, the recommended JRE version for Solaris is 1.1.7.

To quickly determine the Java version, enter:

```
java -version
```

To ensure that the JRE has been correctly installed on your Solaris system, enter the following commands:

```
pkginfo | grep -i java
pkginfo -l SUNWjvrt
```

The following session displays the state of the JRE on a Solaris machine:

```
# pkginfo | grep -i java
system      SUNWdthj       HotJava Browser for Solaris
system      SUNWjvdem      JavaVM demo programs
system      SUNWjvdev      JavaVM developers package, includes javac, javah,
                           and javap
system      SUNWjvjit      Java JIT compiler
system      SUNWjvman      JavaVM man pages
system      SUNWjvrt       JavaVM run time environment

# pkginfo -l SUNWjvrt
   PKGINST:  SUNWjvrt
      NAME:  JavaVM run time environment
  CATEGORY:  system
      ARCH:  sparc
   VERSION:  1.1.7,REV=1999.01.15.10.50
   BASEDIR:  /usr
    VENDOR:  Sun Microsystems, Inc.
      DESC:  JavaVM run time environment, includes java, appletviewer,
             and classes.zip
    PSTAMP:  javavm990205151241
  INSTDATE:  Mar 03 1999 15:26
   HOTLINE:  Please contact your local service provider
    STATUS:  completely installed
     FILES:     67 installed pathnames
                 8 shared pathnames
                 9 directories
                19 executables
             20688 blocks used (approx)
```

If the output of the `pkginfo` command does not indicate level 1.1.7 or the status of this software product is not `completely installed`, you will have to install JRE V1.1.7 on this machine. The CD distribution of WebSphere Performance Pack contains the latest version of the Java Development Kit for Solaris in the directory sun/jdk. Actually, you would not need the full JDK, but only its subset known as JRE, which contains just the Java Virtual Machine (JVM), the Java platform core classes, and supporting files. In other words, the JRE is the smallest set of executables and files that constitute the standard Java platform and it contains only the run-time part of the JDK: no compiler, no debugger, no tools.

4.1.6 AFS Extensions to the Solaris Kernel

On Solaris systems, the AFS extensions are dynamically loaded into the kernel by using the `modload` program. To prepare the system to make AFS system calls, you need to perform the following steps:

1. Create an `afs` entry in slot 105, the slot prior to the `nfs` entry, of the /etc/name_to_sysnum file, as shown in Figure 55 on page 101.

```
nosys        0
rexit        1
fork         2
...
statvfs    103
fstatvfs   104
afs        105
nfs        106
...
```

Figure 55. Creating the AFS Entry in the /etc/name_to_sysnum File

2. Reboot the machine by entering the following command:

```
/usr/sbin/shutdown -y -g0 -i6
```

4.1.7 Disk Space Requirements on Solaris Systems

For a detailed discussion on disk space requirements, please refer to 3.1.5, “Disk Space Requirements on UNIX Systems” on page 57 for AIX systems. These requirements also apply to Solaris systems. In addition, it is not possible to dynamically increase the size of a UFS, such as the /usr file system, on Solaris systems.

4.2 Configuring AFS Partitions on Solaris Systems

Every AFS file server must have a disk partition dedicated to be an AFS partition to store AFS volumes. By convention, the first AFS partition created on an AFS file server machine is called /vicepa, the second partition /vicepb, etc. On a Solaris system, an AFS partition has to be created on a UFS.

To create a UFS to be used as an AFS partition, you need to perform the following operations:

- Create a disk partition
- Construct a new UFS
- Mount this file system
- Create an entry for the new UFS file system in the file system table
- Modify the /sbin/mountall command to recognize the AFS file system

In our environment, the machine is a Sun SPARCstation-20, with 64 MB of RAM and two 1.05 GB hard disks, running Solaris 2.6. The operating system is installed on the default disk with SCSI ID 3. We are creating a disk partition on the second disk, with SCSI ID 1, as an AFS partition.

In the following example, the Solaris machine initially has the following file systems:

```
# df -k
Filesystem            kbytes    used   avail capacity  Mounted on
/dev/dsk/c0t3d0s0     154039   67358   86527    44%    /
/dev/dsk/c0t3d0s6     623246  518096  104458    84%    /usr
/proc                      0       0       0     0%    /proc
fd                         0       0       0     0%    /dev/fd
swap                  193048      16  193032     1%    /tmp
```

The following session demonstrates the steps to create a 64 MB AFS partition:

1. To create a disk partition, use the `format` command to select a disk to be partitioned, as shown in the following session:

```
# format
Searching for disks...done


AVAILABLE DISK SELECTIONS:
       0. c0t1d0 <SUN1.05 cyl 2036 alt 2 hd 14 sec 72>
          /iommu@f,e0000000/sbus@f,e0001000/espdma@f,400000/esp@f,800000/sd@1,0
       1. c0t3d0 <SUN1.05 cyl 2036 alt 2 hd 14 sec 72>
          /iommu@f,e0000000/sbus@f,e0001000/espdma@f,400000/esp@f,800000/sd@3,0
Specify disk (enter its number): 0
selecting c0t1d0
[disk formatted]
Warning: Current Disk has mounted partitions.


FORMAT MENU:
        disk       - select a disk
        type       - select (define) a disk type
        partition  - select (define) a partition table
        current    - describe the current disk
        format     - format and analyze the disk
        repair     - repair a defective sector
        label      - write label to the disk
        analyze    - surface analysis
        defect     - defect list management
        backup     - search for backup labels
        verify     - read and display labels
        save       - save new disk/partition definitions
        inquiry    - show vendor, product and revision
        volname    - set 8-character volume name
        !<cmd>     - execute <cmd>, then return
        quit
format> partition


PARTITION MENU:
        0      - change `0' partition
        1      - change `1' partition
        2      - change `2' partition
        3      - change `3' partition
        4      - change `4' partition
        5      - change `5' partition
        6      - change `6' partition
        7      - change `7' partition
        select - select a predefined table
        modify - modify a predefined partition table
        name   - name the current table
        print  - display the current table
        label  - write partition map and label to the disk
        !<cmd> - execute <cmd>, then return
        quit
format>
```

Figure 56. (Part 1 of 3). Creating a Disk Partition on a Solaris System

```
partition> print
Current partition table (original):
Total disk cylinders available: 2036 + 2 (reserved cylinders)

Part      Tag    Flag     Cylinders        Size            Blocks
  0       root    wm       0 -  755      372.09MB    (756/0/0)   762048
  1       swap    wu     756 - 1011      126.00MB    (256/0/0)   258048
  2     backup    wm       0 - 2035     1002.09MB    (2036/0/0) 2052288
  3 unassigned    wm       0               0         (0/0/0)          0
  4 unassigned    wm       0               0         (0/0/0)          0
  5 unassigned    wm       0               0         (0/0/0)          0
  6 unassigned    wm       0               0         (0/0/0)          0
  7 unassigned    wm       0               0         (0/0/0)          0

partition> 3
Part      Tag    Flag     Cylinders        Size            Blocks
  3 unassigned    wm       0               0         (0/0/0)          0

Enter partition id tag[unassigned]: ?
Expecting one of the following: (abbreviations ok):
        unassigned    boot          root          swap
        usr           backup        stand         var
        home          alternates


Enter partition id tag[unassigned]: stand
Enter partition permission flags[wm]: wm
Enter new starting cyl[0]: 1012
Enter partition size[0b, 0c, 0.00mb, 0.00gb]: 64mb
partition> print
Current partition table (unnamed):
Total disk cylinders available: 2036 + 2 (reserved cylinders)

Part      Tag    Flag     Cylinders        Size            Blocks
  0       root    wm       0 -  755      372.09MB    (756/0/0)   762048
  1       swap    wu     756 - 1011      126.00MB    (256/0/0)   258048
  2     backup    wm       0 - 2035     1002.09MB    (2036/0/0) 2052288
  3      stand    wm    1012 - 1142       64.48MB    (131/0/0)   132048
  4 unassigned    wm       0               0         (0/0/0)          0
  5 unassigned    wm       0               0         (0/0/0)          0
  6 unassigned    wm       0               0         (0/0/0)          0
  7 unassigned    wm       0               0         (0/0/0)          0

partition> label
Ready to label disk, continue? y

partition> quit
```

Figure 57. (Part 2 of 3). Creating a Disk Partition on a Solaris System

```
FORMAT MENU:
        disk       - select a disk
        type       - select (define) a disk type
        partition  - select (define) a partition table
        current    - describe the current disk
        format     - format and analyze the disk
        repair     - repair a defective sector
        label      - write label to the disk
        analyze    - surface analysis
        defect     - defect list management
        backup     - search for backup labels
        verify     - read and display labels
        save       - save new disk/partition definitions
        inquiry    - show vendor, product and revision
        volname    - set 8-character volume name
        !<cmd>     - execute <cmd>, then return
        quit
format> quit
#
```

Figure 58. (Part 3 of 3). Creating a Disk Partition on a Solaris System

2. To construct a new UFS file system, use the `newfs` command, as shown in the following session:

```
# newfs -v /dev/rdsk/c0t1d0s3
newfs: construct a new file system /dev/rdsk/c0t1d0s3: (y/n)? y
mkfs -F ufs /dev/rdsk/c0t1d0s3 132048 72 14 8192 1024 16 10 90 2048 t 0 -1 8 128
/dev/rdsk/c0t1d0s3:     132048 sectors in 131 cylinders of 14 tracks, 72 sectors
        64.5MB in 9 cyl groups (16 c/g, 7.88MB/g, 3776 i/g)
super-block backups (for fsck -F ufs -o b=#) at:
 32, 16240, 32448, 48656, 64864, 81072, 97280, 113488, 129696,
```

Figure 59. Creating a UFS File System on a Solaris System

3. To mount this new UFS file system, use the `mount` command, as shown in the following session:

```
# mkdir /vicepa
# mount /dev/dsk/c0t1d0s3 /vicepa
# df -k
Filesystem            kbytes    used   avail capacity  Mounted on
/dev/dsk/c0t3d0s0     154039   67358   86527    44%    /
/dev/dsk/c0t3d0s6     623246  518096  104458    84%    /usr
/proc                      0       0       0     0%    /proc
fd                         0       0       0     0%    /dev/fd
/dev/dsk/c0t1d0s3      61615       9   61545     1%    /vicepa
swap                  193048      16  193032     1%    /tmp
```

Figure 60. Mounting a UFS File System on a Solaris System

4. To auto-mount this new UFS file system at system boot, create a new entry in the file system table, /etc/vfstab, as shown in Figure 61 on page 106.

```
#device             device              mount      FS     fsck   mount      mount
#to mount           to fsck             point      type   pass   at boot    options

fd                  -                   /dev/fd    fd      -       no          -
/proc               -                   /proc      proc    -       no          -
/dev/dsk/c0t3d0s1   -                   -          swap    -       no          -
/dev/dsk/c0t3d0s0   /dev/rdsk/c0t3d0s0  /          ufs     1       no          -
/dev/dsk/c0t3d0s6   /dev/rdsk/c0t3d0s6  /usr       ufs     1       no          -
/dev/dsk/c0t1d0s3   /dev/rdsk/c0t1d0s3  /vicepa    afs     3       yes         -
swap                -                   /tmp       tmpfs   -       yes         -
```

Figure 61. Adding an AFS Partition Entry in the File System Table /etc/vfstab

5. Append the following line to the file /etc/dfs/fstypes:

   `afs AFS Utilities`

6. To inform the system of the new AFS file system type, modify the `/sbin/mountall` command, as shown in Figure 62 on page 107.

```
...
checkfs() {
        /usr/sbin/fsck -F $2 -m $1  >/dev/null 2>&1

        if [ $? -ne 0 ]
        then
                # Determine fsck options by file system type
                case "$2" in
                ufs)    foptions="-o p"
                        ;;
                afs)    foptions="-o p"
                        ;;
                s5)     foptions="-y -t /var/tmp/tmp$$ -D"
                        ;;
                *)      foptions="-y"
                        ;;
                esac

                echo "The "$3" file system ("$1") is being checked."
                /usr/sbin/fsck -F $2 ${foptions} $1

                ...
        fi
}

...

#
# Read FSTAB, fsck'ing appropriate filesystems:
#
exec < $FSTAB
while  read special fsckdev mountp fstype fsckpass automnt mntopts
do
        ...
        #
        # For fsck purposes, we make a distinction between ufs and other
        # file systems
        #
        if [ "$fstype" = "ufs" -o "$fstype" = "afs" ]; then
                ufs_fscklist="$ufs_fscklist $fsckdev"
                saveentry $fstype "$OPTIONS" $special $mountp
                continue
        fi
        ...
done
...
```

Figure 62. Modifications to the /sbin/mountall Command on a Solaris System

4.3 Installing AFS Files on Solaris Systems

This section describes the procedure to install the necessary AFS files to configure a Solaris machine to be an AFS machine, including the first AFS machine, an additional AFS File Server, and an AFS client.

To install IBM AFS Enterprise File System, the File Sharing component of WebSphere Performance Pack, on Solaris systems, perform the following tasks:

1. Login as root.

2. Insert the IBM Websphere Performance Pack CD-ROM in the CD-ROM drive. The system automatically mounts the WebSphere Performance Pack CD-ROM as /cdrom/websphere.
3. From a command prompt, enter the following command:

   ```
   cd /cdrom/websphere/sun
   ```

4. Start the Java InstallShield installation program by entering:

   ```
   java setup
   ```

Since the installation procedure for Solaris systems is the same as the one for AIX systems, please refer to Step 5 on page 63 through Step 10 on page 66, in 3.3, "Installing AFS Files on AIX Systems" on page 62 to proceed with the installation.

4.3.1 WebSphere AFS Directories

For a detailed discussion on the WebSphere AFS directories, please refer to 2.9.3, "Local Directories" on page 42, 2.10.3, "Local Directories" on page 44 and 3.3.1, "WebSphere AFS Directories" on page 66.

On Solaris systems, the installation program creates a new AFS library directory, /usr/lib/fs/afs, for the Solaris file system checker program, fsck. This program performs file system integrity checks. Notice, however, that the Solaris fsck program discards files from directories that are not part of the UFS file system, including the /vicep*x* directories. These directories are non-standard UNIX directories, containing information that make up AFS volumes. For this reason, the WebSphere AFS installation provides an AFS-aware fsck that correctly handles both AFS and UFS file systems.

Solaris fsck and AFS Partitions

Never run the standard Solaris file system checker program on an AFS File Server. It discards AFS volumes.

The following session illustrates the contents of the new AFS library directory for the AFS-aware fsck program, /usr/lib/fs/afs/fsck.

```
# cd /usr/lib/fs/afs
# ls -l
total 364
lrwxrwxrwx   1 root     other         20 Mar  8 18:20 clri -> /usr/lib/fs/ufs/clri
lrwxrwxrwx   1 root     other         18 Mar  8 18:20 df -> /usr/lib/fs/ufs/df
lrwxrwxrwx   1 root     other         23 Mar  8 18:20 edquota -> /usr/lib/fs/ufs/edquota
lrwxrwxrwx   1 root     other         18 Mar  8 18:20 ff -> /usr/lib/fs/ufs/ff
-rwxr--r--   1 root     other     151512 Apr  1 20:12 fsck
lrwxrwxrwx   1 root     other         20 Mar  8 18:20 fsdb -> /usr/lib/fs/ufs/fsdb
lrwxrwxrwx   1 root     other         23 Mar  8 18:20 fsirand -> /usr/lib/fs/ufs/fsirand
lrwxrwxrwx   1 root     other         21 Mar  8 18:20 fstyp -> /usr/lib/fs/ufs/fstyp
lrwxrwxrwx   1 root     other         23 Mar  8 18:20 labelit -> /usr/lib/fs/ufs/labelit
lrwxrwxrwx   1 root     other         22 Mar  8 18:20 lockfs -> /usr/lib/fs/ufs/lockfs
lrwxrwxrwx   1 root     other         20 Mar  8 18:20 mkfs -> /usr/lib/fs/ufs/mkfs
lrwxrwxrwx   1 root     other         21 Mar  8 18:20 mount -> /usr/lib/fs/ufs/mount
lrwxrwxrwx   1 root     other         22 Mar  8 18:20 ncheck -> /usr/lib/fs/ufs/ncheck
lrwxrwxrwx   1 root     other         21 Mar  8 18:20 newfs -> /usr/lib/fs/ufs/newfs
lrwxrwxrwx   1 root     other         23 Mar  8 18:20 qoutaon -> /usr/lib/fs/ufs/qoutaon
lrwxrwxrwx   1 root     other         20 Mar  8 18:20 quot -> /usr/lib/fs/ufs/quot
lrwxrwxrwx   1 root     other         21 Mar  8 18:20 quota -> /usr/lib/fs/ufs/quota
lrwxrwxrwx   1 root     other         24 Mar  8 18:20 quotaoff -> /usr/lib/fs/ufs/quotaoff
lrwxrwxrwx   1 root     other         24 Mar  8 18:20 repquota -> /usr/lib/fs/ufs/repquota
lrwxrwxrwx   1 root     other         22 Mar  8 18:20 tunefs -> /usr/lib/fs/ufs/tunefs
lrwxrwxrwx   1 root     other         23 Mar  8 18:20 ufsdump -> /usr/lib/fs/ufs/ufsdump
lrwxrwxrwx   1 root     other         26 Mar  8 18:20 ufsrestore -> /usr/lib/fs/ufs/ufsrestore
lrwxrwxrwx   1 root     other         23 Mar  8 18:20 volcopy -> /usr/lib/fs/ufs/volcopy
```

This AFS library directory and the modifications to the /sbin/mountall program described in 4.2, “Configuring AFS Partitions on Solaris Systems” on page 101, allow the system to check correctly both AFS and UFS file systems.

4.3.2 WebSphere AFS Configuration Script

For a detailed discussion on the WebSphere AFS configuration script, refer to 3.3.2, “WebSphere AFS Configuration Script” on page 67.

4.4 Configuring the First AFS Machine on Solaris Systems

Before configuring a Solaris machine to be the first AFS machine, you must complete the following tasks:

- Configure the system to recognize AFS system calls, as demonstrated in 4.1.6, “AFS Extensions to the Solaris Kernel” on page 101.
- Create at least one AFS partition to store AFS volumes, as illustrated in 4.2, “Configuring AFS Partitions on Solaris Systems” on page 101.
- Install the necessary AFS files from the WebSphere Performance Pack CD-ROM distribution to the machine local disk, as shown in 4.3, “Installing AFS Files on Solaris Systems” on page 107.

4.4.1 Using the WebSphere AFS Configuration Script

Since the configuration procedure for Solaris systems is the same as the one for AIX systems, refer to 3.4.1, “Using the WebSphere AFS Configuration Script” on page 68, to proceed with the configuration.

Since authorization checking is disabled during the configuration process, you may safely ignore error messages similar to:

```
bos: no such entry (getting tickets)
bos: running unauthenticated
libprot: no such entry Could not get afs tokens, running unauthenticated
vsu_ClientInit: Could not get afs tokens, running unauthenticated
```

In summary, the AFS configuration script prompts you for the name of the AFS partition, the name of the AFS cell, the afs and admin password, and the size of the AFS disk cache in 1 KB blocks.

In our environment, the name of our first AFS partition is /vicepa, the name of our AFS cell is sashimi.itso.ral.ibm.com, and the size of our AFS disk cache is 40,000 1 KB blocks.

The next several text figures contain the transcript of our configuration session of the first AFS machine on a Solaris system called sun2:

```
# cd /opt/WSPP
# cd afs
# ./afsConfigure.sh
Please choose from the following options
1. Configure this machine as your SCM.
2. Configure this machine as a File Server.
3. Configure this machine as an AFS Client.
4. Quit the program.
Enter choice: 1
This machine is your SCM


Setting up first AFS Server
Loading SUN kernel extensions
nfsd was not found on this machine

Enter the name of your partition: /vicepa
Creating license file

Starting the bosserver

bosserver started successfully

Enter the name you want your cell to be: sashimi.itso.ral.ibm.com

Setting the cell name to sashimi.itso.ral.ibm.com
bos: no such entry (getting tickets)
bos: running unauthenticated

Starting the Authentication Server
bos: no such entry (getting tickets)
bos: running unauthenticated

Enter your AFS and admin user password: XXXXXXXXXX
AFSpassword is XXXXXXXXXX

Creating AFS and ADMIN users...
bos: no such entry (getting tickets)
bos: running unauthenticated
bos: no such entry (getting tickets)
bos: running unauthenticated
bos: no such entry (getting tickets)
bos: running unauthenticated

Starting the Protection Server
bos: no such entry (getting tickets)
bos: running unauthenticated

Adding admin to system:administrators group
libprot: no such entry Could not get afs tokens, running unauthenticated.
User admin has id 1
libprot: no such entry Could not get afs tokens, running unauthenticated.
```

Figure 63. (Part 1 of 3). Configuration of the First AFS Machine on a Solaris System

```
Starting the Volume Location Server
bos: no such entry (getting tickets)
bos: running unauthenticated

Starting the Backup Server
bos: no such entry (getting tickets)
bos: running unauthenticated

Starting the File Server, Volume Server and Salvager
bos: no such entry (getting tickets)
bos: running unauthenticated

Creating the cell's root volume root.afs
vsu_ClientInit: Could not get afs tokens, running unauthenticated.
Volume 536870912 created on partition /vicepa of sun2

Starting the Server Portion of the Update Server
bos: no such entry (getting tickets)
bos: running unauthenticated

Starting the Controller for NTPD
bos: no such entry (getting tickets)
bos: running unauthenticated

Enter the number of 1 KB blocks to devote to the cache: 40000
 U.S. Government Restricted Rights.
====================================
If you are licensing the Software on behalf of the U.S.
Government ("Government"), the following provisions apply
to you. If the Software is supplied to the Department of
Defense ("DoD"), it is classified as "Commercial Computer
Software" under paragraph 252.227-7014 of the DoD Supplement
to the Federal Acquisition Regulations ("DFARS") (or
any successor regulations) and the Government is acquiring
only the license rights granted herein (the license rights
customarily provided to non-Government users).  If the
Software is supplied to any unit or agency of the Government
other than DoD, it is classified as "Restricted Computer
Software" and the Government's rights in the Software are
defined in paragraph 52.227-19 of the Federal Acquisition
Regulations ("FAR") (or any successor regulations) or, in
the case of NASA, in paragraph 18.52.227-86 of the NASA
Supplement in the FAR (or any successor regulations).

afsd: My home cell is 'sashimi.itso.ral.ibm.com'
afsd: Creating '/usr/vice/etc/AFSLog'
CreateCacheFile: Creating cache file '/usr/vice/etc/AFSLog'
SScall(105, 28)=0 afsd: Forking rx callback listener.
afsd: Forking AFS daemon.
afsd: Forking Check Server Daemon
afsd: Forking 3 background daemons.
afsd: Forking rxevent daemon.
```

Figure 64. (Part 2 of 3). Configuration of the First AFS Machine on a Solaris System

```
SScall(105, 28)=0 afsd: Calling AFSOP_CACHEINIT: 2000 stat cache entries,
4000 optimum cache files, 40000 blocks in the cache, flags = 0x0, dcache
entries 800
SScall(105, 28)=0 afsd: Sweeping workstation's AFS cache directory.
SweepAFSCache: Creating '/usr/vice/cache/CacheItems'
CreateCacheFile: Creating cache file '/usr/vice/cache/CacheItems'
SweepAFSCache: Creating '/usr/vice/cache/VolumeItems'
CreateCacheFile: Creating cache file '/usr/vice/cache/VolumeItems'
SweepAFSCache: Creating '/usr/vice/cache/V0'
CreateCacheFile: Creating cache file '/usr/vice/cache/V0'
SweepAFSCache: Creating '/usr/vice/cache/V1'
CreateCacheFile: Creating cache file '/usr/vice/cache/V1'
...
SweepAFSCache: Creating '/usr/vice/cache/V3998'
CreateCacheFile: Creating cache file '/usr/vice/cache/V3998'
SweepAFSCache: Creating '/usr/vice/cache/V3999'
CreateCacheFile: Creating cache file '/usr/vice/cache/V3999'
afsd: 0 out of 4000 data cache files found in sweep 1.
afsd: 4000 out of 4000 data cache files found in sweep 2.
...
afsd: All AFS daemons started.
afsd: Forking trunc-cache daemon.
afsd: Mounting the AFS root on '/afs', flags: 4.

Tokens held by the Cache Manager:

User's (AFS ID 1) tokens for afs@sashimi.itso.ral.ibm.com [Expires Mar 31
19:38]
   --End of list--
All volumeID/name mappings checked.
Volume 536870915 created on partition /vicepa of sun2
Added replication site sun2 /vicepa for volume root.afs
Added replication site sun2 /vicepa for volume root.cell
Released volume root.afs successfully
Released volume root.cell successfully
All volumeID/name mappings checked.
Volume status for vid = 536870913 named root.afs.readonly
Current disk quota is 5000
Current blocks used are 4
The partition has 61537 blocks available out of 61554

Volume status for vid = 536870916 named root.cell.readonly
Current disk quota is 5000
Current blocks used are 2
The partition has 61537 blocks available out of 61554
```

Figure 65. (Part 3 of 3). Configuration of the First AFS Machine on a Solaris System

4.4.2 Verifying the First AFS Machine Configuration

To verify that the configuration of the first AFS machine completed successfully, refer to the procedure in 3.4.2, “Verifying the First AFS Machine Configuration” on page 75.

Recall that the first AFS machine is not only an AFS server, but also an AFS client. To verify that the AFS client is functioning properly, you need to perform the tasks in 4.6.2, “Verifying the AFS Client Configuration” on page 119.

To configure the AFS initialization file to start AFS at system reboot, you need to follow the procedures in 4.7, “Starting AFS at System Reboot on Solaris Systems” on page 119. After completing the AFS initialization file configuration, reboot your machine and verify that all AFS processes are running normally.

4.5 Configuring an AFS File Server Machine on Solaris Systems

To configure a Solaris machine to be an AFS File Server machine in an existing AFS cell, you must have already completed the following tasks:

- Configure the system to recognize AFS system calls, as demonstrated in 4.1.6, “AFS Extensions to the Solaris Kernel” on page 101.
- Create at least one AFS partition to store AFS volumes, as illustrated in 4.2, “Configuring AFS Partitions on Solaris Systems” on page 101.
- Install the necessary AFS files from the WebSphere Performance Pack CD-ROM distribution to the machine’s local disk, as shown in 4.3, “Installing AFS Files on Solaris Systems” on page 107.

4.5.1 Using the WebSphere AFS Configuration Script

Since the configuration procedure for Solaris systems is the same as the one for AIX systems, refer to 3.5.1, “Using the WebSphere AFS Configuration Script” on page 79, to proceed with the configuration.

Since authorization checking is disabled during the configuration process, you may safely ignore error messages similar to:

```
bos: no such entry (getting tickets)
bos: running unauthenticated
libprot: no such entry Could not get afs tokens, running unauthenticated
vsu_ClientInit: Could not get afs tokens, running unauthenticated
```

In summary, the configuration script prompts you for the name of the AFS partition, the name of an existing AFS cell, the host name of your first AFS machine, and the size of the AFS disk cache in 1 KB blocks.

In our environment, the name of the AFS partition on this AFS File Server is /vicepa, the name of our existing AFS cell is sushi.itso.ral.ibm.com, the host name of our first AFS machine is rs60002.itso.ral.ibm.com, and the size of our AFS disk cache is 40,000 1 KB blocks.

The next several text figures contain the transcript of our configuration session of the AFS File Server on a Solaris system called sun:

```
# cd /usr/afs/etc
(Copy or FTP the contents of /usr/afs/etc from your first AFS machine)
# cd /opt/WSPP
# cd afs
# ./afsConfigure.sh
Please choose from the following options
1. Configure this machine as your SCM.
2. Configure this machine as a File Server.
3. Configure this machine as an AFS Client.
4. Quit the program.
Enter choice: 2
This machine is a File server

Setting up File Server
Loading SUN kernel extensions
nfsd was not found on this machine

Enter the name of your partition: /vicepa

Starting the bosserver

bosserver started successfully
Enter the name of your cell: sushi.itso.ral.ibm.com
Enter the hostname of your first AFS Server: rs60002.itso.ral.ibm.com
bos: no such entry (getting tickets)
bos: running unauthenticated

Enter the number of 1 KB blocks to devote to the cache: 40000
 U.S. Government Restricted Rights.
====================================
If you are licensing the Software on behalf of the U.S.
Government ("Government"), the following provisions apply
to you. If the Software is supplied to the Department of
Defense ("DoD"), it is classified as "Commercial Computer
Software" under paragraph 252.227-7014 of the DoD Supplement
to the Federal Acquisition Regulations ("DFARS") (or
any successor regulations) and the Government is acquiring
only the license rights granted herein (the license rights
customarily provided to non-Government users).  If the
Software is supplied to any unit or agency of the Government
other than DoD, it is classified as "Restricted Computer
Software" and the Government's rights in the Software are
defined in paragraph 52.227-19 of the Federal Acquisition
Regulations ("FAR") (or any successor regulations) or, in
the case of NASA, in paragraph 18.52.227-86 of the NASA
Supplement in the FAR (or any successor regulations).
```

Figure 66. (Part 1 of 2). Configuration of an AFS File Server on a Solaris System

```
afsd: My home cell is 'sushi.itso.ral.ibm.com'
afsd: Creating '/usr/vice/etc/AFSLog'
CreateCacheFile: Creating cache file '/usr/vice/etc/AFSLog'
SScall(105, 28)=0 afsd: Forking rx callback listener.
afsd: Forking AFS daemon.
afsd: Forking Check Server Daemon.
afsd: Forking 3 background daemons.
afsd: Forking rxevent daemon.
SScall(105, 28)=0 afsd: Calling AFSOP_CACHEINIT: 2000 stat cache entries,
4000 optimum cache files, 40000 blocks in the cache, flags = 0x0, dcache
entries 800
SScall(105, 28)=0 afsd: Sweeping workstation's AFS cache directory.
SweepAFSCache: Creating '/usr/vice/cache/CacheItems'
CreateCacheFile: Creating cache file '/usr/vice/cache/CacheItems'
SweepAFSCache: Creating '/usr/vice/cache/VolumeItems'
CreateCacheFile: Creating cache file '/usr/vice/cache/VolumeItems'
SweepAFSCache: Creating '/usr/vice/cache/V0'
CreateCacheFile: Creating cache file '/usr/vice/cache/V0'
SweepAFSCache: Creating '/usr/vice/cache/V1'
CreateCacheFile: Creating cache file '/usr/vice/cache/V1'
SweepAFSCache: Creating '/usr/vice/cache/V2'
CreateCacheFile: Creating cache file '/usr/vice/cache/V2'
...
SweepAFSCache: Creating '/usr/vice/cache/V3997'
CreateCacheFile: Creating cache file '/usr/vice/cache/V3997'
SweepAFSCache: Creating '/usr/vice/cache/V3998'
CreateCacheFile: Creating cache file '/usr/vice/cache/V3998'
SweepAFSCache: Creating '/usr/vice/cache/V3999'
CreateCacheFile: Creating cache file '/usr/vice/cache/V3999'
afsd: 0 out of 4000 data cache files found in sweep 1.
afsd: 4000 out of 4000 data cache files found in sweep 2.
...
afsd: All AFS daemons started.
afsd: Forking trunc-cache daemon.
afsd: Mounting the AFS root on '/afs', flags: 4.
```

Figure 67. (Part 2 of 2). Configuration of an AFS File Server on a Solaris System

4.5.2 Verifying the AFS File Server Configuration

To verify that the configuration of an AFS File Server completed successfully, refer to the procedure described in 3.5.2, “Verifying the AFS File Server Configuration” on page 87.

Recall that in WebSphere Performance Pack an AFS File Server is automatically an AFS client. To verify that the AFS client is functioning properly, you need to perform the tasks in 4.6.2, “Verifying the AFS Client Configuration” on page 119.

To configure the AFS initialization file to start AFS at system reboot, you need to follow the procedures in 4.7, “Starting AFS at System Reboot on Solaris Systems” on page 119. After completing the AFS initialization file configuration, reboot your machine and verify that all AFS processes are running normally.

4.6 Configuring an AFS Client Machine on Solaris Systems

To configure a Solaris machine to be an AFS client in an existing AFS cell, you must have already completed the following tasks:

- Configure the system to recognize AFS system calls, as demonstrated in 4.1.6, “AFS Extensions to the Solaris Kernel” on page 101.
- Install the necessary AFS files from the WebSphere Performance Pack CD-ROM distribution to the machine local disk, as shown in 4.3, “Installing AFS Files on Solaris Systems” on page 107.

Then, you can proceed with the instructions in 4.6.1, “Configuring AFS Client with WebSphere AFS Configuration Script” on page 117 and 4.6.2, “Verifying the AFS Client Configuration” on page 119.

4.6.1 Configuring AFS Client with WebSphere AFS Configuration Script

Since the configuration procedure for Solaris systems is the same as the one for AIX systems, refer to 3.6.2, “Verifying the AFS Client Configuration” on page 92, to proceed with the configuration.

In summary, the configuration script prompts you for the name of an existing AFS cell, the AFS admin password, and the size of the AFS disk cache in 1 KB blocks.

In our environment, the name of our existing AFS cell is sushi.itso.ral.ibm.com, and the size of our AFS disk cache is 40,000 1 KB blocks. The next several text figures contain the transcript of our configuration session of an AFS client on a Solaris system called sun3. Notice that we have masked the AFS admin password as xxxxxxxx.

```
# cd /usr/vice/etc
(Copy or FTP the ThisCell and CellServDB file from your first AFS machine)
# cd /opt/WSPP
# cd afs
# ./afsConfigure.sh
Please choose from the following options
1. Configure this machine as your SCM.
2. Configure this machine as a File Server.
3. Configure this machine as an AFS Client.
4. Quit the program.
Enter choice: 3
This machine will be an AFS Client

Setting up AFS Client
Enter your cell name: sushi.itso.ral.ibm.com
Enter you admin password: XXXXXXXX
Loading SUN kernel extensions
nfsd was not found on this machine

Enter the number of 1 KB blocks to devote to the cache: 40000
 U.S. Government Restricted Rights.
====================================
If you are licensing the Software on behalf of the U.S.
Government ("Government"), the following provisions apply
to you. If the Software is supplied to the Department of
Defense ("DoD"), it is classified as "Commercial Computer
Software" under paragraph 252.227-7014 of the DoD Supplement
to the Federal Acquisition Regulations ("DFARS") (or
any successor regulations) and the Government is acquiring
only the license rights granted herein (the license rights
customarily provided to non-Government users).  If the
Software is supplied to any unit or agency of the Government
other than DoD, it is classified as "Restricted Computer
Software" and the Government's rights in the Software are
defined in paragraph 52.227-19 of the Federal Acquisition
Regulations ("FAR") (or any successor regulations) or, in
the case of NASA, in paragraph 18.52.227-86 of the NASA
Supplement in the FAR (or any successor regulations).

afsd: My home cell is 'sushi.itso.ral.ibm.com'
SScall(105, 28)=0 afsd: Forking rx callback listener.
afsd: Forking AFS daemon.
afsd: Forking Check Server Daemon.
afsd: Forking 3 background daemons.
afsd: Forking rxevent daemon.
```

Figure 68. (Part 1 of 2). Configuration of an AFS Client on a Solaris System

```
SScall(105, 28)=0 afsd: Calling AFSOP_CACHEINIT: 2000 stat cache entries,
4000 optimum cache files, 40000 blocks in the cache, flags = 0x0, dcache
entries 800
SScall(105, 28)=0 afsd: Sweeping workstation's AFS cache directory.
SweepAFSCache: Creating '/usr/vice/cache/CacheItems'
CreateCacheFile: Creating cache file '/usr/vice/cache/CacheItems'
SweepAFSCache: Creating '/usr/vice/cache/VolumeItems'
CreateCacheFile: Creating cache file '/usr/vice/cache/VolumeItems'
SweepAFSCache: Creating '/usr/vice/cache/V0'
CreateCacheFile: Creating cache file '/usr/vice/cache/V0'
SweepAFSCache: Creating '/usr/vice/cache/V1'
CreateCacheFile: Creating cache file '/usr/vice/cache/V1'
SweepAFSCache: Creating '/usr/vice/cache/V2'
CreateCacheFile: Creating cache file '/usr/vice/cache/V2'
...
SweepAFSCache: Creating '/usr/vice/cache/V3997'
CreateCacheFile: Creating cache file '/usr/vice/cache/V3997'
SweepAFSCache: Creating '/usr/vice/cache/V3998'
CreateCacheFile: Creating cache file '/usr/vice/cache/V3998'
SweepAFSCache: Creating '/usr/vice/cache/V3999'
CreateCacheFile: Creating cache file '/usr/vice/cache/V3999'
afsd: 0 out of 4000 data cache files found in sweep 1.
afsd: 4000 out of 4000 data cache files found in sweep 2.
...
afsd: All AFS daemons started.
afsd: Forking trunc-cache daemon.
afsd: Mounting the AFS root on '/afs', flags: 4.

Tokens held by the Cache Manager:

User's (AFS ID 1) tokens for afs@sushi.itso.ral.ibm.com [Expires Apr  2
12:11]
   --End of list--
All volumeID/name mappings checked.
```

Figure 69. (Part 2 of 2). Configuration of an AFS Client on a Solaris System

4.6.2 Verifying the AFS Client Configuration

To verify that the configuration of an AFS client completed successfully, refer to the procedure described in 3.6.2, “Verifying the AFS Client Configuration” on page 92.

To configure the AFS initialization file to start AFS at system reboot, you need to follow the procedures in 4.7, “Starting AFS at System Reboot on Solaris Systems” on page 119. After completing the AFS initialization file configuration, reboot your machine and verify that all AFS processes are running normally.

4.7 Starting AFS at System Reboot on Solaris Systems

At system reboot, the system starts AFS by executing the AFS initialization file called /etc/init.d/afs.rc. This file is a script containing commands to:

1. Load the AFS extensions into the kernel
2. Start the BOS Server process for an AFS server

3. Start the AFS Cache Manager for an AFS client

To inform the system about this AFS initialization file, perform the following commands:

```
ln -s ../init.d/afs.rc /etc/rc0.d/K66afs
ln -s ../init.d/afs.rc /etc/rc3.d/S99afs
```

You should make the /etc/init.d/afs.rc file executable. You can do this by entering the command:

```
chmod +x /etc/init.d/afs.rc
```

Figure 70 on page 120 illustrates the configuration of the AFS initialization file.

```
# ls -l /etc/*.d/*afs*
-rwxr-xr-x   1 root     other       3400 Apr 12 20:24 /etc/init.d/afs.rc
lrwxrwxrwx   1 root     other         16 Apr 12 20:26 /etc/rc0.d/K66afs -> ../init.d/afs.rc
lrwxrwxrwx   1 root     other         16 Apr 12 20:26 /etc/rc3.d/S99afs -> ../init.d/afs.rc
```

Figure 70. Listing of the AFS Initialization Files on a Solaris System

This AFS initialization file can be used on any Solaris machine running AFS server and/or client processes.

WebSphere AFS Initialization Files

WebSphere Performance Pack provides a similar AFS initialization file in the directory /opt/WSPP/afs. This file contains control characters that make the file itself unexecutable. Refer to 3.7, “Starting AFS at System Reboot on AIX Systems” on page 94 to see how to solve this problem.

You can also locate other AFS initialization scripts in the directories /usr/vice/etc and /usr/vice/etc/modload.

The next figures contain the listing of the AFS initialization file for a Solaris system, /etc/init.d/afs.rc:

```
#!/bin/sh
#
# afs.rc: rc script for AFS on Solaris 2.x platforms
#
# Install this script as /etc/init.d/afs.rc
# then make links like this:
# ln -s ../init.d/afs.rc /etc/rc0.d/K66afs
# ln -s ../init.d/afs.rc /etc/rc3.d/S99afs
#
CONFIG=/usr/vice/etc/config
AFSDOPT=$CONFIG/afsd.options
PACKAGE=$CONFIG/package.options

LARGE="-stat 2800 -dcache 2400 -daemons 5 -volumes 128"
MEDIUM="-stat 2000 -dcache 800 -daemons 3 -volumes 70"
SMALL="-stat 300 -dcache 100 -daemons 2 -volumes 50"

if [ -f $AFSDOPT ]; then
    OPTIONS=`cat $AFSDOPT`
else
    OPTIONS=$MEDIUM
fi

# Need the commands ps, awk, kill, sleep
PATH=${PATH}${PATH:+:}/sbin:/bin:/usr/bin

killproc() {            # kill the named process(es)
        awkfield2='$2'
        pid=`ps -ef | awk "/$1/ && ! /awk/ {print $awkfield2}"`
        [ "$pid" != "" ] && kill -KILL $pid
}

case $1 in
'start')

#
# Make sure afs exists in /etc/name_to_sysnum
#
if grep -s "afs" /etc/name_to_sysnum > /dev/null; then
    echo "Entry for afs already exists in /etc/name_to_sysnum"
else
    echo "Creating entry for afs in /etc/name_to_sysnum"
    cp /etc/name_to_sysnum /etc/name_to_sysnum.orig
    sed '/nfs/i\
afs                        105' /etc/name_to_sysnum > /tmp/name_to_sysnum
    mv /tmp/name_to_sysnum /etc/name_to_sysnum
    echo "Rebooting now for new /etc/name_to_sysnum to take effect"
    reboot
fi
```

Figure 71. (Part 1 of 3). AFS Initialization File (/etc/init.d/afs.rc) for Solaris Systems

```
#
# Load kernel extensions
#

# nfssrv has to be loaded first
if [ -f /kernel/misc/nfssrv ]; then
        echo "Loading NFS server kernel extensions"
        modload /kernel/misc/nfssrv
else
        echo "/kernel/misc/nfssrv does not exist. Skipping AFS startup."
        exit 1
fi

if [ -f /kernel/fs/afs ]; then
        echo "Loading AFS kernel extensions"
        modload /kernel/fs/afs
else
        echo "/kernel/fs/afs does not exist. Skipping AFS startup."
        exit 1
fi

#
# Start the AFS server processes if a bosserver exists
#

if [ -x /usr/afs/bin/bosserver -a -f /usr/afs/etc/ThisCell \
    -a -f /usr/afs/etc/CellServDB -a -f /usr/afs/etc/KeyFile \
    -a -f /usr/afs/etc/UserList -a -f /usr/afs/etc/License ]; then
        echo "Starting AFS Server processes"
        /usr/afs/bin/bosserver &
        OPTIONS="$OPTIONS -nosettime"
        sleep 30
fi

#
# Check that all of the client configuration files exist
#

for file in /usr/vice/etc/afsd /usr/vice/etc/cacheinfo \
            /usr/vice/etc/ThisCell /usr/vice/etc/CellServDB
do
        if [ ! -f ${file} ]; then
                echo "${file} does not exist. Not starting AFS client."
                exit 1
        fi
done
```

Figure 72. (Part 2 of 3). AFS Initialization File (/etc/init.d/afs.rc) for Solaris Systems

```
#
# Check that the root directory for AFS (/afs)
# and the cache directory (/usr/vice/cache) both exist
#

for dir in `awk -F: '{print $1, $2}' /usr/vice/etc/cacheinfo`
do
        if [ ! -d ${dir} ]; then
                echo "${dir} does not exist. Not starting AFS client."
                exit 2
        fi
done

echo "Starting afsd"
/usr/vice/etc/afsd $OPTIONS

#
# Start AFS inetd services
# (See the AFS Command Ref. for notes on the proper configuration of
inetd.afs)
#
if [ -f /usr/sbin/inetd.afs -a -f /etc/inetd.conf.afs ]; then
        /usr/sbin/inetd.afs /etc/inetd.conf.afs
fi

echo ;;

'stop')

#
# Stop the AFS inetd and server processes
# Note that the afsd processes cannot be killed
#

echo "Killing inetd.afs"
killproc inetd.afs

bosrunning=`ps -ef | awk '/bosserver/ && ! /awk/'`
if [ "${bosrunning}" != "" ]; then
        echo "Shutting down AFS server processes"
        /usr/afs/bin/bos shutdown localhost -localauth -wait
        echo "Killing AFS bosserver"
        killproc bosserver
fi

echo ;;

*)         echo "Invalid option supplied to $0"
        exit 1;;
esac
```

Figure 73. (Part 3 of 3). AFS Initialization File (/etc/init.d/afs.rc) for Solaris Systems

Chapter 5. AFS Installation and Configuration on Windows NT

This section summarizes the installation prerequisites for the AFS component of WebSphere Performance Pack on NT systems. It also describes the installation and configuration procedures, including the first AFS machine, an additional AFS File Server, and an AFS Client.

5.1 Windows NT System Requirements

To prepare your Windows NT system for the installation of AFS, the File Sharing component of WebSphere Performance Pack, you need:

- An Intel 80486 (66 MHz) or Pentium microprocessor PC
- Windows NT Workstation or Server V4.0 with Service Pack 3 or 4
- A minimum of 64 MB of RAM
- A minimum of 128 MB of paging space
- A minimum of 10 MB of available disk space on the default Windows NT system drive to install the AFS files
- A minimum of 20 MB of available disk space on the default Windows NT system drive for the AFS memory mapped file cache for an AFS client

  The default AFS cache size is 20 MB. You can increase the cache size to improve performance.
- At least one NTFS disk partition dedicated to storing AFS data for each AFS File Server

  For AFS to recognize that the disk partition is available to be configured as an AFS partition, the disk partition must not contain any data other than a Windows NT Recycle Bin.
- Java Runtime Environment (JRE) 1.1.6

  At the time of this writing, JRE 1.1.7 still had problems removing programs from Windows NT.
- Administrative privileges

  You need to be a member of the local Administrators group on your Windows NT system in order to install, configure, and start AFS server and client processes.

We also recommend that you read the WebSphere Performance Pack README. This file can be found on the IBM WebSphere Performance Pack for Multiplatforms CD-ROM.

5.1.1 Operating System Version

WebSphere Performance Pack V2.0 includes AFS 3.5 for Windows NT, which is supported on Windows NT Version 4.0 with Service Pack 3 or 4. To examine the Service Pack level on your machine, you need to invoke the Windows NT Diagnostic program. Click the **Start** button on the Windows NT desktop taskbar, point to **Programs**, select the **Administrative Tools (Common)** program group, and click **Windows NT Diagnostics**. In the Windows NT Diagnostic window, select the **Version** tab to display the operating system version and Service Pack level of Windows NT on your machine, as shown in Figure 74 on page 126.

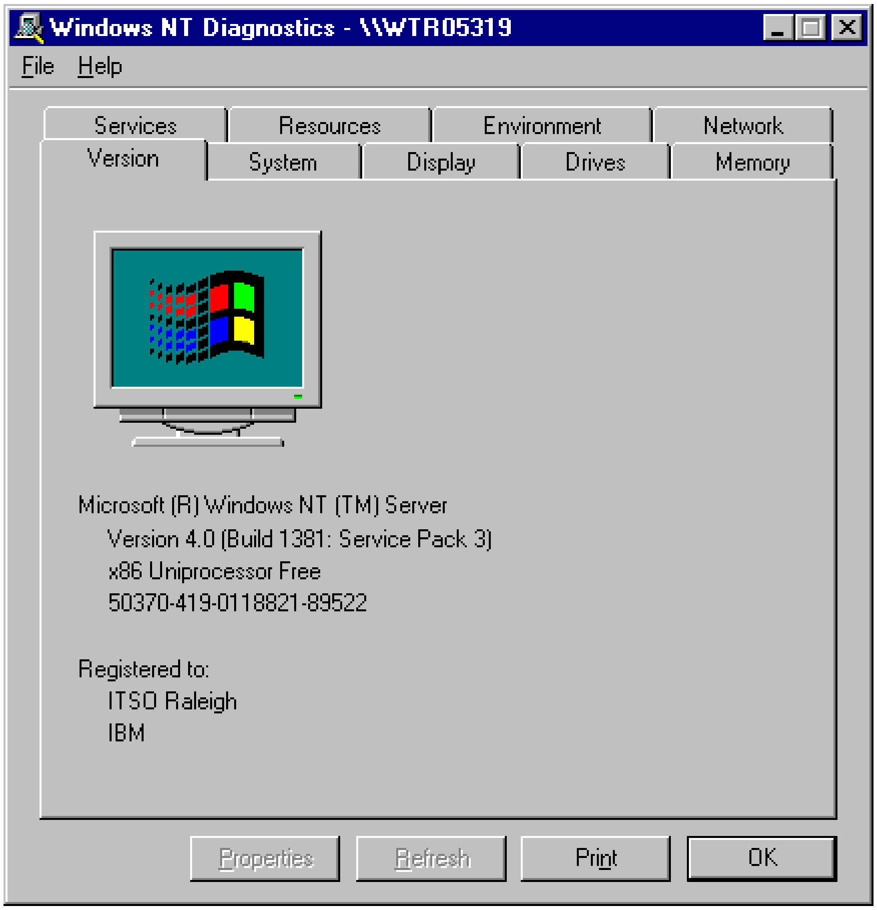

Microsoft (R) Windows NT (TM) Server
Version 4.0 (Build 1381: Service Pack 3)
x86 Uniprocessor Free
50370-419-0118821-89522

Registered to:
ITSO Raleigh
IBM

Figure 74. Displaying the Windows NT Version

5.1.2 Page File Allocation

You need to allocate adequate paging spaces on your Windows NT machine. For performance reasons, the amount of paging space must be at least 10% more than the amount of memory on your machine.

To examine the paging space of your machine, you need to invoke the Windows NT Diagnostic program. Click the **Start** button on the Windows NT desktop taskbar, point to **Programs**, select the **Administrative Tools (Common)** program group, and click **Windows NT Diagnostics**. In the Windows NT Diagnostic window, select the **Memory** tab to display the amount of physical memory and the size of the paging file of your machine, as shown in Figure 75 on page 127.

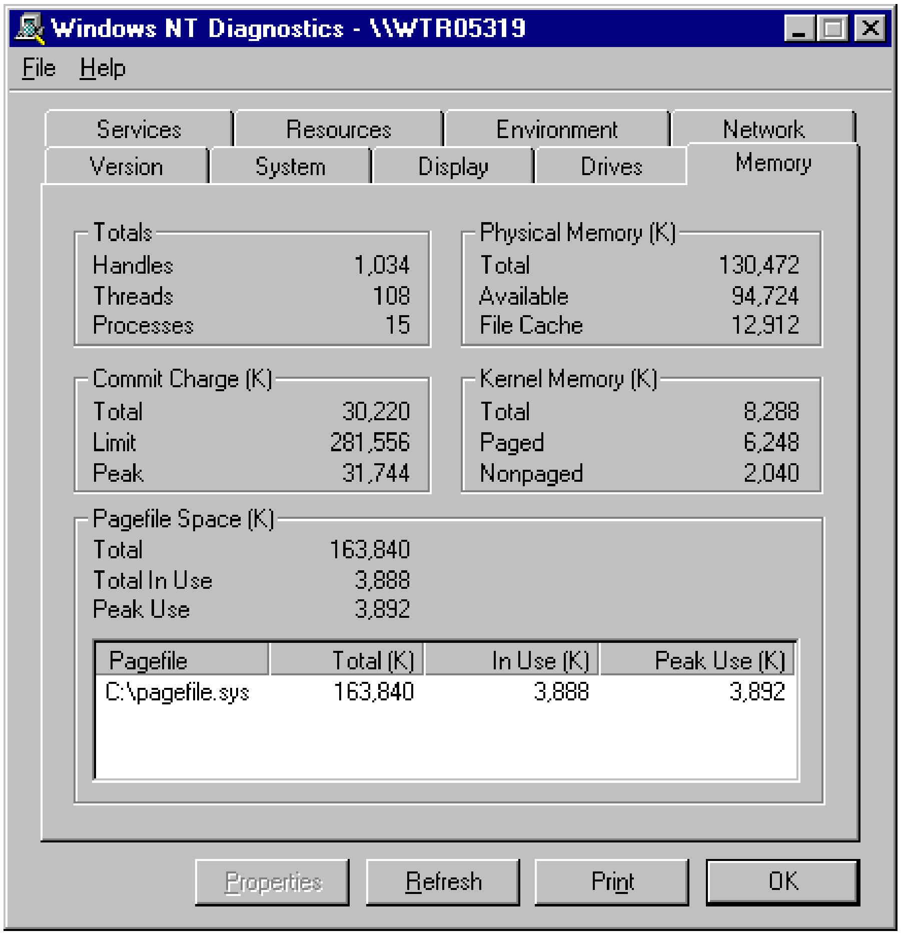

Totals

Handles	1,034
Threads	108
Processes	15

Physical Memory (K)

Total	130,472
Available	94,724
File Cache	12,912

Commit Charge (K)

Total	30,220
Limit	281,556
Peak	31,744

Kernel Memory (K)

Total	8,288
Paged	6,248
Nonpaged	2,040

Pagefile Space (K)

Total	163,840
Total In Use	3,888
Peak Use	3,892

Pagefile	Total (K)	In Use (K)	Peak Use (K)
C:\pagefile.sys	163,840	3,888	3,892

Figure 75. Displaying the Physical Memory and Pagefile Size

To change the size of the paging file, you need to invoke the Control Panel System window. Click the **Start** button on the Windows NT desktop taskbar, point to **Settings**, and click **Control Panel**. From the Windows NT Control Panel, double-click the **System** icon. In the System Properties window, select the **Performance** tab. Then, in the Virtual Memory section, press **Change**, as shown in Figure 76 on page 128.

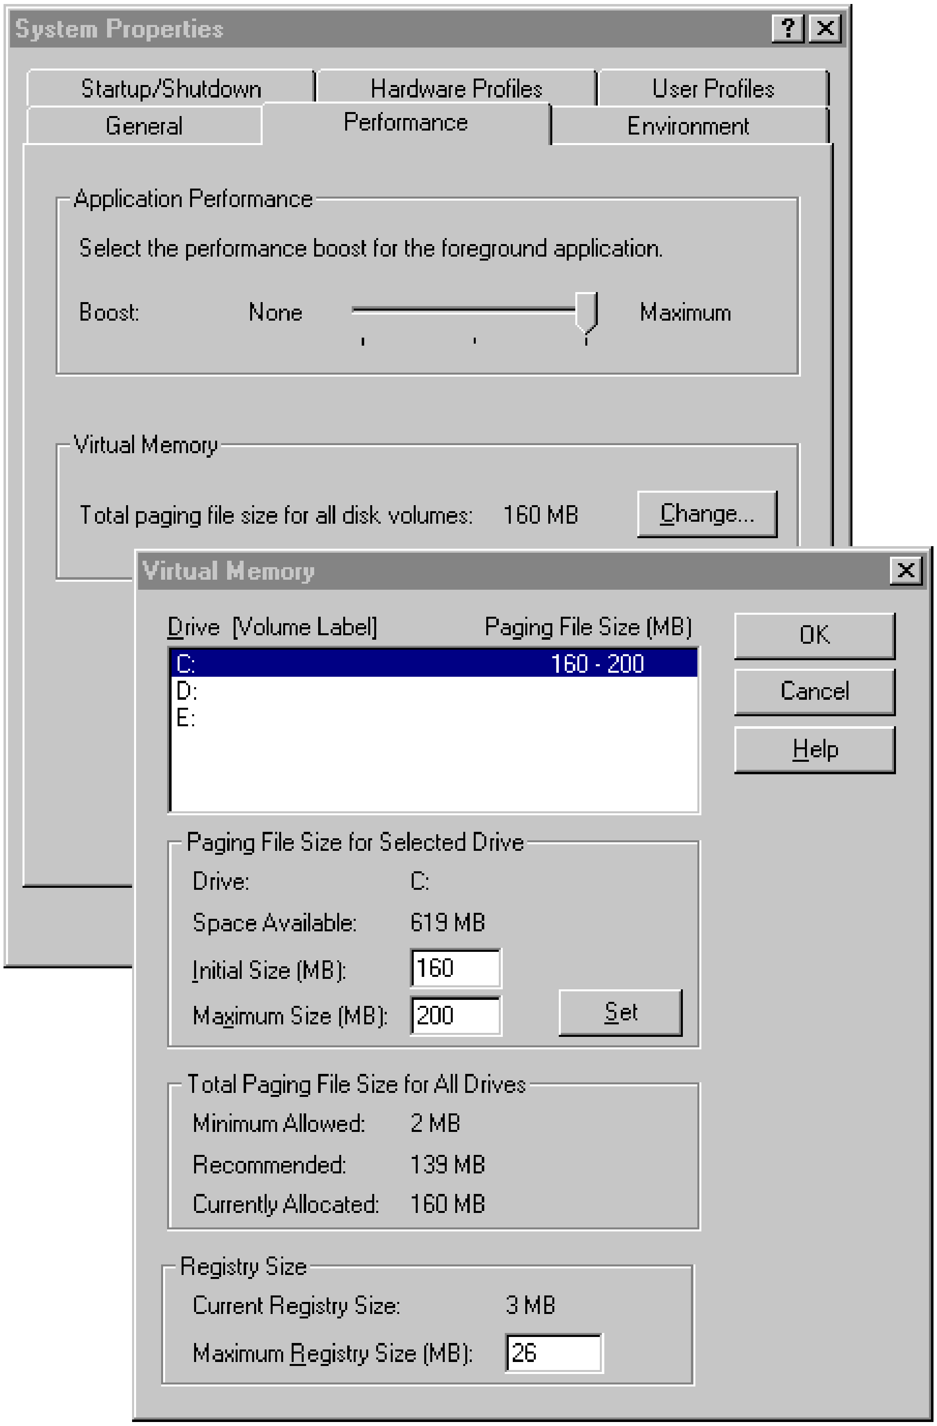

Application Performance

Select the performance boost for the foreground application.

Boost: None — Maximum

Virtual Memory

Total paging file size for all disk volumes: 160 MB

Drive [Volume Label]	Paging File Size (MB)
C:	160 - 200
D:	
E:	

Paging File Size for Selected Drive

Drive:	C:
Space Available:	619 MB
Initial Size (MB):	160
Maximum Size (MB):	200

Total Paging File Size for All Drives

Minimum Allowed:	2 MB
Recommended:	139 MB
Currently Allocated:	160 MB

Registry Size

Current Registry Size:	3 MB
Maximum Registry Size (MB):	26

Figure 76. Displaying the Pagefile Size

5.1.3 Name Resolution Service

We recommend that you use a name service in your environment, such as the Domain Name System (DNS), to resolve the host names of machines. It is difficult to accurately maintain the host names of machines with local host tables.

On Windows NT systems, the `nslookup` command resolves a host name into an Internet address or an Internet address into a host name. By using the `nslookup` command, you can verify that both the forward and reverse name resolution services are functioning correctly.

Figure 77 on page 129 illustrates the usage of the `nslookup` command for a machine with the host name of wtr05319, the IP address of 9.24.104.206, and the DNS domain name of itso.ral.ibm.com.

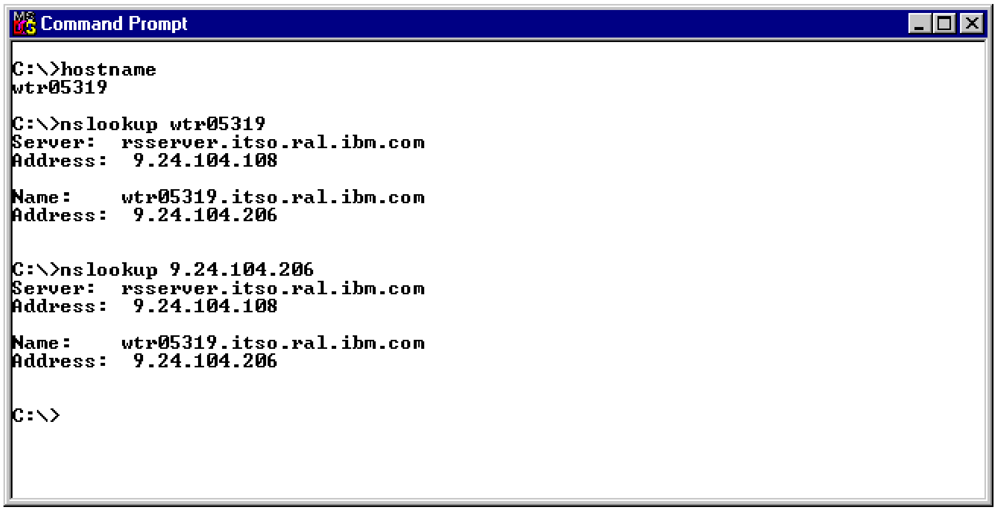

```
C:\>hostname
wtr05319

C:\>nslookup wtr05319
Server:  rsserver.itso.ral.ibm.com
Address:  9.24.104.108

Name:    wtr05319.itso.ral.ibm.com
Address:  9.24.104.206


C:\>nslookup 9.24.104.206
Server:  rsserver.itso.ral.ibm.com
Address:  9.24.104.108

Name:    wtr05319.itso.ral.ibm.com
Address:  9.24.104.206


C:\>
```

Figure 77. Verifying Forward and Reverse Host Name Resolution on an NT System

5.1.4 Network Connection

The Windows NT machine has to be operational on the network. For each machine where you are planning to install and configure AFS, verify network configuration and integrity.

On Windows NT systems, the following commands can assist you in determining your network status:

- The command:

  `ipconfig -all`

  displays the Windows NT IP configuration.

- The command:

  `ping` *`host`*

  sends a network request to a host or gateway.

Figure 78 on page 130 demonstrates the usage of these commands from a Windows NT machine with the IP address 9.24.104.206.

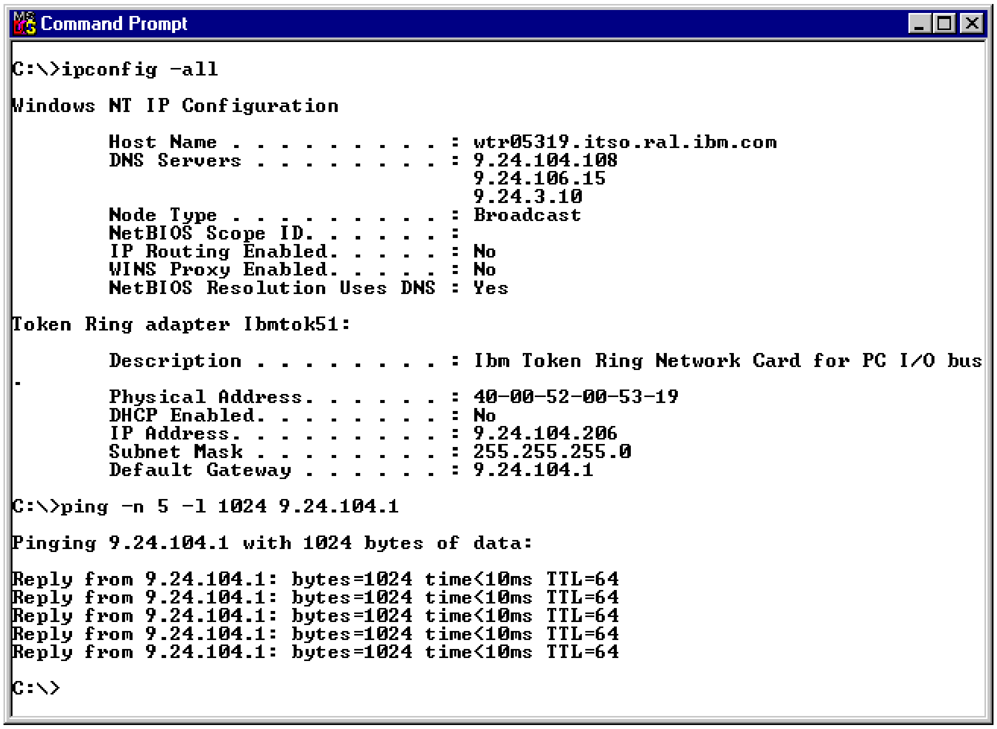

```
C:\>ipconfig -all

Windows NT IP Configuration

        Host Name . . . . . . . . . : wtr05319.itso.ral.ibm.com
        DNS Servers . . . . . . . . : 9.24.104.108
                                      9.24.106.15
                                      9.24.3.10
        Node Type . . . . . . . . . : Broadcast
        NetBIOS Scope ID. . . . . . :
        IP Routing Enabled. . . . . : No
        WINS Proxy Enabled. . . . . : No
        NetBIOS Resolution Uses DNS : Yes

Token Ring adapter Ibmtok51:

        Description . . . . . . . . : Ibm Token Ring Network Card for PC I/O bus
.
        Physical Address. . . . . . : 40-00-52-00-53-19
        DHCP Enabled. . . . . . . . : No
        IP Address. . . . . . . . . : 9.24.104.206
        Subnet Mask . . . . . . . . : 255.255.255.0
        Default Gateway . . . . . . : 9.24.104.1

C:\>ping -n 5 -l 1024 9.24.104.1

Pinging 9.24.104.1 with 1024 bytes of data:

Reply from 9.24.104.1: bytes=1024 time<10ms TTL=64
Reply from 9.24.104.1: bytes=1024 time<10ms TTL=64
Reply from 9.24.104.1: bytes=1024 time<10ms TTL=64
Reply from 9.24.104.1: bytes=1024 time<10ms TTL=64
Reply from 9.24.104.1: bytes=1024 time<10ms TTL=64

C:\>
```

Figure 78. Displaying the Network Interface and Verifying the Network Connection

In addition, you need to verify that the default LAN Adapter (LANA) number is assigned to the active network adapter. The AFS 3.5 NT Client uses the default LANA number of 000.

You can examine the default LANA number setting on your machine from the Control Panel Network window. To invoke this window, click the **Start** button on the Windows NT desktop taskbar, point to **Settings**, and click **Control Panel**. From the Windows NT Control Panel, double-click on the **Network** icon. In the Network window, select the **Services** tab and double-click on the **NetBIOS Interface** entry to display the NetBIOS Configuration window, as shown in Figure 79 on page 131.

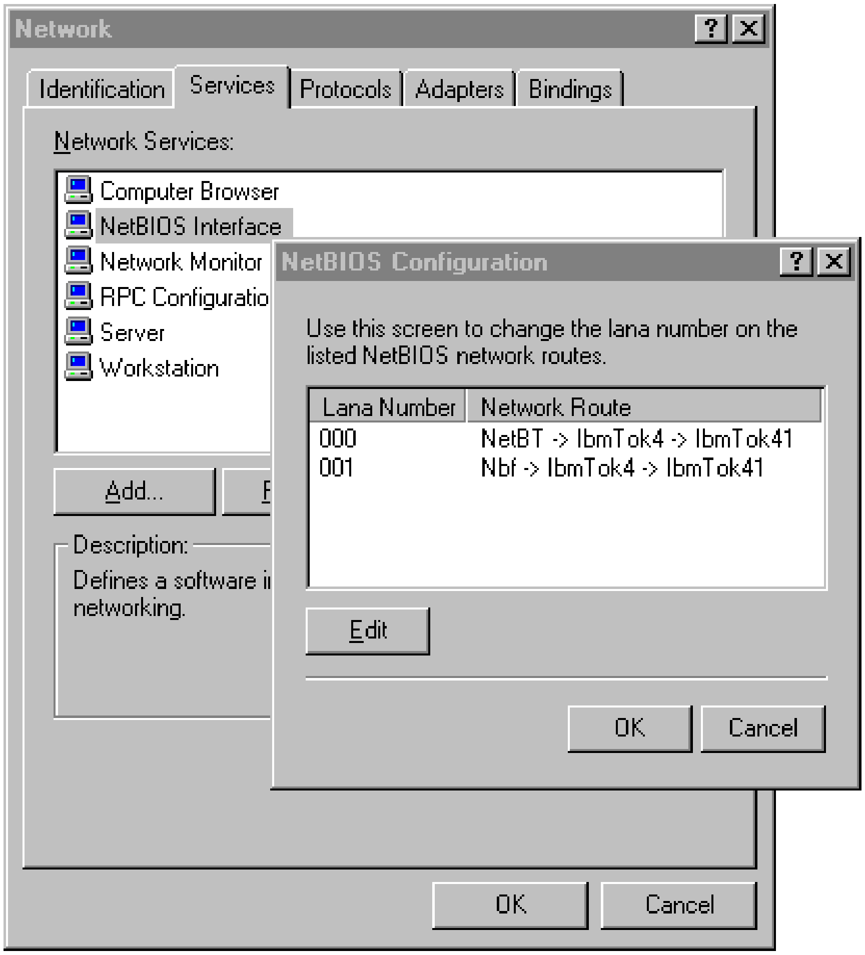

Defines a software i
networking.

Use this screen to change the lana number on the listed NetBIOS network routes.

Lana Number	Network Route
000	NetBT -> IbmTok4 -> IbmTok41
001	Nbf -> IbmTok4 -> IbmTok41

Figure 79. Displaying the LANA Number on an IBM PC Compatible

Figure 79 on page 131 shows the LANA number 000 assigned to the NetBT network service, using the IBM Token-Ring Adapter (IbmTok4), on an IBM Netfinity 3000.

As another example, Figure 80 on page 132 shows the LANA number 000 assigned to the NetBT network service, using the IBM Token-Ring Adapter (IbmTok4), on an IBM ThinkPad 600:

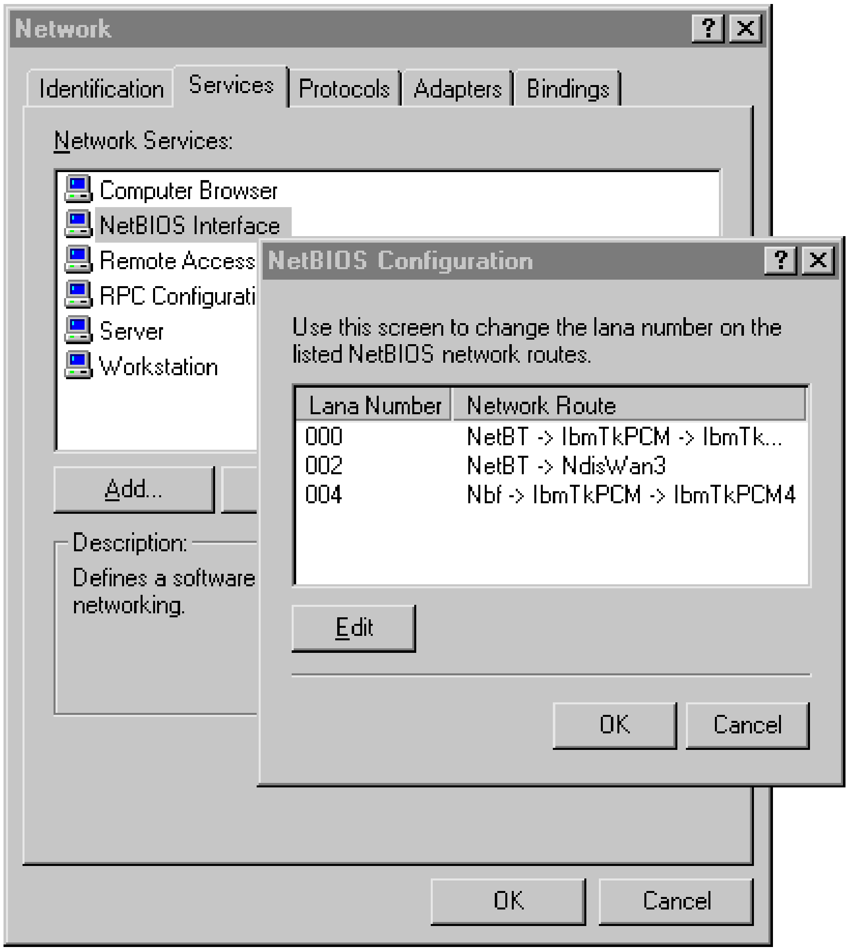

Use this screen to change the lana number on the listed NetBIOS network routes.

Lana Number	Network Route
000	NetBT -> IbmTkPCM -> IbmTk...
002	NetBT -> NdisWan3
004	Nbf -> IbmTkPCM -> IbmTkPCM4

Figure 80. Displaying the LANA Number on an IBM ThinkPad

5.1.5 Time Synchronization

Before configuring additional AFS Servers or Clients on Windows NT machines, manually verify that the clock on each new machine is relatively accurate with respect to the first AFS machine. If the clock is off on a new machine, correct it before continuing. On Windows NT systems, AFS does not provide any time service to synchronize time among AFS NT machines in the AFS cell. We recommend that you use some type of time synchronization mechanism in your environment. Microsoft provides the Time Service (TimeServ) for Windows NT 4.0 on the Microsoft WindowsNT Server Resource Kit CD-ROM. Other tools are available for free downloading from the Web.

5.1.6 JRE Version

The WebSphere Performance Pack installation process uses the Java InstallShield's setup class. You need to ensure that the Java Runtime Environment (JRE) is installed. At the time of this writing, the recommended JRE version for Windows NT is 1.1.6.

To determine the version on your Windows NT system, enter the following command from an MS-DOS Command Prompt window:

```
java -version
```

This program is accepted provided that the bin directory where the Java binaries are installed is correctly added to the Path system environment variable.

The CD distribution of WebSphere Performance Pack contains the latest Java Development Kit (JDK) for Windows NT in the directory nt\jdk. Actually, you would not need the full JDK, but only its subset known as JRE, which contains just the Java Virtual Machine (JVM), the Java platform core classes, and supporting files. In other words, the JRE is the smallest set of executables and files that constitute the standard Java platform and it contains only the run-time part of the JDK: no compiler, no debugger, no tools.

5.1.7 Disk Space Requirements

By default, WebSphere Performance Pack installs the AFS component in the directory C:\WSPP\AFS. You must allocate at least 10 MB of available disk space for this directory.

5.2 Configuring AFS Partitions on Windows NT Systems

Every AFS File Server must have a disk partition dedicated to be an AFS partition to store AFS volumes. By convention, the first AFS partition created on an AFS File Server machine is called /vicepa, the second partition /vicepb, etc. In order for a Windows NT machine to function as an AFS File Server, you need to create at least one NTFS disk partition to be configured as an AFS partition.

There are several ways to partition a disk on Windows NT. This section describes a procedure using Windows NT Disk Administrator tool. The following example assumes that there is enough unpartitioned disk space.

To create an NTFS file system to be used as an AFS partition, you need to perform the following operations:

- Create a primary or extended Windows NT partition
- Create a Windows NT volume on a Windows NT partition
- Format the Windows NT volume as an NTFS file system

The following session demonstrates the steps to create a 1,024 MB NTFS partition to be used as an AFS partition:

1. To invoke the Windows NT Disk Administrator tool, click the **Start** button on your Windows NT desktop taskbar, point to **Programs**, point to **Administrative Tools (Common)**, and click **Disk Administrator**.

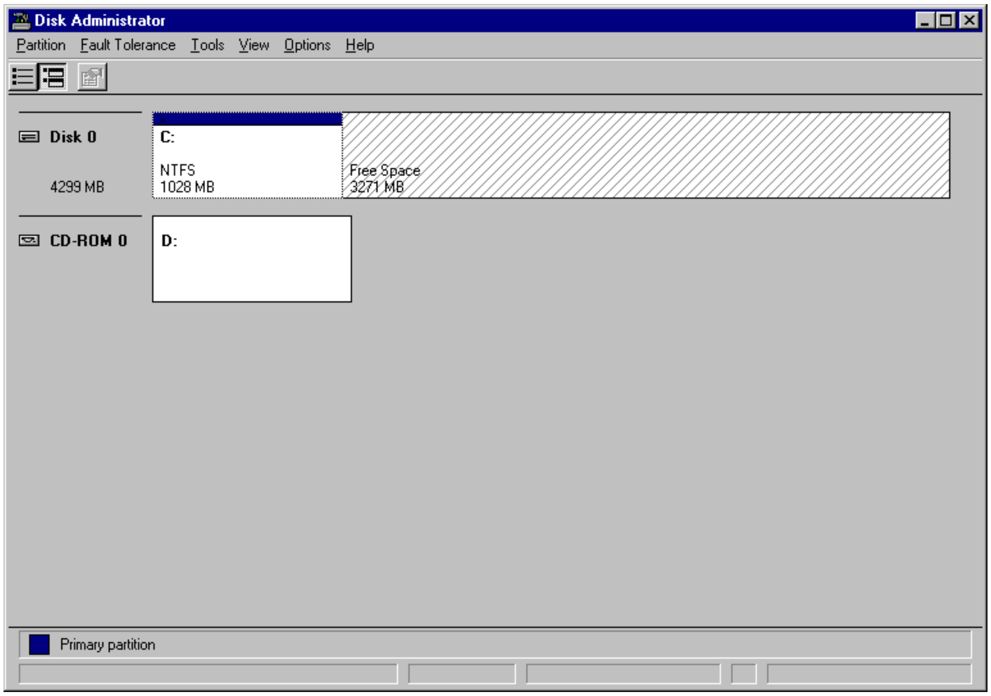

Figure 81. The Windows NT Disk Administrator Window

In this example, the Windows NT machine has one primary partition assigning 1028 MB to the C: drive and 3271 MB of unpartitioned free space left on Disk 0.

2. Click the **Free Space** area to select the unpartitioned disk space. Click the **Partition** button on the Disk Administrator menu bar, and click **Create...**, as shown in Figure 82 on page 135.

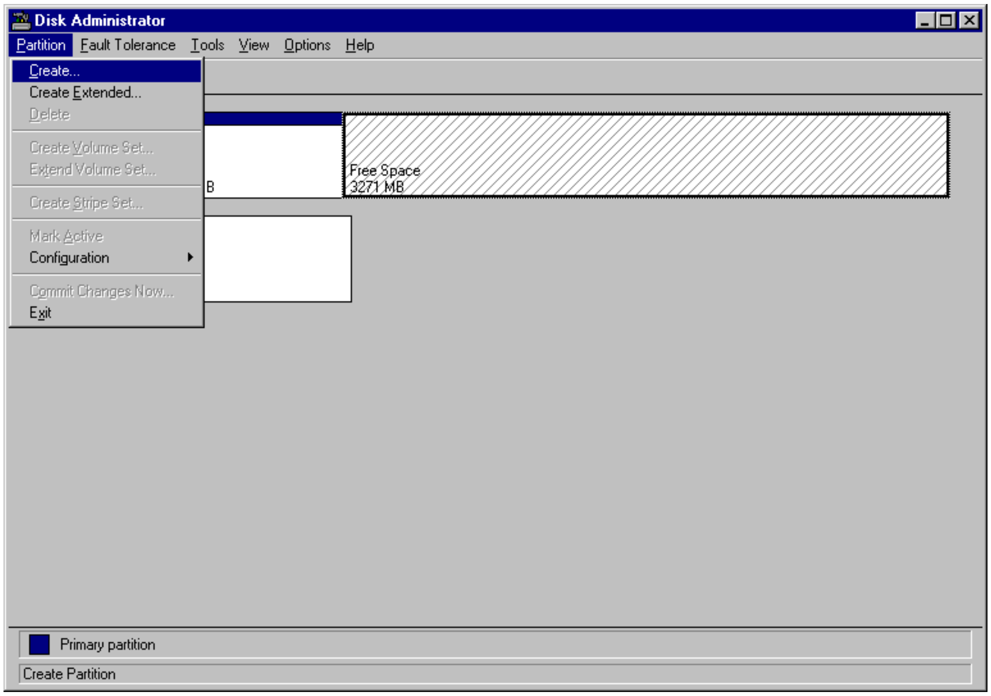

Figure 82. Invoking the Create Partition Option of the Disk Administrator

3. A confirmation message appears as shown in Figure 83 on page 135. Click **Yes** to continue.

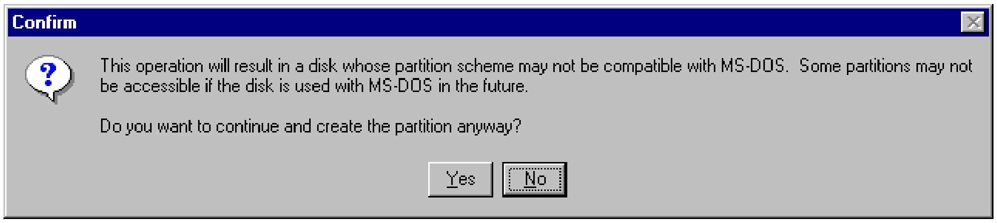

This operation will result in a disk whose partition scheme may not be compatible with MS-DOS. Some partitions may not be accessible if the disk is used with MS-DOS in the future.

Do you want to continue and create the partition anyway?

Figure 83. The Create Partition Confirmation Window

4. In the Create Primary Partition window, enter the desired partition size in MB to create a new disk partition:

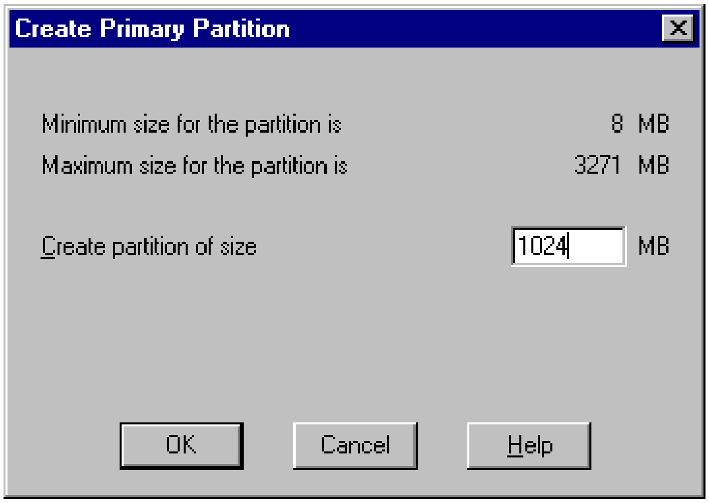
Create Primary Partition

Minimum size for the partition is 8 MB

Maximum size for the partition is 3271 MB

Create partition of size 1024 MB

Figure 84. The Create Primary Partition Window

In this example, we would like to create a 1,024 MB partition. We enter `1024` as shown in Figure 84 on page 136 and click **OK** to continue.

The Disk Administrator assigns the next available logical drive letter to the new unformatted partition, as shown in Figure 85 on page 136.

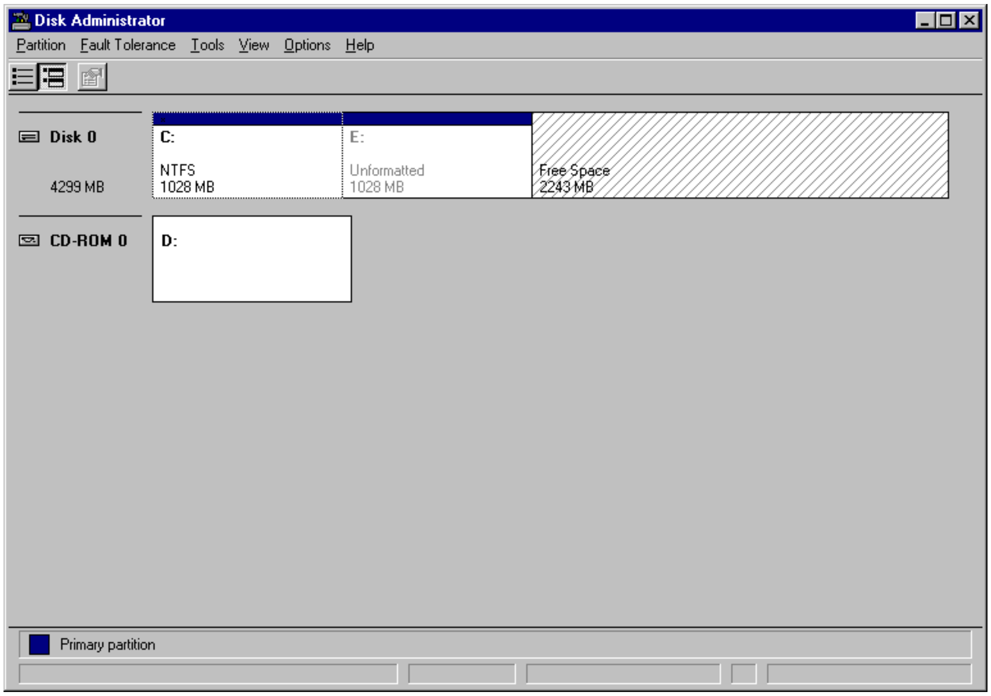

Figure 85. Assigning a Logical Drive Letter to the Unformatted Partition

5. Click the **Partition** button on the Disk Administrator menu bar, and click **Commit Changes Now...** to update the partition table, as shown in Figure 86 on page 137.

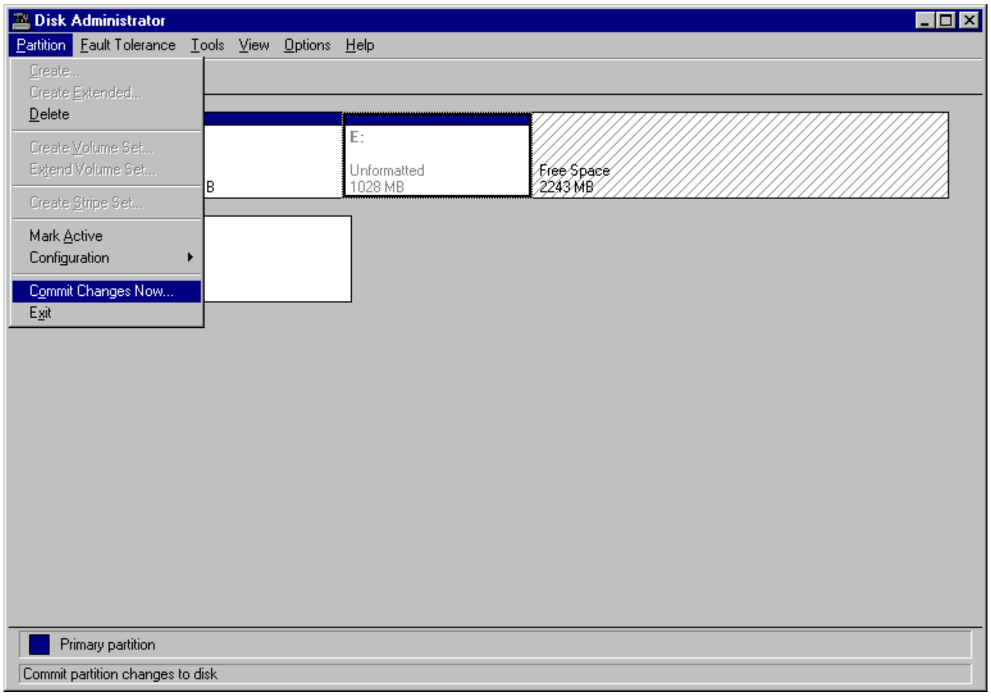

Commit partition changes to disk

Figure 86. Updating the Partition Table

6. A confirmation message appears as shown in Figure 87 on page 137. Click **Yes** to save the changes and continue.

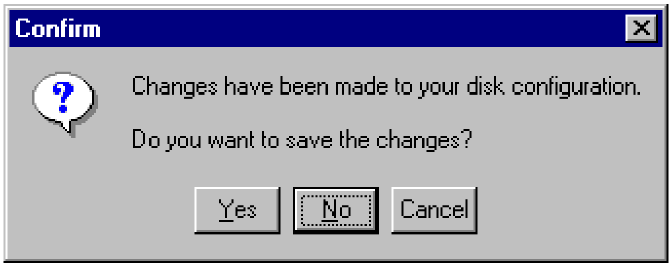

Changes have been made to your disk configuration.

Do you want to save the changes?

Figure 87. The Update Partition Table Confirmation Window

7. After updating the partition table, the Disk Administrator reminds you to update your Emergency Repair Disk, as shown in Figure 88 on page 137. Click **OK** to continue.

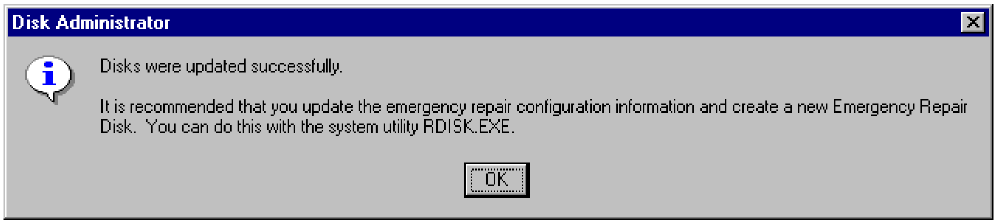

Disks were updated successfully.

It is recommended that you update the emergency repair configuration information and create a new Emergency Repair Disk. You can do this with the system utility RDISK.EXE.

Figure 88. The Emergency Repair Disk Reminder Window

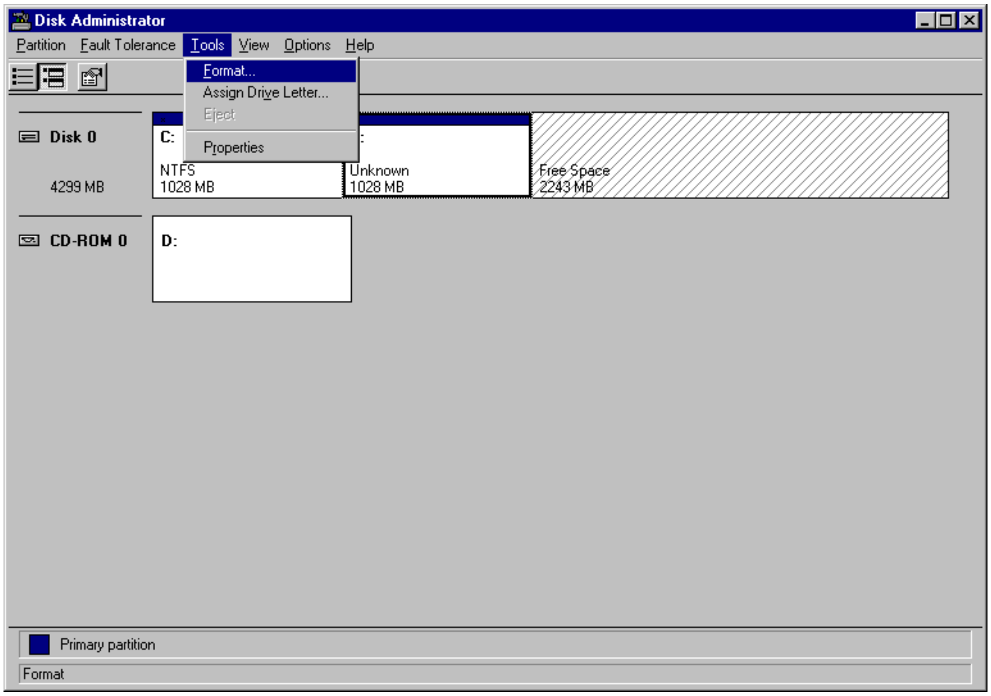

Figure 89. Invoking the Disk Partition Format Process

8. Click on the new partition area to select it. Then, click the **Tools** button on the Disk Administrator menu bar, and click **Format**, as shown in Figure 89 on page 138. The Format window will appear:

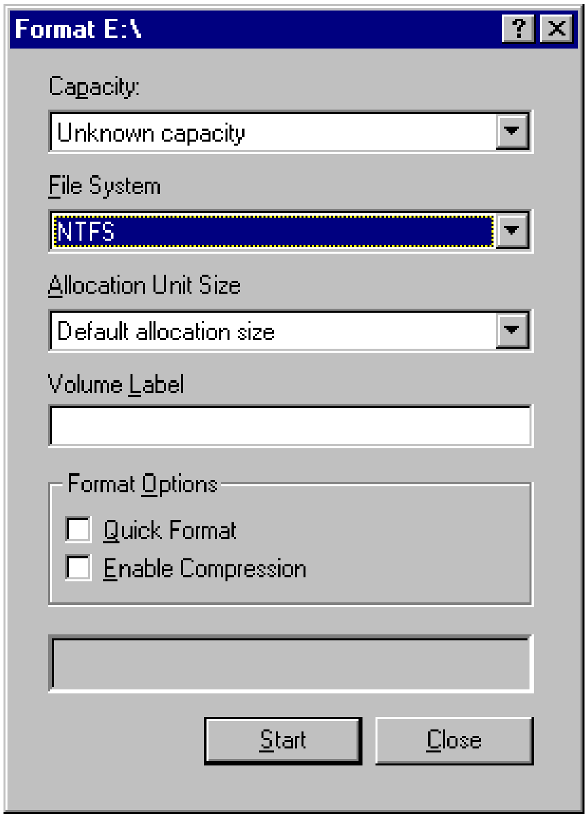

Format E:\

Capacity: Unknown capacity

File System: NTFS

Allocation Unit Size: Default allocation size

Volume Label:

Format Options

□ Quick Format

□ Enable Compression

Figure 90. The Disk Partition Format Window

9. In the Format window, from the File System scroll-down list, select **NTFS** to select NTFS as the file system type to be formatted. Click **Start** to continue.

10. The warning message appears as shown in Figure 91 on page 139. Click **OK** to start formatting the selected disk partition.

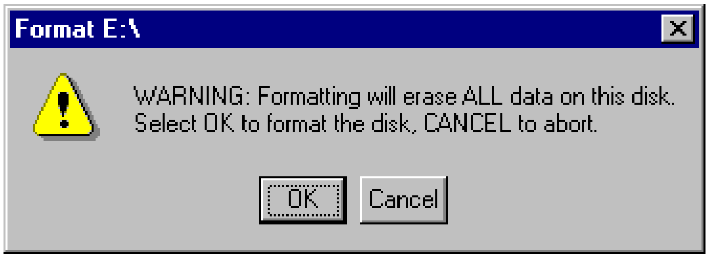

Format E:\

WARNING: Formatting will erase ALL data on this disk.
Select OK to format the disk, CANCEL to abort.

Figure 91. The Disk Partition Format Warning Window

The Disk Administrator displays its progress during the formatting process as shown in Figure 92 on page 140.

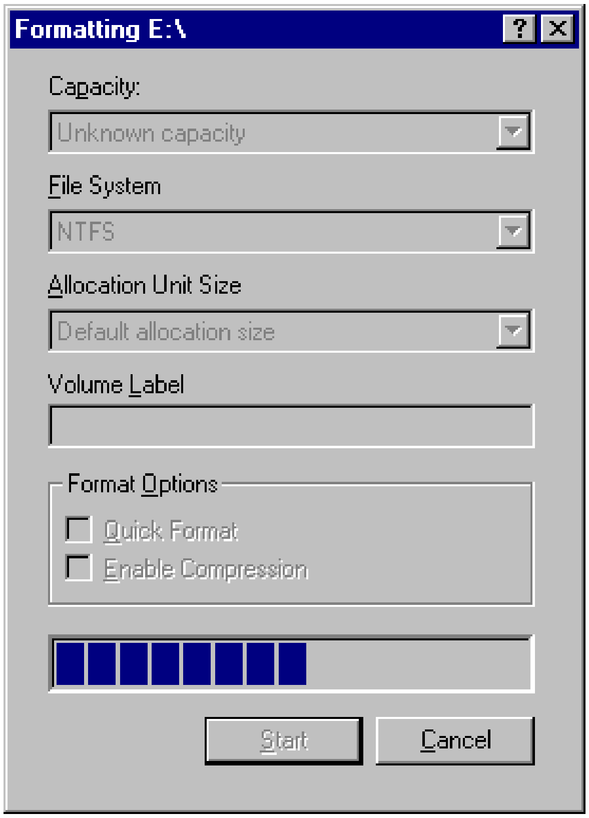

Capacity: Unknown capacity

File System: NTFS

Allocation Unit Size: Default allocation size

Volume Label:

Format Options

□ Quick Format

□ Enable Compression

Figure 92. The Disk Partition Format Progress Window

11.A window appears after successfully completing the formatting of this new NTFS partition, as shown in Figure 93.

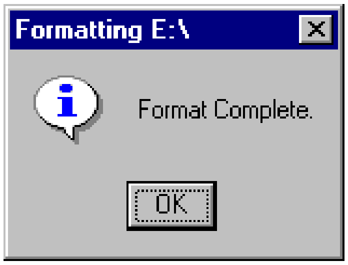

Format Complete.

Figure 93. The Format Complete Window

Click **OK** to continue. At this point, the Disk Administrator displays the new NTFS disk partition to be used as an AFS partition, as shown in Figure 94 on page 141.

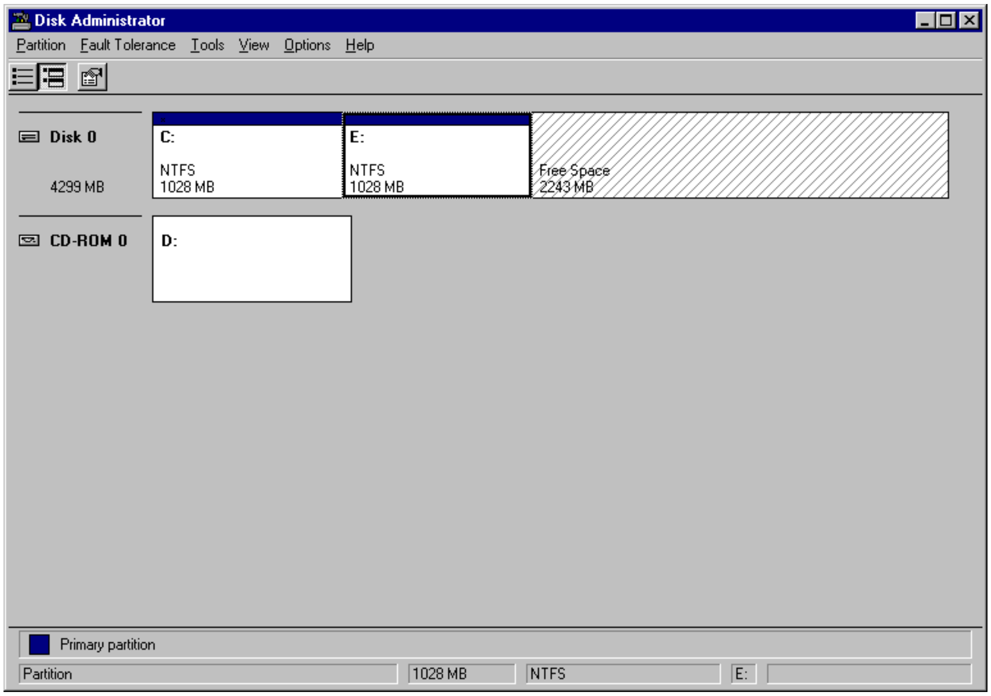

Partition | 1028 MB | NTFS | E:

Figure 94. The New NTFS Disk Partition

5.3 Installing AFS Files on Windows NT Systems

This section describes the procedure to install the necessary AFS files to configure a Windows NT machine to be an AFS machine, including the first AFS machine, an additional AFS Server, an AFS Client, and an AFS Control Center Machine.

To install IBM AFS Enterprise File System – the File Sharing component of WebSphere Performance Pack – on Windows NT systems, perform the following tasks:

1. Log in as an administrative user, who is a member of the local Administrators group on your Windows NT system.
2. Insert the IBM WebSphere Performance Pack for Multiplatforms CD-ROM in the CD-ROM drive.

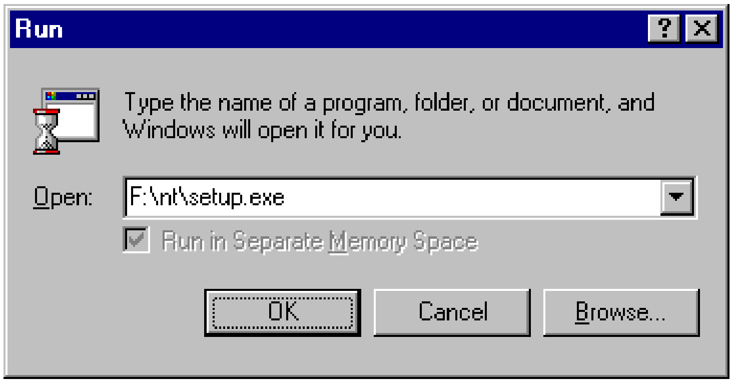

Type the name of a program, folder, or document, and Windows will open it for you.

Open: F:\nt\setup.exe

☑ Run in Separate Memory Space

Figure 95. Starting the WebSphere Performance Pack Setup Program

3. Click the **Start** button on the Windows NT desktop taskbar and select **Run...**. In the Run dialog box, click **Browse...** to search for the WebSphere Performance Pack Setup program. This Setup program, setup.exe, is located in the NT folder of the IBM WebSphere Performance Pack for Multiplatforms CD-ROM. After specifying the path of the Setup program, click **OK** to start the installation process.

   Notice that if JRE 1.1.6 (or higher) is not installed on your machine, you will be prompted to install the JDK from the CD-ROM. Notice that if you have Java 2 SDK, Standard Edition, V1.2 or higher installed on your machine, this will not be detected as a valid JRE by the Setup program, and you will be prompted to install JDK 1.1.6.

   Once Java is installed on your machine, the Setup program can proceed. Another option to run the Setup program once Java is installed on your machine is to access the nt directory of the CD-ROM from a command prompt and launch the command:

   ```
   java setup
   ```

4. The Setup program displays the Welcome window:

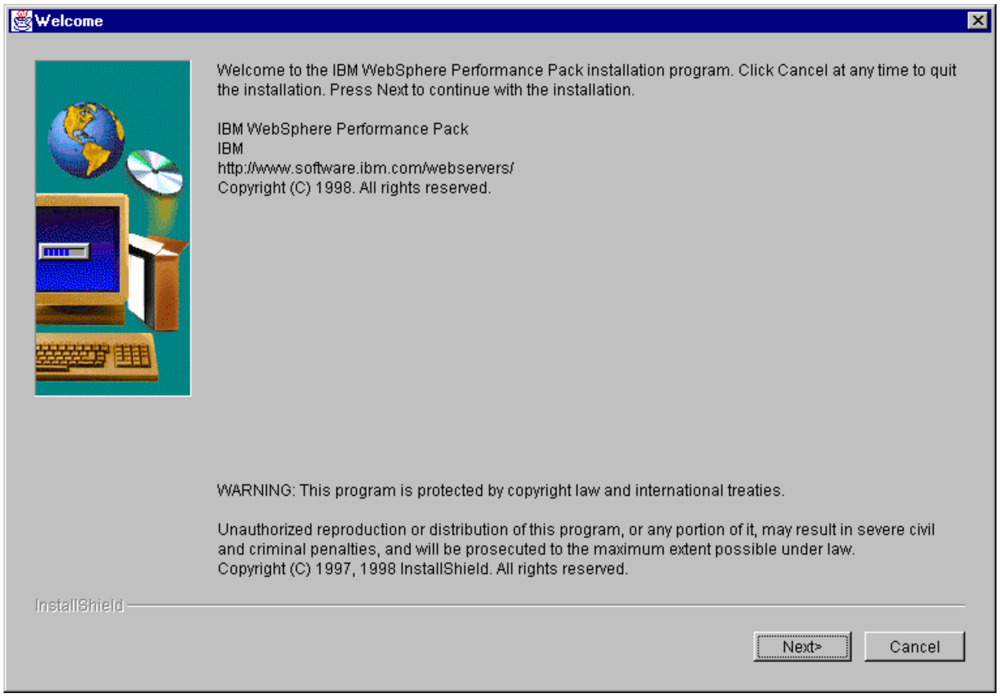
Welcome

Welcome to the IBM WebSphere Performance Pack installation program. Click Cancel at any time to quit the installation. Press Next to continue with the installation.

IBM WebSphere Performance Pack
IBM
http://www.software.ibm.com/webservers/
Copyright (C) 1998. All rights reserved.

WARNING: This program is protected by copyright law and international treaties.

Unauthorized reproduction or distribution of this program, or any portion of it, may result in severe civil and criminal penalties, and will be prosecuted to the maximum extent possible under law.
Copyright (C) 1997, 1998 InstallShield. All rights reserved.

Figure 96. The IBM WebSphere Performance Pack Welcome Window

Click **Next** to continue.

5. The README Information window appears, displaying the latest information on IBM WebSphere Performance Pack. Click **Next** to continue.
6. The Choose Destination Location window appears:

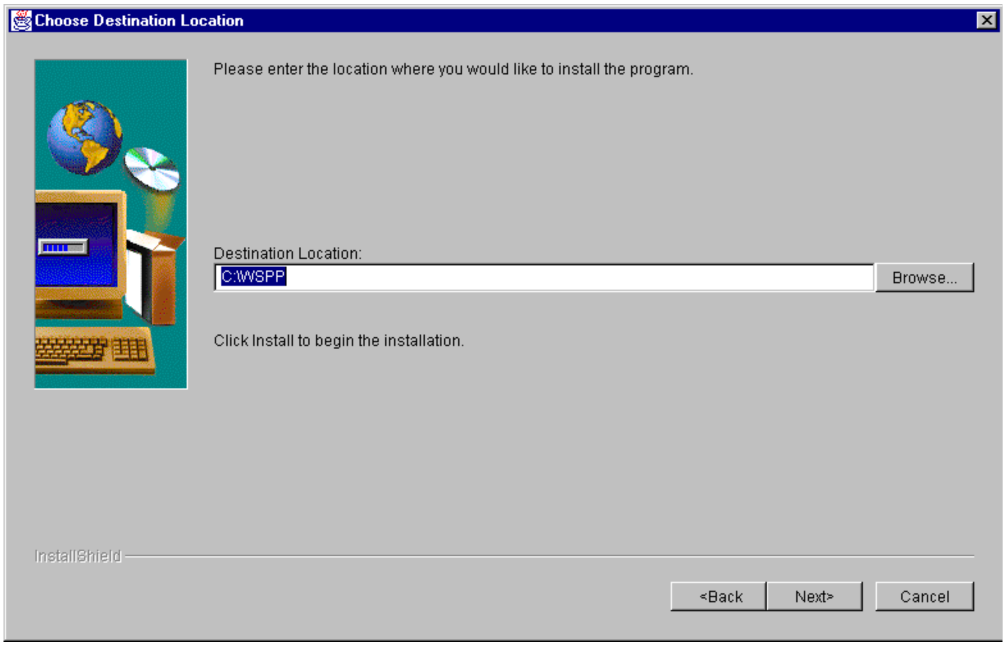
Choose Destination Location

Please enter the location where you would like to install the program.

Destination Location: C:\WSPP

Click Install to begin the installation.

Figure 97. Selecting the Destination Location for WebSphere Performance Pack

You need to choose the folder in which to install the AFS files. The default destination location is C:\WSPP.

If you want to select a different location for installing the AFS files, click **Browse...**. We chose the default destination location, C:\WSPP. The Setup program incorrectly instructs us to click **Install** to begin the installation. There is no Install button.[1] Click **Next** to continue.

If the specified destination location does not exist, a dialog box appears to inquire whether you would like to have it created. If necessary, click **Yes** to create this directory.

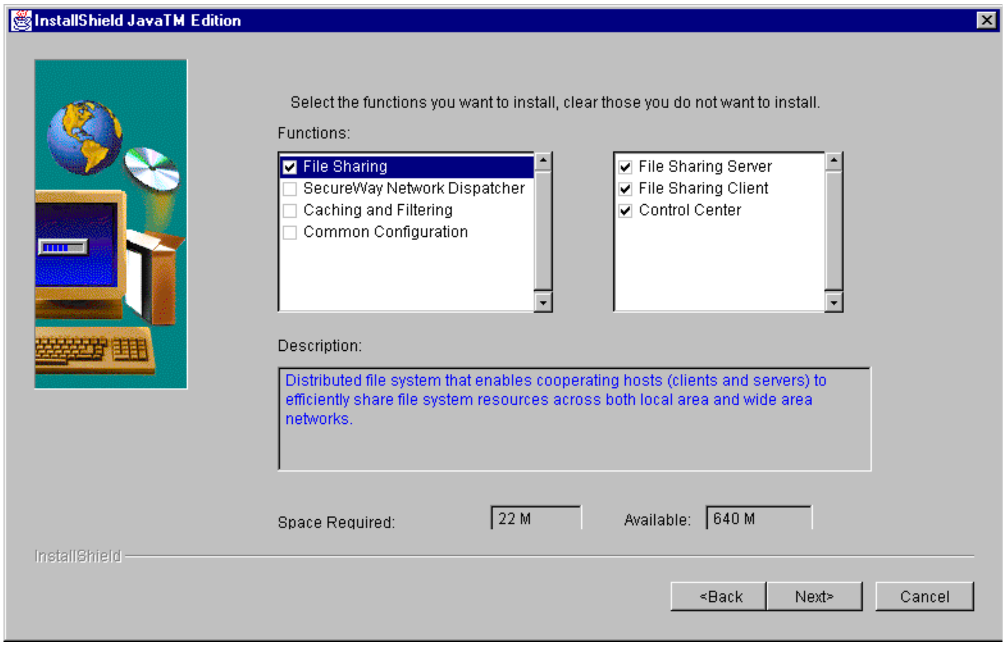

Select the functions you want to install, clear those you do not want to install.

Functions:

☑ File Sharing
□ SecureWay Network Dispatcher
□ Caching and Filtering
□ Common Configuration

☑ File Sharing Server
☑ File Sharing Client
☑ Control Center

Description:

Distributed file system that enables cooperating hosts (clients and servers) to efficiently share file system resources across both local area and wide area networks.

Space Required: 22 M

Available: 640 M

Figure 98. The Performance Pack Component Selection Dialog Box

7. The Setup program displays a dialog box containing the names of the WebSphere Performance Pack components. Select the **File Sharing** check box and a list of File Sharing functions will appear in the adjacent window.

   - If you are installing the first AFS machine, you must select both the **File Sharing Server** and **File Sharing Client**. The File Sharing Server cannot be installed without the File Sharing Client. Select the **File Sharing Server** check box and the **File Sharing Client** check box is set automatically. Optionally, you can also select the **Control Center** function.
   - If you are installing another AFS Server machine, you must select both the **File Sharing Server** and **File Sharing Client**.
   - If you are installing an AFS Client machine, you only have to select the **File Sharing Client**.
   - If you are installing the AFS Control Center, you only have to select the **Control Center**.

   You need to choose one of the above selections and click **Next** to continue.

[1] This is a known Java InstallShield problem. It is not a bug in the WebSphere Performance Pack installation procedure.

8. The Setup program displays a window providing you with the option of replacing the current version of any WebSphere Performance Pack product that has already been installed.

   - If you are installing AFS and there are no other WebSphere Performance Pack products residing on this machine, check **Yes** to replace any copies of the File Sharing components eventually installed in the selected destination location.
   - If you are not sure, check **No** to keep any current copies.

   Then, click **Install** to start the installation of the chosen File Sharing components from the WebSphere Performance Pack for Multiplatforms CD-ROM to local disk. An Installing window appears showing the progress of the installation.

9. The Installation Complete window appears upon successful completion of the extraction of files from the CD-ROM distribution, as shown in Figure 99 on page 145.

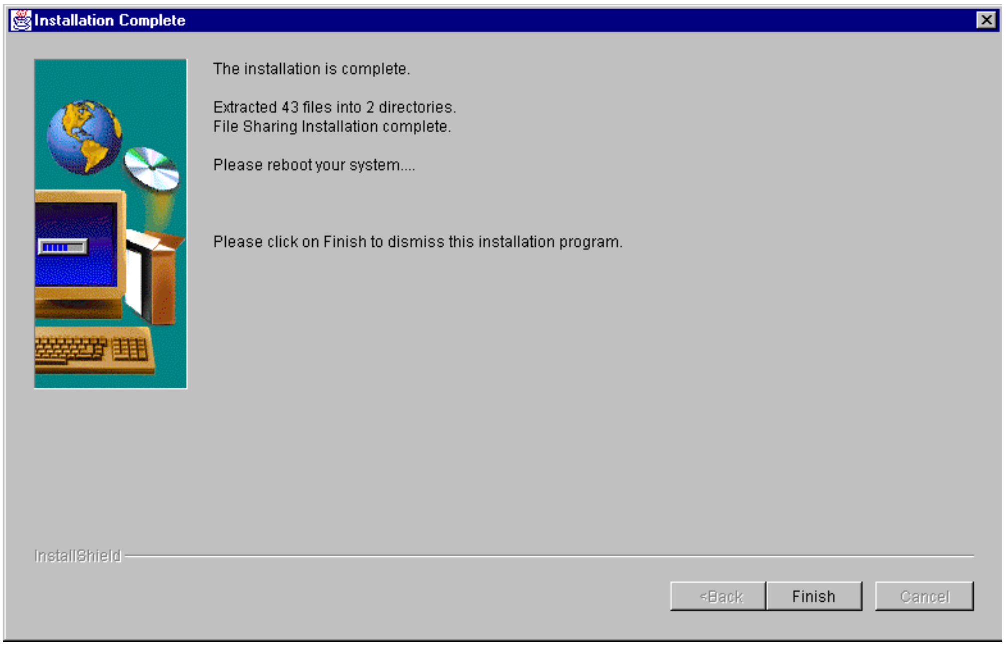

Installation Complete

The installation is complete.

Extracted 43 files into 2 directories.
File Sharing Installation complete.

Please reboot your system....

Please click on Finish to dismiss this installation program.

Figure 99. The Installation Complete Window

   Click **Finish** to complete the installation process.

10. Reboot your machine to apply the system configuration changes to your Windows NT system.

5.3.1 WebSphere AFS Program Group and Directories

After installing the AFS files from the CD-ROM distribution to local disk, the WebSphere Performance Pack Setup program creates a Program Group called IBM WebSphere, as shown in Figure 100 on page 146.

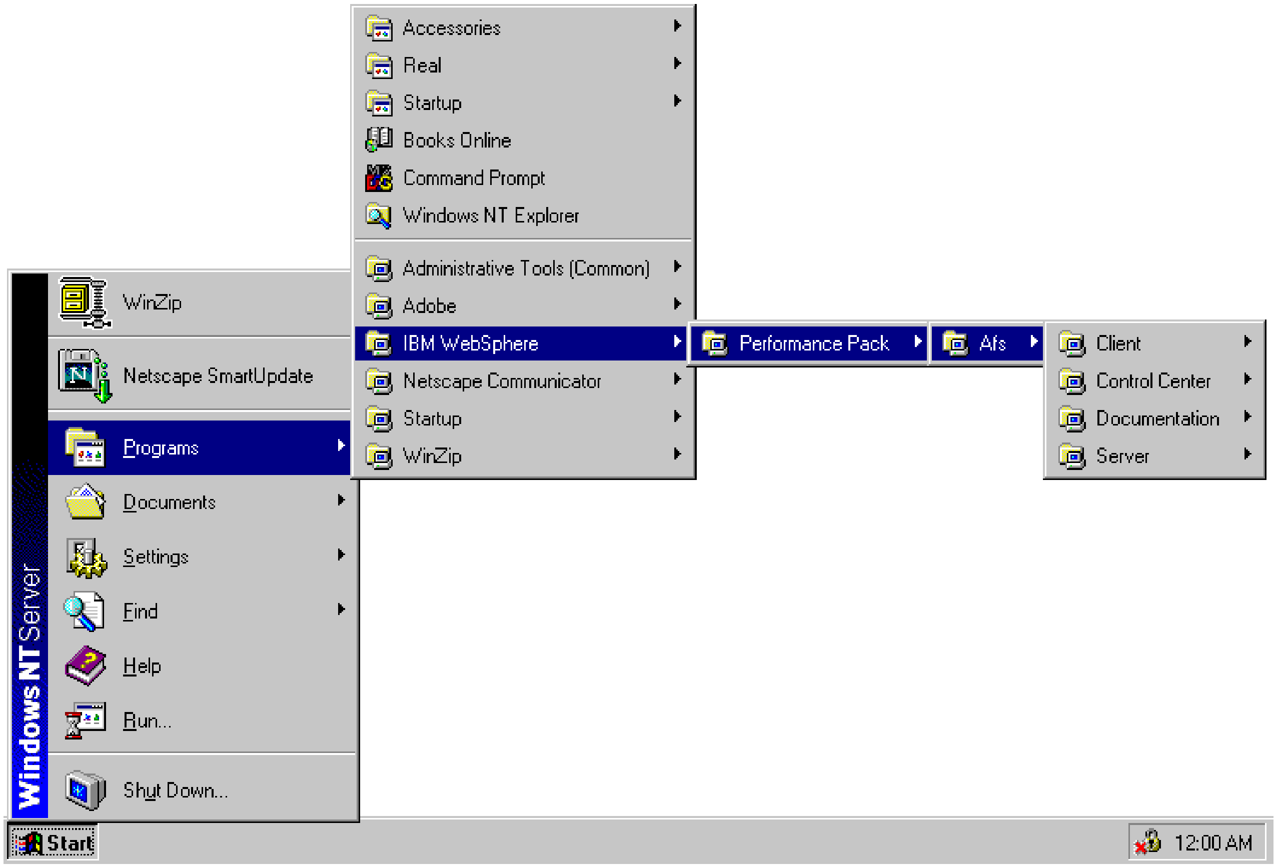

Figure 100. AFS Program Group and Client Icon

In the Windows NT Control Panel, the Setup program creates the AFS Server Configuration, AFS Client Configuration and AFS Control Center applications, depending on which AFS functions you have installed on your system. For a Windows NT machine running only the AFS Client, it creates only the AFS Client Configuration application. For a Windows NT machine configured with only the AFS Control Center, it creates only the AFS Control Center application.

In the Services Control Panel, the Setup program creates IBM AFS Server and IBM AFS Client entries. For a Windows NT machine running only the AFS Client, it creates only an IBM AFS Client entry. For a Windows NT machine configured with only the AFS Control Center, there is no entry created in the Services Control Panel.

For a Windows NT machine configured with the AFS Client, the AFS Client displays an AFS Client icon on the Windows NT desktop taskbar, as shown in Figure 100 on page 146. You can change this default configuration option in the AFS Client Configuration application from the Windows NT Control Panel.

The Setup program creates an AFS subdirectory under the WebSphere Performance Pack installation directory. By default, the WebSphere Performance Pack AFS installation directory is C:\WSPP\AFS. This directory is organized as follows:

- C:\WSPP\AFS\Server

This directory contains AFS server binaries, databases, configuration files, and log files, and is structured through the following paths:

  - C:\WSPP\AFS\Server\usr\afs\bin

    This directory contains AFS binaries, including server executables, system administrator and user commands.

  - C:\WSPP\AFS\Server\usr\afs\etc

    This directory contains common server configuration files, such as a list of database server machines, the local cell name, the server encryption keys, and a list of system administrators in the cell. Initially, this directory is empty.

  - C:\WSPP\AFS\Server\usr\afs\local

    This directory contains local server configuration files specific for each AFS server machine. Initially, this directory is empty.

  - C:\WSPP\AFS\Server\usr\afs\db

    This directory contains database files. Initially, this directory is empty.

  - C:\WSPP\AFS\Server\usr\afs\logs

    This directory contains server log files, including the log file for the AFS Server Configuration Wizard. Initially, this directory is empty.

- C:\WSPP\AFS\Client

  This directory contains AFS client binaries and configuration files.

- C:\WSPP\AFS\Control Center

  This directory contains the User Account Manager and the Server Manager binaries.

- C:\WSPP\AFS\Common

  This directory contains libraries and modules shared by AFS binaries

- C:\WSPP\AFS\Documentation

  This directory contains release notes, information, and documentation

5.3.2 WebSphere AFS Server Configuration Wizard

The AFS component of WebSphere Performance Pack for Windows NT includes the AFS Server Configuration Wizard to facilitate the configuration of the AFS Server on your Windows NT system. The AFS Server Configuration Wizard takes you through the steps of configuring the AFS Server via a sequence of dialog boxes. This Wizard assists you in the following tasks:

- Configure the first AFS machine in a new AFS cell
- Configure additional AFS Server machines of an existing AFS cell

The Wizard provides you with options of configuring an AFS Server machine to be an AFS File Server, an AFS Database Server, and/or an AFS Backup Server.

5.4 Configuring the First AFS Machine on Windows NT Systems

Before you can configure a Windows NT machine to be the first AFS machine, you must complete the following tasks:

- Create at least one dedicated NTFS partition to be configured as an AFS partition to store AFS volumes, as illustrated in 5.2, “Configuring AFS Partitions on Windows NT Systems” on page 133. This NTFS partition must not contain any files other than the Windows NT Recycle Bin.
- Install the necessary AFS files from the WebSphere Performance Pack CD-ROM distribution to the machine local disk, as shown in 5.3, “Installing AFS Files on Windows NT Systems” on page 141.

AFS Logon Message

In the early release of the WebSphere AFS component, at the first login after installing AFS files, the system displays a warning message informing you that the specified realm is unknown. Since you have not configured AFS on this machine yet, click **OK** to continue. In the General Availability (GA) version of the WebSphere AFS component, the system no longer displays this warning message.

To illustrate the configuration of the first AFS machine on a Windows NT system, we select the AFS cell name, moshi.itso.ral.ibm.com. We use an IBM Netfinity 3000 machine called wtr05319.itso.ral.ibm.com, with the IP address of 9.24.104.206, as shown in the following output:

```
>moshi.itso.ral.ibm.com     #IBM ITSO WSPP
9.24.104.206                #wtr05319.itso.ral.ibm.com
```

5.4.1 Using the AFS Server Configuration Wizard

To start the configuration of the first AFS machine on a Windows NT system, perform the following tasks:

1. Log in as an administrative user, who is a member of the local Administrators group on your Windows NT system.
2. To start the AFS Server Configuration Wizard, click the **Start** button on your Windows NT desktop taskbar, point to **Programs**, point to **IBM WebSphere**, point to **Performance Pack**, point to **Afs**, point to **Server**, and click **Configuration Wizard**, as depicted in Figure 101 on page 149.

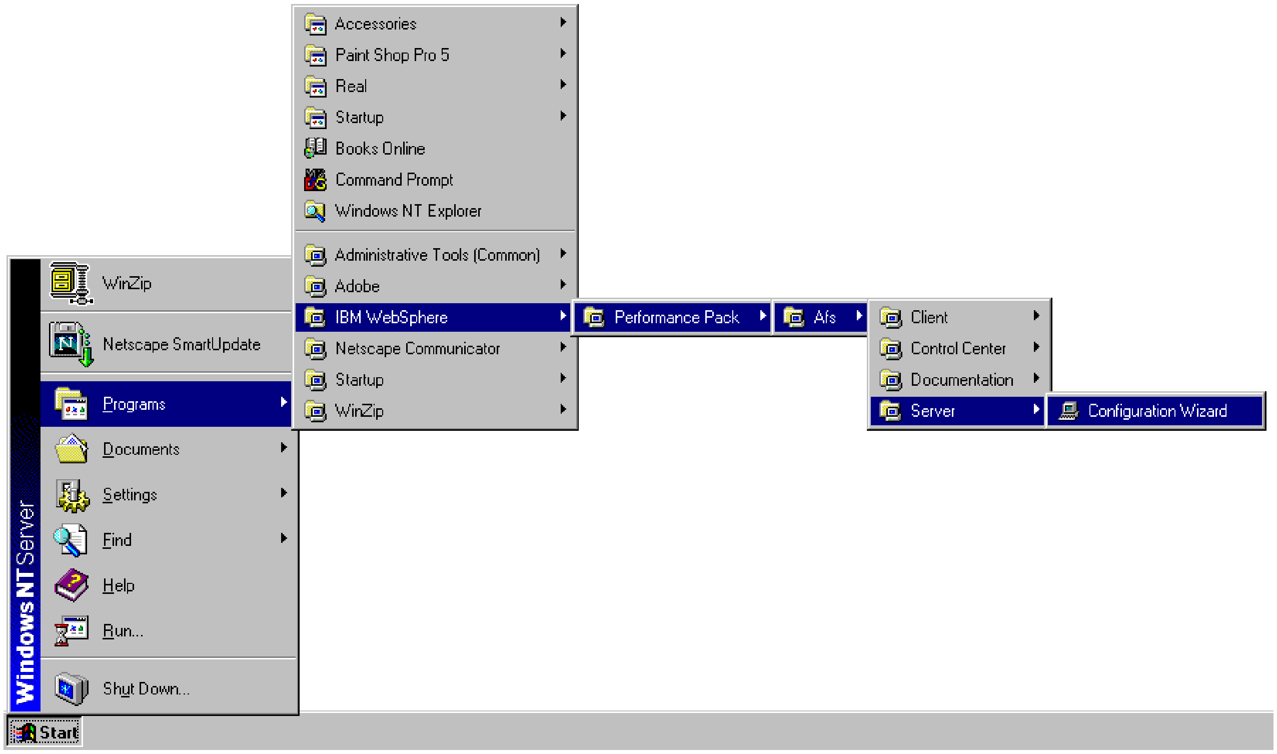

Figure 101. Starting the AFS Server Configuration Wizard

3. The Welcome to the AFS Server Quick Start Wizard window appears. Click **Next** to continue.

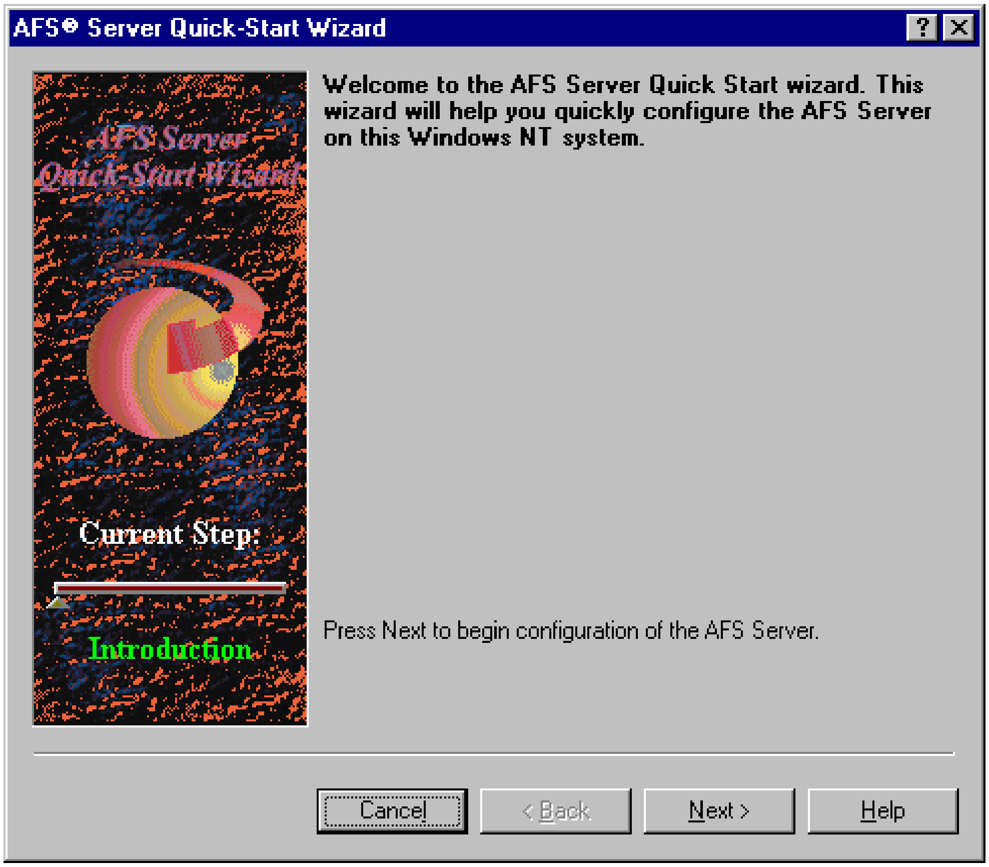
AFS® Server Quick-Start Wizard

Welcome to the AFS Server Quick Start wizard. This wizard will help you quickly configure the AFS Server on this Windows NT system.

Press Next to begin configuration of the AFS Server.

Figure 102. The AFS Server Configuration Wizard Welcome Window

4. In the Cell and Server Information step window, the Wizard provides you with the option of configuring this machine as an AFS server in an existing AFS cell or as the first AFS server in a new AFS cell. Select the **This will be the first server in a new AFS cell** check box. In the Cell Name frame, enter the name of the AFS cell. In this example, the AFS cell name is moshi.itso.ral.ibm.com.

   In the Server Password frame, enter the password for the AFS principal in the Password field. Enter the same password in the Verify password field. For security reasons, the Wizard masks the password string as you type it in.

   The following figure shows our sample configuration:

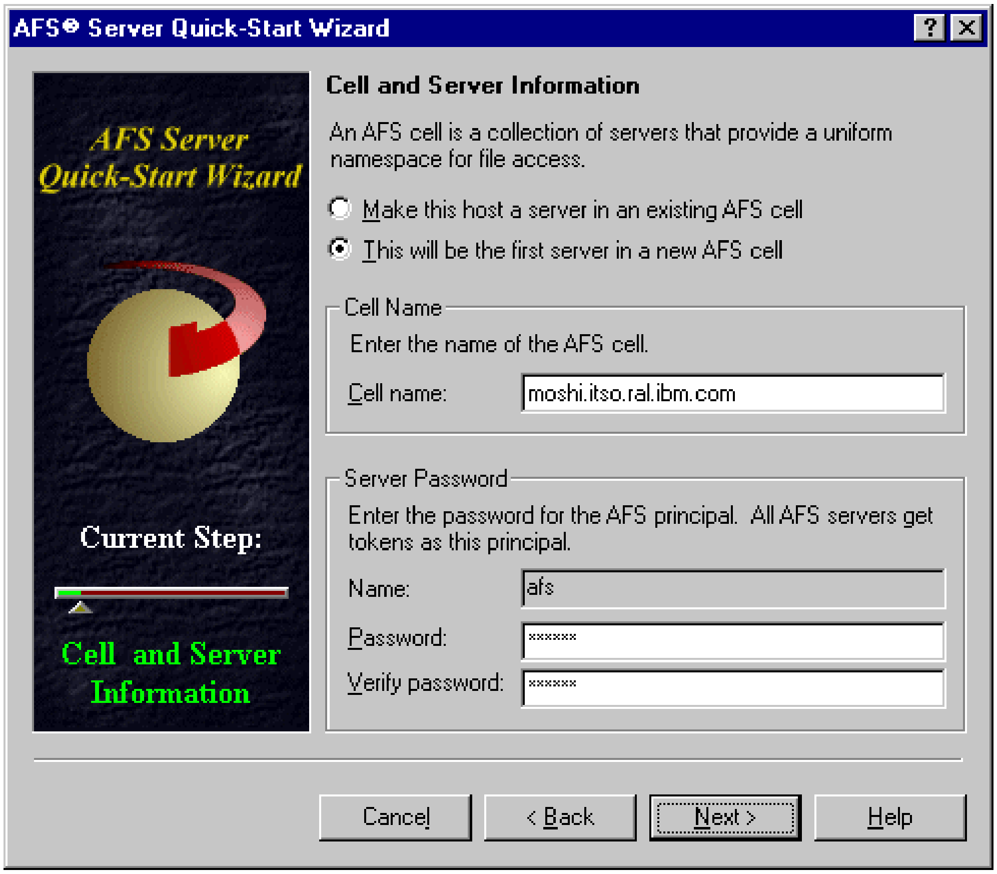
AFS® Server Quick-Start Wizard

Cell and Server Information

An AFS cell is a collection of servers that provide a uniform namespace for file access.

○ Make this host a server in an existing AFS cell

● This will be the first server in a new AFS cell

Cell Name

Enter the name of the AFS cell.

Cell name: moshi.itso.ral.ibm.com

Server Password

Enter the password for the AFS principal. All AFS servers get tokens as this principal.

Name: afs

Password: ******

Verify password: ******

Figure 103. The Cell and Server Information Window

The AFS principal holds the AFS server keys for AFS server processes to perform mutual authentications. You will never authenticate to this account.

Click **Next** to continue.

5. In the Administrative Information window, enter the password for the AFS admin account in the Password field. Enter the same password in the Verify password field. Once again, for security reasons, the Wizard masks the password string, as shown in the following screen:

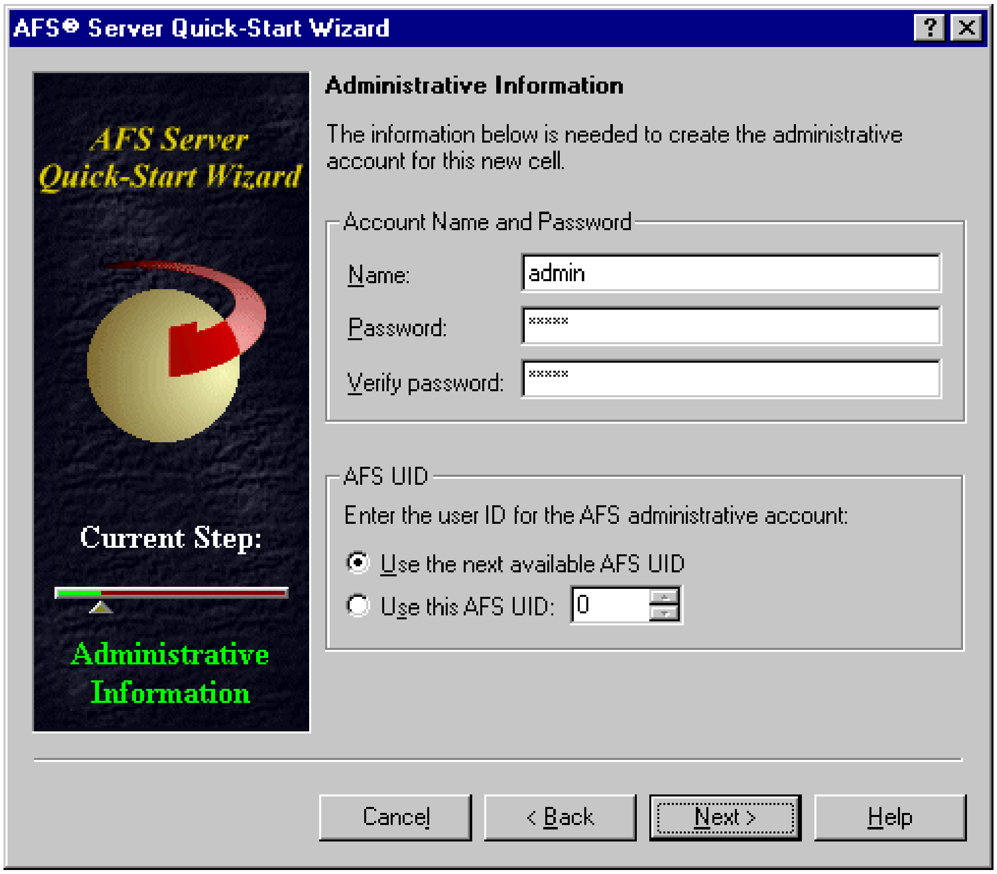

AFS® Server Quick-Start Wizard
AFS Server Quick-Start Wizard
Current Step:
Administrative Information
Cancel
< Back
Next >
Help

Administrative Information

The information below is needed to create the administrative account for this new cell.

Account Name and Password

Name: admin

Password: *****

Verify password: *****

AFS UID

Enter the user ID for the AFS administrative account:

● Use the next available AFS UID

○ Use this AFS UID: 0

Figure 104. The Administrative Information Window

Although the Wizard does not insist that the admin password be different from the afs password, entered in Step 4 on page 150, it is a good idea to choose a unique password for both these system accounts. You need to remember the admin password to authenticate to this account and perform AFS administrative tasks!

The Wizard provides you with the option of using the next available AFS UID number or assigning an explicit AFS UID number for the admin account. The default option is to allow the Wizard to use the next available AFS UID. Click **Next** to continue.

6. In the File Service step window, the Wizard informs you that the AFS File Service will be configured automatically for the first AFS machine:

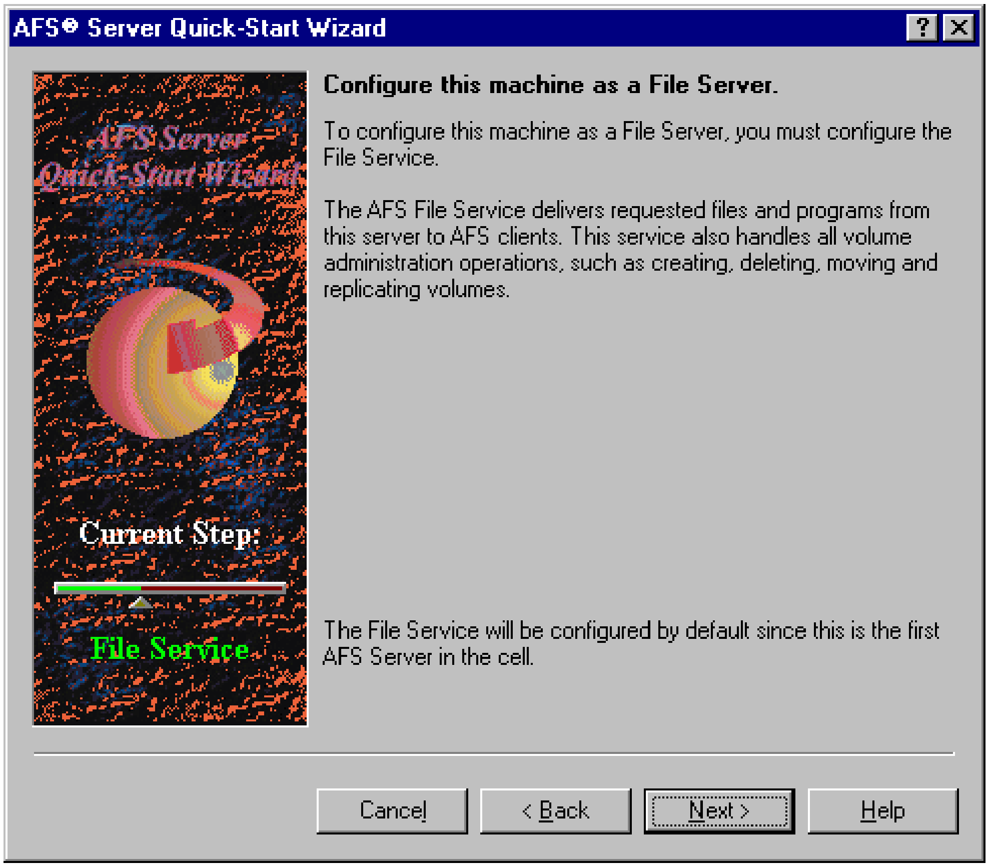

Configure this machine as a File Server.

To configure this machine as a File Server, you must configure the File Service.

The AFS File Service delivers requested files and programs from this server to AFS clients. This service also handles all volume administration operations, such as creating, deleting, moving and replicating volumes.

The File Service will be configured by default since this is the first AFS Server in the cell.

Figure 105. The File Server Configuration Window

Click **Next** to continue.

7. In the Database Service step window, the Wizard informs you that the AFS Database Service will be configured automatically for the first AFS machine:

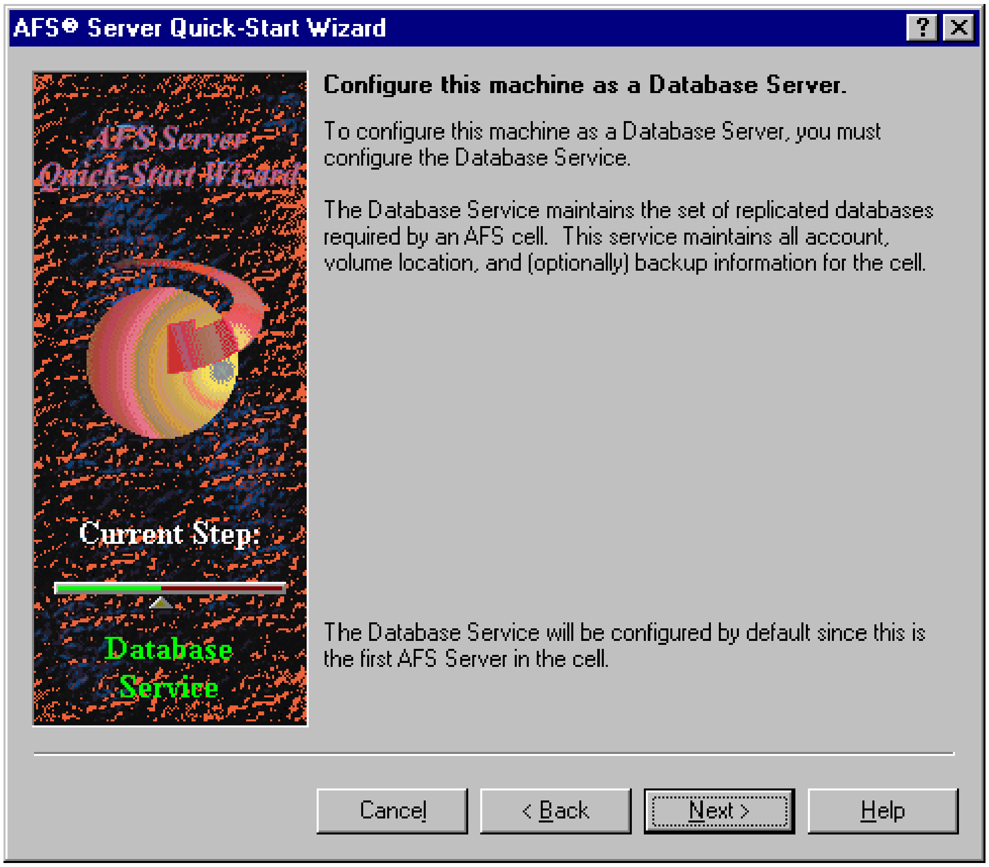

Configure this machine as a Database Server.

To configure this machine as a Database Server, you must configure the Database Service.

The Database Service maintains the set of replicated databases required by an AFS cell. This service maintains all account, volume location, and (optionally) backup information for the cell.

The Database Service will be configured by default since this is the first AFS Server in the cell.

Figure 106. The Database Server Configuration Window

Click **Next** to continue.

8. In the Backup Server step window, the Wizard provides you with the option of configuring this machine as an AFS Backup Server. The default option is to configure the AFS Backup service, as shown in the following figure:

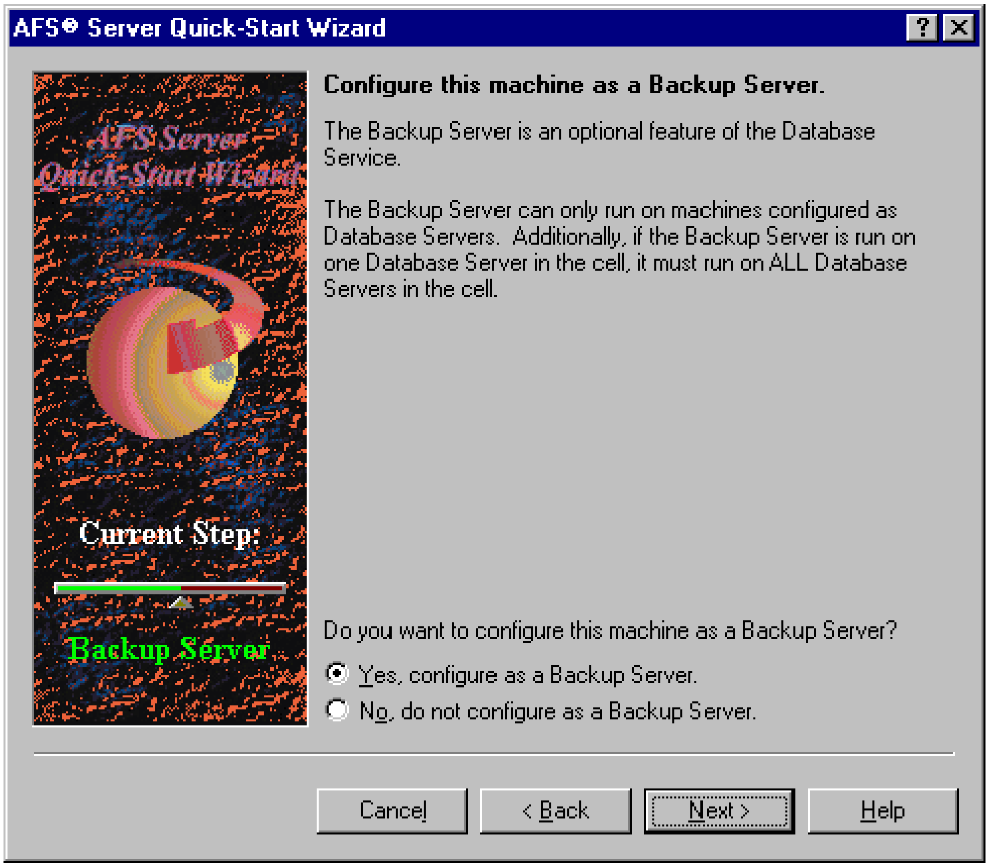

Configure this machine as a Backup Server.

The Backup Server is an optional feature of the Database Service.

The Backup Server can only run on machines configured as Database Servers. Additionally, if the Backup Server is run on one Database Server in the cell, it must run on ALL Database Servers in the cell.

Do you want to configure this machine as a Backup Server?

● Yes, configure as a Backup Server.

○ No, do not configure as a Backup Server.

Figure 107. The Backup Server Configuration Window

Once you have made your decision, click **Next** to continue.

9. In the AFS Partition step window, the Wizard displays all Windows NT logical drives (volumes) that are possibly eligible to be configured as AFS partitions:

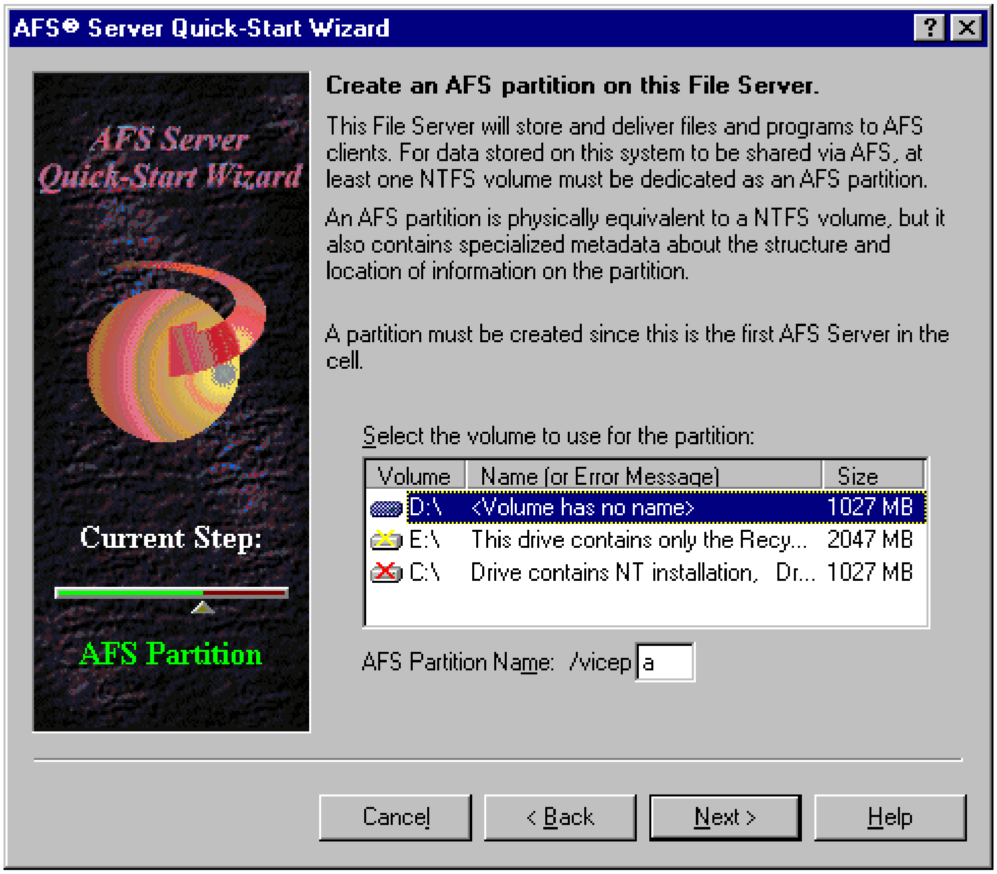

Create an AFS partition on this File Server.

This File Server will store and deliver files and programs to AFS clients. For data stored on this system to be shared via AFS, at least one NTFS volume must be dedicated as an AFS partition.

An AFS partition is physically equivalent to a NTFS volume, but it also contains specialized metadata about the structure and location of information on the partition.

A partition must be created since this is the first AFS Server in the cell.

Select the volume to use for the partition:

Volume	Name (or Error Message)	Size
D:\	<Volume has no name>	1027 MB
E:\	This drive contains only the Recy...	2047 MB
C:\	Drive contains NT installation, Dr...	1027 MB

AFS Partition Name: /vicep a

Figure 108. The Create AFS Partition Window

The Wizard only allows NTFS partitions containing no data or containing only the Windows NT Recycle Bin to be configured as AFS partitions.

In this example, the D drive is an NTFS partition without any data, so is eligible to be configured as an AFS partition. Drive E is also available to be configured as an AFS partition, since this drive contains only the Recycle Bin. On the other hand, the C drive is not eligible to be configured as an AFS partition, because this drive contains the Windows NT operating system. We select **D:\**.

By convention, the first AFS partition on a machine is called /vicepa, the second partition /vicepb, etc. Enter the letter `a` in the AFS Partition Name box to assign /vicepa as the AFS partition name to the selected NTFS partition.

Click **Next** to continue.

10.In the Create Root Volumes step window, the Wizard informs you that the root.afs and root.cell volumes will be created automatically on the first AFS machine:

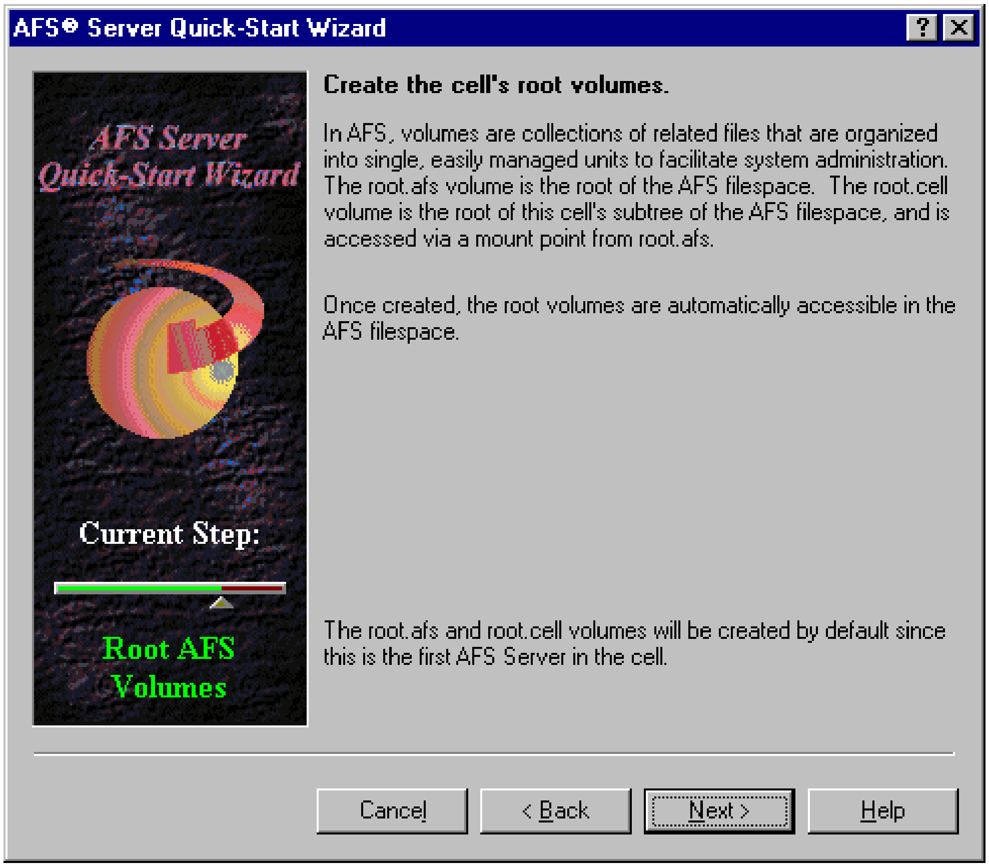

Create the cell's root volumes.

In AFS, volumes are collections of related files that are organized into single, easily managed units to facilitate system administration. The root.afs volume is the root of the AFS filespace. The root.cell volume is the root of this cell's subtree of the AFS filespace, and is accessed via a mount point from root.afs.

Once created, the root volumes are automatically accessible in the AFS filespace.

The root.afs and root.cell volumes will be created by default since this is the first AFS Server in the cell.

Figure 109. The Create Root Volumes Window

Click **Next** to continue.

11. In the Replication step window, the Wizard informs you that the root.afs and root.cell volumes will be replicated automatically on the first AFS machine:

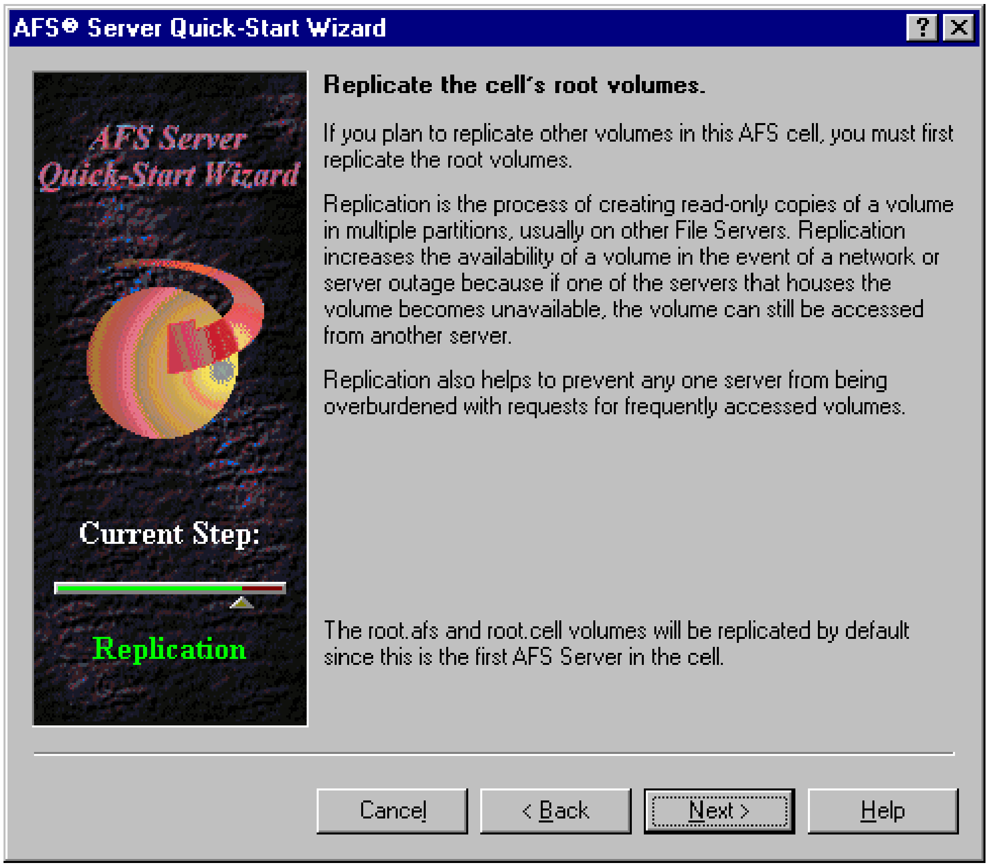

Replicate the cell's root volumes.

If you plan to replicate other volumes in this AFS cell, you must first replicate the root volumes.

Replication is the process of creating read-only copies of a volume in multiple partitions, usually on other File Servers. Replication increases the availability of a volume in the event of a network or server outage because if one of the servers that houses the volume becomes unavailable, the volume can still be accessed from another server.

Replication also helps to prevent any one server from being overburdened with requests for frequently accessed volumes.

The root.afs and root.cell volumes will be replicated by default since this is the first AFS Server in the cell.

Figure 110. The Replicate Root Volumes Window

Click **Next** to continue.

12.In the System Control Service step window, the Wizard provides you with the option of configuring this machine as the System Control Server:

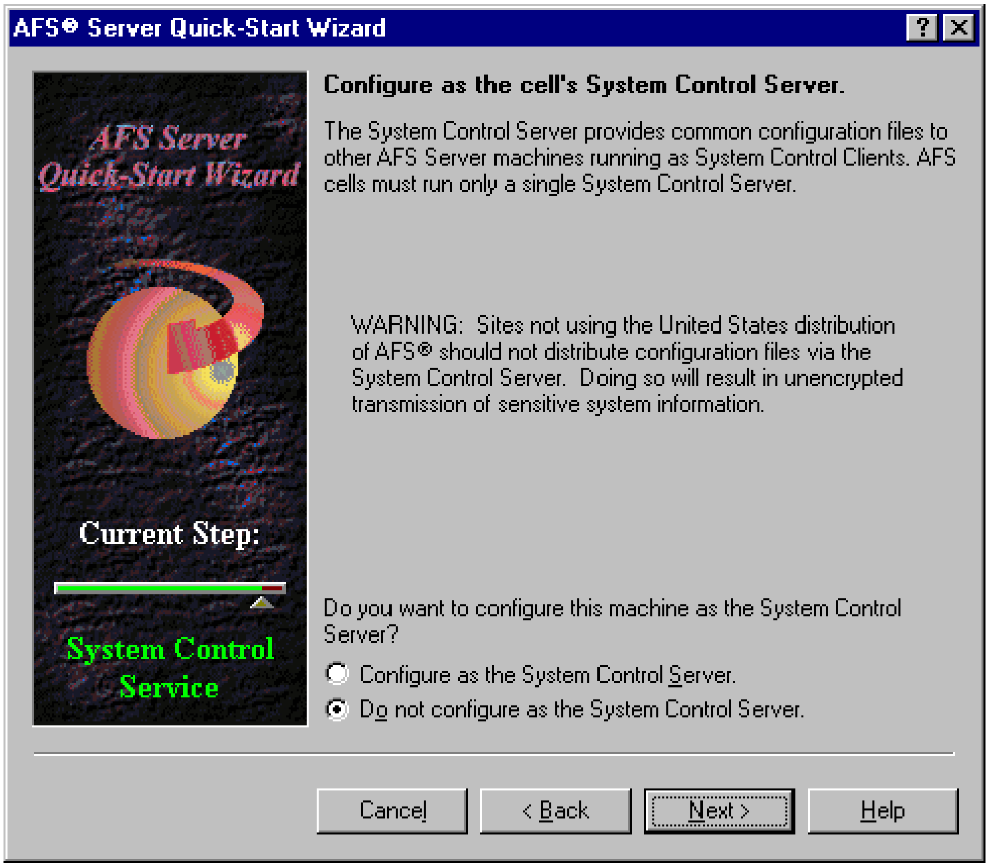

Configure as the cell's System Control Server.

The System Control Server provides common configuration files to other AFS Server machines running as System Control Clients. AFS cells must run only a single System Control Server.

WARNING: Sites not using the United States distribution of AFS® should not distribute configuration files via the System Control Server. Doing so will result in unencrypted transmission of sensitive system information.

Do you want to configure this machine as the System Control Server?

○ Configure as the System Control Server.

● Do not configure as the System Control Server.

Figure 111. The System Control Server Configuration Window

Since IBM WebSphere Performance Pack only provides the international edition of AFS, the System Control Server can only distribute files in unencrypted form. Therefore, we do not recommend that you configure this machine as a System Control Server to distribute AFS server common configuration files. The default option is also to not configure the System Control Service.

Click **Next** to continue.

13.In the Configure the System step window, the Wizard displays all the tasks that will be performed to configure this Windows NT machine as the first AFS machine:

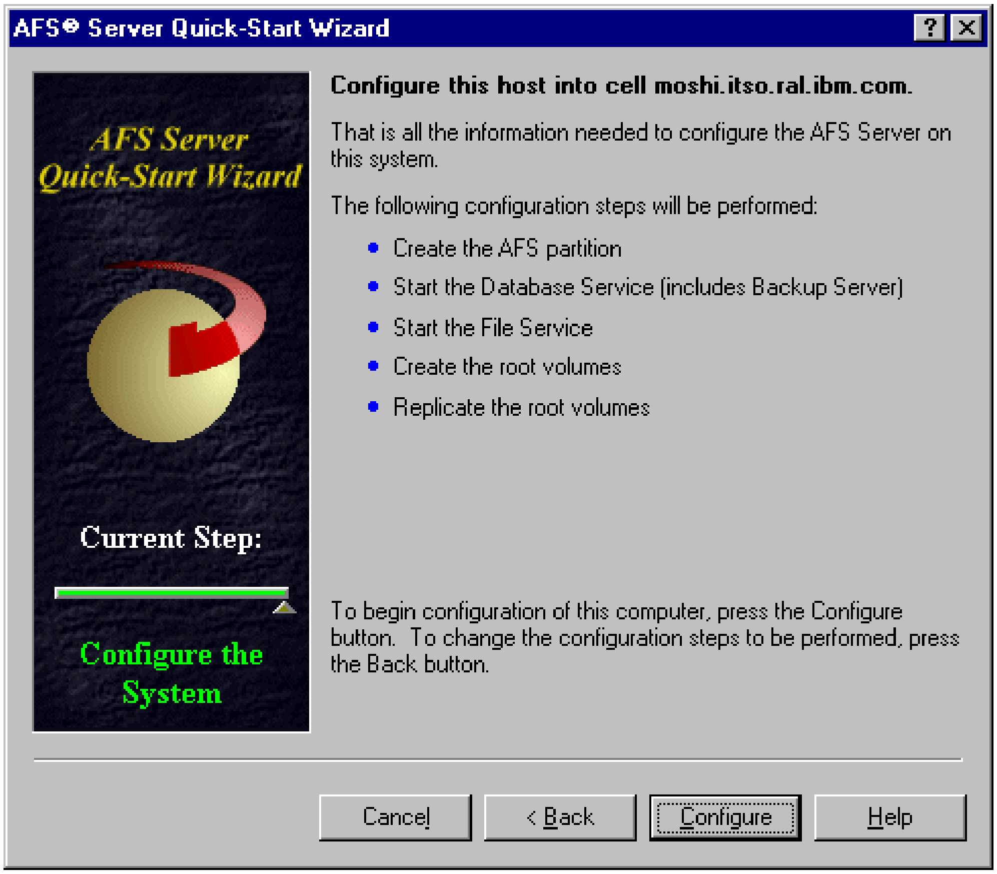

Configure this host into cell moshi.itso.ral.ibm.com.

That is all the information needed to configure the AFS Server on this system.

The following configuration steps will be performed:

- Create the AFS partition
- Start the Database Service (includes Backup Server)
- Start the File Service
- Create the root volumes
- Replicate the root volumes

To begin configuration of this computer, press the Configure button. To change the configuration steps to be performed, press the Back button.

Figure 112. The Server Configuration Window

Click **Configure** to begin the configuration.

14.The progress of the configuration will be displayed at each step, as shown in the following figure:

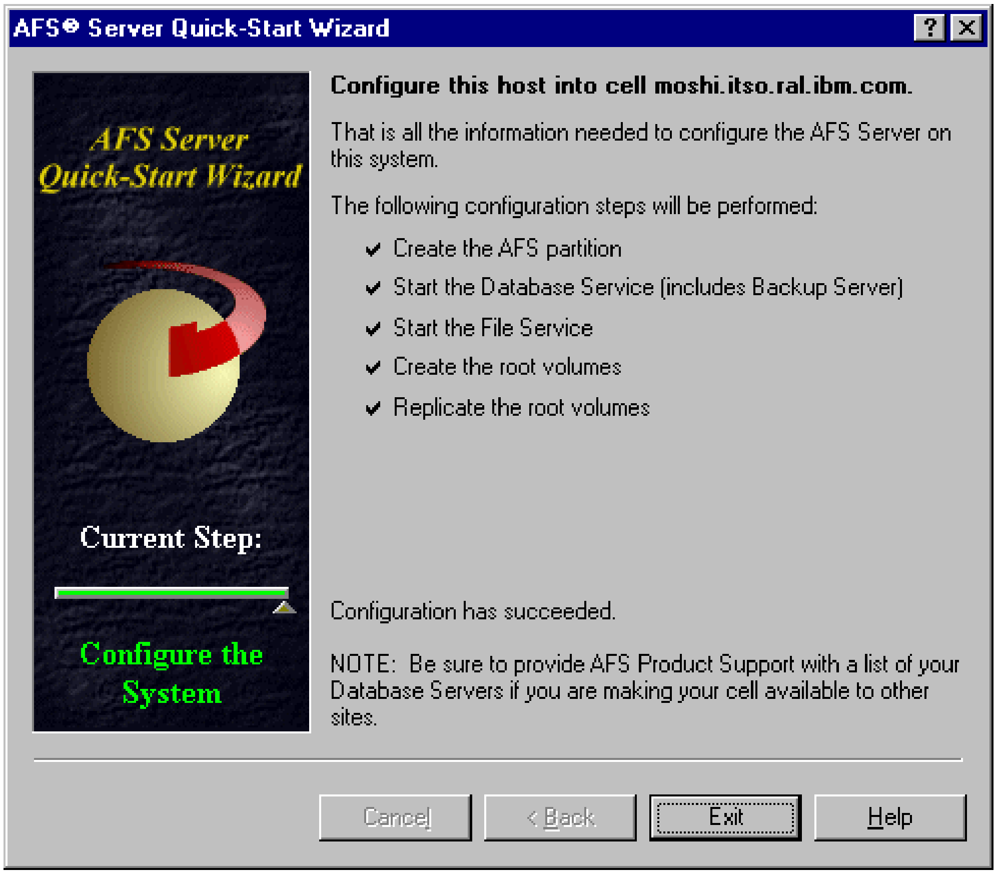

AFS® Server Quick-Start Wizard

Configure this host into cell moshi.itso.ral.ibm.com.

That is all the information needed to configure the AFS Server on this system.

The following configuration steps will be performed:

- ✓ Create the AFS partition
- ✓ Start the Database Service (includes Backup Server)
- ✓ Start the File Service
- ✓ Create the root volumes
- ✓ Replicate the root volumes

Configuration has succeeded.

NOTE: Be sure to provide AFS Product Support with a list of your Database Servers if you are making your cell available to other sites.

Figure 113. The Server Configuration Status Window

After a successful configuration, the Wizard displays the message:

```
Configuration has succeeded.
```

Click **Exit** to complete the configuration of the first AFS machine.

You can find informative messages in the AFS server configuration log file. This log file, named afs_server_config_log.txt, can be found in the directory C:\WSPP\AFS\Server\usr\afs\logs.

The next several figures contain the log of the configuration session of the first AFS machine on a Windows NT system:

```
00:54:42 04/15/99:  Log file open.
00:54:42 04/15/99:  Checking this machine's current configuration...
00:54:42 04/15/99:  Is the AFS Client installed on this machine: Yes.
00:54:42 04/15/99:  The client configuration information on this host is not valid:  (0x4629), the
client is not configured in any cell.
Is there valid configuration information on this machine: No.
00:54:42 04/15/99:  The configuration information on this host is not valid:  (0x460f), the
minimally required server configuration information is missing, unreadable, or invalid.
Getting handles to the cell and the config library.
00:54:42 04/15/99:  Opening a NULL cell handle to use with the client config library handle.
00:54:42 04/15/99:  Getting config handle for the client.
00:54:42 04/15/99:  Opening a NULL cell handle to use with the server config library handle.
00:54:42 04/15/99:  Getting config library handle for the server.
00:54:42 04/15/99:  Does a partition exist on this machine: No.
Configuring server
00:57:53 04/15/99:  Adding an AFS partition on device 'D:' with name '/vicepa'.
00:57:53 04/15/99:  Putting this host in cell 'moshi.itso.ral.ibm.com'.
00:57:53 04/15/99:  Stopping the client.
00:57:54 04/15/99:  Putting the AFS Client in this host's cell.
00:57:54 04/15/99:  Starting the bos server in no auth mode.
00:57:55 04/15/99:  Starting the authentication server.
00:58:01 04/15/99:  Getting credentials in cell 'moshi.itso.ral.ibm.com' as admin 'admin'.
00:58:01 04/15/99:  Getting handles to the cell and the config library.
00:58:01 04/15/99:  Opening a non-NULL cell handle to use with the server config library handle.
00:58:01 04/15/99:  Getting unauthenticated tokens for cell 'moshi.itso.ral.ibm.com'.
00:58:01 04/15/99:  Getting cell handle for cell moshi.itso.ral.ibm.com.
00:58:01 04/15/99:  Getting config library handle for the server.
00:58:01 04/15/99:  Setting the AFS Principal.
00:58:02 04/15/99:  Starting the following servers: Protection   Volume Location   Backup
00:58:02 04/15/99:  Waiting for database servers to reach quorum.
00:58:14 04/15/99:  Setting Admin Principal to 'admin' and UID to 0.
00:58:15 04/15/99:  Getting credentials in cell 'moshi.itso.ral.ibm.com' as admin 'admin'.
00:58:15 04/15/99:  Getting handles to the cell and the config library.
00:58:15 04/15/99:  Opening a non-NULL cell handle to use with the server config library handle.
00:58:15 04/15/99:  Getting tokens in cell moshi.itso.ral.ibm.com for user 'admin'.
00:58:15 04/15/99:  Getting cell handle for cell moshi.itso.ral.ibm.com.
00:58:15 04/15/99:  Getting config library handle for the server.
00:58:15 04/15/99:  Starting the File Server.
00:58:20 04/15/99:  Opening server wtr05319.itso.ral.ibm.com.
00:58:20 04/15/99:  Translating the parition name '/vicepa' to an ID.
00:58:20 04/15/99:  The ID for partition '/vicepa' is 0.
00:58:20 04/15/99:  Stopping the client.
```

Figure 114. (Part 1 of 2). The AFS Server Configuration Log File, afs_server_config_log.txt

```
00:58:21 04/15/99:  Creating volume root.afs on partition 0 with a quota of 5000.
00:58:22 04/15/99:  Volume root.afs was created with an ID of 536870912.
00:58:22 04/15/99:  Starting the AFS Client.
00:58:33 04/15/99:  Putting our tokens into the AFS Client.
00:58:34 04/15/99:  Attempting to create a drive mapping into AFS.
00:58:34 04/15/99:  Attempting to map D: to \\wtr05319-afs\all: failed.
00:58:34 04/15/99:  Attempting to map E: to \\wtr05319-afs\all: failed.
00:58:34 04/15/99:  Attempting to map F: to \\wtr05319-afs\all: failed.
00:58:34 04/15/99:  Attempting to map G: to \\wtr05319-afs\all: succeeded.
00:58:34 04/15/99:  Setting the ACL on root.afs.
00:58:34 04/15/99:  Creating volume root.cell on partition 0 with a quota of 5000.
00:58:35 04/15/99:  Volume root.cell was created with an ID of 536870915.
00:58:35 04/15/99:  Mouting root.cell with a Standard mount point at path
G:\moshi.itso.ral.ibm.com.
00:58:35 04/15/99:  Setting the ACL on root.cell (dir G:\moshi.itso.ral.ibm.com).
00:58:35 04/15/99:  Mounting root.cell with a Read/Write mount point at path
G:\.moshi.itso.ral.ibm.com.
00:58:35 04/15/99:  Creating a read only site for volume root.afs using partition ID 0 and volume
ID 536870912.
00:58:35 04/15/99:  Releasing the root.afs volume using volume ID 536870912.
00:58:36 04/15/99:  Creating a read only site for volume root.cell using partition ID 0 and volume
ID 536870915.
00:58:36 04/15/99:  Releasing the root.cell volume using volume ID 536870915.
00:58:37 04/15/99:  Bos open of server 'wtr05319.itso.ral.ibm.com'.
00:58:37 04/15/99:  Enabling auth checking.
00:58:37 04/15/99:  Closing bos server connection.
00:58:37 04/15/99:  Closing the connection to this server.
00:58:37 04/15/99:  Configuration succeeded.
00:58:37 04/15/99:  Attempting to remove the drive mapping to root.afs: succeeded.
00:58:57 04/15/99:  Closing log file.
```

Figure 115. (Part 2 of 2). The AFS Server Configuration Log File, afs_server_config_log.txt

5.4.2 Verifying the First AFS Machine Configuration

To verify that the configuration of the first AFS machine on your Windows NT system completed successfully, perform the following steps:

1. Display the status of AFS services in the Control Panel Services window. To invoke this window, click the **Start** button on the Windows NT desktop taskbar, point to **Settings**, and click **Control Panel**. From the Windows NT Control Panel, double-click on the **Services** icon. Both the IBM AFS Server and Client service entries must show the Started status, as illustrated in Figure 116 on page 164.

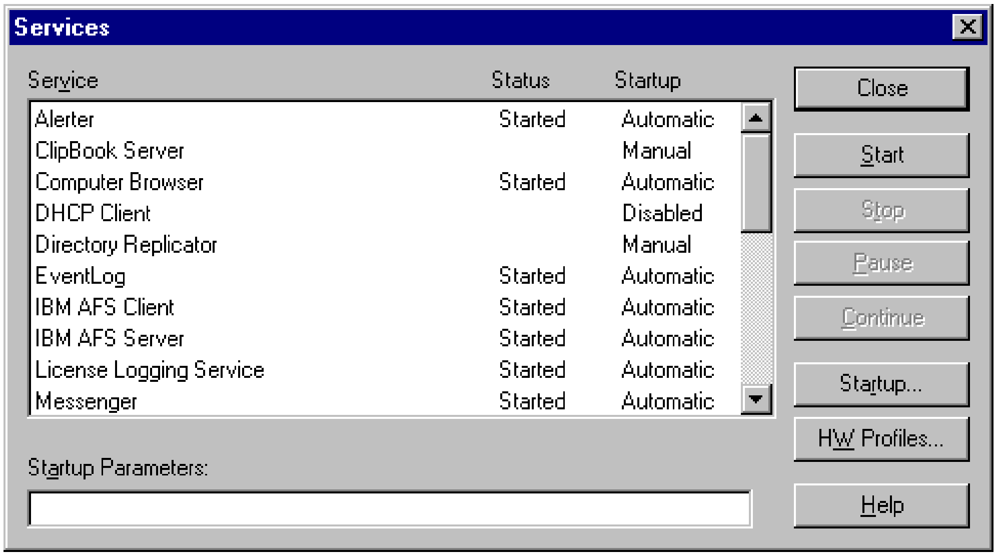

Service	Status	Startup
Alerter	Started	Automatic
ClipBook Server		Manual
Computer Browser	Started	Automatic
DHCP Client		Disabled
Directory Replicator		Manual
EventLog	Started	Automatic
IBM AFS Client	Started	Automatic
IBM AFS Server	Started	Automatic
License Logging Service	Started	Automatic
Messenger	Started	Automatic

Figure 116. The Control Panel Services Window

2. You can display the status of the AFS Server services by using the following methods:

   - The AFS Server Configuration window

     Click the **Start** button on the Windows NT desktop taskbar, point to **Settings**, and click **Control Panel**. From the Windows NT Control Panel, double-click on the **AFS Server Configuration** icon. In the AFS Server Configuration window, click on the **Services** tab to display the running AFS services, as shown in Figure 117 on page 165.

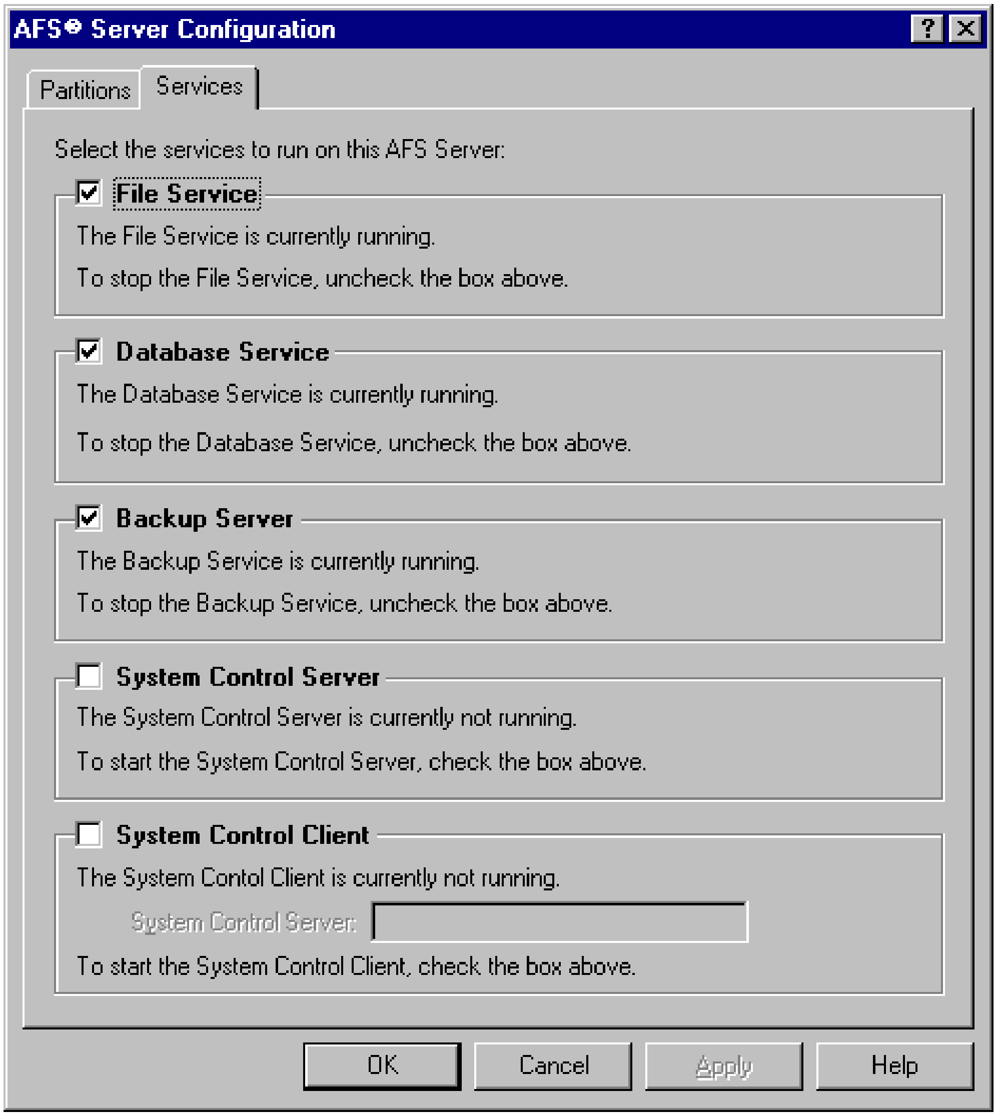

Select the services to run on this AFS Server:

☑ File Service

The File Service is currently running.

To stop the File Service, uncheck the box above.

☑ Database Service

The Database Service is currently running.

To stop the Database Service, uncheck the box above.

☑ Backup Server

The Backup Service is currently running.

To stop the Backup Service, uncheck the box above.

□ System Control Server

The System Control Server is currently not running.

To start the System Control Server, check the box above.

□ System Control Client

The System Contol Client is currently not running.

System Control Server:

To start the System Control Client, check the box above.

Figure 117. Status of AFS Server Services on the First AFS Machine

- The MS-DOS Command Prompt window

  Click the **Start** button on the Windows NT desktop taskbar, point to **Programs**, and click **Command Prompt**. In this window, enter the following command:

  ```
  bos status hostname -long
  ```

  where *`hostname`* is the host name of this first AFS machine.

  Figure 118 on page 166 shows a sample output of the AFS command:

  ```
  bos status
  ```

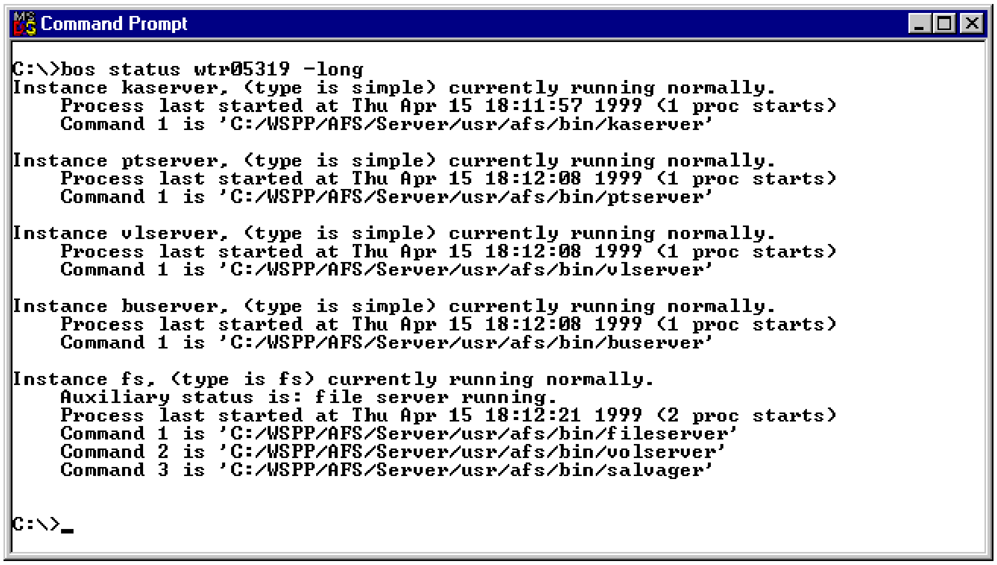

```
C:\>bos status wtr05319 -long
Instance kaserver, (type is simple) currently running normally.
    Process last started at Thu Apr 15 18:11:57 1999 (1 proc starts)
    Command 1 is 'C:/WSPP/AFS/Server/usr/afs/bin/kaserver'

Instance ptserver, (type is simple) currently running normally.
    Process last started at Thu Apr 15 18:12:08 1999 (1 proc starts)
    Command 1 is 'C:/WSPP/AFS/Server/usr/afs/bin/ptserver'

Instance vlserver, (type is simple) currently running normally.
    Process last started at Thu Apr 15 18:12:08 1999 (1 proc starts)
    Command 1 is 'C:/WSPP/AFS/Server/usr/afs/bin/vlserver'

Instance buserver, (type is simple) currently running normally.
    Process last started at Thu Apr 15 18:12:08 1999 (1 proc starts)
    Command 1 is 'C:/WSPP/AFS/Server/usr/afs/bin/buserver'

Instance fs, (type is fs) currently running normally.
    Auxiliary status is: file server running.
    Process last started at Thu Apr 15 18:12:21 1999 (2 proc starts)
    Command 1 is 'C:/WSPP/AFS/Server/usr/afs/bin/fileserver'
    Command 2 is 'C:/WSPP/AFS/Server/usr/afs/bin/volserver'
    Command 3 is 'C:/WSPP/AFS/Server/usr/afs/bin/salvager'


C:\>_
```

Figure 118. Displaying AFS Server Process Status on a Windows NT System

The AFS Server Configuration Services window and the AFS command:

```
bos status
```

display the status of the AFS server processes in the BOS configuration file, C:\WSPP\AFS\Server\usr\afs\local\BosConfig, whose contents are shown in Figure 119 on page 166.

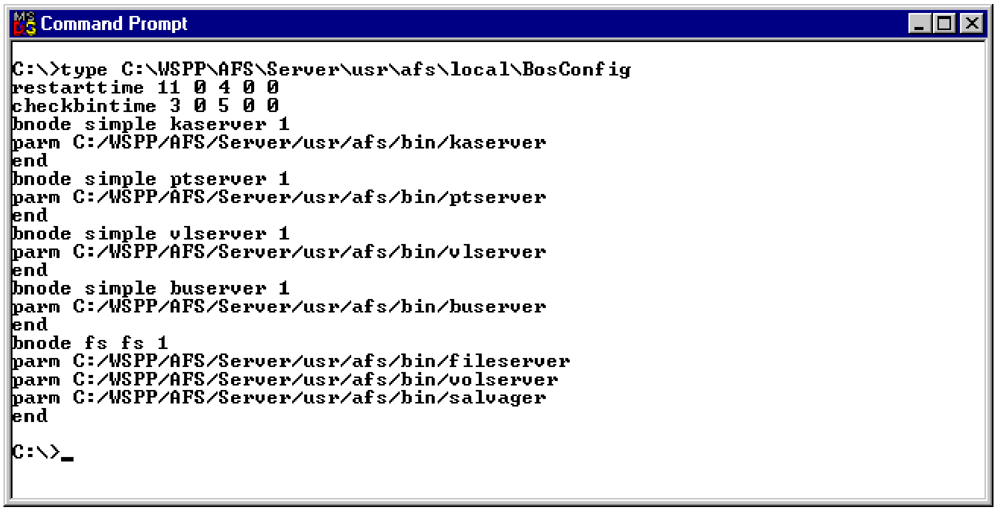

```
C:\>type C:\WSPP\AFS\Server\usr\afs\local\BosConfig
restarttime 11 0 4 0 0
checkbintime 3 0 5 0 0
bnode simple kaserver 1
parm C:/WSPP/AFS/Server/usr/afs/bin/kaserver
end
bnode simple ptserver 1
parm C:/WSPP/AFS/Server/usr/afs/bin/ptserver
end
bnode simple vlserver 1
parm C:/WSPP/AFS/Server/usr/afs/bin/vlserver
end
bnode simple buserver 1
parm C:/WSPP/AFS/Server/usr/afs/bin/buserver
end
bnode fs fs 1
parm C:/WSPP/AFS/Server/usr/afs/bin/fileserver
parm C:/WSPP/AFS/Server/usr/afs/bin/volserver
parm C:/WSPP/AFS/Server/usr/afs/bin/salvager
end

C:\>_
```

Figure 119. Listing of the BOS Configuration File on the First AFS Machine

3. In the AFS Server Configuration window, click on the **Partitions** tab to display the AFS partitions, as shown in Figure 120 on page 167.

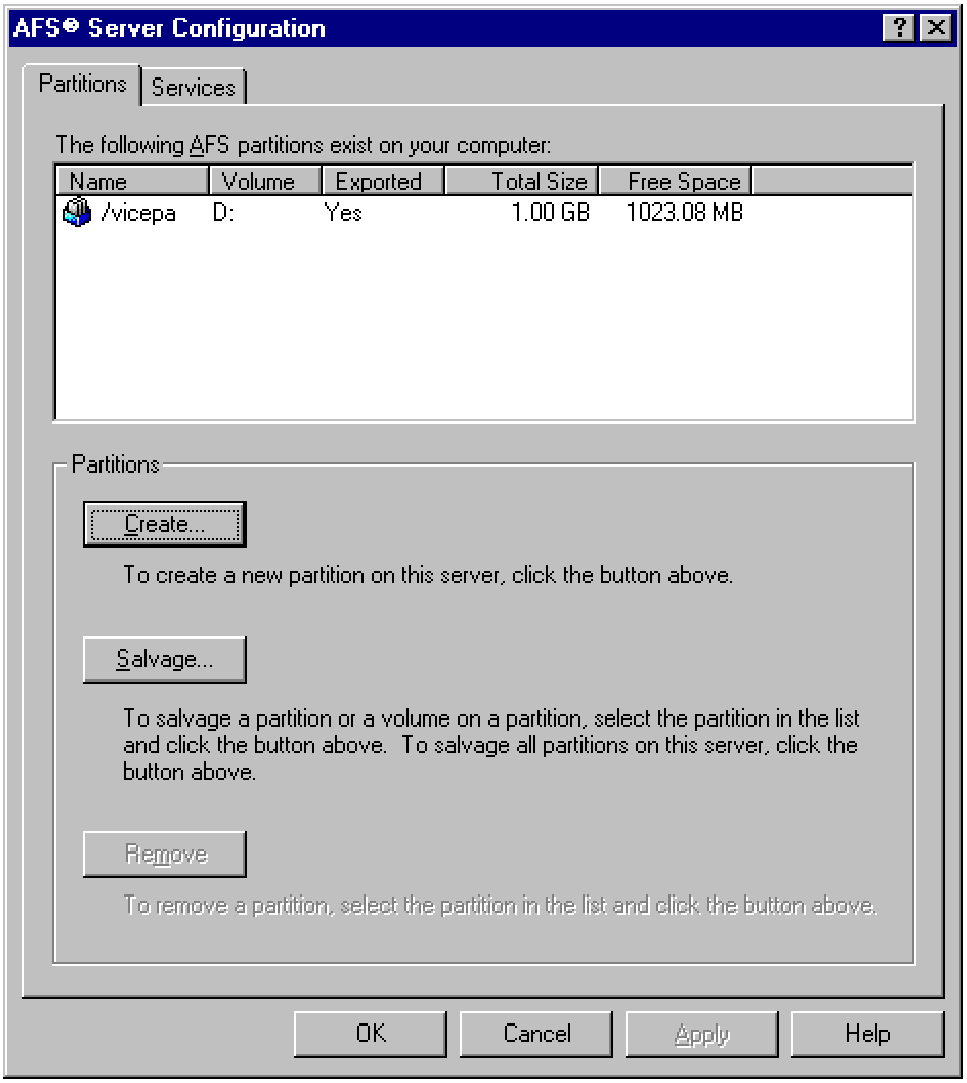

The following AFS partitions exist on your computer:

Name	Volume	Exported	Total Size	Free Space
/vicepa	D:	Yes	1.00 GB	1023.08 MB

Partitions

To create a new partition on this server, click the button above.

To salvage a partition or a volume on a partition, select the partition in the list and click the button above. To salvage all partitions on this server, click the button above.

To remove a partition, select the partition in the list and click the button above.

Figure 120. Listing of AFS Partitions

4. Display the status of the AFS Client from the AFS Client Configuration window. To invoke this window, click the **Start** button on the Windows NT desktop taskbar, point to **Settings**, and click **Control Panel**. From the Windows NT Control Panel, double-click the **AFS Client Configuration** icon. The AFS Client Configuration window displays the status of the AFS Client service and the AFS cell name of the machine, as shown in Figure 121 on page 168.

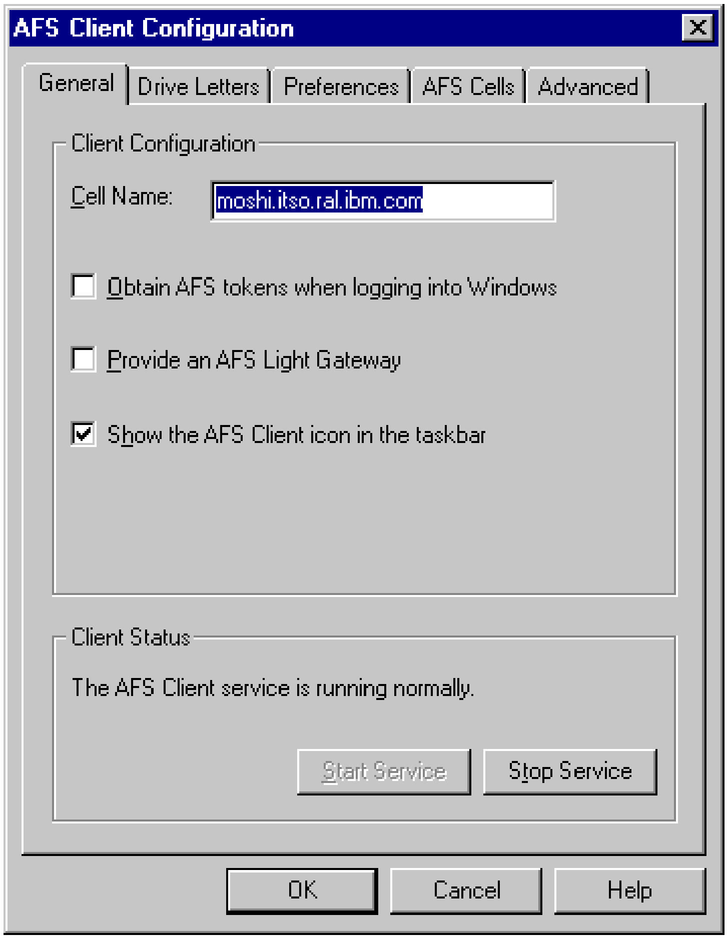

Client Configuration

Cell Name: moshi.itso.ral.ibm.com

☐ Obtain AFS tokens when logging into Windows

☐ Provide an AFS Light Gateway

☑ Show the AFS Client icon in the taskbar

Client Status

The AFS Client service is running normally.

Figure 121. The AFS Client Configuration Window

5. List the running AFS processes from the Windows NT Task Manager window. To invoke this window, position the mouse pointer on the Windows NT desktop taskbar, click the right mouse button, and click **Task Manager...**. The Windows NT Task Manager window appears. Click the **Processes** tab to show the current processes, as shown in the following figure:

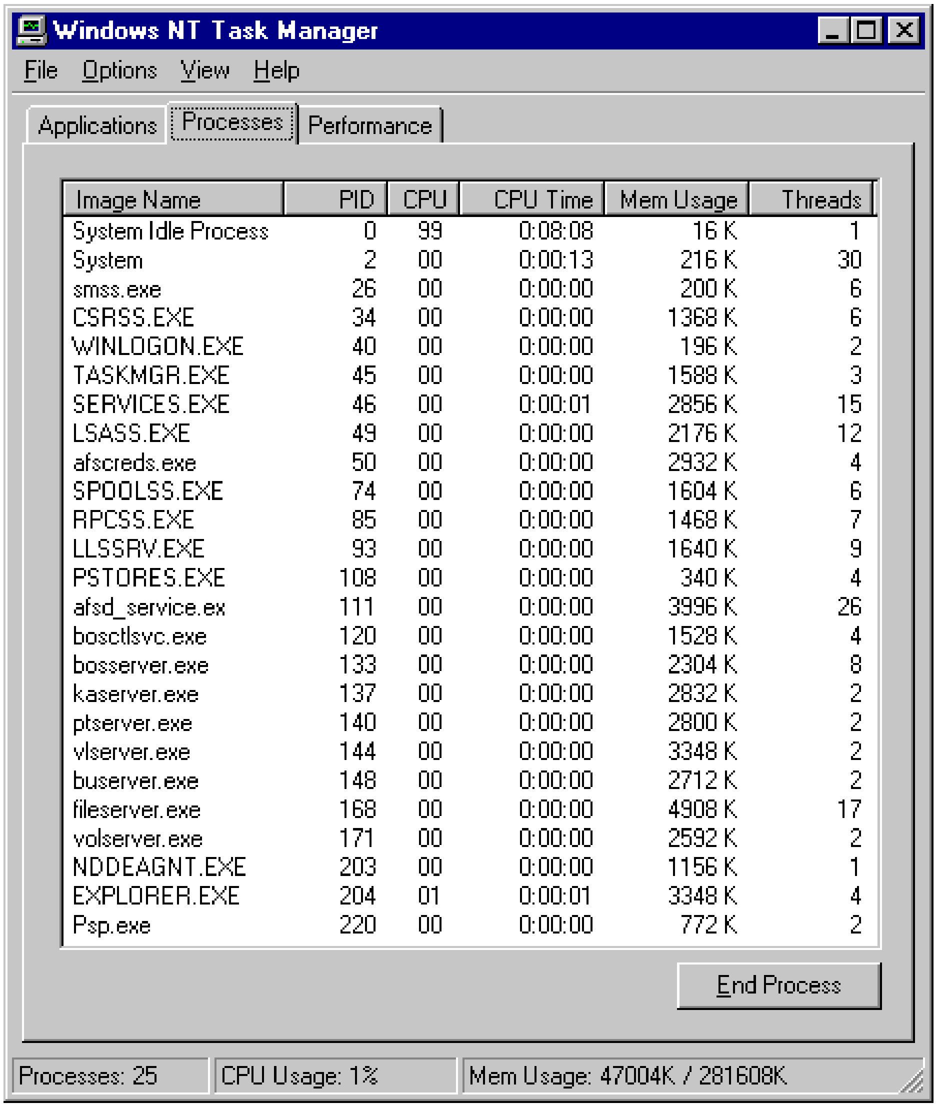

Image Name	PID	CPU	CPU Time	Mem Usage	Threads
System Idle Process	0	99	0:08:08	16 K	1
System	2	00	0:00:13	216 K	30
smss.exe	26	00	0:00:00	200 K	6
CSRSS.EXE	34	00	0:00:00	1368 K	6
WINLOGON.EXE	40	00	0:00:00	196 K	2
TASKMGR.EXE	45	00	0:00:00	1588 K	3
SERVICES.EXE	46	00	0:00:01	2856 K	15
LSASS.EXE	49	00	0:00:00	2176 K	12
afscreds.exe	50	00	0:00:00	2932 K	4
SPOOLSS.EXE	74	00	0:00:00	1604 K	6
RPCSS.EXE	85	00	0:00:00	1468 K	7
LLSSRV.EXE	93	00	0:00:00	1640 K	9
PSTORES.EXE	108	00	0:00:00	340 K	4
afsd_service.ex	111	00	0:00:00	3996 K	26
bosctlsvc.exe	120	00	0:00:00	1528 K	4
bosserver.exe	133	00	0:00:00	2304 K	8
kaserver.exe	137	00	0:00:00	2832 K	2
ptserver.exe	140	00	0:00:00	2800 K	2
vlserver.exe	144	00	0:00:00	3348 K	2
buserver.exe	148	00	0:00:00	2712 K	2
fileserver.exe	168	00	0:00:00	4908 K	17
volserver.exe	171	00	0:00:00	2592 K	2
NDDEAGNT.EXE	203	00	0:00:00	1156 K	1
EXPLORER.EXE	204	01	0:00:01	3348 K	4
Psp.exe	220	00	0:00:00	772 K	2

Processes: 25 | CPU Usage: 1% | Mem Usage: 47004K / 281608K

Figure 122. Listing of AFS Processes on the First AFS Machine

The AFS BOS server process is bosserver.exe.

The AFS database server processes are kaserver.exe, ptserver.exe, vlserver.exe, and buserver.exe. The AFS fs processes are: fileserver.exe and volserver.exe. The AFS client process is afsd_service.exe.

To display the Threads column, as shown in Figure 122 on page 169, click the **View** button of the Windows NT Task Manager menu bar, and click **Select Columns...**. Notice that this selection is allowed only when the **Processes** tab is marked. In the Select Columns window, check the **Thread Count** box and click **OK** to continue. The number of threads per process information is useful in AFS tuning.

5.5 Configuring an AFS File Server Machine on Windows NT Systems

Before you can configure a Windows NT machine to be an AFS File Server to an existing AFS cell, you must complete the following tasks:

- Create at least one dedicated NTFS partition to be configured as an AFS partition to store AFS volumes, as illustrated in 5.2, “Configuring AFS Partitions on Windows NT Systems” on page 133. This NTFS partition must not contain any files other than the Windows NT Recycle Bin.
- Install the necessary AFS files from the WebSphere Performance Pack CD-ROM distribution to the machine local disk, as shown in 5.3, “Installing AFS Files on Windows NT Systems” on page 141.

AFS Logon Message

In the early release of the WebSphere AFS component, at the first login after installing AFS files, the system displays a warning message informing you that the specified realm is unknown. Since you have not configured AFS on this machine yet, click **OK** to continue. In the General Availability (GA) version of the WebSphere AFS component, the system no longer displays this warning message.

The WebSphere Performance Pack Setup program creates a Program Group called IBM WebSphere, containing Performance Pack Afs, as shown in Figure 123 on page 171. It also creates the AFS Server Configuration and the AFS Client Configuration applications in the Windows NT Control Panel. In addition, it creates the IBM AFS Server and IBM AFS Client entries in the Services Control Panel.

The WebSphere AFS component for Windows NT includes the AFS Server Configuration Wizard to assist you in the configuration of the AFS Server on your Windows NT system. The AFS Server Configuration Wizard takes you through the steps of configuring the AFS Server via a sequence of dialog boxes.

5.5.1 Using the AFS Server Configuration Wizard

To start the configuration of an AFS File Server on a Windows NT system, perform the following tasks:

1. Log in as an administrative user, who is a member of the local Administrators group on your Windows NT system.
2. To start the AFS Server Configuration Wizard, click the **Start** button on your Windows NT desktop taskbar, point to **Programs**, point to **IBM WebSphere**, point to **Performance Pack**, point to **Afs**, point to **Server**, and click **Configuration Wizard**, as depicted in Figure 123 on page 171.

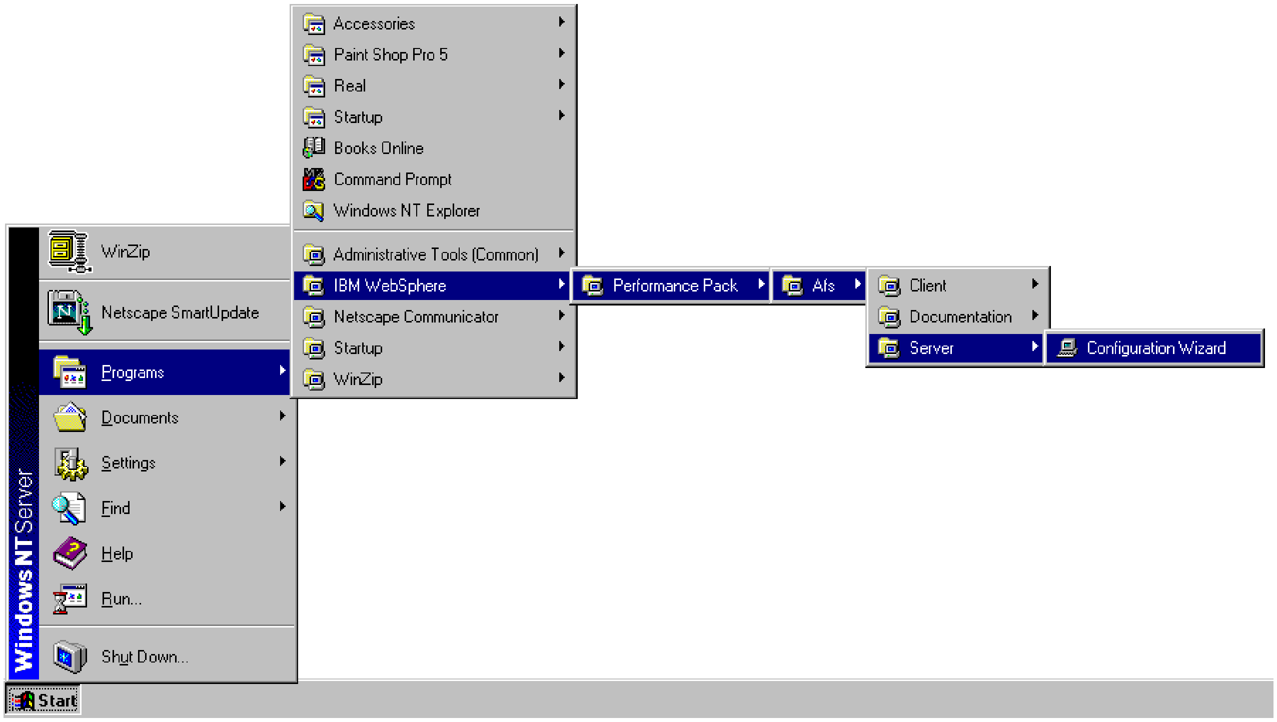

Figure 123. Starting the AFS Server Configuration Wizard

3. The Welcome to the AFS Server Quick Start wizard window appears:

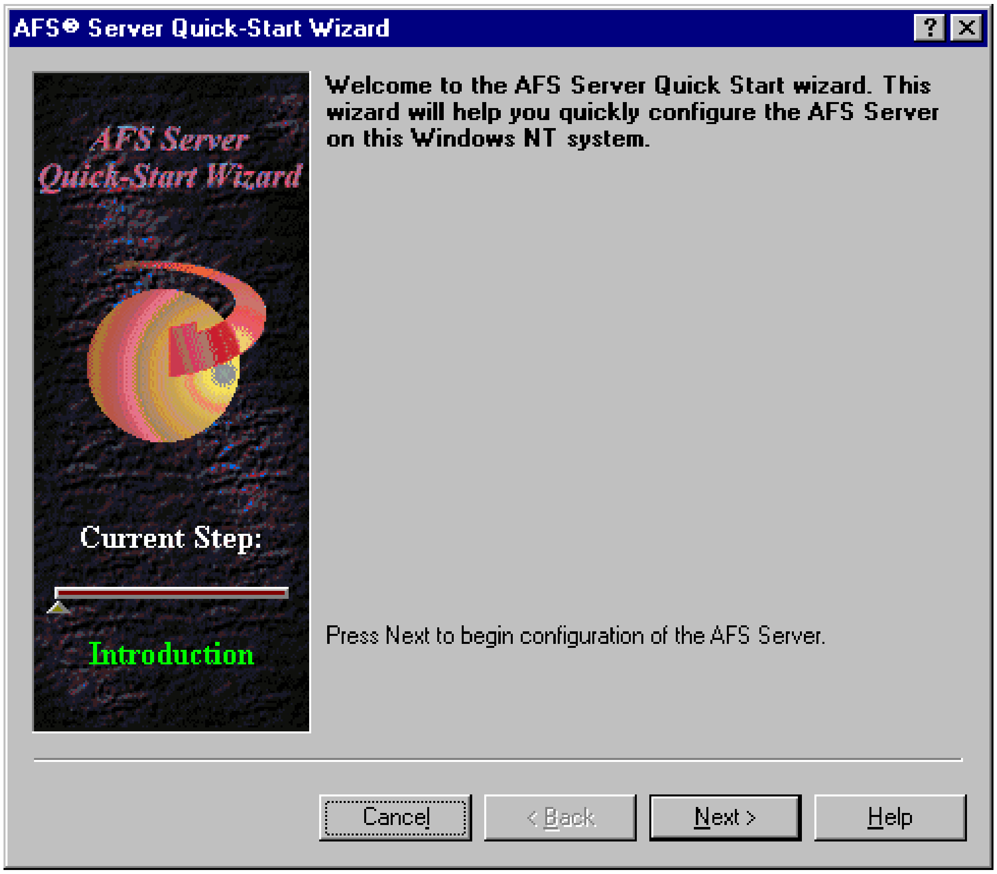

Welcome to the AFS Server Quick Start wizard. This wizard will help you quickly configure the AFS Server on this Windows NT system.

Press Next to begin configuration of the AFS Server.

Figure 124. The AFS Server Configuration Wizard Welcome Window

Click **Next** to continue.

4. In the Cell and Server Information step window, the Wizard provides you with the option of configuring this machine as an AFS server in an existing AFS cell or as the first AFS server in a new AFS cell. Select the **Make this host a server in an existing AFS cell** box. In the Cell Name frame, enter the name of an existing AFS cell that this AFS Server will participate in. In this example, the name of the AFS cell is moshi.itso.ral.ibm.com, as shown in the following figure:

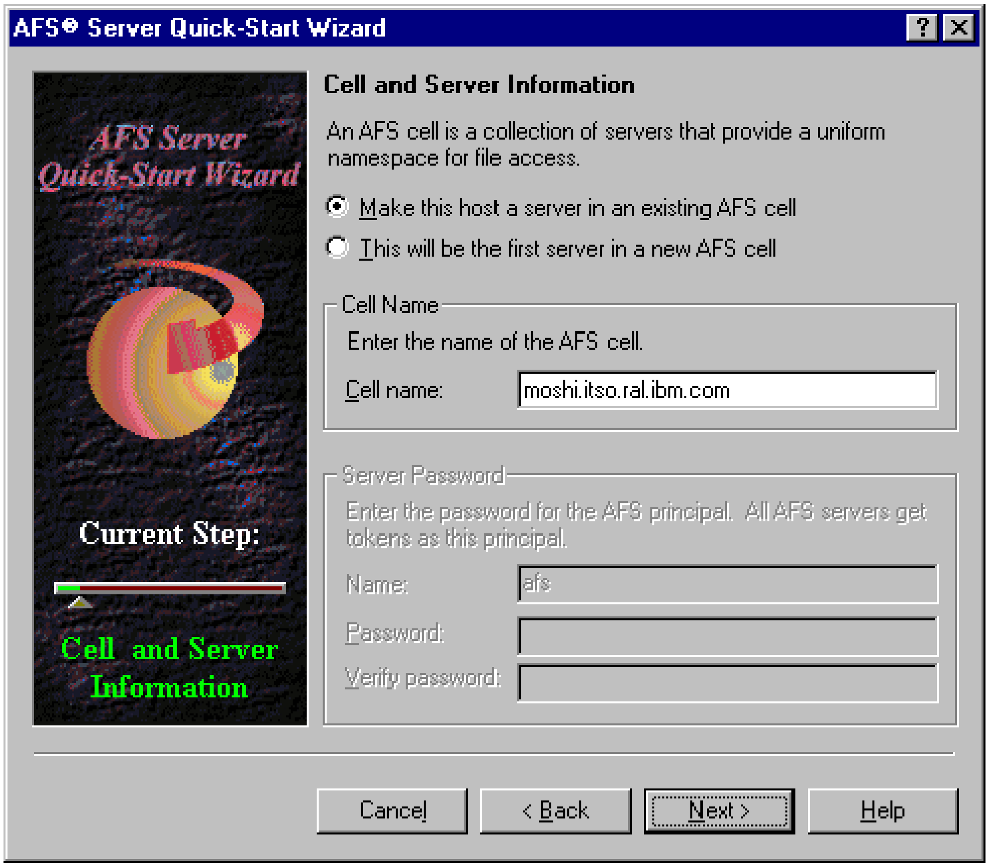

AFS® Server Quick-Start Wizard

Cell and Server Information

An AFS cell is a collection of servers that provide a uniform namespace for file access.

● Make this host a server in an existing AFS cell

○ This will be the first server in a new AFS cell

Cell Name

Enter the name of the AFS cell.

Cell name: moshi.itso.ral.ibm.com

Server Password

Enter the password for the AFS principal. All AFS servers get tokens as this principal.

Name: afs

Password:

Verify password:

Figure 125. The Cell and Server Information Window

Click **Next** to continue.

5. In the Administrative Information step window, enter the password for the AFS admin account in the Password field. In the Existing AFS Server frame, enter the host name of the first AFS machine in your cell. The new AFS Server machine retrieves the server common configuration information from the first AFS machine. The Administrative Information step window is shown in the following figure:

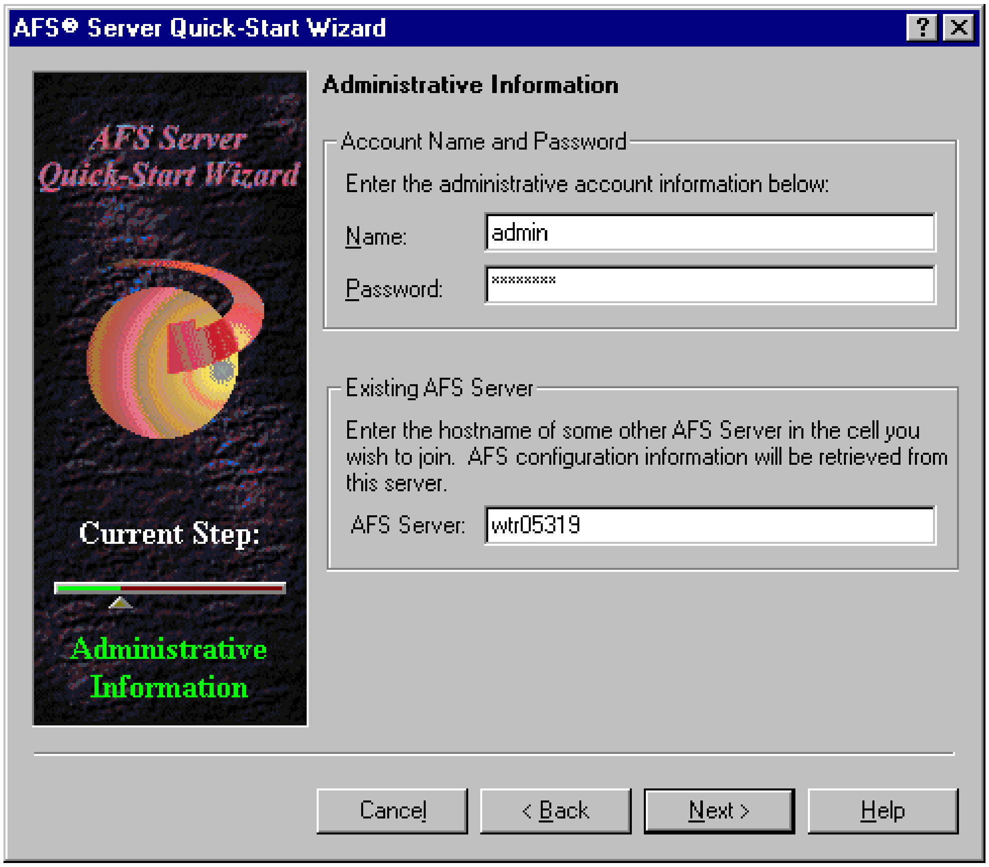
AFS® Server Quick-Start Wizard
AFS Server Quick-Start Wizard
Current Step:
Administrative Information

Administrative Information

Account Name and Password

Enter the administrative account information below:

Name: admin

Password: ********

Existing AFS Server

Enter the hostname of some other AFS Server in the cell you wish to join. AFS configuration information will be retrieved from this server.

AFS Server: wtr05319

Cancel
< Back
Next >
Help

Figure 126. The Administrative Information Window

In this example, the host name of the first AFS machine is wtr05319. Click **Next** to continue.

6. In the File Server step window, the Wizard provides you with the option of configuring this machine as an AFS File Server. The default option is to configure this machine as an AFS File Server, as shown in the following figure:

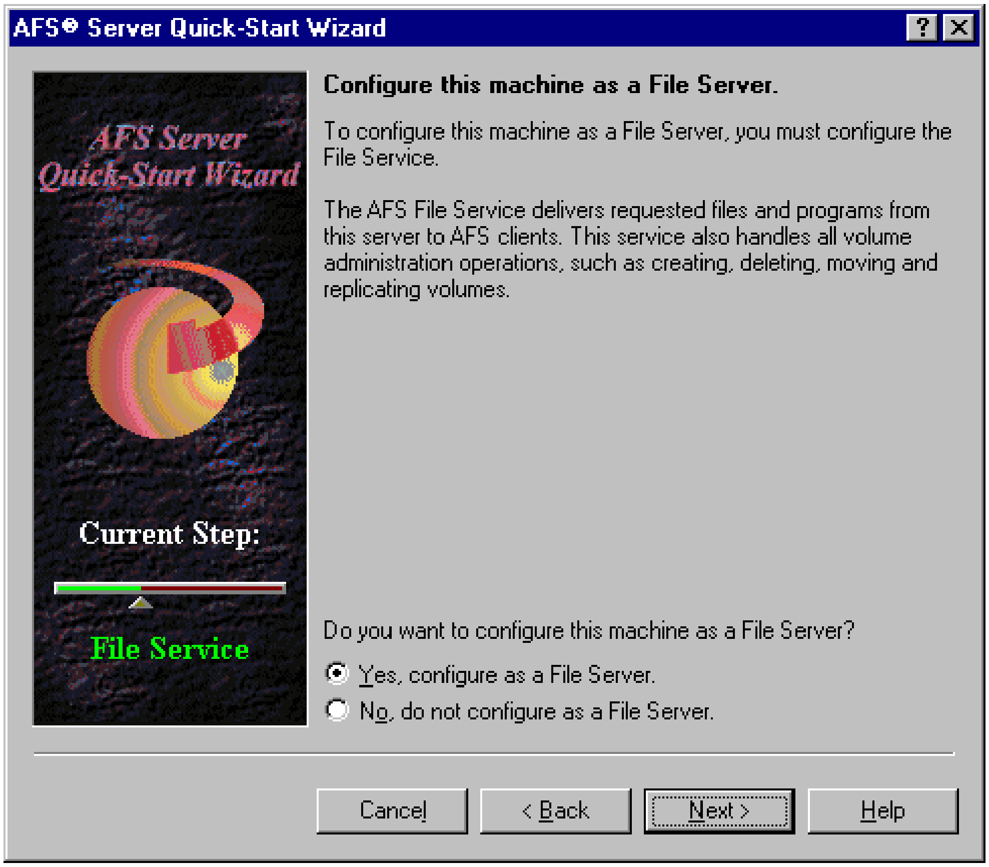

AFS® Server Quick-Start Wizard

Configure this machine as a File Server.

To configure this machine as a File Server, you must configure the File Service.

The AFS File Service delivers requested files and programs from this server to AFS clients. This service also handles all volume administration operations, such as creating, deleting, moving and replicating volumes.

Do you want to configure this machine as a File Server?

● Yes, configure as a File Server.

○ No, do not configure as a File Server.

Figure 127. The File Server Configuration Window

Click **Next** to continue.

7. In the Database Service step window, the Wizard provides you with the option of configuring this machine as an AFS Database Server. The default option is to configure this machine as an AFS Database Server. Since we only want to configure this machine as an AFS File Server, select the **No, do not configure as a Database Server** check box, as shown in the following figure:

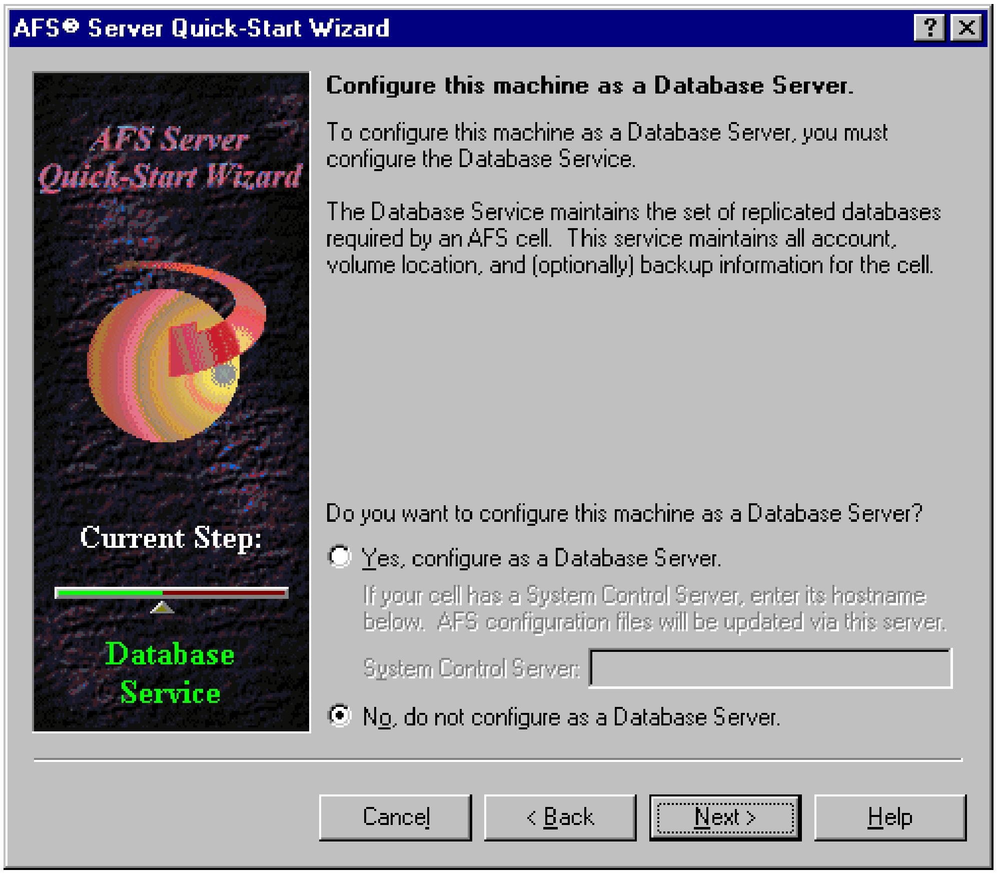

AFS® Server Quick-Start Wizard

Configure this machine as a Database Server.

To configure this machine as a Database Server, you must configure the Database Service.

The Database Service maintains the set of replicated databases required by an AFS cell. This service maintains all account, volume location, and (optionally) backup information for the cell.

Do you want to configure this machine as a Database Server?

○ Yes, configure as a Database Server.

If your cell has a System Control Server, enter its hostname below. AFS configuration files will be updated via this server.

System Control Server:

● No, do not configure as a Database Server.

Figure 128. The AFS Database Server Configuration Window

Click **Next** to continue.

8. In the Backup Server step window, the Wizard informs you that this machine will not be configured as an AFS Backup Server since we chose not to configure this machine as an AFS Database Server. This is shown in the following figure:

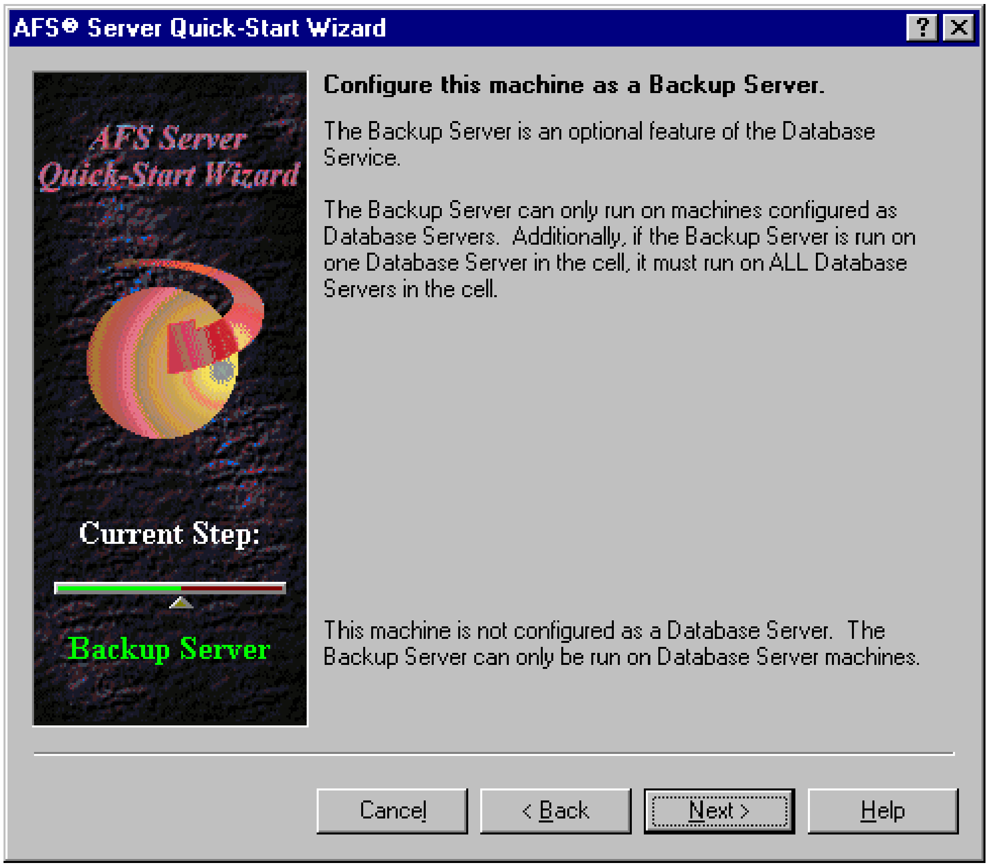

AFS® Server Quick-Start Wizard

Configure this machine as a Backup Server.

The Backup Server is an optional feature of the Database Service.

The Backup Server can only run on machines configured as Database Servers. Additionally, if the Backup Server is run on one Database Server in the cell, it must run on ALL Database Servers in the cell.

This machine is not configured as a Database Server. The Backup Server can only be run on Database Server machines.

Figure 129. The Backup Server Configuration Window

Click **Next** to continue.

9. In the AFS Partition step window, the Wizard displays all available Windows NT logical drives, as shown in the following figure:

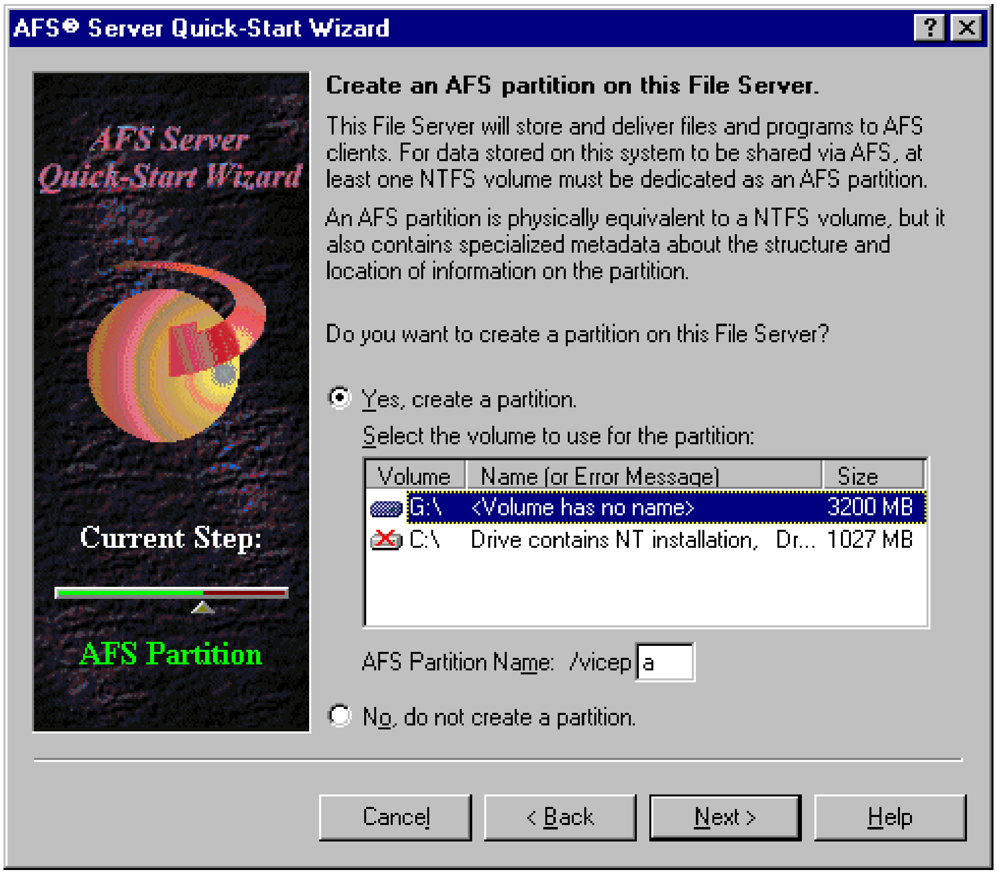

Create an AFS partition on this File Server.

This File Server will store and deliver files and programs to AFS clients. For data stored on this system to be shared via AFS, at least one NTFS volume must be dedicated as an AFS partition.

An AFS partition is physically equivalent to a NTFS volume, but it also contains specialized metadata about the structure and location of information on the partition.

Do you want to create a partition on this File Server?

● Yes, create a partition.

Select the volume to use for the partition:

Volume	Name (or Error Message)	Size
G:\	<Volume has no name>	3200 MB
C:\	Drive contains NT installation, Dr...	1027 MB

AFS Partition Name: /vicep a

○ No, do not create a partition.

Figure 130. The Create AFS Partition Window

Select one of the available NTFS disk partitions to be configured as an AFS partition.

In this example, drive G: is an NTFS partition without any data; for this reason, it is eligible to be configured as an AFS partition. On the other hand, drive C: is not available to be configured as an AFS partition, because this drive contains the Windows NT operating system. We select drive **G:\**.

By convention, the first AFS partition is called /vicepa, the second partition /vicepb, etc. Enter `a` in the AFS Partition Name box to assign /vicepa as the AFS partition name to the selected NTFS partition. Then, click **Next** to continue.

10.In the Root AFS Volumes step window, the Wizard provides you with the option of creating the root volumes in the event that they do not exist. This is shown in the following window:

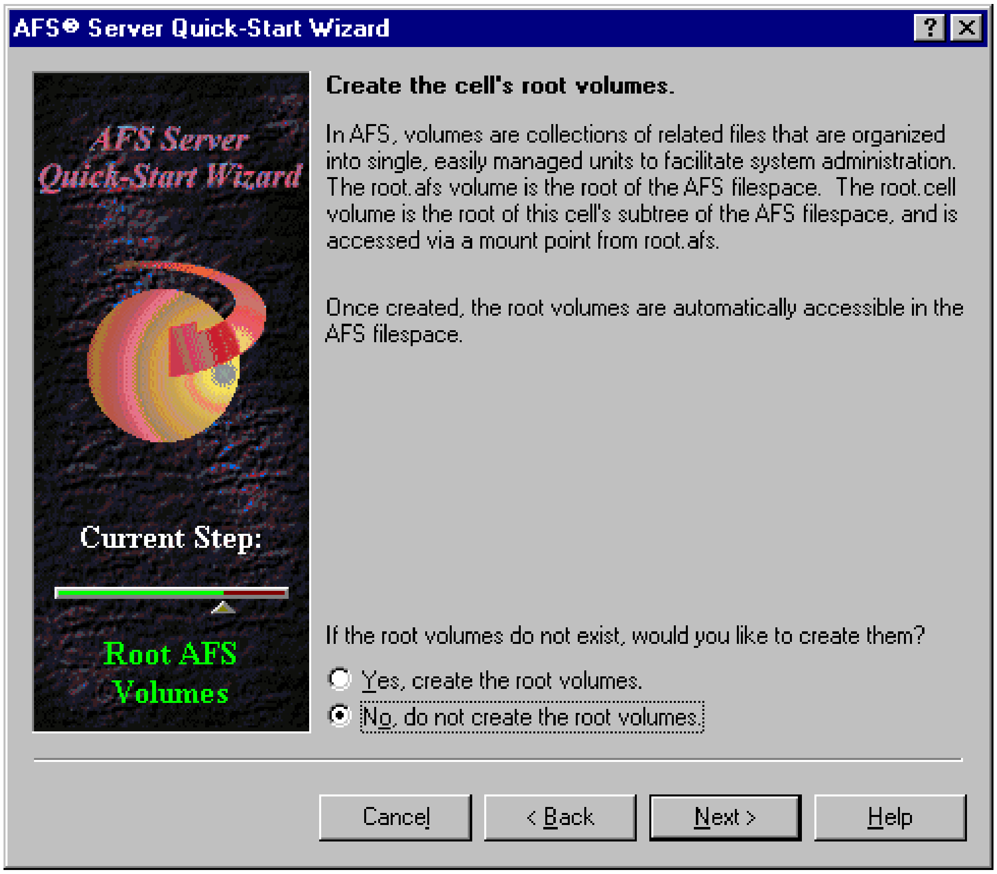

Create the cell's root volumes.

In AFS, volumes are collections of related files that are organized into single, easily managed units to facilitate system administration. The root.afs volume is the root of the AFS filespace. The root.cell volume is the root of this cell's subtree of the AFS filespace, and is accessed via a mount point from root.afs.

Once created, the root volumes are automatically accessible in the AFS filespace.

If the root volumes do not exist, would you like to create them?

○ Yes, create the root volumes.

● No, do not create the root volumes.

Figure 131. The Create Root Volumes Window

Since this machine is to be configured as a new AFS File Server in an existing and operational AFS cell, the root volumes have already been created. Therefore, select the **No, do not create the root volumes** check box. Then, click **Next** to continue.

11.In the Replication step window, the Wizard provides you with the option of replicating the root volumes in the event that they are not replicated, as shown in the following figure:

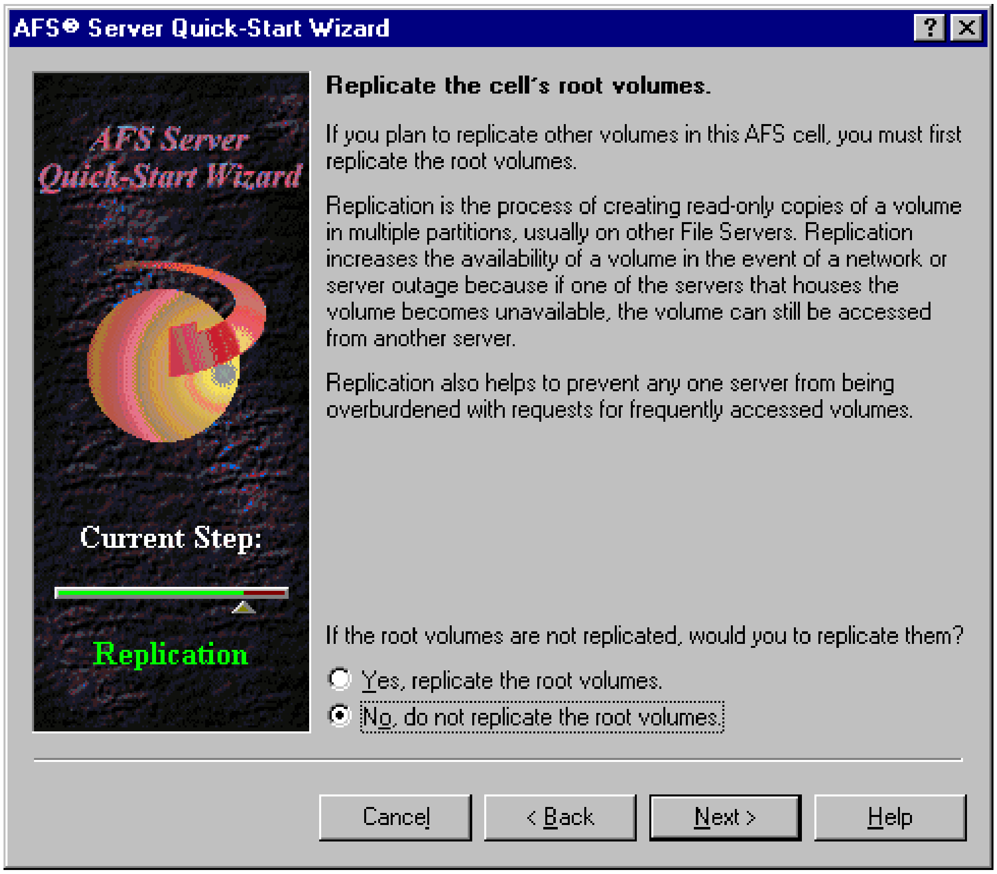

Replicate the cell's root volumes.

If you plan to replicate other volumes in this AFS cell, you must first replicate the root volumes.

Replication is the process of creating read-only copies of a volume in multiple partitions, usually on other File Servers. Replication increases the availability of a volume in the event of a network or server outage because if one of the servers that houses the volume becomes unavailable, the volume can still be accessed from another server.

Replication also helps to prevent any one server from being overburdened with requests for frequently accessed volumes.

If the root volumes are not replicated, would you to replicate them?

○ Yes, replicate the root volumes.

● No, do not replicate the root volumes.

Figure 132. The Replicate Root Volumes Window

Since this machine is a new AFS File Server in an existing and operational AFS cell, the root cell volumes have already been replicated. Therefore, select the **No, do not replicate the root volumes** check box. Then, click **Next** to continue.

12.In the System Control Service step window, the Wizard provides you with the option of configuring this machine as a System Control Server or Client, as shown in the following figure:

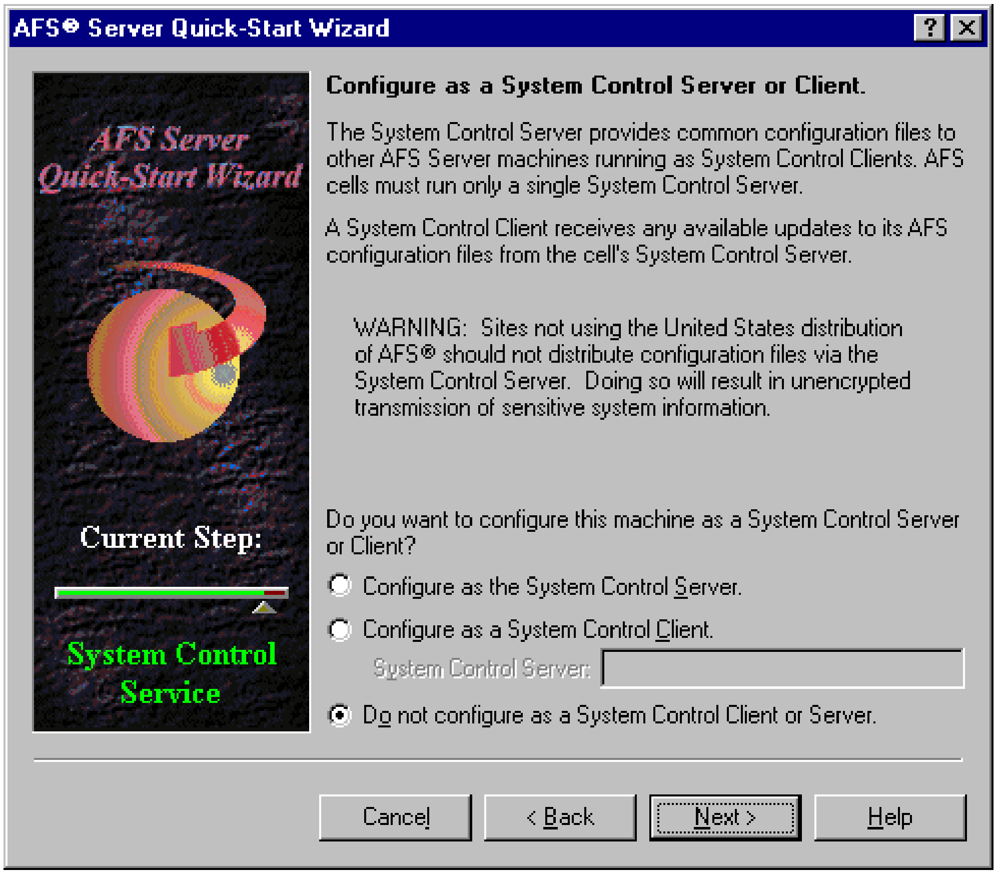
AFS® Server Quick-Start Wizard

Configure as a System Control Server or Client.

The System Control Server provides common configuration files to other AFS Server machines running as System Control Clients. AFS cells must run only a single System Control Server.

A System Control Client receives any available updates to its AFS configuration files from the cell's System Control Server.

WARNING: Sites not using the United States distribution of AFS® should not distribute configuration files via the System Control Server. Doing so will result in unencrypted transmission of sensitive system information.

Do you want to configure this machine as a System Control Server or Client?

○ Configure as the System Control Server.

○ Configure as a System Control Client.

System Control Server:

● Do not configure as a System Control Client or Server.

Figure 133. System Control Configuration Window

In general, the System Control Server is the first AFS machine. Since IBM WebSphere Performance Pack only provides the international edition of AFS, the System Control Server can only distribute files in an unencrypted form. Therefore, we do not recommend that you configure any machine as a System Control Server to distribute AFS server common configuration files.

On the other hand, a System Control Client must communicate with a System Control Server, which in our cell is unavailable if we do not configure any machine as a System Control Server. For this reason, also the option to configure this machine as a System Control Client should not be selected.

Therefore, we select the **Do not configure as a System Control Client or Server** check box. Then, we click **Next** to continue.

13.The Configure the System step window appears, displaying the various steps that will be performed. This is shown in the following figure:

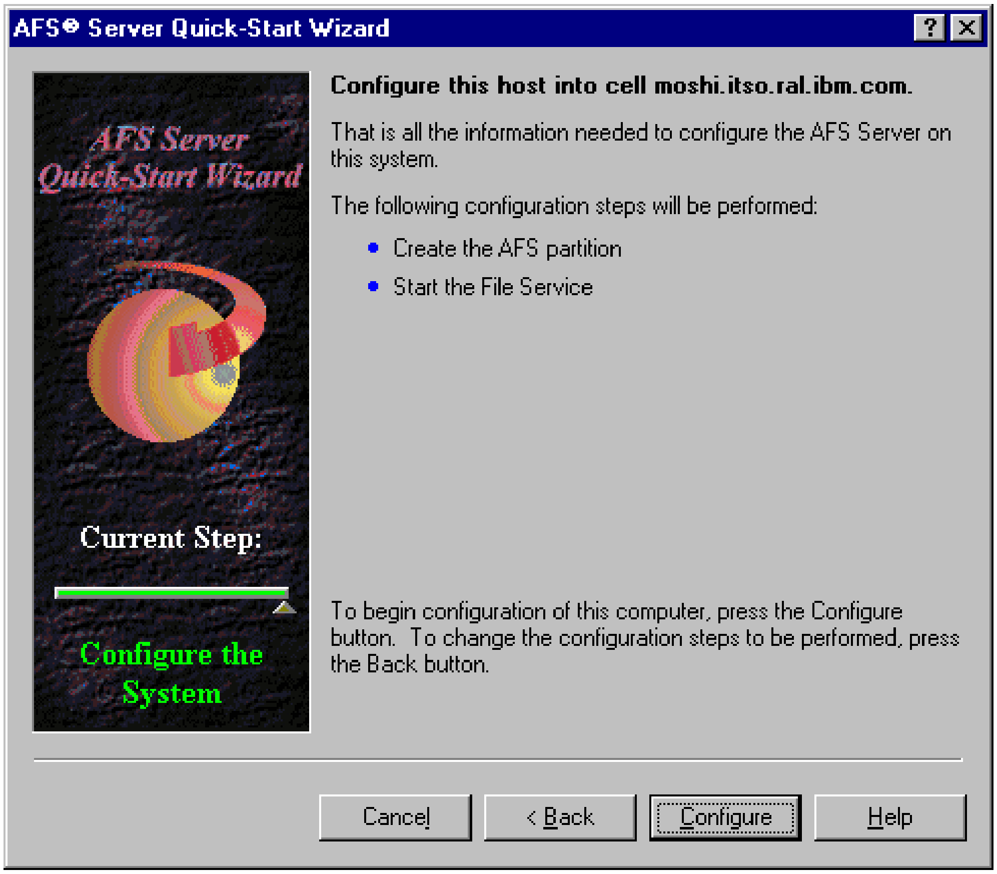

AFS® Server Quick-Start Wizard

Configure this host into cell moshi.itso.ral.ibm.com.

That is all the information needed to configure the AFS Server on this system.

The following configuration steps will be performed:

- Create the AFS partition
- Start the File Service

To begin configuration of this computer, press the Configure button. To change the configuration steps to be performed, press the Back button.

Figure 134. AFS Server Configuration Window

Click **Configure** to continue.

14.The Wizard informs you that the AFS Client will be installed and configured automatically. In WebSphere Performance Pack, an AFS Server machine is always an AFS Client too. Click **OK** to continue, as shown in the following message screen:

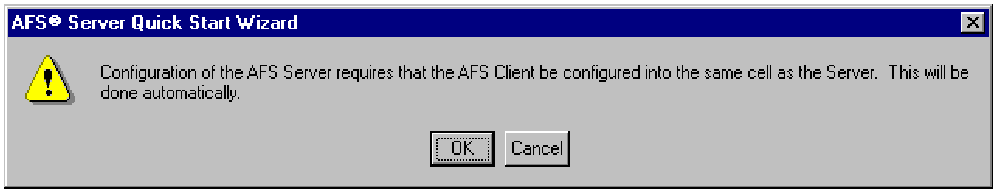

AFS® Server Quick Start Wizard

Configuration of the AFS Server requires that the AFS Client be configured into the same cell as the Server. This will be done automatically.

Figure 135. AFS Server Configuration Wizard Warning Message Window

15.At the end of the process, the Wizard informs you of the success of the configuration. Click **Exit** to complete the configuration of the AFS File Server, as shown in the following figure:

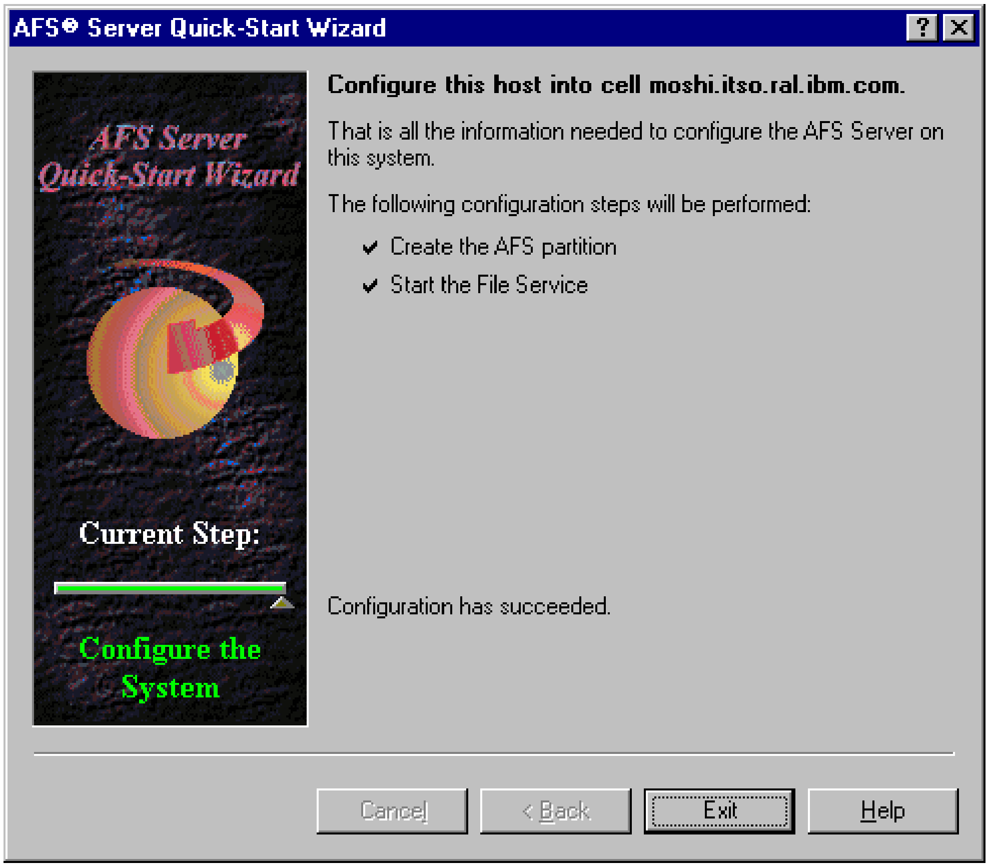

Configure this host into cell moshi.itso.ral.ibm.com.

That is all the information needed to configure the AFS Server on this system.

The following configuration steps will be performed:

- ✓ Create the AFS partition
- ✓ Start the File Service

Configuration has succeeded.

Figure 136. The AFS Server Configuration Window

You can find informative messages in the AFS server configuration log file, afs_server_config_log.txt. By default, this log file can be found in the directory C:\WSPP\AFS\Server\usr\afs\logs.

The next several figures contain the log of the configuration session of an AFS File Server on a Windows NT system:

```
22:07:15 03/25/99:  Log file open.
22:07:15 03/25/99:  Checking this machine's current configuration...
22:07:16 03/25/99:  Is the AFS Client installed on this machine: Yes.
22:07:16 03/25/99:  The client configuration information on this host is not valid:  (0x4629), the
client is not configured in any cell.
Is there valid configuration information on this machine: No.
22:07:16 03/25/99:  The configuration information on this host is not valid:  (0x460f), the
minimally required server configuration information is missing, unreadable, or invalid.
Getting handles to the cell and the config library.
22:07:16 03/25/99:  Opening a NULL cell handle to use with the client config library handle.
22:07:16 03/25/99:  Getting config handle for the client.
22:07:16 03/25/99:  Opening a NULL cell handle to use with the server config library handle.
22:07:16 03/25/99:  Getting config library handle for the server.
22:07:16 03/25/99:  Does a partition exist on this machine: No.
Configuring server
22:35:27 03/25/99:  Stopping the client.
22:35:28 03/25/99:  Getting CellServDB from host wtr05319.
22:35:28 03/25/99:  wtr05319.itso.ral.ibm.com
22:35:28 03/25/99:  Putting the AFS Client in this host's cell.
22:35:28 03/25/99:  Getting credentials in cell 'moshi.itso.ral.ibm.com' as admin 'admin'.
22:35:28 03/25/99:  Getting handles to the cell and the config library.
22:35:28 03/25/99:  Opening a non-NULL cell handle to use with the server config library handle.
22:35:28 03/25/99:  Getting tokens in cell moshi.itso.ral.ibm.com for user 'admin'.
22:35:28 03/25/99:  Getting cell handle for cell moshi.itso.ral.ibm.com.
22:35:28 03/25/99:  Getting config library handle for the server.
22:35:28 03/25/99:  Starting the AFS Client.
22:35:39 03/25/99:  Putting this host in cell 'moshi.itso.ral.ibm.com'.
22:35:39 03/25/99:  Starting the bos server in no auth mode.
22:35:40 03/25/99:  Getting credentials in cell 'moshi.itso.ral.ibm.com' as admin 'admin'.
22:35:40 03/25/99:  Getting handles to the cell and the config library.
22:35:40 03/25/99:  Opening a non-NULL cell handle to use with the server config library handle.
22:35:40 03/25/99:  Getting tokens in cell moshi.itso.ral.ibm.com for user 'admin'.
22:35:40 03/25/99:  Getting cell handle for cell moshi.itso.ral.ibm.com.
22:35:40 03/25/99:  Getting config library handle for the server.
22:35:40 03/25/99:  Setting the AFS Principal.
22:35:40 03/25/99:  Setting Admin Principal to 'admin'.
22:35:40 03/25/99:  Bos open of server 'wtr05320.itso.ral.ibm.com'.
22:35:40 03/25/99:  Enabling auth checking.
22:35:40 03/25/99:  Closing bos server connection.
22:35:40 03/25/99:  Getting credentials in cell 'moshi.itso.ral.ibm.com' as admin 'admin'.
22:35:40 03/25/99:  Getting handles to the cell and the config library.
22:35:40 03/25/99:  Opening a non-NULL cell handle to use with the server config library handle.
22:35:40 03/25/99:  Getting tokens in cell moshi.itso.ral.ibm.com for user 'admin'.
22:35:40 03/25/99:  Getting cell handle for cell moshi.itso.ral.ibm.com.
22:35:40 03/25/99:  Getting config library handle for the server.
22:35:40 03/25/99:  Opening server wtr05320.itso.ral.ibm.com.
22:35:40 03/25/99:  Adding an AFS partition on device 'G:' with name '/vicepa'.
22:35:40 03/25/99:  Starting the File Server.
22:35:45 03/25/99:  Do the root volumes exist: Yes.
Are the root volumes replicated: Yes.
Closing the connection to this server.
22:35:45 03/25/99:  Configuration succeeded.
22:48:15 03/25/99:  Closing log file.
```

Figure 137. The AFS Server Configuration Log File, afs_server_config_log.txt

5.5.2 Verifying the AFS File Server Configuration

After a successful completion of the configuration of the AFS File Server, both AFS Server and Client should be running. In fact, as we have already mentioned, an AFS Server in WebSphere Performance Pack is automatically installed and configured as an AFS Client too.

To verify that the configuration of the AFS File Server on your Windows NT system completed successfully, perform the following steps:

1. Display the status of AFS services in the Control Panel Services window. To invoke this window, click the **Start** button on the Windows NT desktop taskbar, point to **Settings**, and click **Control Panel**. From the Windows NT Control Panel, double-click the **Services** icon. Both the IBM AFS Server and IBM AFS Client service entries should show the Started status, as illustrated in Figure 138 on page 185.

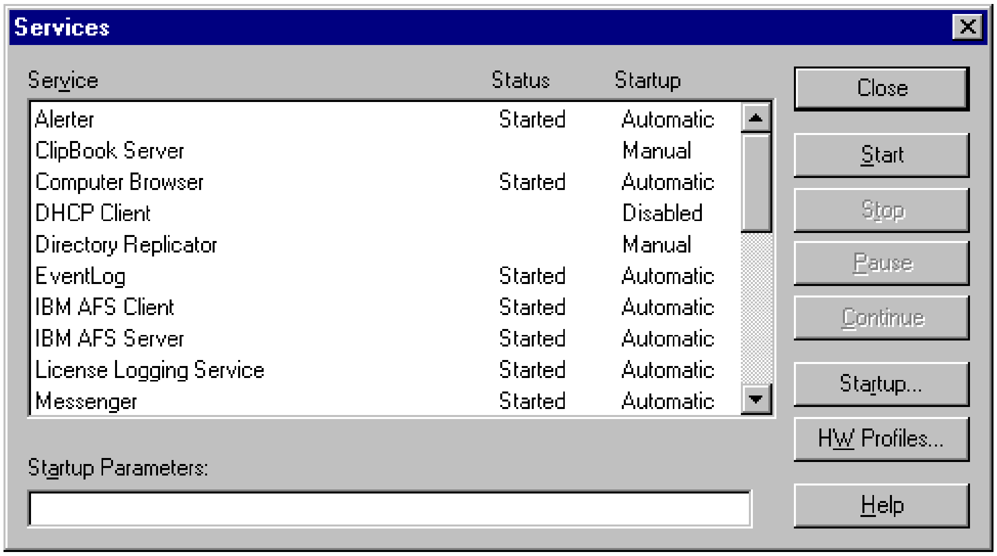

Service	Status	Startup
Alerter	Started	Automatic
ClipBook Server		Manual
Computer Browser	Started	Automatic
DHCP Client		Disabled
Directory Replicator		Manual
EventLog	Started	Automatic
IBM AFS Client	Started	Automatic
IBM AFS Server	Started	Automatic
License Logging Service	Started	Automatic
Messenger	Started	Automatic

Figure 138. The Control Panel Services Window

2. You can display the status of AFS Server services by using the following methods:

   - The AFS Server Configuration window

     Click the **Start** button on the Windows NT desktop taskbar, point to **Settings**, and click **Control Panel**. From the Windows NT Control Panel, double-click the **AFS Server Configuration** icon. In the AFS Server Configuration window, click the **Services** tab to display the running AFS services, as shown in Figure 139 on page 186.

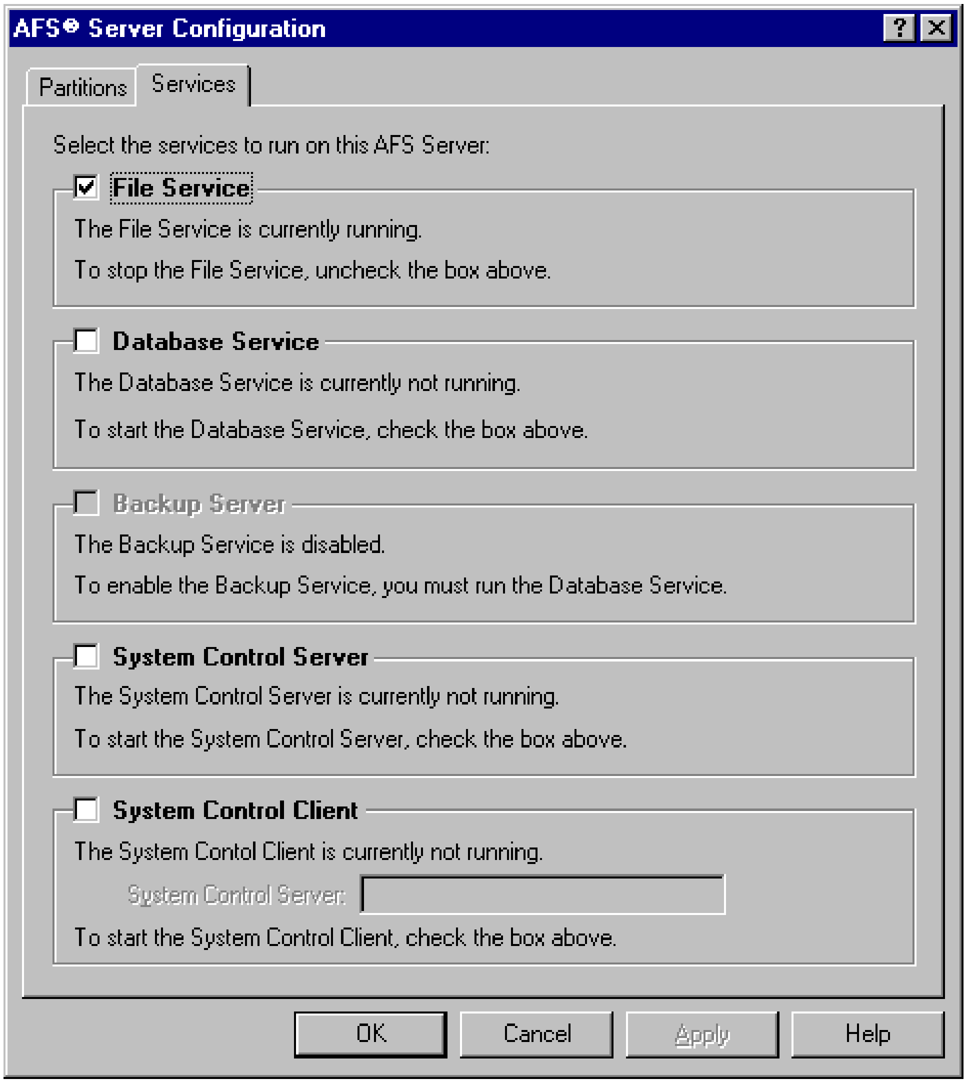

Select the services to run on this AFS Server:

☑ File Service

The File Service is currently running.

To stop the File Service, uncheck the box above.

□ Database Service

The Database Service is currently not running.

To start the Database Service, check the box above.

□ Backup Server

The Backup Service is disabled.

To enable the Backup Service, you must run the Database Service.

□ System Control Server

The System Control Server is currently not running.

To start the System Control Server, check the box above.

□ System Control Client

The System Contol Client is currently not running.

System Control Server:

To start the System Control Client, check the box above.

Figure 139. AFS Server Status on an AFS File Server

- The MS-DOS Command Prompt window

  Click the **Start** button on the Windows NT desktop taskbar, point to **Programs**, and click **Command Prompt**. In this window, enter the following command:

  ```
  bos status hostname -long
  ```

  where *hostname* is the host name of this AFS File Server machine.

  Figure 140 on page 187 shows a sample output of the command above.

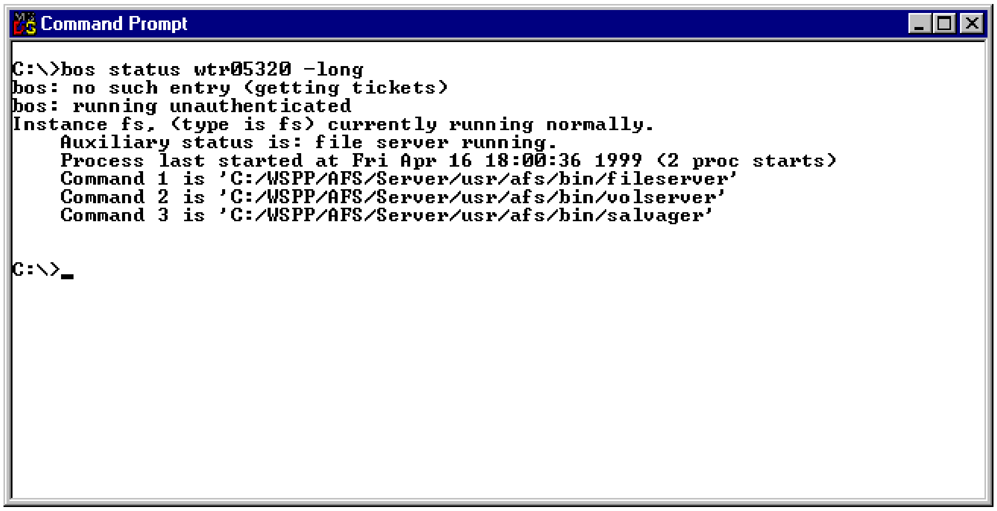

```
C:\>bos status wtr05320 -long
bos: no such entry (getting tickets)
bos: running unauthenticated
Instance fs, (type is fs) currently running normally.
    Auxiliary status is: file server running.
    Process last started at Fri Apr 16 18:00:36 1999 (2 proc starts)
    Command 1 is 'C:/WSPP/AFS/Server/usr/afs/bin/fileserver'
    Command 2 is 'C:/WSPP/AFS/Server/usr/afs/bin/volserver'
    Command 3 is 'C:/WSPP/AFS/Server/usr/afs/bin/salvager'


C:\>_
```

Figure 140. Displaying AFS Server Process Status on a Windows NT System

The AFS Server Configuration Services window and the AFS command

```
bos status
```

display the status of the AFS server processes listed in the BOS configuration file, C:\WSPP\AFS\Server\usr\afs\local\BosConfig. The contents of that file are shown in Figure 141 on page 187.

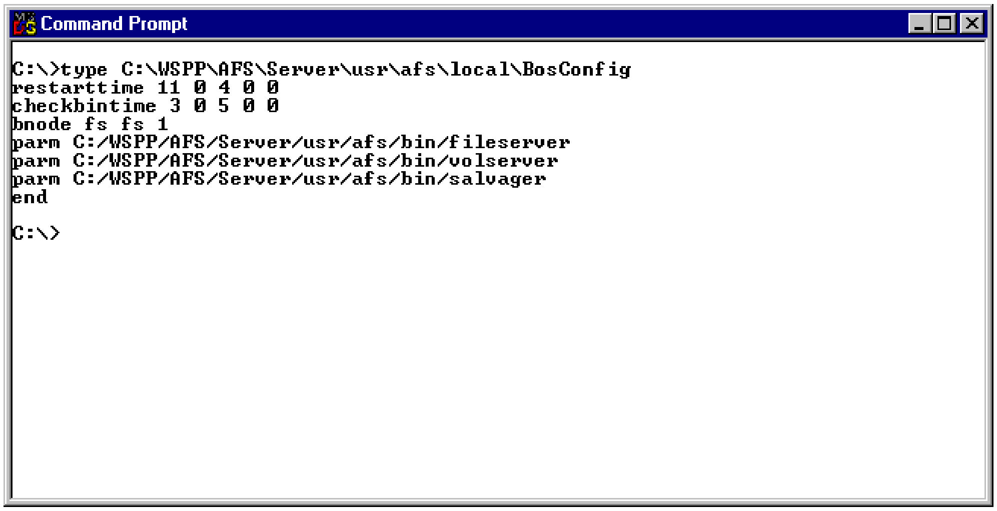

```
C:\>type C:\WSPP\AFS\Server\usr\afs\local\BosConfig
restarttime 11 0 4 0 0
checkbintime 3 0 5 0 0
bnode fs fs 1
parm C:/WSPP/AFS/Server/usr/afs/bin/fileserver
parm C:/WSPP/AFS/Server/usr/afs/bin/volserver
parm C:/WSPP/AFS/Server/usr/afs/bin/salvager
end

C:\>
```

Figure 141. Listing of the BOS Configuration File on an AFS File Server

3. In the AFS Server Configuration window, click the **Partitions** tab to display the AFS partitions, as shown in Figure 142 on page 188.

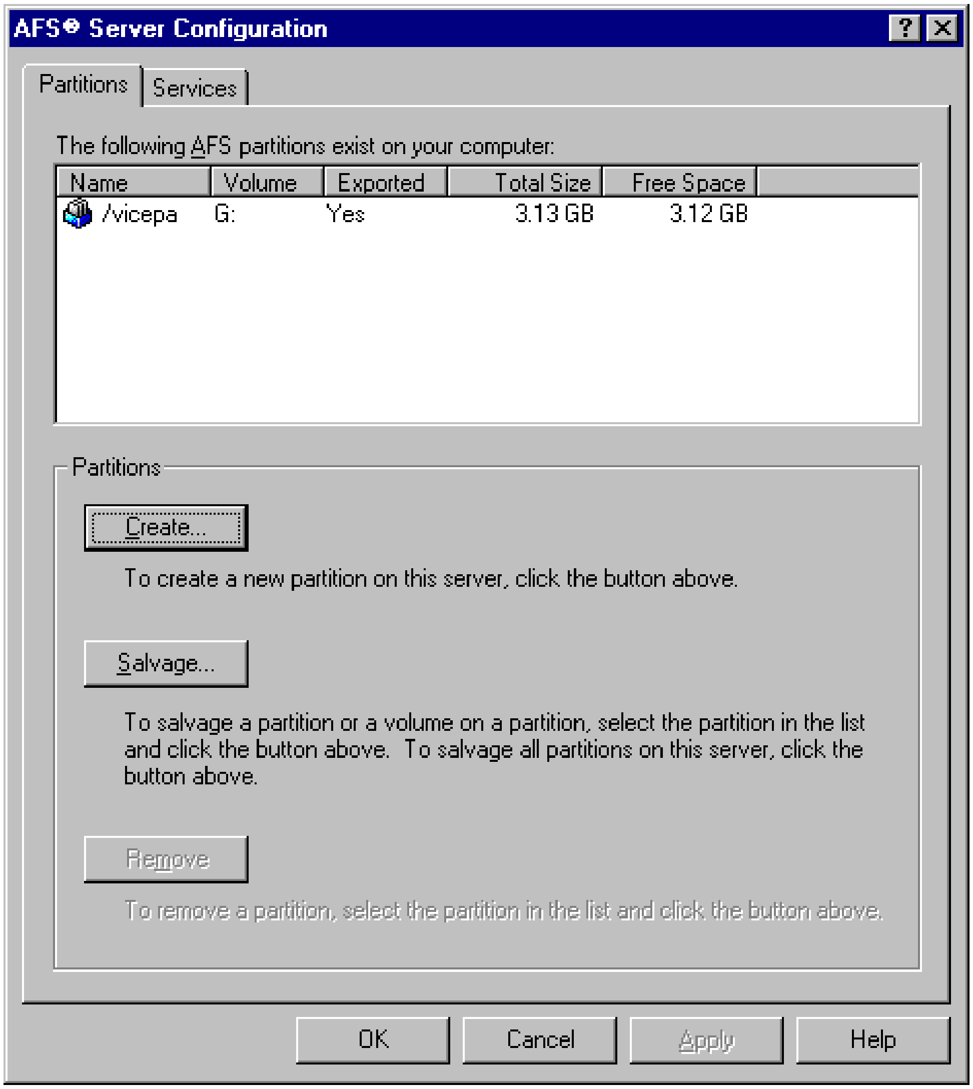

The following AFS partitions exist on your computer:

Name	Volume	Exported	Total Size	Free Space
/vicepa	G:	Yes	3.13 GB	3.12 GB

To create a new partition on this server, click the button above.

To salvage a partition or a volume on a partition, select the partition in the list and click the button above. To salvage all partitions on this server, click the button above.

To remove a partition, select the partition in the list and click the button above.

Figure 142. Listing of AFS Partitions

4. Display the status of the AFS Client from the AFS Client Configuration window. To invoke this window, click the **Start** button on the Windows NT desktop taskbar, point to **Settings**, and click **Control Panel**. From the Windows NT Control Panel, double-click the **AFS Client Configuration** icon. The AFS Client Configuration General window displays the status of the AFS Client service and the AFS cell name of the machine, as shown in Figure 143 on page 189.

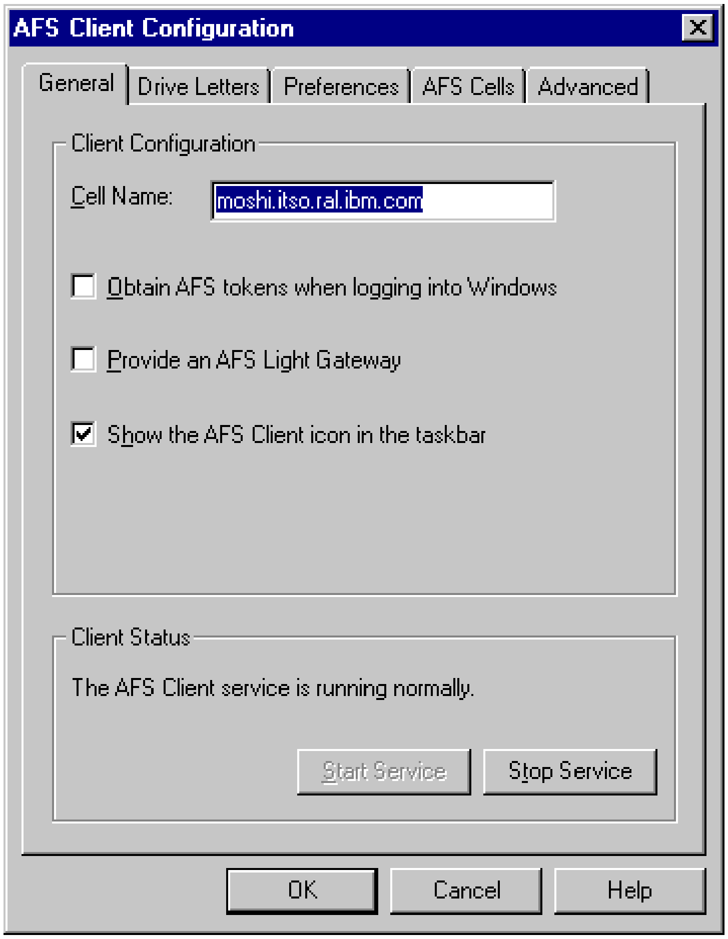

AFS Client Configuration

Client Configuration

Cell Name: moshi.itso.ral.ibm.com

□ Obtain AFS tokens when logging into Windows

□ Provide an AFS Light Gateway

☑ Show the AFS Client icon in the taskbar

Client Status

The AFS Client service is running normally.

Figure 143. The AFS Client Configuration Window

5. List the running AFS processes from the Windows NT Task Manager window. To invoke this window, position the mouse pointer on the Windows NT desktop taskbar, click the right mouse button, and click **Task Manager...**. The Windows NT Task Manager window appears. Click the **Processes** tab to show the current processes, as shown in the following figure:

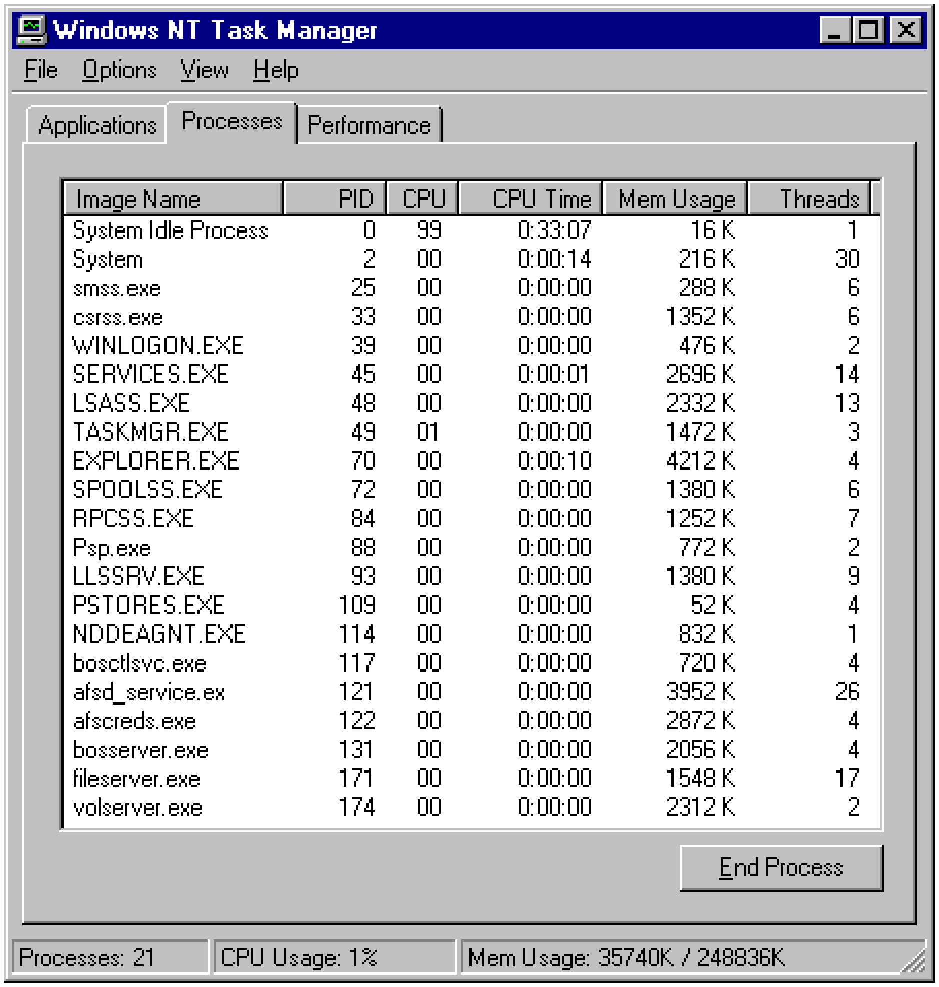

Image Name	PID	CPU	CPU Time	Mem Usage	Threads
System Idle Process	0	99	0:33:07	16 K	1
System	2	00	0:00:14	216 K	30
smss.exe	25	00	0:00:00	288 K	6
csrss.exe	33	00	0:00:00	1352 K	6
WINLOGON.EXE	39	00	0:00:00	476 K	2
SERVICES.EXE	45	00	0:00:01	2696 K	14
LSASS.EXE	48	00	0:00:00	2332 K	13
TASKMGR.EXE	49	01	0:00:00	1472 K	3
EXPLORER.EXE	70	00	0:00:10	4212 K	4
SPOOLSS.EXE	72	00	0:00:00	1380 K	6
RPCSS.EXE	84	00	0:00:00	1252 K	7
Psp.exe	88	00	0:00:00	772 K	2
LLSSRV.EXE	93	00	0:00:00	1380 K	9
PSTORES.EXE	109	00	0:00:00	52 K	4
NDDEAGNT.EXE	114	00	0:00:00	832 K	1
bosctlsvc.exe	117	00	0:00:00	720 K	4
afsd_service.ex	121	00	0:00:00	3952 K	26
afscreds.exe	122	00	0:00:00	2872 K	4
bosserver.exe	131	00	0:00:00	2056 K	4
fileserver.exe	171	00	0:00:00	1548 K	17
volserver.exe	174	00	0:00:00	2312 K	2

Processes: 21 CPU Usage: 1% Mem Usage: 35740K / 248836K

Figure 144. Listing of AFS Processes on an AFS File Server

The AFS BOS server process is bosserver.exe. The AFS fs processes are fileserver.exe, and volserver.exe. The AFS client process is afsd_service.exe. Figure 144 on page 190 shows that all these processes are running correctly.

To display the Threads column in the Windows NT Task Manager, click the **View** button of the menu bar, and then click **Select Columns...**. In the Select Columns window, check the **Thread Count** box and click **OK** to continue. The number of threads per process information is useful in AFS tuning.

5.6 Configuring an AFS Client Machine on Windows NT Systems

Before you can configure a Windows NT machine to be an AFS Client to an existing AFS cell, you must install the necessary AFS files from the IBM WebSphere Performance Pack for Multiplatforms CD to the machine local disk, as shown in 5.3, “Installing AFS Files on Windows NT Systems” on page 141.

The WebSphere Performance Pack Setup program creates a Program Group called IBM WebSphere, containing Performance Pack Afs, as shown in Figure 145 on page 191.

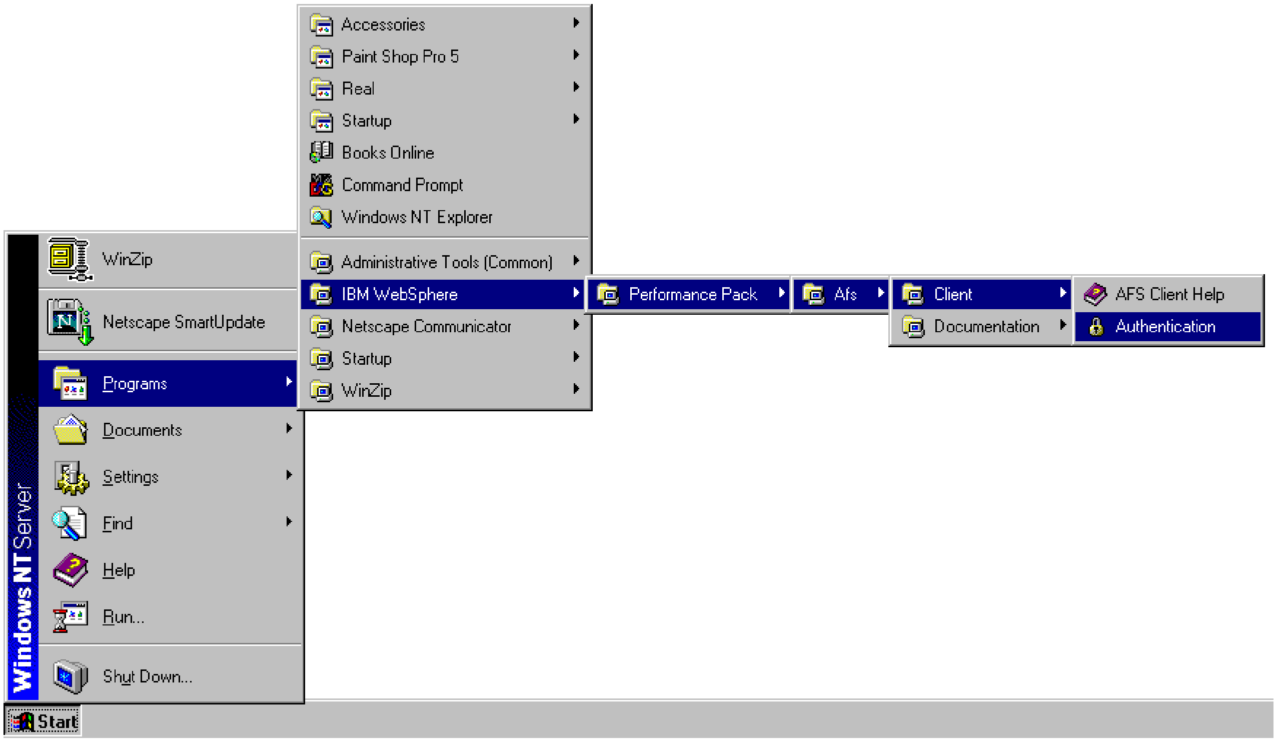

Figure 145. The WebSphere Performance Pack AFS Client

It also creates an AFS Client Configuration application in the Windows NT Control Panel and an IBM AFS Client entry in the Control Panel Services window.

To illustrate the configuration of the AFS Client on a Windows NT system, we use an AFS cell called moshi.itso.ral.ibm.com. This cell has a single AFS Database Server machine, called wtr05319.itso.ral.ibm.com, with the IP address of 9.24.104.206. This AFS Database Server machine is also the first AFS machine (see 5.4, “Configuring the First AFS Machine on Windows NT Systems” on page 147).

5.6.1 AFS Client Configuration Control Panel

To start the configuration of an AFS Client on your Windows NT system, perform the following tasks:

1. Log in as an administrative user, who is a member of the local Administrators group on your Windows NT system.

   After your Windows NT session starts, the system displays an error message indicating that the AFS Client needs to be configured, as shown in Figure 146 on page 192. Click **OK** to dismiss the message.

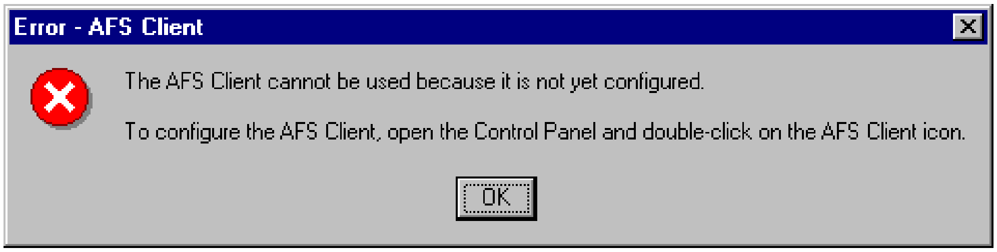

The AFS Client cannot be used because it is not yet configured.

To configure the AFS Client, open the Control Panel and double-click on the AFS Client icon.

Figure 146. AFS Client Error Message

2. To invoke the AFS Client Configuration window, click the **Start** button on the Windows NT desktop taskbar, point to **Settings**, and click **Control Panel**. From the Windows NT Control Panel, double-click the **AFS Client Configuration** icon.

3. The AFS Client Configuration window appears. For a new unconfigured AFS client machine, the AFS cell name is still unknown because it has not been defined yet. This is shown in Figure 147 on page 192.

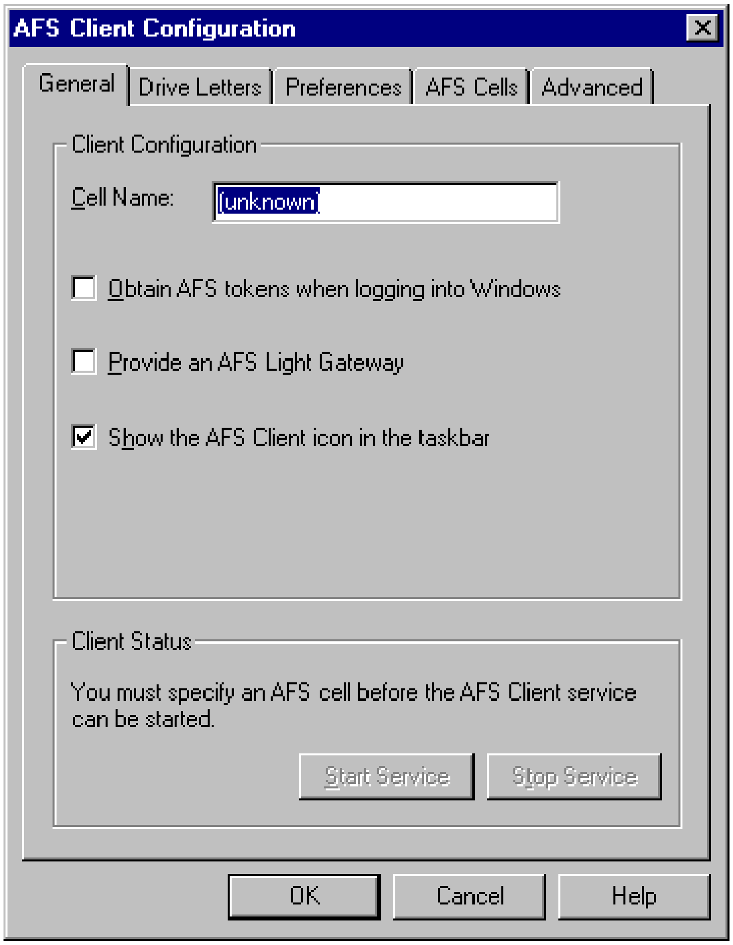

Cell Name: unknown

□ Obtain AFS tokens when logging into Windows

□ Provide an AFS Light Gateway

☑ Show the AFS Client icon in the taskbar

You must specify an AFS cell before the AFS Client service can be started.

Figure 147. The Initial Unconfigured AFS Client Configuration Window

In the General Tab of the AFS Client Configuration window, you can specify the default AFS cell name of your machine. You can also instruct the machine to obtain AFS tokens at login time, to provide an AFS Light Gateway[2], and to show the AFS Client icon in the Windows NT taskbar.

In addition, as soon as you enter the cell name, you can determine whether the AFS Client service is running or not. This is shown in the following figure:

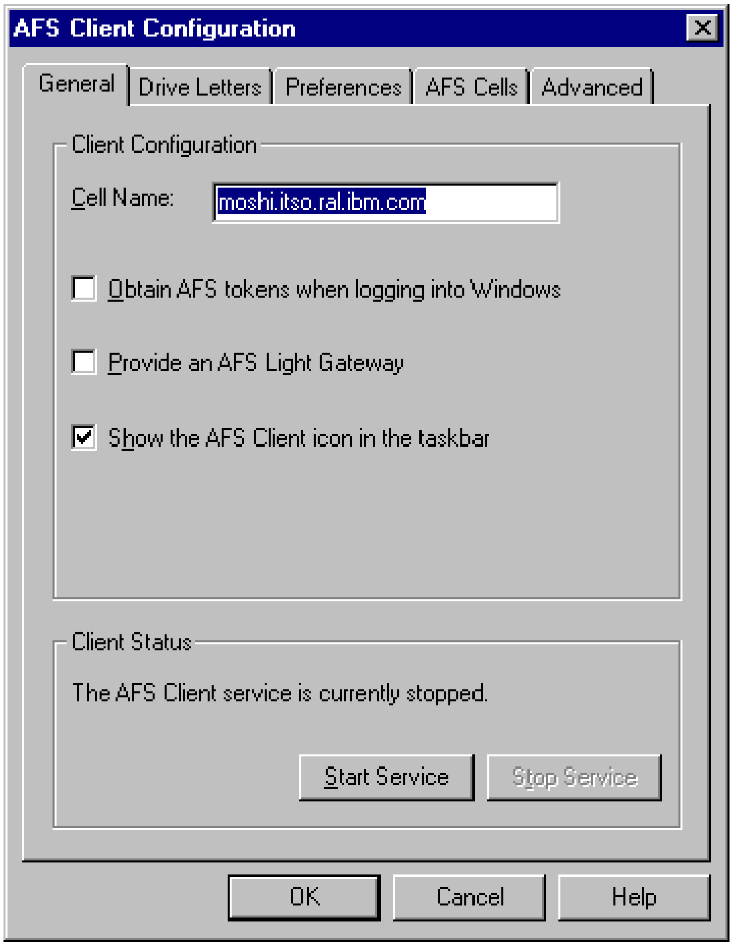

Cell Name: moshi.itso.ral.ibm.com

☐ Obtain AFS tokens when logging into Windows

☐ Provide an AFS Light Gateway

☑ Show the AFS Client icon in the taskbar

The AFS Client service is currently stopped.

Figure 148. Specifying the Default AFS Cell Name

The name of an existing AFS cell that this machine will participate in must be entered in the Cell Name field of the Client Configuration frame. We enter `moshi.itso.ral.ibm.com`, as shown in the figure above.

4. Select the **AFS Cells** tab to verify that the specified AFS cell name exists. You need to locate the specified AFS cell name, moshi.itso.ral.ibm.com, in the list of published AFS cells. The AFS Client configuration program displays the list in alphabetical order, as shown next:

[2] Running an AFS Light Gateway on the Windows NT AFS client allows Windows 95 and Windows 98 machines to access AFS through the AFS Gateway machine.

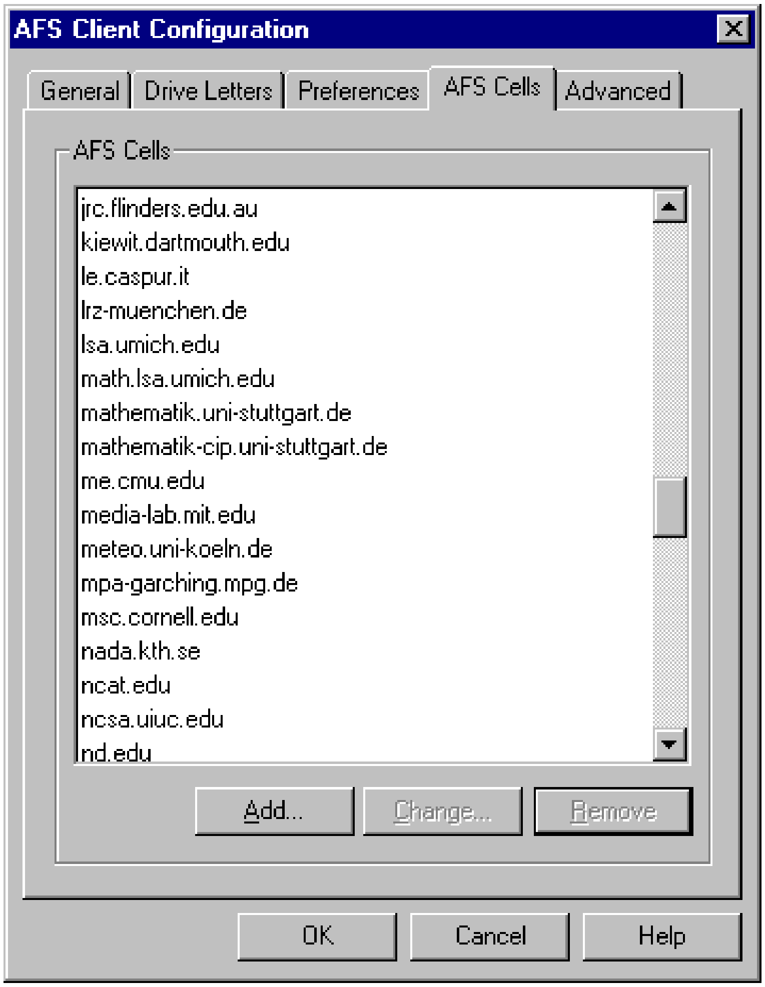

jrc.flinders.edu.au
kiewit.dartmouth.edu
le.caspur.it
lrz-muenchen.de
lsa.umich.edu
math.lsa.umich.edu
mathematik.uni-stuttgart.de
mathematik-cip.uni-stuttgart.de
me.cmu.edu
media-lab.mit.edu
meteo.uni-koeln.de
mpa-garching.mpg.de
msc.cornell.edu
nada.kth.se
ncat.edu
ncsa.uiuc.edu
nd.edu

Figure 149. The AFS Cells Tab

If you cannot find your AFS cell on this list, which is provided by AFS, your AFS cell is not yet defined. Click **Add...** to define your AFS cell.

5. In the New Cell window, you need to provide the name and a description of the new AFS cell, as shown in the following figure:

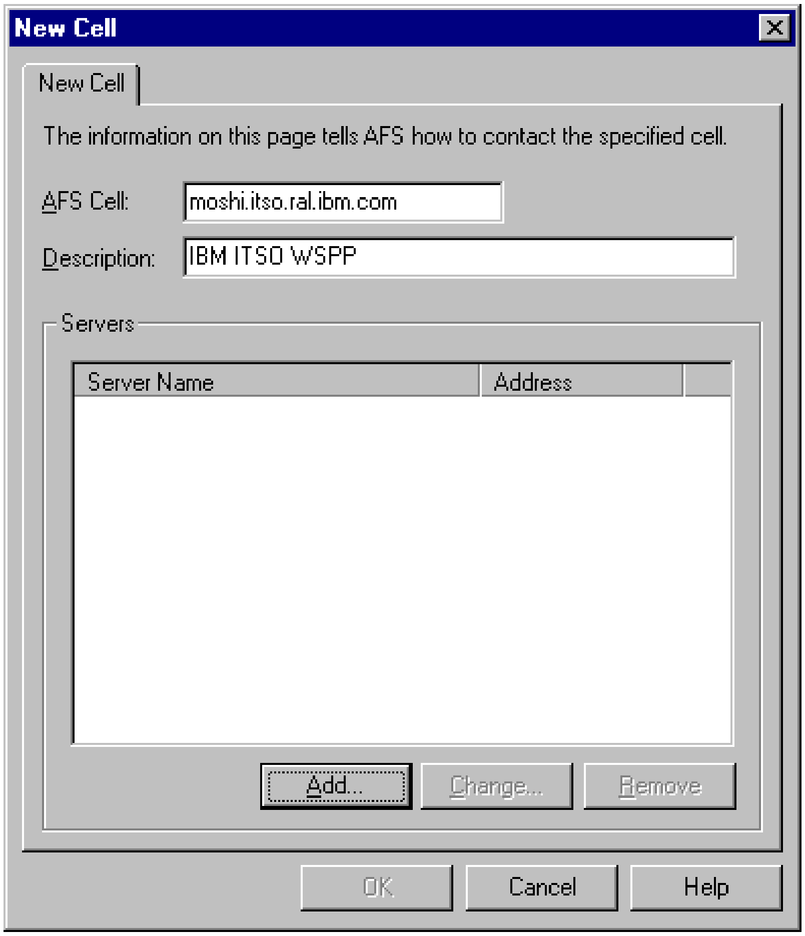

New Cell

The information on this page tells AFS how to contact the specified cell.

AFS Cell: moshi.itso.ral.ibm.com

Description: IBM ITSO WSPP

Figure 150. Defining the New AFS Cell

For example, in the AFS Cell field, we enter `moshi.itso.ral.ibm.com` for the new AFS cell name, and in the Description field, we enter `IBM ITSO WSPP` as a brief description of this cell.

Next, you need to provide the AFS Database Server information for your new AFS cell. In the New Cell window, click **Add...** to enter the host name of an AFS Database Server, which is a member of the specified AFS cell.

6. In the Add Server window, the AFS Client configuration program prompts you for the name of an AFS Volume Location Server in the specified cell, as shown in the following figure:

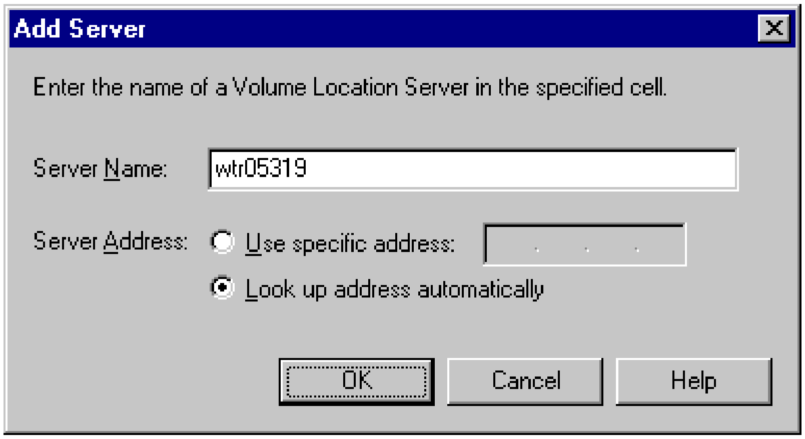

Add Server

Enter the name of a Volume Location Server in the specified cell.

Server Name: wtr05319

Server Address: ○ Use specific address: . . .

● Look up address automatically

Figure 151. Defining an AFS Database Server

An AFS Volume Location Server is a machine that stores and maintains in its physical location an AFS Volume Location Database (VLDB), which is one of the four AFS databases (see 2.5.1, “AFS Servers” on page 33).

In our scenario, the AFS Database Server is the first AFS machine. Therefore, in the Server Name field of the Add Server window, we enter `wtr05319`, the host name of the first AFS machine in our AFS cell, moshi.itso.ral.ibm.com.

The AFS Client configuration program provides you with the option of specifying an IP address or using the Domain Name System (DNS) to look up the IP address automatically. In our environment, all machines are registered to DNS. Therefore, we select the default option to use DNS to resolve host name to IP address.

Click **OK** to add this AFS Server to the specified AFS cell.

7. The AFS Client configuration program returns you to the New Cell window, displaying the information about the specified AFS cell, moshi.itso.ral.ibm.com, as shown in the following figure:

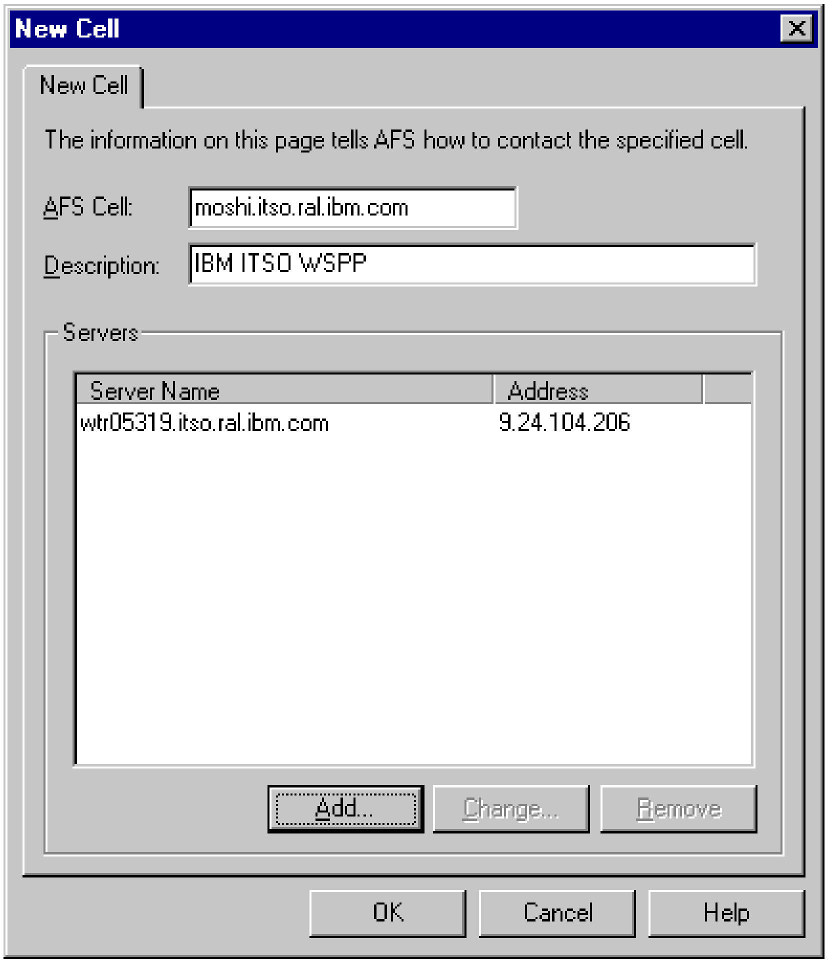

The information on this page tells AFS how to contact the specified cell.

AFS Cell: moshi.itso.ral.ibm.com

Description: IBM ITSO WSPP

Servers

Server Name	Address
wtr05319.itso.ral.ibm.com	9.24.104.206

Figure 152. Displaying the Properties of a New AFS Cell

If there is more than one AFS Database Server in your cell, click **Add...** to define other AFS Database Server machines. Otherwise, click **OK** to add the new AFS cell to the list of known AFS cells.

8. The AFS Client configuration program returns you to the AFS Cells tab of the AFS Client Configuration window, displaying the list of AFS cells with a new entry for the AFS cell we have just added. In this example, the new AFS cell, moshi.itso.ral.ibm.com, is included in the list of AFS cells on our Windows NT machine, as shown next:

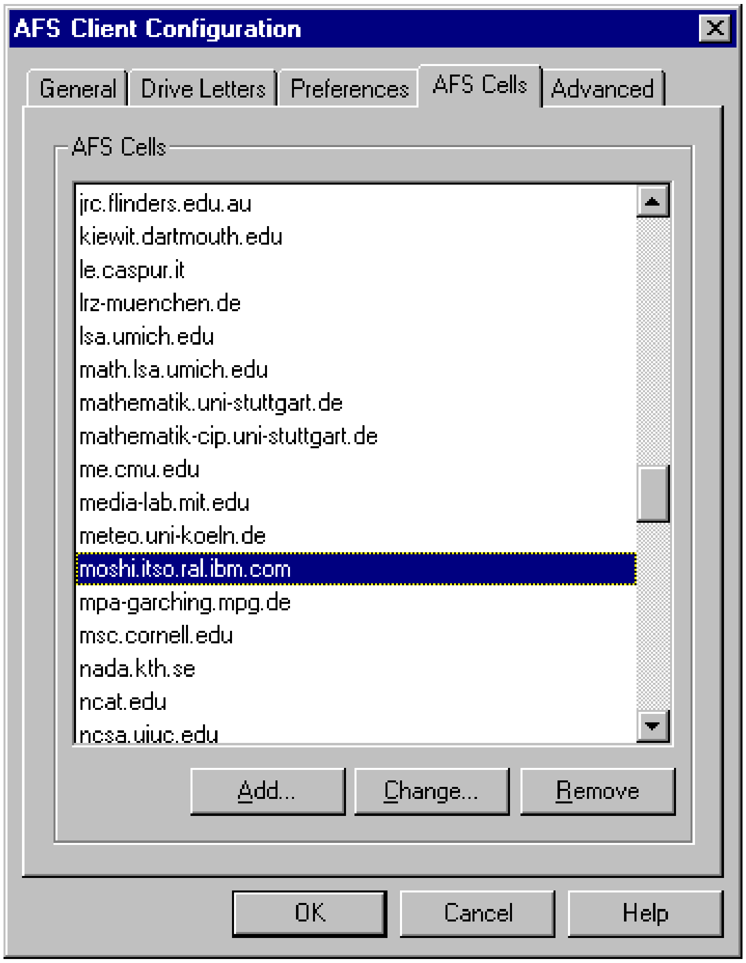

jrc.flinders.edu.au
kiewit.dartmouth.edu
le.caspur.it
lrz-muenchen.de
lsa.umich.edu
math.lsa.umich.edu
mathematik.uni-stuttgart.de
mathematik-cip.uni-stuttgart.de
me.cmu.edu
media-lab.mit.edu
meteo.uni-koeln.de
moshi.itso.ral.ibm.com
mpa-garching.mpg.de
msc.cornell.edu
nada.kth.se
ncat.edu
ncsa.uiuc.edu

Figure 153. Listing of the New AFS Cell

Notice that this process updates the list of AFS Database Server machines for AFS cells stored in the file afsdcell.ini, by default located in the C:\WINNT directory.

9. To examine this new AFS cell, moshi.itso.ral.ibm.com, double-click on its name in the AFS Cells tab of the AFS Client Configuration window. The following window appears:

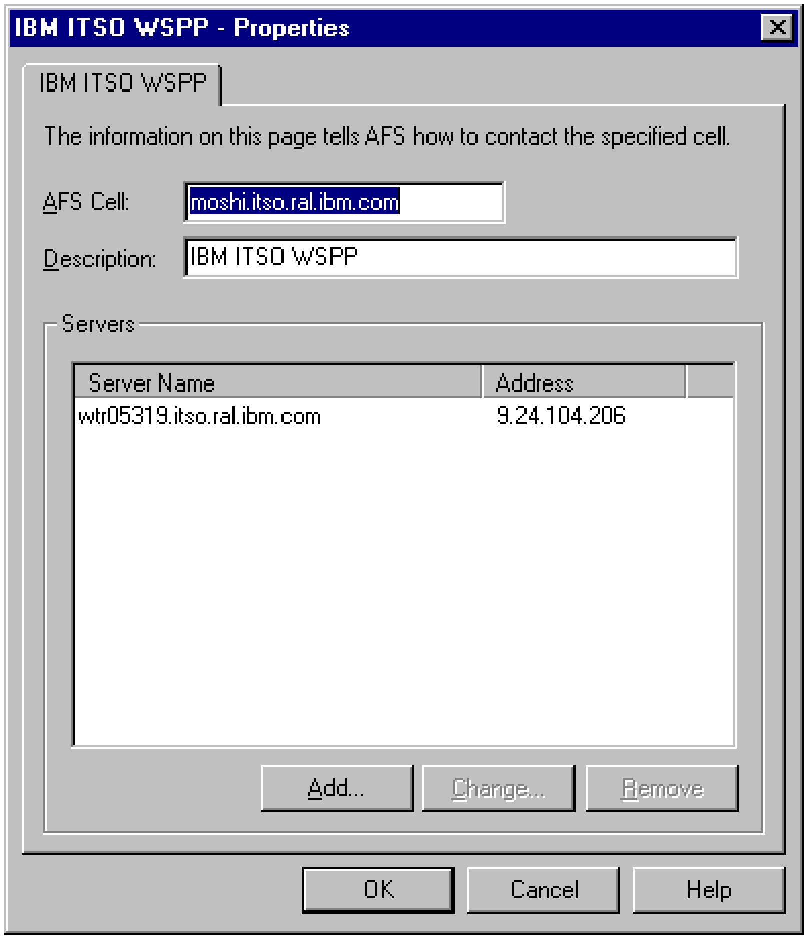

The information on this page tells AFS how to contact the specified cell.

AFS Cell: moshi.itso.ral.ibm.com

Description: IBM ITSO WSPP

Servers

Server Name	Address
wtr05319.itso.ral.ibm.com	9.24.104.206

Figure 154. Displaying the Properties of an Existing AFS Cell

Click **OK** to dismiss the Properties window above and return to the AFS Cells tab of the AFS Client Configuration window.

10.In the AFS Client Configuration window, select the **Advanced** tab to display the default AFS cache configuration parameters, as shown in the following figure:

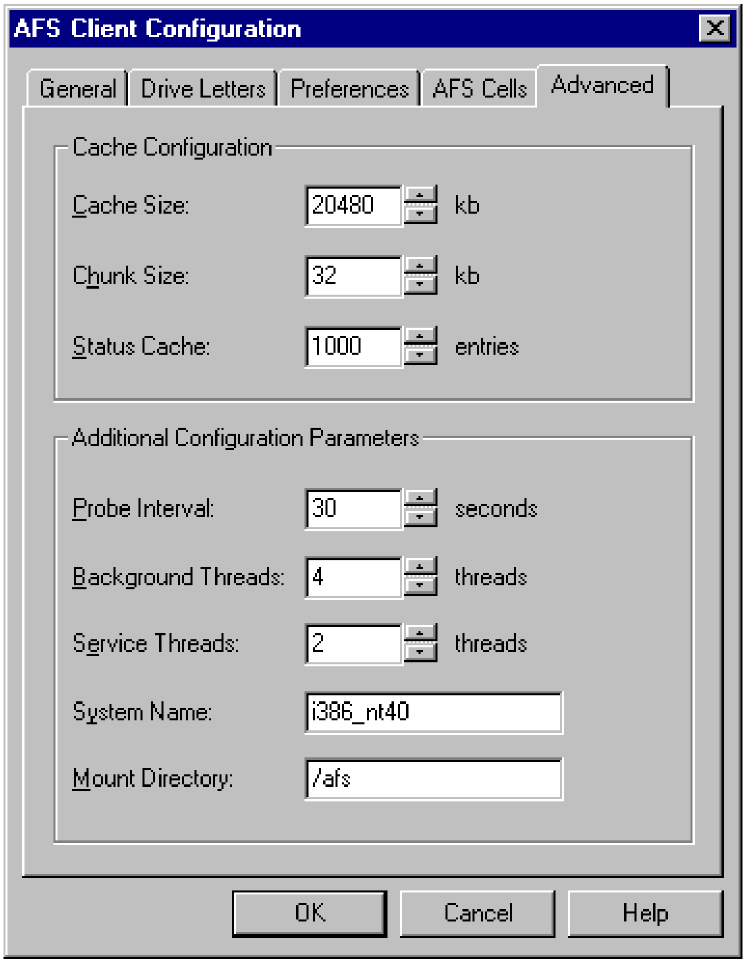

Cache Configuration

Cache Size: 20480 kb

Chunk Size: 32 kb

Status Cache: 1000 entries

Additional Configuration Parameters

Probe Interval: 30 seconds

Background Threads: 4 threads

Service Threads: 2 threads

System Name: i386_nt40

Mount Directory: /afs

Figure 155. Displaying the Default AFS Cache Configuration Parameters

In AFS Version 3.5 for Windows NT, included in WebSphere Performance Pack V2.0, the AFS cache is implemented as a memory mapped file cache, separated from the Windows NT system page file. The default name for this file cache is C:\AFSCache. The default cache size is 20480 KB or 20 MB. In this example, we use the default AFS cache size. You may need to increase the AFS cache size to improve performance.

The other AFS cache configuration parameters will be discussed in Chapter 7, “AFS Client Management” on page 303. In this basic configuration, we can accept the default values, which can also be changed later.

11. Select the **General** tab and click the **Start Service** button to start the AFS Client service, as shown in the following figure:

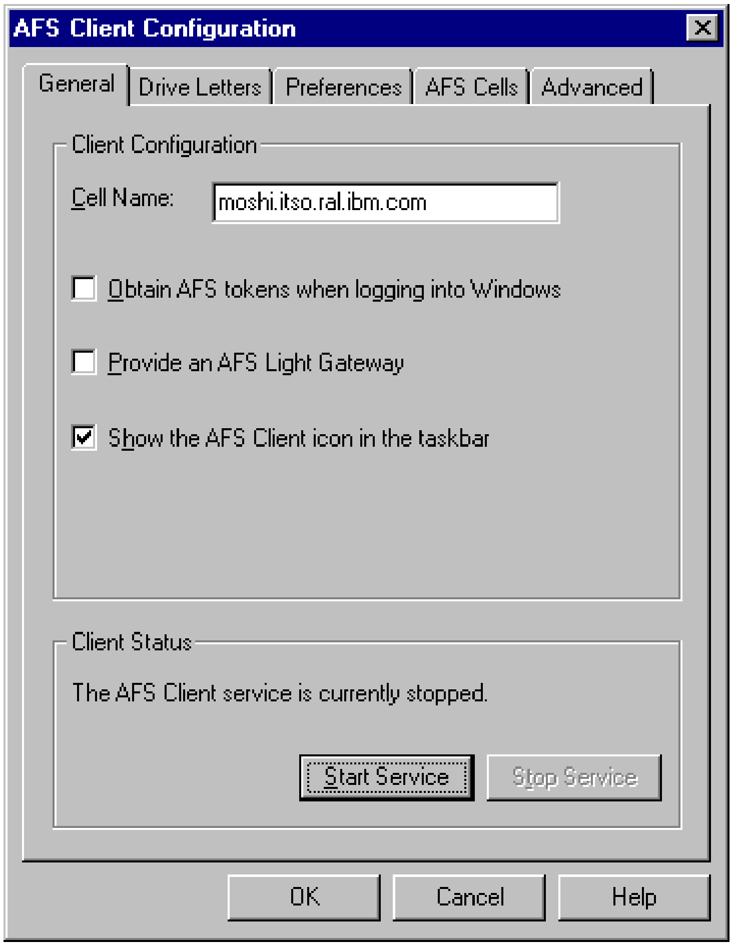

AFS Client Configuration

Client Configuration

Cell Name: moshi.itso.ral.ibm.com

☐ Obtain AFS tokens when logging into Windows

☐ Provide an AFS Light Gateway

☑ Show the AFS Client icon in the taskbar

Client Status

The AFS Client service is currently stopped.

Figure 156. Clicking the Start Service Button

The AFS Client service will display a status window informing you of its progress, as shown in Figure 157 on page 201.

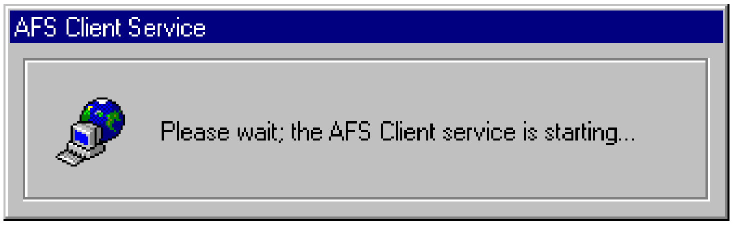

AFS Client Service

Please wait; the AFS Client service is starting...

Figure 157. The AFS Client Service Window

Configuration of the AFS Client is complete, as shown in the following figure:

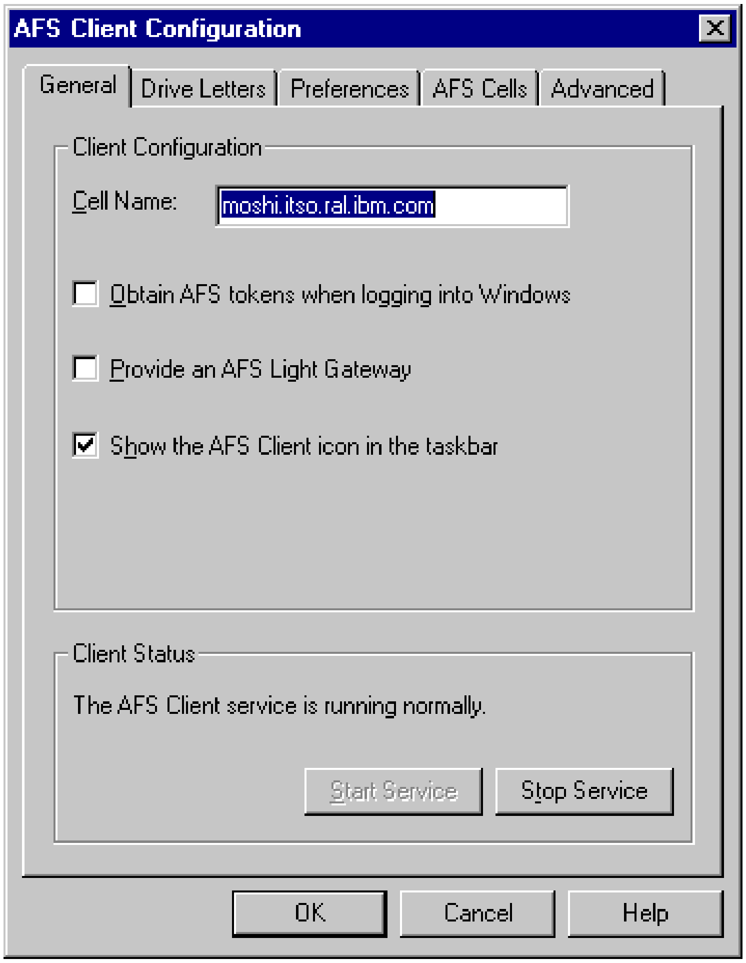

Cell Name: moshi.itso.ral.ibm.com

□ Obtain AFS tokens when logging into Windows

□ Provide an AFS Light Gateway

☑ Show the AFS Client icon in the taskbar

The AFS Client service is running normally.

Figure 158. AFS Client Configuration Window – General Tab

The AFS Client automatically starts when you restart Windows NT, unless you want disable that option in the Windows NT Control Panel Services window.

5.6.2 Mounting the AFS File System

To access the AFS name space, you need to map AFS paths to Windows NT logical drive letters. There are several ways you can do this. In fact, you can use one of the following options:

- The AFS Client Configuration from the Windows NT Control Panel
- The AFS Client from the WebSphere Performance Pack AFS programs group
- The command:

  ```
  net use
  ```

  from an MS-DOS Command Prompt window
- The **Map Network Drive...** option from the Windows NT Explorer Tools menu

5.6.2.1 Mapping AFS with the AFS Client Configuration Application

This section demonstrates the mapping of AFS paths to Windows NT logical drive letters using the AFS Client Configuration application from the Windows NT Control Panel.

1. Click the **Start** button on the Windows NT desktop taskbar, point to **Settings**, and click **Control Panel**. From the Windows NT Control Panel, double-click the **AFS Client Configuration** icon.
2. From the AFS Client Configuration window, select the **Drive Letters** tab, as shown in the following figure:

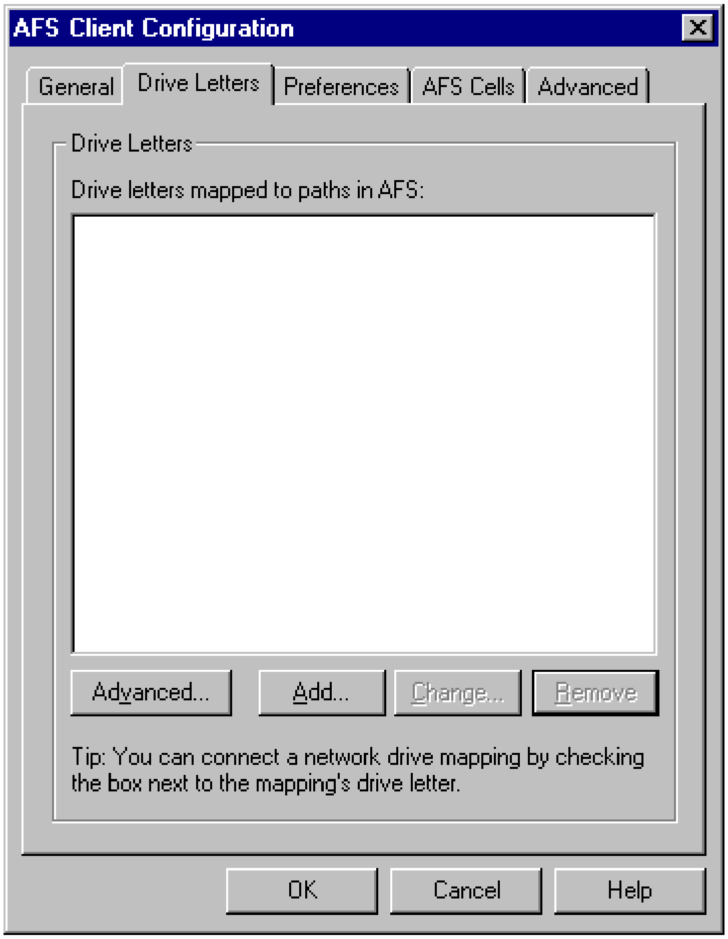
AFS Client Configuration

Drive Letters

Drive letters mapped to paths in AFS:

Tip: You can connect a network drive mapping by checking the box next to the mapping's drive letter.

Figure 159. The Drive Letters Tab of the AFS Client Configuration Window

3. Initially, there is no drive map to any AFS path. Click **Add...** to map a Windows NT logical drive letter to an AFS path.
4. From the Map Drive Letter dialog box, in the Drive Letter field, select the Windows NT logical drive letter to map the AFS path \afs to, as shown in the following figure:

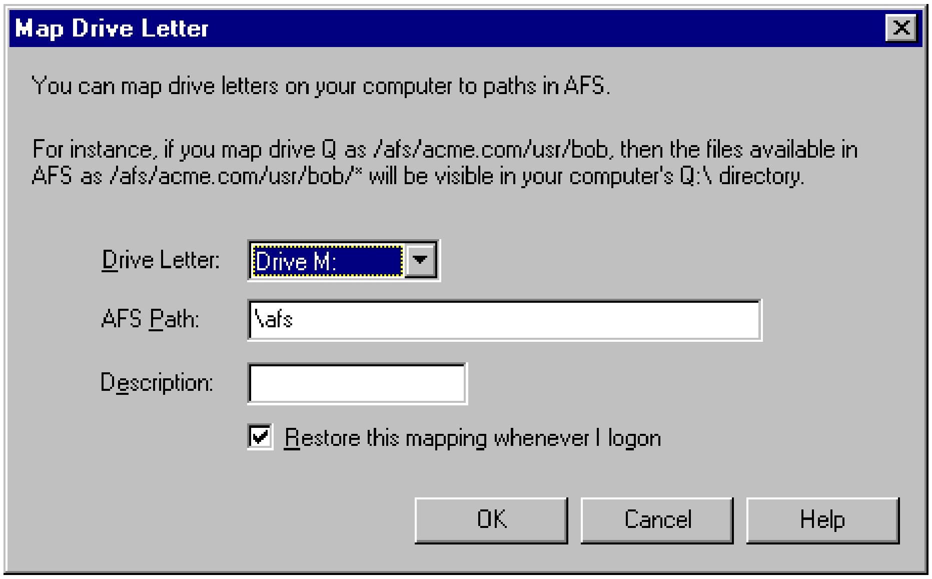

Map Drive Letter

You can map drive letters on your computer to paths in AFS.

For instance, if you map drive Q as /afs/acme.com/usr/bob, then the files available in AFS as /afs/acme.com/usr/bob/* will be visible in your computer's Q:\ directory.

Drive Letter: Drive M:

AFS Path: \afs

Description:

☑ Restore this mapping whenever I logon

Figure 160. Mapping Drive Letter to AFS Path

In this example, we select **Drive M:**. If you want to access this AFS path at every login session, keep the default option **Restore this mapping whenever I logon**. Click **OK** to map the specified AFS path to this drive letter.

5. The AFS Client Configuration program returns you to the Drive Letters tab, displaying the drive mapped to the specified AFS path, as shown in Figure 161 on page 205.

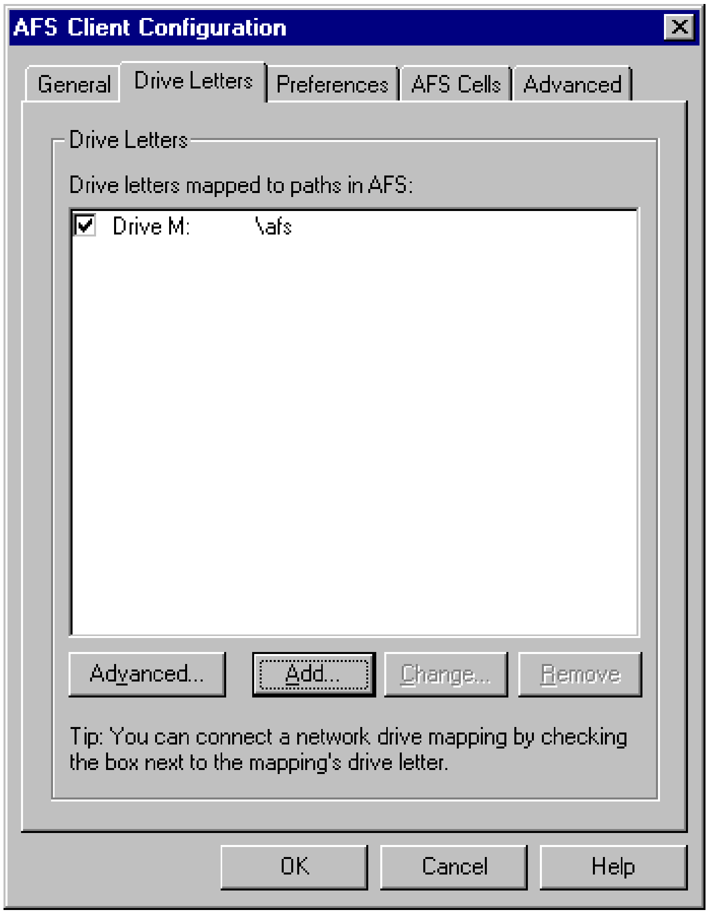

Drive letters mapped to paths in AFS:

☑ Drive M: \afs

Tip: You can connect a network drive mapping by checking the box next to the mapping's drive letter.

Figure 161. Displaying the Drive Letters Mapped to AFS Paths

To display the submounts, click **Advanced...** in the Drive Letters tab of the AFS Client Configuration window. The Submounts screen is automatically brought up, as shown in Figure 162 on page 206. A *submount* is a shared network drive that has been mapped to AFS.

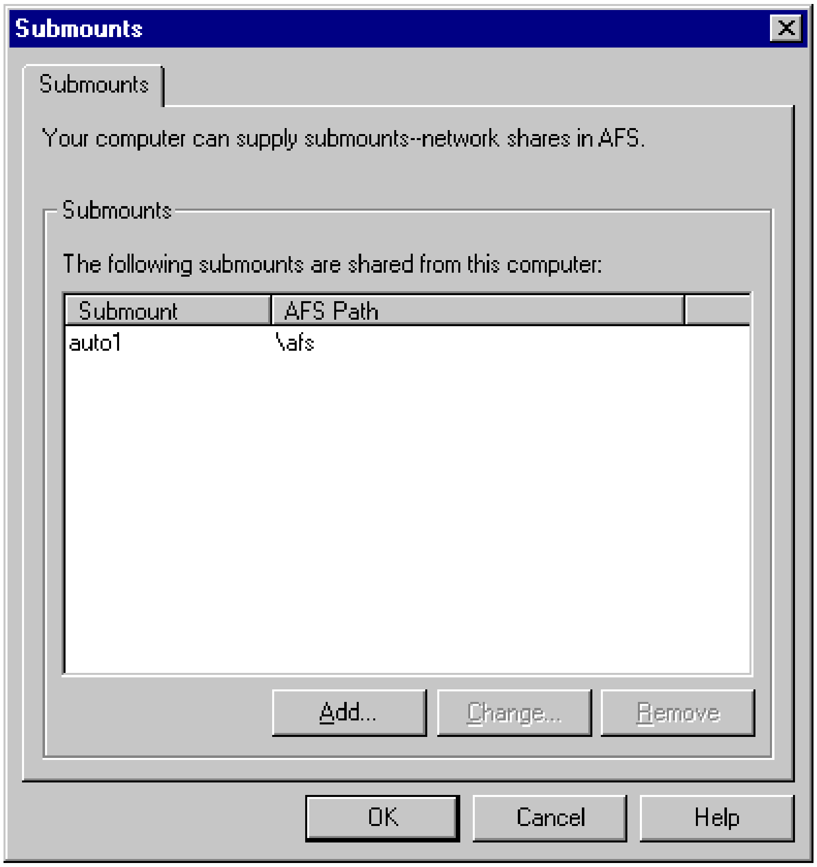

Your computer can supply submounts--network shares in AFS.

The following submounts are shared from this computer:

Submount	AFS Path
auto1	\afs

Figure 162. Displaying the Submounts for AFS Paths

Click **OK** to return to the Drive Letters tab.

5.6.2.2 Mapping AFS with the AFS Client Program

This section demonstrates the mapping of AFS paths to Windows NT logical drive letters using the AFS Client program.

1. You can invoke the AFS Client program by using one of the following methods:
   - Click the **Start** button on your Windows NT desktop taskbar, point to **Programs**, point to **IBM WebSphere**, point to **Performance Pack**, point to **Afs**, point to **Client**, and click **Authentication**, as depicted in Figure 163 on page 207.
   - Alternately, if you selected the **Show the AFS Client icon in the taskbar** check box in the General tab of the AFS Client Configuration window (see Figure 158 on page 202), there is a lock icon on the Windows NT desktop taskbar next to the clock, identifying the AFS Client program, as shown in Figure 163 on page 207. Click this lock icon to start the AFS Client program.

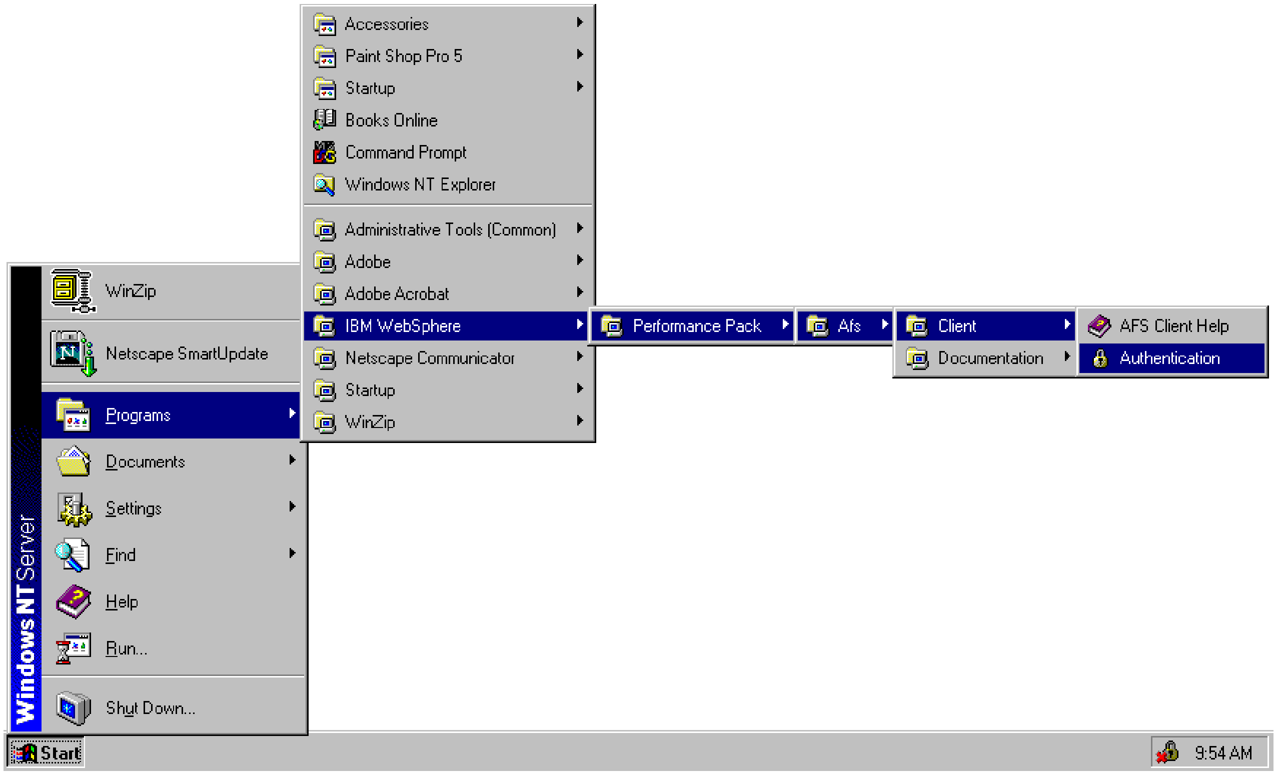

Figure 163. The WebSphere AFS Client Program

2. The AFS Client window appears as shown in Figure 164 on page 207.

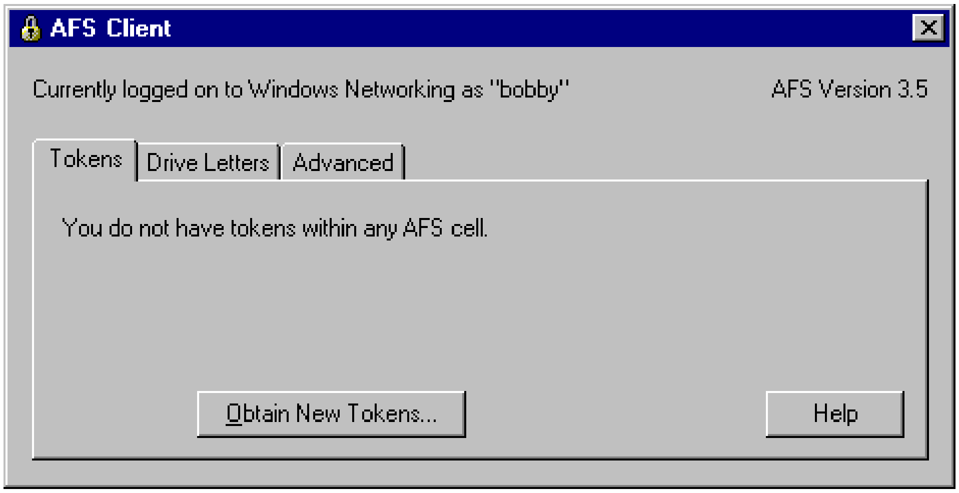

Currently logged on to Windows Networking as "bobby"

AFS Version 3.5

You do not have tokens within any AFS cell.

Figure 164. The AFS Client Program

Select the **Drive Letters** tab from the AFS Client window, as shown in Figure 165 on page 208. Initially, there is no drive map to any AFS path.

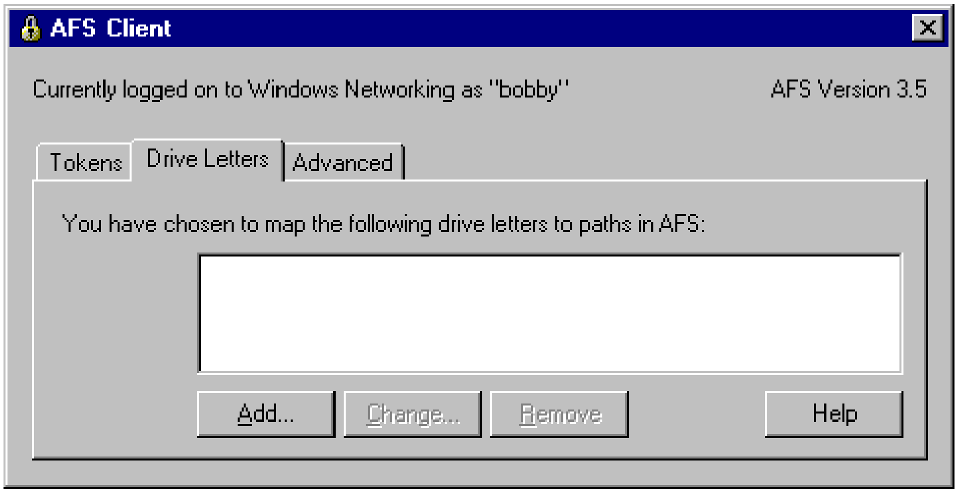

AFS Client

Currently logged on to Windows Networking as "bobby"

AFS Version 3.5

You have chosen to map the following drive letters to paths in AFS:

Figure 165. The Initial Drive Letters Tab of the AFS Client Program

3. Click **Add...** to map a Windows NT logical drive letter to an AFS path.
4. From the Map Drive Letter dialog box, in the Drive Letter field, select the Windows NT logical drive letter to map the AFS path \afs to. In this example, we select **Drive M:**. If you want to access this AFS path at every login session, keep the default option **Restore this mapping whenever I logon**, as shown in Figure 166 on page 208.

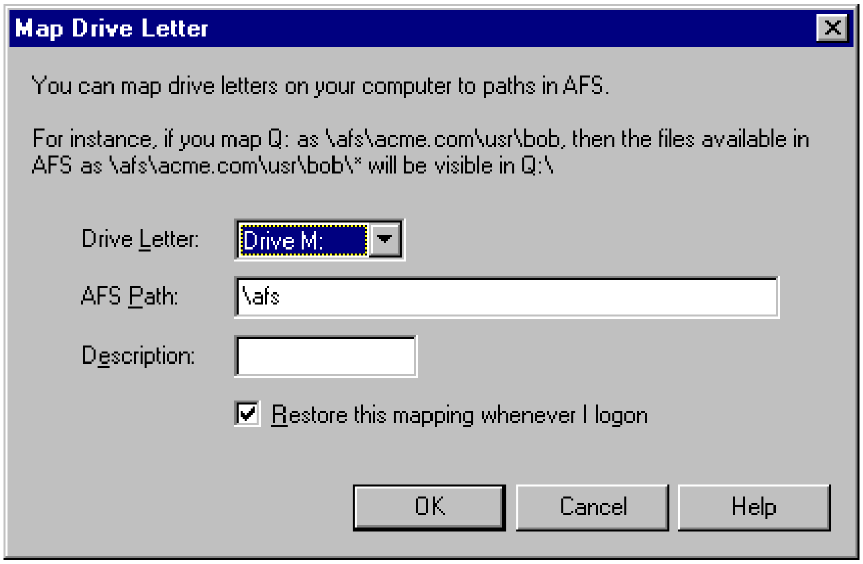

Map Drive Letter

You can map drive letters on your computer to paths in AFS.

For instance, if you map Q: as \afs\acme.com\usr\bob, then the files available in AFS as \afs\acme.com\usr\bob\* will be visible in Q:\

Drive Letter: Drive M:

AFS Path: \afs

Description:

☑ Restore this mapping whenever I logon

Figure 166. Mapping Drive Letter to AFS Path from the AFS Client Program

5. Click **OK** to map the specified AFS path to the drive letter you have selected. The AFS Client program returns you to the Drive Letters tab, displaying the drive mapped to the specified AFS path, as shown in Figure 167 on page 209.

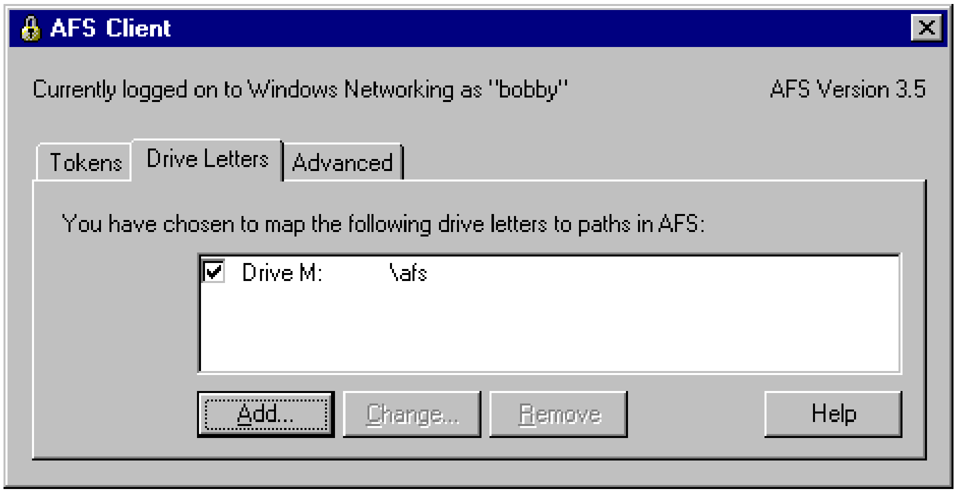

Currently logged on to Windows Networking as "bobby" AFS Version 3.5

You have chosen to map the following drive letters to paths in AFS:

☑ Drive M: \afs

Figure 167. Displaying the Drive Letters Mapped to AFS Paths

5.6.2.3 Mapping AFS with the net use Command

This section demonstrates the mapping of AFS paths to Windows NT logical drive letters using the command:

```
net use
```

From an MS-DOS Command Prompt window, enter:

```
net use drive: \\hostname-afs\submount
```

where:

- *drive* is the drive letter that you want to map the AFS path to.
- *hostname* is the host name of the local Windows NT machine.
- *submount* is the name of the subdirectory to be mounted. Use `auto1` or `all` to access the AFS path /afs.

For example, on the Windows NT machine WTR05212, to map the drive M: to the AFS path /afs, enter:

```
net use M: \\WTR05212-afs\auto1
```

If you want the system to remember the mapping information and reconnect to AFS, use the `/persistent:yes` flag with the command above. In this example, enter:

```
net use M: \\WTR05212-afs\auto1 /persistent:yes
```

Figure 168 on page 210 demonstrates the usage of the above command:

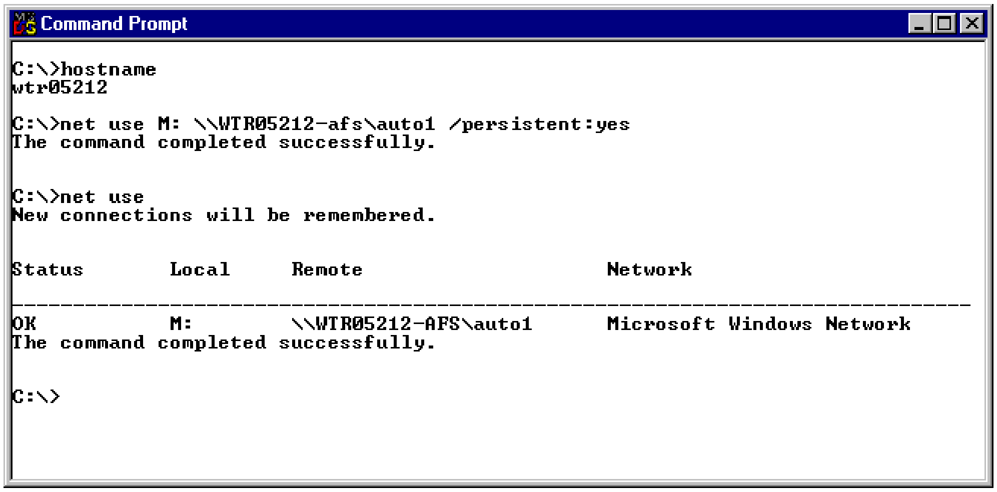

```
C:\>hostname
wtr05212

C:\>net use M: \\WTR05212-afs\auto1 /persistent:yes
The command completed successfully.


C:\>net use
New connections will be remembered.


Status       Local     Remote                    Network

-------------------------------------------------------------------------------
OK           M:        \\WTR05212-AFS\auto1      Microsoft Windows Network
The command completed successfully.


C:\>
```

Figure 168. Mapping Drive Letter to AFS Path Using the net use Command

5.6.2.4 Mapping AFS with the Windows NT Explorer

This section demonstrates the mapping of AFS paths to Windows NT logical drive letters using the Windows NT Explorer.

1. Click the **Start** button on your Windows NT desktop taskbar, point to **Programs**, and click **Windows NT Explorer**.
2. From the Windows NT Explorer window, click the **Tools** button on the Explorer menu bar, and select **Map Network Drive...**, as shown in the following figure:

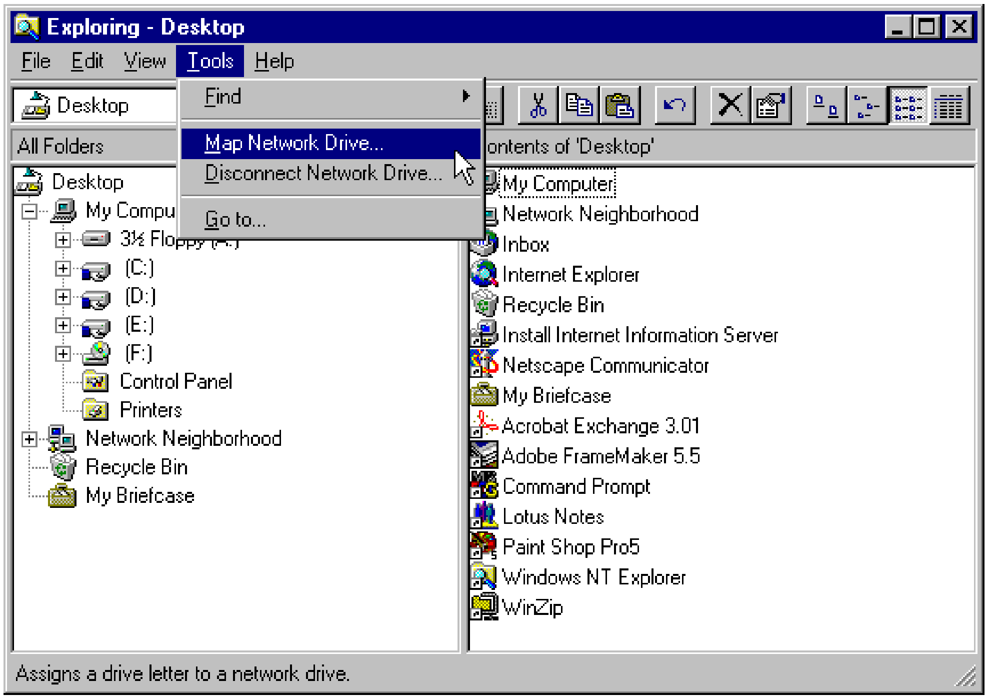

Figure 169. Mapping AFS Paths Using the Windows NT Explorer

3. The Map Network Drive window appears. In the Drive field, select the drive letter to be mapped to an AFS path. In the Path field, specify the path in the following format:

   `\\`*`hostname`*`-afs\`*`submount`*

   where:

   - *`hostname`* is the name of the Windows NT machine
   - *`submount`* is the name of the sub-directory to be mounted. Use `auto1` or `all` to access the AFS path /afs.

   For example, on the Windows NT machine WTR05212, the path to enter is:

   `\\WTR05212-afs\auto1`

   Select the **Reconnect at Logon** check box if you want this mapping to automatically take place at every system reboot. Then, click **OK** to map the drive.

   A sample Map Network Drive window is shown in the following figure:

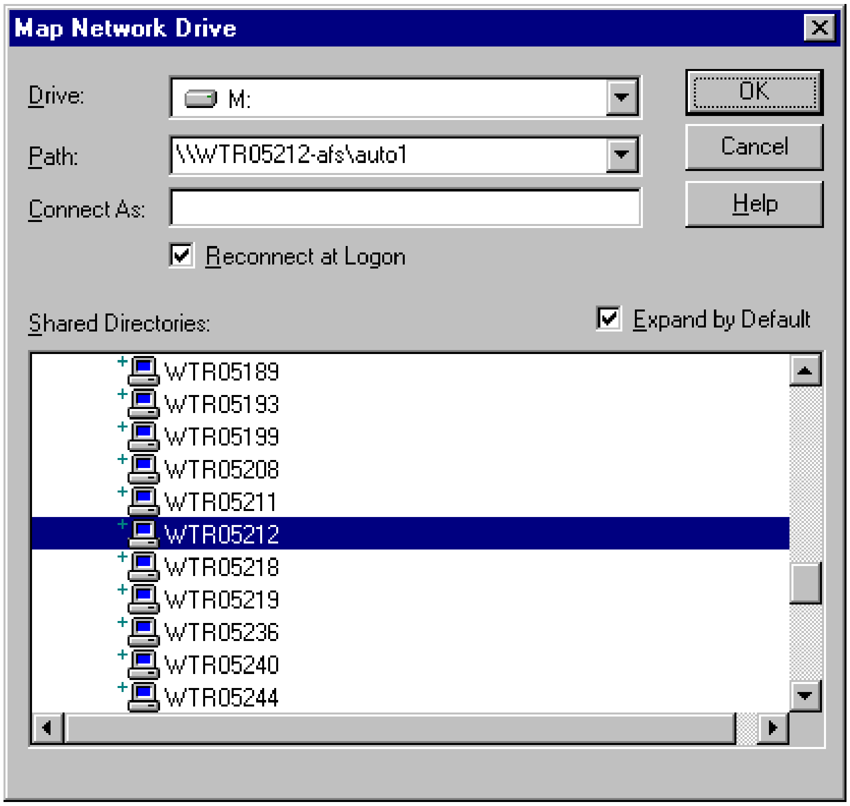

Figure 170. Mapping Drive Letter to AFS Path Using Windows NT Explorer

4. In this example, the name of the AFS cell is moshi.itso.ral.ibm.com. You can access the AFS file space via the drive letter M: using the Windows NT Explorer, as shown in Figure 171 on page 212.

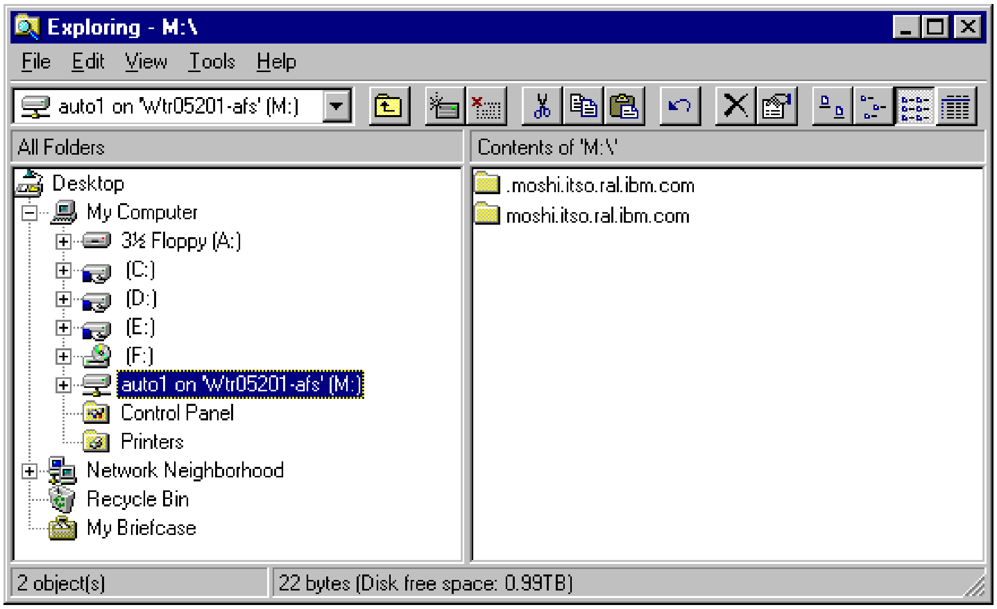

Desktop
My Computer
3½ Floppy (A:)
(C:)
(D:)
(E:)
(F:)
auto1 on 'Wtr05201-afs' (M:)
Control Panel
Printers
Network Neighborhood
Recycle Bin
My Briefcase

.moshi.itso.ral.ibm.com
moshi.itso.ral.ibm.com

2 object(s) | 22 bytes (Disk free space: 0.99TB)

Figure 171. Windows NT Explorer with Drive Letter Mapping to AFS Path

5.6.3 Verifying the AFS Client Configuration

To verify that the configuration of the AFS Client completed successfully, perform the following tasks:

1. Display the status of AFS service in the Control Panel Services window. To invoke this window, click the **Start** button on the Windows NT desktop taskbar, point to **Settings**, and click **Control Panel**. From the Windows NT Control Panel, double-click the **Services** icon. The IBM AFS Client service entry should show the Started status, as illustrated in Figure 172 on page 212.

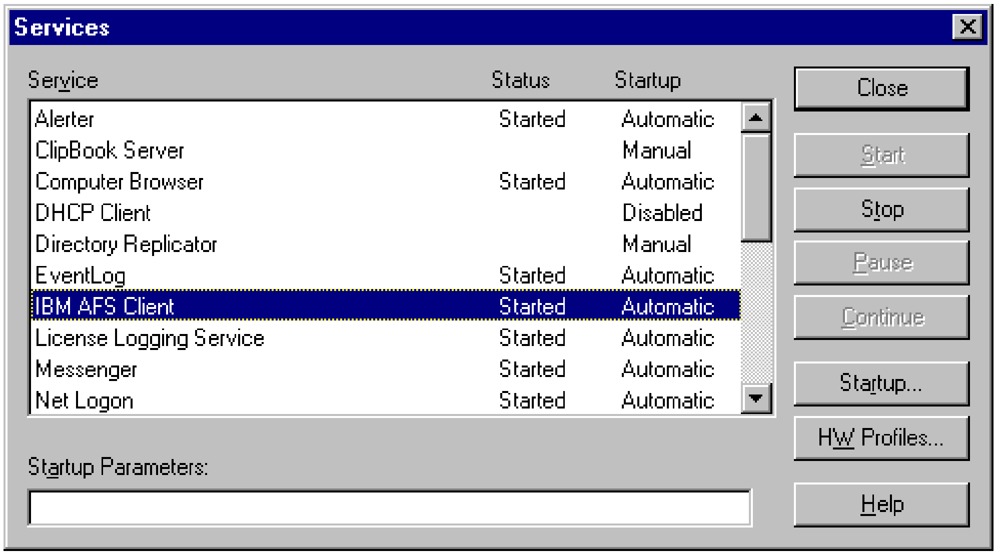

Service	Status	Startup
Alerter	Started	Automatic
ClipBook Server		Manual
Computer Browser	Started	Automatic
DHCP Client		Disabled
Directory Replicator		Manual
EventLog	Started	Automatic
IBM AFS Client	Started	Automatic
License Logging Service	Started	Automatic
Messenger	Started	Automatic
Net Logon	Started	Automatic

Figure 172. IBM AFS Client Service Status Window

2. You can display the status of AFS Client service by using the following methods:

- The AFS Client Configuration program

  Click the **Start** button on the Windows NT desktop taskbar, point to **Settings**, and click **Control Panel**. From the Windows NT Control Panel, double-click the **AFS Client Configuration** icon.

  In the AFS Client Configuration window, select the **General** tab, and you should read a message saying

  ```
  The AFS Client service is running normally
  ```

  This is shown in Figure 173 on page 213.

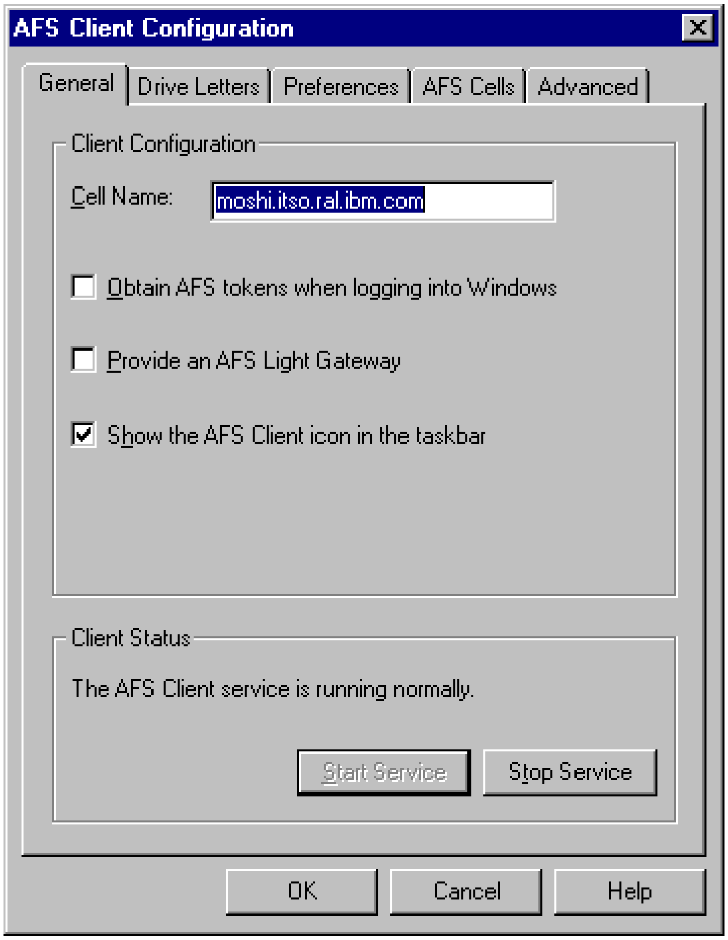

Cell Name: moshi.itso.ral.ibm.com

□ Obtain AFS tokens when logging into Windows

□ Provide an AFS Light Gateway

☑ Show the AFS Client icon in the taskbar

The AFS Client service is running normally.

Figure 173. Displaying the AFS Client Status with the AFS Client Configuration Program

In the AFS Client Configuration window, select the **Drive Letters** tab to display the drive letters mapped to AFS paths, as shown in Figure 174 on page 214.

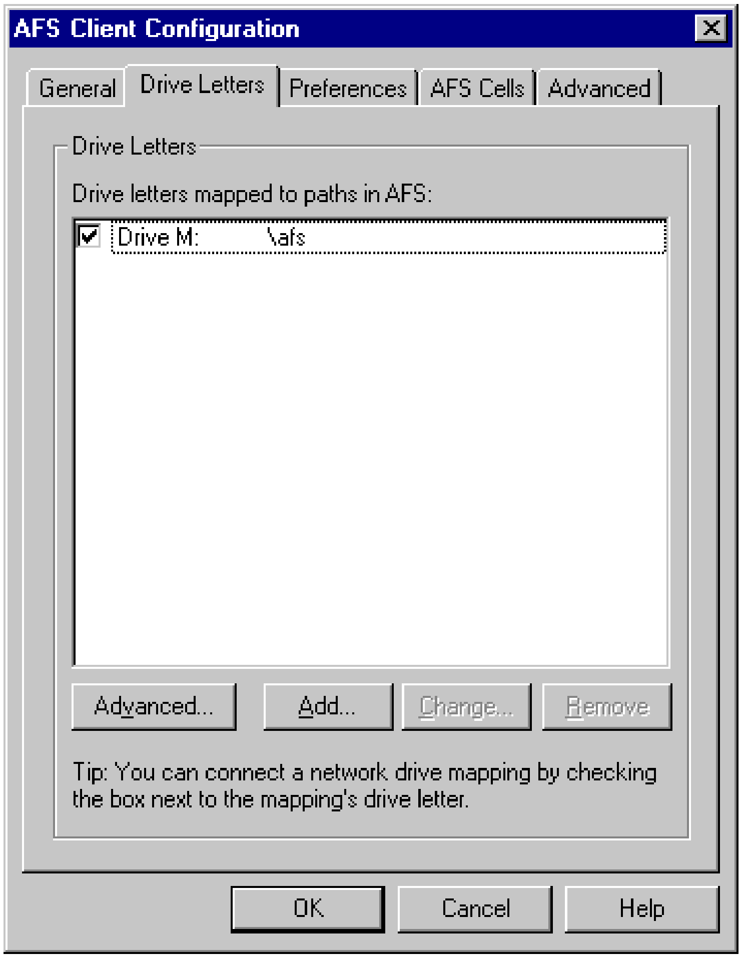

Drive letters mapped to paths in AFS:

☑ Drive M: \afs

Tip: You can connect a network drive mapping by checking the box next to the mapping's drive letter.

Figure 174. Displaying the Drive Letters Mapped to AFS Paths

- The AFS Client program

  You can invoke the AFS Client program by using one of the following methods:

  - Click the **Start** button on your Windows NT desktop taskbar, point to **Programs**, point to **IBM WebSphere**, point to **Performance Pack**, point to **Afs**, point to **Client**, and click **Authentication**, as depicted in Figure 163 on page 207.
  - Alternately, if you select the **Show the AFS Client icon in the taskbar** check box in the General tab of the AFS Client Configuration window, there is a lock icon on the Windows NT desktop taskbar next to the clock, identifying the AFS Client program, as shown in Figure 163 on page 207. Click on this lock icon to start the AFS Client program.

  In the AFS Client window, select the **Advanced** tab to display the status of AFS Client service, as shown in Figure 175 on page 215.

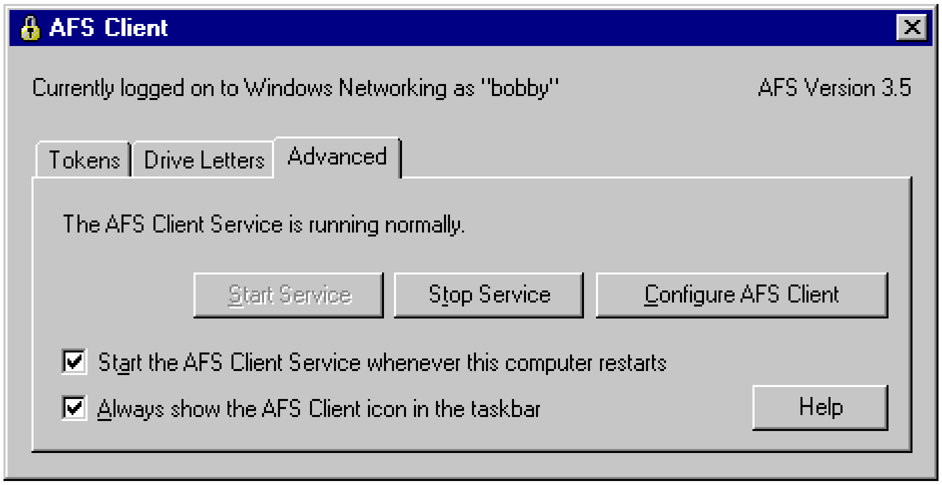

Currently logged on to Windows Networking as "bobby"   AFS Version 3.5

The AFS Client Service is running normally.

☑ Start the AFS Client Service whenever this computer restarts

☑ Always show the AFS Client icon in the taskbar

Figure 175. Displaying the AFS Client Status with the AFS Client Program

In the AFS Client window, select the Drive Letters tab to display the drive letters mapped to AFS paths, as shown in the following window:

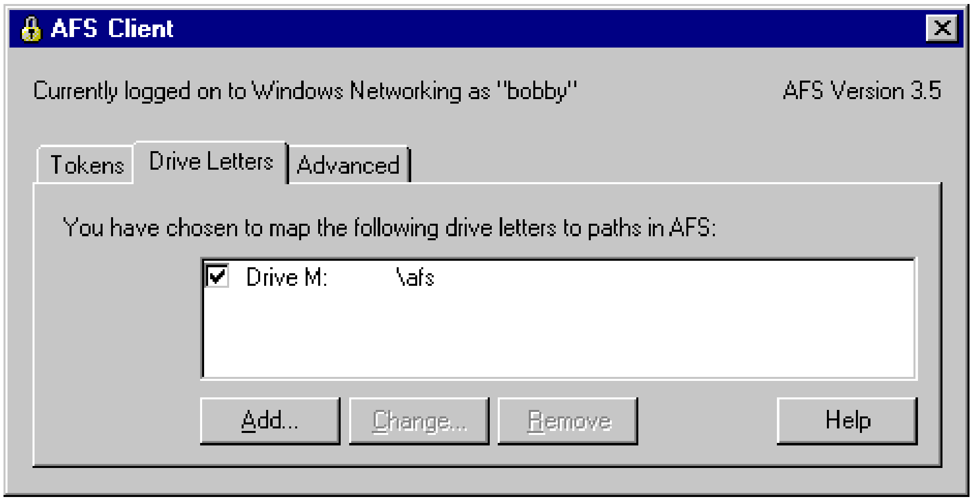

Currently logged on to Windows Networking as "bobby"   AFS Version 3.5

You have chosen to map the following drive letters to paths in AFS:

☑ Drive M: \afs

Figure 176. Displaying the Drive Letters Mapped to AFS Paths

3. List the running AFS processes from the Windows NT Task Manager window. To invoke this window, position the mouse pointer on the Windows NT desktop taskbar, click the right mouse button, and click **Task Manager...**. The Windows NT Task Manager window appears. Select the **Processes** tab to show the current processes. The AFS client process is called afsd_service.exe. Figure 177 on page 216 shows this process running correctly on our Windows NT AFS client system.

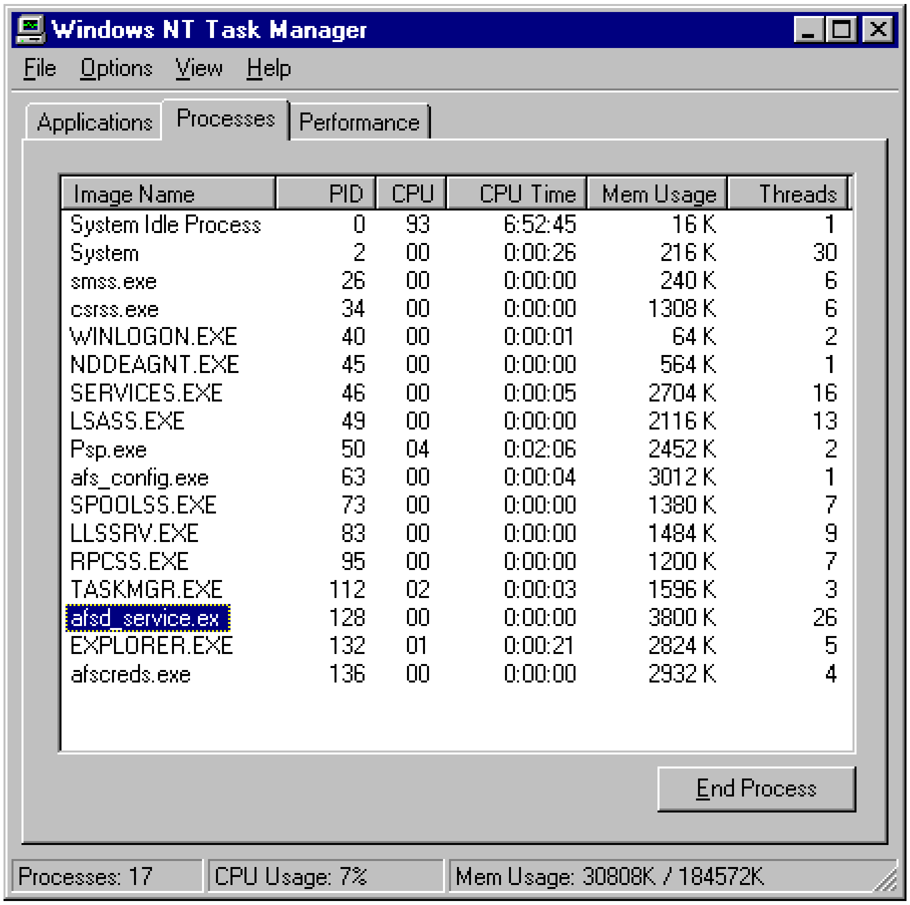

Image Name	PID	CPU	CPU Time	Mem Usage	Threads
System Idle Process	0	93	6:52:45	16 K	1
System	2	00	0:00:26	216 K	30
smss.exe	26	00	0:00:00	240 K	6
csrss.exe	34	00	0:00:00	1308 K	6
WINLOGON.EXE	40	00	0:00:01	64 K	2
NDDEAGNT.EXE	45	00	0:00:00	564 K	1
SERVICES.EXE	46	00	0:00:05	2704 K	16
LSASS.EXE	49	00	0:00:00	2116 K	13
Psp.exe	50	04	0:02:06	2452 K	2
afs_config.exe	63	00	0:00:04	3012 K	1
SPOOLSS.EXE	73	00	0:00:00	1380 K	7
LLSSRV.EXE	83	00	0:00:00	1484 K	9
RPCSS.EXE	95	00	0:00:00	1200 K	7
TASKMGR.EXE	112	02	0:00:03	1596 K	3
afsd_service.ex	128	00	0:00:00	3800 K	26
EXPLORER.EXE	132	01	0:00:21	2824 K	5
afscreds.exe	136	00	0:00:00	2932 K	4

Processes: 17 | CPU Usage: 7% | Mem Usage: 30808K / 184572K

Figure 177. Listing of AFS Client Process

To display the Threads column in the Process tab of the Task Manager window, click on the **View** button of the menu bar, and click on **Select Columns...**. In the Select Columns window, check the **Thread Count** box and click **OK** to continue. The number of threads per process information is useful in AFS tuning.

4. Authenticate to the AFS admin account of the default AFS cell by using the AFS Client program. To invoke this program, click the **Start** button on your Windows NT desktop taskbar, point to **Programs**, point to **IBM WebSphere**, point to **Performance Pack**, point to **Afs**, point to **Client**, and click **Authentication**, as depicted in Figure 163 on page 207.

   The AFS Client program displays the Tokens tab, as shown in Figure 178 on page 217.

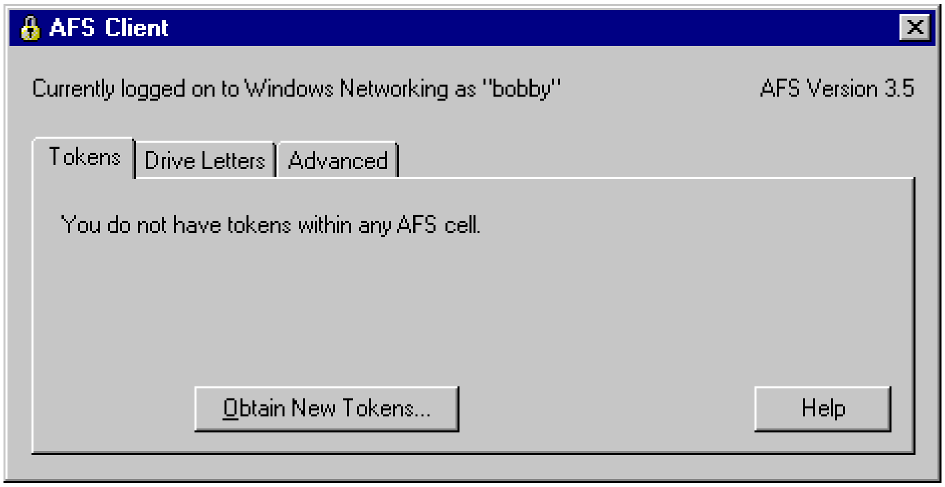
AFS Client

Currently logged on to Windows Networking as "bobby"

AFS Version 3.5

You do not have tokens within any AFS cell.

Figure 178. The AFS Client Program

5. Click the **Obtain New Tokens...** button.
6. In the Obtain New AFS Tokens window, enter the AFS admin user name and password, as shown in Figure 179 on page 217. Then, click **OK** to authenticate yourself as this AFS user.

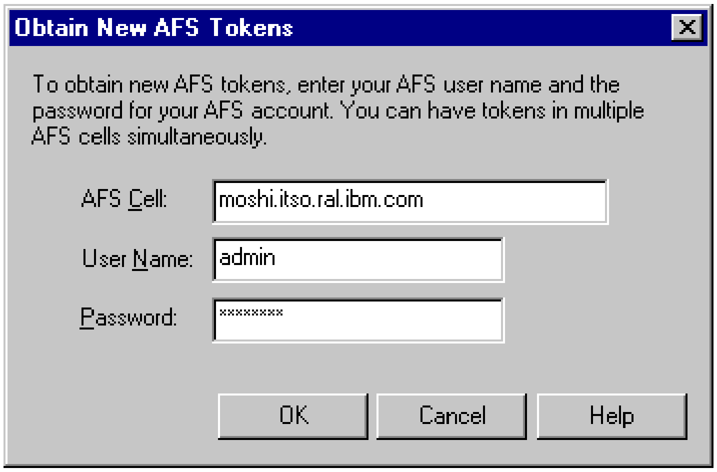
Obtain New AFS Tokens

To obtain new AFS tokens, enter your AFS user name and the password for your AFS account. You can have tokens in multiple AFS cells simultaneously.

AFS Cell: moshi.itso.ral.ibm.com

User Name: admin

Password: *********

Figure 179. Obtaining New AFS Tokens

The AFS Client program displays the AFS tokens as shown in Figure 180 on page 218.

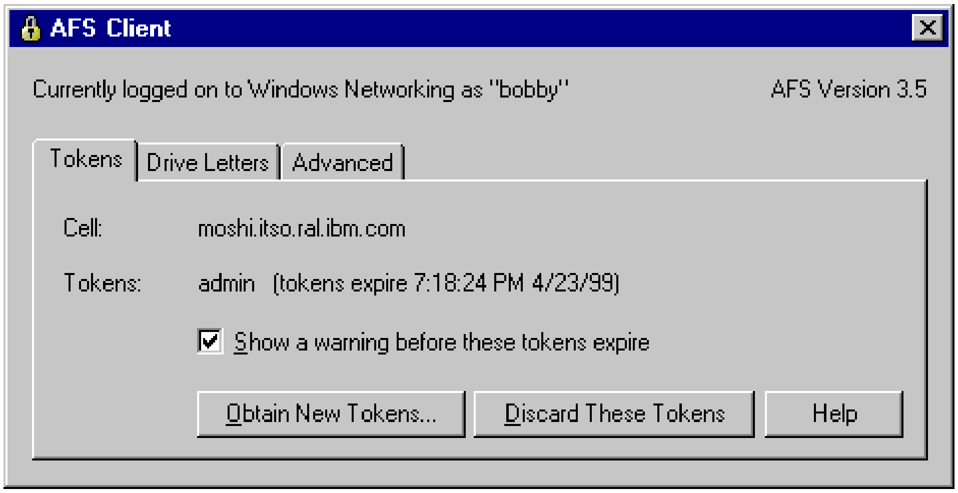

AFS Client

Currently logged on to Windows Networking as "bobby"   AFS Version 3.5

Cell: moshi.itso.ral.ibm.com

Tokens: admin (tokens expire 7:18:24 PM 4/23/99)

☑ Show a warning before these tokens expire

Figure 180. Displaying AFS Tokens

5.7 Configuring the AFS Control Center on Windows NT Systems

Before you can configure a Windows NT machine to be an AFS Control Center to administer existing AFS cells, you must install the necessary AFS files from the WebSphere Performance Pack CD-ROM distribution to the machine's local disk, as shown in 5.3, "Installing AFS Files on Windows NT Systems" on page 141.

The WebSphere Performance Pack Setup program creates a program group called IBM WebSphere, containing Performance Pack Afs, as shown in Figure 181 on page 218. From this program group, you can access the Control Center.

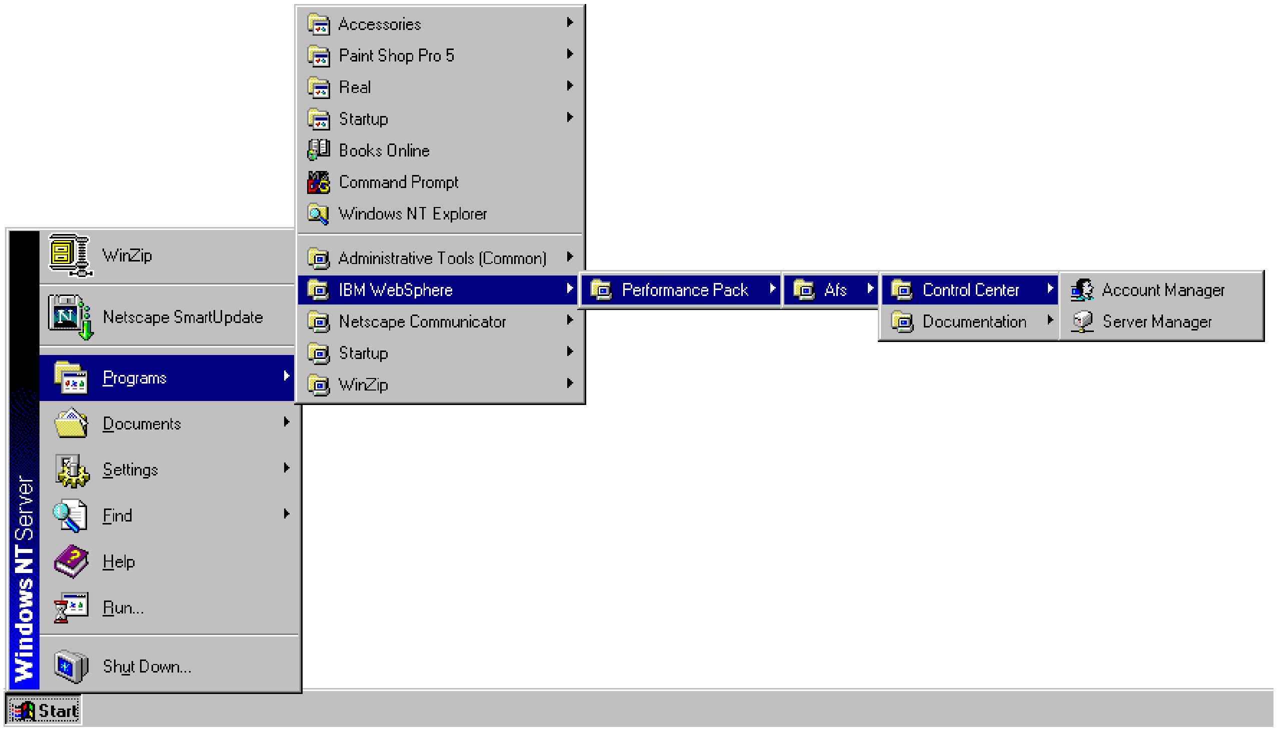

Figure 181. The WebSphere Performance Pack AFS Control Center

The Setup program also creates an AFS Control Center application in the Windows NT Control Panel. There is no AFS entry created in the Services dialog of the Windows NT Control Panel.

To illustrate the configuration of the AFS Control Center on a Windows NT system, we use an AFS cell called moshi.itso.ral.ibm.com. This cell has a single AFS Database Server machine, called wtr05319.itso.ral.ibm.com, with the IP address of 9.24.104.206. This AFS Database Server machine is also the first AFS machine.

To start the configuration of the AFS Control Center on your Windows NT system, perform the following tasks:

1. Log in as an administrative user, who is a member of the local Administrators group on your Windows NT system.
2. To start the AFS Control Center configuration program, click the **Start** button on the Windows NT desktop taskbar, point to **Settings**, and click **Control Panel**. From the Windows NT Control Panel, double-click the **AFS Control Center** icon.
3. In the AFS Control Center Properties window, you can specify the AFS cell to administer.

   Initially, the cell name is unknown, as shown in the following figure:

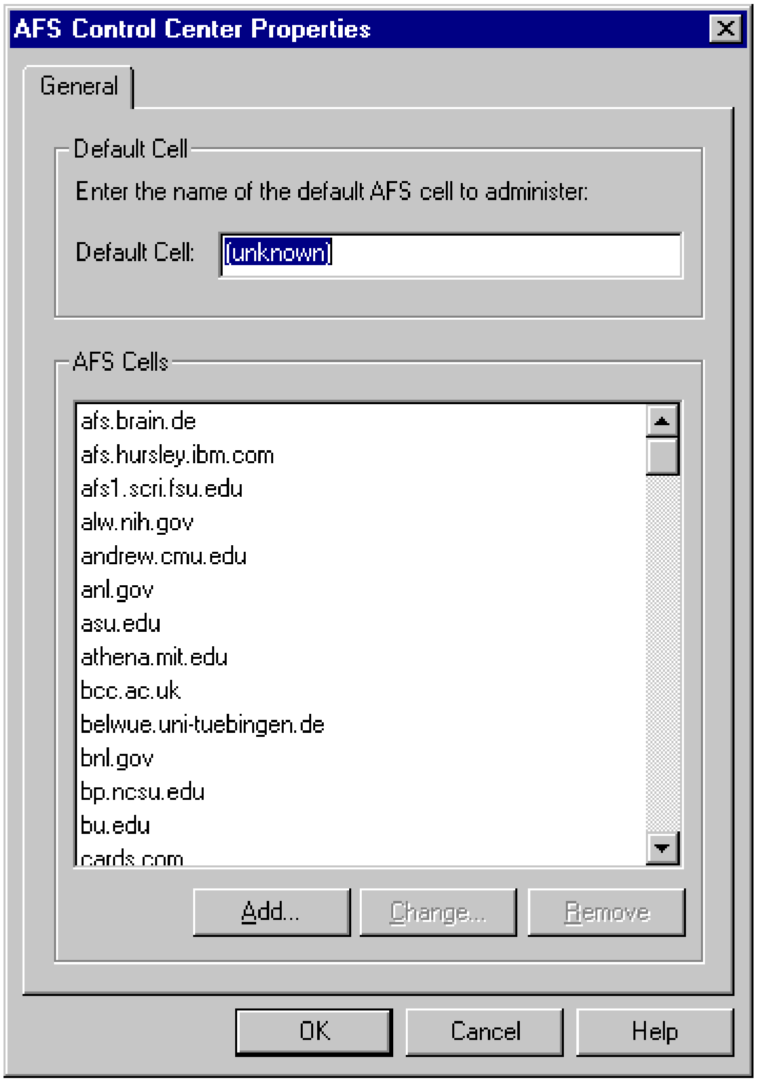

Enter the name of the default AFS cell to administer:

Default Cell: (unknown)

afs.brain.de
afs.hursley.ibm.com
afs1.scri.fsu.edu
alw.nih.gov
andrew.cmu.edu
anl.gov
asu.edu
athena.mit.edu
bcc.ac.uk
belwue.uni-tuebingen.de
bnl.gov
bp.ncsu.edu
bu.edu
cards.com

Figure 182. The AFS Control Center Properties Window

In the Default Cell frame of the AFS Control Center Properties window, enter the name of the existing AFS cell that you want to administer. In this example, we enter `moshi.itso.ral.ibm.com`, which is the name of our AFS cell. This is shown in the following figure:

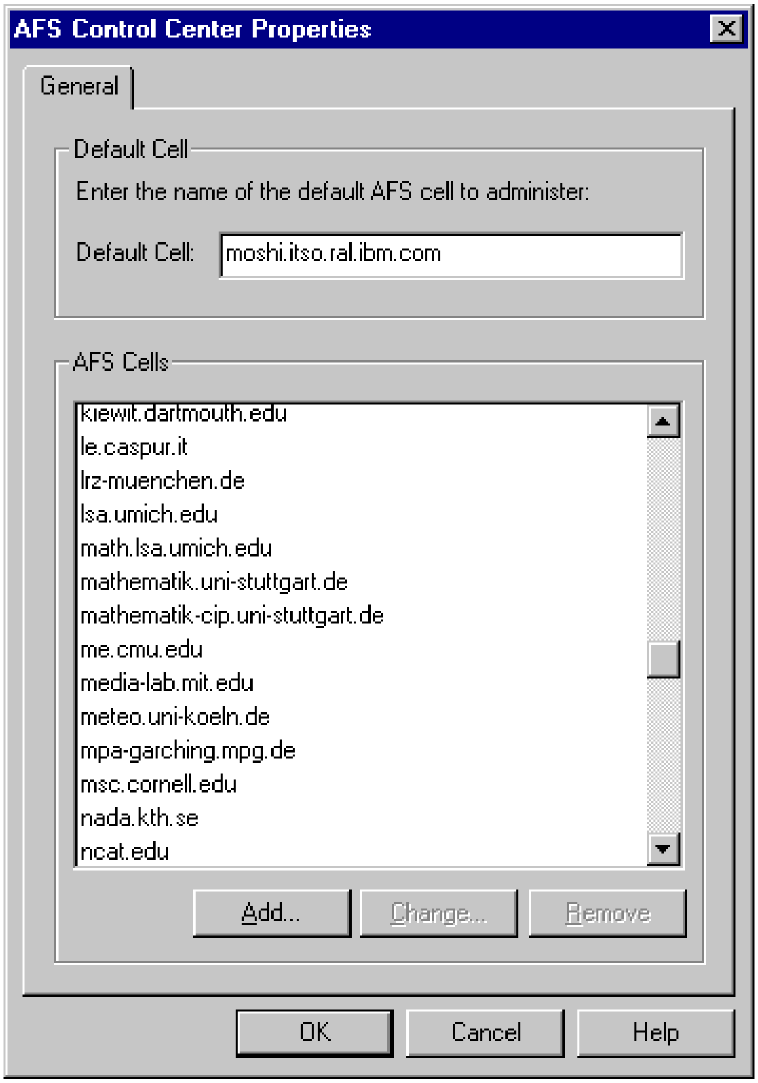

Default Cell

Enter the name of the default AFS cell to administer:

Default Cell: moshi.itso.ral.ibm.com

AFS Cells

kiewit.dartmouth.edu
le.caspur.it
lrz-muenchen.de
lsa.umich.edu
math.lsa.umich.edu
mathematik.uni-stuttgart.de
mathematik-cip.uni-stuttgart.de
me.cmu.edu
media-lab.mit.edu
meteo.uni-koeln.de
mpa-garching.mpg.de
msc.cornell.edu
nada.kth.se
ncat.edu

Figure 183. Specifying the Default AFS Cell Name

4. Next, you need to locate the specified AFS cell name, moshi.itso.ral.ibm.com, in the list of published AFS cells. The AFS Control Center configuration program displays the list in alphabetical order. If you cannot find your AFS cell on this list, your AFS cell is not yet defined. Click **Add...** to define your AFS cell.

5. In the New Cell window, you need to provide the name and a description of the new AFS cell. For example, in the AFS Cell field, enter `moshi.itso.ral.ibm.com` for the new AFS cell name, and in the Description field, enter `IBM ITSO WSPP` as a brief description of this cell. This is shown in the following figure:

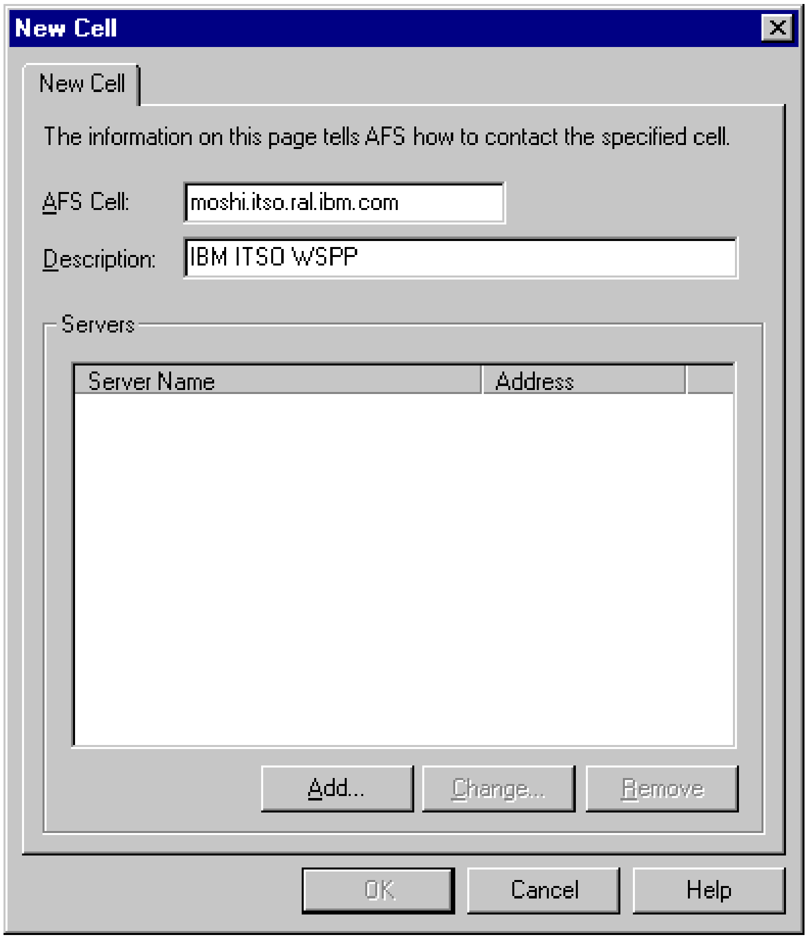

The information on this page tells AFS how to contact the specified cell.

AFS Cell: moshi.itso.ral.ibm.com

Description: IBM ITSO WSPP

Servers

Server Name	Address

Figure 184. Defining a New AFS Cell

6. Next, you need to provide the AFS Database Server information for your new AFS cell. In the New Cell window, click **Add...** to enter the host name of an AFS Database Server of the specified AFS cell.
7. In the Add Server window, the AFS Control Center configuration program prompts you for the name of an AFS Volume Location Server in the specified cell, as shown in the following figure:

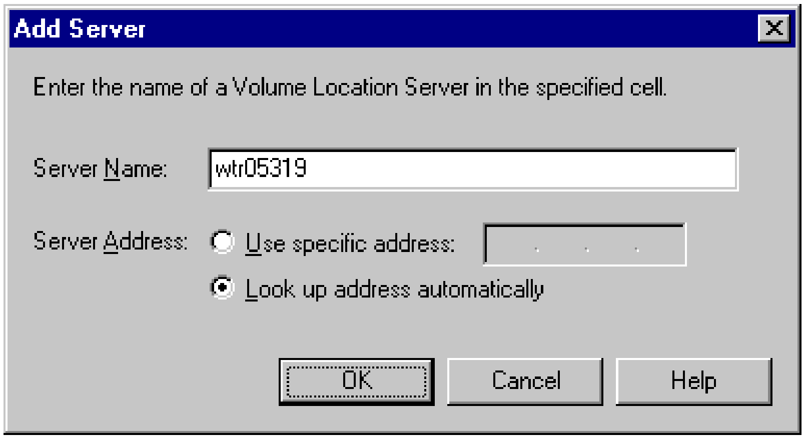

Enter the name of a Volume Location Server in the specified cell.

Server Name: wtr05319

Server Address: ○ Use specific address: . . .
● Look up address automatically

Figure 185. Defining an AFS Database Server

An AFS Volume Location Server is a machine that stores and maintains in its physical location an AFS Volume Location Database (VLDB), which is one of the four AFS databases (see 2.5.1, “AFS Servers” on page 33).

In our scenario, the AFS Database Server is the first AFS machine. Therefore, in the Server Name field of the Add Server window, we enter `wtr05319`, the host name of the first AFS machine in our AFS cell, moshi.itso.ral.ibm.com.

The AFS Client configuration program provides you with the option of specifying an IP address or using the Domain Name System (DNS) to look up the IP address automatically. In our environment, all machines are registered to DNS. Therefore, we select the default option to use DNS to resolve host name to IP address.

8. The AFS Control Center configuration program returns you to the New Cell window, displaying the information about the specified AFS cell, moshi.itso.ral.ibm.com. This is shown in the following figure:

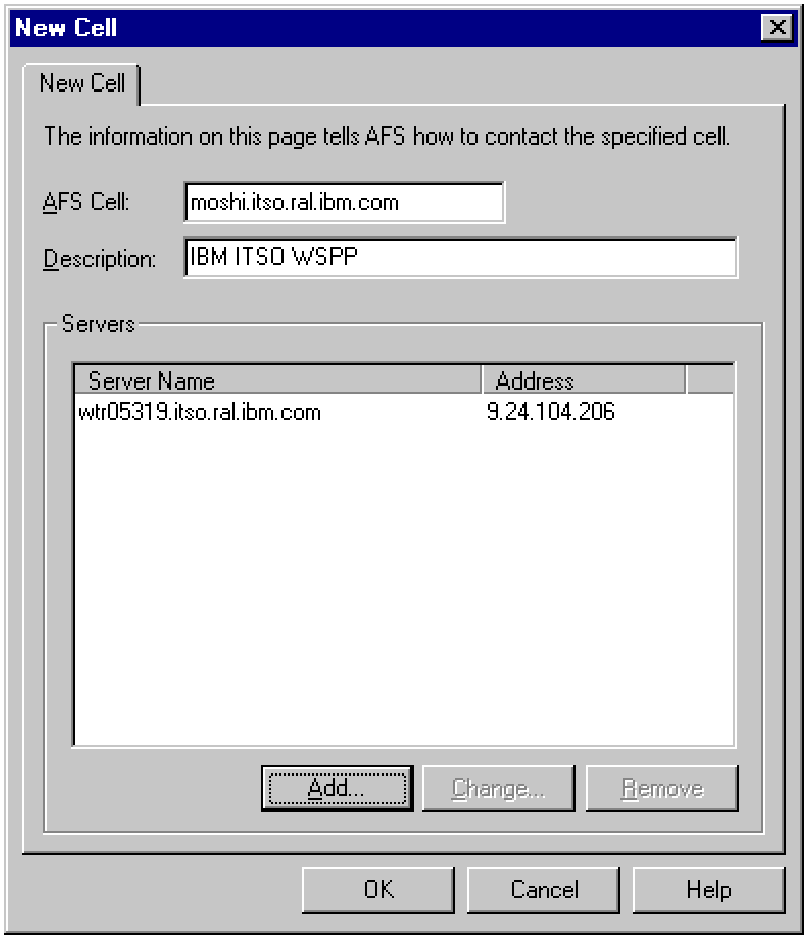

The information on this page tells AFS how to contact the specified cell.

AFS Cell: moshi.itso.ral.ibm.com

Description: IBM ITSO WSPP

Servers

Server Name	Address
wtr05319.itso.ral.ibm.com	9.24.104.206

Figure 186. Displaying the Properties of a New AFS Cell

If there is more than one AFS Database Server in your cell, click **Add...** to define other AFS Database Server machines. Otherwise, click **OK** to add the new AFS cell to the list of known AFS cells.

9. The AFS Control Center configuration program returns you to the General tab, displaying the list of AFS cells with a new entry for the AFS cell you have just entered. In this example, the new AFS cell, moshi.itso.ral.ibm.com, is included in the list of AFS cells on your Windows NT machine, as shown in the following figure:

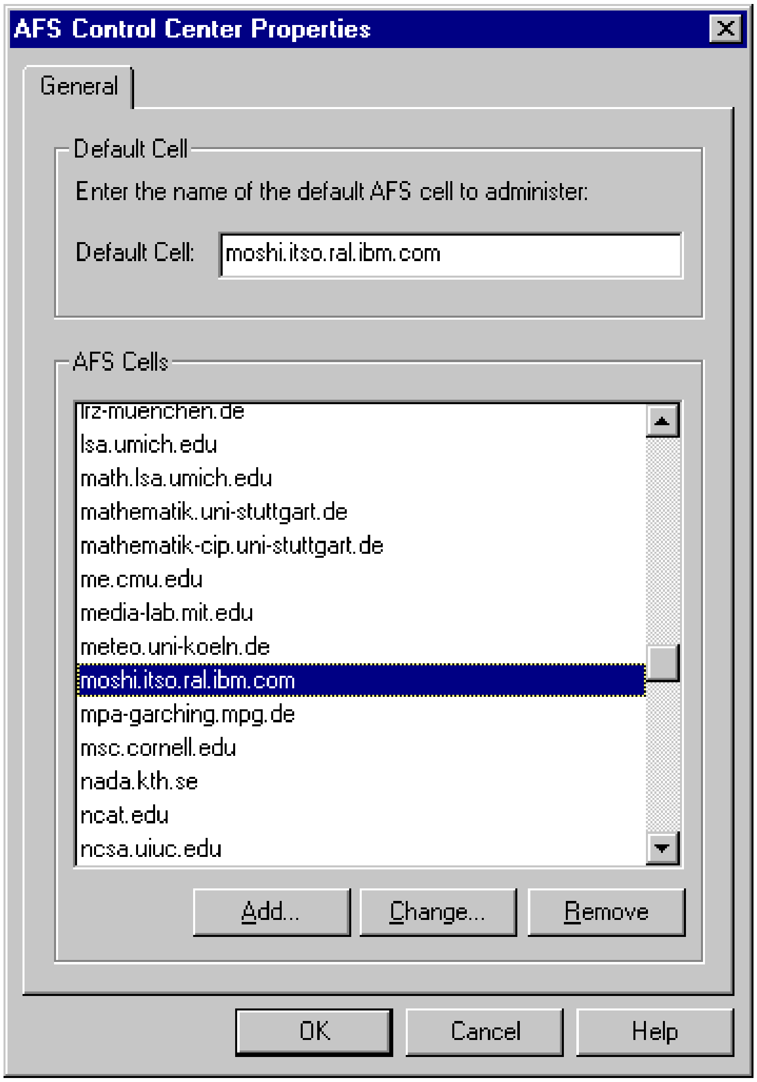

Default Cell

Enter the name of the default AFS cell to administer:

Default Cell: moshi.itso.ral.ibm.com

AFS Cells

lrz-muenchen.de
lsa.umich.edu
math.lsa.umich.edu
mathematik.uni-stuttgart.de
mathematik-cip.uni-stuttgart.de
me.cmu.edu
media-lab.mit.edu
meteo.uni-koeln.de
moshi.itso.ral.ibm.com
mpa-garching.mpg.de
msc.cornell.edu
nada.kth.se
ncat.edu
ncsa.uiuc.edu

Figure 187. Listing of the New AFS Cell

10.This process updates the list of AFS database server machines for AFS cells. By default, this list is stored in the file C:\WINNT\afsdcell.ini. Click **OK** to complete the configuration of the AFS Control Center.

5.8 Starting AFS at System Reboot on Windows NT Systems

After installing the necessary AFS files from the IBM WebSphere Performance Pack for Multiplatforms CD-ROM to the machine's local disk, as shown in 5.3, "Installing AFS Files on Windows NT Systems" on page 141, the WebSphere Performance Pack Setup program creates the necessary AFS entries in the Services Control Panel to automatically start AFS processes at system reboot.

To examine the settings of these AFS entries, click the **Start** button on the Windows NT desktop taskbar, point to **Settings**, and click **Control Panel**. From the Windows NT Control Panel, double-click the **Services** icon. Both the IBM AFS Server and Client service entries show the Automatic Startup setting, as illustrated in Figure 188 on page 226.

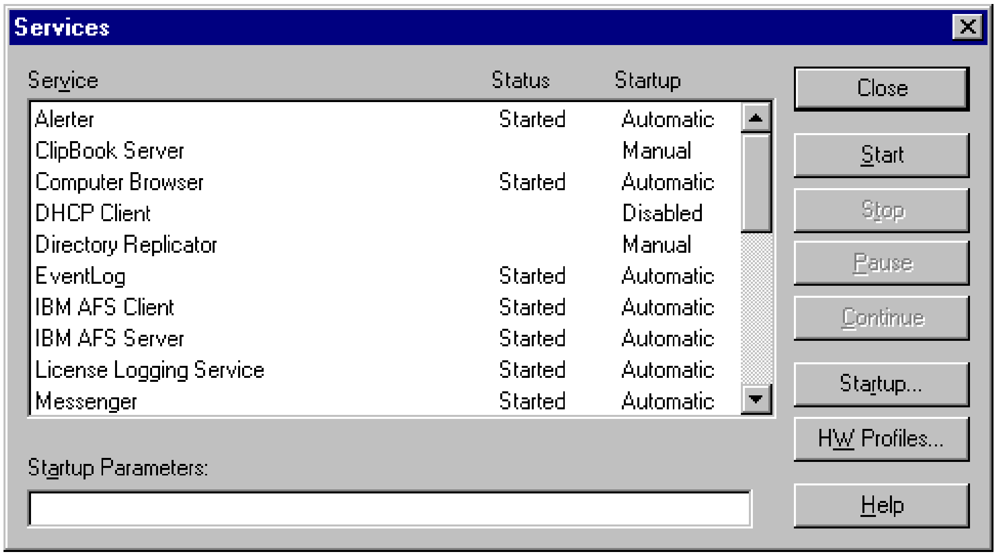

Services

Service	Status	Startup
Alerter	Started	Automatic
ClipBook Server		Manual
Computer Browser	Started	Automatic
DHCP Client		Disabled
Directory Replicator		Manual
EventLog	Started	Automatic
IBM AFS Client	Started	Automatic
IBM AFS Server	Started	Automatic
License Logging Service	Started	Automatic
Messenger	Started	Automatic

Startup Parameters:

Figure 188. The Control Panel Services Window

If you installed only AFS Client, then you see only the IBM AFS Client process in the Services panel.

After rebooting your Windows NT machine, if you log on to your machine before the AFS Client service automatically starts, the AFS Client Wizard will inform you that the AFS Client service has not started, as illustrated in Figure 189 on page 226.

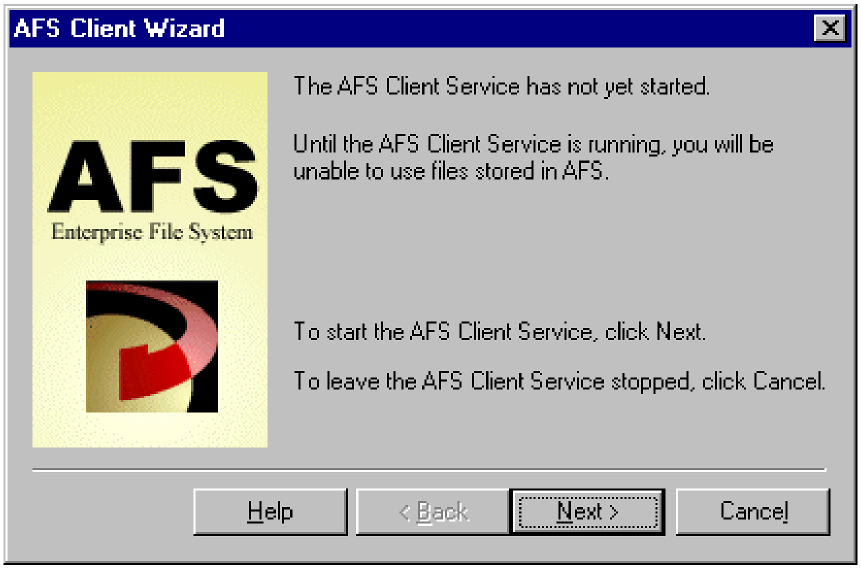

AFS Client Wizard

The AFS Client Service has not yet started.

Until the AFS Client Service is running, you will be unable to use files stored in AFS.

To start the AFS Client Service, click Next.

To leave the AFS Client Service stopped, click Cancel.

Figure 189. The AFS Client Warning Window

Click **Next** to start the AFS Client service.

The AFS Client Wizard provides you with the option of obtaining AFS tokens for a specified AFS cell. The default option is to obtain AFS tokens in the default AFS cell and the user name of the current login session. If you do not have an AFS account or you do not want to authenticate to AFS, select the **Do not obtain**

tokens for any cell check box, as shown in the following figure. Then, click **Next** to continue.

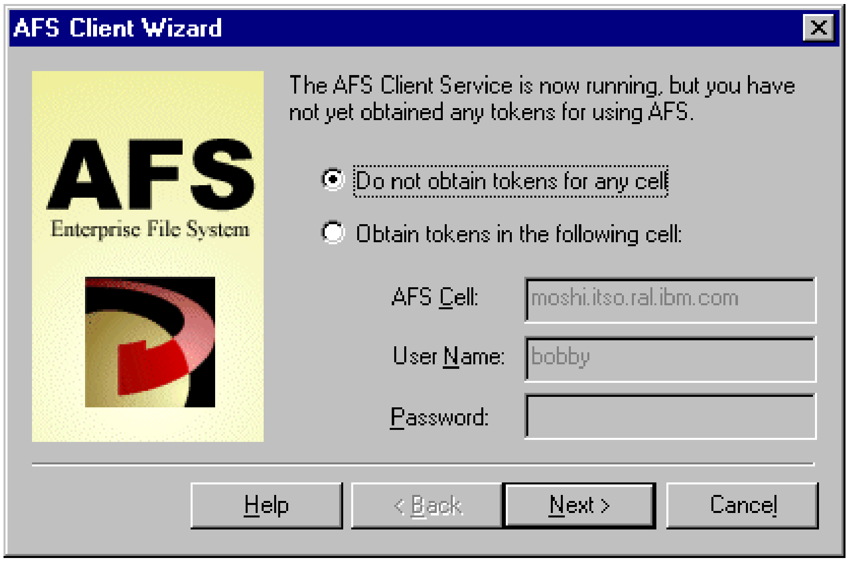

AFS Client Wizard

The AFS Client Service is now running, but you have not yet obtained any tokens for using AFS.

● Do not obtain tokens for any cell

○ Obtain tokens in the following cell:

AFS Cell: moshi.itso.ral.ibm.com

User Name: bobby

Password:

Figure 190. Obtaining AFS Tokens from the AFS Client Wizard

The AFS Client Wizard informs you that the AFS Client is ready for use, as shown in the following figure:

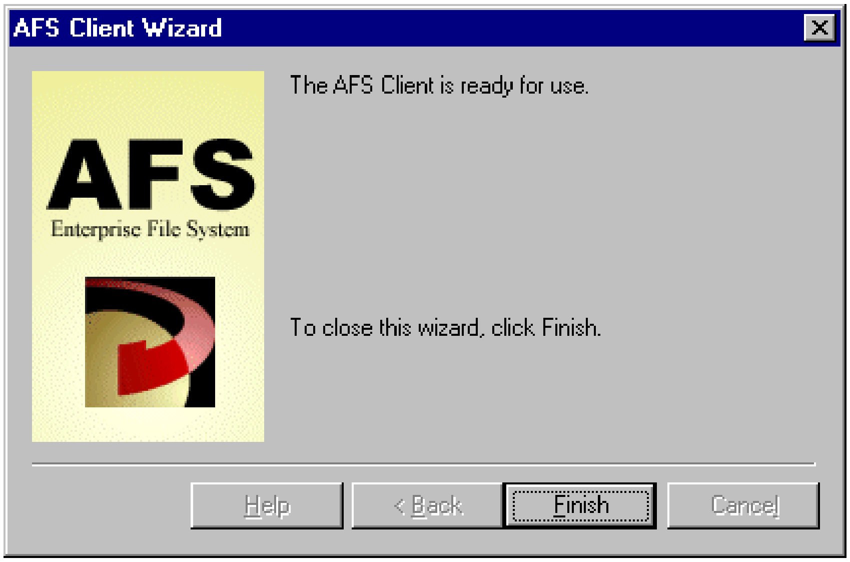

AFS Client Wizard

The AFS Client is ready for use.

To close this wizard, click Finish.

Figure 191. The AFS Client Confirmation Window

You can click **Finish** to close the AFS Client Wizard window.

Chapter 6. AFS Server Management

Administering and monitoring AFS servers will comprise the bulk of an AFS administrator's duties, including not only the file servers but the other AFS servers as well, such as the Authentication Server, the Protection Server and the Volume Location (VL) Server. For AIX, Solaris, and the Windows NT Command Prompt the `bos` command suite is used to manage server processes and the `vos` command suite is used to administer file servers and VL servers in particular. The Protection and Authentication Servers are covered in Chapter 8, “AFS User Account Management” on page 337. For Windows NT there is also the option of using the AFS Control Center Server Manager to administer AFS server processes and File or VL Servers. The Control Center provides a graphic interface to most AFS Server administration tasks.[1] The Server Manager is not limited to controlling Windows NT-based AFS servers either; in a heterogeneous environment the Server Manager can manage AFS servers on AIX and Solaris as well.

This chapter describes the usage of the AFS Control Center Server Manager as well as examples of using the command line equivalents. Information covered in this chapter includes managing and examining server processes, managing partitions and volumes, and ports used by AFS. To illustrate the usage of the AFS Server Manager, we use an AFS cell called moshi.itso.ral.ibm.com. This cell has a single AFS Volume Location (VL) server machine called wtr05319.itso.ral.ibm.com, with an IP address of 9.24.104.206. In WebSphere Performance Pack, this AFS Database Server machine is also called the *first* AFS machine, and it provides AFS file serving as well.

In addition to the first AFS machine, the AFS cell moshi.itso.ral.ibm.com contains a second AFS file server, wtr05320.itso.ral.ibm.com with an IP address of 9.24.104.213.

6.1 Starting the AFS Server Manager

To start the AFS Server Manager component of the AFS Control Center on a Windows NT system, click the **Start** button on your Windows NT desktop taskbar, point to **Programs**, point to **IBM WebSphere**, point to **Performance Pack**, point to **Afs**, point to **Control Center**, and click **Server Manager**, as shown in the following figure:

[1] The AFS Control Center has another component – the Account Manager – which is described in Chapter 8, “AFS User Account Management” on page 337.

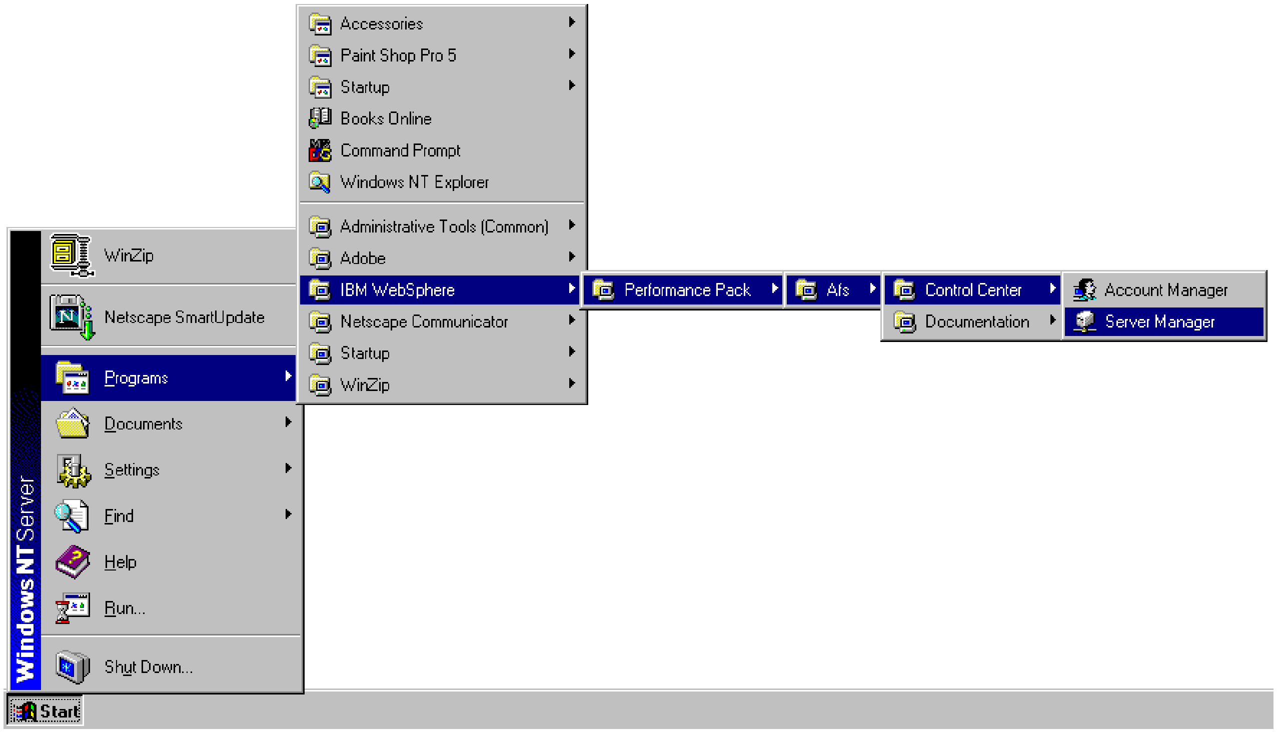

Figure 192. The AFS Control Center

The AFS Server Manager displays the Select Cell panel, as shown in the following figure:

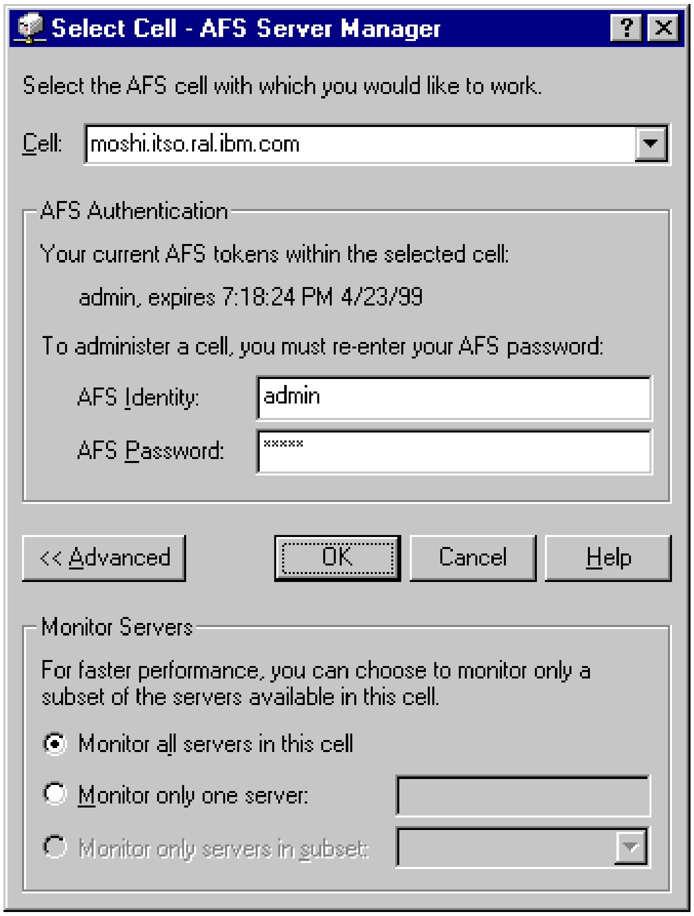
Select Cell - AFS Server Manager

Select the AFS cell with which you would like to work.

Cell: moshi.itso.ral.ibm.com

AFS Authentication

Your current AFS tokens within the selected cell:

admin, expires 7:18:24 PM 4/23/99

To administer a cell, you must re-enter your AFS password:

AFS Identity: admin

AFS Password: *****

Monitor Servers

For faster performance, you can choose to monitor only a subset of the servers available in this cell.

● Monitor all servers in this cell

○ Monitor only one server:

○ Monitor only servers in subset:

Figure 193. Authenticating to the AFS Server Manager

The AFS Server Manager provides you with the option of specifying the AFS cell that you would like to administer. Enter the name of the AFS cell. In this example, the name of the AFS cell is moshi.itso.ral.ibm.com.

You also need to authenticate as an administrator with sufficient privileges in the specified AFS cell. An AFS system administrator has three types of privilege:

- The administrator is included in the file UserList on the local disk of each AFS Server machine. On UNIX systems, this file is located in the directory /usr/afs/etc, and on Windows NT systems, the default location of this file is the directory C:\WSPPAFS\Server\usr\afs\etc. The administrator user can issue AFS `bos`, `vos`, and `backup` commands.
- The `ADMIN` flag is set in the Authentication Database entry for the administrator's account. This user can issue any AFS `kas` command.
- The administrator is a member of the AFS system:administrators group in the Protection Database. The members of this group have explicit administer and lookup rights on every Access Control List (ACL) in the cell and can issue privileged `pts` and `fs` commands.

These privileges can be separately assigned to users. By default, the AFS admin user has *all* of these privileges in order to perform administrative tasks in the AFS cell.

If you authenticate as a regular user, the AFS Server Manager will warn you with the following message:

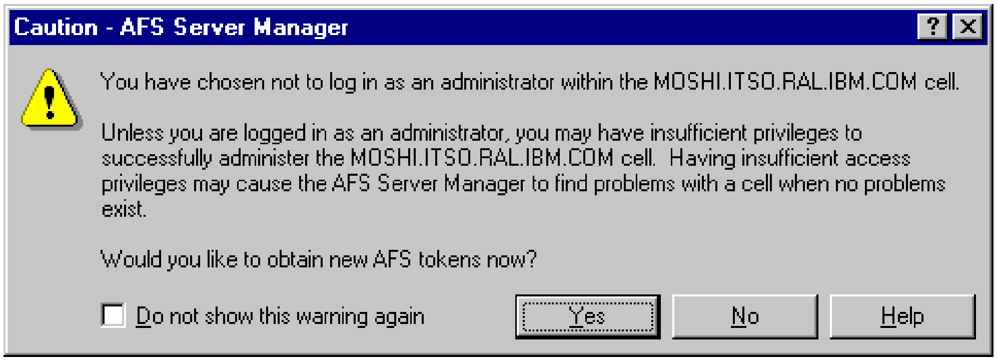

Caution - AFS Server Manager

You have chosen not to log in as an administrator within the MOSHI.ITSO.RAL.IBM.COM cell.

Unless you are logged in as an administrator, you may have insufficient privileges to successfully administer the MOSHI.ITSO.RAL.IBM.COM cell. Having insufficient access privileges may cause the AFS Server Manager to find problems with a cell when no problems exist.

Would you like to obtain new AFS tokens now?

☐ Do not show this warning again

Figure 194. Warning Message for admin Authentication

The AFS Server Manager provides you with the option of obtaining a new AFS token for an administrator. You can click **No** to continue or click **Yes** to return to the Select Cell window of the AFS Server Manager to reauthenticate as an AFS system administrator. From the Select Cell window of the AFS Server Manager, in the AFS Password field of the AFS Authentication frame, enter the password for the AFS admin account. Click the **Advanced** button to choose the AFS servers to monitor. The default option is to monitor all servers in this cell. Click **OK** to continue.

In AIX, Solaris and the Windows NT Command Prompt, authenticate as administrator using the `klog` command, and examine token information using the `tokens` command. An example of this is shown in the following screen:

```
C:\> klog admin
Password:

C:\> tokens

Tokens held by the Cache Manager:

User admin's tokens for afs@moshi.itso.ral.ibm.com [Expires Jun 10 13:42]
   --End of list --

C:\>
```

6.2 Managing AFS Server Processes

In an AFS cell, AFS Server machines can be classified as Database Servers, File Servers, System Control Machines, or Binary Distribution Machines. An AFS Server machine can be configured to provide any combination of these services. In IBM WebSphere Performance Pack, the most common AFS servers are the Database Servers and the File Servers.

A Database Server runs the following AFS server processes: the Authentication Server process (kaserver), the Protection Server process (ptserver), the Volume Location Server process (vlserver), and the Backup Server process (buserver).

A File Server runs the fs collection of AFS server processes consisting of the File Server process (fileserver), the Volume Server process (volserver), and the Salvager process (salvager).

A System Control Machine runs the server portion of the Update Server process (upserver) to distribute master copies of server configuration files to other AFS Servers running the client portion of the Update Server process (upclient). Since WebSphere Performance Pack currently offers the International Edition of AFS, data can only be transferred unencrypted. We recommend that if security is a concern you do not use the Update Server processes to synchronize these system sensitive files but instead use a secure mechanism or install them by hand.

A Binary Distribution Machine runs the server portion of the Update Server process (upserver) to distribute master copies of AFS binaries of a system type to other AFS Servers running the client portion of the Update Server process (upclient).

Each AFS server must run the Basic OverSeer (BOS) Server process (bosserver) to start and monitor the AFS server processes as defined in the BosConfig file. The BOS Server reads the BosConfig file to determine which AFS processes to start and monitor. It stores the starting state of these AFS processes in memory and does not read the BosConfig file again until its next restart. Therefore, it is possible for the memory state of the BOS Server to change independently of the BosConfig file.

6.2.1 Listing AFS Process Status

This section describes how to display the status of AFS server processes, and how to start and stop these processes, on Windows NT, AIX and Solaris. After authenticating to the AFS Server Manager, the AFS Server Manager window displays information about the selected AFS server machines in the specified AFS cell. Click the **Services** tab to list the status of all AFS server processes in the cell as shown in Figure 195 on page 234.

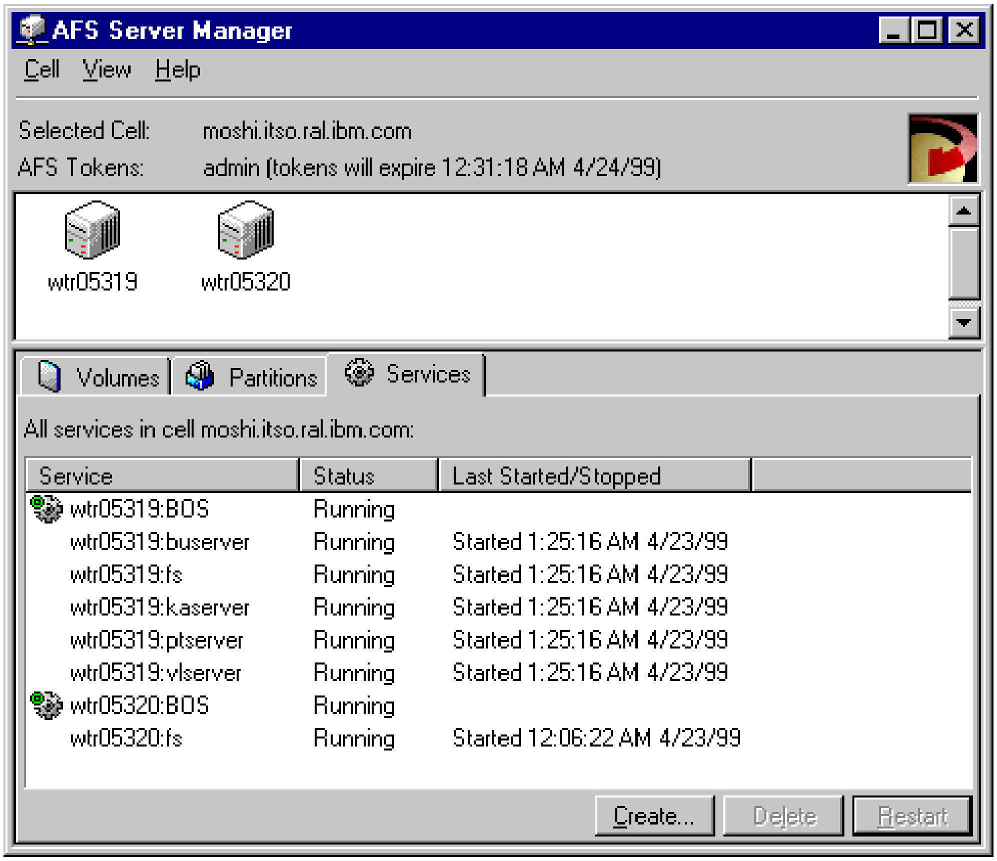

Selected Cell: moshi.itso.ral.ibm.com
AFS Tokens: admin (tokens will expire 12:31:18 AM 4/24/99)

All services in cell moshi.itso.ral.ibm.com:

Service	Status	Last Started/Stopped
wtr05319:BOS	Running	
wtr05319:buserver	Running	Started 1:25:16 AM 4/23/99
wtr05319:fs	Running	Started 1:25:16 AM 4/23/99
wtr05319:kaserver	Running	Started 1:25:16 AM 4/23/99
wtr05319:ptserver	Running	Started 1:25:16 AM 4/23/99
wtr05319:vlserver	Running	Started 1:25:16 AM 4/23/99
wtr05320:BOS	Running	
wtr05320:fs	Running	Started 12:06:22 AM 4/23/99

Figure 195. The Services Tab of the AFS Server Manager

In this example, the AFS server machine, wtr05319, is configured as a Database Server (running the kaserver, ptserver, vlserver, and buserver processes) and a File Server (running the fs trio of processes as well). The AFS server machine, wtr05320, is configured only as a File Server (running only the fs trio of processes).

To accomplish the same task in AIX, Solaris or the Windows NT Command Prompt window, use the command:

```
bos status
```

as shown in the following screen:

```
C:\> bos status -server wtr05319
Instance kaserver, currently running normally.
Instance ptserver, currently running normally.
Instance vlserver, currently running normally.
Instance buserver, currently running normally.
Instance fs, currently running normally.
    Auxiliary status is: file server running.

C:\>
```

Use the `-long` switch to see a more detailed listing of server processes. A sample use of this command is shown in the following figure:

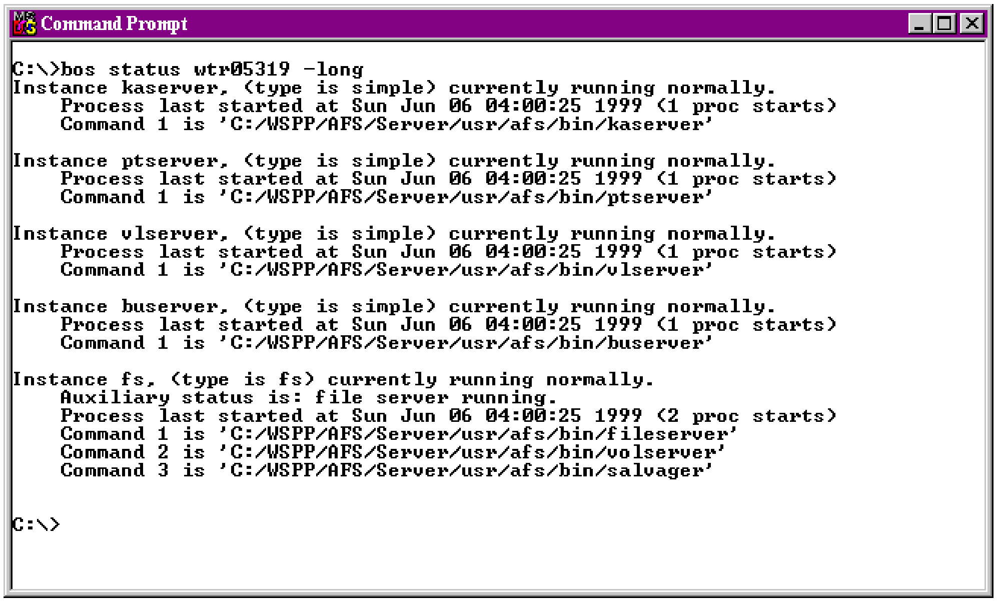

```
C:\>bos status wtr05319 -long
Instance kaserver, (type is simple) currently running normally.
    Process last started at Sun Jun 06 04:00:25 1999 (1 proc starts)
    Command 1 is 'C:/WSPP/AFS/Server/usr/afs/bin/kaserver'

Instance ptserver, (type is simple) currently running normally.
    Process last started at Sun Jun 06 04:00:25 1999 (1 proc starts)
    Command 1 is 'C:/WSPP/AFS/Server/usr/afs/bin/ptserver'

Instance vlserver, (type is simple) currently running normally.
    Process last started at Sun Jun 06 04:00:25 1999 (1 proc starts)
    Command 1 is 'C:/WSPP/AFS/Server/usr/afs/bin/vlserver'

Instance buserver, (type is simple) currently running normally.
    Process last started at Sun Jun 06 04:00:25 1999 (1 proc starts)
    Command 1 is 'C:/WSPP/AFS/Server/usr/afs/bin/buserver'

Instance fs, (type is fs) currently running normally.
    Auxiliary status is: file server running.
    Process last started at Sun Jun 06 04:00:25 1999 (2 proc starts)
    Command 1 is 'C:/WSPP/AFS/Server/usr/afs/bin/fileserver'
    Command 2 is 'C:/WSPP/AFS/Server/usr/afs/bin/volserver'
    Command 3 is 'C:/WSPP/AFS/Server/usr/afs/bin/salvager'


C:\>
```

Figure 196. Output from the bos status Command

6.2.2 Stopping AFS Processes

There may come a time when an AFS administrator will want to stop a server process, either permanently (because the machine will no longer be providing that service) or temporarily to perform administrative tasks. To stop an AFS server process, select the desired AFS process by highlighting it. Position the pointer on the name of the selected AFS process and click the right mouse button to display a pop-up menu. From the pop-up menu, select **Stop** to invoke the option to stop the selected AFS process. In the example shown in the following figure, we want to stop the AFS File Server process (corresponding to the fs trio of processes) on the machine wtr05319:

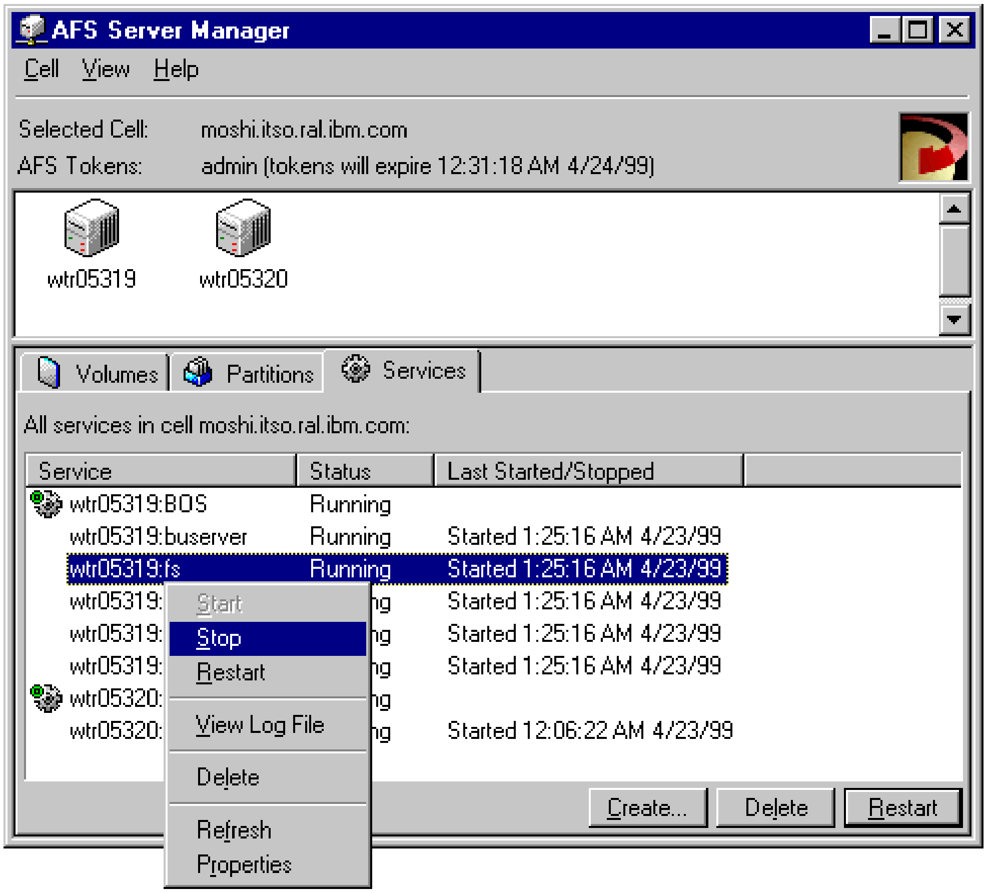

Selected Cell: moshi.itso.ral.ibm.com
AFS Tokens: admin (tokens will expire 12:31:18 AM 4/24/99)

All services in cell moshi.itso.ral.ibm.com:

Service	Status	Last Started/Stopped
wtr05319:BOS	Running	
wtr05319:buserver	Running	Started 1:25:16 AM 4/23/99
wtr05319:fs	Running	Started 1:25:16 AM 4/23/99
wtr05319:	ng	Started 1:25:16 AM 4/23/99
wtr05319:	ng	Started 1:25:16 AM 4/23/99
wtr05319:	ng	Started 1:25:16 AM 4/23/99
wtr05320:	ng	
wtr05320:	ng	Started 12:06:22 AM 4/23/99

Figure 197. Stopping an AFS Service from the AFS Server Manager

The Stop Service window appears providing you with the options of permanently or temporarily stopping the selected AFS service:

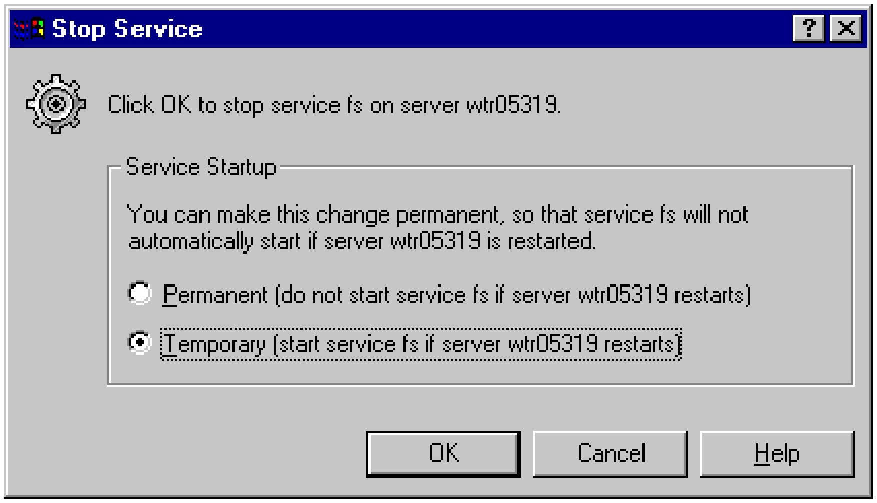

Click OK to stop service fs on server wtr05319.

Service Startup

You can make this change permanent, so that service fs will not automatically start if server wtr05319 is restarted.

○ Permanent (do not start service fs if server wtr05319 restarts)

● Temporary (start service fs if server wtr05319 restarts)

Figure 198. Stop AFS Service Dialog Box

If this machine will no longer be providing the selected AFS service, you want to permanently stop this process by selecting the **Permanent** check box. This action sets the status flag of this service to `NotRun` in both the memory state of the BOS Server and the BosConfig file. From an MS-DOS Command Prompt window or a UNIX shell prompt, you can use the AFS command:

```
bos stop
```

to permanently stop a process.

If you need to stop the selected AFS service to perform some administrative tasks, you only need to temporarily stop this process by selecting the **Temporary** check box. This action sets the status flag of this service to `NotRun` in the memory state of the BOS Server. From an MS-DOS Command Prompt window or a UNIX shell prompt, you can use the AFS command:

```
bos shutdown
```

to temporarily stop a process.

In this example, we select the **Temporary** check box to temporarily stop the AFS File Server fs trio of processes on the AFS server machine, wtr05319. Click **OK** to stop the selected AFS service.

In this example, after stopping the AFS File Server (fs) process on the AFS server machine wtr05319, the status of this service shows `Stopped` and the time this service stopped, as shown in Figure 199 on page 237.

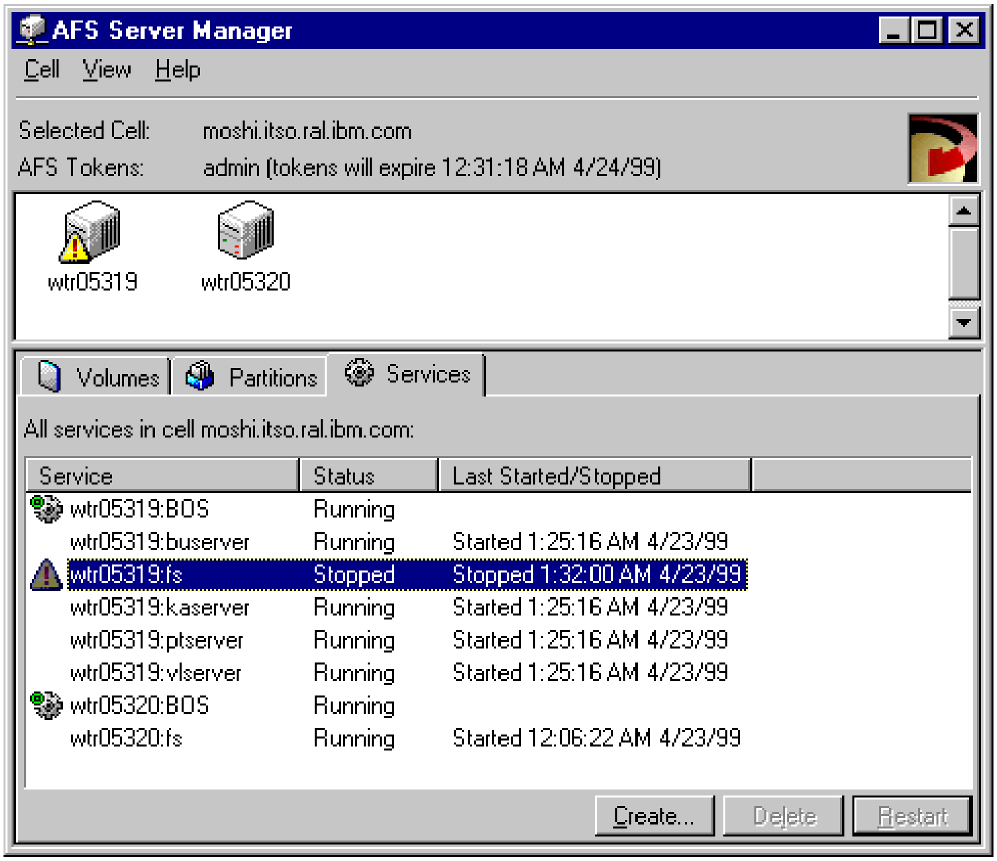

Selected Cell: moshi.itso.ral.ibm.com
AFS Tokens: admin (tokens will expire 12:31:18 AM 4/24/99)

All services in cell moshi.itso.ral.ibm.com:

Service	Status	Last Started/Stopped
wtr05319:BOS	Running	
wtr05319:buserver	Running	Started 1:25:16 AM 4/23/99
wtr05319:fs	Stopped	Stopped 1:32:00 AM 4/23/99
wtr05319:kaserver	Running	Started 1:25:16 AM 4/23/99
wtr05319:ptserver	Running	Started 1:25:16 AM 4/23/99
wtr05319:vlserver	Running	Started 1:25:16 AM 4/23/99
wtr05320:BOS	Running	
wtr05320:fs	Running	Started 12:06:22 AM 4/23/99

Figure 199. The Services Tab with Stopped AFS Service

Note that the AFS Server Manager also displays a warning icon on the AFS Server machine with the stopped process.

6.2.3 Starting AFS Processes

To start an AFS server process, select the desired AFS process by highlighting it. Position the pointer on the name of the selected AFS process and click the right mouse button to display a pop-up menu. From the pop-up menu, select **Start** to invoke the option to start the selected AFS process. In this example, we want to start the AFS File Server fs trio of processes on the machine wtr05319, as shown in the following figure:

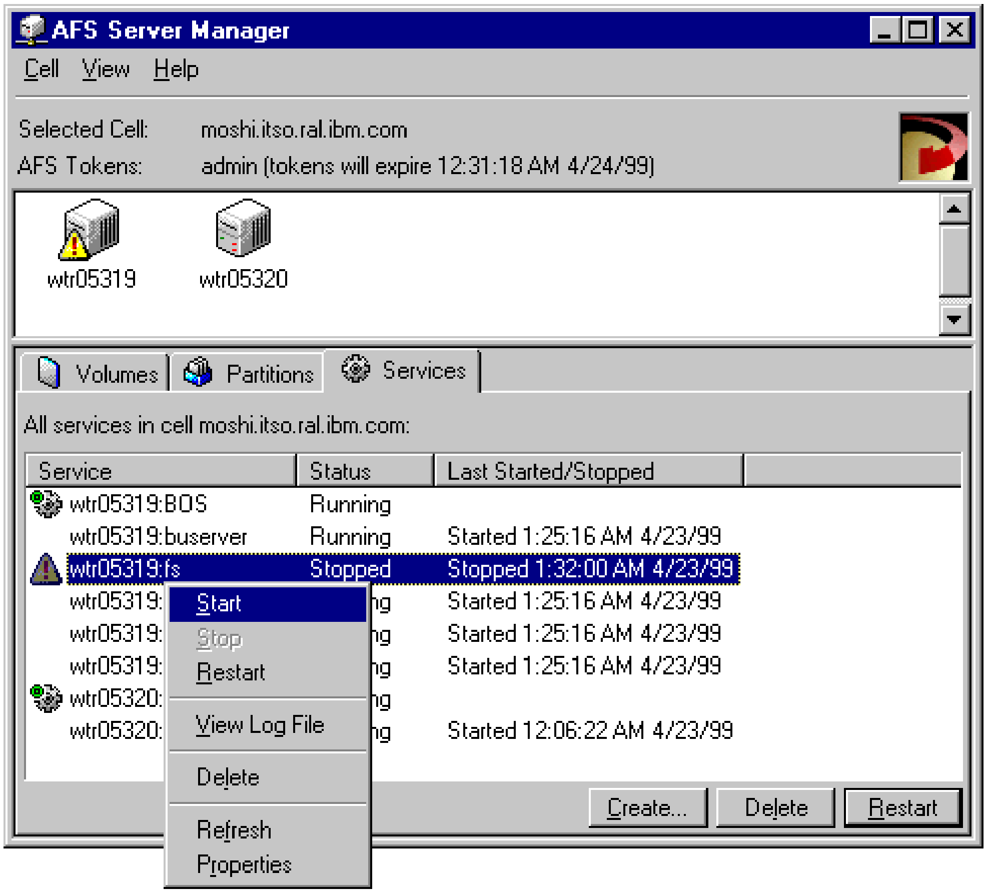

Selected Cell: moshi.itso.ral.ibm.com
AFS Tokens: admin (tokens will expire 12:31:18 AM 4/24/99)

All services in cell moshi.itso.ral.ibm.com:

Service	Status	Last Started/Stopped
wtr05319:BOS	Running	
wtr05319:buserver	Running	Started 1:25:16 AM 4/23/99
wtr05319:fs	Stopped	Stopped 1:32:00 AM 4/23/99
wtr05319:	ng	Started 1:25:16 AM 4/23/99
wtr05319:	ng	Started 1:25:16 AM 4/23/99
wtr05319:	ng	Started 1:25:16 AM 4/23/99
wtr05320:	ng	
wtr05320:	ng	Started 12:06:22 AM 4/23/99

Figure 200. Starting an AFS Service from the AFS Server Manager

The Start Service window appears providing you with the options of permanently or temporarily starting the selected AFS service:

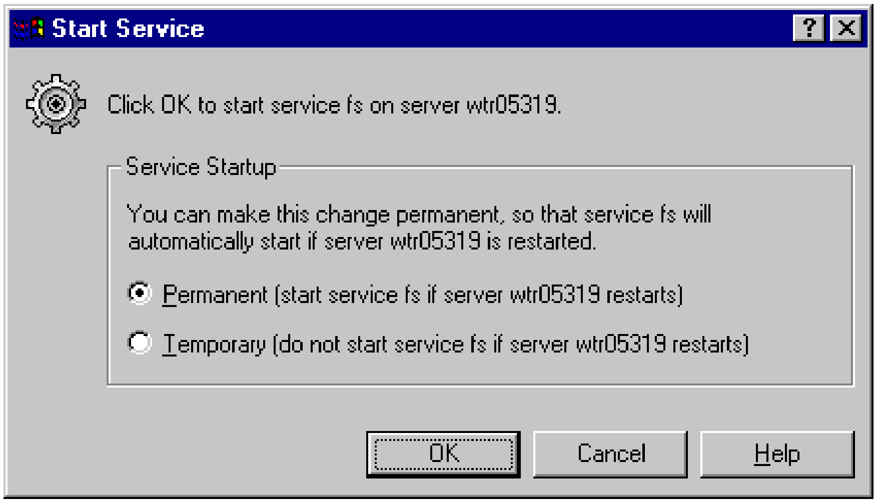

Click OK to start service fs on server wtr05319.

Service Startup

You can make this change permanent, so that service fs will automatically start if server wtr05319 is restarted.

● Permanent (start service fs if server wtr05319 restarts)

○ Temporary (do not start service fs if server wtr05319 restarts)

Figure 201. The Start AFS Service Dialog Box

If this machine will continue to provide the indicated AFS service, you want to permanently start this process by selecting the **Permanent** check box. This action sets the status flag of this service to `Run` in both the memory state of the BOS Server and the BosConfig file. From an MS-DOS Command Prompt window or a UNIX shell prompt, you can use the AFS command:

```
bos start
```

to permanently start a process.

If you need to start the selected AFS service to test run a server process, you only need to temporarily start this process by selecting the **Temporary** check box. This action sets the status flag of this service to `Run` in the memory state of the BOS Server. From an MS-DOS Command Prompt window or a UNIX shell prompt, you can use the AFS command:

```
bos startup
```

to temporarily start a process.

In this example, we select the **Permanent** check box to permanently start the stopped AFS File Server fs trio of processes, which we had previously stopped, on the AFS Server machine wtr05319. Click **OK** to start the selected AFS service.

In this example, after starting the AFS File Server (fs) process on the AFS server machine wtr05319, the status of this service shows `Running` and the time this service started, as shown in Figure 202 on page 240.

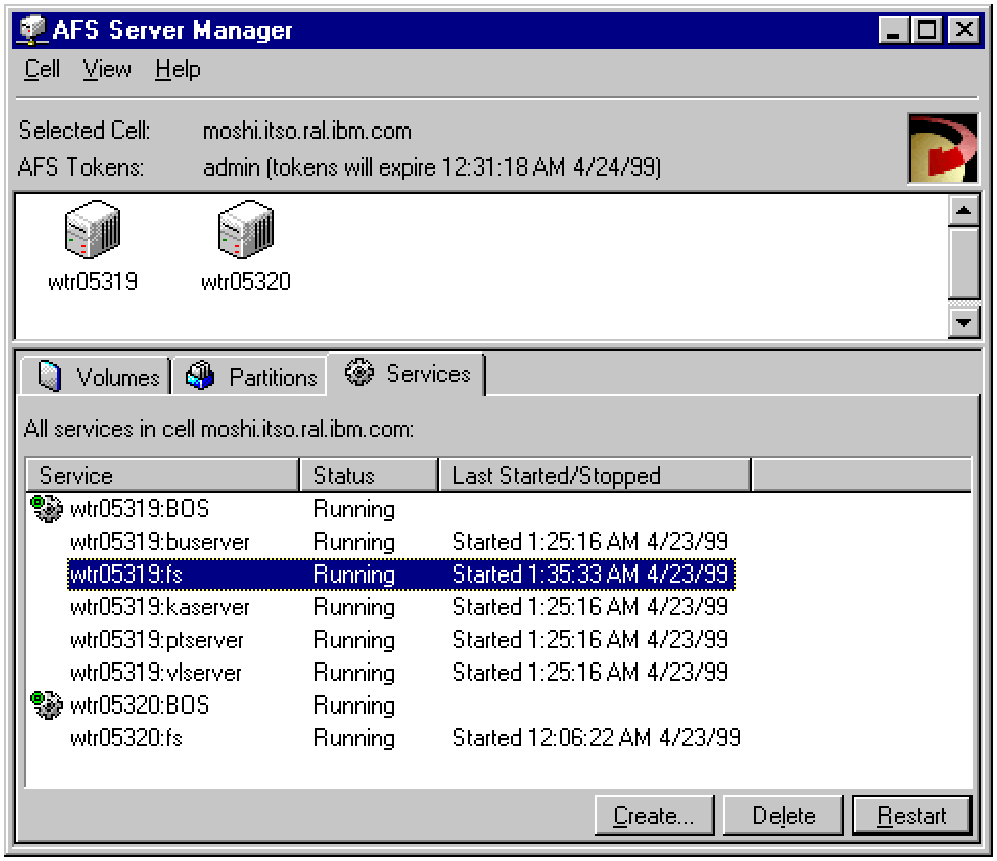

Selected Cell: moshi.itso.ral.ibm.com
AFS Tokens: admin (tokens will expire 12:31:18 AM 4/24/99)

All services in cell moshi.itso.ral.ibm.com:

Service	Status	Last Started/Stopped
wtr05319:BOS	Running	
wtr05319:buserver	Running	Started 1:25:16 AM 4/23/99
wtr05319:fs	Running	Started 1:35:33 AM 4/23/99
wtr05319:kaserver	Running	Started 1:25:16 AM 4/23/99
wtr05319:ptserver	Running	Started 1:25:16 AM 4/23/99
wtr05319:vlserver	Running	Started 1:25:16 AM 4/23/99
wtr05320:BOS	Running	
wtr05320:fs	Running	Started 12:06:22 AM 4/23/99

Figure 202. The Normal Services Tab of the AFS Server Manager

6.2.4 Examining AFS Process Information

To examine information about an AFS server process, select the desired AFS process by highlighting it. Position the pointer on the name of the selected AFS process and click the right mouse button to display a pop-up menu. From the pop-up menu, select **Properties** to invoke the option to display information about the selected AFS process. In this example, we want to show properties of the AFS File Server fs trio of processes on the machine wtr05319:

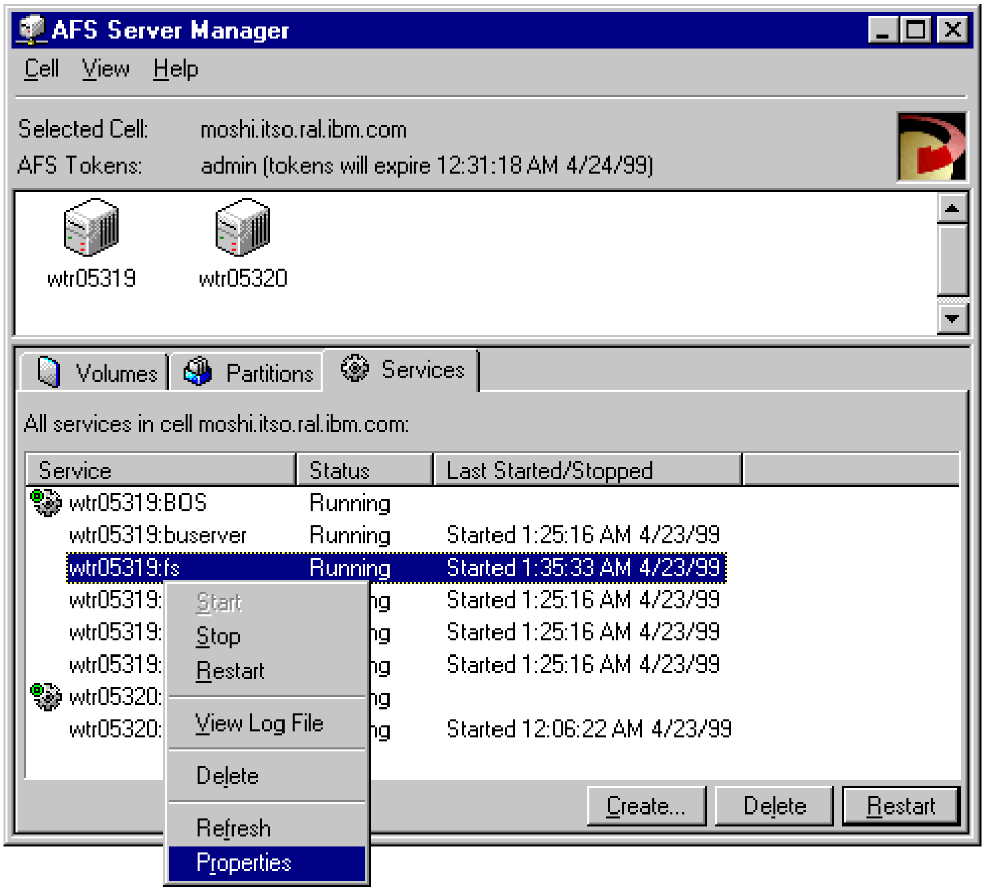

Selected Cell: moshi.itso.ral.ibm.com
AFS Tokens: admin (tokens will expire 12:31:18 AM 4/24/99)

All services in cell moshi.itso.ral.ibm.com:

Service	Status	Last Started/Stopped
wtr05319:BOS	Running	
wtr05319:buserver	Running	Started 1:25:16 AM 4/23/99
wtr05319:fs	Running	Started 1:35:33 AM 4/23/99
wtr05319:	ng	Started 1:25:16 AM 4/23/99
wtr05319:	ng	Started 1:25:16 AM 4/23/99
wtr05319:	ng	Started 1:25:16 AM 4/23/99
wtr05320:	ng	
wtr05320:	ng	Started 12:06:22 AM 4/23/99

Figure 203. Viewing the Properties of an AFS Service

The AFS Server Manager displays the Service Properties window:

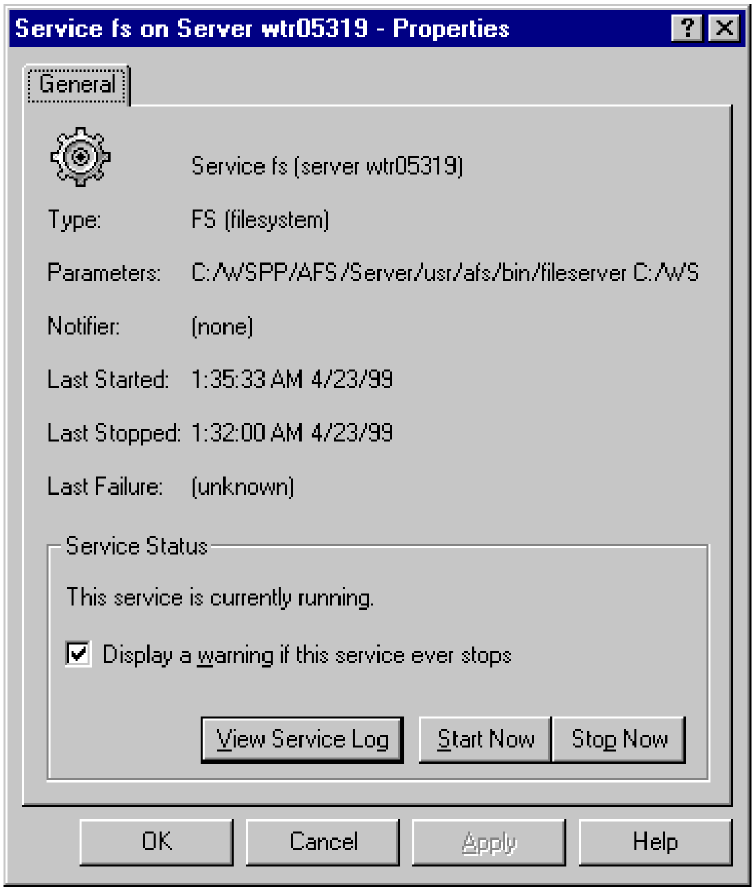

Service fs (server wtr05319)

Type: FS (filesystem)

Parameters: C:/WSPP/AFS/Server/usr/afs/bin/fileserver C:/WS

Notifier: (none)

Last Started: 1:35:33 AM 4/23/99

Last Stopped: 1:32:00 AM 4/23/99

Last Failure: (unknown)

Service Status

This service is currently running.

☑ Display a warning if this service ever stops

Figure 204. The AFS Service Properties Window

In this example, the AFS Server Manager shows information about the AFS File Server process on the machine wtr05319. Again, to get information on server processes on AIX, Solaris and the Windows NT Command Prompt use the command:

```
bos status
```

To view the log of this AFS service, click the **View Service Log** button, and a window similar to the following is displayed:

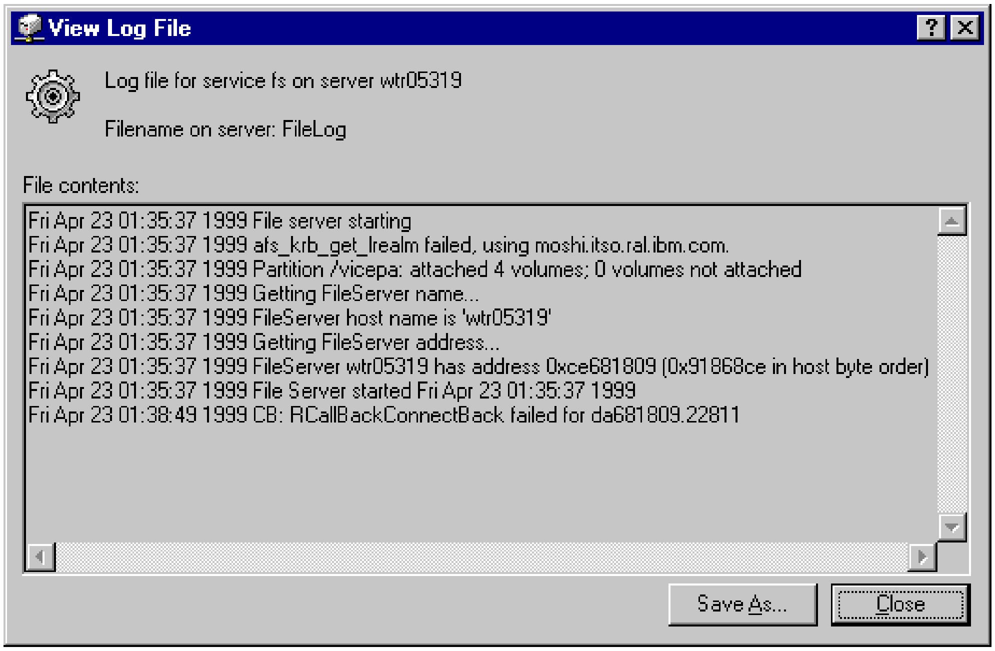

Log file for service fs on server wtr05319

Filename on server: FileLog

File contents:

```
Fri Apr 23 01:35:37 1999 File server starting
Fri Apr 23 01:35:37 1999 afs_krb_get_lrealm failed, using moshi.itso.ral.ibm.com.
Fri Apr 23 01:35:37 1999 Partition /vicepa: attached 4 volumes; 0 volumes not attached
Fri Apr 23 01:35:37 1999 Getting FileServer name...
Fri Apr 23 01:35:37 1999 FileServer host name is 'wtr05319'
Fri Apr 23 01:35:37 1999 Getting FileServer address...
Fri Apr 23 01:35:37 1999 FileServer wtr05319 has address 0xce681809 (0x91868ce in host byte order)
Fri Apr 23 01:35:37 1999 File Server started Fri Apr 23 01:35:37 1999
Fri Apr 23 01:38:49 1999 CB: RCallBackConnectBack failed for da681809.22811
```

Figure 205. The AFS Service Log

In this example, the AFS Server Manager displays the log of the AFS File Server process on the machine wtr05319. Click **Close** to return to the Service Properties window. From the Service Properties window, click **OK** to return to the main window of the AFS Server Manager.

For AIX, Solaris and the Windows NT Command Prompt use the command:

```
bos getlog
```

A sample usage is shown in the following session screen:

```
C:\> bos getlog -server wtr05319 -file FileLog
Fetching log file 'FileLog'...
Sun Jun 06 04:00:26 1999 File server starting
Sun Jun 06 04:00:26 1999 afs_krb_get_lrealm failed, using moshi.itso.ral.ibm.com
un Jun 06 04:00:27 1999 VL_RegisterAddrs rpc failed; will retry periodically (co
de=5376, err=2)
Sun Jun 06 04:00:29 1999 Partition /vicepa: attached 6 volumes; 0 volumes not at
tached
Sun Jun 06 04:00:29 1999 Partition /vicepb: attached 0 volumes; 0 volumes not at
tached
Sun Jun 06 04:00:29 1999 Getting FileServer name...
Sun Jun 06 04:00:29 1999 FileServer host name is 'wtr05319'
Sun Jun 06 04:00:29 1999 Getting FileServer address...
Sun Jun 06 04:00:29 1999 FileServer wtr05319 has address 0xce681809 (0x91868ce i
n host byte order)
Sun Jun 06 04:00:29 1999 File Server started Sun Jun 06 04:00:29 1999
```

The valid log file choices for this command are:

- AuthLog, the Authentication Server (kaserver) log file
- BackupLog, the Backup Server (buserver) log file
- BosLog, the BOS Server (bosserver) log file

- FileLog, the File Server (fileserver) log file
- SalvageLog, the Salvager (salvager) log file
- VLLog, the Volume Location (VL) Server (vlserver) log file
- VolserLog, the Volume Server (volserver) log file

6.2.5 Restarting AFS Processes

To restart an AFS server process, select the desired AFS process by highlighting it. Position the pointer on the name of the selected AFS process and click the right mouse button to display a pop-up menu. From the pop-up menu, select **Restart** to invoke the option to restart the selected AFS process. In this example, we want to restart the AFS Volume Location Server process (vlserver) on the machine wtr05319:

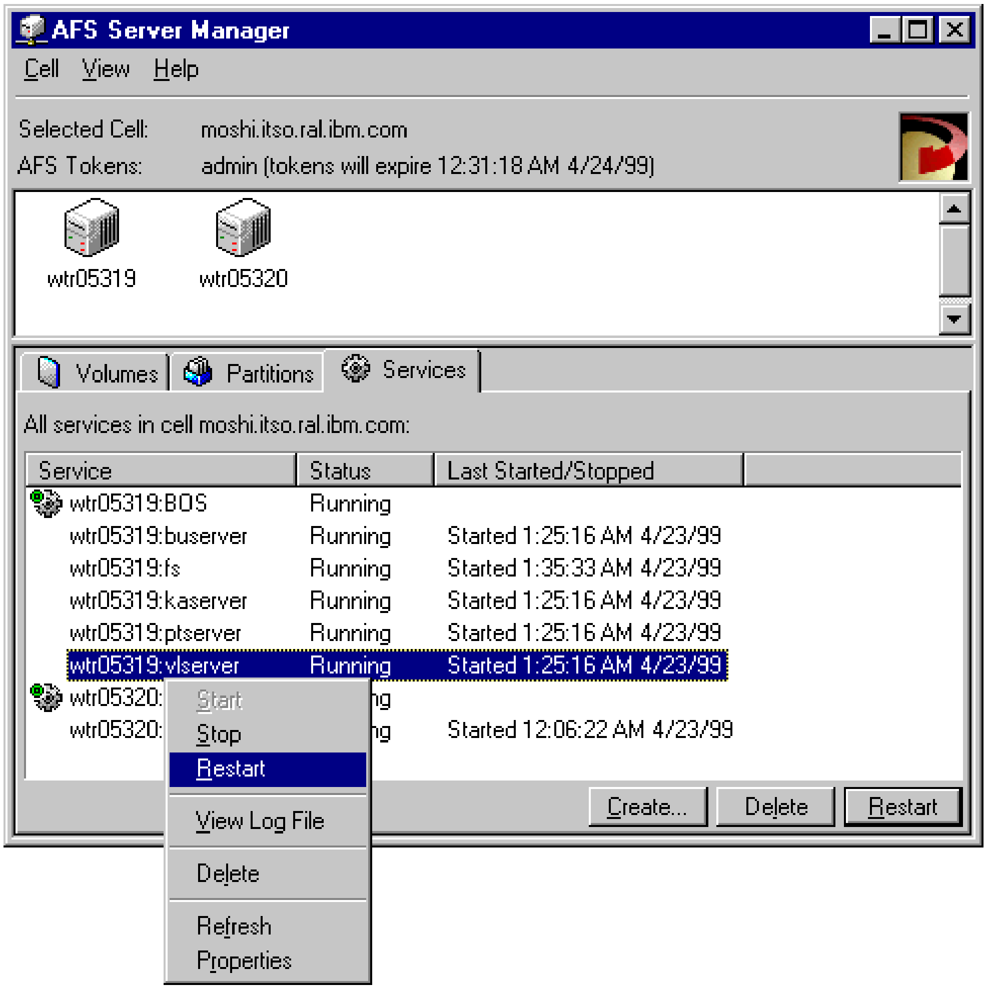

Selected Cell: moshi.itso.ral.ibm.com
AFS Tokens: admin (tokens will expire 12:31:18 AM 4/24/99)

All services in cell moshi.itso.ral.ibm.com:

Service	Status	Last Started/Stopped
wtr05319:BOS	Running	
wtr05319:buserver	Running	Started 1:25:16 AM 4/23/99
wtr05319:fs	Running	Started 1:35:33 AM 4/23/99
wtr05319:kaserver	Running	Started 1:25:16 AM 4/23/99
wtr05319:ptserver	Running	Started 1:25:16 AM 4/23/99
wtr05319:vlserver	Running	Started 1:25:16 AM 4/23/99
wtr05320:	ng	
wtr05320:	ng	Started 12:06:22 AM 4/23/99

Figure 206. Restarting an AFS Service

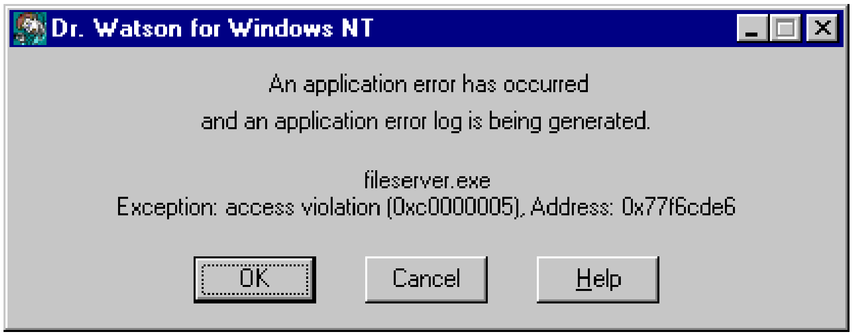

An application error has occurred
and an application error log is being generated.

fileserver.exe
Exception: access violation (0xc0000005), Address: 0x77f6cde6

Figure 207. An Error Message after Restarting an AFS Service

As of this writing selecting **Restart** generates a Dr. Watson error, which is clearly a bug in the AFS Server Manager. A dump is generated, and clicking **OK** returns control to the Server Manager. Interestingly it is always the fileserver.exe process that generates the exception error, no matter what process is being restarted. Also, this error has no *apparent* effect on the server: the processes restart, as requested, and the server is returned to a normal working state with no ill effects. This bug will show up if you installed the AFS code shipped in WebSphere performance pack V2.0, but a fix has been developed, and no problems will arise if you install AFS 3.5 Patch 1 for Windows NT (see `http://www.transarc.com`).

The equivalent command in AIX, Solaris and the Windows NT Command Prompt is the following:

```
bos restart
```

This command restarts all of the processes controlled by the BOS server. A sample usage is shown in the following session screen:

```
C:\> bos restart -server wtr05319 -all

C:\>
```

6.3 Managing AFS Partitions

Volumes containing collections of AFS files and directories are stored on the partitions of AFS File Server machines. AFS partitions are mapped to the underlying operating system disk partitions, for example /dev/hd4 on AIX, /dev/dsk/c0t1d0s3 on Solaris, and D: on Windows NT. This implies that the disk must already have been partitioned and a native operating system file system placed on the partition:

- Journaled file system (JFS) for AIX
- UNIX file system (UFS) for Solaris
- NTFS for Windows NT

6.3.1 Creating AFS Partitions on Windows NT

Before you can configure a partition to house AFS volumes, you must verify that the partition is an NTFS partition. To do this, open the **Disk Administrator** from the Start menu, under **Administrative Tools**. In this example you see there are three NTFS partitions: drives C:, D: and E:. However, drive C: cannot be formatted as an AFS partition.

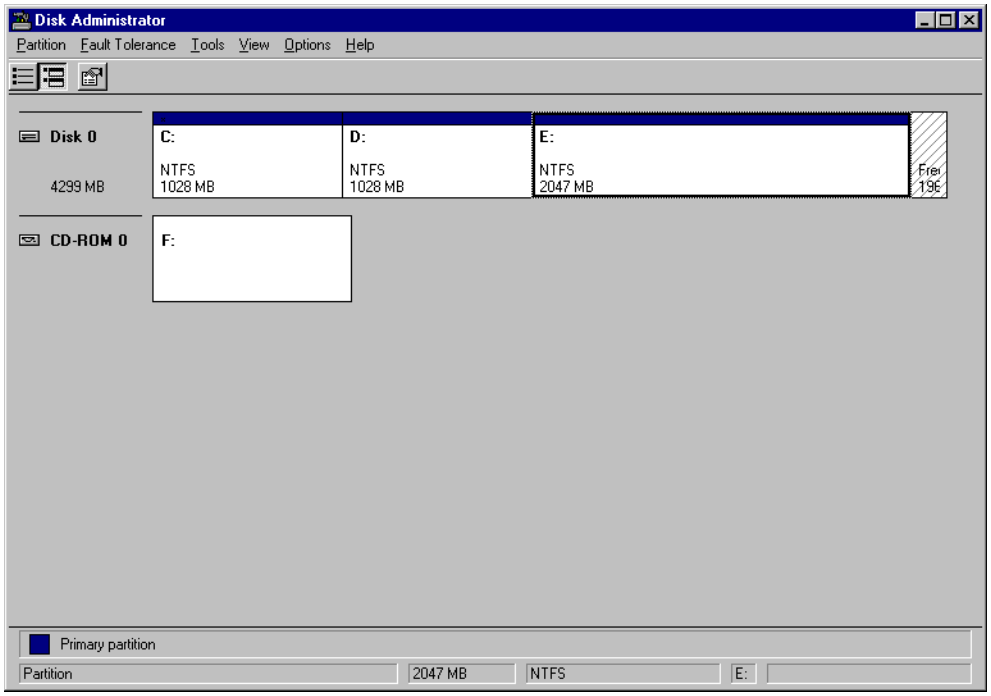

Figure 208. Displaying NTFS Partitions to be Configured as AFS Partitions

To configure an AFS partition on one of the NTFS partitions, open the Control Panel and double-click the **AFS Server Configuration** icon to launch the AFS Server Configuration window, shown in the following figure:

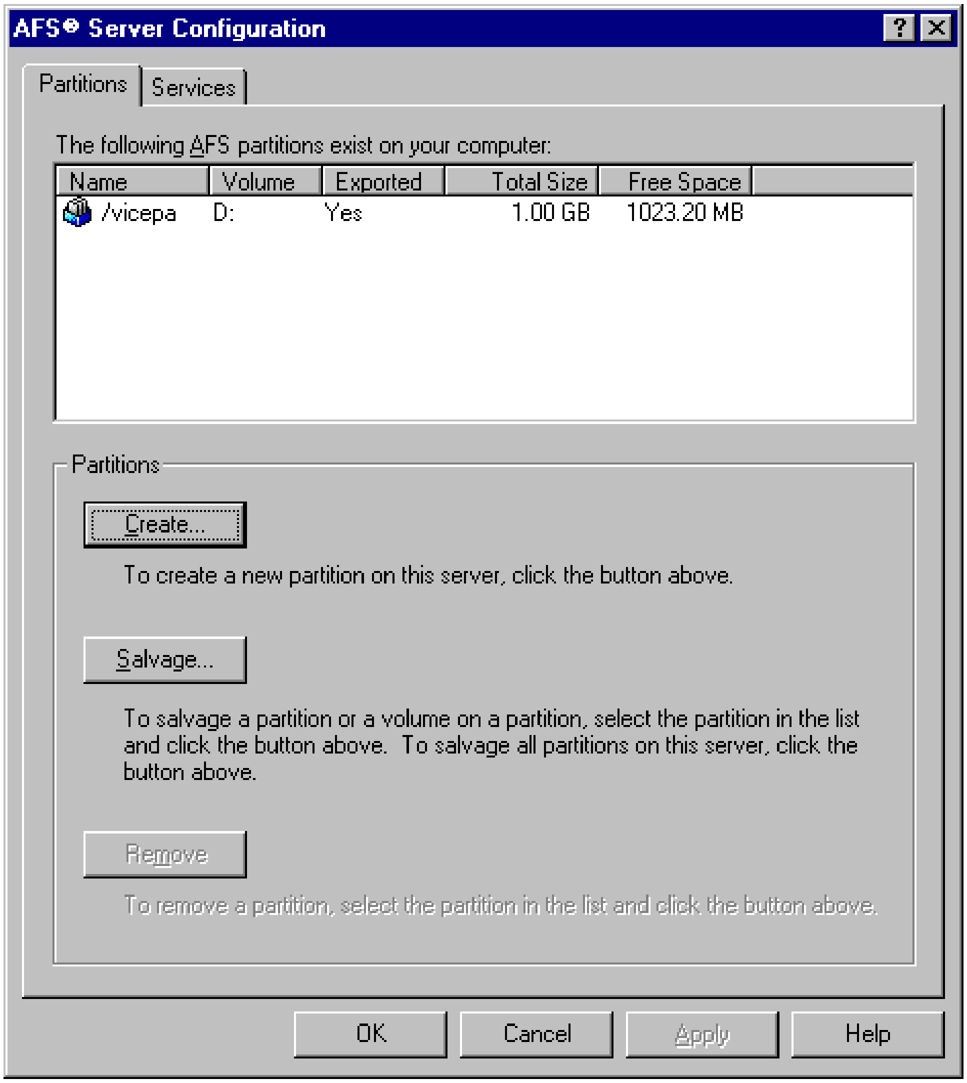
AFS® Server Configuration

The following AFS partitions exist on your computer:

Name	Volume	Exported	Total Size	Free Space
/vicepa	D:	Yes	1.00 GB	1023.20 MB

To create a new partition on this server, click the button above.

To salvage a partition or a volume on a partition, select the partition in the list and click the button above. To salvage all partitions on this server, click the button above.

To remove a partition, select the partition in the list and click the button above.

Figure 209. The Partitions Tab of the AFS Server Configuration

In the AFS Server Configuration window, click the **Create...** button to create a new AFS partition. The Create AFS Partition window appears:

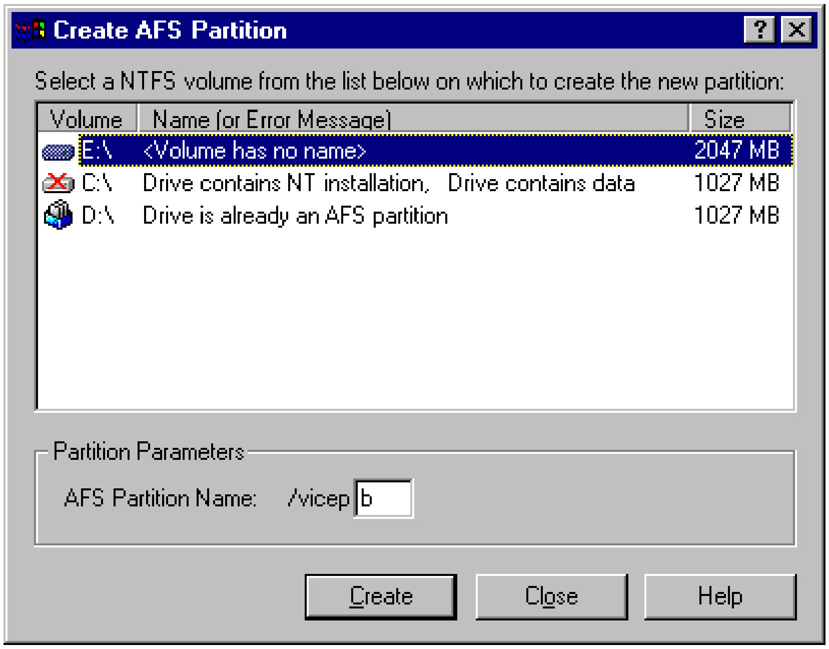

Create AFS Partition

Select a NTFS volume from the list below on which to create the new partition:

Volume	Name (or Error Message)	Size
E:\	<Volume has no name>	2047 MB
C:\	Drive contains NT installation, Drive contains data	1027 MB
D:\	Drive is already an AFS partition	1027 MB

Partition Parameters

AFS Partition Name: /vicep b

Figure 210. Selecting an NTFS Partition to Create a New AFS Partition

The Create AFS Partition window displays information about all NTFS volumes on the server. Notice that drive C: is not able to be configured as an AFS partition because it already contains NT installation files, and drive D: is already an AFS partition. In this example we selected drive E: from the list. In the field for Partition Parameters, you must assign a name to the AFS partition. Since volume names are in the form /vicep*x*, where *x* is a letter in the range from a to z and from aa to iv, just fill in the value for *x*. After you have selected a drive letter and specified a name for the AFS partition, click the **Create** button.

A confirmation window will appear to let you know that the AFS partition was successfully created:

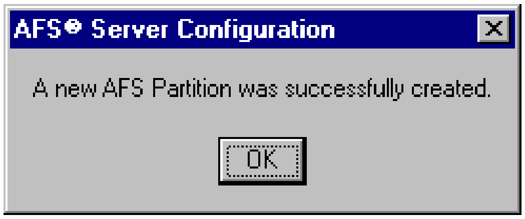

AFS® Server Configuration

A new AFS Partition was successfully created.

Figure 211. Confirmation Window

The Create AFS Partition window will be updated to indicate that drive E: is now an AFS partition:

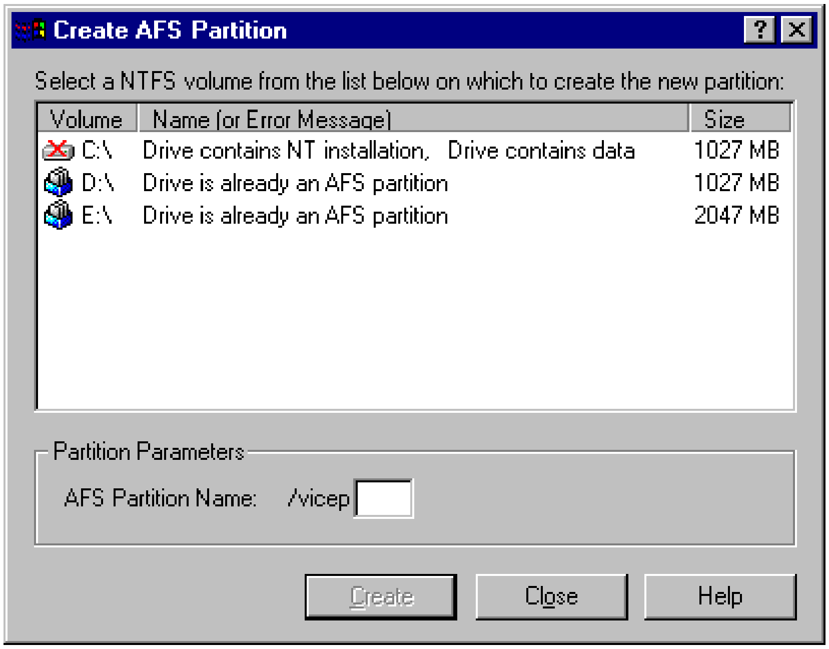

Select a NTFS volume from the list below on which to create the new partition:

Volume	Name (or Error Message)	Size
C:\	Drive contains NT installation, Drive contains data	1027 MB
D:\	Drive is already an AFS partition	1027 MB
E:\	Drive is already an AFS partition	2047 MB

Partition Parameters

AFS Partition Name: /vicep

Figure 212. Displaying the New AFS Partition

Click the **Close** button to close this window.

When you return to the AFS Server Configuration window, you will notice that the newly created AFS partition is not exported yet:

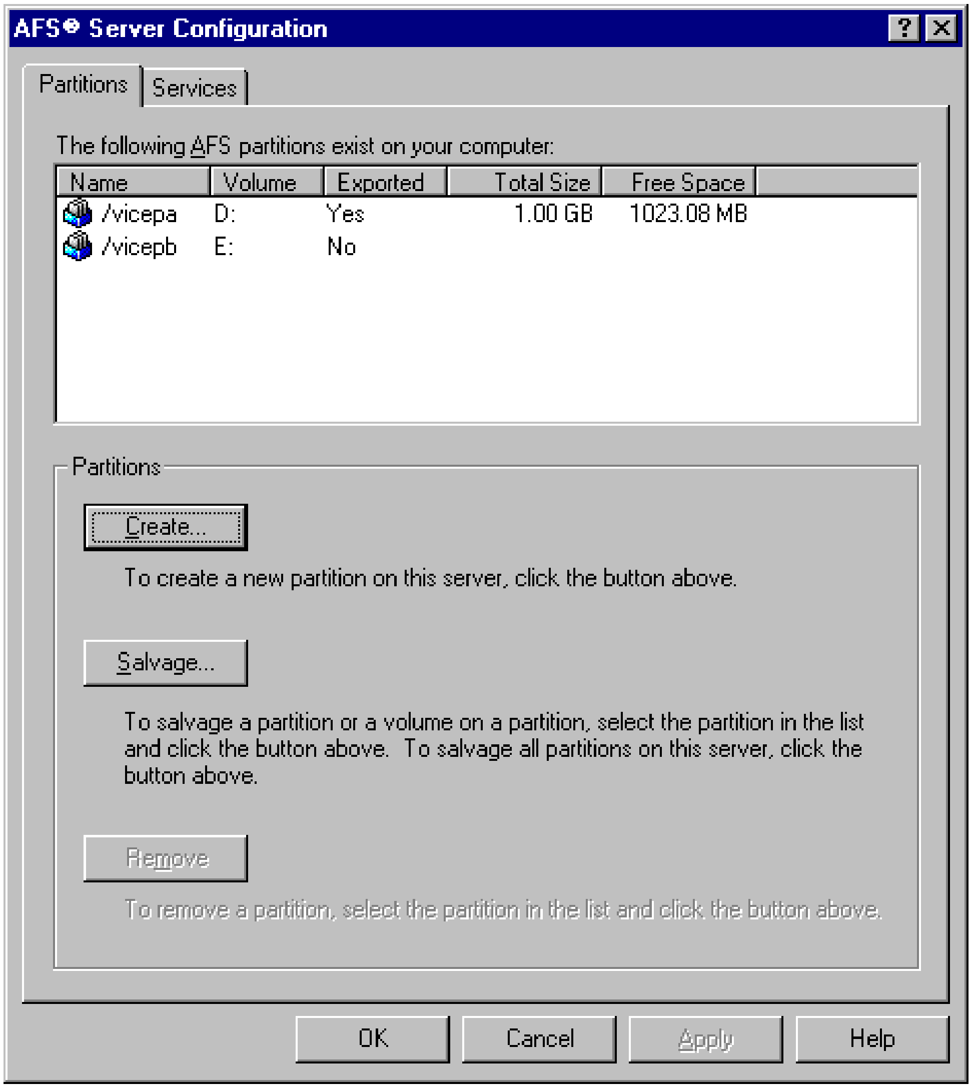

The following AFS partitions exist on your computer:

Name	Volume	Exported	Total Size	Free Space
/vicepa	D:	Yes	1.00 GB	1023.08 MB
/vicepb	E:	No		

To create a new partition on this server, click the button above.

To salvage a partition or a volume on a partition, select the partition in the list and click the button above. To salvage all partitions on this server, click the button above.

To remove a partition, select the partition in the list and click the button above.

Figure 213. Displaying the New Unexported AFS Partition

This means that the AFS File Server cannot store AFS volumes on the partition. To fix this, you must restart the File Server collection of processes, fs. See 6.2.5, “Restarting AFS Processes” on page 244, for instructions on restarting server processes.

Until you restart the fs collection of processes, the new AFS partition, /vicepb, will not appear in the AFS Server Manager’s Partition tab:

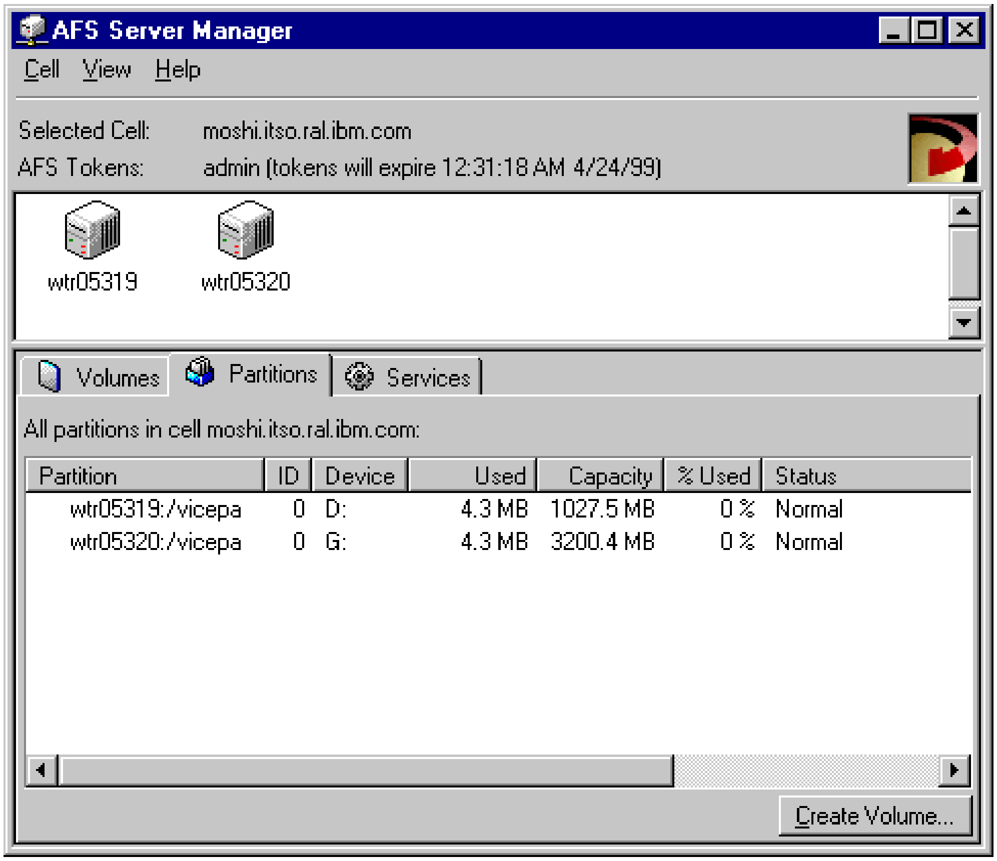

Selected Cell: moshi.itso.ral.ibm.com
AFS Tokens: admin (tokens will expire 12:31:18 AM 4/24/99)

All partitions in cell moshi.itso.ral.ibm.com:

Partition	ID	Device	Used	Capacity	% Used	Status
wtr05319:/vicepa	0	D:	4.3 MB	1027.5 MB	0 %	Normal
wtr05320:/vicepa	0	G:	4.3 MB	3200.4 MB	0 %	Normal

Figure 214. The Partitions Tab before Stop/Start the AFS File Server

As shown in the following figure, the new AFS partition, /vicepb, does appear after you restart the fs collection of processes:

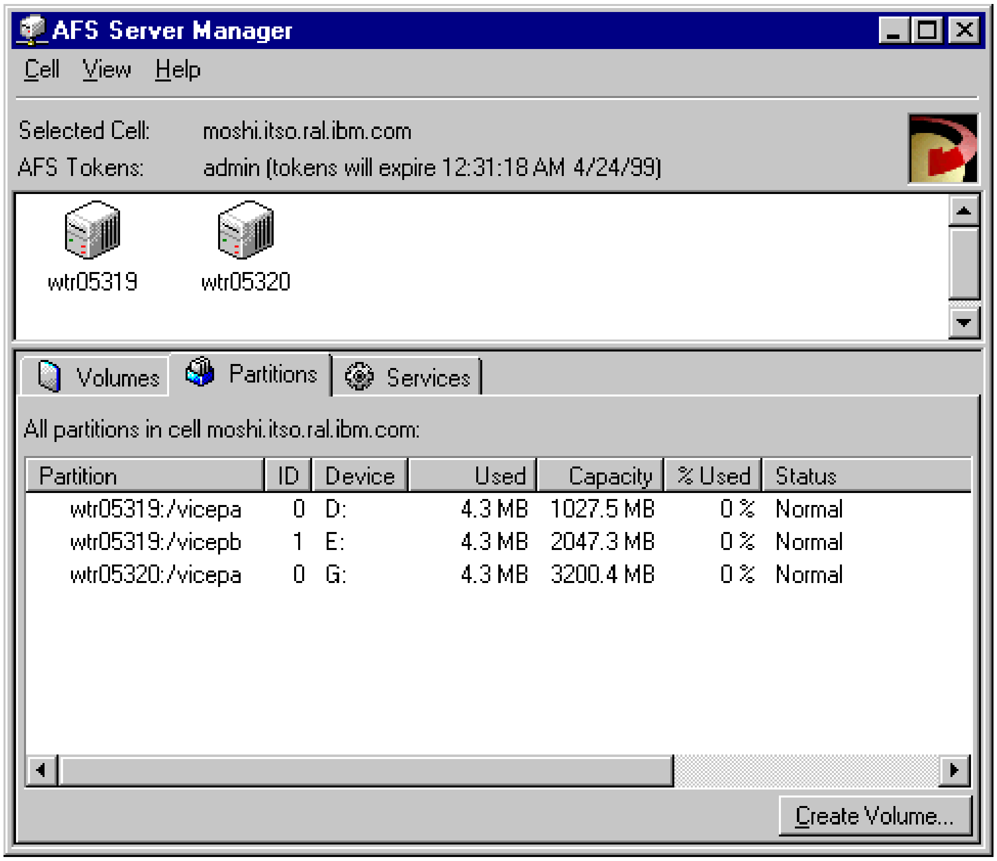

Selected Cell: moshi.itso.ral.ibm.com
AFS Tokens: admin (tokens will expire 12:31:18 AM 4/24/99)

All partitions in cell moshi.itso.ral.ibm.com:

Partition	ID	Device	Used	Capacity	% Used	Status
wtr05319:/vicepa	0	D:	4.3 MB	1027.5 MB	0 %	Normal
wtr05319:/vicepb	1	E:	4.3 MB	2047.3 MB	0 %	Normal
wtr05320:/vicepa	0	G:	4.3 MB	3200.4 MB	0 %	Normal

Figure 215. The Partitions Tab after Stop/Start the AFS File Server

Also notice that when you return to the AFS Server Configuration window, the newly created AFS partition is now exported:

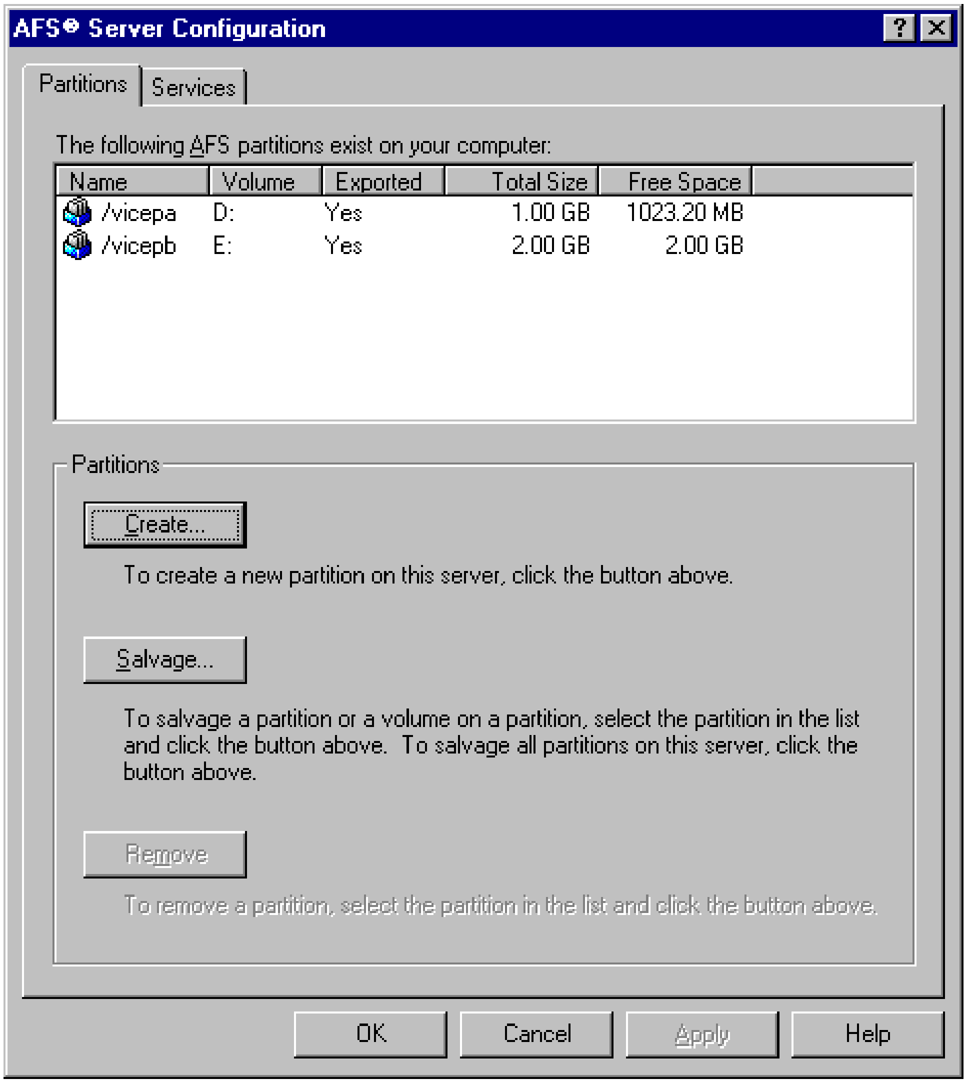

The following AFS partitions exist on your computer:

Name	Volume	Exported	Total Size	Free Space
/vicepa	D:	Yes	1.00 GB	1023.20 MB
/vicepb	E:	Yes	2.00 GB	2.00 GB

To create a new partition on this server, click the button above.

To salvage a partition or a volume on a partition, select the partition in the list and click the button above. To salvage all partitions on this server, click the button above.

To remove a partition, select the partition in the list and click the button above.

Figure 216. Displaying the New Exported AFS Partition

6.3.2 Creating AFS Partitions on AIX and Solaris

For AIX and Solaris there are several steps to set up AFS partitions:[2]

1. Become the local superuser root on the machine, if you are not already, by issuing the su command:

```
% su root
Password:
```

2. Decide how many AFS partitions to divide the new disk into and the names of the directories on which to mount them. Recall that AFS partition names are of the format /vicep*x*, where *x* is a letter in the range from a to z and from aa to iv, for example /vicepf, /vicepz, /vicepdr and /vicepiv. To display the names of the existing server partitions on the machine, issue the command:

```
vos listpart
```

[2] Notice however that there is no equivalent way to do this in the Windows NT Command Prompt window.

However, you should include the `-localauth` flag because you are logged in as the local superuser root. Therefore, the full command should look like the following:

```
vos listpart file_server_name -localauth
```

3. Create each directory in the server's root directory where the vice partitions will be mounted.

   ```
   mkdir /vicepx
   ```

4. Using a text editor, create an entry in the machine's file systems registry file for each new disk partition, mapping its device name to the directory you created in the previous step. For AIX this file is /etc/filesystems and for Solaris it is /etc/vfstab.
5. If the operating system requires that you shut off the machine to install a new disk, issue the command:

   ```
   bos shutdown
   ```

   to shut down all AFS server processes other than the BOS Server. This terminates safely when you shut off the machine. The full command should be:

   ```
   bos shutdown machine_name -localauth [-wait]
   ```

6. If necessary, shut off the machine. Install and format the new disk according to the instructions provided by the disk and operating system vendors. If necessary, edit the disk's partition table to reflect the changes you made to the file's system registry file in Step 4 on page 254; consult the operating system documentation for instructions.
7. If you shut off the machine in Step 6 on page 254, turn it on. Otherwise, issue the command:

   ```
   bos restart
   ```

   to restart the fs process, forcing it to recognize the new set of server partitions. The command syntax is shown below:

   ```
   bos restart machine_name fs -localauth
   ```

8. Issue the command:

   ```
   bos status
   ```

   to verify that all server processes are running correctly. The syntax of this command is shown below:

   ```
   bos status machine_name
   ```

6.3.3 Salvaging AFS Partitions

The Salvager returns consistency to the AFS file system's data, and is roughly equivalent to running `fsck` on UNIX and `chkdsk` on Windows NT. The Salvager is run automatically at startup time if the bosserver process detects that AFS server processes were not properly shut down. It may also be necessary to invoke the Salvager manually if certain volume operations are interrupted or if other sources of inconsistencies appear, such as disk corruption, controller errors, etc.

6.3.3.1 Salvaging on Windows NT

To invoke the Salvager on Windows NT, open the Control Panel and double-click the **AFS Server Configuration** icon to access the AFS Server Configuration window:

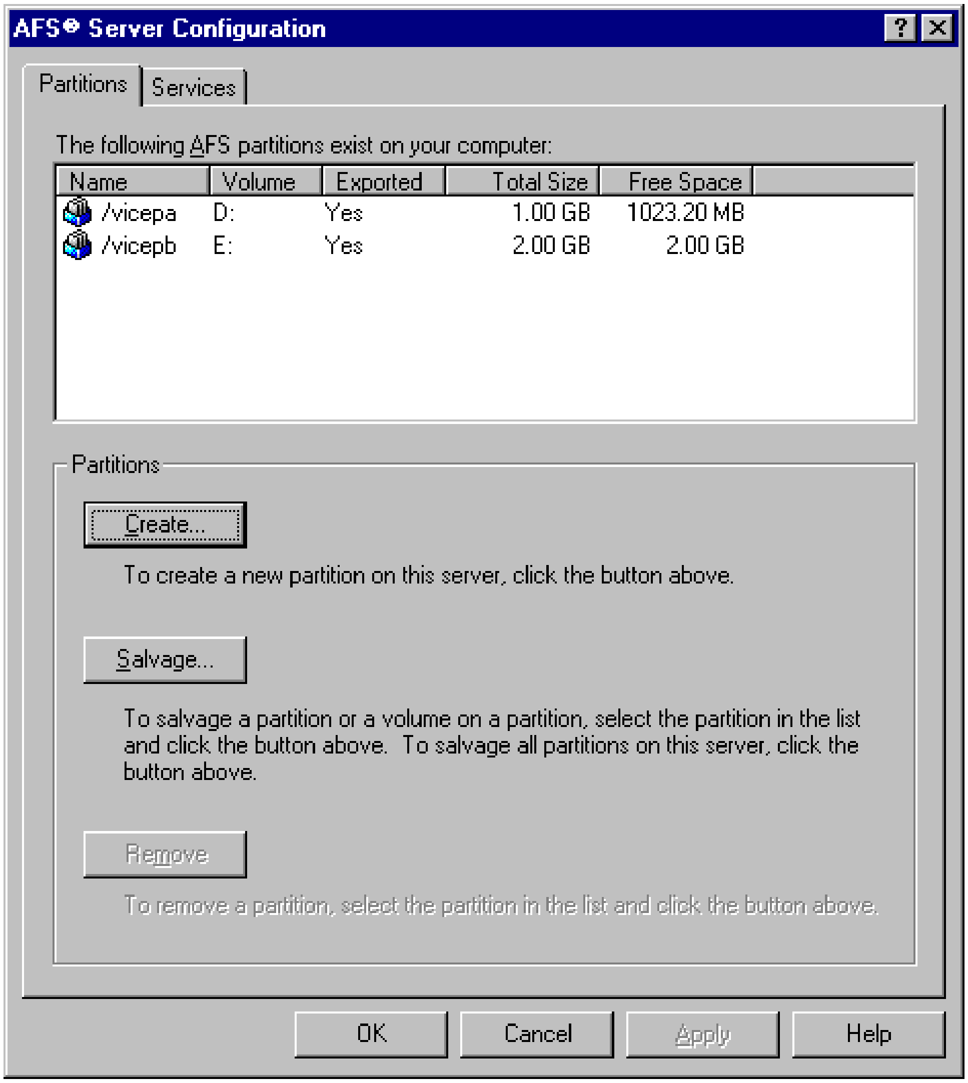

The following AFS partitions exist on your computer:

Name	Volume	Exported	Total Size	Free Space
/vicepa	D:	Yes	1.00 GB	1023.20 MB
/vicepb	E:	Yes	2.00 GB	2.00 GB

Partitions

To create a new partition on this server, click the button above.

To salvage a partition or a volume on a partition, select the partition in the list and click the button above. To salvage all partitions on this server, click the button above.

To remove a partition, select the partition in the list and click the button above.

Figure 217. Server Configuration Window - No Selection

At this point an administrator can either select a particular volume to salvage, an entire partition to salvage, or all volumes and partitions on the server. In our example we only have one partition with two volumes anyway, so we opt to salvage the entire server. In the AFS Server Configuration window we click the **Salvage...** button to access the Salvage window:

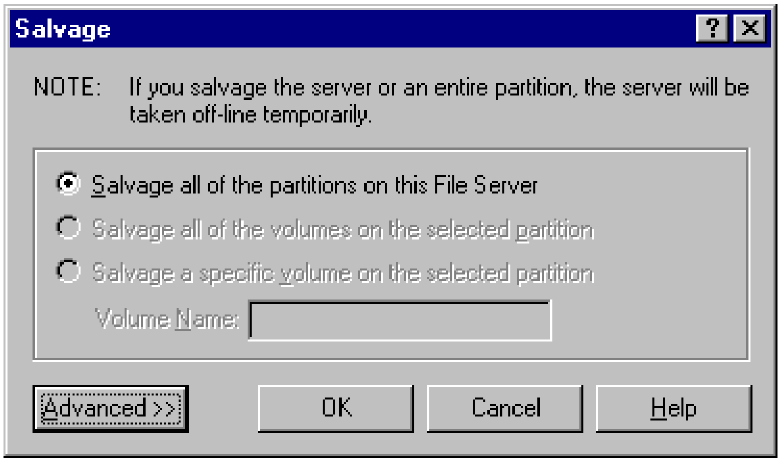

NOTE: If you salvage the server or an entire partition, the server will be taken off-line temporarily.

● Salvage all of the partitions on this File Server

○ Salvage all of the volumes on the selected partition

○ Salvage a specific volume on the selected partition

Volume Name:

Figure 218. Salvage Entire File Server

Note that we are only given the option of salvaging all partitions on the server because we did not select a specific volume or partition in the previous step. And because we are salvaging all partitions and volumes on this server we are also warned that the server will be taken offline temporarily. In the case of salvaging individual volumes or partitions only that volume or partition is taken offline.

Click **OK** when ready to start. You will be prompted for the AFS administrator password and then the salvage will start. The Salvager log file will be displayed and results output during the salvage run:

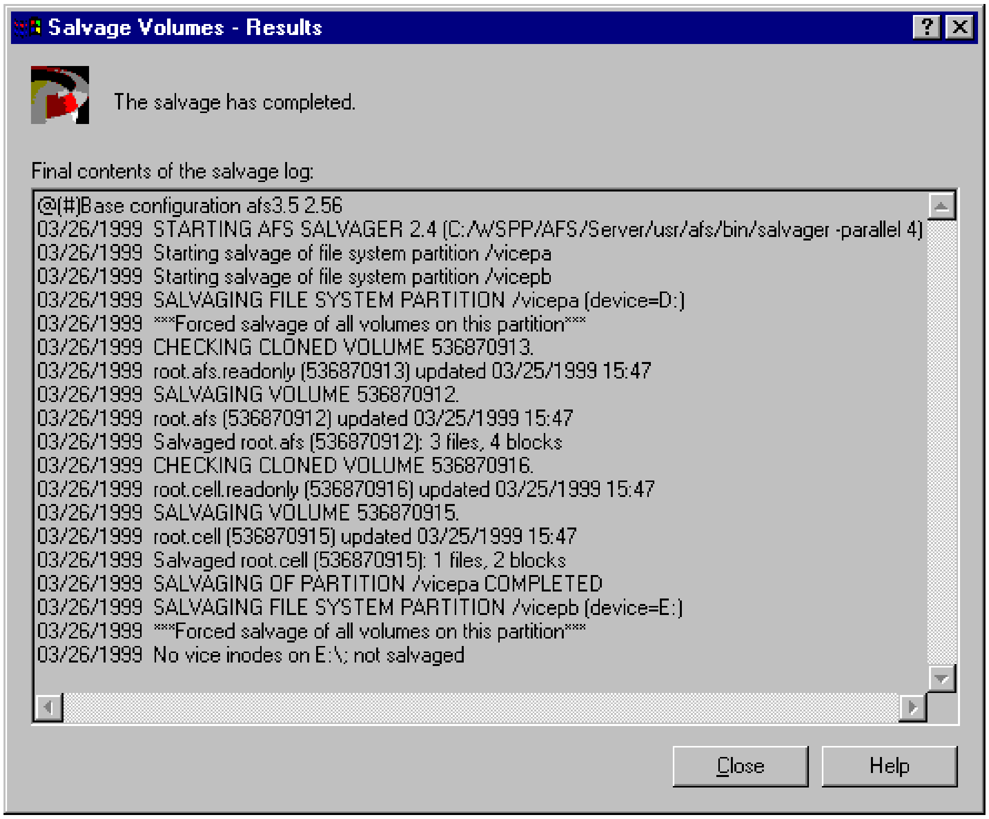

The salvage has completed.

Final contents of the salvage log:

```
@(#)Base configuration afs3.5 2.56
03/26/1999  STARTING AFS SALVAGER 2.4 (C:/WSPP/AFS/Server/usr/afs/bin/salvager -parallel 4)
03/26/1999  Starting salvage of file system partition /vicepa
03/26/1999  Starting salvage of file system partition /vicepb
03/26/1999  SALVAGING FILE SYSTEM PARTITION /vicepa (device=D:)
03/26/1999  ***Forced salvage of all volumes on this partition***
03/26/1999  CHECKING CLONED VOLUME 536870913.
03/26/1999  root.afs.readonly (536870913) updated 03/25/1999 15:47
03/26/1999  SALVAGING VOLUME 536870912.
03/26/1999  root.afs (536870912) updated 03/25/1999 15:47
03/26/1999  Salvaged root.afs (536870912): 3 files, 4 blocks
03/26/1999  CHECKING CLONED VOLUME 536870916.
03/26/1999  root.cell.readonly (536870916) updated 03/25/1999 15:47
03/26/1999  SALVAGING VOLUME 536870915.
03/26/1999  root.cell (536870915) updated 03/25/1999 15:47
03/26/1999  Salvaged root.cell (536870915): 1 files, 2 blocks
03/26/1999  SALVAGING OF PARTITION /vicepa COMPLETED
03/26/1999  SALVAGING FILE SYSTEM PARTITION /vicepb (device=E:)
03/26/1999  ***Forced salvage of all volumes on this partition***
03/26/1999  No vice inodes on E:\; not salvaged
```

Figure 219. Results of Salvaging File Server

The Salvager will notify you when it has completed its run:

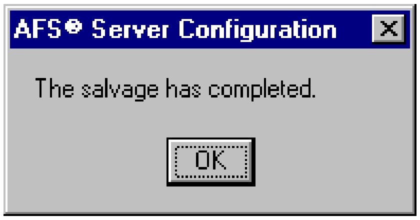

The salvage has completed.

Figure 220. Salvage Completion Notice

Check the log output for any errors or warning messages.

To salvage an individual partition (and all of the volumes on it) select that partition in the AFS Server Configuration window and click the **Salvage...** button:

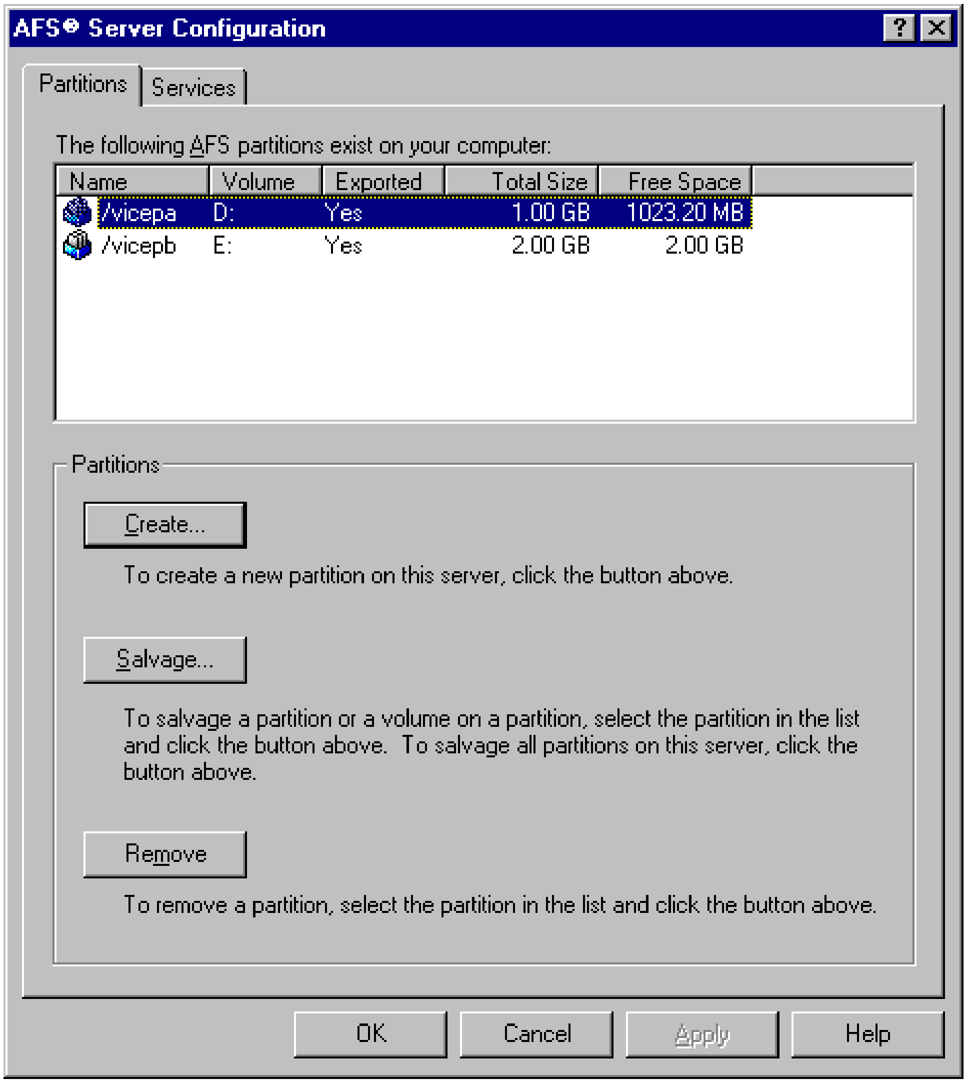

The following AFS partitions exist on your computer:

Name	Volume	Exported	Total Size	Free Space
/vicepa	D:	Yes	1.00 GB	1023.20 MB
/vicepb	E:	Yes	2.00 GB	2.00 GB

To create a new partition on this server, click the button above.

To salvage a partition or a volume on a partition, select the partition in the list and click the button above. To salvage all partitions on this server, click the button above.

To remove a partition, select the partition in the list and click the button above.

Figure 221. Selecting a Partition to Salvage

Click the button to salvage all volumes on the selected partition and click **OK**:

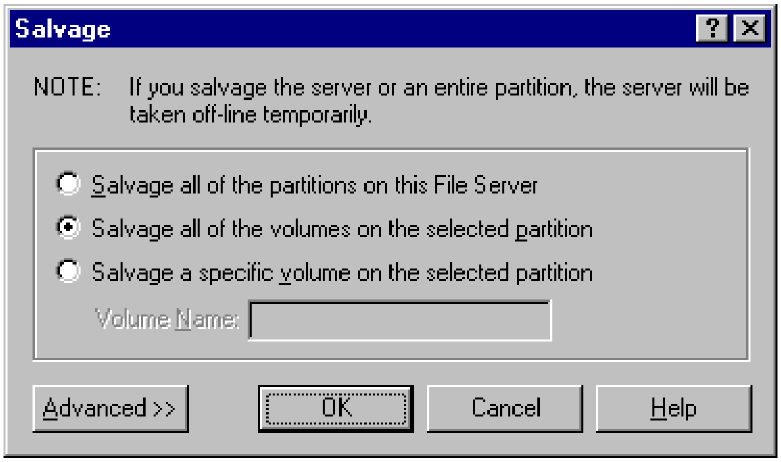

NOTE: If you salvage the server or an entire partition, the server will be taken off-line temporarily.

- ○ Salvage all of the partitions on this File Server
- ● Salvage all of the volumes on the selected partition
- ○ Salvage a specific volume on the selected partition

Volume Name:

Figure 222. Salvaging All Volumes on Partition

Again, this will temporarily make the partition and the volumes it contains unavailable to users.

The administrator will be prompted for the AFS admin password. The results are output to the Salvager log window:

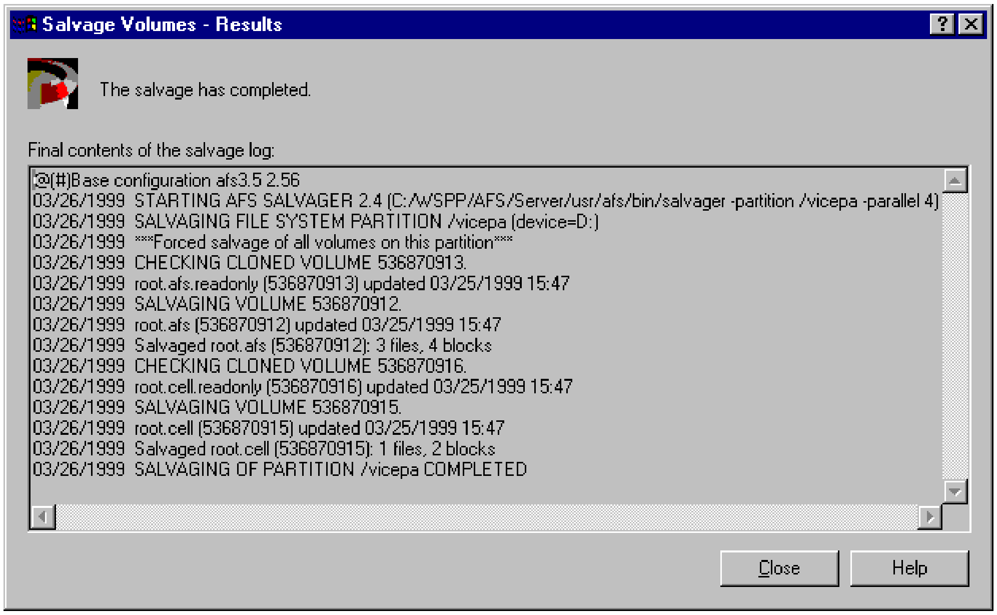

The salvage has completed.

Final contents of the salvage log:

```
@(#)Base configuration afs3.5 2.56
03/26/1999  STARTING AFS SALVAGER 2.4 (C:/WSPP/AFS/Server/usr/afs/bin/salvager -partition /vicepa -parallel 4)
03/26/1999  SALVAGING FILE SYSTEM PARTITION /vicepa (device=D:)
03/26/1999  ***Forced salvage of all volumes on this partition***
03/26/1999  CHECKING CLONED VOLUME 536870913.
03/26/1999  root.afs.readonly (536870913) updated 03/25/1999 15:47
03/26/1999  SALVAGING VOLUME 536870912.
03/26/1999  root.afs (536870912) updated 03/25/1999 15:47
03/26/1999  Salvaged root.afs (536870912): 3 files, 4 blocks
03/26/1999  CHECKING CLONED VOLUME 536870916.
03/26/1999  root.cell.readonly (536870916) updated 03/25/1999 15:47
03/26/1999  SALVAGING VOLUME 536870915.
03/26/1999  root.cell (536870915) updated 03/25/1999 15:47
03/26/1999  Salvaged root.cell (536870915): 1 files, 2 blocks
03/26/1999  SALVAGING OF PARTITION /vicepa COMPLETED
```

Figure 223. Results of Salvaging a Partition

Finally, to salvage an individual volume, select the partition for that volume in the AFS Server Configuration window (see Figure 221 on page 258), and click the **Salvage...** button.

In order to salvage an individual volume on a partition, the salvage program takes the *volume ID number* rather than the volume name as an argument. To get the volume ID number, use the command:

```
vos examine
```

which is valid in a Windows NT Command Prompt window as well as on AIX or Solaris. In our example we wish to salvage the volume root.cell, so the command

```
vos examine
```

we issue will be related to that volume:

```
C:\>vos examine root.cell
root.cell                         536870915 RW          3 K  On-line
    wtr05319.itso.ral.ibm.com /vicepa
    RWrite  536870915 ROnly  536870916 Backup          0
    MaxQuota       5000 K
    Creation    Fri Apr 16 15:16:40 1999
    Last Update Sun Apr 25 18:48:55 1999
    0 accesses in the past day (i.e., vnode references)

    RWrite: 536870915     ROnly: 536870916
    number of sites -> 2
       server wtr05319.itso.ral.ibm.com partition /vicepa RW Site
       server wtr05319.itso.ral.ibm.com partition /vicepa RO Site

C:\>
```

For the volume root.cell the volume ID number is 536870915 and so that is the value we enter for the Volume Name field:

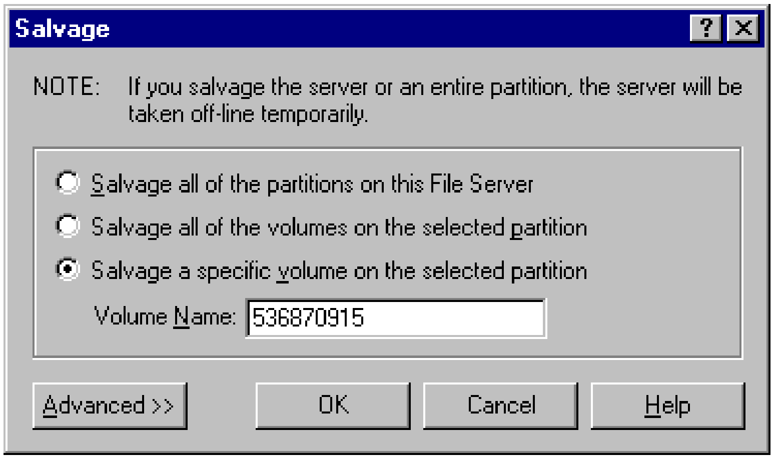

NOTE: If you salvage the server or an entire partition, the server will be taken off-line temporarily.

○ Salvage all of the partitions on this File Server

○ Salvage all of the volumes on the selected partition

● Salvage a specific volume on the selected partition

Volume Name: 536870915

Figure 224. Salvaging an Individual Volume

When the salvage is finished the results will be output to the Salvager log and displayed:

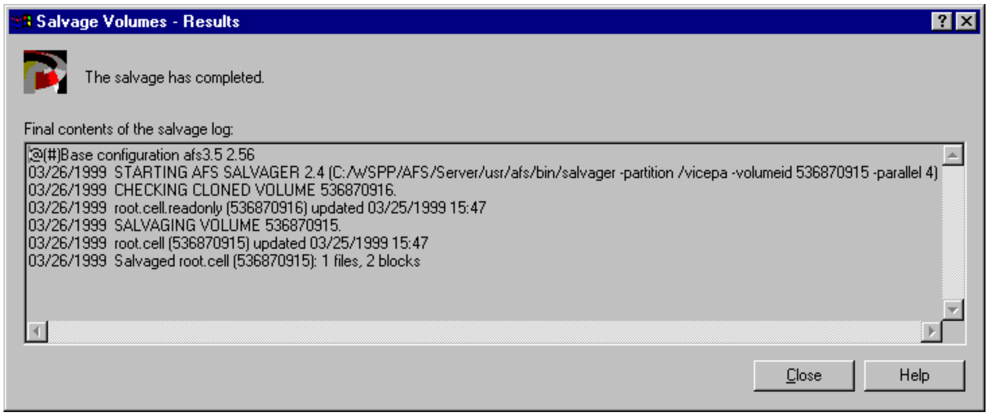

Final contents of the salvage log:

```
@(#)Base configuration afs3.5 2.56
03/26/1999  STARTING AFS SALVAGER 2.4 (C:/WSPP/AFS/Server/usr/afs/bin/salvager -partition /vicepa -volumeid 536870915 -parallel 4)
03/26/1999  CHECKING CLONED VOLUME 536870916.
03/26/1999  root.cell.readonly (536870916) updated 03/25/1999 15:47
03/26/1999  SALVAGING VOLUME 536870915.
03/26/1999  root.cell (536870915) updated 03/25/1999 15:47
03/26/1999  Salvaged root.cell (536870915): 1 files, 2 blocks
```

Figure 225. Results of Salvaging an Individual Volume

6.3.3.2 Salvaging in AIX, Solaris and the Windows NT Command Prompt

To run the Salvager in AIX, Solaris and Windows NT Command Prompt, use the command:

```
bos salvage
```

In the following examples, we salvage an entire File Server and then a partition on a different File Server:

```
C:\> bos salvage -server wtr05319 -all
bos: shutting down fs.
Starting salvage.
bos: waiting for salvage to complete.
bos: salvage completed
bos: restarting fs.

C:\> bos salvage -server wtr05320 -partition vicepa
bos: shutting down fs.
Starting salvage.
bos: salvage completed
bos: restarting fs.

C:\>
```

The output is written to SalvageLog, which can be viewed with the command:

```
bos getlog -server fileserver -file SalvageLog
```

In this final example we salvage the root.cell volume:

```
C:\> bos salvage -server wtr05319 -partition vicepa -volume user -showlog
Starting salvage.
bos: salvage completed
SalvageLog:
@(#)Base configuration afs3.5 3.0
06/08/1999  STARTING AFS SALVAGER 2.4 (C:/WSPP/AFS/Server/usr/afs/bin/salvager /vicepa 536870931)
06/08/1999  CHECKING CLONED VOLUME 536870932.
06/08/1999  user.readonly (536870932) updated 04/25/1999 18:54
06/08/1999  SALVAGING VOLUME 536870931.
06/08/1999  user (536870931) updated 04/25/1999 18:54
06/08/1999  Salvaged user (536870931): 2 files, 3 blocks

C:\>
```

Note that we have to specify the server and partition where the volume is located. Also note that, unlike the Server Configuration window on Windows NT, we can identify the volume by name; we do not have to resort to using only the volume ID number. Finally, we use the `-showlog` switch to show the results to have the Salvager results written to both the SalvageLog and standard output.

6.3.4 Removing an AFS Partition

There may come a time in the life of the AFS cell when you will want to remove AFS partitions, perhaps to decommission a File Server or to replace a disk, or even to reclaim the partition for use by the local operating system.

Exported AFS partitions cannot be removed while the File Service is running. You must first stop the fs trio of processes before removing any partition. Refer back to 6.2.2, “Stopping AFS Processes” on page 235 in this chapter to review how to stop an AFS server process. On a File Server, once the fs trio of processes has been stopped, the partitions are no longer being exported and data is no longer being served. At this point you can delete partitions.

6.3.4.1 Removing AFS Partitions on Windows NT

On Windows NT, partitions are deleted using the AFS Server Configuration Control Panel. Note that when you launch this application, the partitions that are listed for this server are marked as being not exported:

(

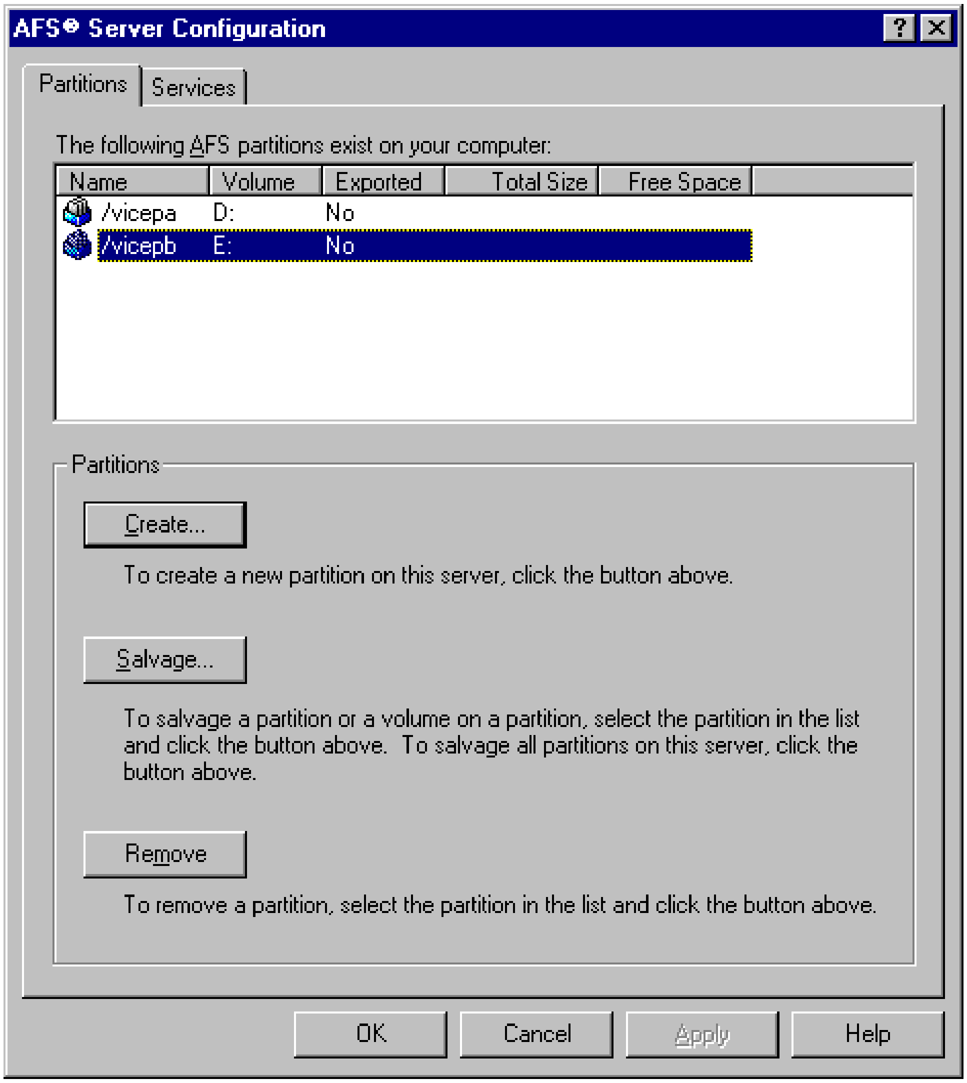

The following AFS partitions exist on your computer:

Name	Volume	Exported	Total Size	Free Space
/vicepa	D:	No		
/vicepb	E:	No		

To create a new partition on this server, click the button above.

To salvage a partition or a volume on a partition, select the partition in the list and click the button above. To salvage all partitions on this server, click the button above.

To remove a partition, select the partition in the list and click the button above.

Figure 226. Selecting a Partition for Removal

Select the partition you wish to delete and click **OK**. A dialog box will ask you to confirm the deletion; click **Yes**.

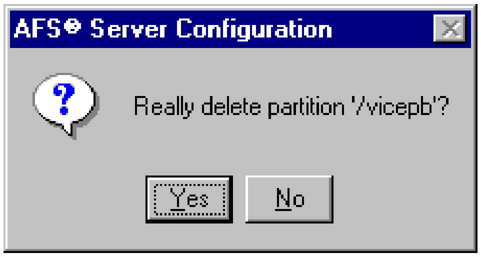

Really delete partition '/vicepb'?

Figure 227. Removal Confirmation

The AFS partition is now deleted:

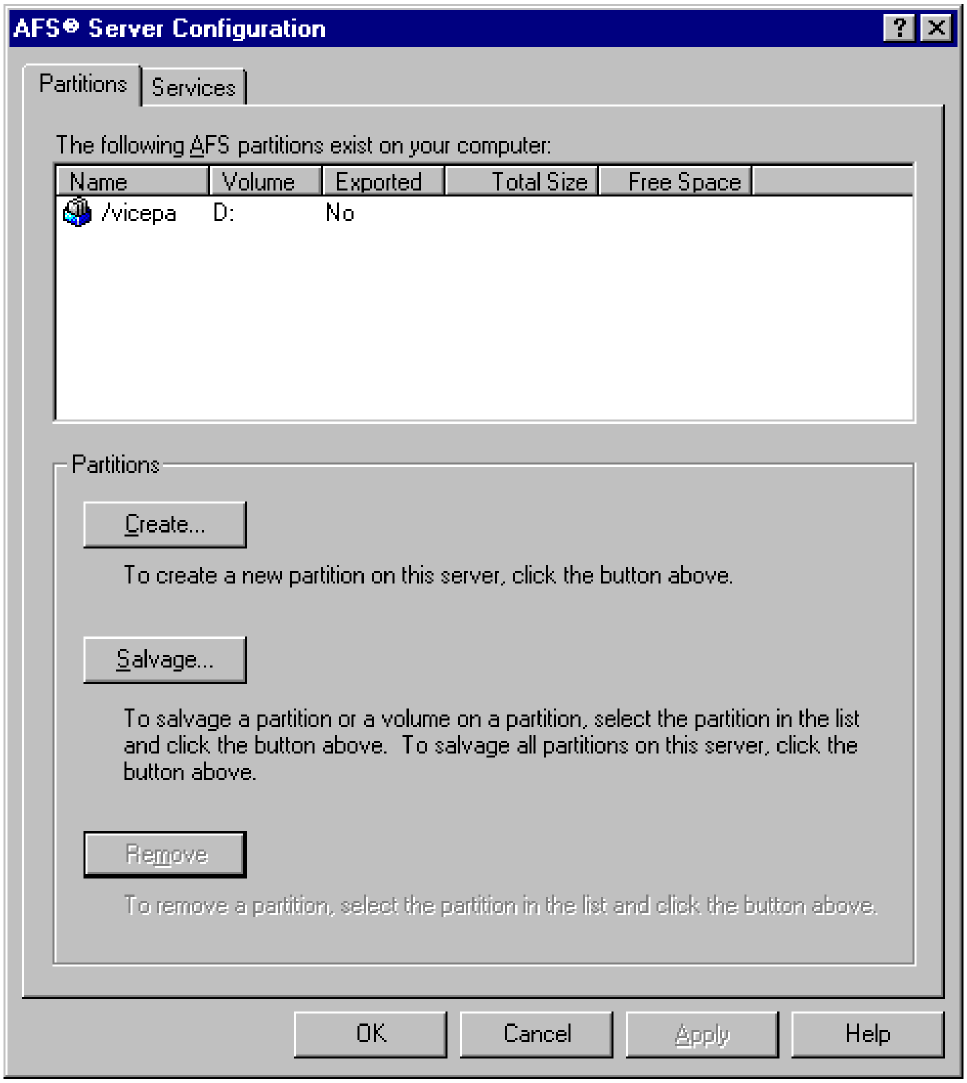

The following AFS partitions exist on your computer:

Name	Volume	Exported	Total Size	Free Space
/vicepa	D:	No		

To create a new partition on this server, click the button above.

To salvage a partition or a volume on a partition, select the partition in the list and click the button above. To salvage all partitions on this server, click the button above.

To remove a partition, select the partition in the list and click the button above.

Figure 228. The Partition Is Gone!

This does not delete the Windows NT partition on that drive, however; there is still an E: drive available to Windows NT in this case. This drive is just unavailable for use by the AFS File Server on this machine. Note also that the remaining partition is still unexported. Restart the fs trio of processes in order to export the partition (see 6.2.5, “Restarting AFS Processes” on page 244).

6.3.4.2 Removing AFS Partitions on AIX and Solaris

On AIX and Solaris, the following steps are necessary to remove an AFS partition from a File Server:[3]

1. Issue the command:

   ```
   vos listvol
   ```

   to list the volumes housed on each partition of each disk you are about to remove, in preparation for removing them or moving them to other partitions. The syntax for this command is the following:

   ```
   vos listvol machine_name [partition_name]
   ```

[3] Notice that there are no equivalent commands for the Windows NT Command Prompt.

2. Move any volume you wish to retain in the file system to another partition. You can only move ReadWrite volumes. For more detailed instructions, and for instructions on moving ReadOnly and Backup volumes, see 6.4.7, “Moving AFS Volumes” on page 298. The command used to move volumes is:

   ```
   vos move
   ```

   and its syntax is the following:

   ```
   vos move volume source_server source_partition dest_source
   dest_partition
   ```

3. Become the local superuser root on the machine, if you are not already, by issuing the command:

   ```
   su root
   ```

   You will be prompted to enter root’s password.

4. Issue the `umount` command, repeating it for each partition on the disk to be removed:

   ```
   umount /dev/partition_block_device_name
   ```

5. Using a text editor, remove or comment out each partition's entry from the machine's file systems registry file. On AIX this file is /etc/filesystems, and on Solaris this file is /etc/vfstab.

6. Remove the /vicep*x* directory associated with each partition, by issuing the following command:

   ```
   rmdir /vicepx
   ```

7. If the operating system requires that you shut off the machine to remove a disk, issue the command:

   ```
   bos shutdown
   ```

   to shut down all AFS server processes other than the BOS Server, which terminates safely when you shut off the machine. The syntax for this command is:

   ```
   bos shutdown machine_name -localauth [-wait]
   ```

8. If necessary, shut off the machine. Remove the disk according to the instructions provided by the disk and operating system vendors. If necessary, edit the disk's partition table to reflect the changes you made to the file’s system registry file in Step 5 on page 265; consult the operating system documentation for instructions.

9. If you shut off the machine in Step 8 on page 265, turn it on. Otherwise, issue the command:

   ```
   bos restart
   ```

   to restart the fs trio of processes, forcing it to recognize the new set of server partitions. The syntax of this command is:

   ```
   bos restart machine_name fs -localauth
   ```

10. Issue the command:

    ```
    bos status
    ```

    to verify that all server processes are running correctly. The full syntax of this command is shown below:

    ```
    bos status machine_name
    ```

6.4 Managing AFS Volumes

Recall from 2.11, “AFS Volume Concepts” on page 46 that in AFS a volume can be thought of as being a unit of administration for moving, replicating, and backing up collections of related files and directories. Volumes are created in AFS partitions, and consist of at least one root directory and possibly files and other subdirectories; in turn this root directory is mounted into the cell’s AFS directory tree. Once you have created partitions for use by AFS, the typical steps that follow are to create a mount point for that volume, set access control lists (ACLs) for the volume and its directories, and if applicable replicate the volume.

6.4.1 Examining AFS Mount Points

Recall that mount points attach volumes into the AFS directory tree. When you create a volume, you create an entry for the volume in the Volume Location Database (VLDB) and you create the volume on the File Server machine’s AFS partition. For example, in Windows NT you do not attach it directly into the AFS directory tree but instead must create a mount point via the Windows NT Explorer window’s pop-up AFS menu. When you view the AFS directory tree from the Windows NT Explorer window, mount points appear as directories.

The WebSphere Performance Pack installation process creates two initial volumes and mount points for you and automatically replicates them. These two initial volumes and mount points are important because they are the root of the AFS directory tree, and this carries some important considerations. The volume root.afs is *the* root of the AFS directory tree, and the volume root.cell is the volume mounted as the cell name underneath the AFS root. In our example root.cell is mounted as .moshi.itso.ral.ibm.com. Both the root.afs and root.cell volumes are automatically replicated for you, and the replica of root.cell is mounted as moshi.itso.ral.ibm.com. Now we take a closer look at these mount points.

To examine a mount point in Windows NT, select the directory that is really the mount point itself and right click to display a Windows NT Explorer pop-up menu. From there, select **AFS** to access the AFS menu, then select **Mount Point**, and finally select **Show...**, as indicated in the following figure:

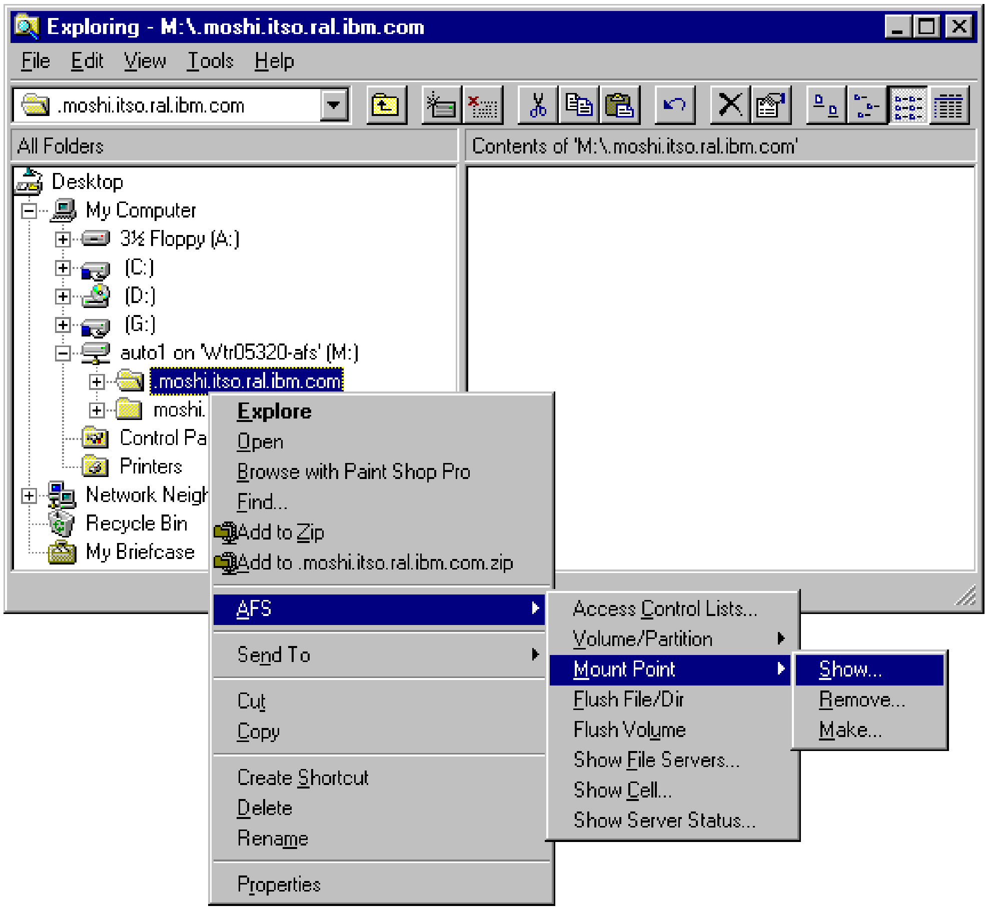

Figure 229. Showing the AFS Mount Point of the Read/Write AFS Cell Path

The Mount Points window appears.

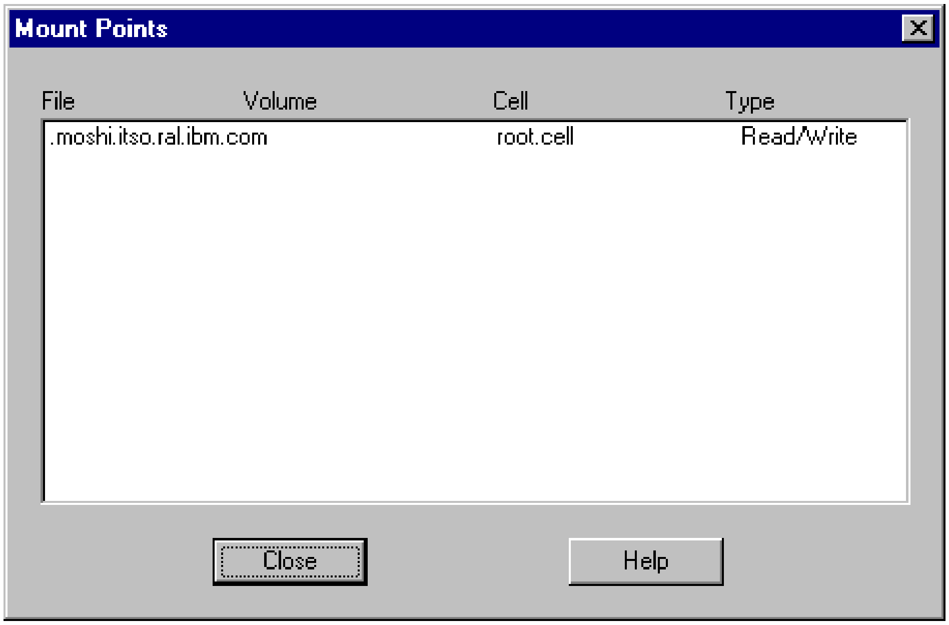

File	Volume	Cell	Type
.moshi.itso.ral.ibm.com		root.cell	Read/Write

Figure 230. Mount Point Information for .moshi.itso.ral.ibm.com

In this example, we selected the directory called **.moshi.itso.ral.ibm.com** from the Windows NT Explorer window. When we examine the information in the Mount Points window we see that this mount point references the volume *root.cell* and that it is a ReadWrite mount point. A ReadWrite mount point is significant in that it signals a client's Cache Manager to ignore replicated volumes and only use the ReadWrite version of a volume. As per the Volume Traversal Rules (see 2.11.2, "Types of Volumes" on page 47 and 2.11.3, "Volume Traversal Rules" on page 47) this implies that any client traversing the AFS directory tree below this ReadWrite mount point will *only* ever see ReadWrite volumes and will never make use of the replicas of volumes mounted below this point.

For AIX, Solaris and the Windows NT Command Prompt use the command:

```
fs lsmount
```

to examine mount points. This is illustrated in the following session screen:

```
C:\>fs lsmount m:\\.moshi.itso.ral.ibm.com
'm:\\.moshi.itso.ral.ibm.com' is a mount point for volume '%root.cell'

C:\>fs lsmount m:\\moshi.itso.ral.ibm.com
'm:\\moshi.itso.ral.ibm.com' is a mount point for volume '#root.cell'
```

Note in this example that the .moshi.itso.ral.ibm.com mount point references %root.cell and moshi.itso.ral.ibm.com references #root.cell. The percent sign (%) indicates a ReadWrite mount point and the pound sign (#) marks a regular mount point. An AFS client traversing moshi.itso.ral.ibm.com would access data from a root.cell.readonly volume (if one exists – see 2.11.4, "Naming Suggestions" on page 50), whereas a client traversing .moshi.itso.ral.ibm.com would access data only from root.cell.

Next, we want to learn more about the root.cell volume.

On Windows NT, we select **.moshi.itso.ral.ibm.com** from the Windows NT Explorer window and access the AFS menu again. This time, we go to the Volume/Partition menu and select **Properties**, as shown in the following figure:

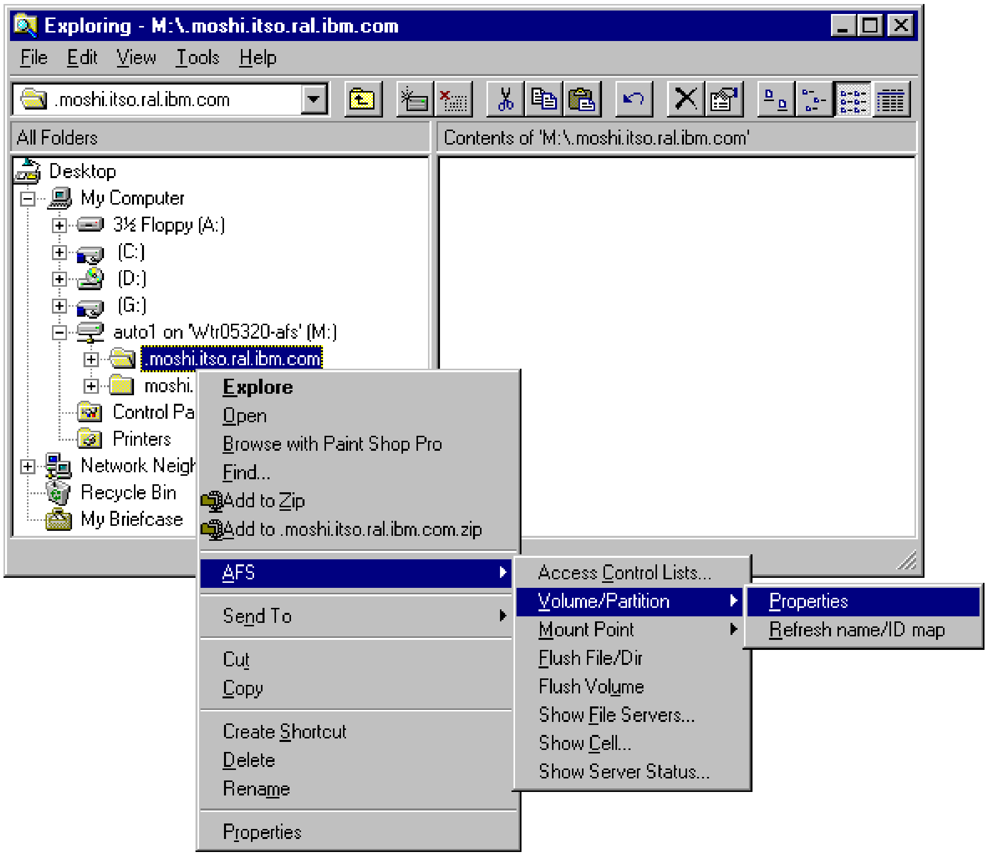

Figure 231. Selecting Volume Properties

The Volume Information window appears:

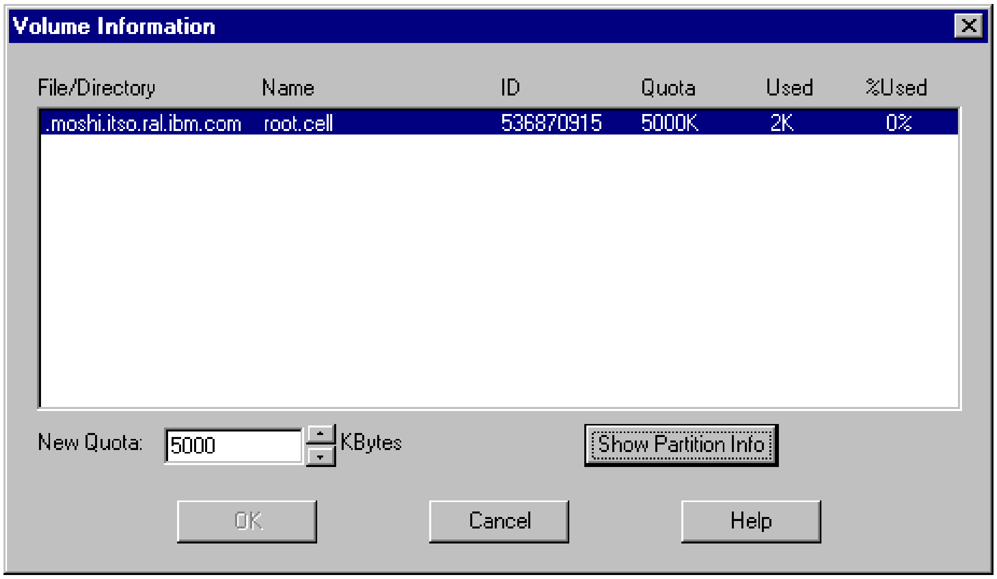

File/Directory	Name	ID	Quota	Used	%Used
.moshi.itso.ral.ibm.com	root.cell	536870915	5000K	2K	0%

New Quota: 5000 KBytes

Figure 232. Volume Information for root.cell

In this example, we see the ID, quota, amount of quota used, and percentage of the quota used for the root.cell volume. To see information about the partition that houses this volume, click the **Show Partition Info** button. The Partition Info window appears:

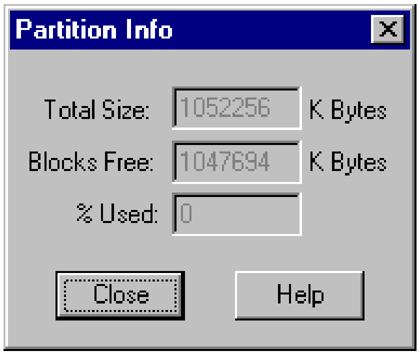

Total Size: 1052256 K Bytes
Blocks Free: 1047694 K Bytes
% Used: 0

Figure 233. Partition Information

In the Partition Info window, we see the size of the partition, the number of blocks that are free, and the percentage of the partition in use.

In the previous example, we examined information about the .moshi.itso.ral.ibm.com directory. We learned that it is actually a ReadWrite mount point referencing the root.cell volume. In the next example, we will examine information about the moshi.itso.ral.ibm.com directory. As before, we select the name of the directory from the Windows NT Explorer window, access the AFS menu, select **Mount Points**, and select **Show...**. The Mount Points window appears:

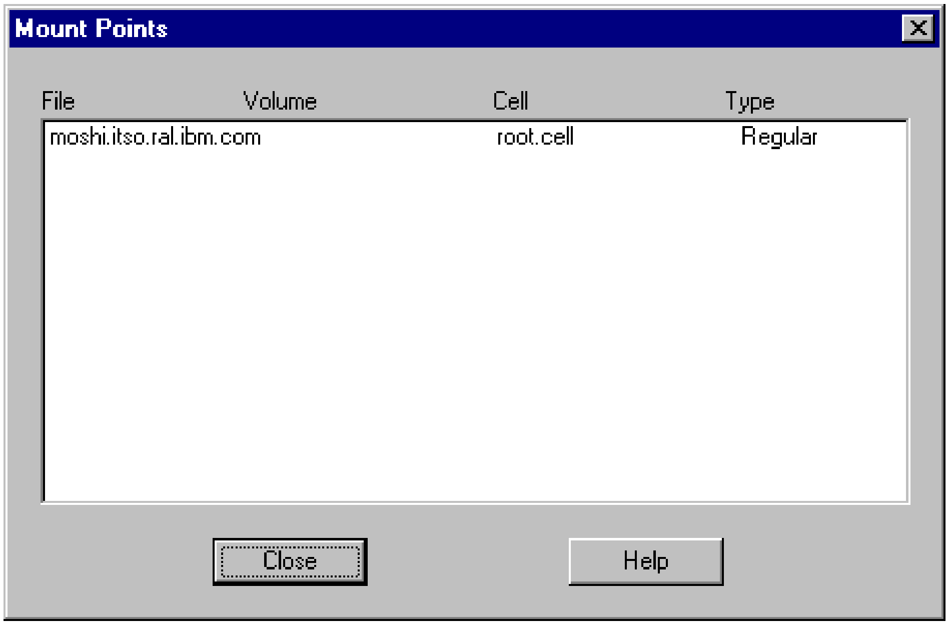

File	Volume	Cell	Type
moshi.itso.ral.ibm.com		root.cell	Regular

Figure 234. Examining Mount Points for moshi.itso.ral.ibm.com

Note that the mount point moshi.itso.ral.ibm.com is flagged as being a Regular mount point. A Regular mount point is typical (see 2.11.3, “Volume Traversal Rules” on page 47); it tells the Cache Manager to use a replica of this volume if available, and use the ReadWrite if no replica has been created. Since root.cell is replicated, any AFS client accessing the moshi.itso.ral.ibm.com directory will be served data from a root.cell.readonly volume.

Getting volume information for the moshi.itso.ral.ibm.com mount point confirms that we are using a replicated version of this volume:

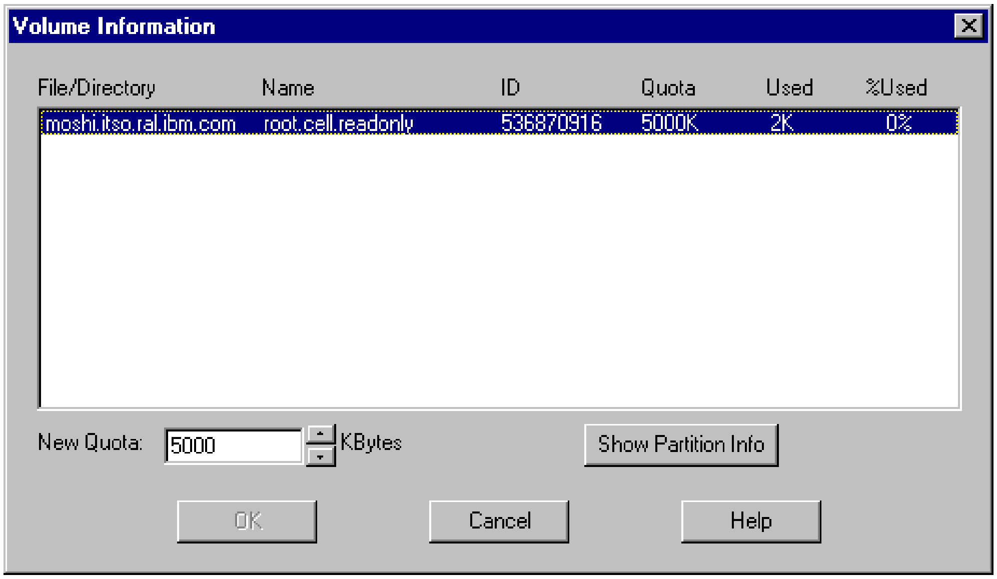

File/Directory	Name	ID	Quota	Used	%Used
moshi.itso.ral.ibm.com	root.cell.readonly	536870916	5000K	2K	0%

New Quota: 5000 KBytes

Figure 235. Volume Information for root.cell.readonly

On AIX, Solaris, and the Windows NT Command Prompt, these same tasks are done using several different commands:

- To list a mount point use the command:

  ```
  fs lsmount
  ```

- To list a volume's quota use the command:

  ```
  fs listquota
  ```

- To set a volume's quota use the command:

  ```
  fs setquota
  ```

- To examine a volume, use the command:

  ```
  vos examine
  ```

- To get partition information, use the commands:

  ```
  vos partinfo
  fs diskinfo
  ```

In the following example we examine the mount point at /afs/moshi.itso.ral.ibm.com/user/bobby, and then change the quota for the volume mounted there:

```
C:\>fs lsmount m:\moshi.itso.ral.ibm.com\user\bobby
'm:\moshi.itso.ral.ibm.com\user\bobby' is a mount point for volume '#user.bobby'

C:\>fs diskfree m:\moshi.itso.ral.ibm.com\user\bobby
Volume Name           kbytes  used     avail     %used
user.bobby           3277256    4504 3272752        0%

C:\>fs listq m:\moshi.itso.ral.ibm.com\user\bobby
Volume Name             Quota    Used    % Used   Partition
user.bobby               5000       2        0%          0%

C:\>fs setq m:\moshi.itso.ral.ibm.com\user\bobby 8000

C:\>fs listq m:\moshi.itso.ral.ibm.com\user\bobby
Volume Name             Quota    Used    % Used   Partition
user.bobby               8000       2        0%          0%
```

We then examine the volume referenced by the mount point and get information about the AFS partition on which it sits:

```
C:\>vos examine user.bobby
user.bobby                        536870934 RW          2 K  On-line
    wtr05320.itso.ral.ibm.com /vicepa
    RWrite  536870934 ROnly          0 Backup         0
    MaxQuota       8000 K
    Creation    Sun Apr 25 14:11:34 1999
    Last Update Sun Apr 25 14:11:34 1999
    5 accesses in the past day (i.e., vnode references)

    RWrite: 536870934
    number of sites -> 1
       server wtr05320.itso.ral.ibm.com partition /vicepa RW Site

C:\>vos partinfo -server wtr05320 -partition vicepa
Free space on partition /vicepa: 3272752 K blocks out of total 3277256
```

6.4.2 Creating AFS Volumes

Now we are ready to start adding additional volumes to our AFS cell. Again, volumes are administrative units of data sets. You may assign volumes to individual users, create volumes for particular software packages or for storing particular sets of related data, etc. Examples of volumes might include user.dumbo, user.dumbo.src, i386_nt40.msoffice, wspp.data, wspp.applications, wte.cache.1, and so on.

The typical steps involved in creating AFS volumes include:

1. Select a partition with free space.
2. Create a volume on that partition.
3. Create a mount point for the volume in AFS.
4. Set an ACL on the mounted volume.
5. Set a quota on the mounted volume.
6. If applicable, replicate the volume and release the data to the replicas.

We will cover each of these steps in turn.

In Windows NT, to create an AFS volume on an AFS partition select in Server Manager the desired AFS partition by highlighting it. Position the pointer on the name of the selected AFS partition and click the right mouse button to display a pop-up menu. From the pop-up menu, select **Create Volume...**, as shown in the following figure:

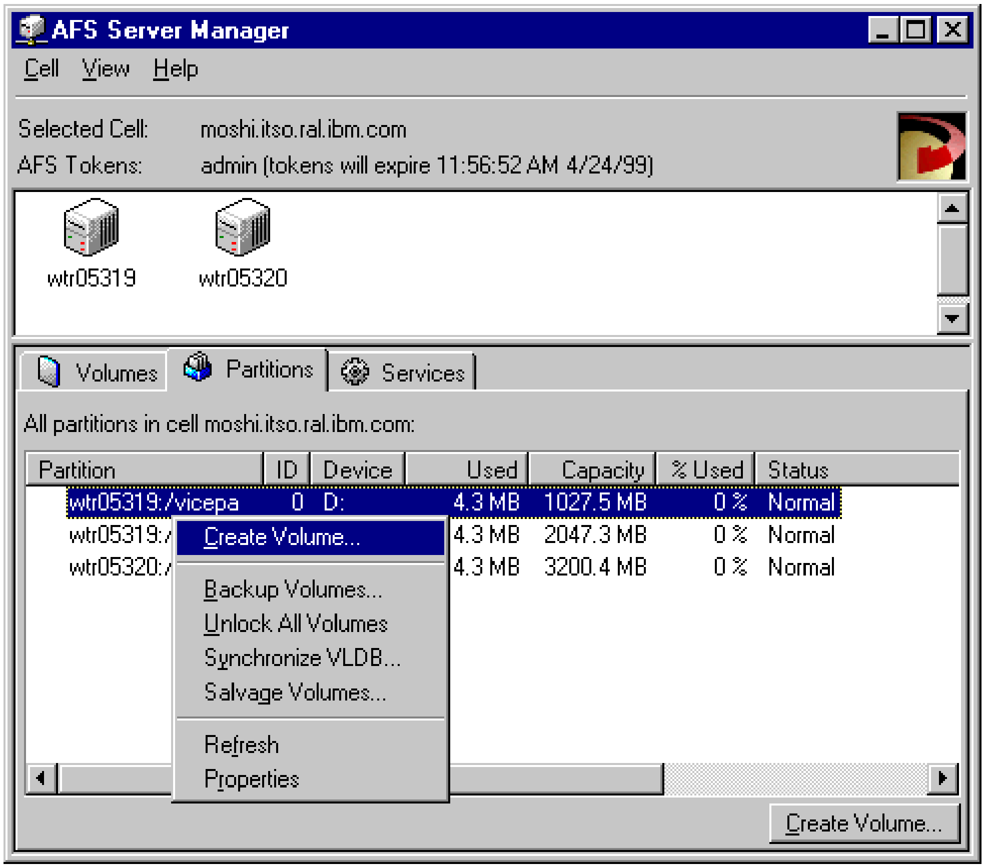

Selected Cell: moshi.itso.ral.ibm.com
AFS Tokens: admin (tokens will expire 11:56:52 AM 4/24/99)

All partitions in cell moshi.itso.ral.ibm.com:

Partition	ID	Device	Used	Capacity	% Used	Status
wtr05319:/vicepa	0	D:	4.3 MB	1027.5 MB	0 %	Normal
wtr05319:/			4.3 MB	2047.3 MB	0 %	Normal
wtr05320:/			4.3 MB	3200.4 MB	0 %	Normal

Figure 236. Creating a Volume in AFS Server Manager

This way, you will access the Create Volume window:

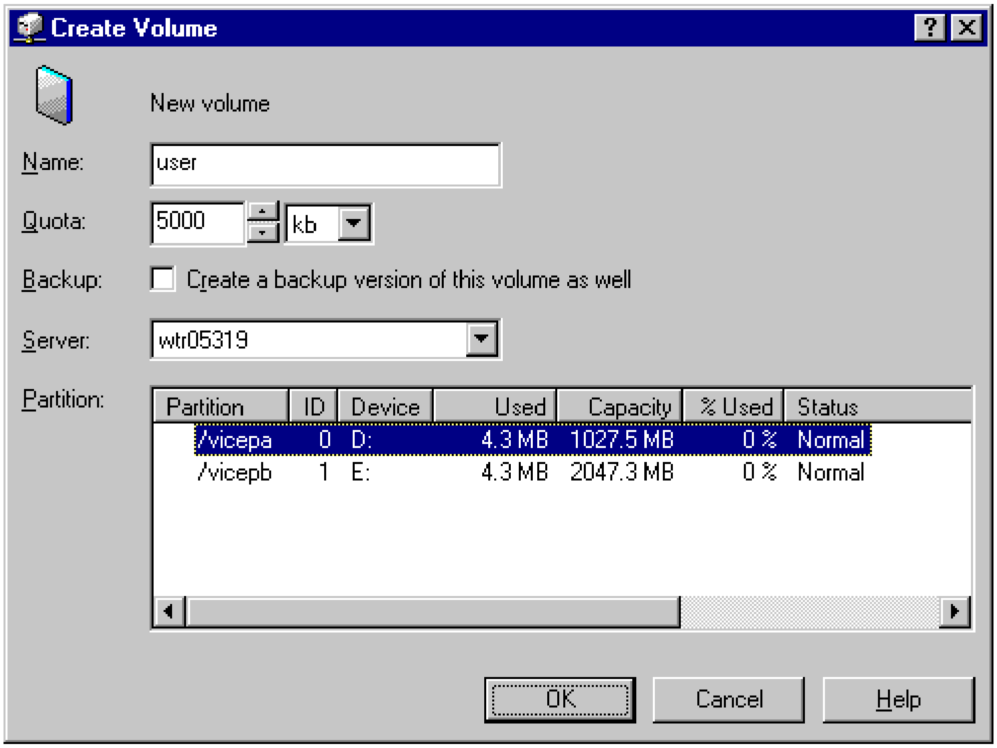

Name: user

Quota: 5000 kb

Backup: □ Create a backup version of this volume as well

Server: wtr05319

Partition:

Partition	ID	Device	Used	Capacity	% Used	Status
/vicepa	0	D:	4.3 MB	1027.5 MB	0 %	Normal
/vicepb	1	E:	4.3 MB	2047.3 MB	0 %	Normal

Figure 237. Creating an AFS Volume

In this example, we indicate that we want to create a volume called user. The initial quota, or maximum amount of space this volume can consume, is set to 5000 kilobytes which can be reset at a later time. This volume will be created on the server wtr05319 partition /vicepa. After specifying this information, click the **OK** button to create the volume.

When you return to the AFS Server Manager, select the **Volumes** tab. Notice that the user volume appears in the tree structure:

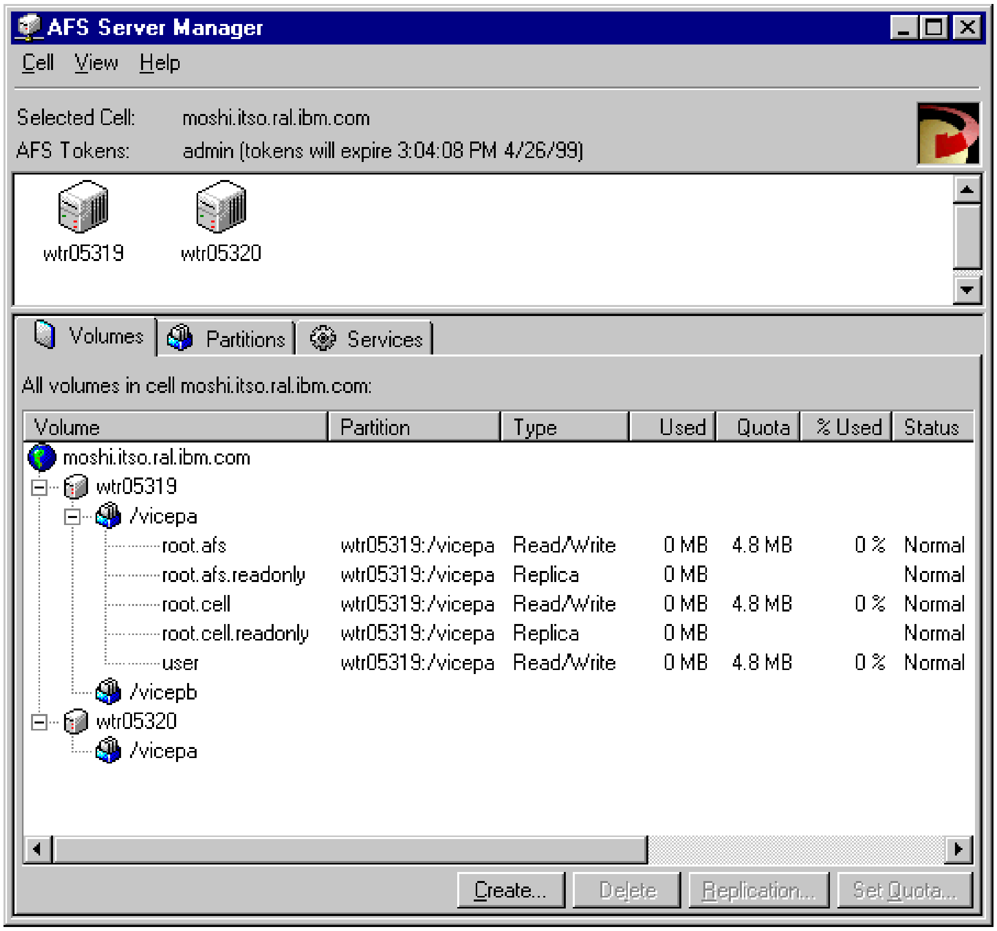

Selected Cell: moshi.itso.ral.ibm.com
AFS Tokens: admin (tokens will expire 3:04:08 PM 4/26/99)

All volumes in cell moshi.itso.ral.ibm.com:

Volume	Partition	Type	Used	Quota	% Used	Status
moshi.itso.ral.ibm.com						
wtr05319						
/vicepa						
root.afs	wtr05319:/vicepa	Read/Write	0 MB	4.8 MB	0 %	Normal
root.afs.readonly	wtr05319:/vicepa	Replica	0 MB			Normal
root.cell	wtr05319:/vicepa	Read/Write	0 MB	4.8 MB	0 %	Normal
root.cell.readonly	wtr05319:/vicepa	Replica	0 MB			Normal
user	wtr05319:/vicepa	Read/Write	0 MB	4.8 MB	0 %	Normal
/vicepb						
wtr05320						
/vicepa						

Figure 238. Displaying AFS Volumes

Next, we create a volume called *user.bobb*y on the server wtr05319 partition /vicepa. From the Server Manager’s Volumes tab, we click the **Create...** button to access the Create Volume window again:

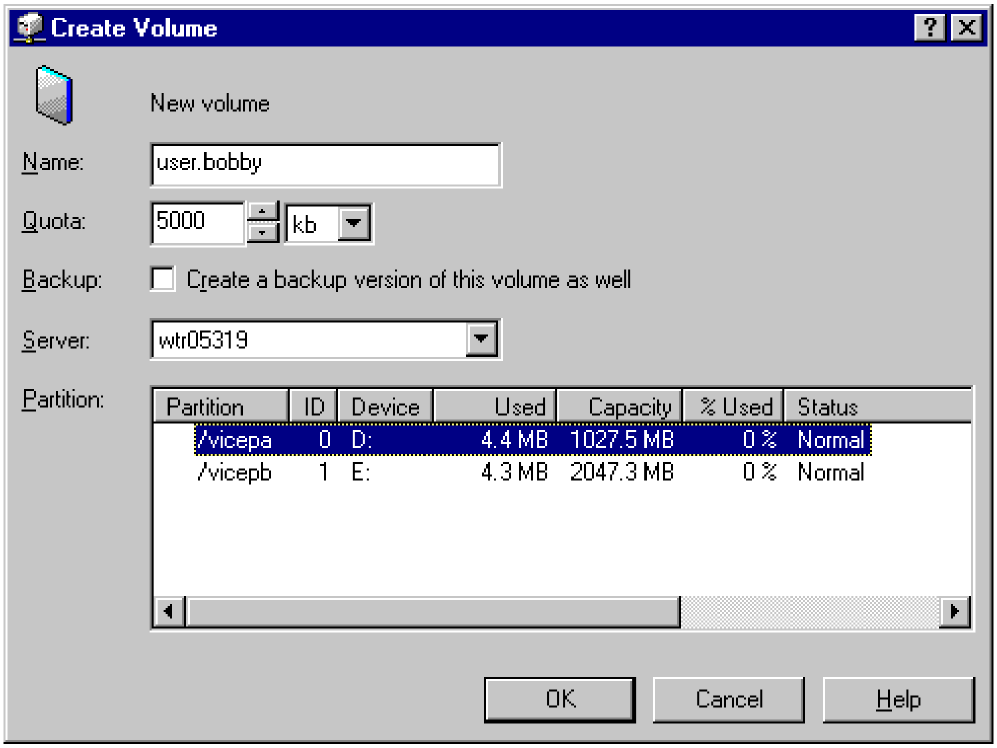

Name: user.bobby

Quota: 5000 kb

Backup: □ Create a backup version of this volume as well

Server: wtr05319

Partition:

Partition	ID	Device	Used	Capacity	% Used	Status
/vicepa	0	D:	4.4 MB	1027.5 MB	0 %	Normal
/vicepb	1	E:	4.3 MB	2047.3 MB	0 %	Normal

Figure 239. Creating the user.bobby Volume

After specifying the information for this volume, click the **OK** button. The *user.bobby* volume now appears in the tree structure in the Server Manager's Volumes tab, as shown in the following figure:

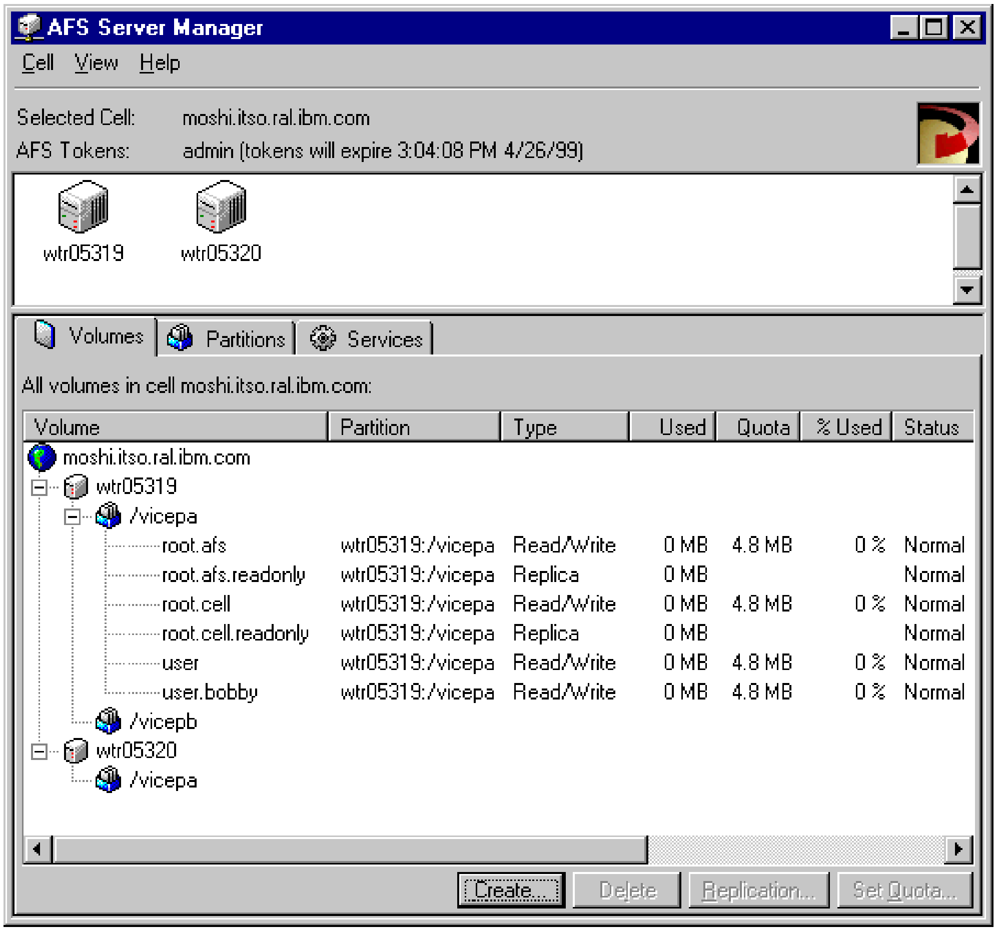

Selected Cell: moshi.itso.ral.ibm.com
AFS Tokens: admin (tokens will expire 3:04:08 PM 4/26/99)

All volumes in cell moshi.itso.ral.ibm.com:

Volume	Partition	Type	Used	Quota	% Used	Status
moshi.itso.ral.ibm.com						
wtr05319						
/vicepa						
root.afs	wtr05319:/vicepa	Read/Write	0 MB	4.8 MB	0 %	Normal
root.afs.readonly	wtr05319:/vicepa	Replica	0 MB			Normal
root.cell	wtr05319:/vicepa	Read/Write	0 MB	4.8 MB	0 %	Normal
root.cell.readonly	wtr05319:/vicepa	Replica	0 MB			Normal
user	wtr05319:/vicepa	Read/Write	0 MB	4.8 MB	0 %	Normal
user.bobby	wtr05319:/vicepa	Read/Write	0 MB	4.8 MB	0 %	Normal
/vicepb						
wtr05320						
/vicepa						

Figure 240. The user.bobby Volume

To create volumes in AIX, Solaris and the Windows NT Command Prompt, use the command:

```
vos create
```

We first use the command:

```
vos partinfo
```

to look for open space on the file server. A sample scenario is shown in the following session screen:

```
C:\>vos partinfo wtr05319
Free space on partition /vicepa: 1047610 K blocks out of total 1052256
Free space on partition /vicepb: 2092044 K blocks out of total 2096482

C:\>vos create -server wtr05319 -partition vicepb -name user.tom -maxquota 7000
Volume 536870940 created on partition /vicepb of wtr05319
```

6.4.3 Creating AFS Mount Points

Once volumes have been created, they must be mounted into the AFS directory tree. Volumes themselves contain at least a root directory, and mount points are really nothing more than pointers to this root directory of the named volume.

When traversing AFS directories and crossing mount points, it is the job of a client's Cache Manager to consult the VLDB to learn the location of the particular volume to which a given mount point refers.

In the following examples we will be creating mount points for the user and user.bobby volumes we created. However the volumes for /afs and for /afs/moshi.itso.ral.ibm.com directories (root.afs and root.cell respectively) have already been replicated by the WebSphere Performance Pack installation procedure, and so we specifically have to mount our new volumes down the ReadWrite path. Recall from the discussion in 2.11.3, "Volume Traversal Rules" on page 47 that a Cache Manager will always prefer a ReadOnly version of a volume. If we want to make any changes to root.afs and root.cell (such as creating mount points or subdirectories within them), because they are already replicated we must specifically access the ReadWrite versions of these volumes.

The mount point for the explicit ReadWrite mount point of root.cell in our example is .moshi.itso.ral.ibm.com. Note the convention of prefacing the mount point with a dot (.).

On Windows NT, to create an AFS mount point using Windows NT Explorer, select the directory in which the mount point is to be created, right-click to bring up the pop-up menu, and select sequentially **AFS**, **Mount Point**, and then **Make...**, as shown in the following figure:

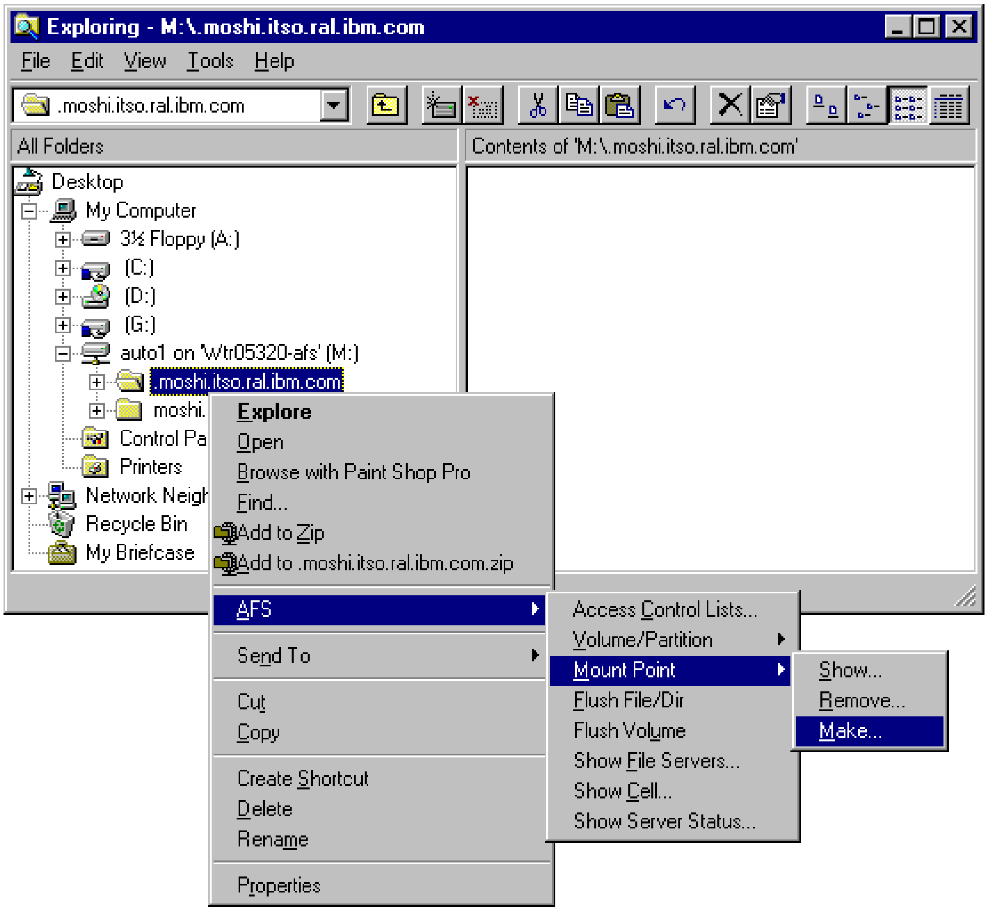

Figure 241. Making a Mount Point in Explorer

This will bring up the Make Mount Point dialog box:

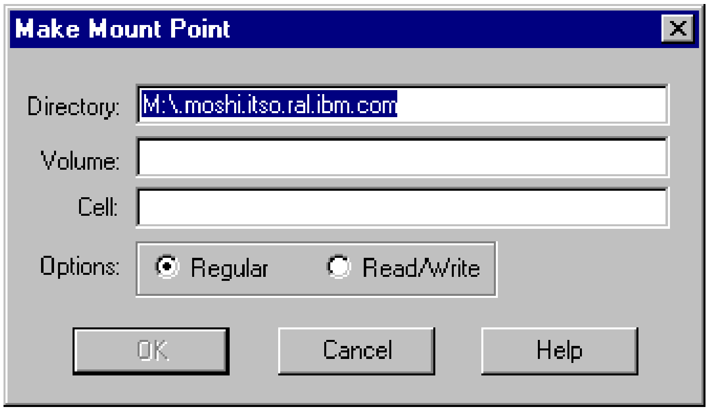

Directory: M:\.moshi.itso.ral.ibm.com

Volume:

Cell:

Options: ● Regular ○ Read/Write

Figure 242. Making Mount Point Window

The Directory field must be filled in with the location of the mount point, which in this case is the directory we had initially selected in the Windows NT Explorer window. The Volume field should contain the name of the volume being mounted into the AFS directory tree by the mount point. The Cell field will indicate the volume's cell location (which defaults to the workstation's home cell). Almost invariably Regular will be the desired mount option; see 2.11.2, "Types of Volumes" on page 47 and 2.11.3, "Volume Traversal Rules" on page 47 for when it is appropriate to set the Read/Write option.

We enter the volume name `user`, and because it is in the local cell and is a regular mount point we accept the defaults, as shown in the following figure, and click **OK**:

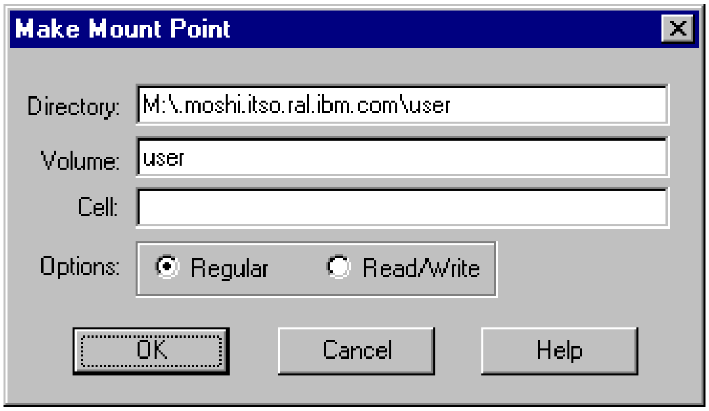

Directory: M:\.moshi.itso.ral.ibm.com\user

Volume: user

Cell:

Options: ● Regular ○ Read/Write

Figure 243. Configuring the Mount Point

In Windows Explorer we now see the volume user appear as a subdirectory below .moshi.itso.ral.ibm.com, and so the volume is mounted:

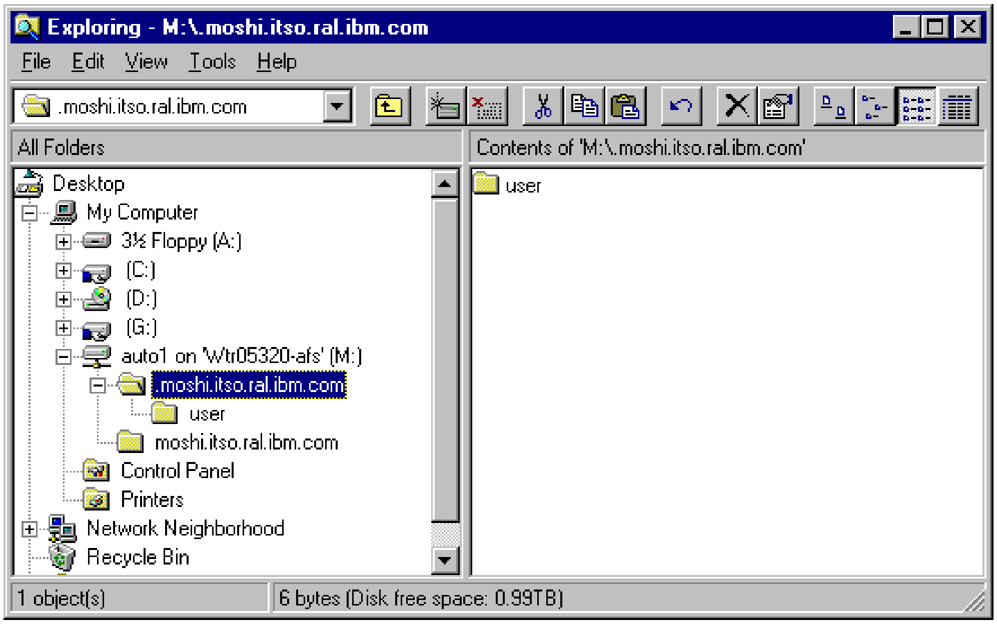

Figure 244. Exploring .moshi.itso.ral.ibm.com

Note however that user does *not* yet appear under moshi.itso.ral.ibm.com. Again this is because moshi.itso.ral.ibm.com and .moshi.itso.ral.ibm.com are different mount points for the same volume, root.cell. In detail, .moshi.itso.ral.ibm.com is the explicitly mounted ReadWrite version of root.cell, whereas moshi.itso.ral.ibm.com is a regular mount point for root.cell. And because root.cell is replicated, a client's Cache Manager will actually obtain data for moshi.itso.ral.ibm.com from root.cell.readonly. In order for our mount point to show up under moshi.itso.ral.ibm.com, we will eventually need to *release* our changes from the ReadWrite volume to the replicas.

First, however, we will create another mount point. This time we want to mount the volume user.bobby within the volume user. To do this, we right-click **user** in the Windows NT Explorer window, and from the pop-up menu, we select **AFS**, **Mount Point** and **Make...**, as shown in the following figure:

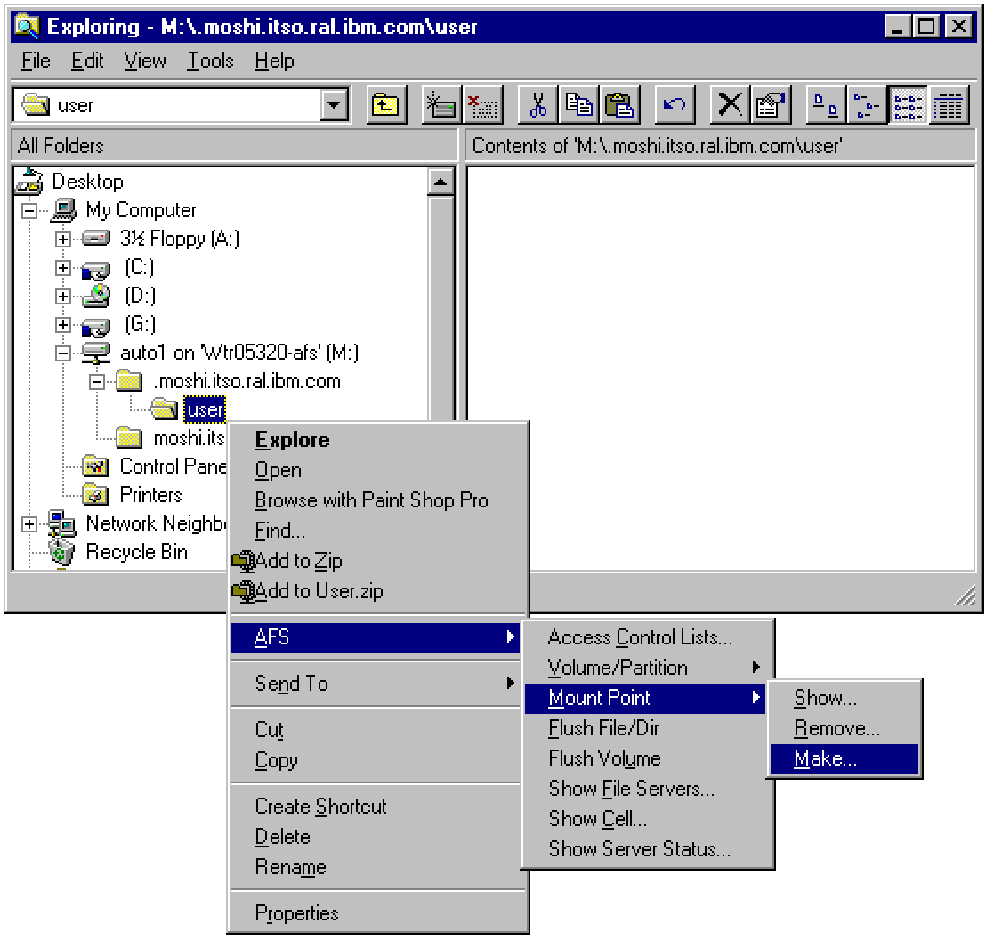

Figure 245. Creating a Mount Point in Explorer

This command brings up the Make Mount Point dialog box, which we fill in as indicated in the following figure:

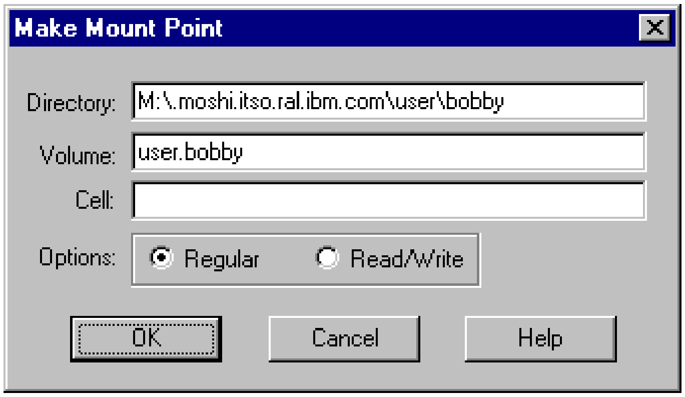

Make Mount Point

Directory: M:\.moshi.itso.ral.ibm.com\user\bobby

Volume: user.bobby

Cell:

Options: ● Regular ○ Read/Write

Figure 246. Making a Mount Point for user.bobby

In the dialog box we need to specify the directory location of the mount point and the volume name. Note that this mount point differs slightly from the mount point we created for user in that the volume name does not *exactly* correspond with the mount point name. That is, we are creating a mount point called bobby for the volume user.bobby. There need not be any correlation between the mount point name and the volume name, and we could just as easily have created a mount point called phambo for user.bobby. However, for administrative reasons, it's generally best to keep some correspondence between volume and mount point names.

We now see that the mount point bobby appears as a subdirectory under user:

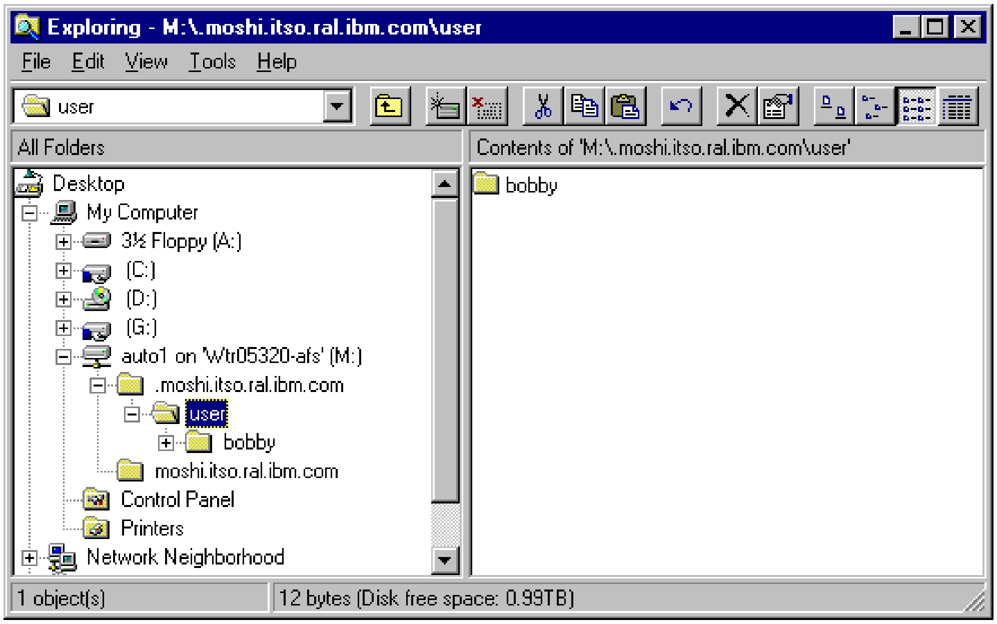

Figure 247. The bobby Mount Point

Again note that neither user nor bobby appear yet under moshi.itso.ral.ibm.com.

To create mount points in AIX, Solaris and the Windows NT Command Prompt use the command:

```
fs mkmount
```

The command:

```
fs lsmount
```

can be used to examine mount points.

In this example we attempt to create a mount point for the volume user.tom but we receive an error. This is because both root.cell and user are replicated, and so our client is pointed to a ReadOnly volume. Next we try to make the mount point down the ReadWrite path (.moshi.itso.ral.ibm.com) and the command succeeds. The entire scenario is shown in the following session screen:

```
C:\> fs mkmount -dir m:\moshi.itso.ral.ibm.com\user\tom -vol user.tom
fs: You can not change a backup or readonly volume

C:\> fs mkmount -dir m:\.moshi.itso.ral.ibm.com\user\tom -vol user.tom
```

To verify that the command has succeeded we do a directory listing down the ReadWrite path to where we installed the mount point.

```
C:\>dir m:\.moshi.itso.ral.ibm.com\user
 Volume in drive M is AFS
 Volume Serial Number is 0000-04D2

 Directory of m:\.moshi.itso.ral.ibm.com\user

06/08/99  07:03p        <DIR>          .
06/08/99  07:03p        <DIR>          ..
04/25/99  06:54p        <DIR>          bobby
06/08/99  07:03p        <DIR>          tom
               4 File(s)          4,118 bytes
                      1,099,511,626,752 bytes free
```

Now we do a directory listing down the normal traversal path and see that our new mount point has not appeared yet. The mount point for user.tom has been created in the ReadWrite version of user, but this change has not yet been pushed out to the replica(s) (user.readonly):

```
C:\>dir m:\moshi.itso.ral.ibm.com\user
 Volume in drive M is AFS
 Volume Serial Number is 0000-04D2

 Directory of m:\moshi.itso.ral.ibm.com\user

04/25/99  06:54p        <DIR>          .
04/25/99  06:54p        <DIR>          ..
04/25/99  06:54p        <DIR>          bobby
               3 File(s)          4,108 bytes
                      1,099,511,626,752 bytes free
```

We will cover this in 6.4.6, “Releasing AFS Volumes” on page 293.

Finally, as an exercise we look at the mount points for volumes root.cell and user. Notice that the volumes referenced by \moshi.itso.ral.ibm.com and \moshi.itso.ral.ibm.com\user are prefaced with the pound sign (`#`). This symbol indicates that this mount point is a normal mount point, which means that it will supply a ReadOnly volume if it can and a ReadWrite if it cannot. Conversely, the percent symbol (`%`) – used for root.cell, mounted at .moshi.itso.ral.ibm.com – is a ReadWrite mount point and it will *only* supply a ReadWrite volume.

The following session screen shows the entire scenario:

```
C:\>fs lsmount m:\moshi.itso.ral.ibm.com
'm:\moshi.itso.ral.ibm.com' is a mount point for volume '#root.cell'

C:\>fs lsmount m:\.moshi.itso.ral.ibm.com
'm:\.moshi.itso.ral.ibm.com' is a mount point for volume '%root.cell'

C:\>fs lsmount m:\moshi.itso.ral.ibm.com\user
'm:\moshi.itso.ral.ibm.com\user' is a mount point for volume '#user'
```

6.4.4 Setting AFS Access Control Lists

In our examples for Windows NT, we have so far created two volumes and mounted them within our AFS directory tree. Typically (though not necessarily) an administrator will next set ACLs on the newly created volumes in Windows NT Explorer. In the next demonstration we will set an ACL on the user mount point but ignore the bobby mount point for the sake of expediency. See 2.6.6, “Protecting AFS Files and Directories” on page 38 for a brief discussion on ACL permissions in AFS.

In Windows NT Explorer, to set an ACL on a directory in AFS (either a mount point or a normal subdirectory) select the directory and right-click to access the AFS menu. Then click **Access Control Lists...**, as indicated in the following figure:

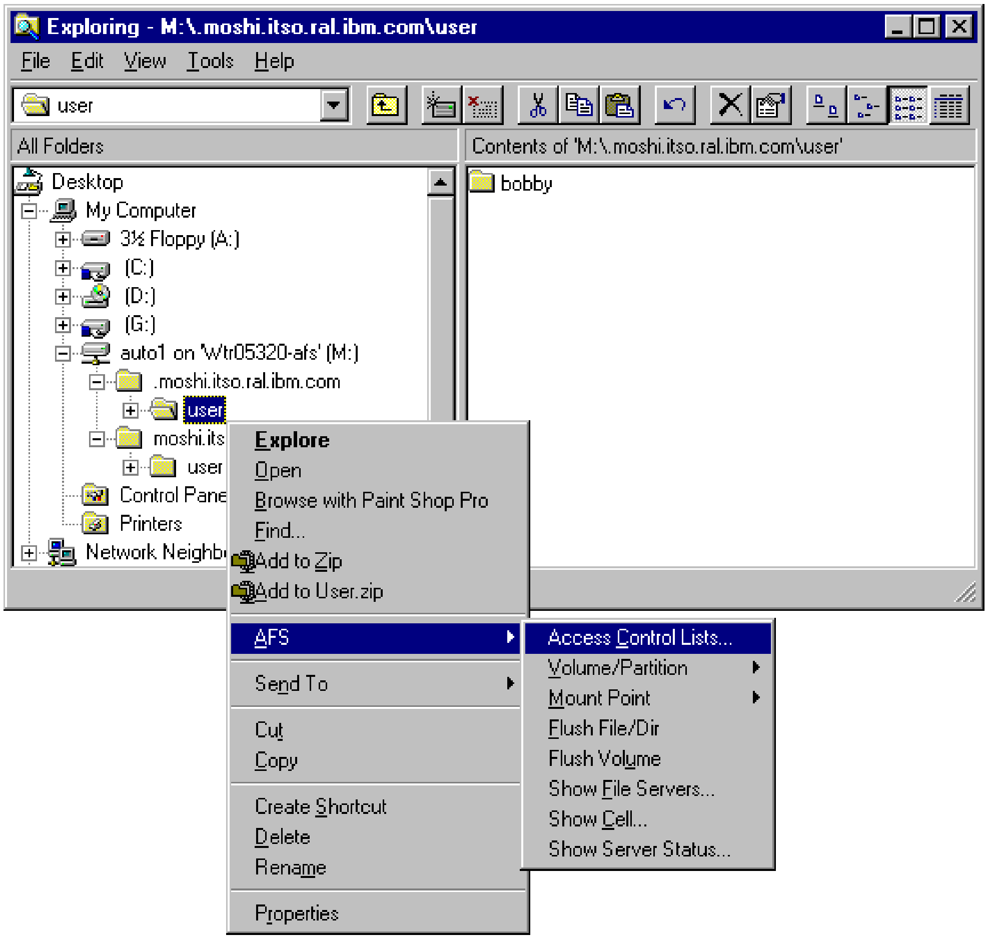

Figure 248. Selecting Access Control Lists Command

This will bring up the Set AFS ACL dialog box:

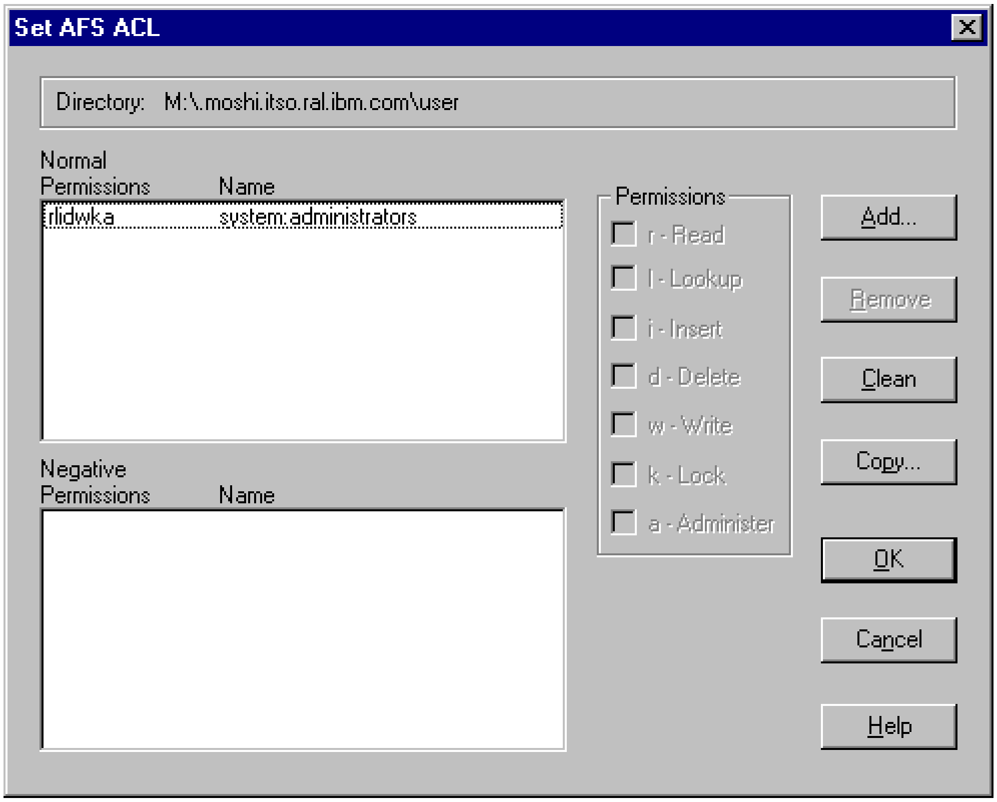

Directory: M:\.moshi.itso.ral.ibm.com\user

Normal

Permissions	Name
rlidwka	system:administrators

Negative

Permissions	Name

Permissions

☐ r - Read
☐ l - Lookup
☐ i - Insert
☐ d - Delete
☐ w - Write
☐ k - Lock
☐ a - Administer

Figure 249. The ACL Window

In this example, we are looking at the ACL set for the user subdirectory, which again is a mount point for the volume user.

As you can see, the only permissions set for this directory are currently for the group system:administrators to have full permissions. To change this particular entry, click it and then subtract or add permissions under the Permissions box. Because this is a new volume, however, we are more interested in adding a new ACL entry, so we click **Add...** to bring up the Add ACL Entry dialog box, as shown in the following figure:

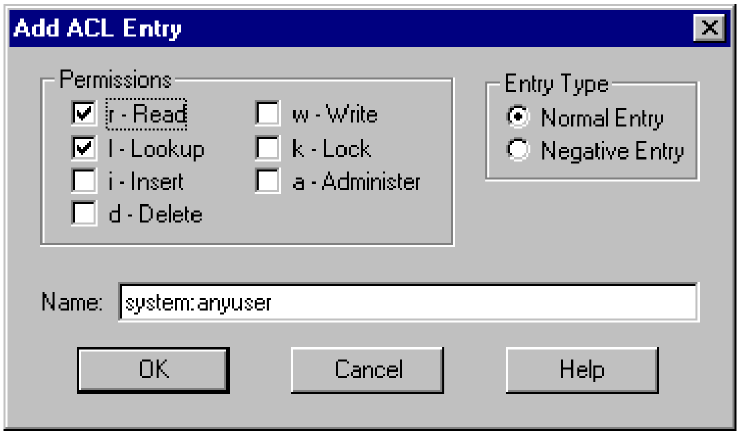

Permissions

☑ r - Read
☑ l - Lookup
☐ i - Insert
☐ d - Delete
☐ w - Write
☐ k - Lock
☐ a - Administer

Entry Type

● Normal Entry
○ Negative Entry

Name: system:anyuser

Figure 250. Adding an ACL Entry

Enter the name of the group or the user at the bottom of the dialog box and select which permissions you want this entry to have. In our example we give the group system:anyuser (which is essentially the default for *anyone*, authenticated or not, accessing this directory) read (r) and lookup (l) permissions (see 2.6.6, “Protecting AFS Files and Directories” on page 38). Then we press the **OK** button. Clicking **OK** returns us to the original Set AFS ACL window, displaying the new ACL entry:

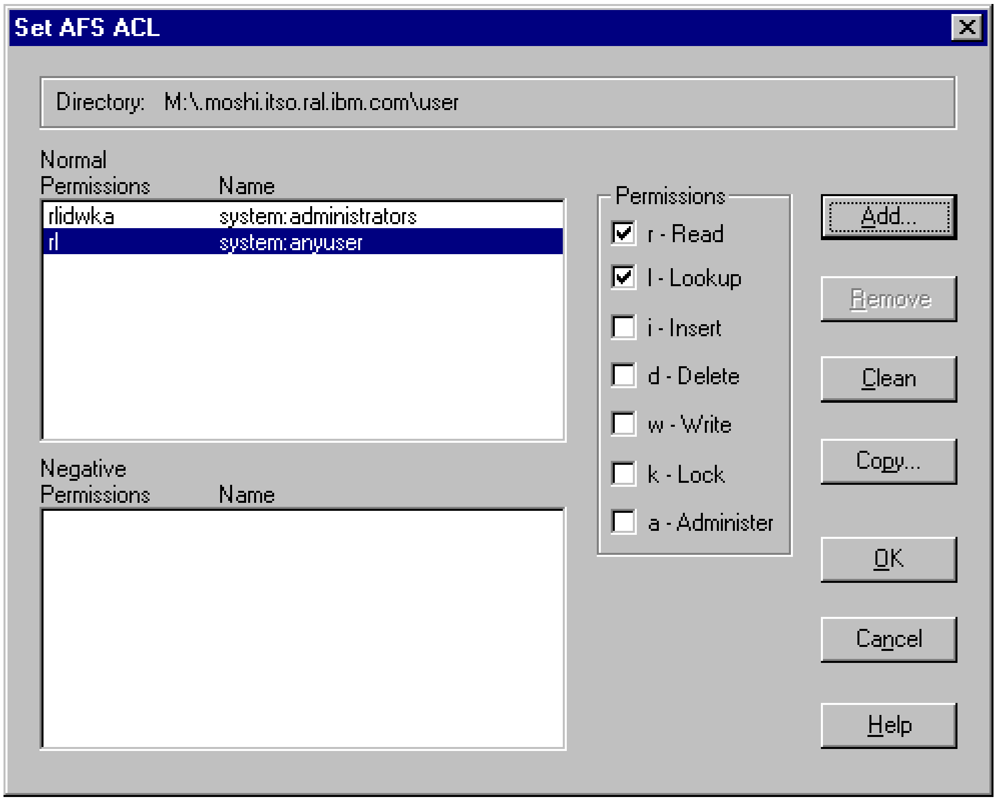

Directory: M:\.moshi.itso.ral.ibm.com\user

Normal Permissions	Name
rlidwka	system:administrators
rl	system:anyuser

Negative Permissions	Name

Permissions

☑ r - Read
☑ l - Lookup
☐ i - Insert
☐ d - Delete
☐ w - Write
☐ k - Lock
☐ a - Administer

Figure 251. Added ACL Entry

Click **OK** when finished making changes to the ACL. Note that *negative permissions* are a denial of specific permissions.

In AIX, Solaris, and the Windows NT Command Prompt, to manage AFS ACLs, use the commands:

```
fs listacl
fs setacl
```

In the following example, we examine the default ACL currently set on \moshi.itso.ral.ibm.com\user\tom, and then grant user tom full access and deny user bobby any access:

```
C:\> fs listacl m:\moshi.itso.ral.ibm.com\user\tom
Access list for m:\moshi.itso.ral.ibm.com\user\tom is
Normal rights:
  system:administrators rlidwka

C:\> fs setacl m:\moshi.itso.ral.ibm.com\user\tom -acl tom all

C:\> fs setacl m:\moshi.itso.ral.ibm.com\user\tom -acl bobby all -

C:\> fs listacl m:\moshi.itso.ral.ibm.com\user\tom
Access list for m:\moshi.itso.ral.ibm.com\user\tom is
Normal rights:
  system:administrators rlidwka
  tom rlidwka
Negative rights:
  bobby rlidwka
```

6.4.5 Replicating AFS Volumes

AFS volumes that are read often but in which the data seldom changes are replicated across multiple file servers for high availability and for load balancing. For example a volume that contains software that is used by many clients but is only changed when new versions are released should be replicated, whereas user volumes that are used throughout the day would not be good candidates for AFS replication.

In our examples we have created several volumes, mounted them within the ReadWrite version of the root.cell volume, and have also set ACLs. We now consider whether the volumes should be replicated and whether we need to release updates to existing replicas. We created the user volume and mounted it specifically as an intermediary between the moshi.itso.ral.ibm.com directory (which is a mount point for root.cell) and volumes such as user.bobby. This implies that the user volume itself changes only when we add new mount points to it, and so it is a candidate for replication. The volume user.bobby, on the other hand, is probably going to be associated with an AFS user ID and so will be used frequently. If this usage involves writing data, then the volume does not lend itself to replication, and so we choose not to replicate it.

To replicate an AFS volume in Windows NT, launch the AFS Server Manager. Select the AFS volume to be replicated by highlighting it. Position the pointer on the name of the selected AFS volume and click the right mouse button to display a pop-up menu. From the pop-up menu, select **Replication...**, as shown in the following window:

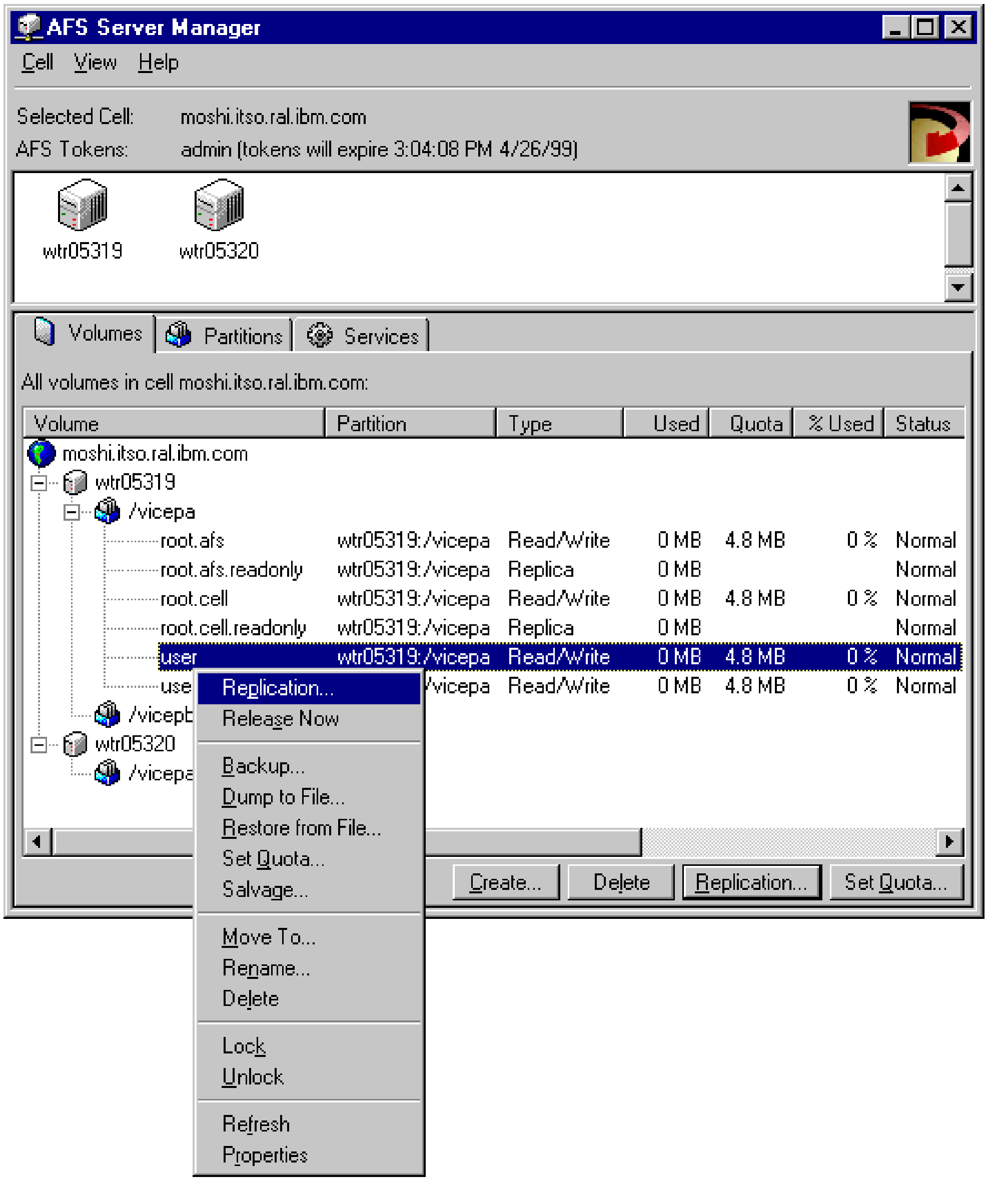

Selected Cell: moshi.itso.ral.ibm.com
AFS Tokens: admin (tokens will expire 3:04:08 PM 4/26/99)

All volumes in cell moshi.itso.ral.ibm.com:

Volume	Partition	Type	Used	Quota	% Used	Status
moshi.itso.ral.ibm.com						
wtr05319						
/vicepa						
root.afs	wtr05319:/vicepa	Read/Write	0 MB	4.8 MB	0 %	Normal
root.afs.readonly	wtr05319:/vicepa	Replica	0 MB			Normal
root.cell	wtr05319:/vicepa	Read/Write	0 MB	4.8 MB	0 %	Normal
root.cell.readonly	wtr05319:/vicepa	Replica	0 MB			Normal
user	wtr05319:/vicepa	Read/Write	0 MB	4.8 MB	0 %	Normal
use	/vicepa	Read/Write	0 MB	4.8 MB	0 %	Normal
/vicepb						
wtr05320						
/vicepa						

Figure 252. Administrative Options to AFS Volumes

The Replication window is displayed. In the Replication window, click the **Create...** button to add a replication site. The Replicate Volume window appears, as shown in the following figure:

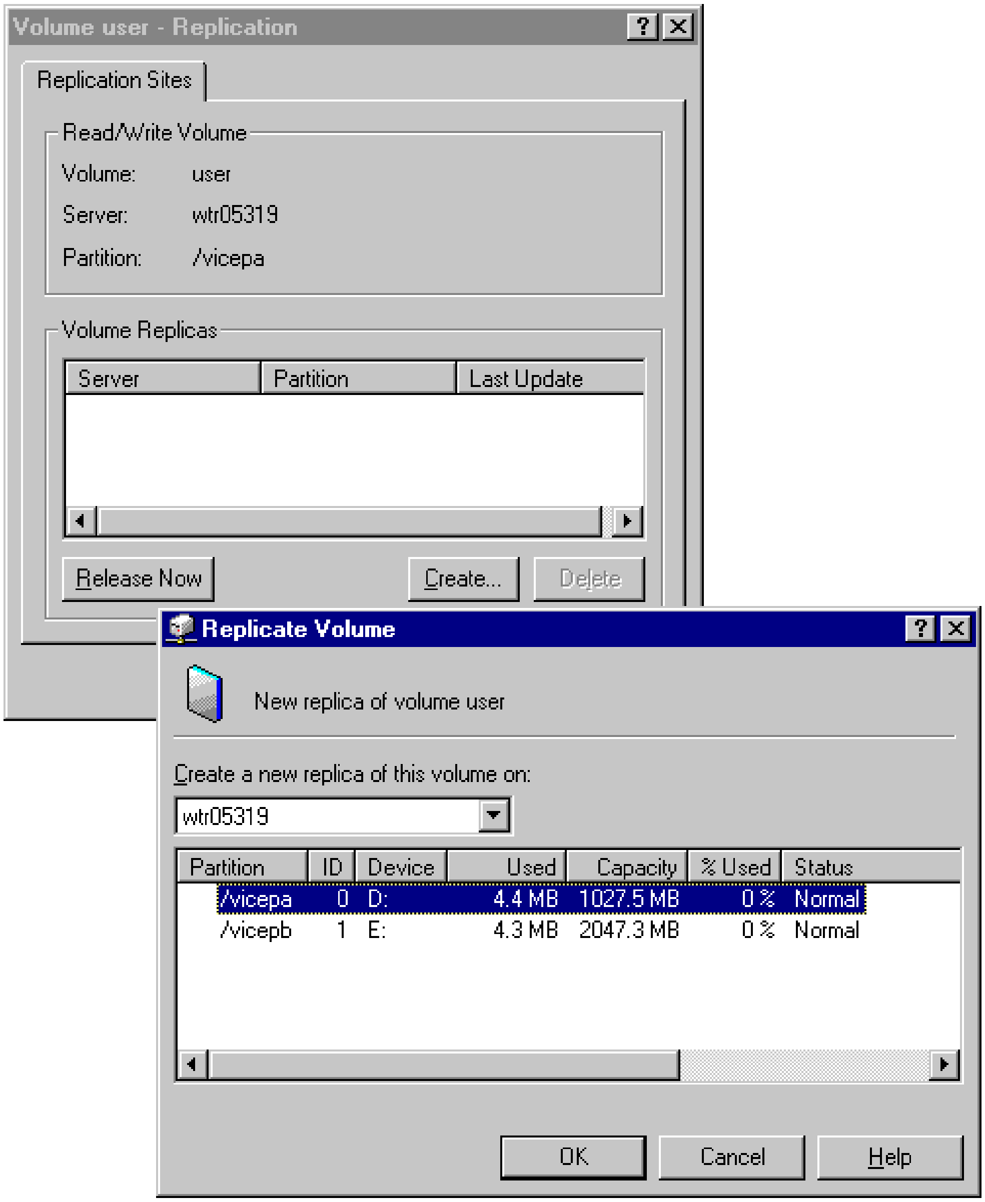

Volume: user

Server: wtr05319

Partition: /vicepa

Server	Partition	Last Update

New replica of volume user

Create a new replica of this volume on: wtr05319

Partition	ID	Device	Used	Capacity	% Used	Status
/vicepa	0	D:	4.4 MB	1027.5 MB	0 %	Normal
/vicepb	1	E:	4.3 MB	2047.3 MB	0 %	Normal

Figure 253. Replicating an AFS Volume

Use the pull-down menu to specify the name of the server from the list of available file servers in the cell (Windows NT or otherwise). Select the name of the partition on that server where you would like to create the replica. After you have specified the server and partition, click **OK** to return to the Replication window:

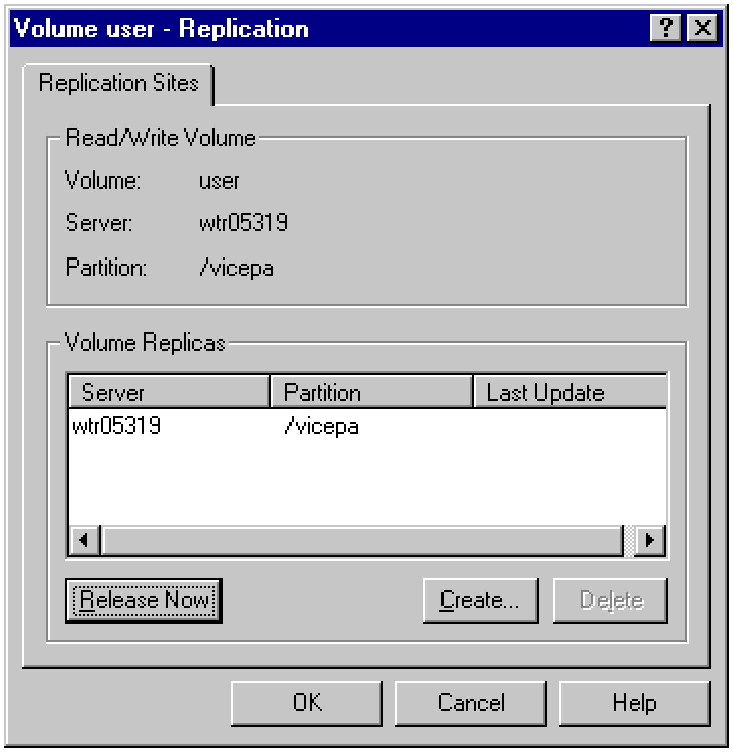
Volume user - Replication

Replication Sites

Read/Write Volume

Volume: user

Server: wtr05319

Partition: /vicepa

Volume Replicas

Server	Partition	Last Update
wtr05319	/vicepa	

Figure 254. Displaying the Replication Sites

The Replication window now includes information about the replication site. This is information that is stored in the VLDB. To actually create the ReadOnly volume on the server, you must *release* the volume. Releasing a volume pushes changes made from a ReadWrite volume to the ReadOnly versions.

Click the **Release Now** button to create the volume. The Release Volume window will appear:

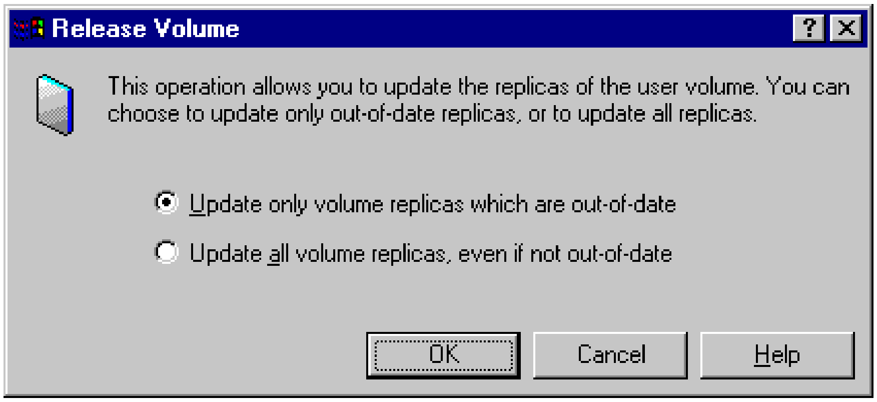
Release Volume

This operation allows you to update the replicas of the user volume. You can choose to update only out-of-date replicas, or to update all replicas.

● Update only volume replicas which are out-of-date

○ Update all volume replicas, even if not out-of-date

Figure 255. Releasing the AFS Volume

In the Release Volume window, you must specify whether you want to update all replicas, or only those that are out-of-date. In this example, we select the check box to indicate that we want to update only those replicas which are out-of-date. The difference between these options is that updating out-of-date replicas pushes data only to those replicas that did not receive an update (this means that they are new or were offline at the last update), whereas updating all replicas exhaustively pushes out data to all replicas without any distinction. Typically you would update all replicas when making a change to the ReadWrite, and use out-of-date for adding new replicas.

When we return to the AFS Server Manager's Volume tab, notice that the user.readonly volume now appears:

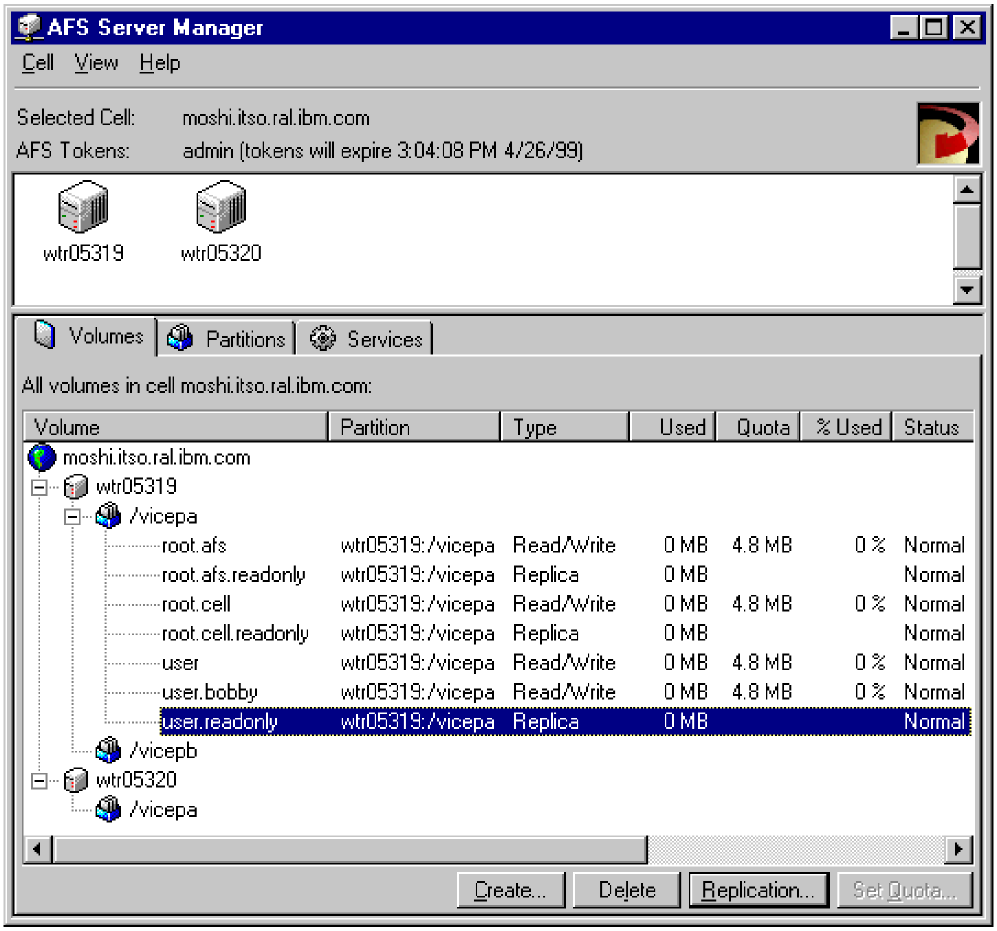

Selected Cell: moshi.itso.ral.ibm.com
AFS Tokens: admin (tokens will expire 3:04:08 PM 4/26/99)

All volumes in cell moshi.itso.ral.ibm.com:

Volume	Partition	Type	Used	Quota	% Used	Status
moshi.itso.ral.ibm.com						
wtr05319						
/vicepa						
root.afs	wtr05319:/vicepa	Read/Write	0 MB	4.8 MB	0 %	Normal
root.afs.readonly	wtr05319:/vicepa	Replica	0 MB			Normal
root.cell	wtr05319:/vicepa	Read/Write	0 MB	4.8 MB	0 %	Normal
root.cell.readonly	wtr05319:/vicepa	Replica	0 MB			Normal
user	wtr05319:/vicepa	Read/Write	0 MB	4.8 MB	0 %	Normal
user.bobby	wtr05319:/vicepa	Read/Write	0 MB	4.8 MB	0 %	Normal
user.readonly	wtr05319:/vicepa	Replica	0 MB			Normal
/vicepb						
wtr05320						
/vicepa						

Figure 256. Displaying the AFS Volumes

This is the replica that we just created. Note however that this replica is not yet accessible to clients via Windows NT Explorer or the command line, as its parent AFS directory (moshi.itso.ral.ibm.com) has not updated its own replicas yet.

To create replicated volumes in AIX, Solaris, or the Windows NT Command Prompt, use the command:

```
vos addsite
```

The following session screen shows a sample scenario:

```
C:\>vos addsite -server wtr05320 -partition vicepa -id root.cell
Added replication site wtr05320 /vicepa for volume root.cell

C:\>vos addsite -server wtr05320 -partition vicepa -id root.afs -verbose
Adding a new site ... done
Added replication site wtr05320 /vicepa for volume root.afs
```

After adding the replication sites, we still need to release the volumes. We cover the way to do this in the next section.

6.4.6 Releasing AFS Volumes

As seen in the previous sections, when a new replica of an existing volume is created, the volume itself must be released before the replica is seen. This is also true of any changes to data in replicated volumes, such as adding or deleting files and directories, adding or changing ACLs, or even creating and deleting mount points. The changes are made to the original ReadWrite volume and then the changes are released to the replicated versions of the volume. In our example, not only did we replicate the user volume but we had also created the mount point for user in root.cell, which itself is a replicated volume. Thus, we need to release our changes to the replicas for both root.cell and user.

In this example we know we have made changes in the root.cell volume: we have added the mount point for user, set an ACL on it, and created mount points within user. Before these changes will be seen by AFS clients, however, they must be released to the replicas of the root.cell volume.

To release changes to replicas in Windows NT, in the Server Manager select the ReadWrite volume to be released and right-click to bring up the volume's pop-up menu. Then, select **Release Now**, as shown in the following figure:

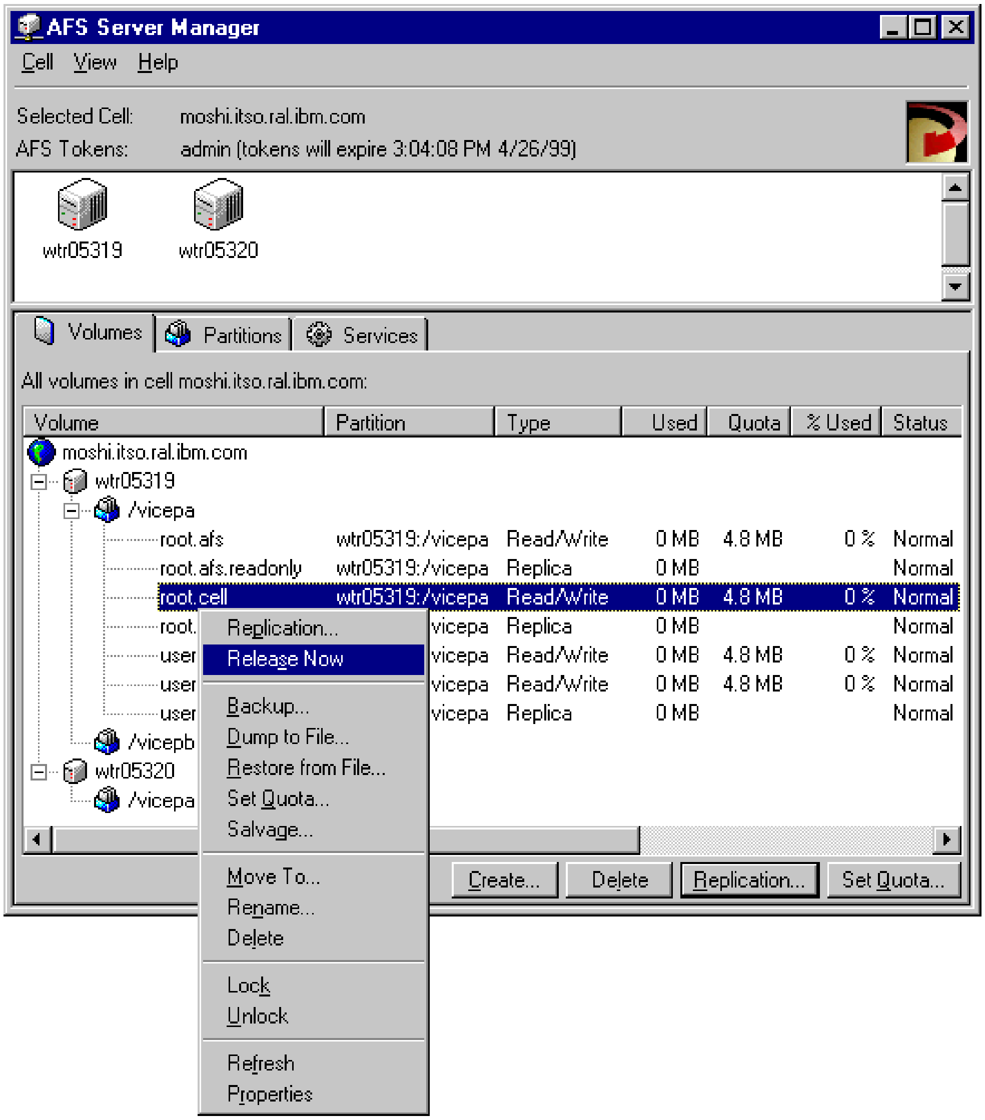

Selected Cell: moshi.itso.ral.ibm.com
AFS Tokens: admin (tokens will expire 3:04:08 PM 4/26/99)

All volumes in cell moshi.itso.ral.ibm.com:

Volume	Partition	Type	Used	Quota	% Used	Status
moshi.itso.ral.ibm.com						
wtr05319						
/vicepa						
root.afs	wtr05319:/vicepa	Read/Write	0 MB	4.8 MB	0 %	Normal
root.afs.readonly	wtr05319:/vicepa	Replica	0 MB			Normal
root.cell	wtr05319:/vicepa	Read/Write	0 MB	4.8 MB	0 %	Normal
root.	vicepa	Replica	0 MB			Normal
user	vicepa	Read/Write	0 MB	4.8 MB	0 %	Normal
user	vicepa	Read/Write	0 MB	4.8 MB	0 %	Normal
user	vicepa	Replica	0 MB			Normal
/vicepb						
wtr05320						
/vicepa						

Figure 257. Starting the Release Command

You will be prompted for the manner in which updates are released:

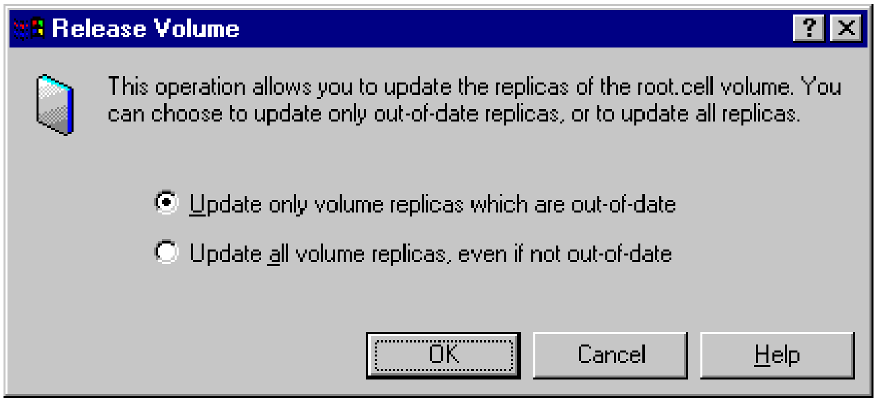
Release Volume

This operation allows you to update the replicas of the root.cell volume. You can choose to update only out-of-date replicas, or to update all replicas.

● Update only volume replicas which are out-of-date

○ Update all volume replicas, even if not out-of-date

Figure 258. Release Volume Window

We select **Update only volume replicas which are out-of-date**, which updates new or out-of-date replicas. Click **OK** when your selection is done.

Note that in Windows NT Explorer we now see the user mount point underneath moshi.itso.ral.ibm.com:

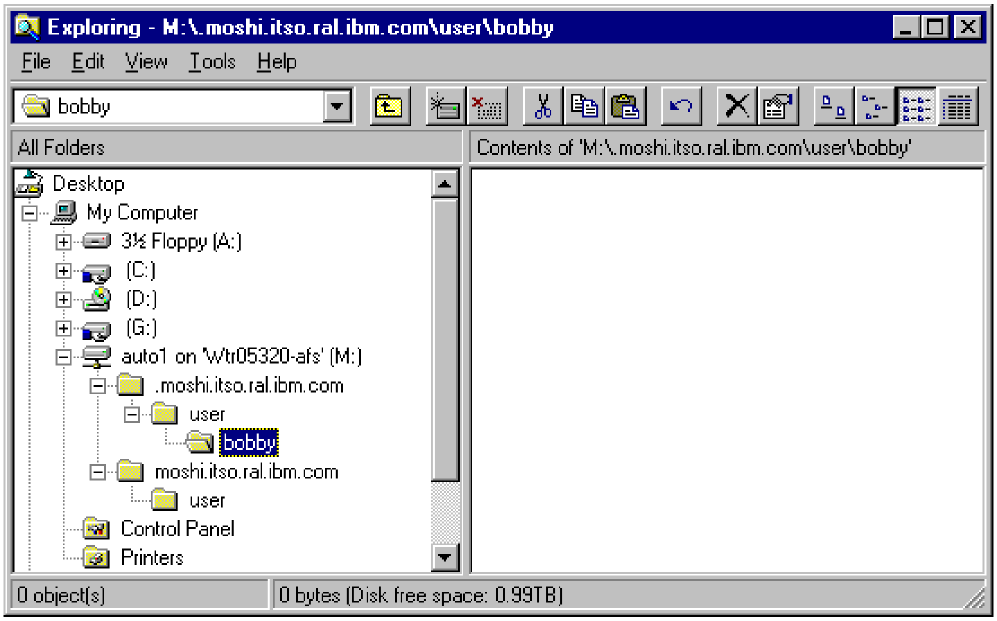
Exploring - M:\.moshi.itso.ral.ibm.com\user\bobby

Figure 259. Updated Directory Tree

Releasing the changes to the root.cell volume pushed a new copy of the volume data from root.cell to root.cell.readonly. Note however that the bobby mount point fails to show up underneath user. The reason for this is because user itself is a replicated volume, and the bobby mount point has been added to the ReadWrite version of user. Therefore, user also needs to have this change released to its replicas. For this reason, in the AFS Server Manager we select the user volume and right click to bring up the pop-up menu. Then, we select the **Release Now** command, as shown in the following screen:

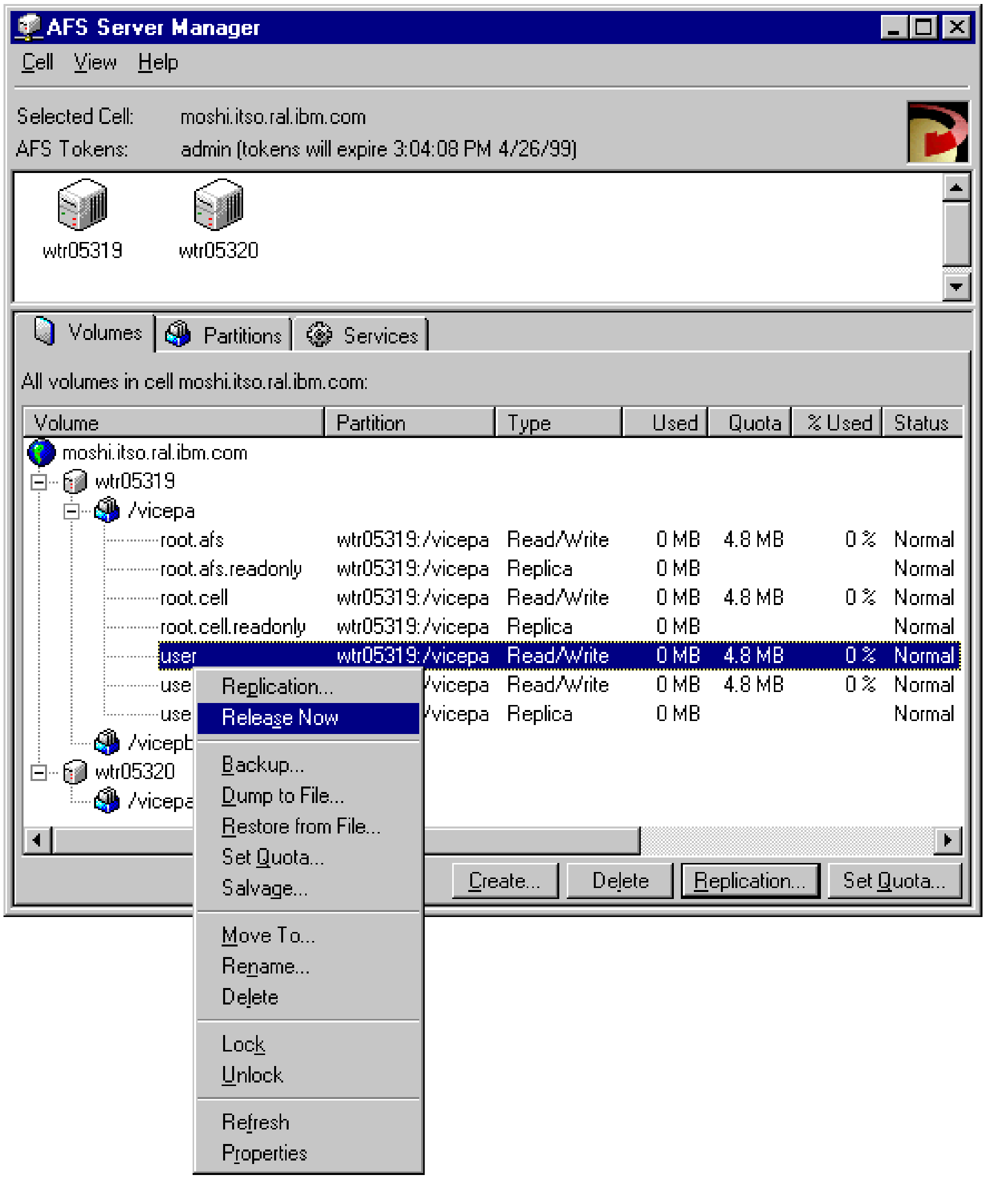

Selected Cell: moshi.itso.ral.ibm.com

AFS Tokens: admin (tokens will expire 3:04:08 PM 4/26/99)

All volumes in cell moshi.itso.ral.ibm.com:

Volume	Partition	Type	Used	Quota	% Used	Status
moshi.itso.ral.ibm.com						
wtr05319						
/vicepa						
root.afs	wtr05319:/vicepa	Read/Write	0 MB	4.8 MB	0 %	Normal
root.afs.readonly	wtr05319:/vicepa	Replica	0 MB			Normal
root.cell	wtr05319:/vicepa	Read/Write	0 MB	4.8 MB	0 %	Normal
root.cell.readonly	wtr05319:/vicepa	Replica	0 MB			Normal
user	wtr05319:/vicepa	Read/Write	0 MB	4.8 MB	0 %	Normal
use	/vicepa	Read/Write	0 MB	4.8 MB	0 %	Normal
use	/vicepa	Replica	0 MB			Normal
/vicepb						
wtr05320						
/vicepa						

Figure 260. Releasing the user Volume

This releases the volume data again to the user volume’s replicas. After releasing the volume, we now see the complete path:

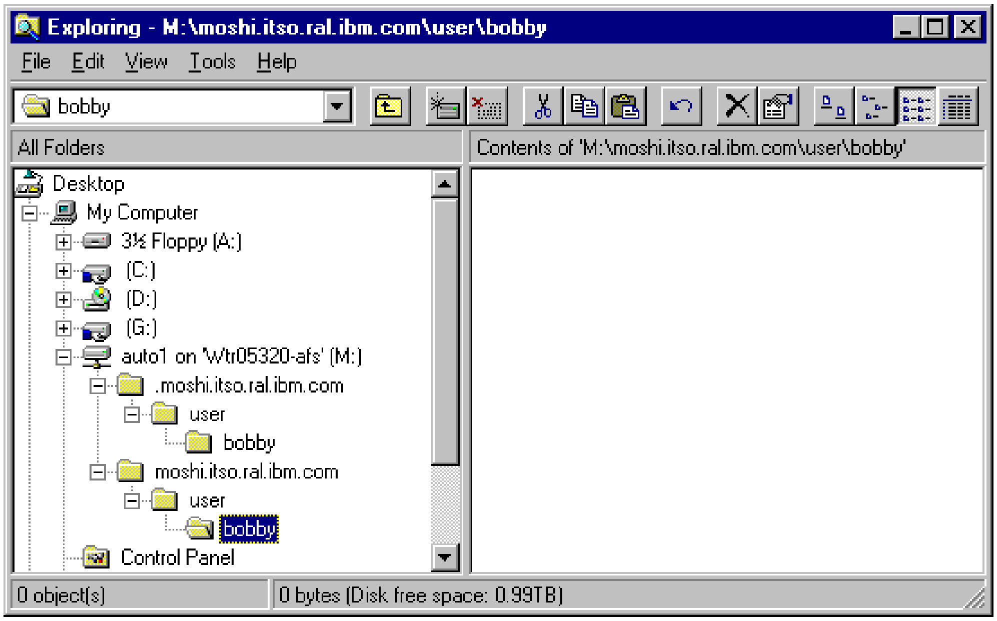
Exploring - M:\moshi.itso.ral.ibm.com\user\bobby
File Edit View Tools Help
bobby
All Folders
Contents of 'M:\moshi.itso.ral.ibm.com\user\bobby'
Desktop
My Computer
3½ Floppy (A:)
(C:)
(D:)
(G:)
auto1 on 'Wtr05320-afs' (M:)
.moshi.itso.ral.ibm.com
user
bobby
moshi.itso.ral.ibm.com
user
bobby
Control Panel
0 object(s)
0 bytes (Disk free space: 0.99TB)

Figure 261. Results of Releasing the user Volume

In AIX, Solaris and the Windows NT Command Prompt, use the command:

```
vos release
```

to release changes from a ReadWrite volume to its replicas, or to build new replicas after using the command

```
vos addsite
```

In this example we release changes to replicas of root.cell:

```
C:\> vos release -id root.cell
Released volume root.cell successfully
```

We do the same for root.afs, demonstrating the use of the `-verbose` switch, as shown in the following session screen:

```
C:\>vos release -id root.afs -verbose

root.afs
    RWrite: 536870912     ROnly: 536870913
    number of sites -> 3
       server wtr05319.itso.ral.ibm.com partition /vicepa RW Site
       server wtr05319.itso.ral.ibm.com partition /vicepa RO Site
server wtr05320.itso.ral.ibm.com partition /vicepa RO Site  -- Not released
This is a complete release of the volume 536870912
Recloning RW volume ...
Creating new volume 536870913 on replication site wtr05320.itso.ral.ibm.com: done.
Starting ForwardMulti from 536870913 to 536870913 on wtr05320.itso.ral.ibm.com (full release).
updating VLDB ... done
Released volume root.afs successfully
```

6.4.7 Moving AFS Volumes

One of AFS's important administrative advantages is that data can be transparently moved around between File Servers with no downtime and little effect on users. An administrator simply identifies the volume to be moved and its destination, and AFS handles all of the details of moving the data and handling client access to the data during the move.

To move volumes in Windows NT, start the AFS Server Manager and select the volume to be moved. Right-click to bring up the pop-up menu, and select **Move To...**. In this case we will be moving the volume user.bobby:

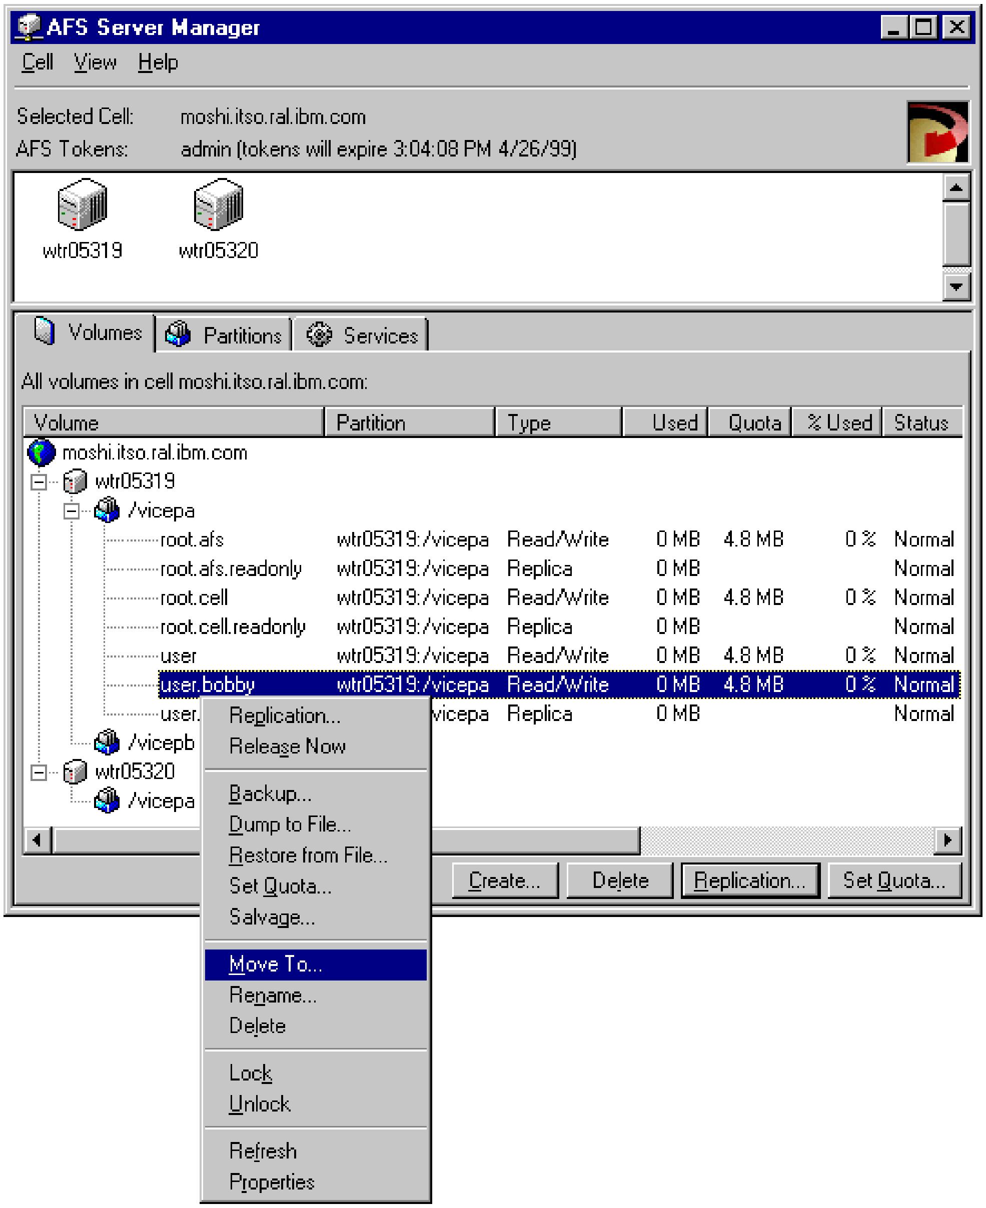

Selected Cell: moshi.itso.ral.ibm.com
AFS Tokens: admin (tokens will expire 3:04:08 PM 4/26/99)

All volumes in cell moshi.itso.ral.ibm.com:

Volume	Partition	Type	Used	Quota	% Used	Status
moshi.itso.ral.ibm.com						
wtr05319						
/vicepa						
root.afs	wtr05319:/vicepa	Read/Write	0 MB	4.8 MB	0 %	Normal
root.afs.readonly	wtr05319:/vicepa	Replica	0 MB			Normal
root.cell	wtr05319:/vicepa	Read/Write	0 MB	4.8 MB	0 %	Normal
root.cell.readonly	wtr05319:/vicepa	Replica	0 MB			Normal
user	wtr05319:/vicepa	Read/Write	0 MB	4.8 MB	0 %	Normal
user.bobby	wtr05319:/vicepa	Read/Write	0 MB	4.8 MB	0 %	Normal
user.	vicepa	Replica	0 MB			Normal
/vicepb						
wtr05320						
/vicepa						

Figure 262. Moving an AFS Volume

In the Move Volume window, select the destination File Server and the destination vice partition. In this example we are moving user.bobby to the /vicepa partition on File Server wtr05320:

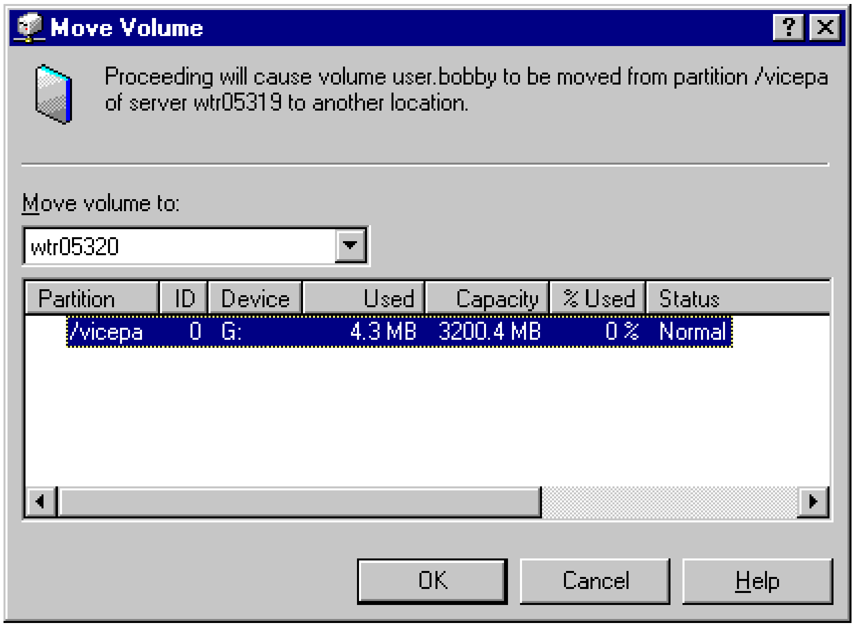

Proceeding will cause volume user.bobby to be moved from partition /vicepa of server wtr05319 to another location.

Move volume to: wtr05320

Partition	ID	Device	Used	Capacity	% Used	Status
/vicepa	0	G:	4.3 MB	3200.4 MB	0 %	Normal

Figure 263. Specifying a New AFS Server and Partition to Move an AFS Volume

Click **OK** when you are done. Depending on the size of the volume and the amount of data it contains, the move should take anywhere from a few seconds to several minutes. Note in this example that user.bobby now appears on a new file server:

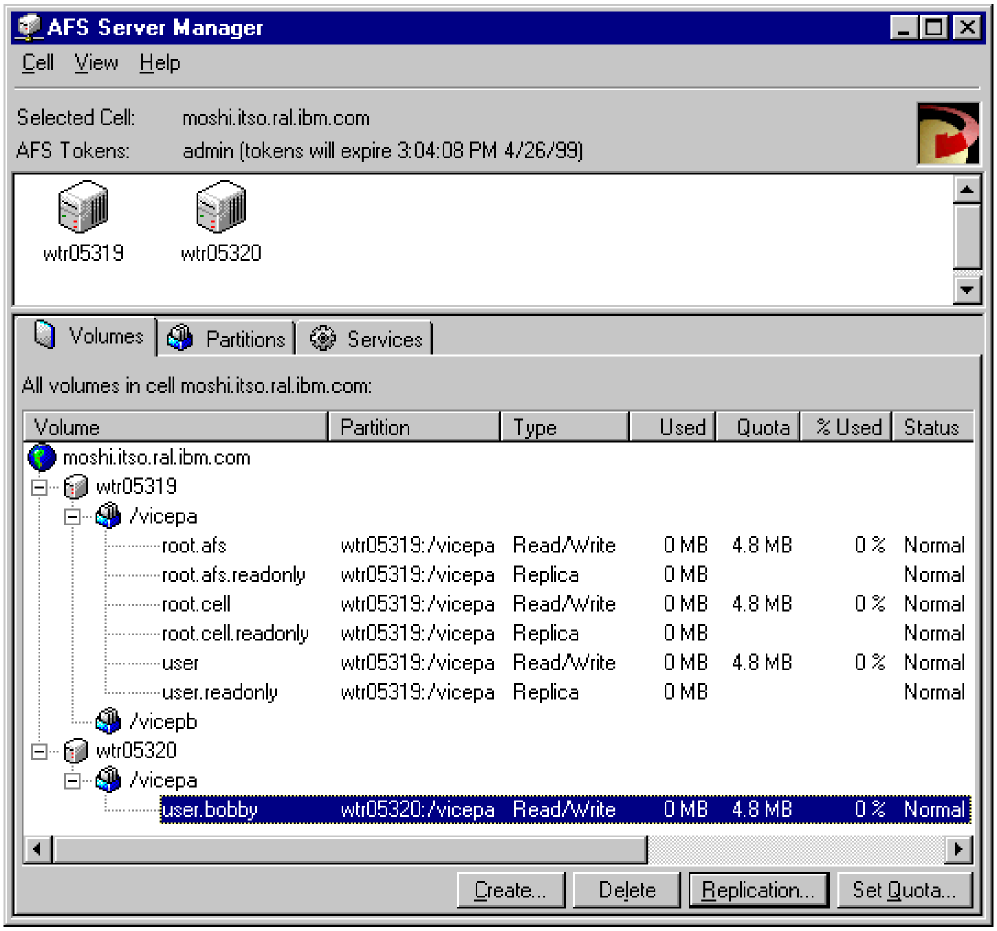

Selected Cell: moshi.itso.ral.ibm.com
AFS Tokens: admin (tokens will expire 3:04:08 PM 4/26/99)

All volumes in cell moshi.itso.ral.ibm.com:

Volume	Partition	Type	Used	Quota	% Used	Status
moshi.itso.ral.ibm.com						
wtr05319						
/vicepa						
root.afs	wtr05319:/vicepa	Read/Write	0 MB	4.8 MB	0 %	Normal
root.afs.readonly	wtr05319:/vicepa	Replica	0 MB			Normal
root.cell	wtr05319:/vicepa	Read/Write	0 MB	4.8 MB	0 %	Normal
root.cell.readonly	wtr05319:/vicepa	Replica	0 MB			Normal
user	wtr05319:/vicepa	Read/Write	0 MB	4.8 MB	0 %	Normal
user.readonly	wtr05319:/vicepa	Replica	0 MB			Normal
/vicepb						
wtr05320						
/vicepa						
user.bobby	wtr05320:/vicepa	Read/Write	0 MB	4.8 MB	0 %	Normal

Figure 264. Displaying the New Location of a Moved AFS Volume

The effect on any user trying to simultaneously access user.bobby would have been minimal: the user might see a brief delay when accessing data, but otherwise their client's Cache Manager transparently learns of user.bobby's new location and immediately resumes fetching data from the new file server.

In AIX, Solaris, and the Windows NT Command Prompt, use the command:

```
vos move
```

to move volumes.

In the following example we move user.tom from one file server to another:

```
C:\> vos move user.tom wtr05319 vicepb wtr05320 vicepa -verbose
Starting transaction  on source volume 536870940 ... done
Cloning source volume 536870940 ... done
Ending the transaction on the source volume 536870940 ... done
Starting transaction on the cloned volume 536870946 ... done
Creating the destination volume 536870940 ... done
Dumping from clone 536870946 on source to volume 536870940 on destination ... do
ne
Doing the incremental dump from source to destination for volume 536870940 ...
done
Deleting old volume 536870940 on source ... done
Starting transaction on the cloned volume 536870946 ... done
Deleting the clone 536870946 ... done
Volume 536870940 moved from wtr05319 /vicepb to wtr05320 /vicepa
```

We recommend always using the `-verbose` switch when moving volumes in order to see what steps are being completed. This is especially important if the move is interrupted for any reason, as the administrator may have to reissue the command above or take other diagnostic steps (unlock the VLDB entry for the volume, for example).

6.5 UDP Ports Used by AFS

As an addendum to this chapter on AFS Server Management, we offer this table of port numbers used by AFS. AFS's communications are entirely UDP-based and use the following UDP ports:

Table 3. UDP Ports Used by AFS Services

Port	Service
7000	fileserver
7001	Cache Manager callback service
7002	ptserver
7003	vlserver (VLDB)
7004	kaserver
7005	volserver (volume management)
7007	bosserver
7008	upserver
7009	AFS/NFS Translator
7020	AFS backup coordinator
7021	AFS backup buserver
7025-7032	AFS backup tape controllers
7101	xstat
2106	fs monitor port
next available port	pts, kas, fs, klog, etc.

Ports 7000-7032 are dedicated for server communications, but the clients use the *next available* port, which means that there are no dedicated AFS client UDP port numbers. To enable AFS access, you need at minimum to open UDP ports 1024 and above.

An AFS server listens on well-defined endpoints (7000 and 7002-7009 for the basic services, 7020-7023 for the Backup System). An AFS client's Cache Manager communicates into port 7001. However, the remaining AFS utilities use the next available port, so it is not possible to predict what port they will use, except that it will not be one of the ones reserved in /etc/services. These utilities include klog (for getting authentication tokens) and all of the command suites (fs, pts, etc).

Chapter 7. AFS Client Management

After the initial AFS cell has been configured and the various servers (File Server machines and Database Server machines) have been set up, the AFS clients must be configured. For Windows NT this can be done through a Control Panel interface; for AIX and Solaris the client configuration must be done via the command line. In this chapter we demonstrate AFS client configuration and also touch upon some configuration and client tuning parameters.

7.1 AFS Client Management for Windows NT

Almost all configuration of the AFS client on Windows NT is done exclusively through the AFS Client Configuration Control Panel. What follows is an example of an AFS client initial configuration and management, going through the following steps:

1. Bring up the AFS Client Configuration Control Panel and configure the client into the cell.
2. Specify a known Database Server and add it to the list of known cells.
3. Mount AFS and map a Windows NT drive to AFS.
4. Set the cache size for the AFS client.
5. Start the AFS client service.

7.1.1 Initial Steps of AFS Client Configuration

We begin this example by bringing up the AFS Client Configuration Control Panel and assigning the client to a cell. This can be done through the following steps:

1. Click **Start** on the taskbar, go to **Settings**, and select **Control Panel**.
2. Double-click the **AFS Client Configuration** icon of the Control Panel. The AFS Client Configuration dialog box appears:

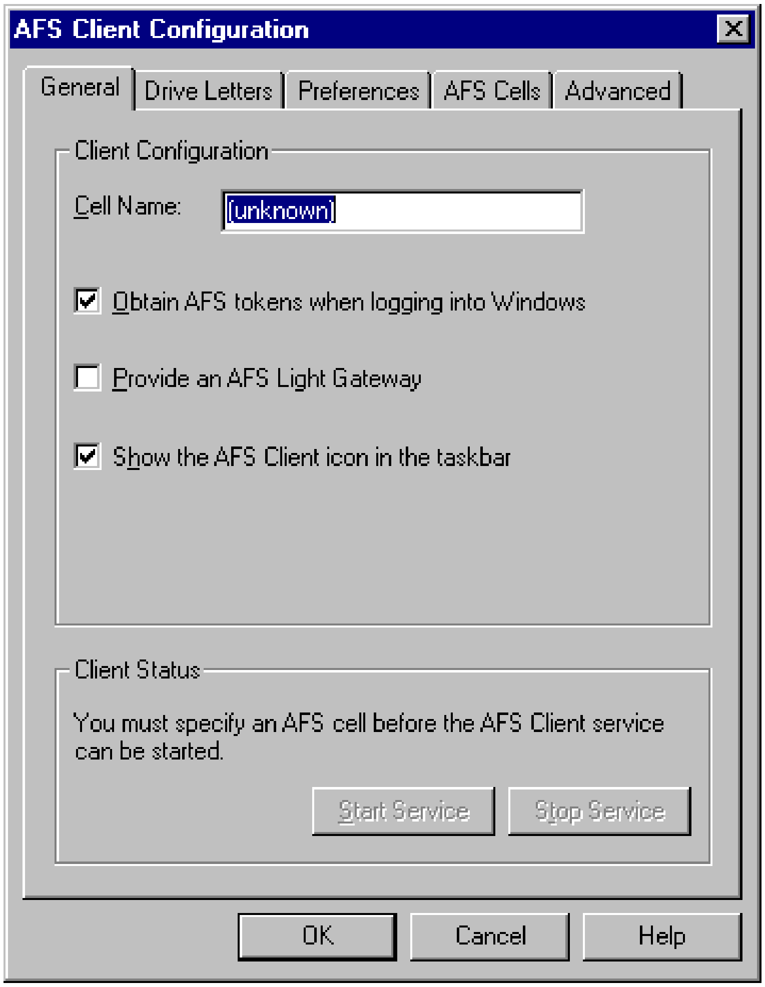

Cell Name: (unknown)

☑ Obtain AFS tokens when logging into Windows

□ Provide an AFS Light Gateway

☑ Show the AFS Client icon in the taskbar

You must specify an AFS cell before the AFS Client service can be started.

Figure 265. An Unconfigured AFS Client Configuration Dialog Box

In the AFS Client Configuration window, under the General tab, you can specify the local AFS cell of your client machine. You can also instruct the machine to allow users to obtain AFS tokens at login time (this process is otherwise known as *integrated logon* and will be discussed in 8.5, “AFS Authentication and Integrated Login” on page 362), to provide an AFS Light Gateway, or to show the AFS Client icon in the Windows NT taskbar. Running an AFS Light Gateway on the Windows NT AFS client allows Windows 95 and Windows 98 machines to access AFS through the gateway machine.

You can also determine whether the AFS Client service is currently running by looking at the Service buttons at the bottom of the screen. If the **Stop Service** button is not grayed out then the AFS Client service is running.

3. Since you want to configure this machine to be an AFS Client of an existing AFS cell, you need to specify the name of the AFS cell in which this client will be participating. Type the name of the AFS cell in the Cell Name field of the Client Configuration frame. In this example, the AFS cell name is *moshi.itso.ral.ibm.com*.

   Since we have yet to create any AFS user accounts we also uncheck the **Obtain AFS tokens when logging into Windows** option for the time being.

At the end of our customization, the General tab of the AFS Client Configuration window appears as shown in the following figure:

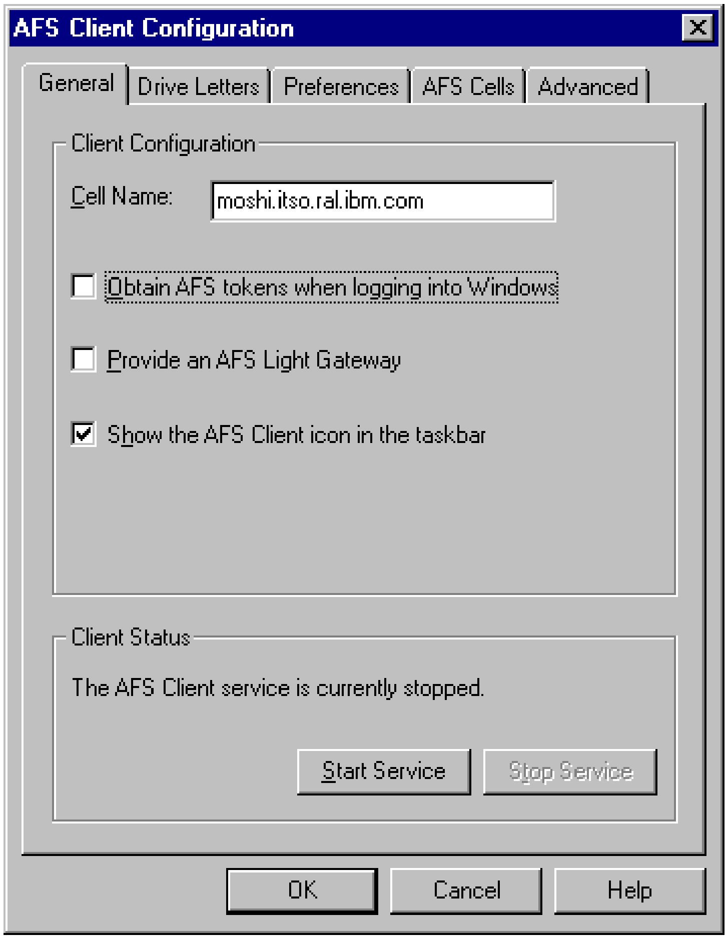

Cell Name: moshi.itso.ral.ibm.com

□ Obtain AFS tokens when logging into Windows

□ Provide an AFS Light Gateway

☑ Show the AFS Client icon in the taskbar

The AFS Client service is currently stopped.

Figure 266. Specify the AFS Cell Name

Although we have assigned the client to an AFS cell, we still need to add information about the cell to the client's list of known cells (stored in C:\WINNT\afsdcell.ini). In particular we need to add the name of the cell and a Database Server to the list. Recall from 2.5.1, "AFS Servers" on page 33 and 6.2, "Managing AFS Server Processes" on page 232 that an AFS Database Server runs a number of database services, including the Volume Location Database (VLDB) Server process, the Protection Server process, and the Authentication Server process (as well as the Backup Server process if configured).

1. Select the **AFS Cells** tab in the AFS Client Configuration window to show the list of known cells (this window may be blank or as in our example have one or more cells). The list of AFS cell names (if there are any) will be listed in alphabetical order; look for the cell to which this client has been assigned.

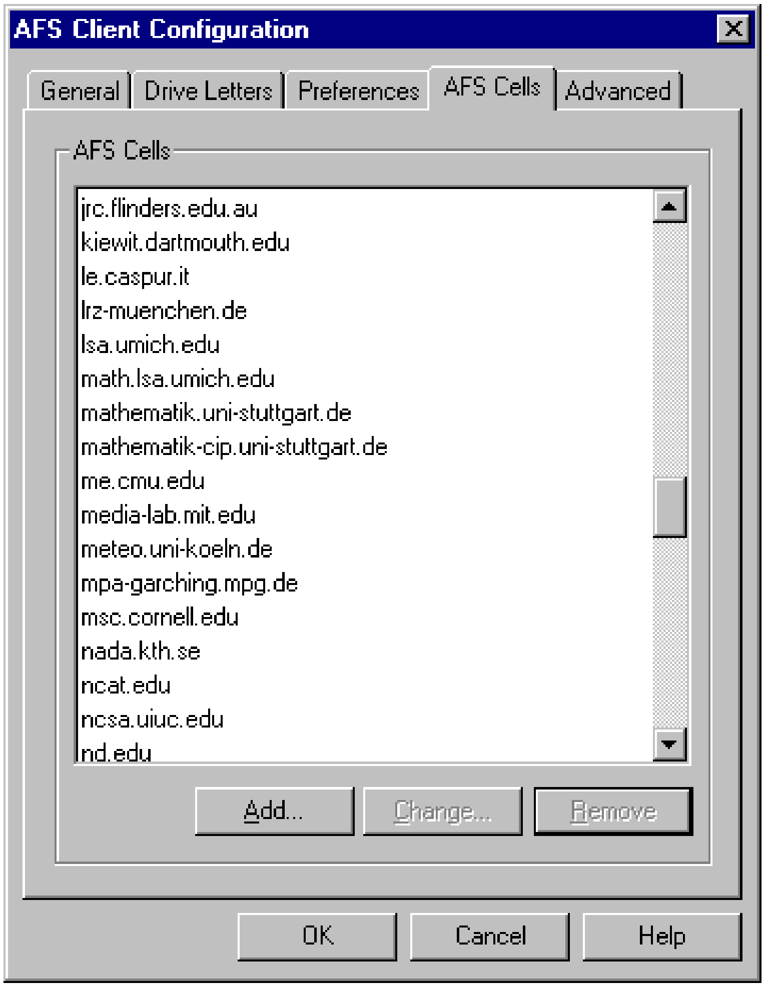

jrc.flinders.edu.au
kiewit.dartmouth.edu
le.caspur.it
lrz-muenchen.de
lsa.umich.edu
math.lsa.umich.edu
mathematik.uni-stuttgart.de
mathematik-cip.uni-stuttgart.de
me.cmu.edu
media-lab.mit.edu
meteo.uni-koeln.de
mpa-garching.mpg.de
msc.cornell.edu
nada.kth.se
ncat.edu
ncsa.uiuc.edu
nd.edu

Figure 267. The AFS Cells Tab

2. In our example we are searching for the AFS cell moshi.itso.ral.ibm.com, which in fact has not yet been defined in this client’s afsdcell.ini file. Click **Add...** to display the New Cell dialog box to add a new AFS cell. In the **AFS Cell** field, type the name of the new AFS cell name, moshi.itso.ral.ibm.com. In the Description field, enter a brief description of this cell, as shown in the following figure:

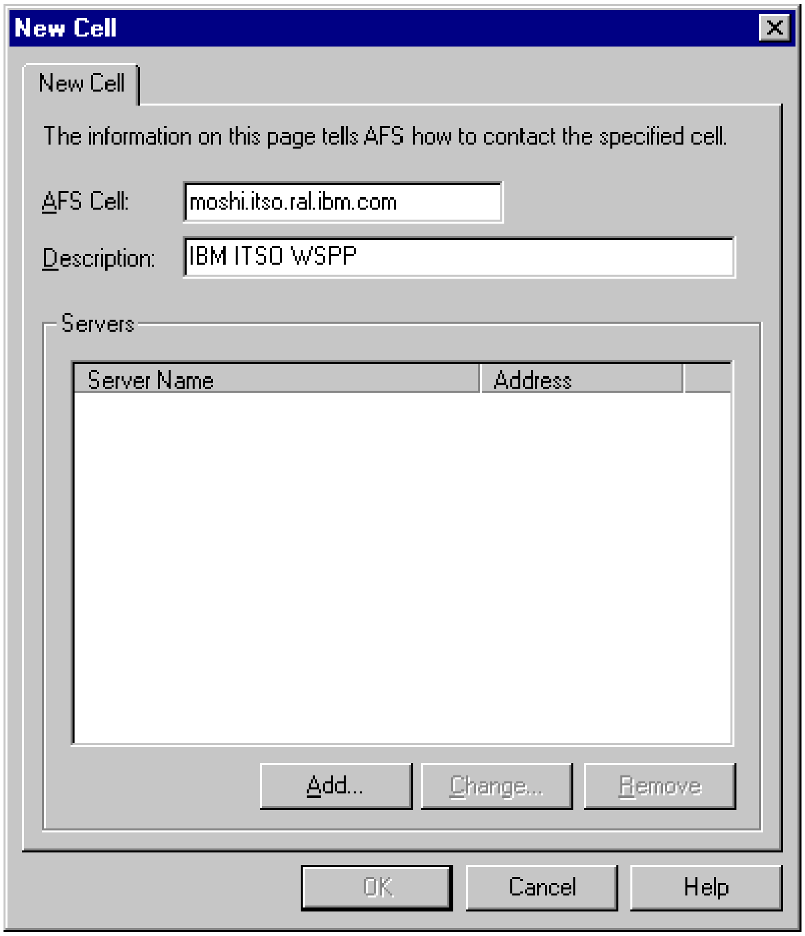
New Cell
New Cell
The information on this page tells AFS how to contact the specified cell.

AFS Cell: moshi.itso.ral.ibm.com

Description: IBM ITSO WSPP

Figure 268. Define the New AFS Cell Name and Description

3. Then click **Add...** to bring up the Add Server dialog box. From the Add Server window, in the Server Name field, enter the name of a Database Server for the AFS cell. In this example, wtr05319 is the name of the Database Server machine in cell moshi.itso.ral.ibm.com.

   Note that when installing WebSphere Performance Pack, the first machine to have the AFS File Sharing component installed will automatically be configured as the first AFS File Server, the first Database Server, and also as an AFS client.

   For the Server Address, you can enter a specific IP address or use the Domain Name System (DNS) to look up the address automatically. If the server machine has been registered with DNS, check the **Look up address automatically** option to use DNS to resolve the host names.

   The following figure shows the configuration in our sample system:

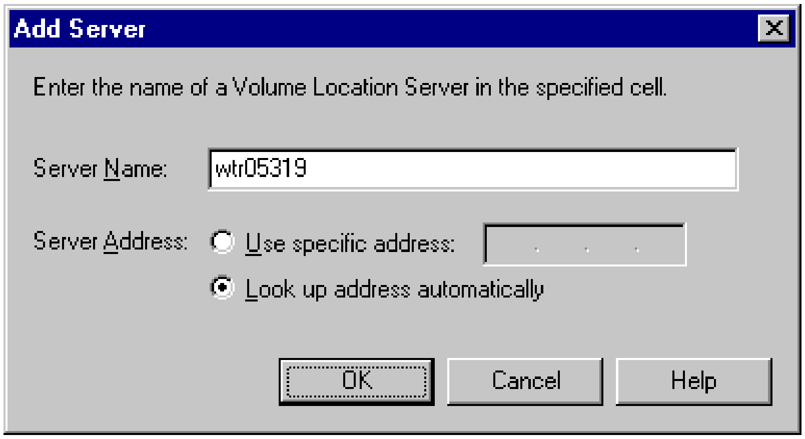

Add Server

Enter the name of a Volume Location Server in the specified cell.

Server Name: wtr05319

Server Address: ○ Use specific address: . . .

● Look up address automatically

Figure 269. Define the Name of a Database Server in the Cell

Click **OK** to add this server to the specified AFS cell.

4. The AFS Client configuration program returns you to the New Cell window, showing the information about the new AFS cell being added:

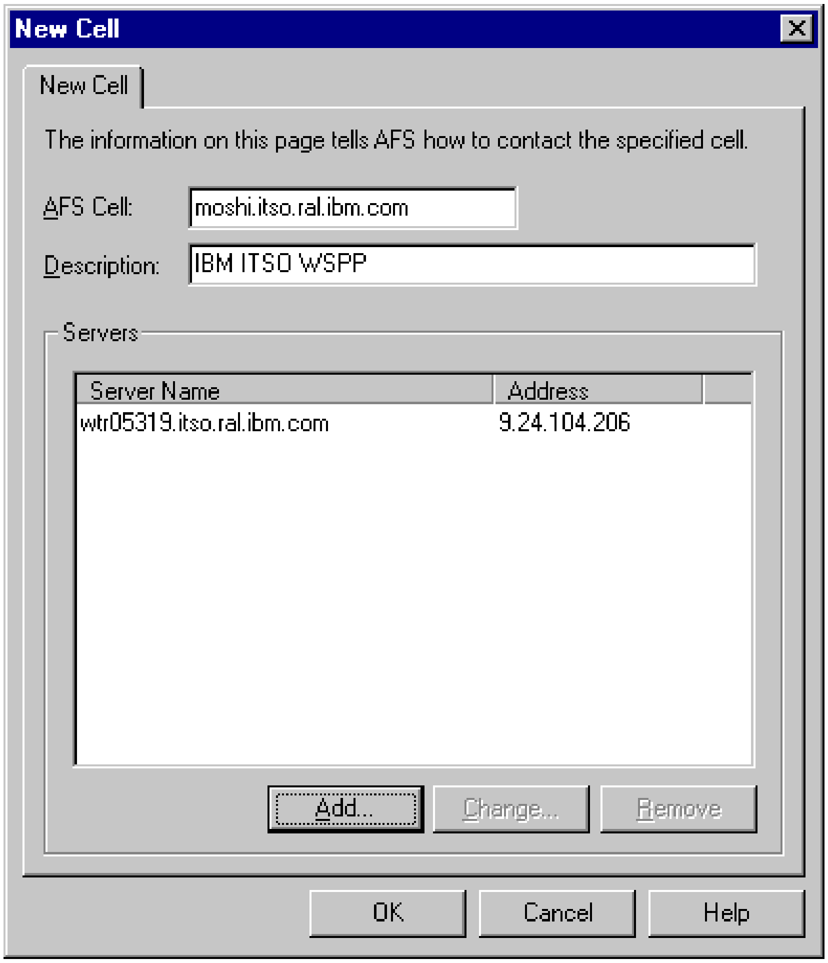

New Cell

The information on this page tells AFS how to contact the specified cell.

AFS Cell: moshi.itso.ral.ibm.com

Description: IBM ITSO WSPP

Servers

Server Name	Address
wtr05319.itso.ral.ibm.com	9.24.104.206

Figure 270. The New AFS Cell Information

If there is more than one Database Server in your cell you should add them to the list of known Volume Location (VL) servers in the cell by repeating the **Add...** process. Otherwise, click **OK** to finally add the new AFS cell to the list of known AFS cells.

Note that the specified AFS cell *moshi.itso.ral.ibm.com* has now been added to the list of known AFS cells on this particular Windows NT client.

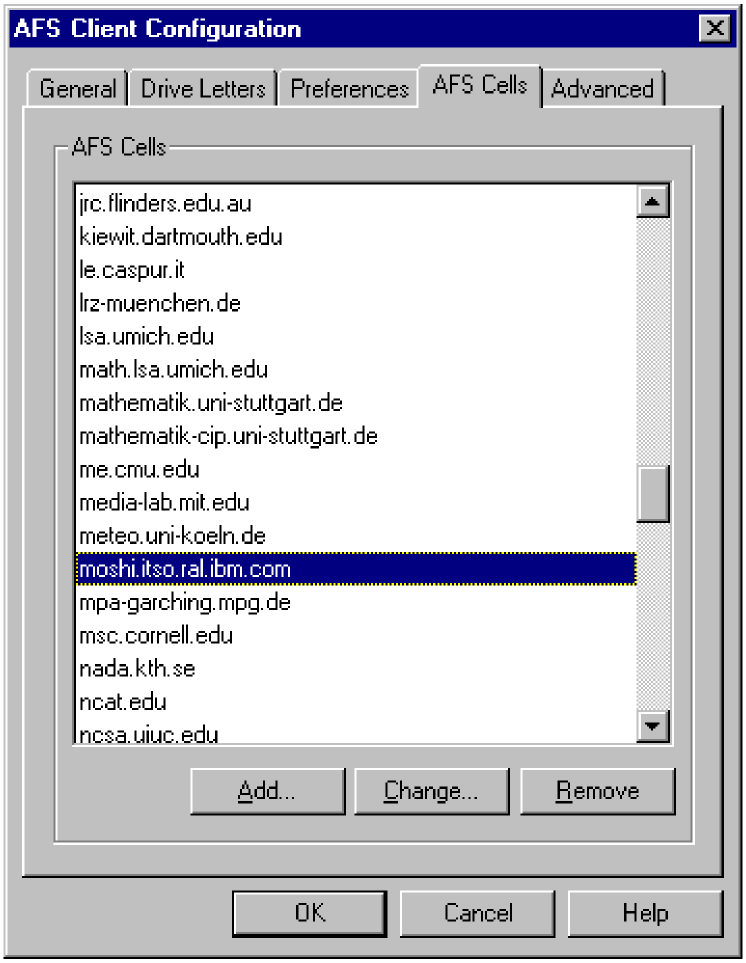

jrc.flinders.edu.au
kiewit.dartmouth.edu
le.caspur.it
lrz-muenchen.de
lsa.umich.edu
math.lsa.umich.edu
mathematik.uni-stuttgart.de
mathematik-cip.uni-stuttgart.de
me.cmu.edu
media-lab.mit.edu
meteo.uni-koeln.de
moshi.itso.ral.ibm.com
mpa-garching.mpg.de
msc.cornell.edu
nada.kth.se
ncat.edu
ncsa.uiuc.edu

Figure 271. The New AFS Cell in the AFS Cell List

Again, this process updates the list of AFS Database Server machines for local and foreign cells stored in the file C:\WINNT\afsdcell.ini.

The next steps in our initial client configuration are intended to set the size of the AFS client cache:

1. Select the **Advanced** tab to display the default AFS cache configuration. The default Cache Size value is set to be 20480. In AFS Version 3.5 for Windows NT, the AFS cache is implemented as a single cache file, C:\AFSCache, which is mapped into memory. In this example we opt to use the default AFS cache size, though in practice you may need to increase the AFS cache size to improve performance:

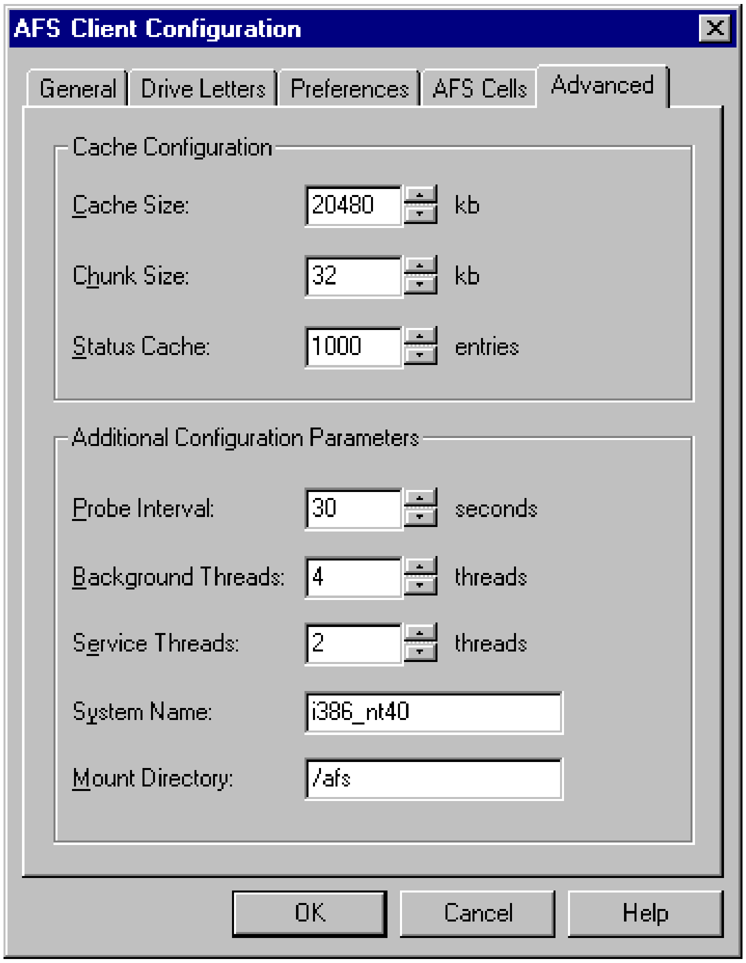

Cache Configuration

Cache Size: 20480 kb

Chunk Size: 32 kb

Status Cache: 1000 entries

Additional Configuration Parameters

Probe Interval: 30 seconds

Background Threads: 4 threads

Service Threads: 2 threads

System Name: i386_nt40

Mount Directory: /afs

Figure 272. The AFS Cache Configuration Parameters

The other AFS cache configuration parameters will be discussed in 7.3, “Managing the AFS Client Cache” on page 318 and 7.4.3, “AFS Client Tuning” on page 332.

2. Click **OK** when done.

Finally we start the AFS Client service, which runs as `IBM AFS Client` in the Services Control Panel. To do this, we select the **General** tab in the AFS Client Configuration window and click **Start Service** to start the AFS Client service. This is shown in the following window:

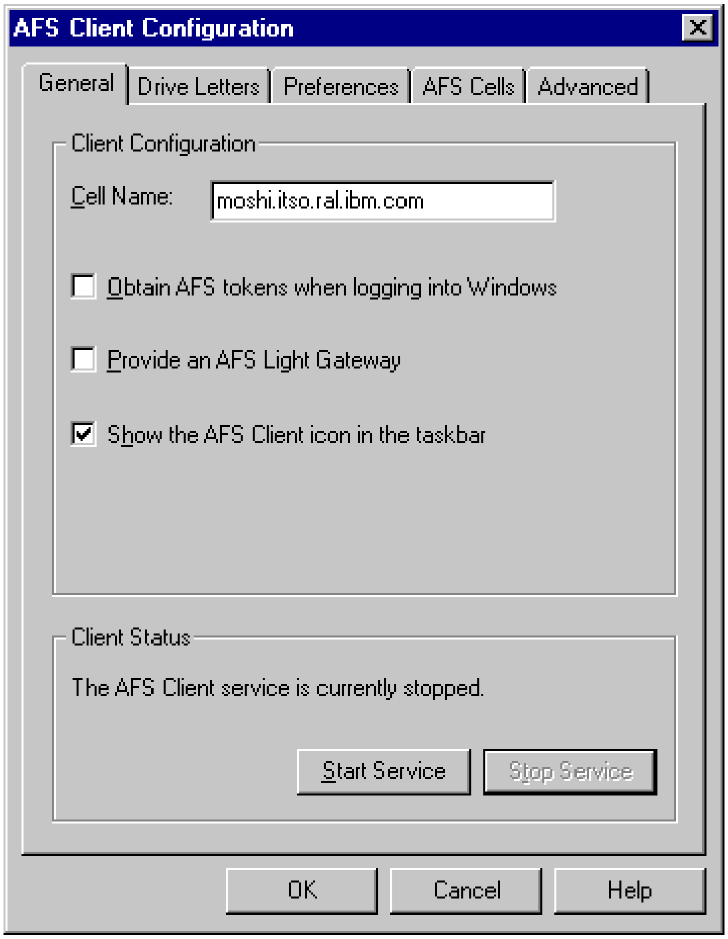

Client Configuration

Cell Name: moshi.itso.ral.ibm.com

☐ Obtain AFS tokens when logging into Windows

☐ Provide an AFS Light Gateway

☑ Show the AFS Client icon in the taskbar

Client Status

The AFS Client service is currently stopped.

Figure 273. Starting the AFS Client Service

The AFS Client service will display a status window informing you of its progress:

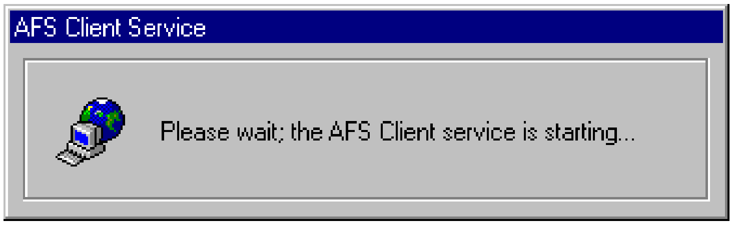

Please wait; the AFS Client service is starting...

Figure 274. The AFS Client Service Window

Configuration of the AFS Client is complete. If you want to stop the AFS Client service, go back to the AFS Client Configuration window, and click the button **Stop Service**, which is enabled when the AFS Client service is running:

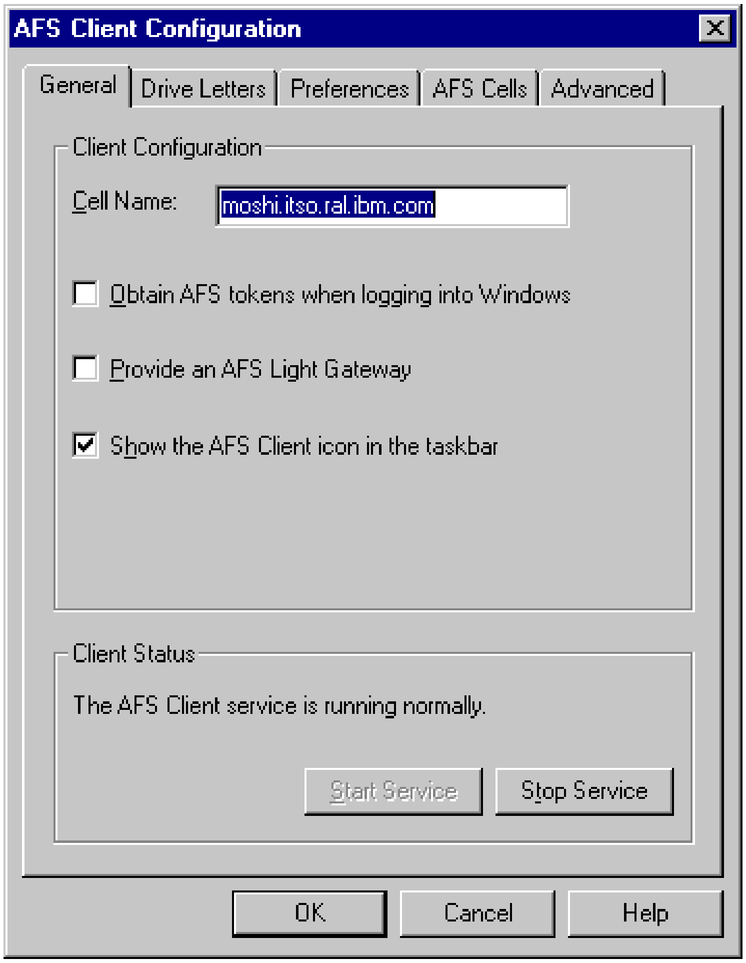

Cell Name: moshi.itso.ral.ibm.com

□ Obtain AFS tokens when logging into Windows

□ Provide an AFS Light Gateway

☑ Show the AFS Client icon in the taskbar

The AFS Client service is running normally.

Figure 275. Clicking the Stop Service Button

Notice that the AFS Client service automatically starts when you restart Windows NT, unless you disable that option in the Windows NT Control Panel Services window during configuration. This configuration option is also built into the AFS Client GUI; the user must access the AFS Client Property box from the AFS Client icon in the taskbar, and select the **Advanced** tab. The option appears as follows: **Start the AFS Client Service whenever this computer restarts**.

7.1.2 Mounting the AFS File System

Although we are now running the AFS client service, we still cannot access data in AFS file space. To access the AFS file space, you need to map AFS paths to drive letters. There are three methods you can use to map AFS directory space to drive letters:

1. The AFS Client Configuration Control Panel
2. The command:

   ```
   net use
   ```

   in a Windows NT Command Prompt window
3. The **Map Network Drive...** option in the Windows NT Explorer.

In this section we demonstrate with an example how to use each method.

7.1.2.1 AFS Client Configuration from the Control Panel

Click **Start** on the Windows NT taskbar, go to **Settings**, and select **Control Panel**. Click the **AFS Client Configuration** icon. Select the **Drive Letters** tab. Initially, there is no drive map to any AFS path:

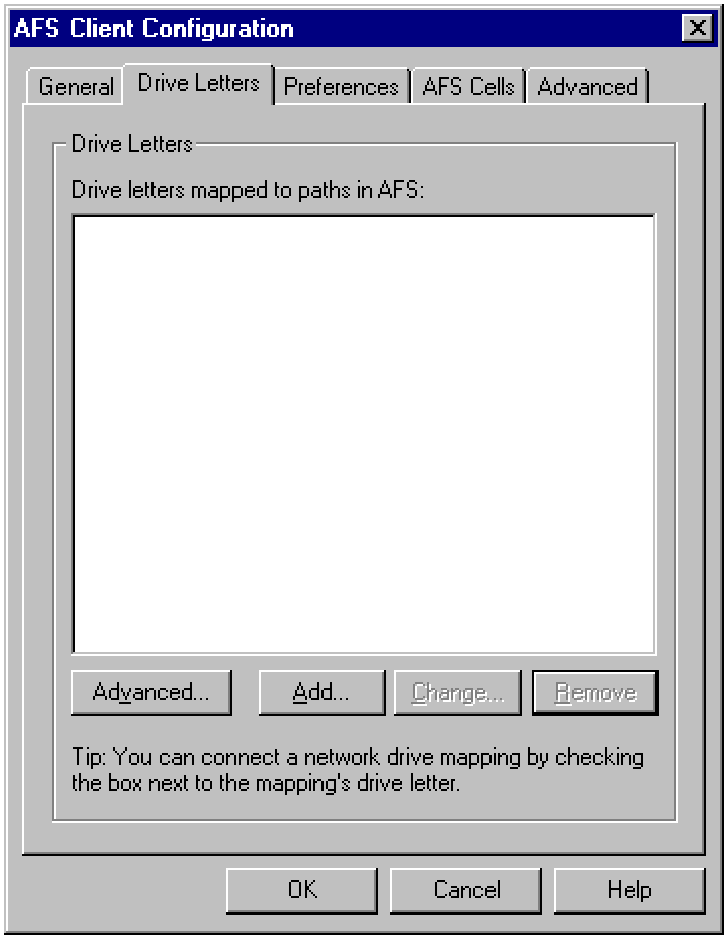

Tip: You can connect a network drive mapping by checking the box next to the mapping's drive letter.

Figure 276. The Drive Letters Tab of the AFS Client Configuration Window

Click **Add** to go to the Add Drive Letter Window. Select the drive letter that you would like to map the AFS path /afs to. In this example, we use drive M:. Check the **Restore this mapping whenever I logon** box to be able to access this AFS path each login session. These settings are shown in the following figure:

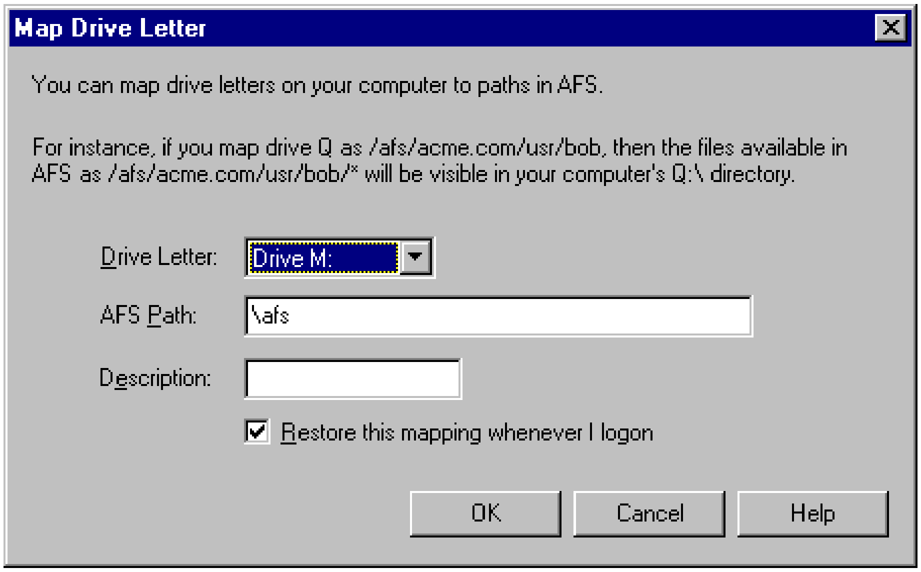

Map Drive Letter

You can map drive letters on your computer to paths in AFS.

For instance, if you map drive Q as /afs/acme.com/usr/bob, then the files available in AFS as /afs/acme.com/usr/bob/* will be visible in your computer's Q:\ directory.

Drive Letter: Drive M:

AFS Path: \afs

Description:

☑ Restore this mapping whenever I logon

Figure 277. Mapping Drive Letter to AFS Path

Click **OK** to complete mapping the drive. The Drive Letters tab now displays the drive to be mapped to the specified AFS path:

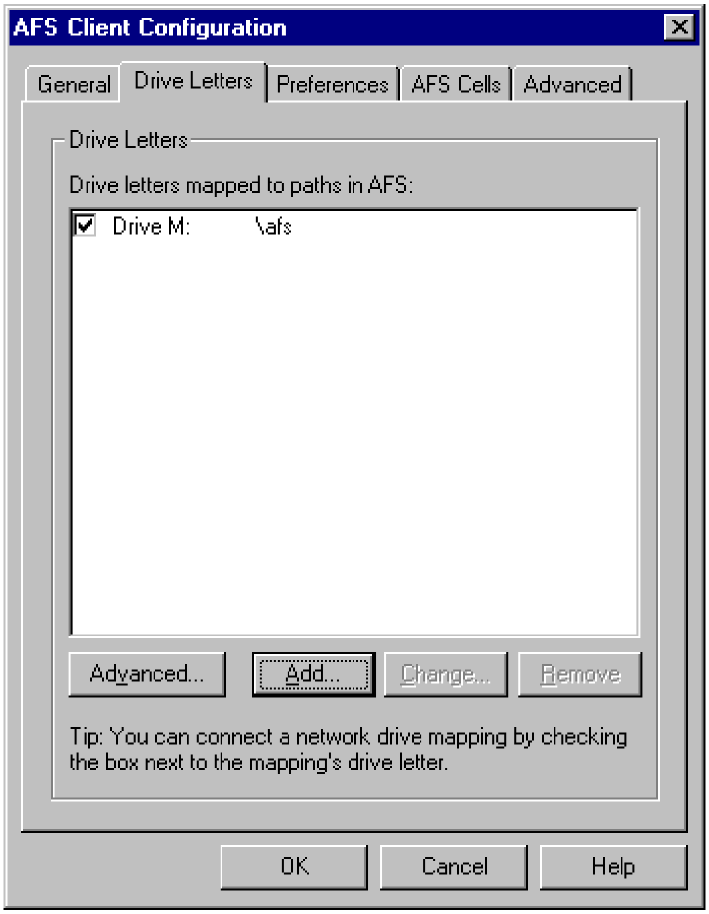

AFS Client Configuration

Drive Letters

Drive letters mapped to paths in AFS:

☑ Drive M: \afs

Tip: You can connect a network drive mapping by checking the box next to the mapping's drive letter.

Figure 278. Drive Letters Mapped to AFS Paths

Clicking on the **Advanced...** button in the window above brings up the Submounts window. In AFS terms, a *submount* is a shared network drive that has been mapped to AFS:

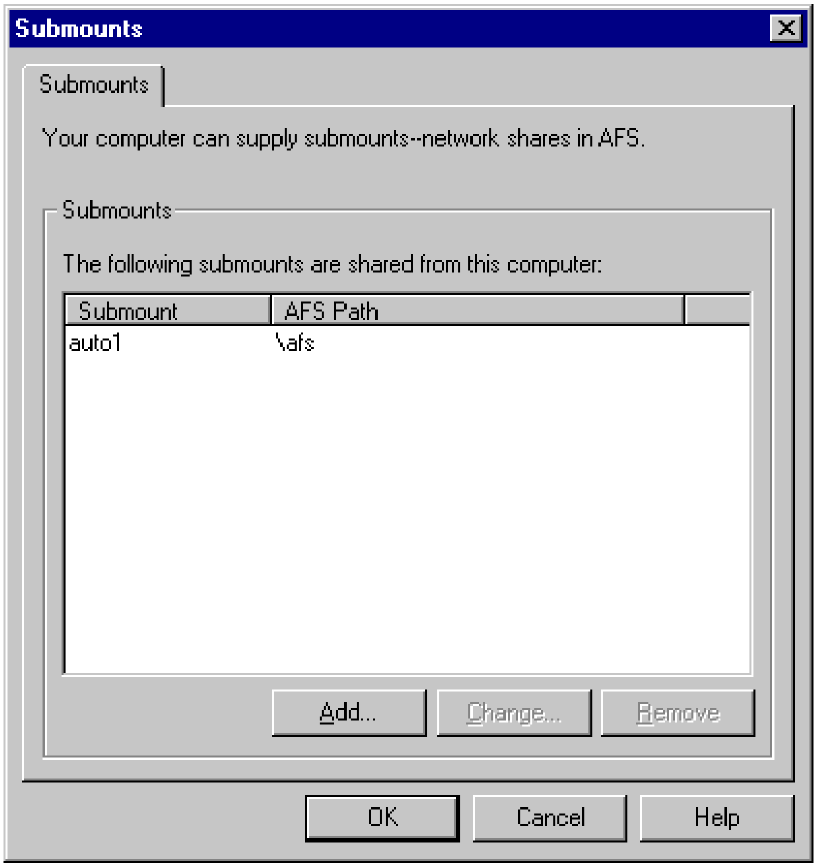

Your computer can supply submounts--network shares in AFS.

The following submounts are shared from this computer:

Submount	AFS Path
auto1	\afs

Figure 279. Submount Example

7.1.2.2 The net use Command

This command, that you should enter from an MS-DOS Command Prompt window, has the following syntax:

```
net use drive: \\hostname-afs\submount
```

where:

- *drive* is the drive letter that you want to map the AFS path to.
- *hostname* is the host name of the Windows NT machine.
- *submount* is the name of the subdirectory to be mounted.

  Use either `all` or `auto1` as the submount name under which the AFS path /afs will be mounted. However, notice that `all` and `auto1` are not fully equivalent. The name all is reserved to mean that you are mounting the file space. The name auto1 is the name of a submount that you have already created.

For example, on the Windows NT machine WTR05201, to map drive M: to the AFS path /afs, enter:

```
net use M: \\WTR05201-afs\all
```

If you enter:

```
net use M: \\WTR05201-afs\auto1
```

this command will only map drive M: to /afs if the auto1 submount is mapped to the location /afs.

If you want the system to remember the mapping information and reconnect to AFS, use the `/persistent:yes` flag with the command above. In this example, you should enter:

```
net use M: \\WTR05201-afs\all /persistent:yes
```

7.1.2.3 Windows NT Explorer

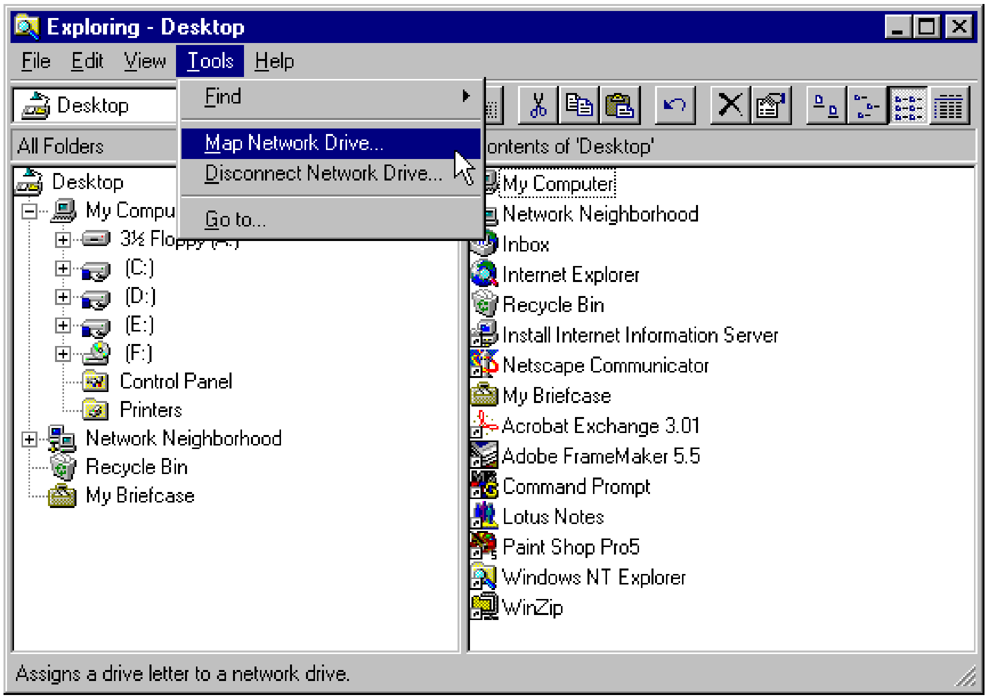

Assigns a drive letter to a network drive.

Figure 280. Map Network Drive from Windows NT Explorer

From a Windows NT Explorer window, click **Tools** and select **Map Network Drive...**.

After the Map Network Drive window is displayed, select the drive to be mapped. Enter the AFS path in the format, `\\`*`hostname`*`-afs\all` as in 7.1.2.2, “The net use Command” on page 315. In this example, the AFS path is \\WTR05201-afs\all. Check the **Reconnect at Logon** box and click **OK** to map the drive. These settings are shown in the following figure:

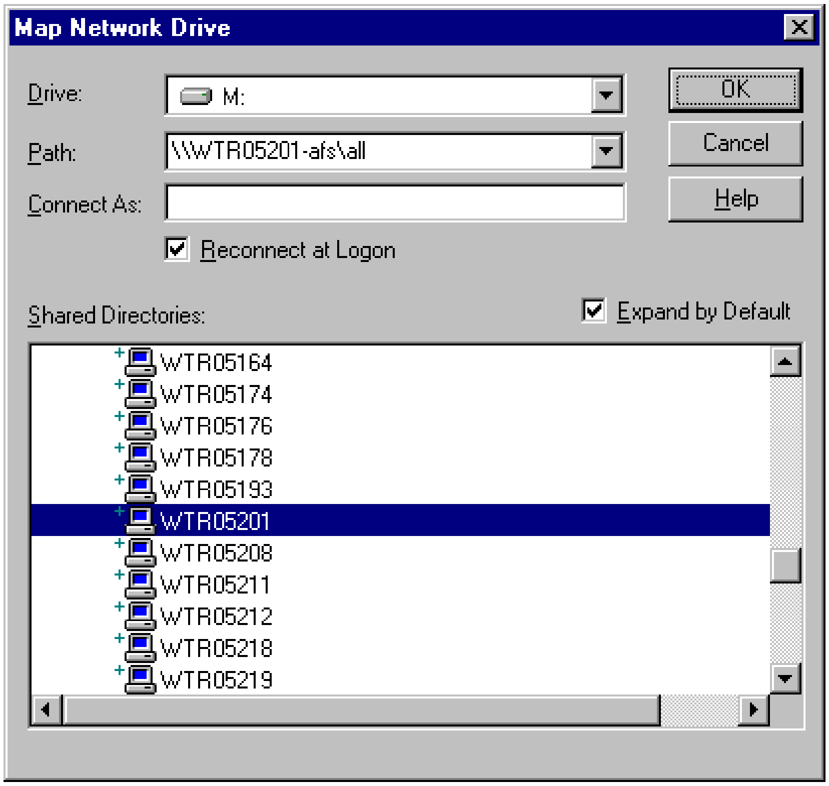

Drive: M:

Path: \\WTR05201-afs\all

Connect As:

☑ Reconnect at Logon

Shared Directories: ☑ Expand by Default

- WTR05164
- WTR05174
- WTR05176
- WTR05178
- WTR05193
- WTR05201
- WTR05208
- WTR05211
- WTR05212
- WTR05218
- WTR05219

Figure 281. Mapping Network Drive

The AFS Client is now configured in the specified AFS cell, moshi.itso.ral.ibm.com. You can access the AFS file space via the selected drive mapping, M:, using Windows NT Explorer:

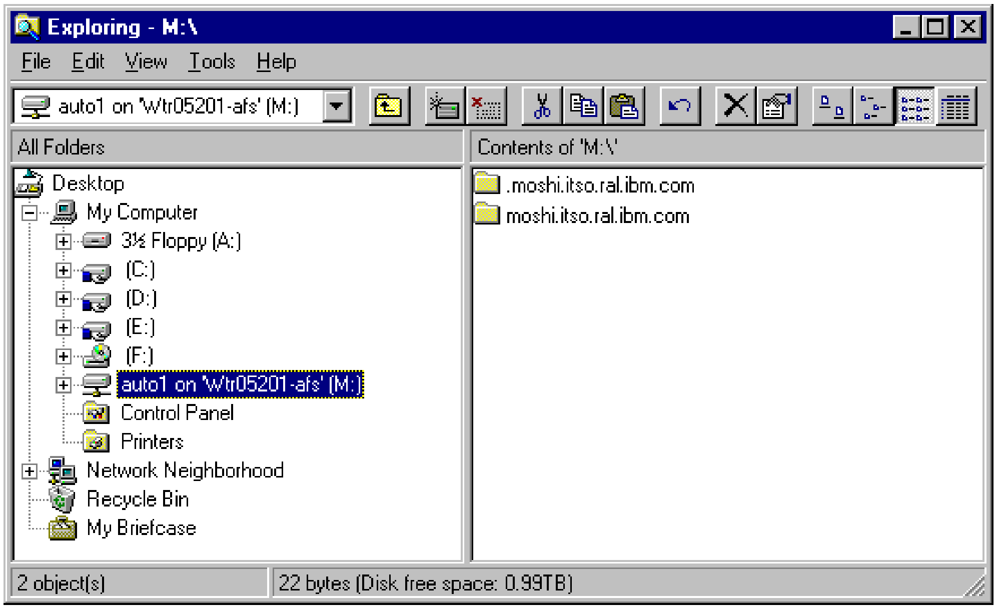

auto1 on 'Wtr05201-afs' (M:)

Contents of 'M:\'

- .moshi.itso.ral.ibm.com
- moshi.itso.ral.ibm.com

2 object(s) | 22 bytes (Disk free space: 0.99TB)

Figure 282. Windows NT Explorer with Drive Letter Mapping to AFS Path

7.2 AFS Client Management for AIX and Solaris

Unlike Windows NT, the management of AIX and Solaris AFS clients must be done through a command line interface. For most of the tasks and options available in the Windows NT AFS Client Configuration Control Panel there is a corresponding command or file to edit in UNIX. For instructions on installing and initially configuring AFS clients on UNIX, refer to 3.6, “Configuring an AFS Client Machine on AIX Systems” on page 89 for AIX and 4.6, “Configuring an AFS Client Machine on Solaris Systems” on page 117 for Solaris. Briefly:

- To change a client’s cell membership edit the file /usr/vice/etc/ThisCell, which contains the name of the cell to which the AFS client belongs.
- To add or delete cells and VL servers about which the local AFS client is aware, edit the file /usr/vice/etc/CellServDB. Below is an example of entries in a CellServDB file:

```
% more /usr/vice/etc/CellServDB
>ies.auc.dk                # Aalborg Univ., Inst. of Electronic Systems, Denmark
130.225.51.28                   #idefix.kom.auc.dk
>fusione.it                     #Assoc. Fusione/Euratom, ENEA, Frascati, Italy
192.107.51.81                   #efrs03.frascati.enea.it
>umbc.edu                       #University of Maryland, Baltimore Co.
130.85.6.33                     #chardonnay.umbc.edu
130.85.1.13                     #burgundy.umbc.edu
130.85.60.12                    #chablis.umbc.edu
>dsi.uniroma1.it           # Univ. Rome-1, Dept. of Computer Science
151.100.17.2                    #afscinfo.dsi.uniroma1.it
>ovpit.indiana.edu         # OVPIT at Indiana University
129.79.207.195                  #noir.ovpit.indiana.edu
>wu-wien.ac.at             # University of Economics, Vienna, Austria
137.208.3.33                    #afsdb1.wu-wien.ac.at
137.208.7.4                     #afsdb2.wu-wien.ac.at
```

- To change the size of the AFS client cache (effective with the next reboot) edit the last field in the file /usr/vice/etc/cacheinfo. An example of a cacheinfo file is shown in the following session screen:

```
% cat /usr/vice/etc/cacheinfo
/afs:/usr/vice/cache:90000
```

- To change the AFS client’s behavior on bootup in AIX, edit the file /etc/rc.afs.
- To change the AFS client’s behavior on bootup in Solaris, edit the file /etc/init.d/afs (note that the files /etc/rc3.d/S99afs and /etc/rc0.d/K66afs are linked to it).

7.3 Managing the AFS Client Cache

In this section we address the AFS client cache, demonstrating how to change the cache size and how to flush data from the cache, and we look more closely at data consistency and validation in the AFS client cache.

7.3.1 Setting the Client Cache Size

The size of the AFS client cache is probably the single most important factor in determining the performance of AFS. Factors that determine the appropriate cache size for a given client include the number of users working on the machine, the size of the files with which they usually work, and (for a memory cache on UNIX clients) the number of processes that usually run on the machine. The higher the demand from these factors, the larger the cache needs to be to maintain good performance. Disk caches smaller than 10 MB do not generally perform well. Machines serving multiple users usually perform better with a cache of at least 60 to 70 MB. The point at which enlarging the cache further does not really improve performance depends on the factors mentioned above and is difficult to predict. An AFS client's cache can be initially set at installation and then later adjusted (upward) dynamically, up to near the maximum partition size set aside for it.

7.3.1.1 Windows NT

On Windows NT the cache can be sized either using the AFS Client Configuration Control Panel or using commands in the Command Prompt windows. To set the cache size using the graphical user interface (GUI), open up the Control Panel double-click **AFS Client Configuration**, and select the **Advanced** tab. In our example the cache is currently set to 40 MB, as you can see from the following figure:

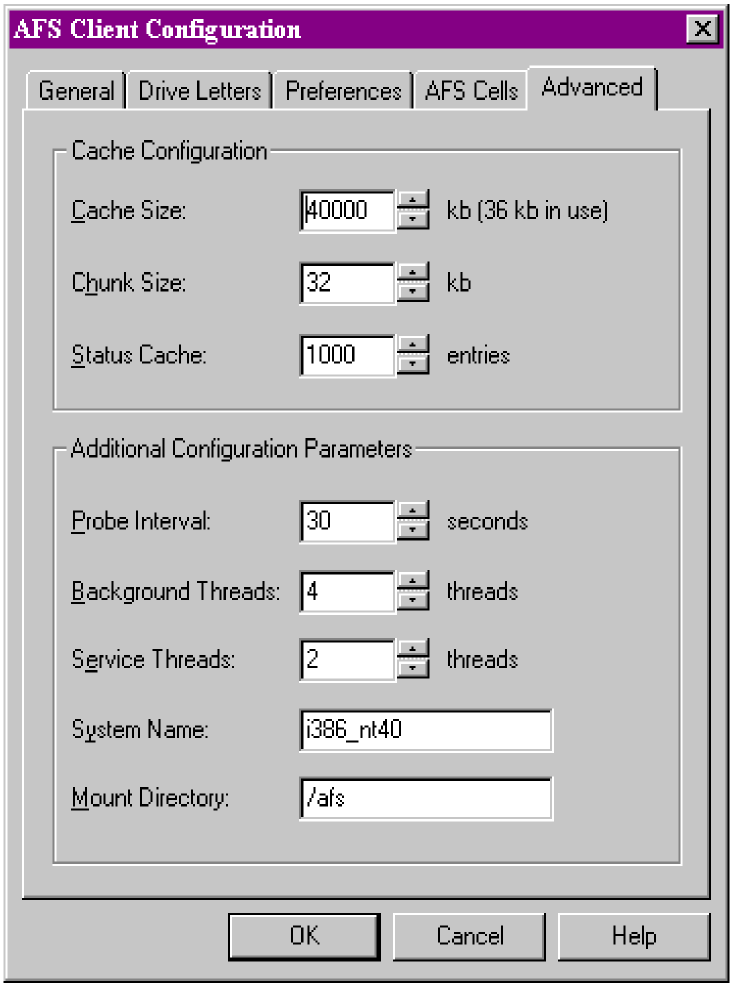

Cache Configuration

Cache Size: 40000 kb (36 kb in use)

Chunk Size: 32 kb

Status Cache: 1000 entries

Additional Configuration Parameters

Probe Interval: 30 seconds

Background Threads: 4 threads

Service Threads: 2 threads

System Name: i386_nt40

Mount Directory: /afs

Figure 283. Client Cache Configuration

Click in the **Cache Size** field and change the cache size to a desired setting. Note that if you *reduce* the cache size you will have to stop and restart the IBM AFS Client service in order for the change to take effect. In our example we are increasing the cache size to 60 MB, as shown in the following figure:

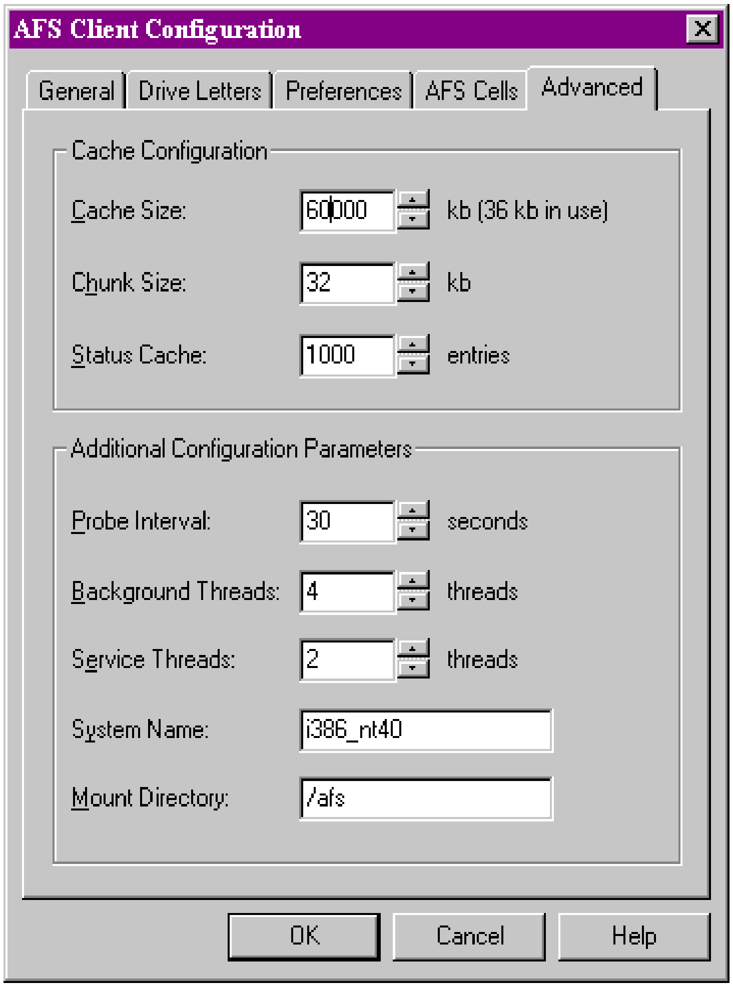

Cache Configuration

Cache Size: 60000 kb (36 kb in use)

Chunk Size: 32 kb

Status Cache: 1000 entries

Additional Configuration Parameters

Probe Interval: 30 seconds

Background Threads: 4 threads

Service Threads: 2 threads

System Name: i386_nt40

Mount Directory: /afs

Figure 284. Client Cache Configuration after Change

Click **OK** to close the window. To verify that the cache size has been changed, we have before and after results, taken in the Command Prompt window, which show the size of the AFS cache file C:\AFSCache:

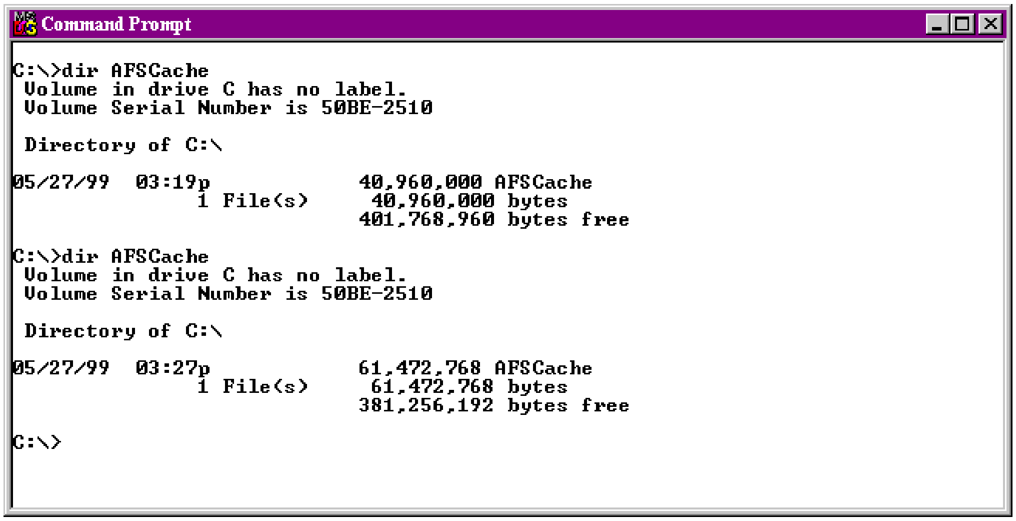

```
C:\>dir AFSCache
 Volume in drive C has no label.
 Volume Serial Number is 50BE-2510

 Directory of C:\

05/27/99  03:19p            40,960,000 AFSCache
               1 File(s)     40,960,000 bytes
                            401,768,960 bytes free

C:\>dir AFSCache
 Volume in drive C has no label.
 Volume Serial Number is 50BE-2510

 Directory of C:\

05/27/99  03:27p            61,472,768 AFSCache
               1 File(s)     61,472,768 bytes
                            381,256,192 bytes free

C:\>
```

Figure 285. Before and after Size of AFSCache File

7.3.1.2 AIX, Solaris and Windows NT Command Prompt

For AIX, Solaris and the Windows NT Command Prompt, there are several commands for getting and setting cache parameters. Use the command:

```
fs getcacheparms
```

to get rudimentary cache statistics, for example:

```
C:\> fs getcacheparms
AFS using 36 of the cache's available 60032 1K byte blocks.
```

Use the command

```
fs setcachesize
```

to dynamically change the cache size. Again, the Cache Manager will not accept a reduction in cache size while running, so attempting to set a value lower than the current cache size will result in an error:

```
C:\> fs setcache -blocks 30000
fs: Invalid argument.

C:\> fs setcache -blocks 60200
New cache size set.
```

Notice that the command:

```
fs setcachesize 0
```

does not actually set the cache size to 0. Instead, it sets the cache size back to its value in the cacheinfo file.

Finally, on AIX and Solaris the /usr/vice/etc/cacheinfo file should be edited to reflect any permanent changes to the cache size. The machine will have to be rebooted in order to pick up this change.

7.3.2 Flushing Cached Data

There may be times when you will want to flush data from an AFS Client's cache, typically because you know the data has changed on the File Server and you wish to pick up those changes immediately.

Use the command:

```
fs checkvolumes
```

to force the Cache Manager to discard its cached knowledge of volume locations. It does this by discarding its internal mappings of volume names to IDs, which in turn forces the Cache Manager on data access to consult the VLDB to relearn the location of volumes. The command for AIX, Solaris, and the Windows NT Command Prompt is shown in the following session screen:

```
% fs checkvolumes
All volumeID/name mappings checked.
```

The two commands for flushing user data from the AFS Client cache are:

```
fs flush
fs flushvolume
```

Again, these are the same for AIX, Solaris and the Windows NT Command Prompt:

1. The first of the two commands will flush from the cache all data and status information *for the named directory or file*, which in turn will force the Cache Manager to refetch the data from the server the next time it is accessed.
2. In contrast, the second of the two commands will flush all data *from the volume referenced by the named directory or file*.

In the example below, we first flush data for the /afs/moshi.itso.ral.ibm.com/user directory, and then we flush the volume referenced by the directory (or mount point) /afs/moshi.itso.ral.ibm.com/user/bobby, which in this case is the user.bobby volume:

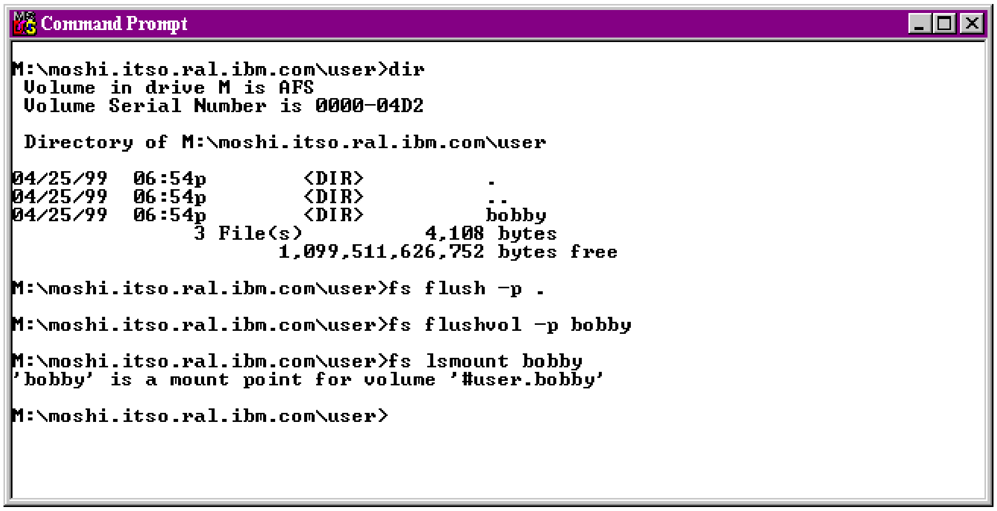

```
M:\moshi.itso.ral.ibm.com\user>dir
 Volume in drive M is AFS
 Volume Serial Number is 0000-04D2

 Directory of M:\moshi.itso.ral.ibm.com\user

04/25/99  06:54p        <DIR>          .
04/25/99  06:54p        <DIR>          ..
04/25/99  06:54p        <DIR>          bobby
               3 File(s)         4,108 bytes
                      1,099,511,626,752 bytes free

M:\moshi.itso.ral.ibm.com\user>fs flush -p .

M:\moshi.itso.ral.ibm.com\user>fs flushvol -p bobby

M:\moshi.itso.ral.ibm.com\user>fs lsmount bobby
'bobby' is a mount point for volume '#user.bobby'

M:\moshi.itso.ral.ibm.com\user>
```

Figure 286. Cache Flushing Commands

Note that there is no output from either one of the cache flushing commands.

On Windows NT, users can right-click AFS files and directories, access the AFS menu, and then flush files and directories or flush volumes.

7.3.3 A Closer Look at AFS Cached Data Validation

Normally Web browsers locally cache files that have been accessed recently, but they cannot guarantee that they have the most recent version or that the current cached copy has been changed back on the server that stores the original copy. The browser could contact the origin server to determine if its copy is still good, but this requires more network traffic and delay. AFS uses a proactive approach, since the AFS File Server machine notifies that there is a new version to all potentially interested clients.

When an AFS user, from an AFS client machine, requests a file from the AFS File Server machine, the Cache Manager first checks whether that file is already present in the local cache in the AFS client machine. This is true if that file has already been requested. If the file is not there, then the Cache Manager requests it from the AFS File Server Machine, gets the file, stores it in its local cache and serves the file to the application that has requested it (for example, a text editor).

The AFS File Server machine keeps a list of all the AFS client machines that have requested and cached data associated with a certain file in their AFS cache (either RAM or disk). When that particular file is changed by a user or process that is entitled to do so, the AFS File Server machine directly contacts all the AFS Clients that requested that file in the past, informing them that newer data is now available. This process is also known as *callback*.

The AFS client's kernel can safely return cached data to the application until it is notified by callback. When callback is received, the client must retrieve a fresh copy of the data since the cache is out of date.

Notice that if an AFS client receives callback from the AFS server, it marks the cached data as invalid. However, to reduce useless network traffic, the AFS client will retrieve a fresh copy of the file *only* when the client really needs it.

This particular distributed approach places a lot of responsibility on the client but speeds up the user's response. Most of the data is read from the local disk unless the original file is changed. Network traffic is reduced since local caching requires fewer file accesses and new copies of cached files are requested only if needed.

We also want to point out that *last save wins.* This means that if two users have opened the same file, the first user saves changes, then the second user saves the changes as well, then the second user overwrites changes made by the first user. AFS is not a revision control system.

In this section we describe to you an experience we had on the AFS environment we set up, in order to understand how AFS really works. We wanted to study the exchange of information between the AFS File Server machine and the AFS client machines when a file that the client requests has been modified. To do this we used the `tcpdump` command, launched on a first AFS machine, which in this case was the RISC/6000 rs600030.

First, we want to show you the particulars of the AFS environment we used in this experience:

Table 4. The Particulars of Our AFS Environment

Workstation	First AFS Machine	AFS Client	AFS Client
Operating System	AIX 4.3	AIX 4.3	Windows NT 4.0
IP address	9.24.104.97	9.24.104.124	9.24.104.223
Hostname	rs600030	rs600012	wtr05195

We logged into the Windows NT AFS Client machine wtr05195 as the AFS user afsuser. The AFS user afsuser was the owner of the home directory /afs/itso.ral.ibm.com/afsuser.

We list now the ports that are involved in this exchange of information, as they appear in the file /etc/services:

- port 7000

  This port is used by the file server fs trio of processes on the AFS File Server machine.

- port 7001

  This port is used on the AFS client machines for AFS callback.

- port 7003

  This port is used on the AFS server machine by the Volume Location Server process vlserver.

We followed this sequence of operations:

1. From the Windows NT AFS client machine wtr05195, we opened a text file located in /afs/itso.ral.ibm.com/afsuser.
2. We launched the `tcpdump` command on the AFS server, to monitor the traffic generated by the wtr05195 machine and received on the token-ring network interface tr0 of the AFS server:

   ```
   tcpdump -I -nt -i tr0 src 9.24.104.223
   ```

3. We modified and saved the text file we had opened.

This is the output we got as soon as we saved the file:

```
9.24.104.223.7001 > 9.24.104.97.7000: udp 44
9.24.104.223.7001 > 9.24.104.97.7000: udp 636
9.24.104.223.7001 > 9.24.104.97.7000: udp 58
9.24.104.223.7001 > 9.24.104.97.7000: udp 456
9.24.104.223.7001 > 9.24.104.97.7000: udp 57
9.24.104.223.7001 > 9.24.104.97.7000: udp 253
9.24.104.223.7001 > 9.24.104.97.7000: udp 57
```

You can see from the above screen that the wtr05195 machine, from its port 7001, informs the AFS Server listening on port 7000, that the text file has been modified. This is the first step of the callback process.

Of course, the AFS server immediately informs all the AFS clients that in the past had requested that particular file that the next time they will need that file they should retrieve a fresh copy of it, since the cached copy is out-of-date. We wanted to monitor this exchange of information too.

So, first we opened the file from the other AFS client machine that was part of our environment, rs600012, to ensure that such a file was stored in the cache of this client too. Then we repeated the same sequence of operations just described, but this time we issued the `tcpdump` command with different parameters, since we wanted to study all the network traffic exchanged between the three machines, and involving port 7001:

```
tcpdump -I -nt -i tr0 port 7001
```

We show you the output we got divided in two parts. The first part demonstrates once again that the Windows NT AFS client wtr05195 informs the AFS server that the file has been modified. This exchange of information is generated by the AFS client machine through a callback:

```
9.24.104.223.7001 > 9.24.104.97.7000: udp 291
9.24.104.97.7000 > 9.24.104.223.7001: udp 136
9.24.104.97.7000 > 9.24.104.223.7001: udp 136
9.24.104.223.7001 > 9.24.104.97.7000: udp 57
```

The second part of the output shows how the AFS server also informs the other AFS client rs600012 that the file it had retrieved and cached is now out-of-date:

```
9.24.104.97.7000 > 9.24.104.124.7001: udp 36
9.24.104.97.7000 > 9.24.104.124.7001: udp 36
9.24.104.124.7001 > 9.24.104.97.7000: udp 61
```

Through the above callback, the Cache Manager on the AFS client machine rs600012 has been made aware that the file it retrieved and cached in the past has now been changed, and it will be necessary to retrieve a fresh copy of the same file if some users or processes request that particular file. We wanted to study this process closely, again using the `tcpdump` command. So, we launched the following command on the AFS server machine rs600030:

```
tcpdump -I -nt -i tr0 src 9.24.104.124
```

Then we logged in as afsuser on the AIX machine rs600012, and we tried to open the file we had previously modified on wtr05195. This is the output we got on rs600030 after launching the above command:

```
9.24.104.124.7001 > 9.24.104.97.7003: udp 48
9.24.104.124.7001 > 9.24.104.97.7000: udp 44
9.24.104.124.7001 > 9.24.104.97.7000: udp 52
9.24.104.124.7001 > 9.24.104.97.7003: udp 61
9.24.104.124.7001 > 9.24.104.97.7000: udp 61
```

First of all, the AFS client machine rs600012, from its port 7001, contacts the Volume Location Server process, vlserver, listening on port 7003 on the AFS server machine. The reason for this call is that the Cache Manager knows that the file has been changed, and needs to know from the Volume Location Server the exact location of the file. In fact, the Volume Location Server maintains the Volume Location Database, which is aware of which File Server machine in the cell houses the volume that contains the file you are searching for.

After this first call, the AFS client machine knows the exact location of the file, and physically retrieves a fresh copy of it from the File Server machine.

7.4 Other AFS Client Issues

In this section we address some other AFS Client considerations, including cell membership, server preferences and client performance tuning.

7.4.1 Determining Cell Membership

For AFS environments that have access to multiple cells there are two commands for determining cell membership. To determine to which cell an AFS client belongs use the command:

```
fs wscell
```

A sample usage of this command is shown in the following session screen:

```
C:\> fs wscell
This workstation belongs to cell 'moshi.itso.ral.ibm.com'
```

To determine in which cell a directory or file resides use the command:

```
fs whichcell
```

A sample usage of this command is shown in the following session screen:

```
C:\> fs whichcell -p M:\moshi.itso.ral.ibm.com\user
File M:\moshi.itso.ral.ibm.com\user lives in cell 'moshi.itso.ral.ibm.com'
```

Note that these commands work for AIX, Solaris, and the Windows NT Command Prompt.

7.4.2 Examining and Setting Server Preferences

When it comes to selecting a VLDB replica or replicated File Server data, AFS clients use internal rankings of the various VL servers and File Servers to pick which machine to contact. For VL servers, a client will assign a value of 10,000 to each VL server listed in the CellServDB file, and then add a random value between 0 to 126 for a final ranking; VL servers are then ranked for preference from lowest to highest score.

For File Servers, the AFS Client will compare the File Servers' IP addresses to its own and compute a score based on the following algorithm:

- If the AFS client itself is a File Server, its base score is 5,000.
- If the AFS File Server is on the same subnetwork as the client, its base score is 20,000.

- If the AFS File Server is on the same network as the client, its base score is 30,000.
- If the AFS File Server is on a different network, or if the Cache Manager cannot otherwise obtain network information about the File Server, its base score is 40,000.

Once the base score is assigned to a File Server, the client further randomly adds a number between 0 and 14 in order to try to differentiate machines with the same base score. The results are the client's rankings of the File Servers about which it knows, with lower scores preferred. For both VL Servers and File Servers, the client will contact those machines with lowest scores first, and if those machines are not available, the client will subsequently work its way through higher-ranked machines.

The corollary to this ranking scheme is that clients can be *biased* to prefer particular VL servers or file servers. By setting specific ranking scores, clients will favor particular servers for VLDB and File Server access. This would be especially useful for directing clients to take advantage of network topologies and corresponding AFS Server locations, for example limiting the number of network hops an AFS client's communications must make to VL servers or file servers.

On the AFS client for Windows NT you can examine and set VL Server and File Server preferences through the AFS Client Configuration Control Panel. Under the Control Panel open up the AFS Client Configuration and click the **Preferences** tab. As you can see, the **Show File Server preferences** radio button is automatically checked:

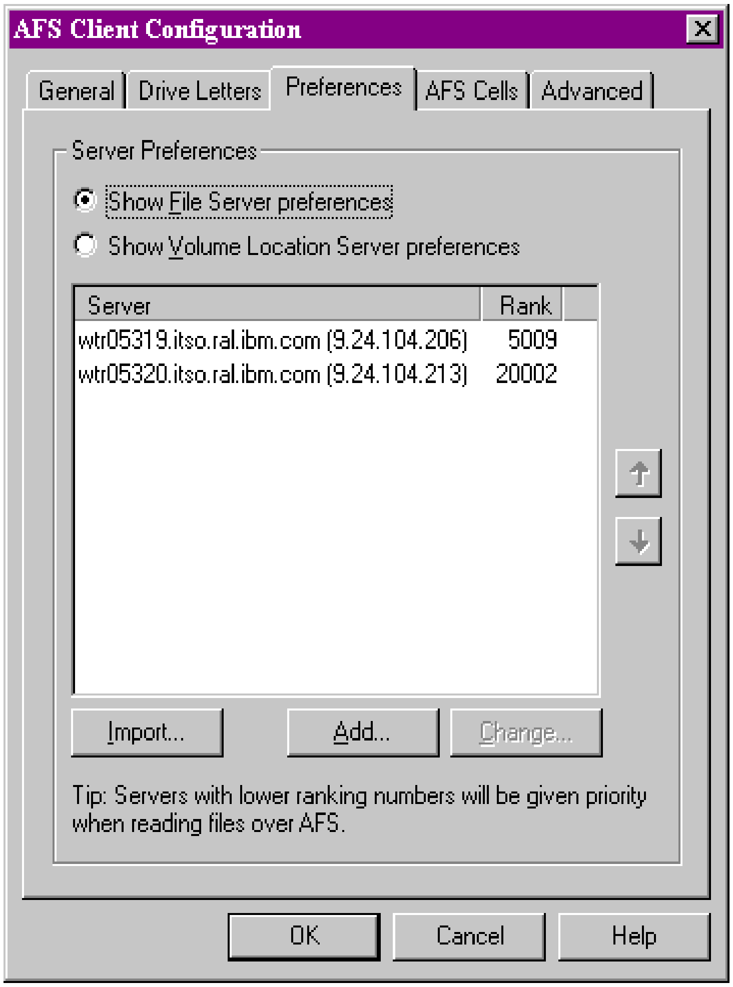

Server Preferences

● Show File Server preferences

○ Show Volume Location Server preferences

Server	Rank
wtr05319.itso.ral.ibm.com (9.24.104.206)	5009
wtr05320.itso.ral.ibm.com (9.24.104.213)	20002

Tip: Servers with lower ranking numbers will be given priority when reading files over AFS.

Figure 287. File Server Preferences

In our example we have two File Servers available, and because machine wtr05319.itso.ral.ibm.com has a score close to 5000, we can infer that we are in fact working on wtr05319. Furthermore because wtr05320.itso.ral.ibm.com has a score of 20,002, we deduce that it is on the same subnetwork as wtr05319 (as shown by their IP addresses as well). To specifically bias the client using server rankings, click the server whose preference you wish to change and click the **Change...** button. In our example we will change the rank of wtr05319.itso.ral.ibm.com:

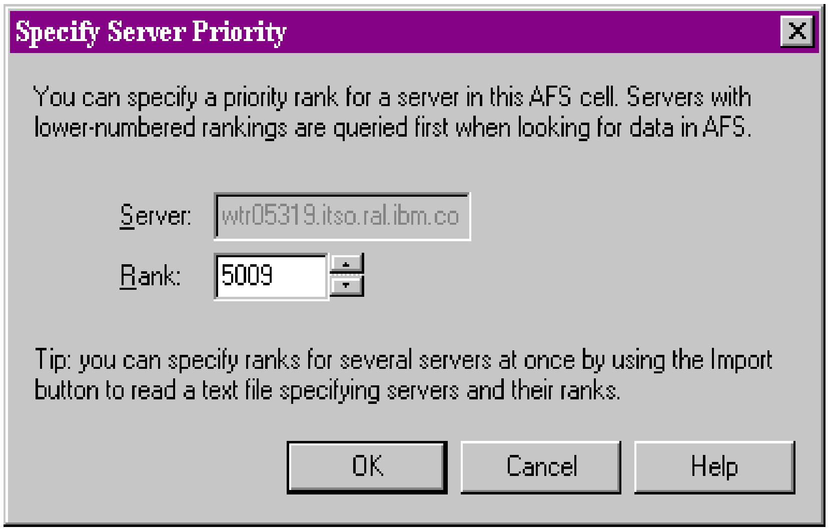
Specify Server Priority

You can specify a priority rank for a server in this AFS cell. Servers with lower-numbered rankings are queried first when looking for data in AFS.

Server: wtr05319.itso.ral.ibm.co

Rank: 5009

Tip: you can specify ranks for several servers at once by using the Import button to read a text file specifying servers and their ranks.

Figure 288. Changing File Server Rank

Set the desired rank and click **OK**. In our example we change the original rank 5,009 to 10,000. This change is reflected in the Preferences tab of the AFS Client Configuration window:

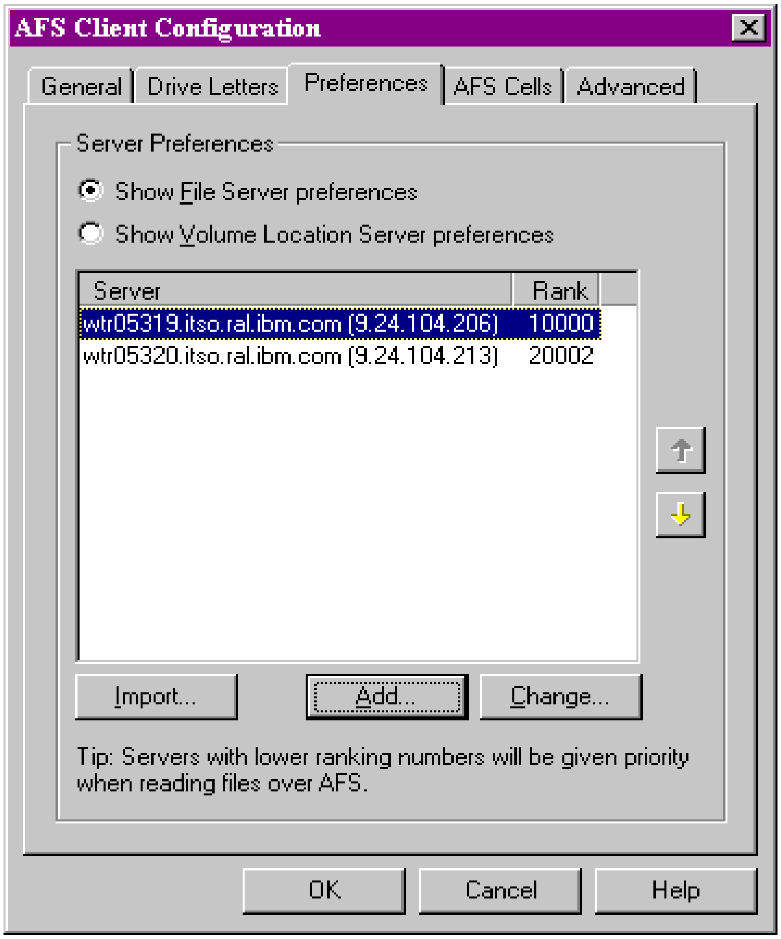
AFS Client Configuration

Server Preferences

● Show File Server preferences

○ Show Volume Location Server preferences

Server	Rank
wtr05319.itso.ral.ibm.com (9.24.104.206)	10000
wtr05320.itso.ral.ibm.com (9.24.104.213)	20002

Tip: Servers with lower ranking numbers will be given priority when reading files over AFS.

Figure 289. Changed File Server Rank

Note that with only two File Servers in our example there is no real effect from the change we made. This change is also persistent; in other words, it will survive a restart of the IBM AFS client service or a reboot. Use the **Add...** button to add a File Server to the ranking list, and use the **Import...** button to import a file of rankings, where the format is, per line:

```
file_server_name desired_numerical_rank
```

The two fields above are separated by some space.

To display the VL Server rankings click the other Server Preferences radio button, **Show Volume Location Server preferences**, and use the **Change...**, **Add...** and **Import...** buttons to make any changes:

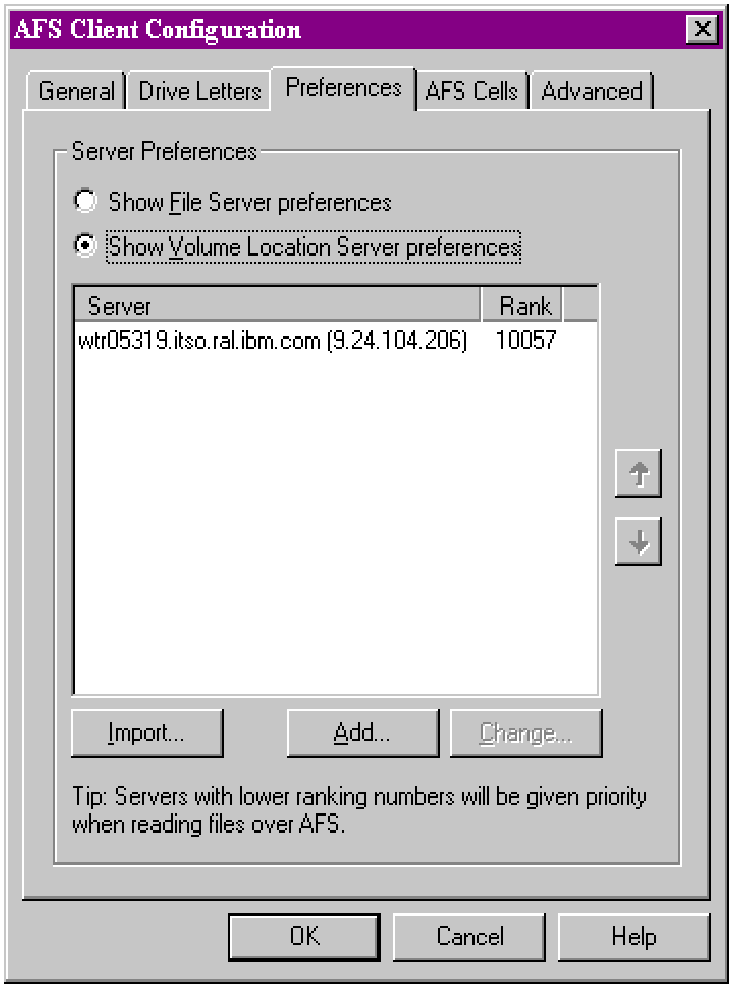

Server	Rank
wtr05319.itso.ral.ibm.com (9.24.104.206)	10057

Tip: Servers with lower ranking numbers will be given priority when reading files over AFS.

Figure 290. VL Server Preferences

On AIX, Solaris, and the Windows NT Command Prompt use the commands:

```
fs getserverprefs
fs setserverprefs
```

to get and set server preferences.

Notice that using the command:

fs setserverprefs

only temporarily sets the server preferences; that is the preferences are lost on a reboot. However this command can take input from a file that contains a list of servers and desired ranks or it can take input from standard input. The command could then be inserted into a script run at boot time to set the server preferences from a file, or from another program or script that for example outputs a list of the closest servers based on the results of running ping.

In the following example we get file server preferences, both as host names and as IP addresses; we then follow with VL Server preferences. Finally we demonstrate how to use an input file to set file server preferences and we display both the results and the contents of the input file.

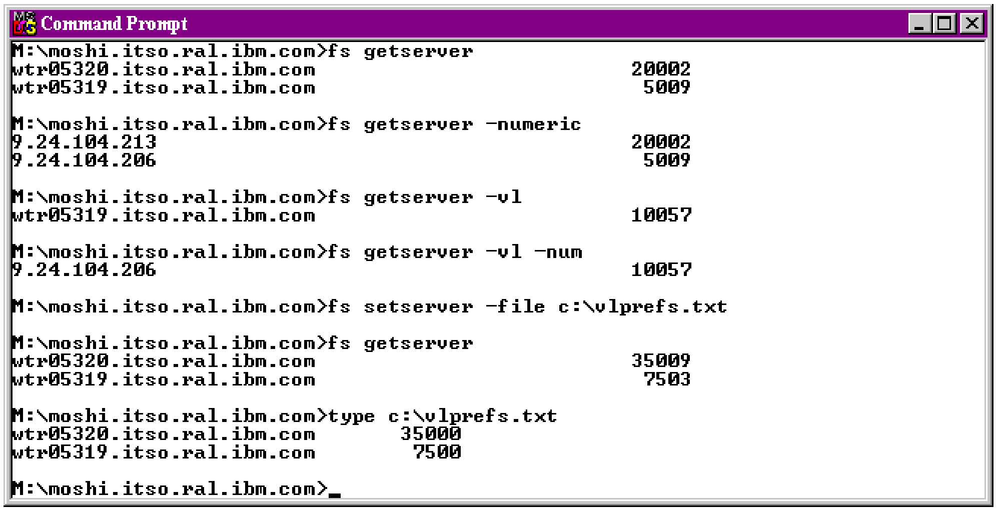

```
M:\moshi.itso.ral.ibm.com>fs getserver
wtr05320.itso.ral.ibm.com                          20002
wtr05319.itso.ral.ibm.com                           5009

M:\moshi.itso.ral.ibm.com>fs getserver -numeric
9.24.104.213                                       20002
9.24.104.206                                        5009

M:\moshi.itso.ral.ibm.com>fs getserver -vl
wtr05319.itso.ral.ibm.com                          10057

M:\moshi.itso.ral.ibm.com>fs getserver -vl -num
9.24.104.206                                       10057

M:\moshi.itso.ral.ibm.com>fs setserver -file c:\vlprefs.txt

M:\moshi.itso.ral.ibm.com>fs getserver
wtr05320.itso.ral.ibm.com                          35009
wtr05319.itso.ral.ibm.com                           7503

M:\moshi.itso.ral.ibm.com>type c:\vlprefs.txt
wtr05320.itso.ral.ibm.com       35000
wtr05319.itso.ral.ibm.com        7500

M:\moshi.itso.ral.ibm.com>_
```

Figure 291. Using the Command Prompt to Set Preferences

7.4.3 AFS Client Tuning

AFS offers a number of parameters for tuning client performance and behavior. One such parameter – server preferences – has already been discussed (see 7.4.2, “Examining and Setting Server Preferences” on page 327). A more detailed discussion of AFS performance and stability tuning is beyond the scope of this book but we will at least mention some points in passing:

- The AFS client for Windows NT provides a few tuning parameters in the AFS Client Configuration Control Panel, which are accessed by selecting the **Advanced** tab:

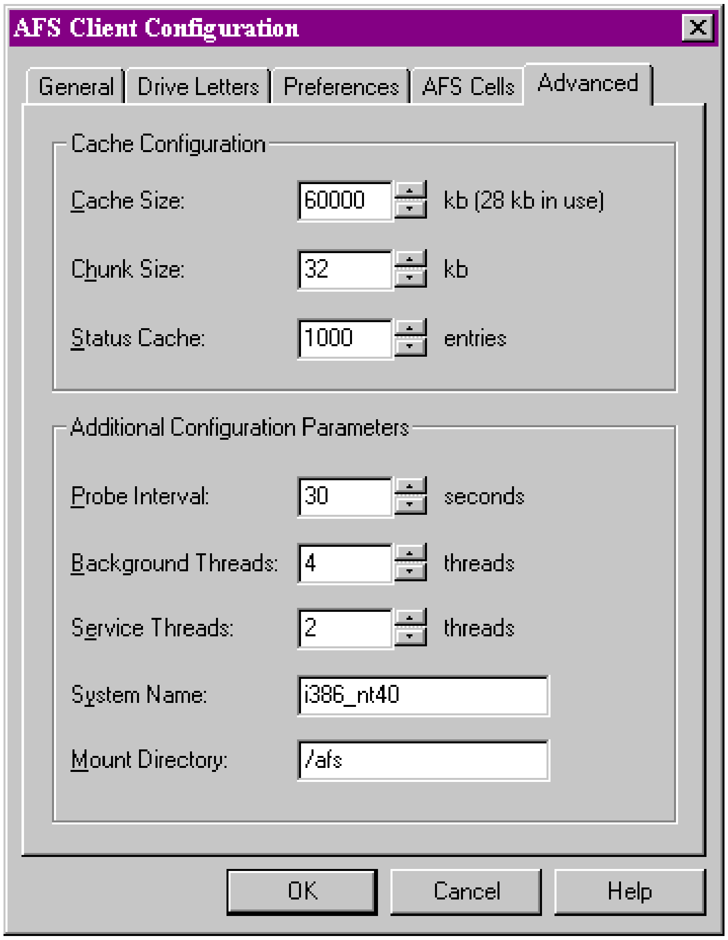

Cache Configuration

Cache Size: 60000 kb (28 kb in use)

Chunk Size: 32 kb

Status Cache: 1000 entries

Additional Configuration Parameters

Probe Interval: 30 seconds

Background Threads: 4 threads

Service Threads: 2 threads

System Name: i386_nt40

Mount Directory: /afs

Figure 292. Advanced Client Parameters

The following parameters affect performance:

- The Cache Size of course determines how much AFS data can be cached locally on the client. The obvious consideration in determining the cache size is the amount of available disk space on the client, but the cache size should not be made extraordinarily large. As mentioned before factors such as the number of users on the machine and the type of data being accessed should be considered when determining cache size. For example, a Windows NT workstation being used for office applications might want to have a cache no larger than 30-60 MB, but a workstation being used for CAD/CAM applications and drawings may benefit from a much larger cache.
- The Chunk Size determines the size of the data chunks sent from the File Server. A small chunk size requires more memory usage by the client and may lead to performance problems, but a larger chunk size may waste cache space if a lot of the cached files are smaller than the chunk size setting. Typically this setting is adjusted upward for both large data files (for example, CAD/CAM drawings) and fast network technology (ATM, 100 Mbps Ethernet, etc.).
- The Status Cache setting determines the number of cached files about which the Cache Manager retains status information. Adjust this setting upward if the client is caching many small files as opposed to fewer large files.

- The Probe Interval sets how often the client contacts the File Servers to make sure they are still responding.
- The Background Threads handle background tasks for the Cache Manager.
- The Service Threads handle all initial requests for data; applications could possibly be delayed while waiting for a service thread to complete.

With AIX and Solaris, there are more parameters for tuning AFS client performance than there are for Windows NT, either via the AFS Client Configuration Control Panel or the Command Prompt. The first consideration is the AFS client cache itself, as it can be used as either a disk-based cache or, if sufficient RAM is available, as a memory-based cache. Also the AFS client cache on UNIX systems can typically be set larger than what Windows NT can accommodate, and for cases of large file and data set access, cache sizes of 500 MB to 1 GB or more are possible. Moreover, the Cache Manager process, afsd, accepts a number of different arguments that influence its behavior and performance. Some of these include:

- `-memcache` initializes a memory cache rather than a disk cache.
- `-blocks` is the number of 1024-byte blocks in the cache and so overrides the cache size specified in the /usr/vice/etc/cacheinfo file.
- `-files` sets the number of V-files created in the cache directory, which overrides the default that is calculated automatically by afsd. Upon startup afsd will set the number of V-files used in the cache depending on whether a memory cache or disk cache is used. If a memory cache is used the number of chunks (V-files) will simply be the cache size divided by the chunk size. For a disk cache the number of chunks will be the largest of the following:
  - 100
  - (*cachesize* / *chunksize*) x 1.5
  - *cachesize* / 10,240
- `-stat` is the number of entries stored in memory for recording status information about files in the cache; the default is 300.
- `-daemons` sets the number of background daemons to run on the client. These daemons improve performance by doing prefetching and background writing of saved data. The default value is 2, which is adequate for a machine with approximately 5 concurrent users. Values greater than 6 are not generally more effective than 6.
- `-chunksize` sets the size of chunks (or V-files) in the client cache, expressed as a power of 2; the default is 16 for disk cache (2^{16} bytes = 64 KB) and 13 for RAM cache (2^{13} bytes = 8 KB). This will also be the size of the data chunks that the client requests from File Servers.
- `-dcache` sets the number of dcache entries in memory, which are used to store information about cache chunks. For a disk cache this overrides the default, which is half the number of V-files; its use is not recommended for memory caches.
- `-volumes` is the number of memory structures allocated for storing volume location information; the default is 50.
- `-biods` sets the number of VM daemons dedicated to performing I/O operations on machines running AIX with virtual memory (VM) integration.

Changing or setting any one of these parameters on an AFS client is liable to affect performance, so care must be taken to ensure proper performance monitoring and evaluation. For more information on AFS client (and AFS server) performance tuning see the *AFS System Administrator's Guide* and the *AFS File Server Operations Tuning and Troubleshooting Guide*, available on the IBM Transarc Lab's Product Support Web site at `http://www.transarc.com/Support/`.

Chapter 8. AFS User Account Management

The final stage in setting up an AFS cell is to add users and groups to the AFS environment. In a typical WebSphere Performance Pack scenario, AFS clients are used by Web authors to update Web content, server administrators to manage the file system and servers, or Web Traffic Express (WTE) proxy servers to share the same cache through the Remote Cache Access (RCA) protocol.

AFS for Windows NT includes the AFS Account Manager tool as a graphic interface for user and group management. For AIX, Solaris, and the Windows NT Command Prompt, the `pts` and `kas` command suites are used to accomplish the same tasks.

AFS for AIX and Solaris also comes with the `uss` command suite, which completes many user account creation tasks with a single command and can even utilize a template file for creating accounts in bulk. This command suite is not available on Windows NT.

Notice that, when you use the command:

```
uss add
```

to create a single account, you usually have a template file that indicates what the account should include (for example, a volume, mount point, login files, etc.). Without this template file, the command above just creates entries in the Authentication and Protection Database. When you use the command:

```
uss bulk
```

to create multiple files, you need an input file that indicates which accounts you want to create and delete. For each account being created, you can also use a template file (same one as you used before) to define what the account will have.

AFS user and password information is stored in the Authentication Database and served by the Authentication Server, which is any AFS server running the kaserver process. Users authenticate against the Authentication Database to obtain tokens for authorized access to data in AFS. The Protection Server (ptserver) is responsible for checking authorization when users access data and for also storing group membership information. But before any users can authenticate, they must have an account in the Authentication Database. Typically the components of an AFS user account include:

- A Protection Database entry, for example `bobby`
- An Authentication Database entry which must match the Protection Database entry, for example `bobby`
- A user volume, for example user.bobby
- A home directory in which the volume is mounted, for example /afs/moshi.itso.ral.ibm.com/usr/bobby
- Ownership of the home directory and full permissions on its ACL
- An entry in the local password file (the UNIX /etc/passwd password file, a Windows NT domain, or equivalent) of each machine the user needs to log into

- Optionally, standard files and subdirectories that you as an administrator have provided to make the account more useful (such as .login and .profile on UNIX, and a user subdirectory tree in C:\WINNT\Profiles on Windows NT)

Note that a user account in AFS does not necessarily need anything other than an entry in the Protection and Authentication databases. This means that an account can be created specifically for authentication and authorization purposes, but not have any volume or login information associated with it. With WebSphere Performance Pack, for example, it may be useful to set up directories in AFS that restrict access to Web servers or application servers, or to specific subsets of servers (intranet vs extranet). In these cases, accounts would be set up for the machines for authentication purposes but would not be used for interactive login sessions.

8.1 AFS Users on Windows NT Using the Account Manager

We begin our examples by bringing up the AFS Account Manager in Windows NT to examine current user accounts and to add, change or delete existing accounts. Under the Start menu select **Programs**, then **IBM WebSphere**, then **Performance Pack**, then **Afs**, then **Control Center**, and finally **Account Manager**, as shown in the following figure:

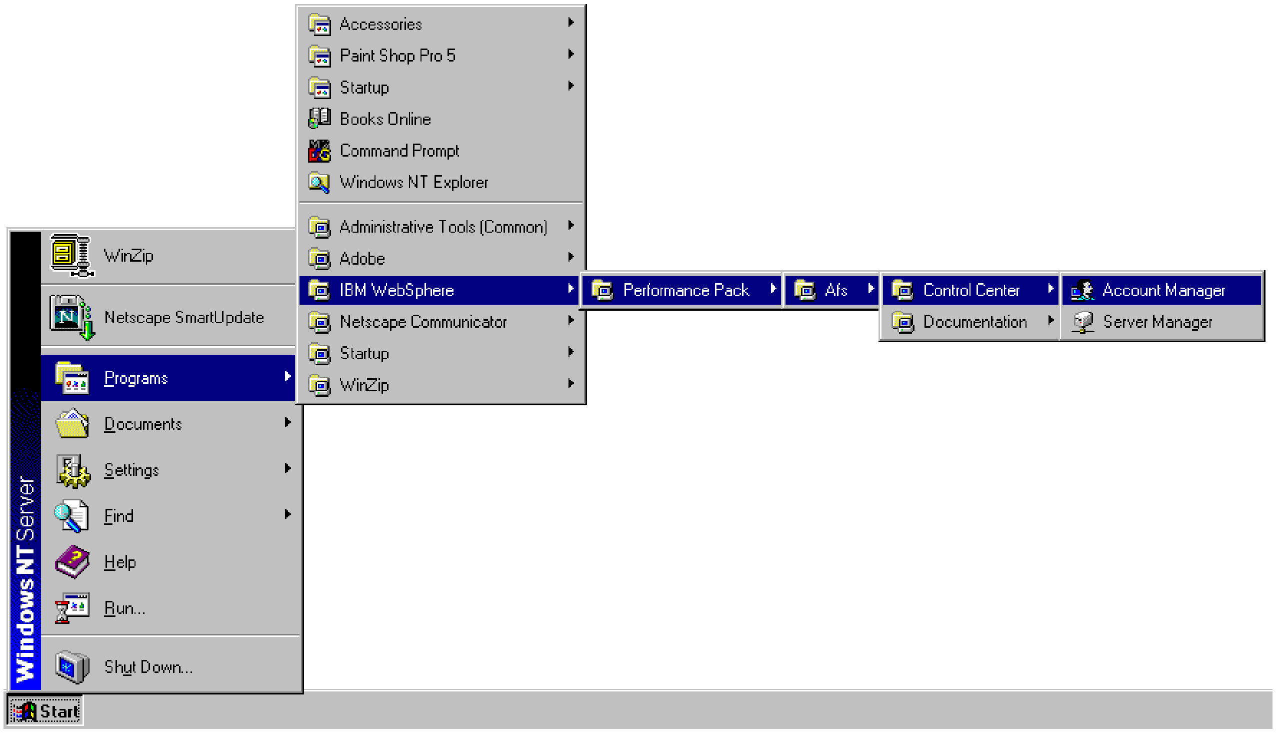

Figure 293. Launching the AFS Account Manager

In this case we have launched the AFS Account Manager on a machine in the moshi.itso.ral.ibm.com cell. After the AFS component of WebSphere Performance Pack is first installed, there are several default accounts automatically created, most notably the admin account for performing AFS administrative tasks. The other existing default accounts are for authentication of various services within AFS and are not meant to be interactive.

Note some of the options available for user accounts, including password policies and account expiration:

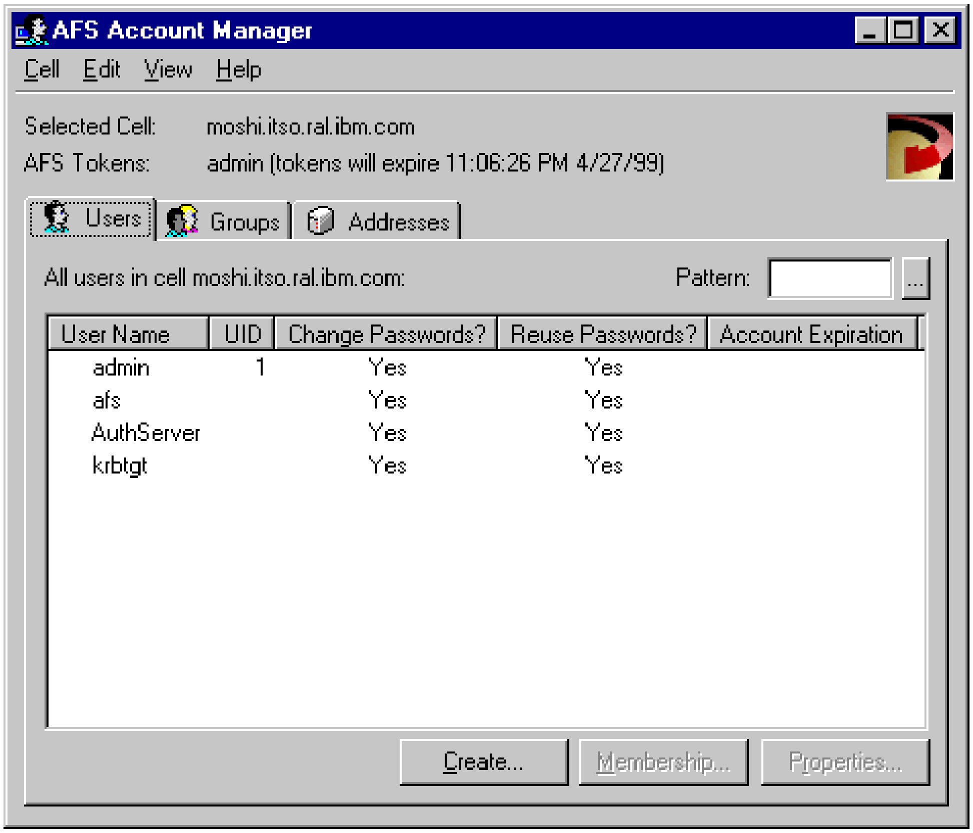

Selected Cell: moshi.itso.ral.ibm.com
AFS Tokens: admin (tokens will expire 11:06:26 PM 4/27/99)

All users in cell moshi.itso.ral.ibm.com:

User Name	UID	Change Passwords?	Reuse Passwords?	Account Expiration
admin	1	Yes	Yes	
afs		Yes	Yes	
AuthServer		Yes	Yes	
krbtgt		Yes	Yes	

Figure 294. List of Initial Accounts in AFS Cell

In this example we will create an account called wspp that will be an *authorization only* AFS account; this means that, rather than be an interactive user account with an associated AFS volume, it will be used primarily as an owner of groups of other user accounts.

Click on the **Create...** button to add a new AFS user account and enter the account name and a password for the new account. Unless the account is to be mapped to an existing UNIX account, select **Automatically use the next available user ID** (the user ID to be selected appears in parentheses). Because no groups have been created other than existing default groups, ignore the **Groups...** button for the moment. A sample configuration is shown in the following figure:

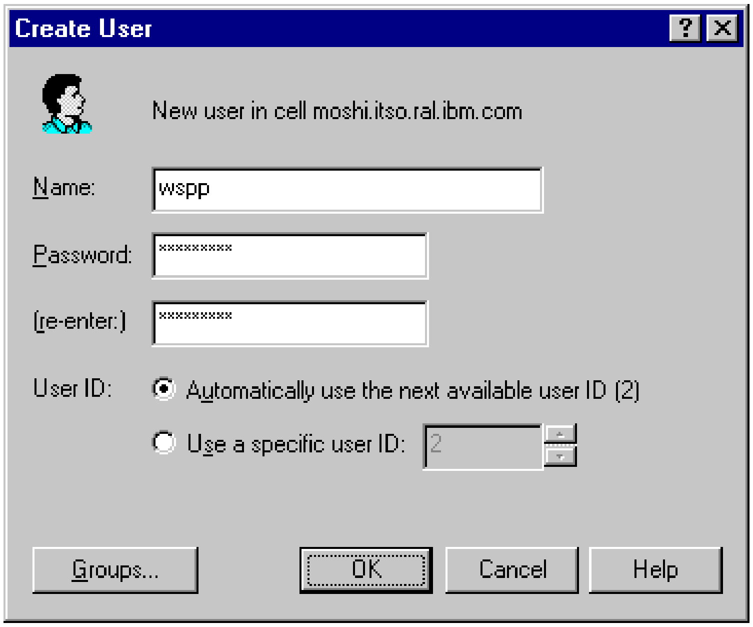

New user in cell moshi.itso.ral.ibm.com

Name: wspp

Password: **********

(re-enter:) **********

User ID: ● Automatically use the next available user ID (2)

○ Use a specific user ID: 2

Figure 295. User Account Creation

Click **OK** when finished; this will return you to the AFS Account Manager window. As you can see, the wspp user shows up now as one of the defined users:

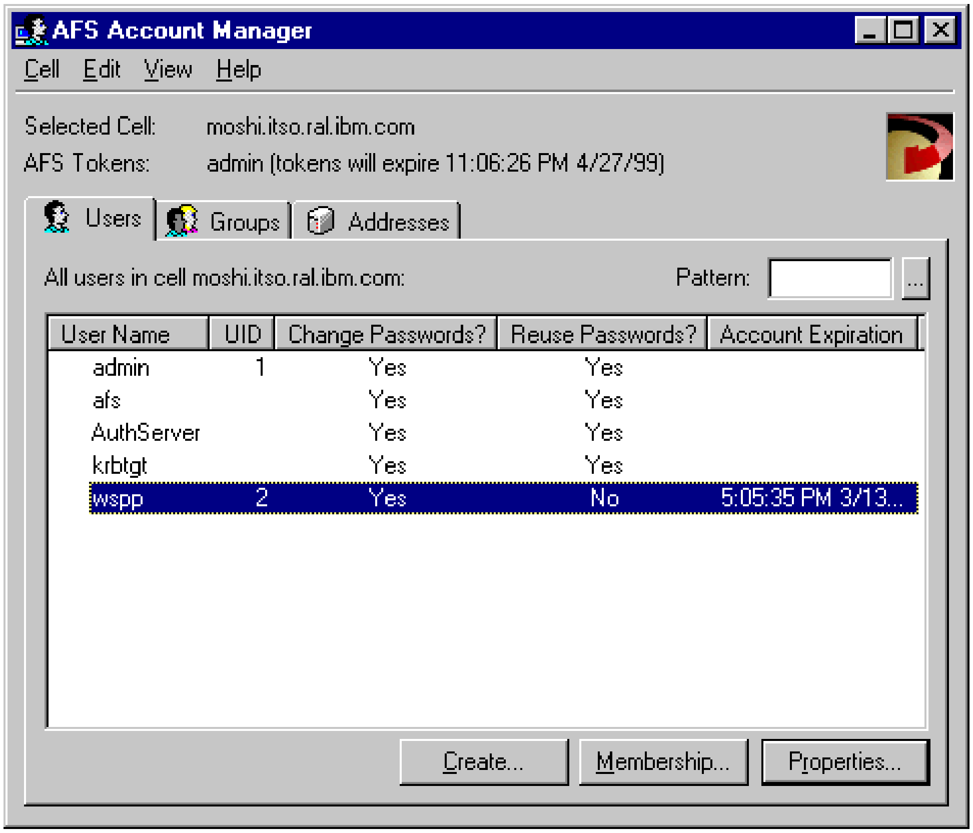

Selected Cell: moshi.itso.ral.ibm.com

AFS Tokens: admin (tokens will expire 11:06:26 PM 4/27/99)

All users in cell moshi.itso.ral.ibm.com:

User Name	UID	Change Passwords?	Reuse Passwords?	Account Expiration
admin	1	Yes	Yes	
afs		Yes	Yes	
AuthServer		Yes	Yes	
krbtgt		Yes	Yes	
wspp	2	Yes	No	5:05:35 PM 3/13...

Figure 296. User Account Created

Note the default account properties that are entered for the account: an account expiration is set and password reuse is turned off. Turning password reuse off prevents users from reusing their previous 20 passwords.

To change account properties such as account expiration, failed login attempts and password policy select the desired account to change and click **Properties...**. This will bring up the Properties window for that specific AFS account:

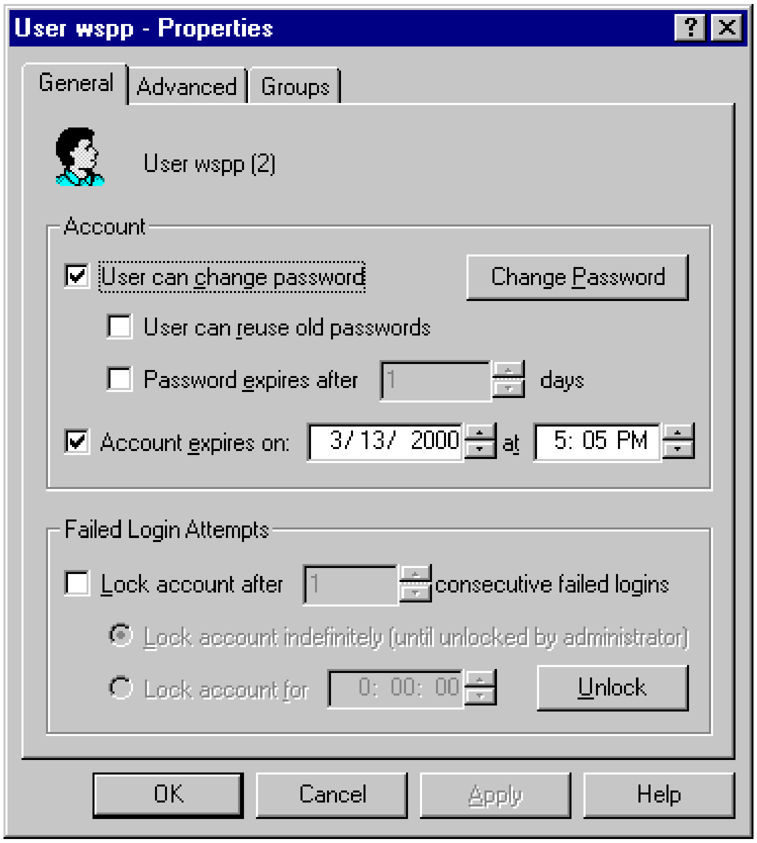
User wspp - Properties

User wspp (2)

Account

☑ User can change password

□ User can reuse old passwords

□ Password expires after 1 days

☑ Account expires on: 3/13/ 2000 at 5: 05 PM

Failed Login Attempts

□ Lock account after 1 consecutive failed logins

● Lock account indefinitely (until unlocked by administrator)

○ Lock account for 0: 00: 00

Figure 297. User Account General Properties

Again, note that though we have created an AFS user account called wspp, we have created no volume for this user. Someone could login and authenticate as wspp, but unless AFS directory protections are set up to grant this user access, their use of the AFS directory tree is limited.

In our next example we create a new AFS user bobby. Recall that in 6.4.2, “Creating AFS Volumes” on page 273, we had as an example created a user.bobby volume and mounted it under /afs/moshi.itso.ral.ibm.com; we are now creating a user to associate with this home volume. In typical practice, user accounts, home volumes and directories are either created simultaneously or user accounts are created first.

To create this new user account go into the Account Manager and click the **Create...** button. This will bring up the Create User window:

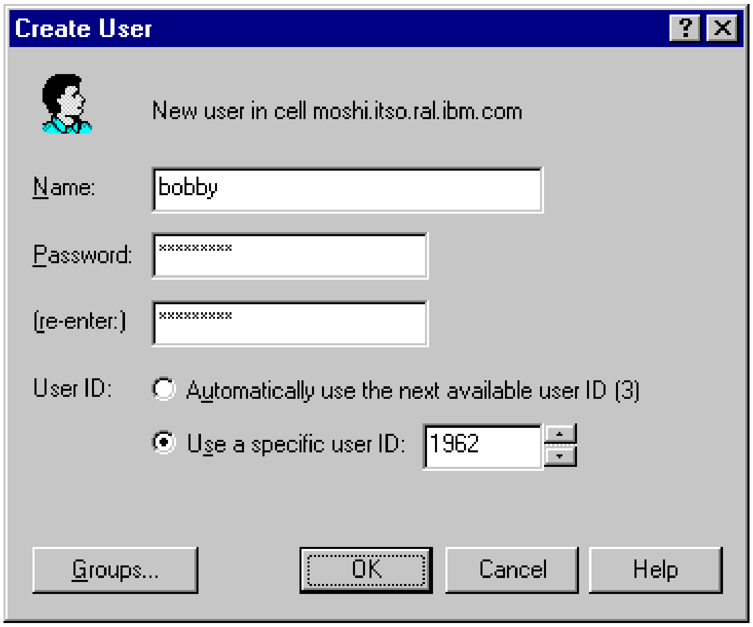

Name: bobby

Password: **********

(re-enter:) **********

User ID: ○ Automatically use the next available user ID (3)

● Use a specific user ID: 1962

Figure 298. Creating Account for User bobby

One thing we do differently with this account creation example is that, this time, we specify a user ID. If you look back at Figure 295 on page 340, you will notice that when we created the wspp user we had set the option to automatically use the next available user ID.

Using a specific user ID, as we are doing now for the user bobby, is necessary in the case where a user may login also to an AIX or Solaris machine. It is *highly recommended* that AFS user IDs and UNIX user IDs (found in the /etc/passwd file) be kept synchronized. One important reason for this is that AFS uses AFS user IDs to assign ownership of directories and files in AFS space, but UNIX commands such as

```
ls -l
```

will automatically compare the AFS user ID values to the user ID values in the /etc/passwd file, and print out the user name found there. If there is no direct correspondence between UNIX user IDs and AFS user IDs, a command such as:

```
ls -l
```

could for example report incorrect file and directory ownership. Windows NT users do not have this consideration, but the AFS administrator has to bear this in mind.

In general, if a UNIX account already exists for a user, then when creating the AFS account, you should assign the user the same user ID number. If no UNIX account exists (yet), allow the account creation process to automatically assign the user ID number. In the case of Windows NT, the only consideration is that in order for integrated login to work, the AFS user name and password *must* match the Windows NT user name and password, whether the Windows NT account is in the local Security Registry on the Windows NT machine or in a Windows NT domain.

In the case of user bobby we specifically assign this user the AFS user ID 1962, which corresponds to the UNIX user ID. In fact this user will possibly be accessing AFS through AIX or Solaris machines. The results of our account creation are shown below:

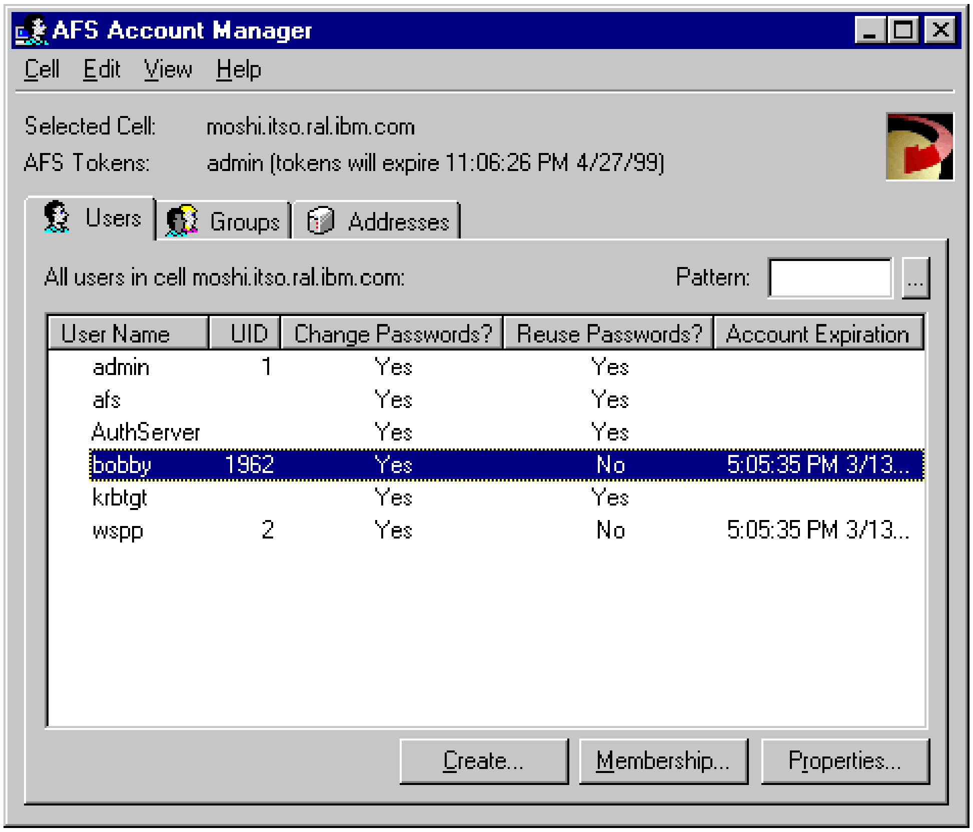

Selected Cell: moshi.itso.ral.ibm.com
AFS Tokens: admin (tokens will expire 11:06:26 PM 4/27/99)

All users in cell moshi.itso.ral.ibm.com:

User Name	UID	Change Passwords?	Reuse Passwords?	Account Expiration
admin	1	Yes	Yes	
afs		Yes	Yes	
AuthServer		Yes	Yes	
bobby	1962	Yes	No	5:05:35 PM 3/13...
krbtgt		Yes	Yes	
wspp	2	Yes	No	5:05:35 PM 3/13...

Figure 299. Account for User bobby

Again, specifically assign AFS user IDs in the case where UNIX user IDs exist on UNIX machines in the cell; otherwise, for a brand new environment or one in which no UNIX user IDs exist, allow the user creation process (via AFS Account Manager on Windows NT, the kas or pts command suites on AIX, Solaris, and the Windows NT command prompt, and the uss command suite on AIX and Solaris) to assign AFS user IDs.

8.2 AFS Users in AIX, Solaris and the Windows NT Command Prompt

To create AFS user accounts on AIX or Solaris (not in the Windows NT Command Prompt window), an administrator can use the uss command suite to create the user account, user volume, and a mount point to the user volume, all with one command. For example:

```
% uss add -user tom -realname "Tom Menner" -pass tjones -server wtr05319 \
-partition vicepa -mount /afs/moshi.itso.ral.ibm.com/usr/tom -uid 2341
```

This command creates a user named tom, assigns his password as `tjones`, and creates a volume called user.tom on server wtr05319, partition /vicepa; this volume is mounted in /afs/moshi.itso.ral.ibm.com/usr as tom. Enter the command:

```
uss add -h
```

to see the full list of options that can be set.

Alternatively, user accounts can be created on AIX, Solaris, and the Windows NT Command Prompt window using the commands:

```
pts createuser
kas create
```

The following session screen shows how to create the AFS user account tom from the command prompt:

```
C:\> pts createuser -name tom -id 2341
User testme has id 2341

C:\> kas create tom -admin admin
Administrator's (admin_user) password:
initial_password:
Verifying, please re-enter initial_password:

C:\>
```

The `pts` command suite in this case is used to create the user entry in the Protection Database, and the `kas` command suite is used to create the corresponding entry in the Authentication Database.

Note in the above examples that we specify a user ID for the account we create. Recall from the discussion in 8.1, “AFS Users on Windows NT Using the Account Manager” on page 338, that when creating new user accounts both the AFS Account Manager on Windows NT and the account creation commands will either select the next available user ID number for an account or they will let you specify a number. In practice it is *highly recommended* that AFS user ID numbers correspond to the UNIX user ID numbers, but if UNIX user IDs are not a factor, then you can safely allow the account creation process to automatically assign AFS user IDs.

Use the commands:

```
kas examine
kas setfields
```

to examine and change account information in the Authentication Database. These commands can be launched on AIX, Solaris and the Windows NT Command Prompt window.

In the following session screen, we change tom’s user account to allow reuse of old passwords (notice that this could be a security risk):

```
C:\>kas examine -name tom -admin_username admin
Administrator's (admin) Password:

User data for tom (ADMIN+NOSEAL)
  key (0) cksum is 3156620375, last cpw: Mon May 24 17:21:38 1999
  password will never expire.
  An unlimited number of unsuccessful authentications is permitted.
  entry expires on Wed May 24 17:21:27 2000.  Max ticket lifetime 25.00 hours.
  last mod on Mon May 24 17:21:38 1999 by admin
  don't permit password reuse

C:\>kas setfields -name tom -reuse yes -admin_username admin
Administrator's (admin) Password:

C:\>
```

8.3 Managing AFS Groups

ACL entries on AFS directories can contain two types of entities: *users* and *groups* of users and machines. By default, AFS users are allowed to define their own groups, which are of the form *username*:*groupname* (for example, bobby:friends). Administrators can define these groups as well as prefix-less groups of the form *groupname*, for example accounting. Finally, there are three default system groups automatically defined in AFS:

- system:anyuser is the group of all users, authenticated or not, that can access the AFS cell. This group does not actually have members assigned to it but instead generically references all users able to access the cell. This group might be placed on the access control list (ACL) of an external public FTP directory, for example.
- system:authuser is the group of all users that specifically have authenticated to the local cell. This means that the users of this groups have AFS tokens for the local cell. As with system:anyuser, this group does not have specific members; instead any authenticated user (or, in other words, any user with AFS tokens) is implicitly a member of this group.
- system:administrators is the group of users with AFS administrative privileges. This group is automatically assigned with administer and lookup rights to every ACL in AFS.

Groups are defined in the Protection Database and are administered using the AFS Account Manager on Windows NT or using the `pts` suite of commands on AIX, Solaris, and the Windows NT Command Prompt window.

8.3.1 How to Use the Windows NT Account Manager

In this example we will create wspp:web-servers, a group that is to be owned by the wspp user account we created in 8.1, “AFS Users on Windows NT Using the Account Manager” on page 338. Begin by launching the AFS Account Manager from the IBM WebSphere menu tree. Once in the Account Manager, you need to make sure that you are authenticated as the wspp user account. The Account Manager will not let you create a group owned by someone else; for example, even though you are authenticated as admin and belong to the system:administrators group, you cannot just create a group called wspp:web-servers. Nonetheless, you can create the group web-servers and then subsequently transfer its ownership to wspp.

In this case, however, we will simply change our authentication to that of wspp. We are already authenticated as admin but we want to discard that authentication and instead obtain tokens for the wspp user account. In the AFS Account Manager select the Cell menu and select **Authentication...**, as shown in the following window:

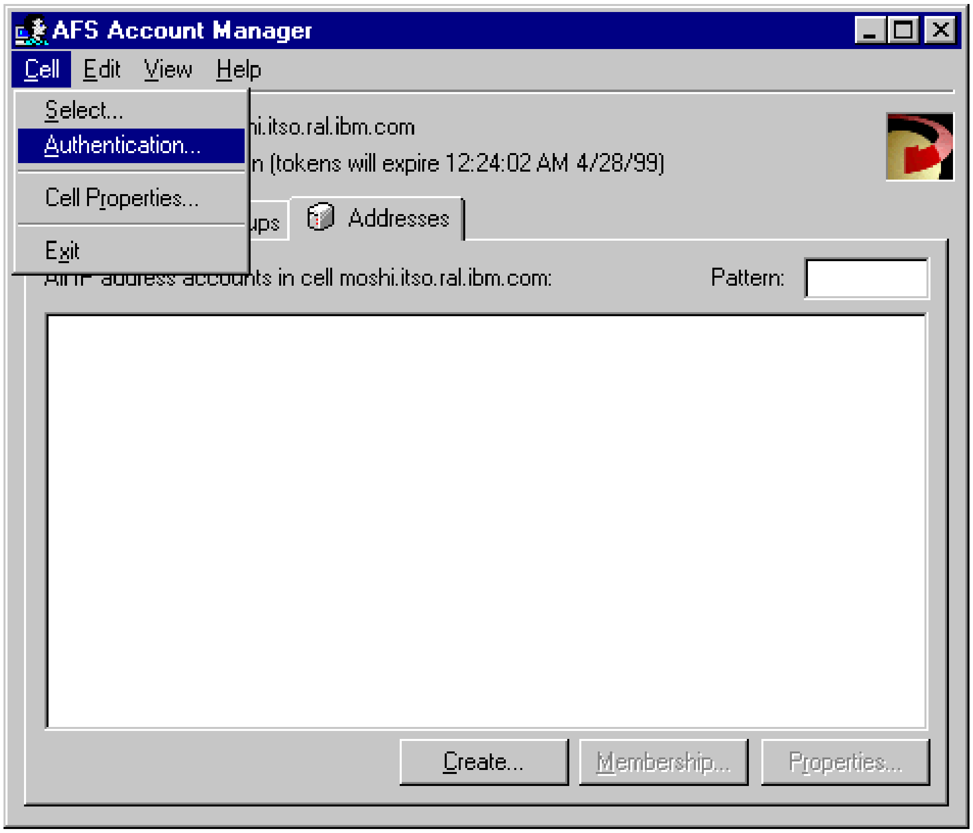

hi.itso.ral.ibm.com

n (tokens will expire 12:24:02 AM 4/28/99)

All IP address accounts in cell moshi.itso.ral.ibm.com:

Figure 300. Authenticating in the AFS Account Manager

This will bring up the AFS Authentication - AFS Account Manager window:

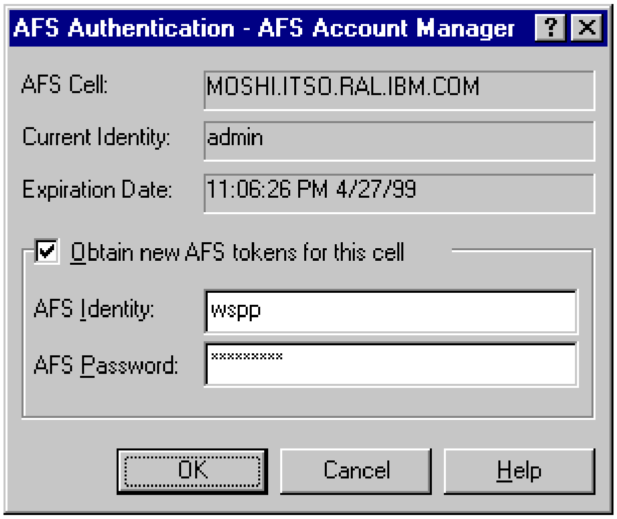

AFS Cell: MOSHI.ITSO.RAL.IBM.COM

Current Identity: admin

Expiration Date: 11:06:26 PM 4/27/99

☑ Obtain new AFS tokens for this cell

AFS Identity: wspp

AFS Password: **********

Figure 301. AFS Account Manager Authentication

As you can see, we are already authenticated as admin so we click **Obtain new AFS tokens for this cell** and enter the AFS identity wspp and its password and click **OK**.

After authenticating, we click the **Groups** tab in the AFS Account Manager window, and because we are authenticated as wspp, and do not have administrative privileges, we do not see any of the system groups that already exist, such as system:administrators:

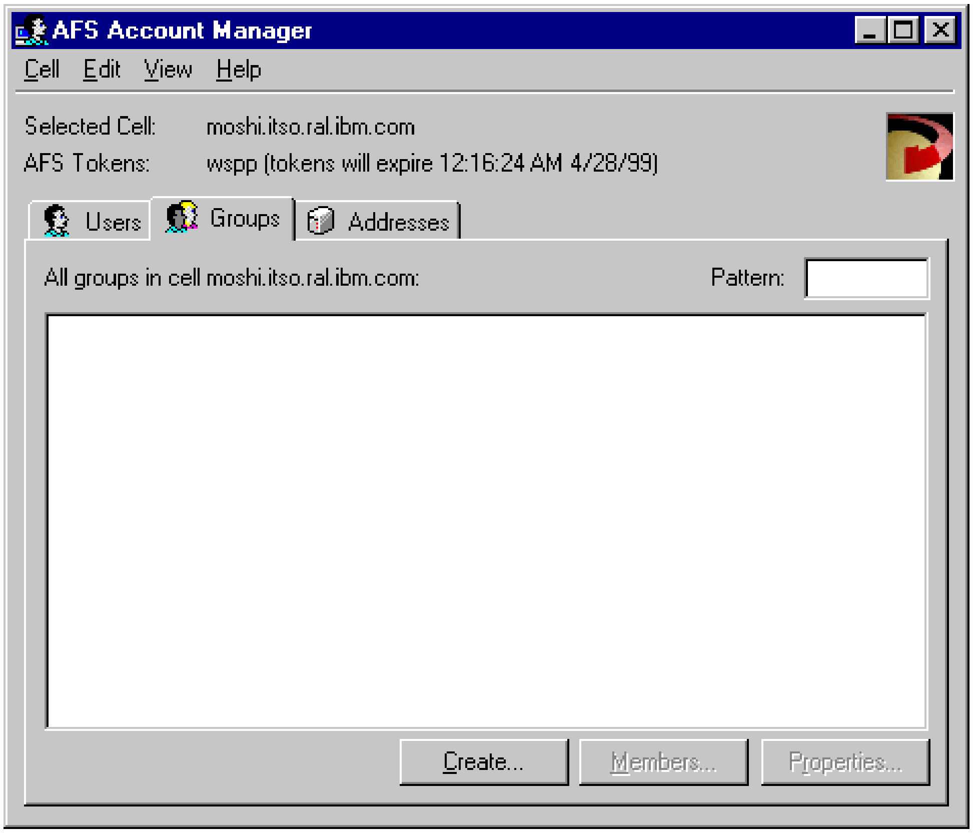

Selected Cell: moshi.itso.ral.ibm.com

AFS Tokens: wspp (tokens will expire 12:16:24 AM 4/28/99)

All groups in cell moshi.itso.ral.ibm.com:

Pattern:

Figure 302. Groups in the AFS Account Manager

Click the **Create...** button to bring up the Create Group dialog box and create a new group. In this case we are creating a group whose owner is the user wspp, so the group name will be wspp:web-servers.

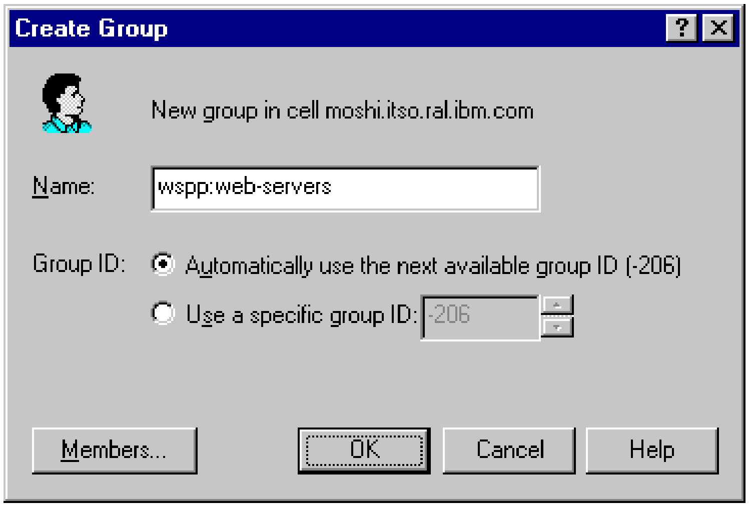

New group in cell moshi.itso.ral.ibm.com

Name: wspp:web-servers

Group ID: ● Automatically use the next available group ID (-206)

○ Use a specific group ID: -206

Figure 303. Create Group Window

Typically you should let the group creation process automatically assign a group ID (which will be a negative integer), but there is otherwise no correlation between AFS groups and groups in UNIX or Windows NT. If we had other user accounts created we could assign members now by clicking the **Members...** button, but in this case we click **OK**.

The new group is now listed in the AFS Account Manager, and wspp appears as the owner of the group. Note that there are currently no members assigned to the group:

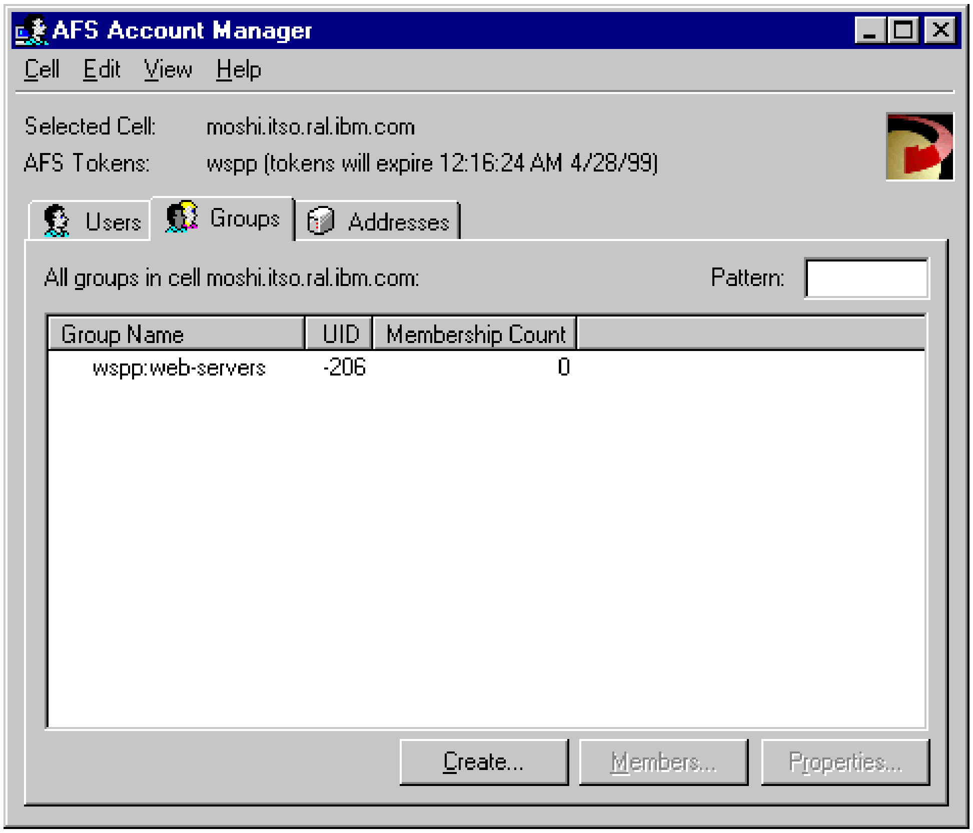

Selected Cell: moshi.itso.ral.ibm.com
AFS Tokens: wspp (tokens will expire 12:16:24 AM 4/28/99)

All groups in cell moshi.itso.ral.ibm.com:

Pattern:

Group Name	UID	Membership Count
wspp:web-servers	-206	0

Figure 304. The Newly Created Group

We finish this example by reauthenticating as the AFS admin identity by clicking the Cell menu and selecting **Authentication...**, and re-entering the user account information for admin:

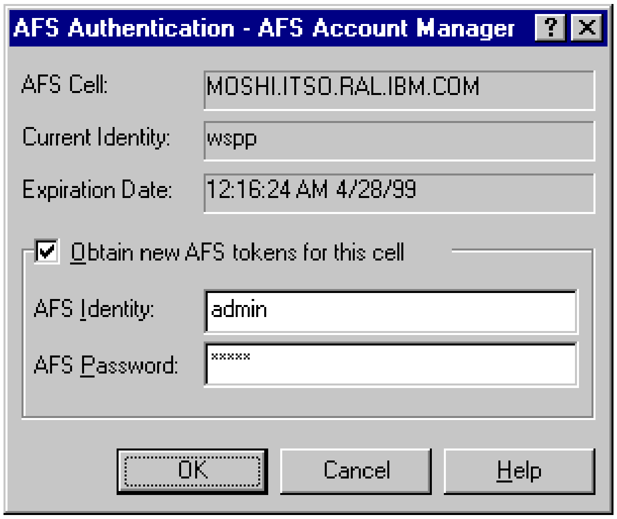

AFS Cell: MOSHI.ITSO.RAL.IBM.COM

Current Identity: wspp

Expiration Date: 12:16:24 AM 4/28/99

☑ Obtain new AFS tokens for this cell

AFS Identity: admin

AFS Password: *****

Figure 305. Reauthenticating as Admin

The wspp AFS user account now owns a group called wspp:web-servers.

8.3.2 AFS Groups in AIX, Solaris and Windows NT Command Prompt

In AFS, the `pts` command suite is used for creating and managing groups via the command line. The commands are the same for AIX, Solaris and the Windows NT Command Prompt.

To create a group use the command:

```
pts creategroup
```

To add members to the group use the command:

```
pts adduser
```

Unlike the Account Manager, an administrator can specify on the command line the owner of the group, as shown in the following example:

```
C:\> klog admin
Password:

C:\> tokens
Tokens held by the Cache Manager:
User admin's tokens for afs@moshi.itso.ral.ibm.com [Expires Jun 07 15:35]
   --End of list --

C:\>pts creategroup -name wspp:wte-admins -owner wspp
group wspp:wte-servers has id -211

C:\>pts adduser -user bobby -group wspp:wte-admins

C:\>pts membership -nameorid wspp:wte-admins
Members of wspp:wte-servers (id: -211) are:
bobby

C:\>
```

Use the `pts` commands `examine` and `listowned` to get group information. Also, `pts` in the Windows NT Command Prompt, has a `listentries` command to get a complete list of groups (and users). An example of this is shown in the following session screen:

```
C:\> pts examine wspp:wte-admins
Name: wspp:wte-servers, id: -211, owner: wspp, creator: admin,
  membership: 1, flags: S-M--, group quota: 0.

C:\> pts listowned wspp
Groups owned by wspp (id: 2) are:
  wspp:wte-admins
  wspp:web-servers

C:\>pts listentries -groups
Name                           ID  Owner Creator
system:administrators        -204   -204    -204
system:backup                -205   -204    -204
system:anyuser               -101   -204    -204
system:authuser              -102   -204    -204
wspp:web-servers             -206      2       2
wspp:wte-admins              -211      2       1

C:\>
```

8.4 Managing Machine Accounts in AFS

A feature of AFS that might be of particular interest to WebSphere Performance Pack administrators is the ability to control access to AFS directories *by machine*. ACLs on AFS directories can have entries for groups comprised of machine IP addresses, thus controlling access to AFS directories on a per-machine basis.

These machine ACL entries must be done in groups. This means that machine IP addresses cannot directly be listed on ACLs, but instead are assigned to a group, and that group in turn is placed on the ACL. In order for a machine to be assigned to a group, however, it must first have an AFS account associated with it.

When creating an AFS account for a machine, the IP address is specified to be the user entry. This entry can also represent multiple machines at once using wildcards, according to the following rules:

- *W.X.Y.Z* represents a single machine, for example 192.12.108.240.
- *W.X.Y.*0 matches all machines whose IP addresses start with the first three numbers. For example, 192.12.108.0 matches both 192.12.108.119 and 192.12.108.120, but does not match 192.12.105.144.
- *W.X.*0.0 matches all machines whose IP addresses start with the first two numbers. For example, the address 192.12.0.0 matches both 192.12.106.23 and 192.12.108.120, but does not match 192.5.30.95.
- *W.*0.0.0 matches all machines whose IP addresses start with the first number in the specified address. For example, the address 192.0.0.0 matches both 192.5.30.95 and 192.12.108.120, but does not match 138.5.63.52.
- Do not define a machine entry with the name 0.0.0.0 to match every machine. The system:anyuser group is equivalent.

8.4.1 Managing Machine Accounts on Windows NT

In the following example, we will create some machine user accounts and assign them to a group. That group could then be placed on an AFS directory ACL, which would in turn control all access from a machine to a particular AFS directory.

Begin by starting the AFS Account Manager and clicking the **Addresses** tab. Note at the top of the AFS Account Manager window that we have tokens for the admin user account, which provides us with the necessary administrative privileges:

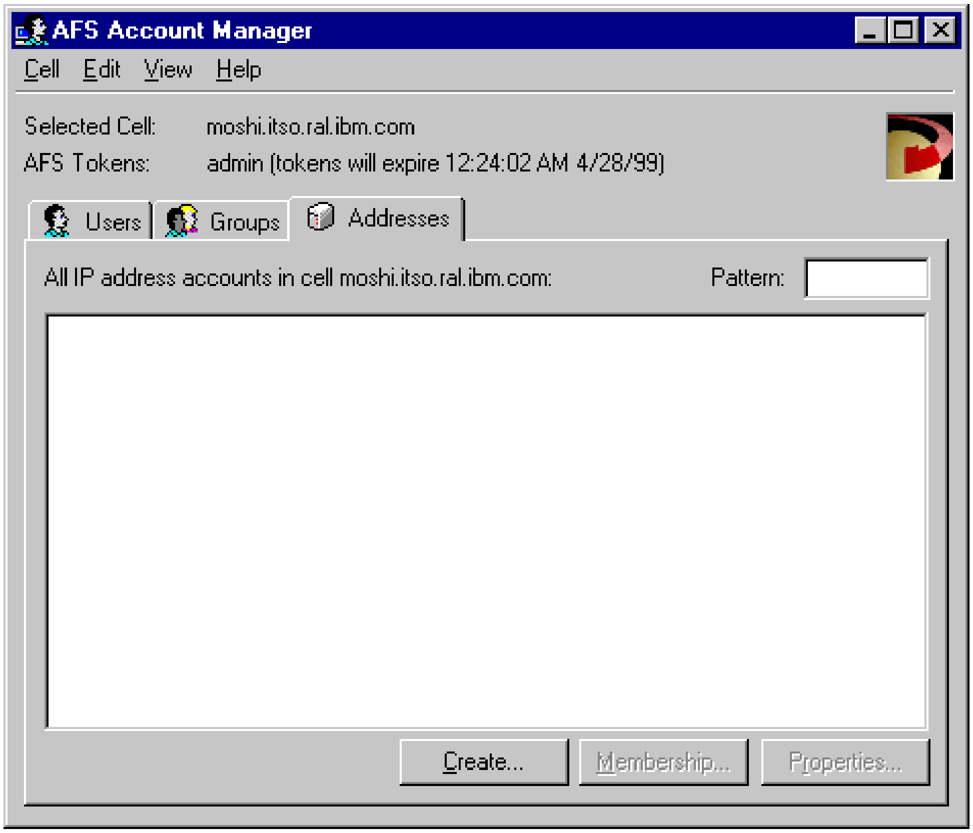

Selected Cell: moshi.itso.ral.ibm.com

AFS Tokens: admin (tokens will expire 12:24:02 AM 4/28/99)

All IP address accounts in cell moshi.itso.ral.ibm.com:

Pattern:

Figure 306. Addresses Tab in the AFS Account Manager

Click the **Create...** button to bring up the Create IP Address Account window:

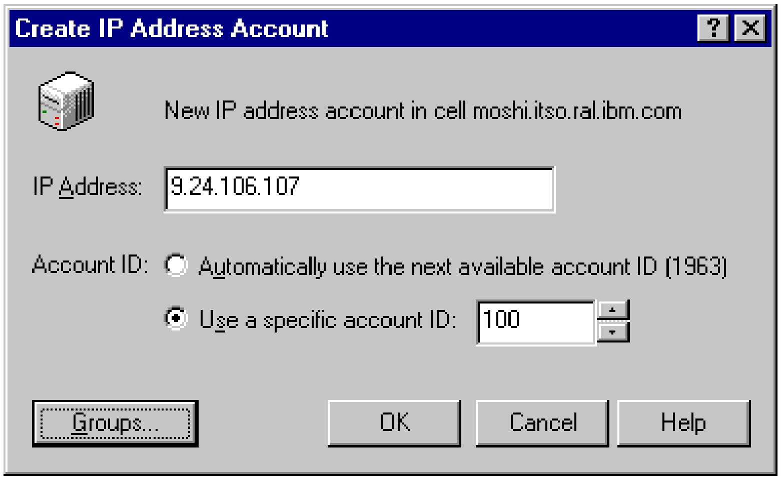
Create IP Address Account

New IP address account in cell moshi.itso.ral.ibm.com

IP Address: 9.24.106.107

Account ID: ○ Automatically use the next available account ID (1963)

● Use a specific account ID: 100

Figure 307. Create IP Address Account Window

In this window, enter the IP address of the machine for which you are creating this address account. Typically, an administrator would let the Account Manager select an account ID number, but in this case we decide we want all IP address accounts to range from 100 to 200 (for example, for administrative purposes); for this reason, we enter `100` as a specific account ID; we could just as easily have specified some other account ID, or let the creation process assign an account ID. Also we will be assigning this IP address account to a group, so we click the **Groups...** button. The Create IP Address Account - Advanced window is brought up:

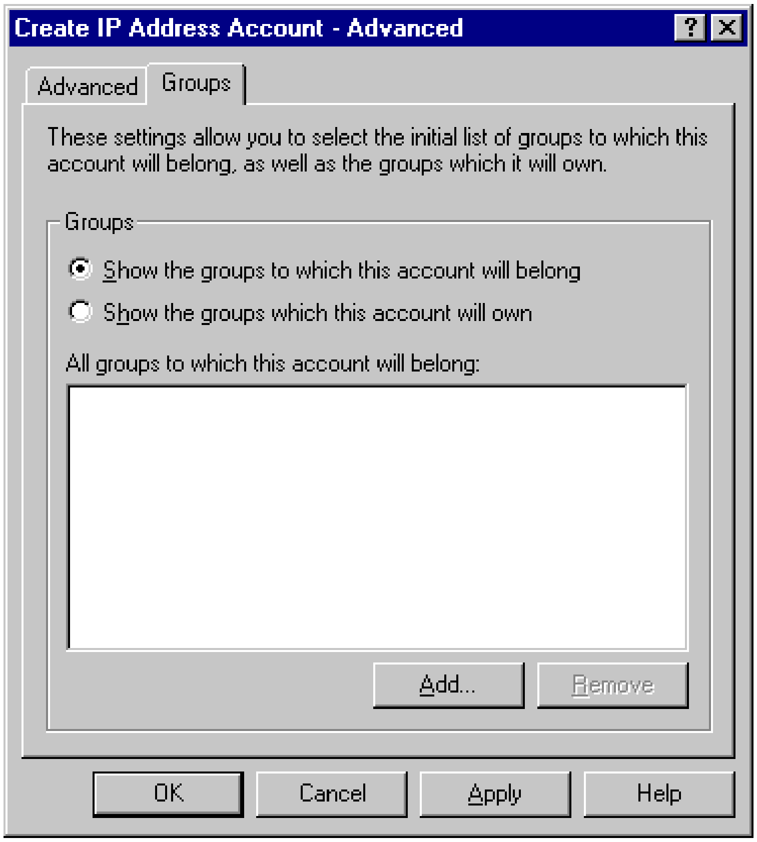

These settings allow you to select the initial list of groups to which this account will belong, as well as the groups which it will own.

Groups

● Show the groups to which this account will belong

○ Show the groups which this account will own

All groups to which this account will belong:

Figure 308. Create IP Address Account - Advanced Window

In the window above, we can assign any group memberships we want this account to have. Click **Add...** and from the list of displayed groups select the group to which this account should be added. This process will be repeated multiple times if you wish to add the account to multiple groups. In our example, we add this IP address account to the wspp:web-servers group, as shown in the following window:

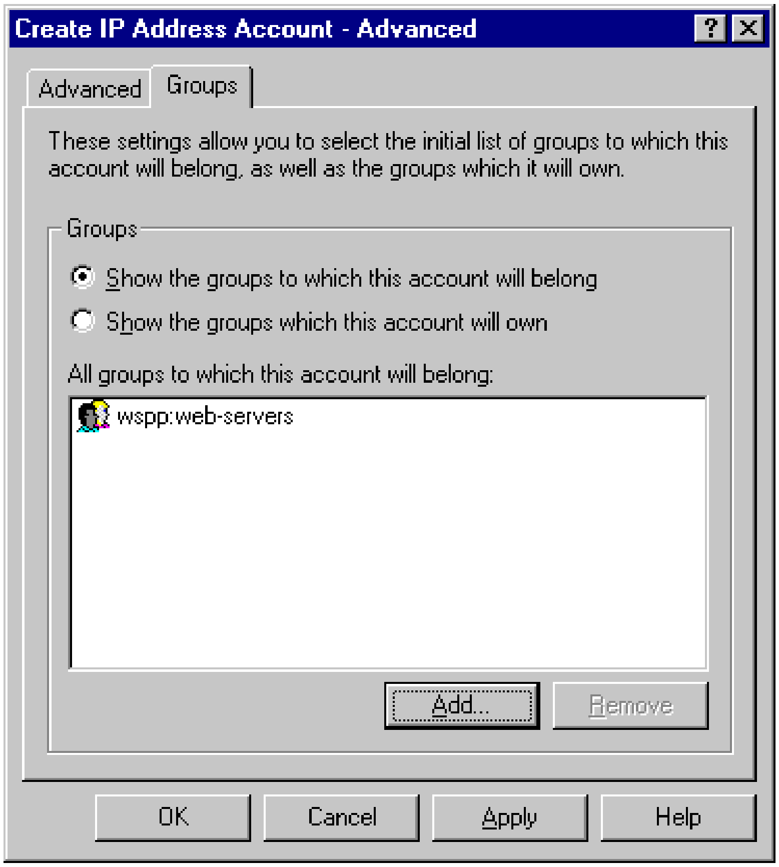
Create IP Address Account - Advanced

These settings allow you to select the initial list of groups to which this account will belong, as well as the groups which it will own.

Groups

● Show the groups to which this account will belong

○ Show the groups which this account will own

All groups to which this account will belong:

wspp:web-servers

Figure 309. IP Address Account Added to Group

When finished adding the IP address account to groups, click the **OK** button, and then click **OK** also in the Create IP Address Account window. As you can see, the IP address account has been created and now it is registered in the AFS Account Manager window:

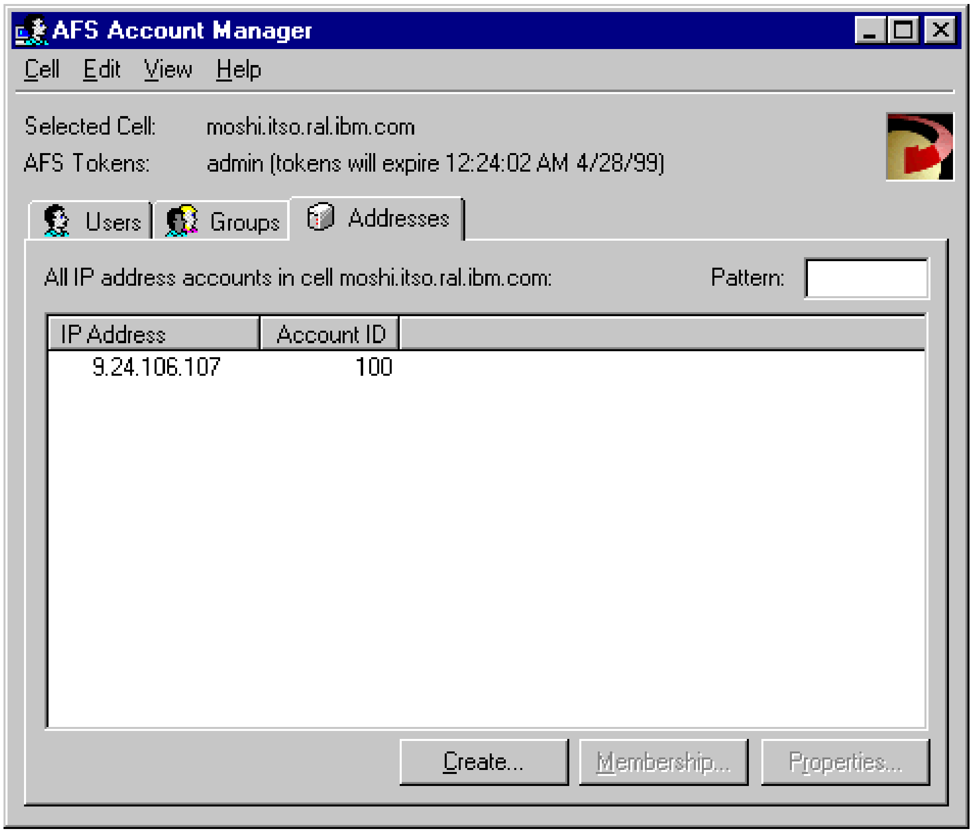

Selected Cell: moshi.itso.ral.ibm.com

AFS Tokens: admin (tokens will expire 12:24:02 AM 4/28/99)

All IP address accounts in cell moshi.itso.ral.ibm.com:

Pattern:

IP Address	Account ID
9.24.106.107	100

Figure 310. IP Address Account Added

Clicking the **Groups** tab shows that the wspp:web-servers group has a member added to it, which is the IP address account we have just created:

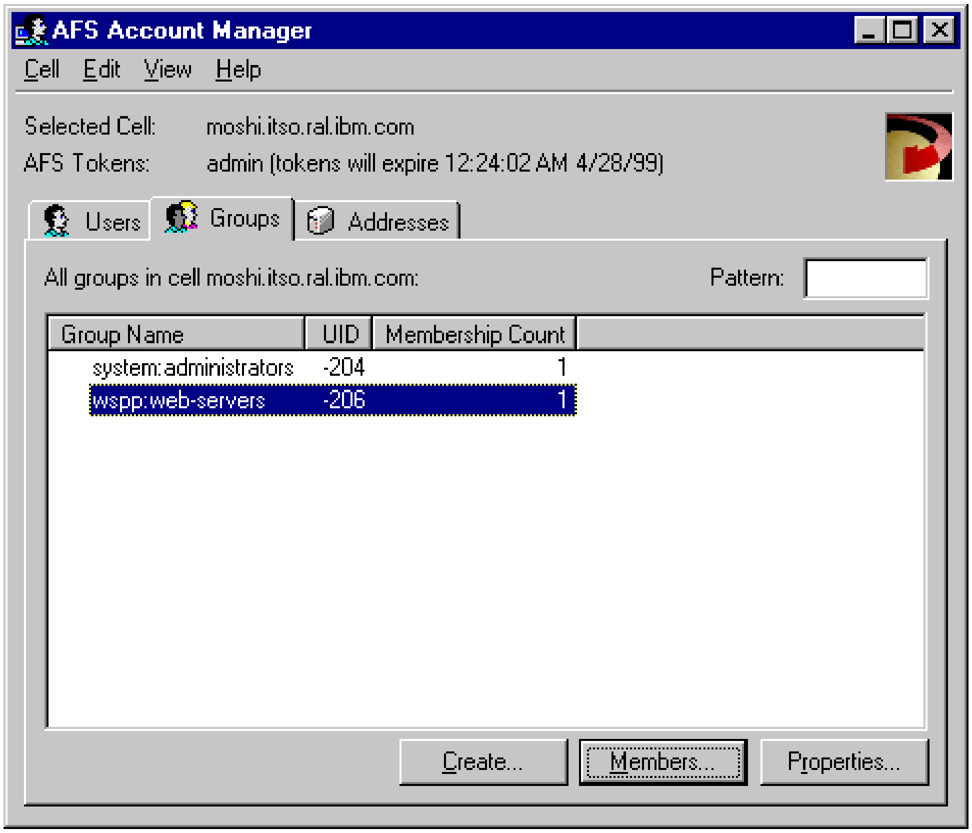

Selected Cell: moshi.itso.ral.ibm.com
AFS Tokens: admin (tokens will expire 12:24:02 AM 4/28/99)

All groups in cell moshi.itso.ral.ibm.com:

Group Name	UID	Membership Count
system:administrators	-204	1
wspp:web-servers	-206	1

Figure 311. The Groups Window with a New Member

Selecting the wspp:web-servers group entry and clicking **Members...** brings up the membership list for the group, which in this case includes the IP address account we created:

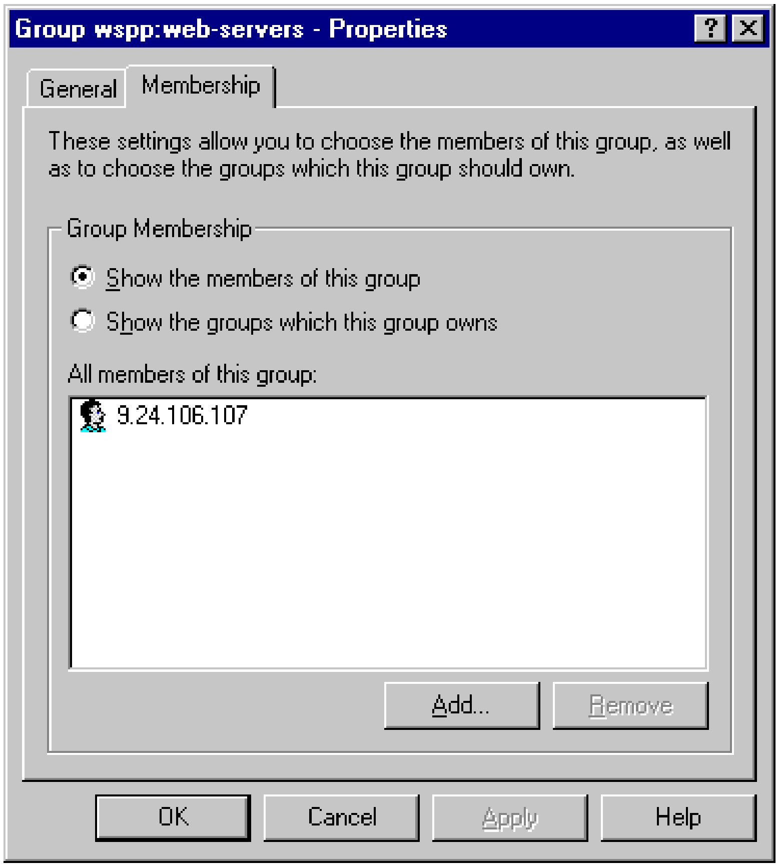

Group wspp:web-servers - Properties

These settings allow you to choose the members of this group, as well as to choose the groups which this group should own.

Group Membership

● Show the members of this group

○ Show the groups which this group owns

All members of this group:

9.24.106.107

Figure 312. Membership Information

Clicking the **General** tab of the same window brings up a list of properties that can be changed for the group. Using this window, you can change ownership and set permissions for group tasks, such as viewing group information and changing memberships:

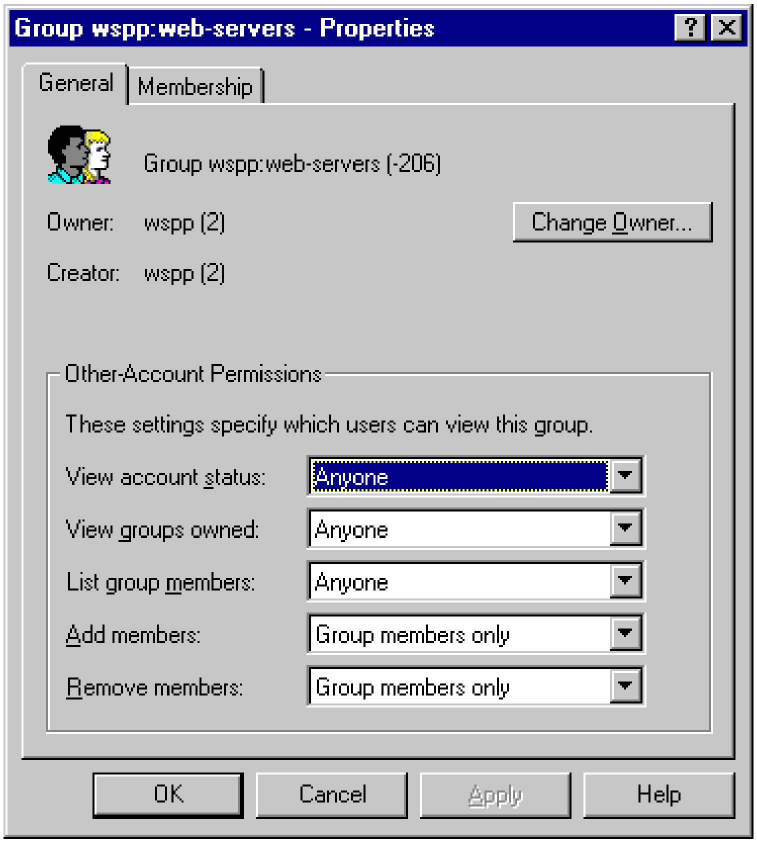

Group wspp:web-servers (-206)

Owner: wspp (2)

Creator: wspp (2)

Other-Account Permissions

These settings specify which users can view this group.

View account status: Anyone

View groups owned: Anyone

List group members: Anyone

Add members: Group members only

Remove members: Group members only

Figure 313. General Properties of Groups

In our example, we go on to create two additional IP address accounts, this time using wildcard values in the fourth octet of the IP address:

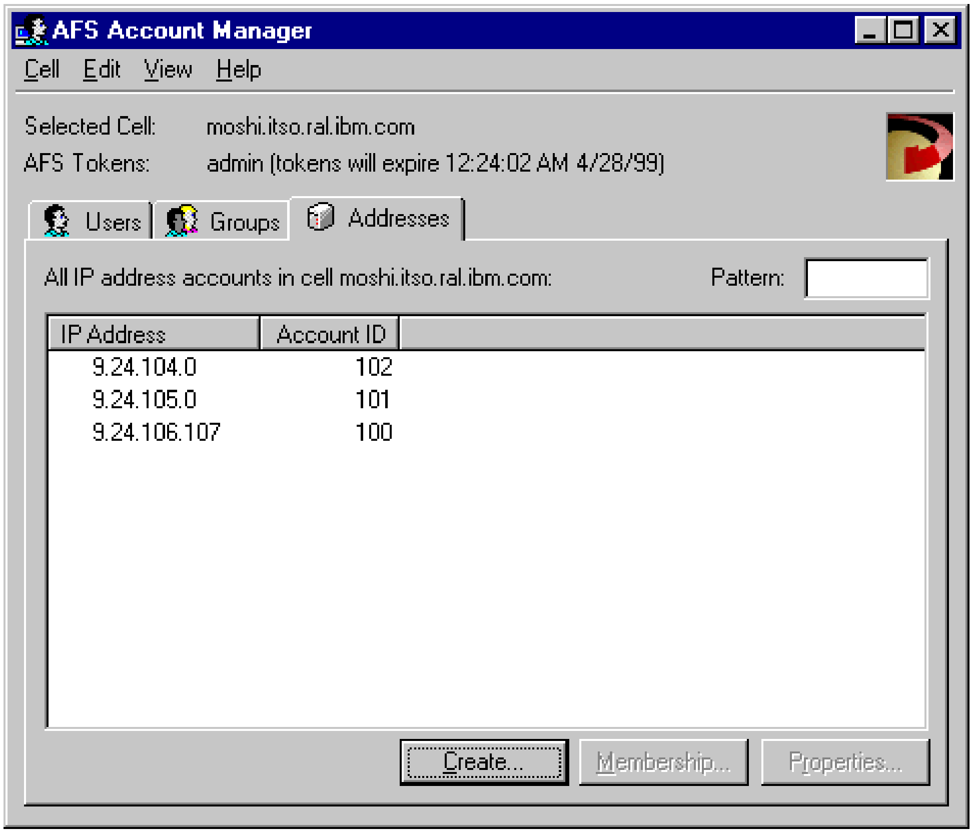

Selected Cell: moshi.itso.ral.ibm.com

AFS Tokens: admin (tokens will expire 12:24:02 AM 4/28/99)

All IP address accounts in cell moshi.itso.ral.ibm.com:

IP Address	Account ID
9.24.104.0	102
9.24.105.0	101
9.24.106.107	100

Figure 314. Additional IP Address Accounts

In essence, we have created user accounts whose members are machines on the 9.24.104 and 9.24.105 subnets respectively by using the number 0 in the fourth octet of the address. To do this we clicked **Create . . .** , and then we typed each IP address as shown in Figure 307 on page 353.

We have also added these new IP address accounts to the wspp:web-servers group as well. Upon opening the Group wspp:web-servers - Properties window, we see the following:

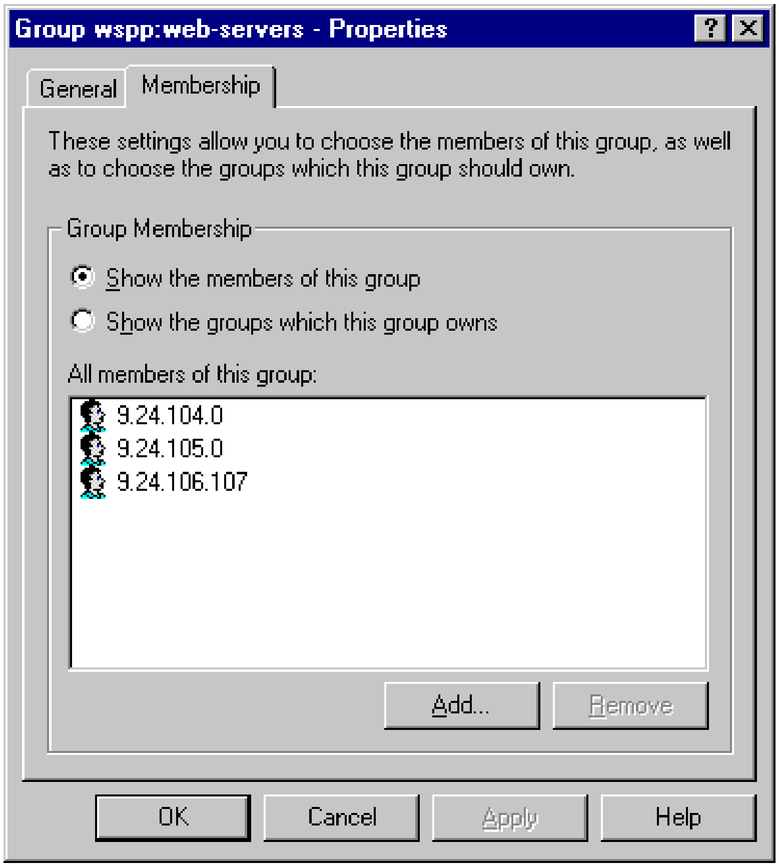

These settings allow you to choose the members of this group, as well as to choose the groups which this group should own.

● Show the members of this group

○ Show the groups which this group owns

All members of this group:

9.24.104.0
9.24.105.0
9.24.106.107

Figure 315. Additional IP Address Accounts in wspp:web-servers Group

We now have a group called wspp:web-servers whose members are three IP address accounts; one account is a specific machine and the other two are entries for entire subnets. This group can now be attached to ACL entries to control access by these machines to AFS directories. For example, you could store Web pages in a directory called /afs/moshi.itso.ral.ibm.com/public/html, and use ACLs to control access to and from this directory for members of the wspp:web-servers group.

8.4.2 Machine Accounts Management through the Command Line

To create IP address accounts in AIX, Solaris, and the Windows NT Command Prompt, simply use the command:

```
pts createuser
```

to create the account, and the command:

```
pts adduser
```

to add the accounts to groups.

An example is shown in the following session screen:

```
C:\>pts createuser -name 9.24.103.0 -id 103
User 9.24.103.0 has id 103

C:\>pts adduser -user 9.24.103.0 -group wspp:web-servers

C:\>
```

8.5 AFS Authentication and Integrated Login

AFS users have two options for authenticating to AFS in order to receive tokens. First, after logging in to the machine's operating system, a user can run a program to authenticate themselves against the Authentication Server. On AIX, Solaris and the Windows NT Command Prompt, a user issues the `klog` command, as shown in the following session screen:

```
C:\> klog tom
Password:

C:\> tokens

Tokens held by the Cache Manager:

User tom's tokens for afs@moshi.itso.ral.ibm.com [Expires Jun 07 19:19]
   --End of list --

C:\>
```

On Windows NT, users also have the option of authenticating through a graphic interface. Under the Start menu, select **Program Files**, then **IBM WebSphere**, then **Performance Pack**, then **Afs**, then **Client**, and finally **Authentication** (alternatively, right-click the **AFS Client** icon in the Windows NT taskbar and select the **Show Tokens...** menu option). This will bring up the AFS Client window:

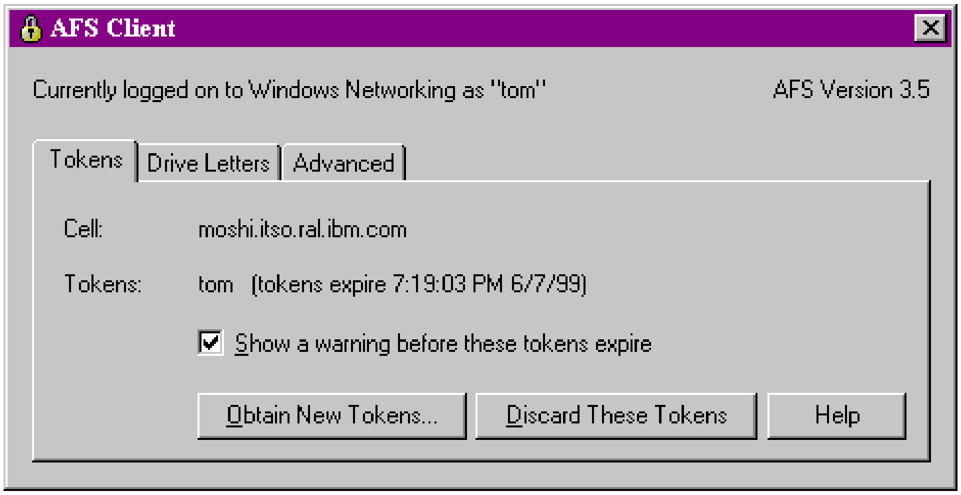

Currently logged on to Windows Networking as "tom"

AFS Version 3.5

Cell: moshi.itso.ral.ibm.com

Tokens: tom (tokens expire 7:19:03 PM 6/7/99)

☑ Show a warning before these tokens expire

Figure 316. The AFS Client Window

If the user currently has any tokens, their information will be displayed here. To authenticate and obtain tokens click the **Obtain New Tokens...** button. This will bring up a dialog box into which the user can enter user name and password for authentication:

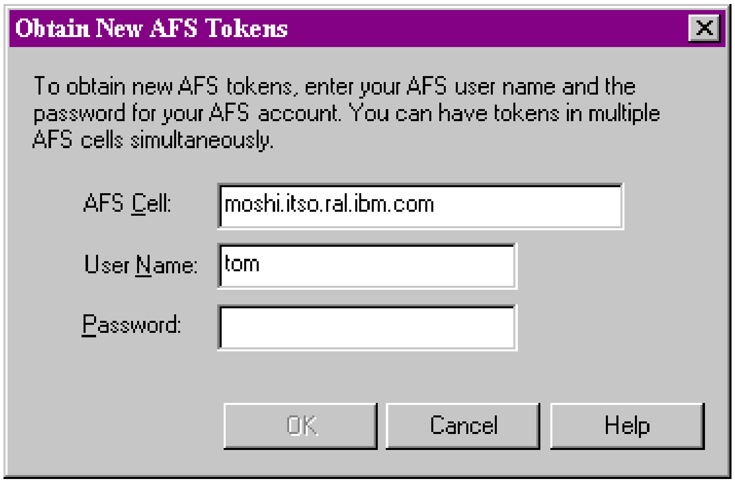

To obtain new AFS tokens, enter your AFS user name and the password for your AFS account. You can have tokens in multiple AFS cells simultaneously.

AFS Cell: moshi.itso.ral.ibm.com

User Name: tom

Password:

Figure 317. Authenticating to AFS

The second method for user authentication requires the AFS administrator to set up integrated login, which allows a user to obtain AFS tokens when first logging in to an AFS client machine. When a user first logs in to an AIX, Solaris or Windows NT workstation, they are simultaneously authenticated against the AFS Authentication Server and are granted a token for the AFS cell, and this works regardless of the client and server operating systems; in other words, they need not be the same platform. Setting this up on clients generally requires some effort on the part of the administrator, though, and for two of the platforms supported by WebSphere Performance Pack we go into more detail below.

8.5.1 AFS Integrated Login on AIX

First, make sure that the file afs_dynamic_auth is installed on every AFS client machine which will have integrated login. After the original installation of the File Sharing client component of IBM WebSphere Performance Pack, this file is by default located in the directory /usr/vice/etc and is also executable.

Next, you have to modify the files /etc/security/user and /etc/security/login.cfg. Before making these modifications, we recommend that you save copies of the original files with different names. For example, we entered the following two commands:

```
cp /etc/security/user /etc/security/user.original
cp /etc/security/login.cfg /etc/security/login.cfg.original
```

In the following figure, we show you the two sections, `default` and `root`, that we modified in the /etc/security/user file. The modifications we made are highlighted in boldface text:

```
default:
        admin = false
        login = true
        su = true
        daemon = true
        rlogin = true
        sugroups = ALL
        admgroups =
        ttys = ALL
        auth1 = SYSTEM
        auth2 = NONE
        tpath = nosak
        umask = 022
        expires = 0
        SYSTEM = "AFS OR (AFS[UNAVAIL] OR compat[SUCCESS])"
        registry = DCE
        logintimes =
        pwdwarntime = 0
        account_locked = false
        loginretries = 0
        histexpire = 0
        histsize = 0
        minage = 0
        maxage = 0
        maxexpired = -1
        minother = 0
        minlen = 0
        mindiff = 0
        maxrepeats = 8
        dictionlist =
        pwdchecks =

root:
        registry = files
        admin = true
        SYSTEM = "compat"
        loginretries = 0
        account_locked = false
```

Note

If you follow the directions to enable the AFS login on AIX found in the guide *AFS Installation Guide* (shipped with the product or downloadable from `http://www.transarc.com/Support`), you will see a small difference from what we have described here. The difference is that, for AIX machines that are only AFS clients, they recommend you add the following line in the `default` section of the /etc/security/user file:

```
SYSTEM = "AFS OR AFS [UNAVAIL] AND compat [SUCCESS]"
```

The problem is this modification does not work.

The correct line to add for AIX 4.2.1 and AIX 4.3 should be instead:

```
SYSTEM = "AFS OR (AFS[UNAVAIL] OR compat[SUCCESS])"
```

Although this appears to be a small difference, it is necessary to properly enable AFS Integrated Login.

Notice that if the machine is both an AFS and a DCE client, the SYSTEM variable in the default section of the /etc/security/user file should be set as follows:

```
SYSTEM = "DCE or DCE [UNAVAIL] OR AFS OR AFS [UNAVAIL] AND compat[SUCCESS]"
```

In the file /etc/security/login.cfg, we added the following at the bottom of the file:

```
DCE:
program = /usr/vice/etc/afs_dynamic_auth

AFS:
program = /usr/vice/etc/afs_dynamic_auth
retry = 3
timeout = 30
retry_delay = 10
```

When the AFS user enters the AIX command `login`, the system does not execute the standard login program but rather the AFS authentication program afs_dynamic_auth instead, which leads the user to log in to AIX and AFS simultaneously.

Note that it is not necessary for user's AIX and AFS passwords to be the same when the AFS Integrated Login on AIX is activated. In fact, it is the afs_dynamic_auth authentication program that takes precedence in terms of authentication. As long as the user provides the correct AFS password, that user is authenticated directly to both AFS and AIX, and the authentication program afs_dynamic_auth will not even consult the AIX password file /etc/security/passwd.

8.5.2 Interactions between the AFS Login and the AIX Login

After completing the procedure we described to enable AFS Integrated Login, an AFS user logging in is automatically authenticated to AFS. As long as the user provides the correct AFS password, the AFS login never consults the local password file. In this situation, root is the only non-AFS user that can log into the AIX machine without being authenticated to the AFS file system. All other users must pass through the AFS login and so only authorized AFS users can log into the AFS file system and the AIX machine. We want to show you now how you can configure your AIX system to accept other AIX users, other than root, that are not AFS users.

On our AFS client AIX machine, we created through `smitty` a new AIX user named madtest. To do this, we simply entered, as root, the following command:

```
smitty users
```

Then we selected **Add a User**. In the window that was displayed, we entered `madtest` as the name of the new user, and left all the other fields at their default values. As soon as you press the Enter key, madtest is a new AIX user added to your system and you will find a new entry for madtest in the file /etc/passwd. However, this user still falls under the default user category, for which the AFS login takes precedence on the AIX login. You cannot even set a password for madtest. We tried with both `smitty` and the `passwd` command, but no entry was added to the /etc/security/passwd file. For this reason, you cannot log in as madtest yet, since no password has been defined for this user.

After this step, we opened the file /etc/security/user with the vi editor, by entering:

```
vi /etc/security/user
```

Then we modified the entry for madtest. We added the two lines that are highlighted in boldfaced text in the following screen:

```
madtest:
         admin = false
         SYSTEM = "compat"
         registry = files
```

This way the new user madtest does not fall under the default user category, for which the AFS login takes precedence. You can see that we have modified the entry for madtest so that now it has become similar to the entry for root.

We issued again, as root, the command:

```
passwd madtest
```

This time we were able to set a password for madtest and to see that password scrambled in the /etc/security/passwd file. After this step, we successfully logged in as madtest. Notice that when we logged in as madtest, it was the AIX login that authenticated us, rather than the AFS login. An AFS entry named madtest does not even exist. To be sure that the AFS login had not come into the scene, we entered the command:

```
/usr/afs/bin/tokens
```

We saw the following output, which indicates the user has no tokens:

```
Tokens held by the Cache Manager:

   --End of list--
```

We want to show you now how you can use the local password file /etc/security/passwd to still control access even with AFS users:

1. Entering an asterisk * in the password field prevents the user from either authenticating with AFS or logging into the machine's AIX file system. This is very useful in case of an emergency, if you want to prevent certain users from logging into the machine.
2. Entering in the password field a string of any length other than the standard thirteen characters (typically, you will enter a simple `X`) means that if the user does not provide the correct AFS password, they will also be not allowed to log into the local machine's AIX file system. This is appropriate if you want to make sure that only people with AFS accounts in your cell can log into your machines.
3. Putting an actual scrambled password in the password field by using the `passwd` command means that if the user does not provide the correct AFS password, he or she can still log in to the local AIX file system by providing the correct AIX password. In this case the AFS and AIX passwords should differ, or an incorrect AFS password would prevent the user from authenticating into the AIX system as well.

We successfully experimented with all three cases above. We then edited again the file /etc/security/user, to remove the modification we made to the madtest entry. In this way, madtest fell again into the default user category, for which the AFS login takes precedence. Then we defined a new AFS account, still named madtest. In particular, the AFS and AIX user IDs for madtest were the same, but the passwords were different. We tried to log in as madtest by entering the correct AFS password and the AFS login authenticated us on both AFS and AIX, without checking the entry for madtest in the etc/security/passwd file. We also saw that a token had been produced on the AFS client machine rs600012 for madtest.

We then discarded that token, and logged out. When we logged in again as madtest, this time we entered the password for the AIX authentication. This time the AFS login program did not authenticate us, since it found an incorrect password, but since we had entered the correct AIX password at least we were able to access the AIX machine as a normal AIX user, as shown in the following screen:

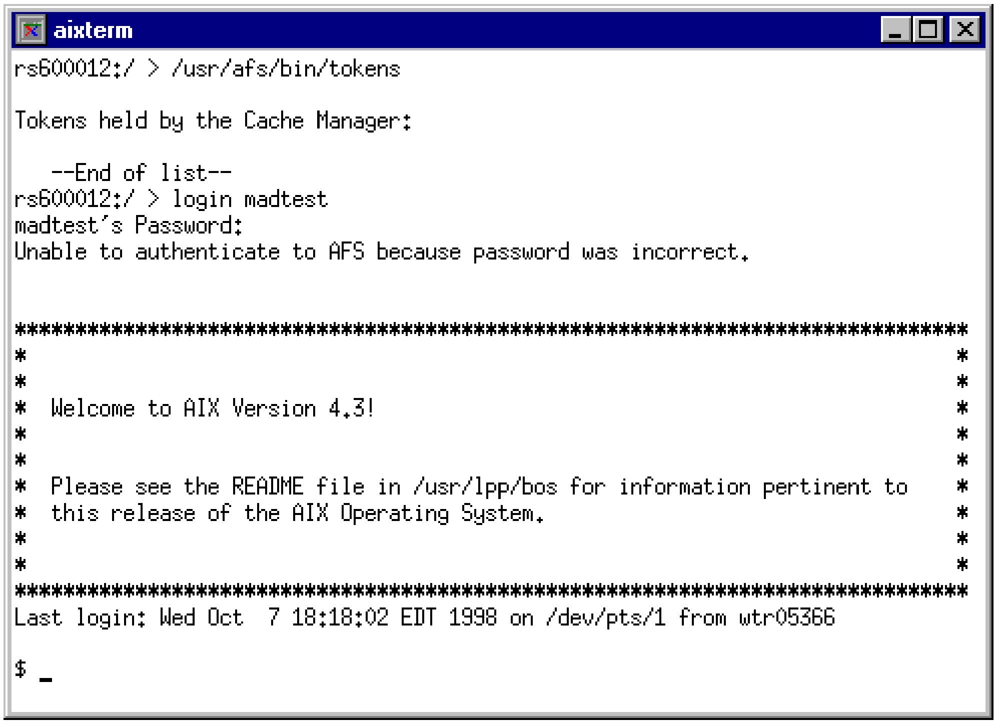

```
rs600012:/ > /usr/afs/bin/tokens

Tokens held by the Cache Manager:

   --End of list--
rs600012:/ > login madtest
madtest's Password:
Unable to authenticate to AFS because password was incorrect.


*******************************************************************************
*                                                                             *
*                                                                             *
*  Welcome to AIX Version 4.3!                                                *
*                                                                             *
*                                                                             *
*  Please see the README file in /usr/lpp/bos for information pertinent to    *
*  this release of the AIX Operating System.                                  *
*                                                                             *
*                                                                             *
*******************************************************************************
Last login: Wed Oct  7 18:18:02 EDT 1998 on /dev/pts/1 from wtr05366

$ _
```

Figure 318. Authenticating on AIX but Not on AFS

As you can see from the above figure, before logging in, no tokens were held by the Cache Manager. On entering the correct AIX password, we were not able to authenticate to AFS, since that password was considered incorrect for AFS. The following error message appeared:

```
Unable to authenticate to AFS because password was incorrect
```

However, after that, the password was compared with the scrambled password stored in the /etc/security/password file, and we were able to authenticate to the AIX system.

8.5.3 AFS Integrated Login on Solaris

AFS Integrated Login for Solaris involves installation and configuration of Pluggable Authentication Modules (PAM), which is beyond the scope of this redbook. See the *AFS Installation Guide* for instructions on how to set up AFS integrated login for Solaris. The software copy of this guide comes with the product and can also be downloaded from `http://www.transarc.com/Support`.

8.5.4 AFS Integrated Login on Windows NT

To set up AFS integrated login on Windows NT the administrator must first make sure that the Windows NT user name and password match the AFS user name and password. The Windows NT user account can either be defined on the local client or in a Windows NT domain, but in either case the name and password must match their equivalents in AFS.

Once this is done open the AFS Client Configuration Control Panel. On the very first screen (under the **General** tab) you will see the option **Obtain AFS tokens**

when logging in to Windows. Simply check this box, and integrated login will take effect upon reboot.

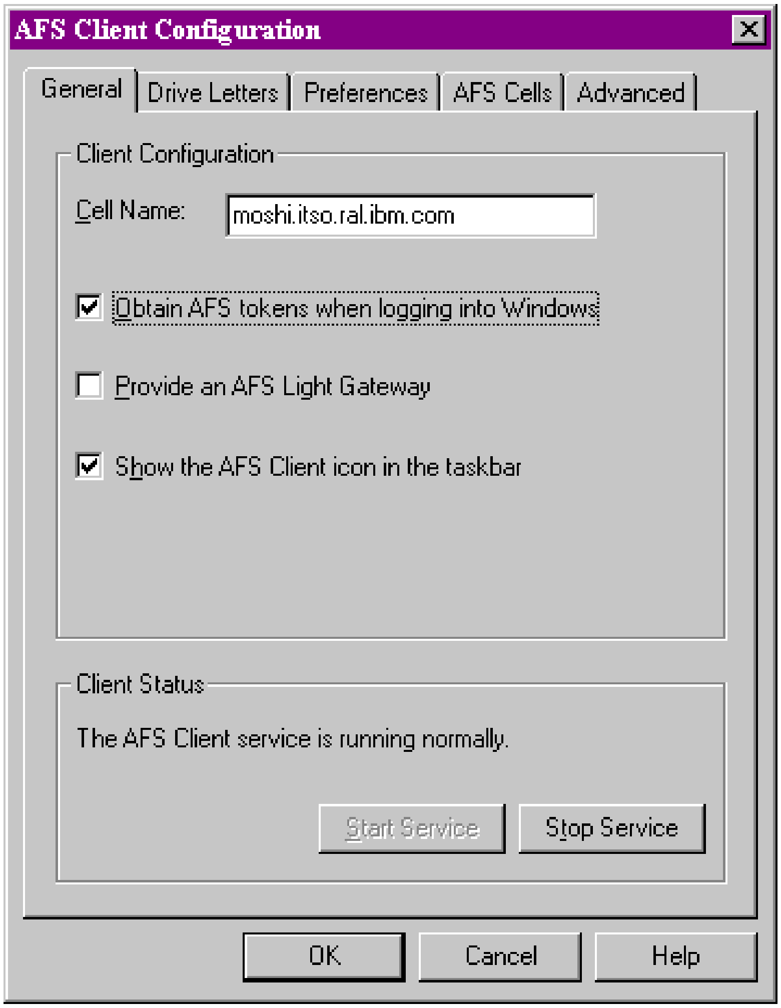

Cell Name: moshi.itso.ral.ibm.com

☑ Obtain AFS tokens when logging into Windows

□ Provide an AFS Light Gateway

☑ Show the AFS Client icon in the taskbar

The AFS Client service is running normally.

Figure 319. Setting Up Integrated Login on Windows NT

Part 2. WebSphere Performance Pack Component Integration

Chapter 9. Content Based Routing

Content Based Routing (CBR) is an entirely new component within this version of SecureWay Network Dispatcher (ND). The purpose of CBR is to route client requests to specific servers based on the content of the URL request.

In this chapter we will show some scenarios using CBR. See *IBM WebSphere Performance Pack - Load Balancing with IBM SecureWay Network Dispatcher,* SG24-5858 for more information on how CBR works.

9.1 Installation of the CBR Function

The CBR component of ND Version 2.1 is supported on three operating systems: IBM AIX 4.2.1 or later, Microsoft Windows NT 4.0 and Sun Solaris 2.6 or later. Refer to the appropriate platform section of *IBM WebSphere Performance Pack - Load Balancing with IBM SecureWay Network Dispatcher,* SG24-5858 for details on how to use the Java InstallShield on your respective platform to install ND.

When you reach the point where you choose which ND component to install, select these three to install CBR on your machine:

- **Content Based Routing Runtime**
- **Content Based Routing Administration**
- **Content Based Routing License**

The following figure shows these three items as being selected:

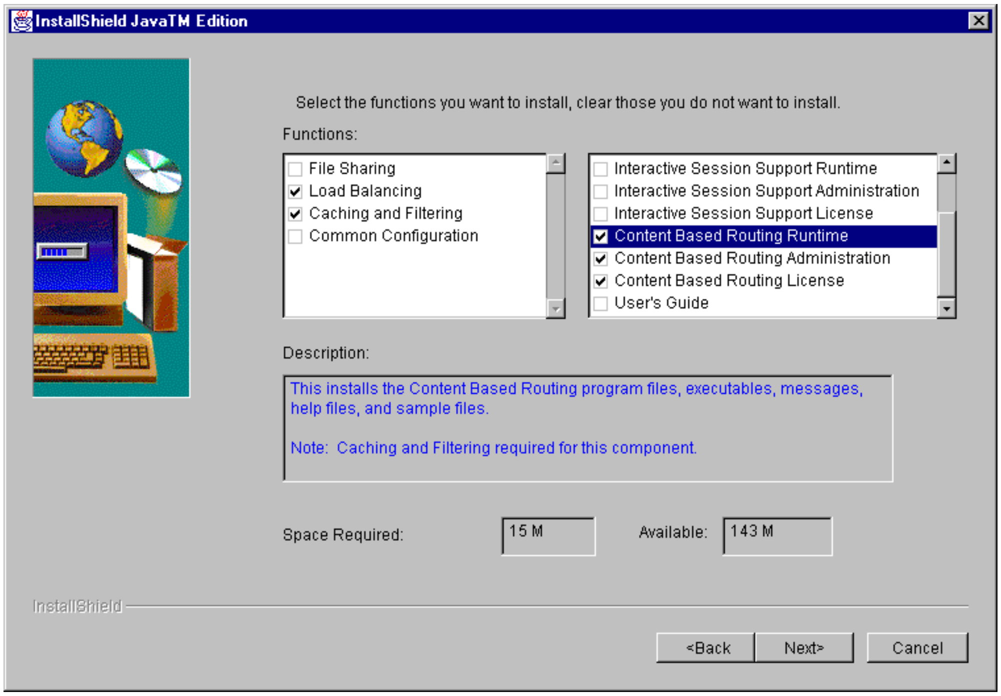

InstallShield JavaTM Edition

Select the functions you want to install, clear those you do not want to install.

Functions:

☐ File Sharing
☑ Load Balancing
☑ Caching and Filtering
☐ Common Configuration

☐ Interactive Session Support Runtime
☐ Interactive Session Support Administration
☐ Interactive Session Support License
☑ Content Based Routing Runtime
☑ Content Based Routing Administration
☑ Content Based Routing License
☐ User's Guide

Description:

This installs the Content Based Routing program files, executables, messages, help files, and sample files.

Note: Caching and Filtering required for this component.

Space Required: 15 M

Available: 143 M

Figure 320. Java InstallShield WebSphere Performance Pack Component Selection

When you select the **Content Based Routing Runtime** ND component, the **Caching and Filtering** WebSphere Performance Pack Version 2 component is automatically selected to be installed, as CBR cannot be installed without it.

The next screen asks if you would like to use the WTE proxy server to cache documents. If you click on the **Caching** checkbox, you will be asked to fill in three blanks:

- Cache Root Directory
- Cache Access Log
- Cache Size

The following figure shows this on the Windows NT platform:

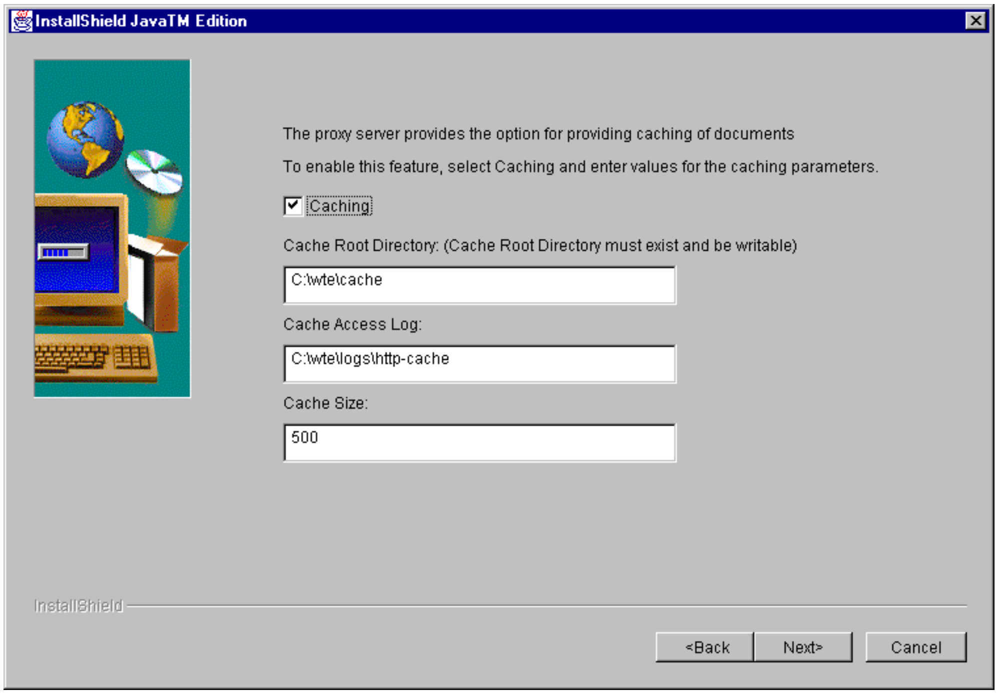

InstallShield JavaTM Edition

The proxy server provides the option for providing caching of documents

To enable this feature, select Caching and enter values for the caching parameters.

☑ Caching

Cache Root Directory: (Cache Root Directory must exist and be writable)

C:\wte\cache

Cache Access Log:

C:\wte\logs\http-cache

Cache Size:

500

Figure 321. CBR Install Process Asks for Caching Details

In our case, we wanted to use the caching function of the WTE proxy server, so we checked the **Caching** item in the dialog window, and entered the values shown in Figure 321 on page 374. 9.3.1, “WTE Configuration Overview” on page 376 explains how we selected these values. WTE and CBR can be installed with caching disabled for simplicity and then caching can be enabled and configured at a later time.

After we clicked on the **Next** button, CBR and WTE were installed.

9.1.1 Installation Locations

If CBR is installed from the WebSphere Performance Pack Version 2 InstallShield as demonstrated in 9.1, “Installation of the CBR Function” on page 373, the WTE component is also automatically installed. Alternatively, if CBR is installed from the SecureWay Network Dispatcher InstallShield, then WTE is not automatically installed and must be installed manually.

Note that in each case, the default base install directories will be different on Windows NT as shown in the following table:

Table 5. CBR and WTE Default Installation Base Directories

Platform	WebSphere Performance Pack components	
	WTE	CBR
AIX	/usr/lpp/internet/server_root	/usr/lpp/nd/cbr
Solaris	/opt/internet/server_root	/opt/nd/cbr
Windows NT WSPP InstallShield	C:\WSPP\WWW\	C:\WSPP\IBM\nd\cbr
Windows NT WTE or CBR InstallShield	C:\WWW	C:\Program Files\IBM\nd\cbr

9.2 Configuration of the CBR Function

In order to perform the configuration you must be the root user on AIX and Solaris or, if the ND server is installed on Windows NT, a member of the Administrators group.

Configuration of CBR cannot take place until WTE has been configured to start the CBR subprocess. We performed CBR configuration in three steps:

1. Because the CBR server is a subprocess of WTE, the first step we performed was to confirm that the WTE server was functional. Optionally, during this step, the WTE server can be configured with values different from the defaults set at installation time (for example, caching can be enabled).
2. Once this is done, four lines must be added to the WTE configuration file to enable CBR functionality. The WTE configuration file is:
   - C:\WINNT\ibmproxy.conf on Windows NT
   - /etc/ibmproxy.conf on AIX and Solaris

   Notice that the WTE server must be stopped and restarted for the changes to take effect.
3. Once the modifications have been made to the WTE configuration file, and the WTE server has been restarted using the modified configuration file, only then can CBR configuration be done. CBR configuration is very similar in style to Dispatcher configuration, in that one or more clusters are usually configured with their associated ports and servers. Similar to rule-based load balancing, rules are then defined, but in the case of CBR the type of the rule is Content. There are three methods of configuring CBR available:
   - `cbrcontrol` commands can be issued from the command line. As well, you can make use of the `cbrcontrol` *persistent session.* The persistent session is a limited shell like utility that can be used to enter `cbrcontrol` commands. Use the `cbrcontrol` command without any parameters to start the persistent session. You will receive the

     ```
     cbrcontrol >>
     ```

     prompt to which you can respond by entering `cbrcontrol` commands without the `cbrcontrol` keyword.

     See *SecureWay Network Dispatcher User's Guide Version 2.1 for Solaris, Windows NT and AIX*, GC31-8496 for `cbrcontrol` command usage information.

- The cbrcontrol configuration commands can be placed in a script or command file that can be executed as a batch job. When the file is run the commands are executed in sequence. A CBR configuration file can be loaded into the execution environment with the command:

  ```
  cbrcontrol file load configfilename
  ```

  See Figure 328 on page 388 for an example of a file containing cbrcontrol commands.

- The ND graphical user interface (GUI) can be used to perform all of the configuration steps that can be performed with cbrcontrol commands.

In addition, all three methods of configuration can be performed by a user on a machine other than the one that is being configured. This new feature is referred to as Remote Authenticated Administration and is explained further in *IBM WebSphere Performance Pack - Load Balancing with IBM SecureWay Network Dispatcher,* SG24-5858.

The three configuration steps above are described in the next section.

9.3 CBR Scenario

In this scenario, we demonstrate how to use CBR by configuring a basic environment. In 9.3.1, “WTE Configuration Overview” on page 376 we give an overview of how to configure WTE. In 9.3.2, “WTE Configuration File CBR Modifications” on page 377 we give details on how to change the WTE configuration file to enable the CBR subprocess. Then in 9.3.3, “CBR Configuration” on page 379 we use the ND GUI to perform CBR configuration of our cluster environment and add the Content type rules. Following this, we show how the rules influence which server is selected to service our request.

9.3.1 WTE Configuration Overview

The WTE component of WebSphere Performance Pack Version 2 contains functionality to work with the CBR component of ND Version 2.1. The steps required to start the WTE server are different on AIX and Solaris than on Windows NT. However, on all these platforms, WTE can be started with all of the default values without making any modifications to the configuration file.

We chose to modify the installed default configuration of WTE to make two changes. The first change was to modify the log used by the WTE proxy server (the proxy server is the basic function of the WTE server that is started if no user configuration is done after WTE installation) and the second change was to enable caching. *IBM WebSphere Performance Pack - Load Balancing with IBM SecureWay Network Dispatcher,* SG24-5858, contains information on WTE including details on how to configure its different functions.

Following a summary list of tasks that we performed to configure WTE on our server:

1. We created a WTE admin user ID and password for use in the following steps.
2. We created a separate file system or directory (/wte on UNIX or C:\wte in the case of Windows NT) to contain the WTE cache root directory and log files. On AIX, we found it necessary to change its ownership with the following command:

```
chown nobody.nobody /wte
```

3. We used the WTE Configuration and Administration Forms, accessible by directing a Web browser to the machine where WTE is installed, to change the location of the proxy log file. The only item we changed was the proxy server log file.
4. In the Configuration and Administration Forms, clicking **Submit** caused the ibmproxy.conf file to be updated with our new log file, but **Submit** or even **Restart Server** (the bar icon at the top of the window) did not result in these changes being used by the server. So we had to stop and restart the WTE server.
5. At this point, we were able to start the Netscape Navigator browser on our client machine. We selected **Preferences** from the Edit menu. Then, we clicked **Advanced** and **Proxies** and in the Proxies panel, we selected the **Manual Proxy Configuration** radio button. In the HTTP field of the Manual Proxy Configuration window, we entered the IP address of the machine where we had installed and configured the WTE server, with the corresponding port field set to `80`.
6. We made a request for a Web page that was available in another subnet that we had connectivity to. At this point we were presented with the Web page and verified from the Server Activity Monitor, accessible from the WTE Configuration and Administration Forms, that the request had been proxied by our server.
7. The process of configuring the caching component is very similar to this and was done by selecting the **Cache Configuration** item in the WTE Configuration and Administration Forms window.

9.3.2 WTE Configuration File CBR Modifications

When WTE is installed, the WTE configuration file, ibmproxy.conf, contains many lines of preconfigured WTE directives. It is necessary to add four lines to this configuration file to enable WTE to start the CBR server subprocess. In the configuration file, we searched for the word `ServerInit` and discovered that `ServerInit` was only contained in the default ibmproxy.conf file as a comment. At the end of the `ServerInit` comment section, we added the four lines.

It is important to note that the first of the four lines is very long and therefore wraps over the right side of the line to the next line. The parameters on the first line are as follows:

- Client keys directory
- Server key directory
- Install path
- Class path
- RMI port
- Log directory
- Save directory

The four lines are shown in the following screens. In each of the screens, it appears that the `CBR_CLIENT_KEYS_DIRECTORY` item is on a new line, but this is due to a formatting problem in this publication. In fact, the `CBR_CLIENT_KEYS_DIRECTORY` item is on the same line as the `ServerInit` directive. On the `ServerInit` line there is one space after the `ServerInit` keyword and one space before

CBR_CLIENT_KEYS_DIRECTORY. Do not include the new line character or any spaces within the parameter string on this line.

Some of the fields have variable-like names included on the line, and others do not. The fields are comma delimited. Following are the four lines that are necessary to configure CBR. The first figure contains the lines we used on our Windows NT system in C:\WINNT\ibmproxy.conf:

```
ServerInit C:\WSPP\IBM\nd\cbr\lib\libndcbr.dll:ndServerInit
CBR_CLIENT_KEYS_DIRECTORY=C:\WSPP\IBM\nd\admin\keys\cbr,CBR_SERVER_KEYS_DIRECTORY=C:\WSPP\IBM\nd\cbr\key,
END_INSTALL_PATH=C:\WSPP\IBM\nd,C:\WSPP\IBM\nd\cbr\lib;C:\WSPP\IBM\nd\cbr\lib\ibmcbr.jar;C:\WSPP\IBM\nd\a
dmin\lib\ChartRuntime.jar,11099,C:\WSPP\IBM\nd\cbr\logs\,C:\WSPP\IBM\nd\cbr\configurations\

PreExit C:\WSPP\IBM\nd\cbr\lib\libndcbr.dll:ndPreExit

PostExit C:\WSPP\IBM\nd\cbr\lib\libndcbr.dll:ndPostExit

ServerTerm C:\WSPP\IBM\nd\cbr\lib\libndcbr.dll:ndServerTerm
```

Note that if you installed CBR with the ND InstallShield rather than the WebSphere Performance Pack InstallShield, the path contained in these lines will have to be changed from \WSPP\IBM to \Progra~1\IBM. You are required to type the DOS format of the Program Files directory because there can be no spaces within the path elements on these lines.

The changes to the WTE configuration file /etc/ibmproxy.conf on AIX are shown in the following screen:

```
ServerInit /usr/lpp/nd/cbr/lib/libndcbr.so:ndServerInit
CBR_CLIENT_KEYS_DIRECTORY=/usr/lpp/nd/admin/keys/cbr,CBR_SERVER_KEYS_DIRECTORY=/usr/lpp/nd/cbr/key,END_I
NSTALL_PATH=/usr/lpp/nd,/usr/lpp/nd/cbr/lib:/usr/lpp/nd/cbr/lib/ibmcbr.jar:/usr/lpp/nd/admin/lib/ChartRun
time.jar,11099,/usr/lpp/nd/cbr/logs/,/usr/lpp/nd/cbr/configurations/

PreExit /usr/lpp/nd/cbr/lib/libndcbr.so:ndPreExit

Postexit /usr/lpp/nd/cbr/lib/libndcbr.so:ndPostExit

ServerTerm /usr/lpp/nd/cbr/lib/libndcbr.so:ndServerTerm
```

The changes to the WTE configuration file /etc/ibmproxy.conf on Solaris are shown in the following screen:

```
ServerInit /opt/nd/cbr/lib/libndcbr.so:ndServerInit
CBR_CLIENT_KEYS_DIRECTORY=/opt/nd/admin/keys/cbr,CBR_SERVER_KEYS_DIRECTORY=/opt/nd/cbr/key,END_INSTALL_PA
TH=/opt/nd,/opt/nd/cbr/lib:/opt/nd/cbr/lib/ibmcbr.jar:/opt/nd/admin/lib/ChartRuntime.jar,11099,/opt/nd/cb
r/logs/,/opt/nd/cbr/configurations/

PreExit /opt/nd/cbr/lib/libndcbr.so:ndPreExit

Postexit /opt/nd/cbr/lib/libndcbr.so:ndPostExit

ServerTerm /opt/nd/cbr/lib/libndcbr.so:ndServerTerm
```

After adding these lines to the WTE configuration file and saving the file, there is another configuration issue you should take care of:

- On Windows NT, add the following to the Path system environment variable:

  ```
  C:\jdk1.1.7\bin;C:\Program Files\nd\cbr\lib
  ```

  Notice here that we are assuming that Version 1.1.7 of the Java Development Kit (JDK) was installed in the default installation directory C:\jdk1.1.7.
- On AIX, add the following to the LIBPATH system environment variable:

  ```
  /usr/jdk_base/lib:/usr/jdk_base/lib/aix/native_threads:/usr/lpp/nd/cbr/lib
  ```
- On Solaris, add the following to the LD_LIBRARY_PATH system environment variable:

  ```
  /opt/jre1.1.7/lib/sparc/native_threads:/opt/nd/cbr/lib
  ```

Now, you can restart WTE.

Starting the WTE Server on AIX Systems

On AIX, it is advisable to use the `httpd` command to start the WTE server in a CBR environment. The System Resource Control (SRC) command:

```
startsrc -s httpd
```

does not start WTE correctly in the CBR case because WTE cannot find the library files necessary to configure and start CBR. This can be verified by looking at the WTE error log file.

9.3.3 CBR Configuration

Once the ibmproxy.conf file has been modified with the four lines to enable CBR, and the WTE proxy server has been restarted, then CBR configuration can begin.

In this scenario we defined two clusters, A and B, each with two servers. The two servers in cluster A were AIX systems and the two servers in cluster B were Solaris systems. Each cluster already had one rule defined by us, which load balanced requests for a particular page to the two servers. In this section we describe the steps necessary to define two additional CBR rules.

9.3.3.1 Network Environment

The following figure is a graphical representation of our scenario environment:

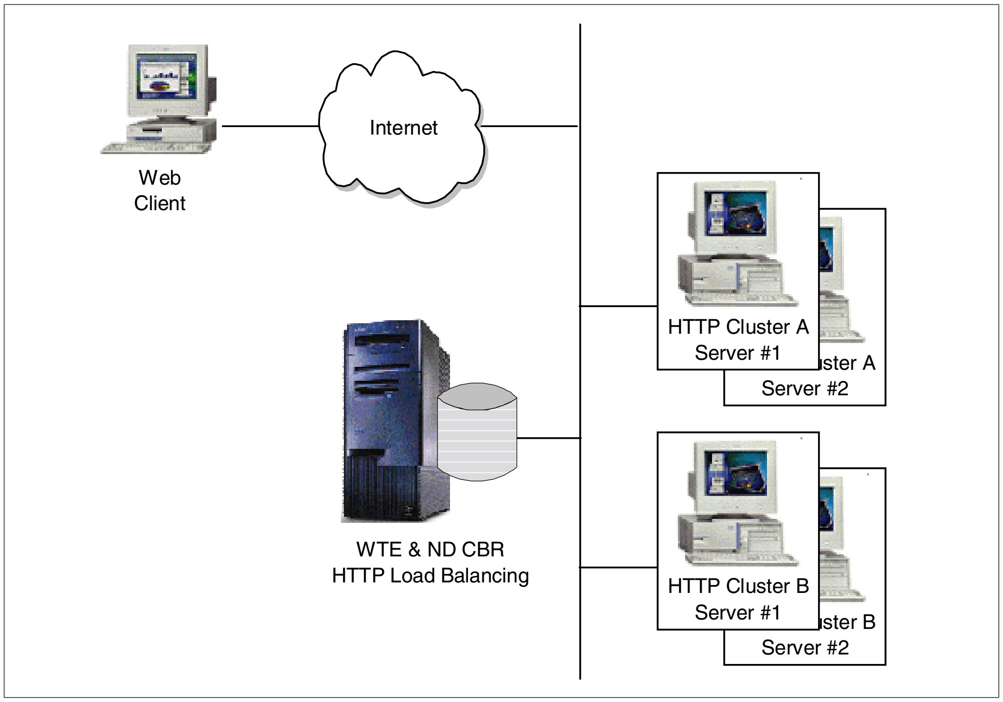

Figure 322. CBR Basic Scenario Environment

A summary of the software and network configuration of the environment where we performed our test is reported in the following table:

Table 6. CBR Basic Scenario - Hardware, Software and Network Configuration

Service		IP Address	Operating System
WTE & CBR server		9.67.133.67	Windows NT Server 4.0
Cluster A IP address 9.67.133.18	Web server 1	9.67.133.75	Solaris 2.6
	Web server 2	9.67.133.77	Solaris 2.6
Cluster B IP address 9.67.134.221	Web server 1	9.67.131.151	AIX 4.3.1
	Web server 2	9.67.131.153	AIX 4.3.1

9.3.3.2 Cluster, Port, Server and Rule Configuration

Once the CBR-enabled proxy server was running, we used the ND GUI to configure our CBR cluster. We used the `ndadmin` command to start the GUI and after right-clicking the **Content Based Routing** component, we selected **Connect to Host** from the pop-up menu. This is an effect of the new Remote Authenticated Administration feature which enables the ND component configuration to be done on a remote client machine. In this case, we were performing this configuration on the Dispatcher machine itself, however the connection still had to take place.

If, at the point where you try to connect to the host to do configuration either through the ND GUI or with the `cbrcontrol` commands, you receive a key related error message, then it is possible that there is an error in the four lines that were added to your ibmproxy.conf file.

We then started the Executor, added our clusters, and to each cluster a port and two servers. As the following figure demonstrates, for our testing purposes we had already added a rule named sunpage for cluster A and a rule named aixpage for cluster B. Following this configuration, the ND GUI appeared as follows:

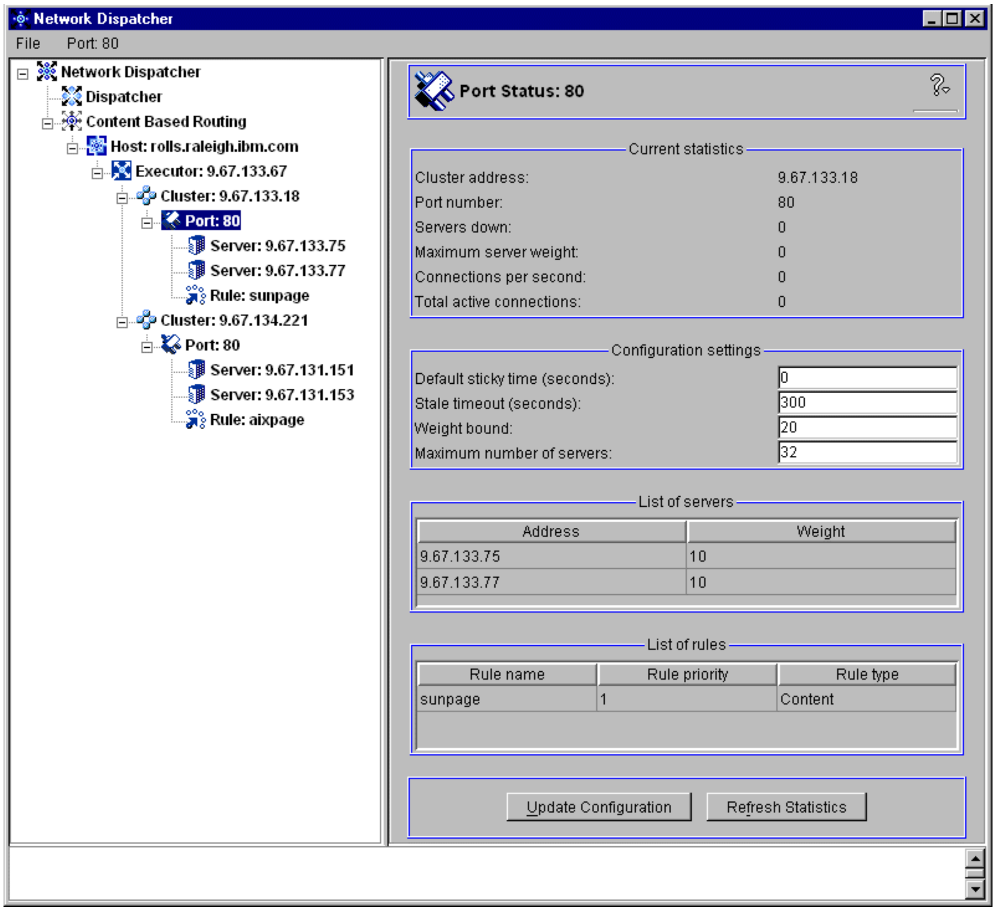

Current statistics

Cluster address:	9.67.133.18
Port number:	80
Servers down:	0
Maximum server weight:	0
Connections per second:	0
Total active connections:	0

Configuration settings

Default sticky time (seconds):	0
Stale timeout (seconds):	300
Weight bound:	20
Maximum number of servers:	32

List of servers

Address	Weight
9.67.133.75	10
9.67.133.77	10

List of rules

Rule name	Rule priority	Rule type
sunpage	1	Content

Figure 323. ND GUI Showing CBR Cluster, Port and Server Information

We will show you now details of how to add two more rules to the 9.67.133.18 cluster. To configure the content rules that CBR uses to determine which servers to load balance the request amongst, we right-clicked the **Port:80** item in the navigation portion of the GUI and selected **Add Rule...** from the pop-up menu, as shown in the following window:

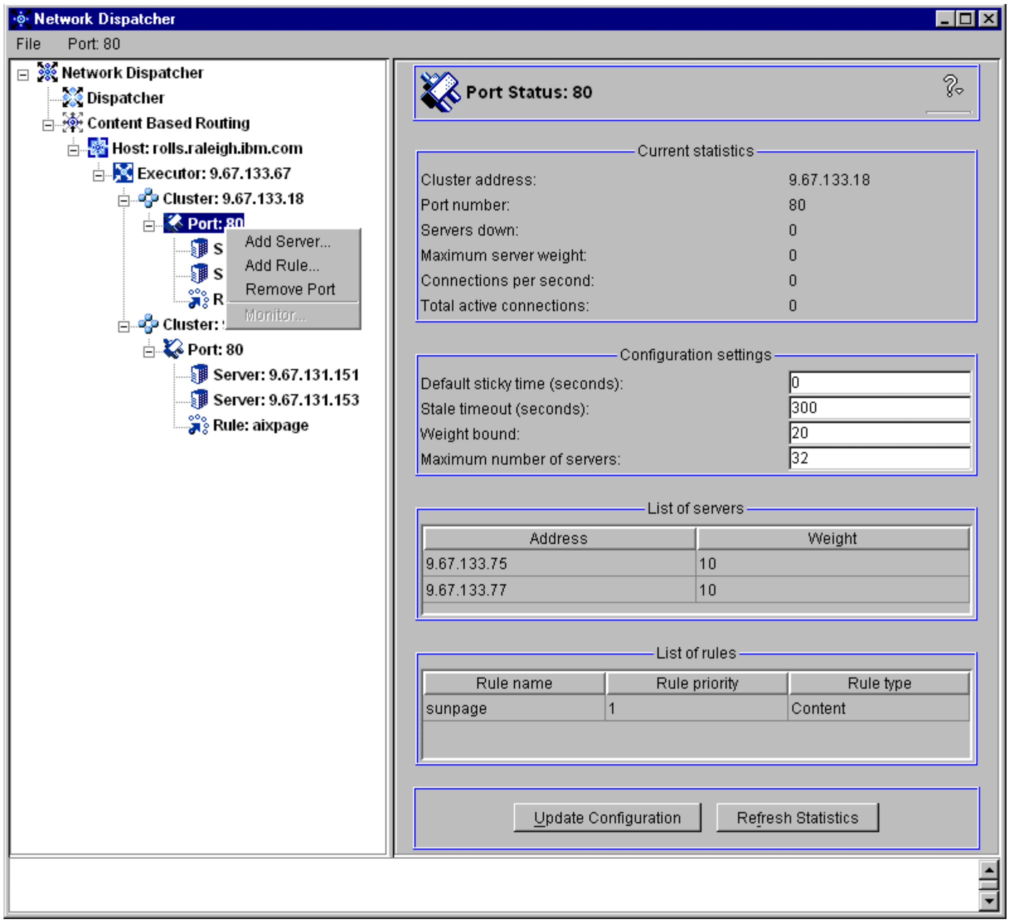

Current statistics

Cluster address:	9.67.133.18
Port number:	80
Servers down:	0
Maximum server weight:	0
Connections per second:	0
Total active connections:	0

Configuration settings

Default sticky time (seconds):	0
Stale timeout (seconds):	300
Weight bound:	20
Maximum number of servers:	32

List of servers

Address	Weight
9.67.133.75	10
9.67.133.77	10

List of rules

Rule name	Rule priority	Rule type
sunpage	1	Content

Figure 324. ND GUI Showing Add Rule Menu Item

Of course you should plan the logic that you want CBR to follow before you start adding rules to your configuration.

Notice that the network interface card on the CBR machine must be aliased to all the cluster addresses defined in the configuration above. This is the same thing you would do in a Dispatcher scenario. However, unlike a Dispatcher scenario, the loopback adapters on the TCP server machines must not be aliased to any cluster address.

Cluster Aliasing Note

The network interface card of the CBR machine must be aliased to all the cluster addresses used in the configuration. However, the loopback interface on the TCP server machines does not need to be aliased; for those who are familiar with this kind of aliasing operations in a Dispatcher scenario, it may seem unusual that similar aliasing is not done with CBR.

The reason for this is that when the packet is sent from the client machine, its destination IP address is the IP address of one of the clusters. When the packet arrives at the CBR machine, the packet is accepted because the network interface card on the CBR machine has been aliased to that cluster address. WTE receives the request and offers CBR the opportunity to examine its clusters and rules for a match. If a match is found, URL name translation is done.

If and when the proxy server sends the request to a clustered TCP server (the selection of which was done by CBR), WTE proxies a new request to the TCP server with the modified URL, and a destination IP address that is the IP address of the TCP server machine. When the TCP server replies, its response goes back to the CBR server, and is cached by WTE if caching is enabled and if WTE caching algorithms require caching of that particular page. For this reason, there is no requirement to alias the cluster address to the loopback interface on the TCP server machine.

On the contrary, Dispatcher keeps the cluster address as the destination address of the packet, and identifies the TCP servers through their Media Access Control (MAC) addresses. For this to work, the TCP servers need to have the cluster address aliased on their loopback interface. An advantage of this is that the server’s response can flow directly to the client, without any need for it to be proxied by the Dispatcher. This would not be a good solution with CBR though, because it would not allow caching.

Another positive effect of this is that a CBR cluster does not have the same restriction as a Dispatcher cluster with regard to the location of the TCP servers relative to the Dispatcher or CBR server. With CBR, because WTE proxies the request to the servers, there is no requirement for the servers to be located on the same LAN as the CBR machine.

Our next step was to create a content rule to direct all requests to cluster 9.67.133.18 that contained the string `productx.html` in the URL to server 9.67.133.75. To accomplish this, in the **Add Rule** dialog window, we filled in the fields as follows:

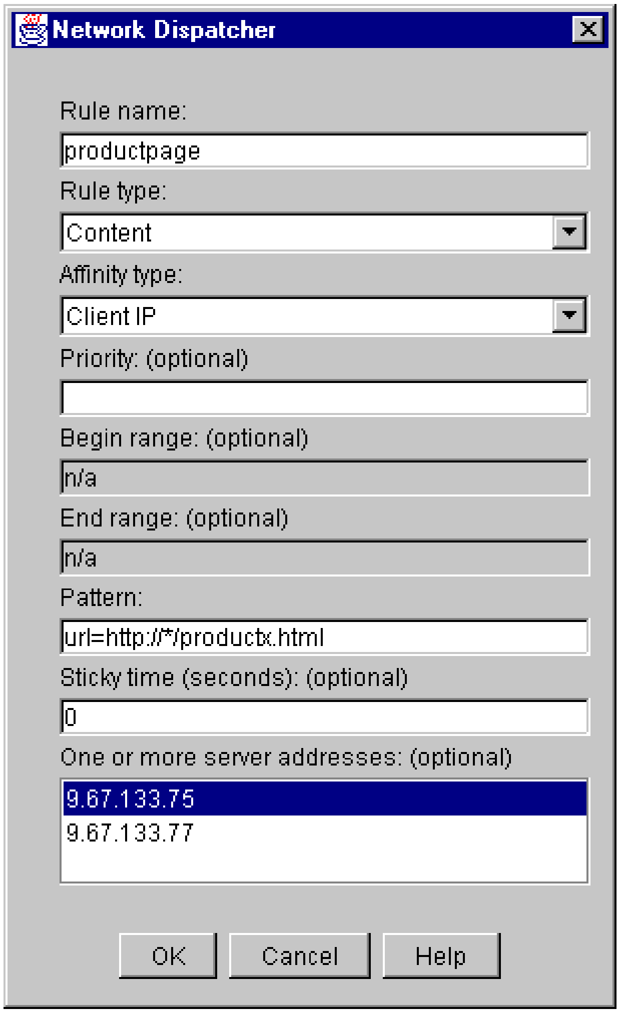

Rule name: productpage

Rule type: Content

Affinity type: Client IP

Priority: (optional)

Begin range: (optional) n/a

End range: (optional) n/a

Pattern: url=http://*/productx.html

Sticky time (seconds): (optional) 0

One or more server addresses: (optional)
9.67.133.75
9.67.133.77

Figure 325. Add Rule Dialog Window for the productpage Rule

- We assigned the name `productpage` to this rule.
- We selected the type of the rule to be **Content**. This is the new rule type added for use by the CBR component.
- The Affinity type field contains one of two possible values: **Client IP** or **Cookie**. Client IP affinity as used in the Dispatcher component can also be used with CBR. The cookie affinity feature applies only to the CBR component and provides a new way to make clients *sticky* to a particular server. This function is enabled by setting the sticky time of a rule to a positive number, and setting the affinity to **Cookie**.

  We left the default value of **Client IP** in the Affinity type field. In the next rule we demonstrate setting cookie affinity and sticky time.
- We also did not put a value in the priority field for the rule. Priorities establish the order in which rules will be reviewed. This parameter accepts integer values. If you do not specify the priority of the first rule you add, CBR will set it by default to `1`. When a subsequent rule is added, by default its priority is calculated to be 10 plus the current lowest priority of any existing rule. For example, assume you have an existing rule whose priority is 30. You add a new rule and set its priority at `25` (which is a higher priority than 30). Then you

add a third rule without setting a priority. The priority of the third rule is calculated to be as 30 + 10 = 40, and so on.

- The Begin range parameter and the End range parameters are not used on Content type rules.
- The Pattern field is used to define the pattern of characters that CBR will match against each client request.

  The pattern must not contain any spaces and can make use of the special characters listed in the following table:

Table 7. Special Characters Allowed in the Pattern Field

Character	Function
`*`	Matches 0 to *x* of any character
`(`	Used for logic grouping
`)`	Used for logic grouping
`&`	Logical AND
`|`	Logical OR
`!`	Logical NOT

The reserved keywords shown in the following table must always be followed by the equal (=) sign:

Table 8. Keywords Followed by the Equal (=) Sign

Keyword	Value
`client`	Client IP address
`url`	URL in request
`protocol`	Protocol section of URL
`path`	Path section of URL
`refer`	Referred URL (quality of service)
`user`	User ID section of URL

In our case, we made use of the `url` reserved word and the wildcard (*) character. We specified the pattern as:

```
url=http://*/productx.html
```

This rule specifies that any URL that is a request for the page productx.html will be a match.

Other examples of valid rule patterns are:

- `url=http://*/*.gif`
- `client=9.32.*`
- `(path=index/*.gif&protocol=http)|(client=9.1.2.3)`
- `!(path=*.jpeg)`

- The Sticky time field is used along with the Affinity type field. We demonstrate the use of this in our next rule.
- The last item on the dialog window shown in Figure 325 on page 384 is a scrolled list of server addresses to optionally choose from. In this case we chose the server whose IP address was 9.67.133.75, meaning that we wanted

this server to serve the client requests containing the string productx.html in the URL.

We then clicked the **OK** button in the dialog window shown in Figure 325 on page 384.

The next rule was added to send requests to both servers in the 9.67.133.18 cluster for client requests for the page purchase.html. In this case, however, we specified an Affinity type of **Cookie** and a Sticky time of `30` seconds. This means that when a client request to the cluster is made and the URL matches the rule (that is, contains the string `purchase.html`), CBR would load balance the request to the best server of the two and then all subsequent HTTP requests from that client to this cluster address would also go to that server for a period of 30 seconds. We called this rule `purchase`. The dialog window in this case appeared as follows:

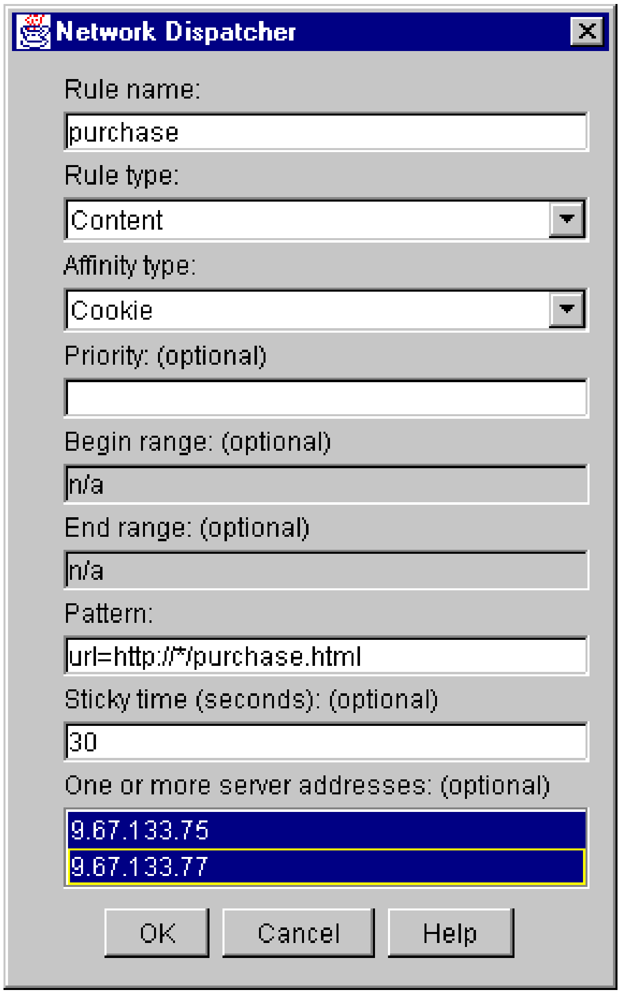

Network Dispatcher

Rule name: purchase

Rule type: Content

Affinity type: Cookie

Priority: (optional)

Begin range: (optional) n/a

End range: (optional) n/a

Pattern: url=http://*/purchase.html

Sticky time (seconds): (optional) 30

One or more server addresses: (optional)

9.67.133.75
9.67.133.77

Figure 326. Add Rule Dialog Window for the productpage Rule

We then clicked the **OK** button in the dialog window above. The GUI reported the configuration was updated with the added rules, as shown in the following figure:

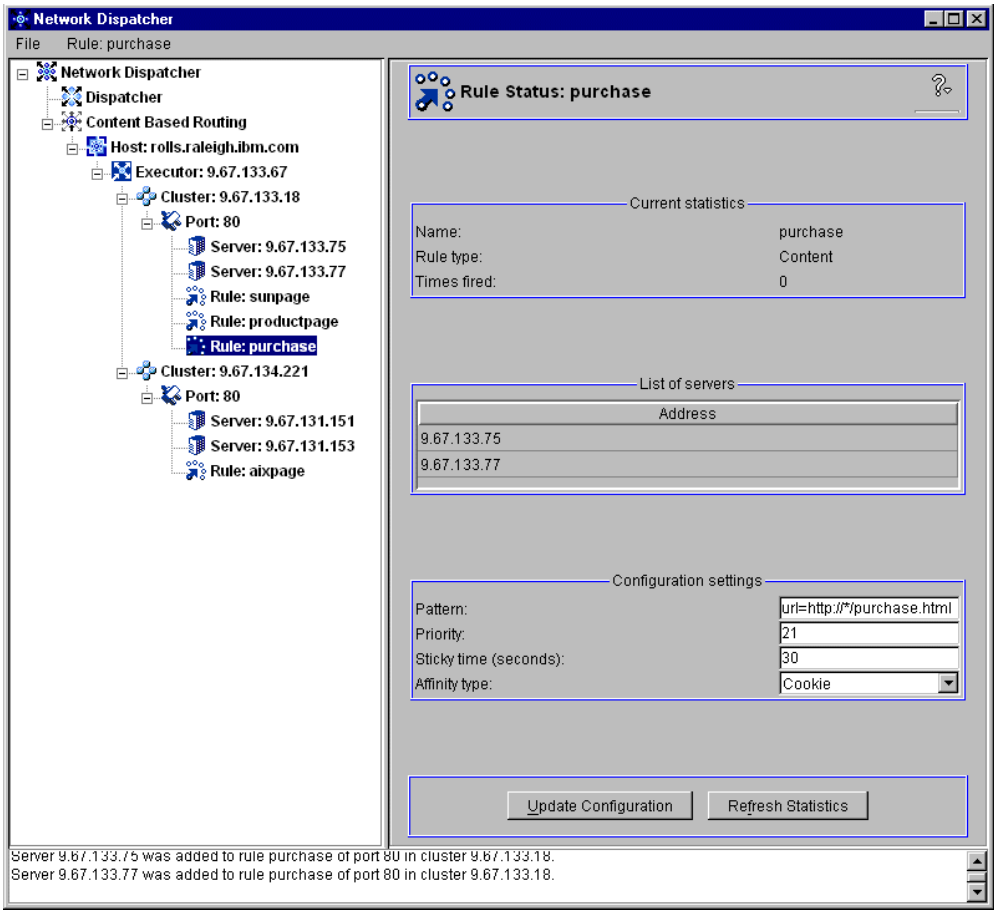

Current statistics

Name:	purchase
Rule type:	Content
Times fired:	0

List of servers

Address
9.67.133.75
9.67.133.77

Configuration settings

Pattern:	url=http://*/purchase.html
Priority:	21
Sticky time (seconds):	30
Affinity type:	Cookie

```
Server 9.67.133.75 was added to rule purchase of port 80 in cluster 9.67.133.18.
Server 9.67.133.77 was added to rule purchase of port 80 in cluster 9.67.133.18.
```

Figure 327. ND GUI Showing the Newly Added productpage Rule

9.3.4 CBR Manager and Advisors

As with the Dispatcher, starting the Manager and using the Advisors is optional. The Manager can be activated with the command:

```
cbrcontrol manager start
```

The typical advisor that is used with CBR is the HTTP advisor, which can be launched by entering the following command:

```
cbrcontrol advisor start http port
```

where `port` is the port configured for the cluster.

If you have configured any Advisors, you must change the Manager proportions to allow the Advisor information to be included in the load balancing decisions. To do this, you should use the command:

```
cbrcontrol manager proportions
```

For simplicity, in the scenarios described in this chapter, we did not make use of the CBR Manager or Advisors functions.

9.3.5 Saving the Configuration

Once we finished configuring CBR, we saved the configuration to a file. To do this from the ND GUI, we right-clicked **Host** and selected **Save Configuration File As...**.

In the resulting pop-up window, we were prompted to enter the name of the configuration file where we wanted to save this information. We entered the file name landon.cfg in the Save Configuration pop-up window. By default this file is placed in the directory configurations under *installbase*, where *installbase* is for the CBR component and varies by operating system. See Table 5 on page 375 for a list of the installbase locations.

The configuration file is saved in ASCII format and contains the list of commands that would be necessary to reconfigure your environment. Following is the content of the landon.cfg file:

```
cbrcontrol cluster add 9.67.133.18

cbrcontrol port add 9.67.133.18:80

cbrcontrol server add 9.67.133.18:80:9.67.133.77

cbrcontrol server add 9.67.133.18:80:9.67.133.75
cbrcontrol rule add 9.67.133.18:80:sunpage type content pattern url=http://*/sunpage.html priority 1 beginrange 0
endrange 0
cbrcontrol rule useserver 9.67.133.18:80:sunpage 9.67.133.75
cbrcontrol rule useserver 9.67.133.18:80:sunpage 9.67.133.77
cbrcontrol rule set 9.67.133.18:80:sunpage stickytime 30

cbrcontrol rule add 9.67.133.18:80:productpage type content pattern url=http://*/productx.html priority 11
beginrange 0 endrange 0
cbrcontrol rule useserver 9.67.133.18:80:productpage 9.67.133.75

cbrcontrol rule add 9.67.133.18:80:purchase type content pattern url=http://*/purchase.html priority 21 beginrange 0
endrange 0
cbrcontrol rule useserver 9.67.133.18:80:purchase 9.67.133.75
cbrcontrol rule useserver 9.67.133.18:80:purchase 9.67.133.77
cbrcontrol rule set 9.67.133.18:80:purchase stickytime 30
cbrcontrol rule set 9.67.133.18:80:purchase affinity cookie

cbrcontrol cluster add 9.67.134.221

cbrcontrol port add 9.67.134.221:80

cbrcontrol server add 9.67.134.221:80:9.67.131.153

cbrcontrol server add 9.67.134.221:80:9.67.131.151

cbrcontrol rule add 9.67.134.221:80:aixpage type content pattern url=http://*/aixpage.html priority 1 beginrange 0
endrange 0
cbrcontrol rule useserver 9.67.134.221:80:aixpage 9.67.131.151
cbrcontrol rule useserver 9.67.134.221:80:aixpage 9.67.131.153
```

Figure 328. Configuration File landon.cfg

Examination of the saved configuration file is interesting because it contains each of the `cbrcontrol` commands that could also be entered from the command line to configure the same environment. If you chose to configure CBR with commands, you can save your configured environment with the command:

`cbrcontrol file save` *`configurationfilename`*

You can then subsequently reload the configuration file with the command:

`cbrcontrol file load` *`configurationfilename`*

9.3.6 Scenario Results

We placed simple HTML files in the document root directories on each of our Web servers. On the Web server with IP address 9.67.133.75 we placed a file named productx.html and another called purchase.html. On the Web server with IP address 9.67.133.77 we placed a different file also called productx.html. Recall that we had defined a rule specifying that client requests containing the string `productx.html` in the URL be load balanced between both of these servers. Each of the files uniquely identified the server that it was located on and the name of the page.

Client Configuration

In a CBR scenario, the client machine does not need any special configuration. In particular, client browsers must not be set to redirect all the requests to the CBR machine. For example, in a real-life situation, it would not be appropriate to require end users to reconfigure their Web browsers before accessing a Web site that uses CBR. The use of CBR in a Web site is completely transparent to end users.

9.3.6.1 Client IP Affinity Demonstration

For our first request, with our browser we requested the productx.html page from our cluster address and received the following page:

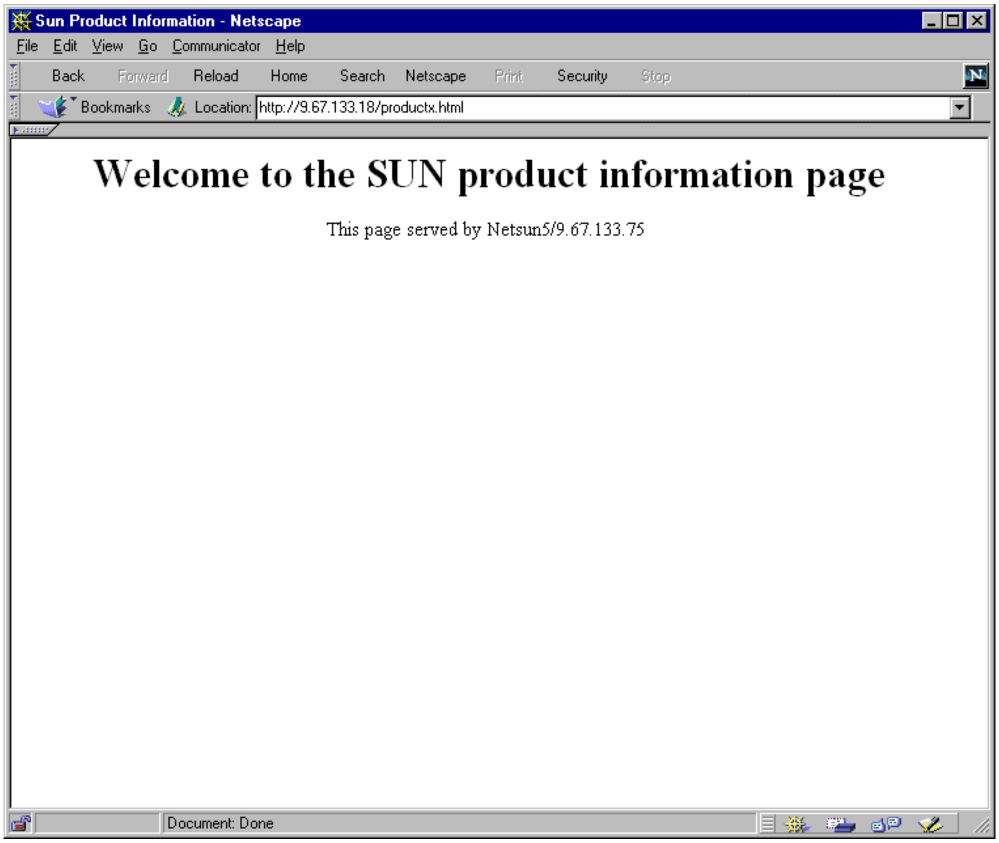

Sun Product Information - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://9.67.133.18/productx.html

Welcome to the SUN product information page

This page served by Netsun5/9.67.133.75

Figure 329. Sun Product Request

To verify that our productpage rule was matched by this request, we used the `cbrcontrol` persistent session command

```
rule rep ::
```

The command line version of this is:

```
cbrcontrol rule report cluster:port:rule
```

but as with other `cbrcontrol` commands, short forms of the keywords can be used. In order to specify that you would like the report to include all clusters, ports and rules, use only the two colon (::) delimiters on the command line. The command and its output are shown in the following figure:

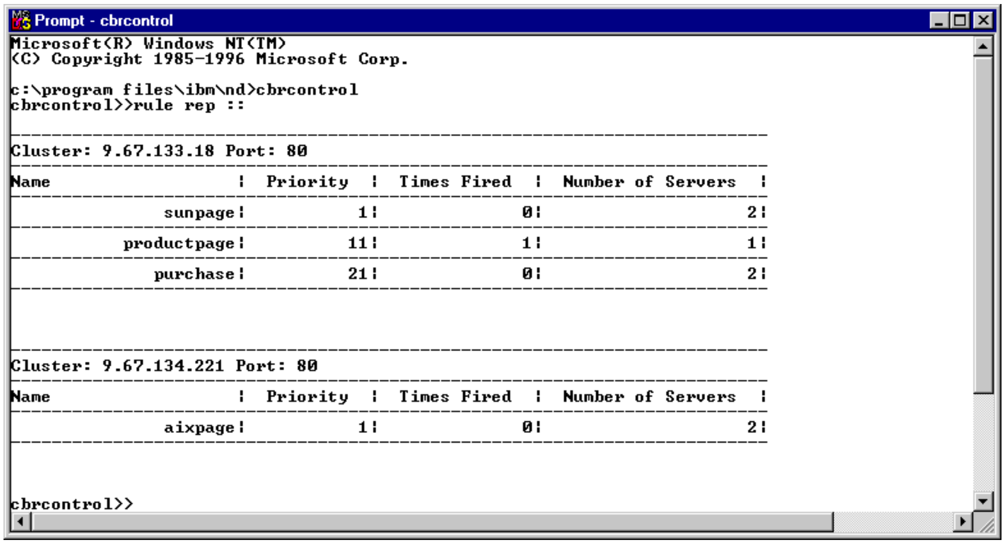

```
Microsoft(R) Windows NT(TM)
(C) Copyright 1985-1996 Microsoft Corp.

c:\program files\ibm\nd>cbrcontrol
cbrcontrol>>rule rep ::

--------------------------------------------------------------------------
Cluster: 9.67.133.18 Port: 80
--------------------------------------------------------------------------
Name                   |  Priority  |  Times Fired  |  Number of Servers  |
--------------------------------------------------------------------------
               sunpage|           1|              0|                    2|
--------------------------------------------------------------------------
           productpage|          11|              1|                    1|
--------------------------------------------------------------------------
              purchase|          21|              0|                    2|
--------------------------------------------------------------------------


--------------------------------------------------------------------------
Cluster: 9.67.134.221 Port: 80
--------------------------------------------------------------------------
Name                   |  Priority  |  Times Fired  |  Number of Servers  |
--------------------------------------------------------------------------
               aixpage|           1|              0|                    2|
--------------------------------------------------------------------------


cbrcontrol>>
```

Figure 330. Rule Report Showing productpage Rule Fired One Time

Since the sticky time is set to 0 seconds, Web requests from the same client should not stick to the same Web server, and subsequent requests should normally load balance between the two Web servers defined in cluster A. If the sticky time were set to a positive number of seconds, CBR would choose one of the two servers the first time a client request arrives that matches the productpage rule, and then would redirect all the client requests coming from the same IP address to the same server until the sticky time expires.

In this case, however, we forced the CBR server to redirect all the requests to the Web server having IP address 9.67.133.75, because this was the only server available in the productpage rule.

9.3.6.2 Cookie Affinity Demonstration

The next test we made was to demonstrate the use of cookie affinity.

Web cookies are simple pieces of information passed between the client Web browser and the Web server during an HTTP transaction. Cookies do not contain any information about the client that the server does not already know and they cannot do anything on the client machine that the client itself cannot already do, provided the browser is within the specifications. Cookies were introduced as an answer to a fundamental problem of the Web's underlying HTTP 1.0 protocol: the lack of a state, or a persistent connection.

The first time a client accesses a Web site that serves cookies, the chosen server sends a cookie to the client browser, identifying the server and some information about which URLs the cookie is good for. The next time the client visits one of those URLs, the browser includes the cookie in its request. As long as the client's future requests contain the cookie, and each request arrives within the sticky time interval, the client will maintain affinity with the initial server.

Prior to making a request for the purchase page that will trigger the 30 second affinity rule, we examined the contents of the local cookie file (cookies.txt) in the browser directory on our client machine and it was empty:

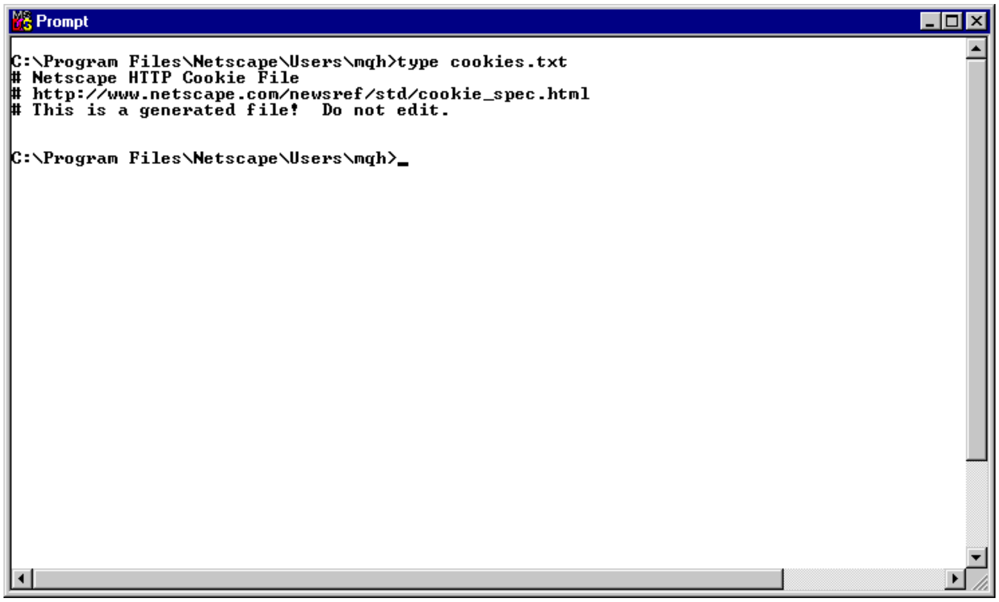

```
C:\Program Files\Netscape\Users\mqh>type cookies.txt
# Netscape HTTP Cookie File
# http://www.netscape.com/newsref/std/cookie_spec.html
# This is a generated file!  Do not edit.


C:\Program Files\Netscape\Users\mqh>_
```

Figure 331. Empty Cookie File Before Request Was Made

Recall that we had enabled a 30 second cookie affinity on our purchase rule by setting the sticky time to `30`, and setting the affinity to **Cookie** when we created the rule*.* This also could have been set with the command:

```
cbrcontrol rule set
```

Once a server is selected by CBR to respond to our request, subsequent requests were also responded to by the same server.

In the 30 second period we made three requests for `http://9.67.133.18/purchase.html`. Each request was responded to by the same server:

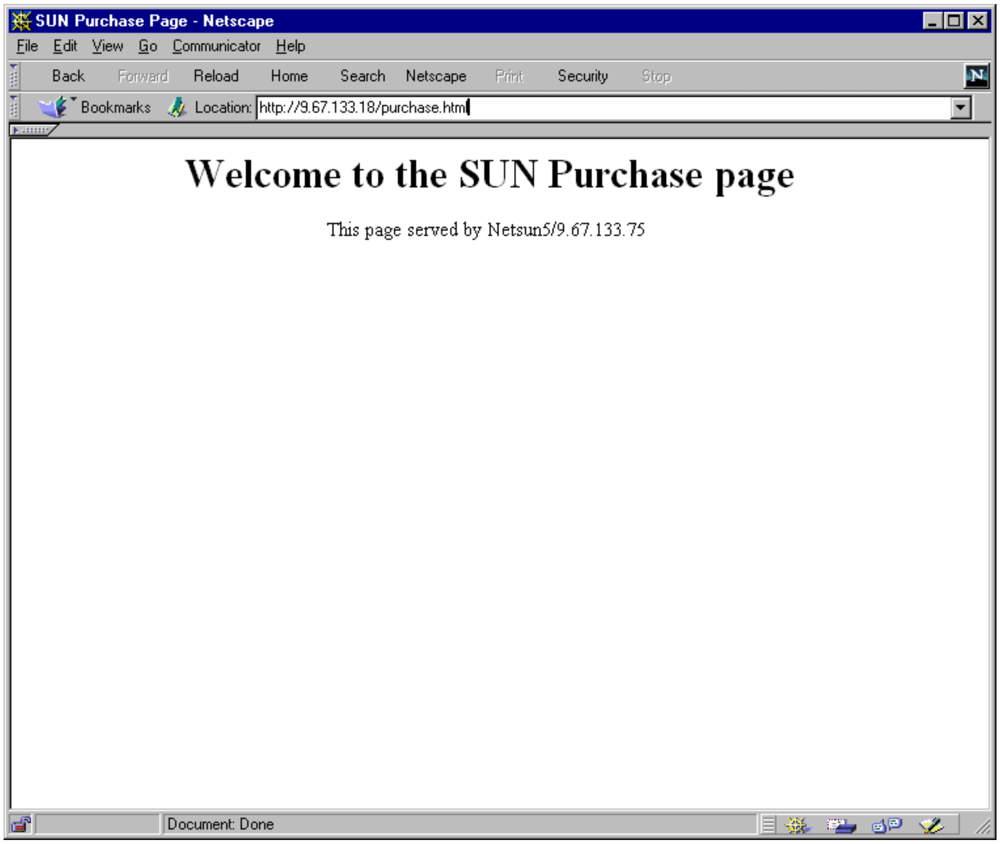
SUN Purchase Page - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://9.67.133.18/purchase.html

Welcome to the SUN Purchase page

This page served by Netsun5/9.67.133.75

Figure 332. Sun Purchase Using a Cookie

We verified that the cookie was set on our browser machine by examining the cookie.txt file:

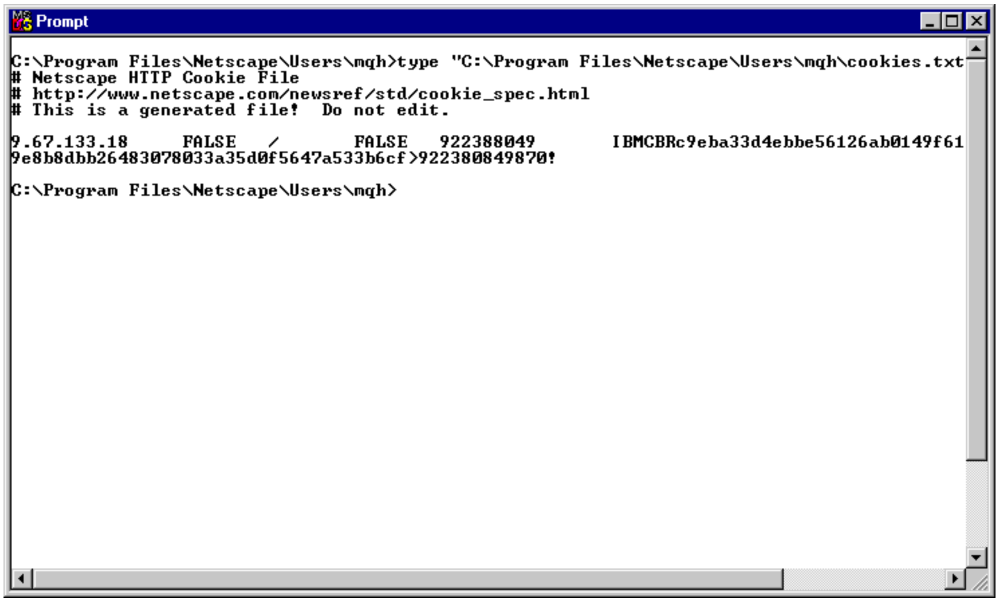

```
C:\Program Files\Netscape\Users\mqh>type "C:\Program Files\Netscape\Users\mqh\cookies.txt
# Netscape HTTP Cookie File
# http://www.netscape.com/newsref/std/cookie_spec.html
# This is a generated file!  Do not edit.

9.67.133.18     FALSE   /       FALSE   922388049       IBMCBRc9eba33d4ebbe56126ab0149f61
9e8b8dbb26483078033a35d0f5647a533b6cf>922380849870!

C:\Program Files\Netscape\Users\mqh>
```

Figure 333. Cookie File Containing one Cookie after Request Was Made

Even though the contents of the cookie are not understandable by us, they are meaningful to the Web server and browser.

After the 30 second sticky time expired, we saw that the CBR server was allowed to choose the other Web server to satisfy the request from the client.

9.4 WTE CacheByIncomingUrl Directive

A new WTE directive has been added for use with CBR in the ibmproxy.conf configuration file. This directive, `CacheByIncomingUrl`, specifies whether to use the incoming URL or the outgoing URL as the basis for generating cache file names. The values of this directive can be `on` or `off`:

- If `CacheByIncomingUrl` is set to `on`, the incoming URL will be used to generate the cache file name. In other words, when this directive is set to `on`, WTE keeps the original URL and uses it to cache the page that it gets back from the ND-changed URL.
- On the other hand, if `off` is specified, CBR rule matching and load balancing will be done on the incoming URL and the resulting URL will be used to generate the cache name. In this case, WTE uses the ND-changed URL to cache the page.

The `CacheByIncomingUrl` directive's default value is `off` in the ibmproxy.conf file.

Notice that `CacheByIncomingUrl` is a *hidden directive* of WTE. In other words, it is possible to alter its value only by manually editing the configuration file ibmproxy.conf; there is no way to change the value of this directive through the Configuration and Administration Forms.

Notice also that a Web page is cached only if WTE decides it should be. WTE does this by looking at the `Expires` tag. If there is not one, it estimates an expiration time by the `Last-Modified` parameter. You should keep this in mind if you are testing CBR caching with sample pages you have just created and which do not contain `Expires` header information. In this case, recently created pages would not be cached, since they would be interpreted by WTE as frequently changed pages, therefore resulting in an apparent failure of the CBR caching function.

Chapter 10. Remote Cache Access

It is not uncommon for a proxy server to receive more traffic than it can handle. One solution is to distribute the traffic between multiple, load balanced proxy servers. However, in this situation, the content of one cache is likely to overlap with the contents of the other caches. Besides unnecessary redundancy, this situation also requires additional bandwidth, because each copy in each server is fetched fresh from the origin. This problem can be minimized by chaining a hierarchy of proxies together, but this still results in additional traffic passing through a given server, and each additional link in the chain adds latency and increases the possibility of a failure. Though useful, proxy chaining does not solve the problem as it does not allow sharing the cache across multiple systems.

IBM Web Traffic Express (WTE) Version 2, the Caching and Filtering component of IBM WebSphere Performance Pack Version 2, has a feature called Remote Cache Access (RCA), which permits a number of WTE servers to share their cache. Eliminating duplicate objects results in bandwidth savings because:

- Objects are not fetched multiple times.
- A larger, combined logical cache yields a higher hit rate.

RCA, which was not available in the previous release of IBM WebSphere Performance Pack, allows multiple proxy servers to cooperate to form a cache array. RCA uses the Cache Array Routing Protocol (CARP) to determine which peer in the array should process the incoming request. If the requested file is not in the proxy’s cache, it will query the other peers in the array to determine which proxy might have the object cached. Therefore, using RCA, multiple proxy servers can distribute the cache contents across their combined, logical cache to improve hit rates and reduce redundancy of cached content.

To run a RCA, WTE demands a shared file system. This can be IBM AFS Enterprise File System (AFS), DFS, NFS, the Windows NT file system, or any other, but the cache directories must have read and write permissions for all the cache array members.

10.1 How RCA Works

In the WTE cache array, all the machines know the existence of the others, and all run a special hash function that maps every URL to an address in the cache. An address in the cache is defined as the pair:

```
(array_member, cache_file_name)
```

This way, the cache space is the combination of all members (WTE servers) and all cache directories. Each URL has a unique location in the cache, and the hash function allows each member to compute the member that owns any given URL.

If a request is made to one proxy, this proxy will first check its local cache. If the object is not found, then it will compute the owner of the URL and ask the owner if it has the object already cached. If so, the first proxy will directly access the cache space of the other proxy, retrieve the requested object and serve it to the requesting client. If not, the requested object will be retrieved directly from the remote content server.

The following diagram shows a graphical representation of an RCA flow:

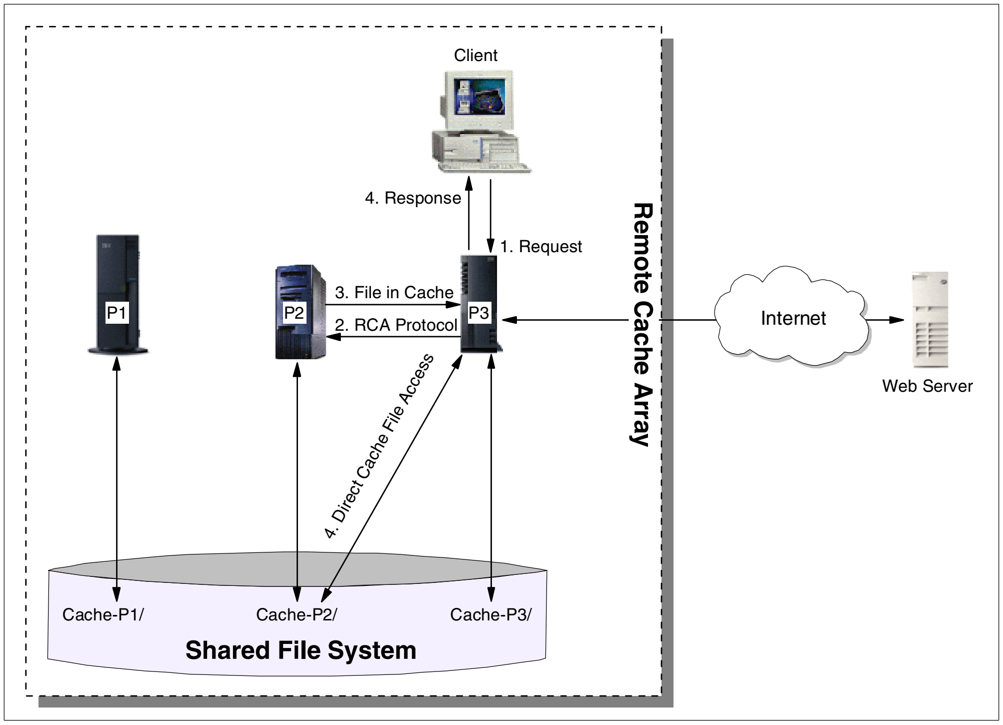

Figure 334. RCA Information Flow

The information flow on RCA is as follows:

1. The client makes a URL request to one of the proxies. In our example, this proxy is P3. Of course, the configuration of the figure above would make much more sense if a load balancing machine were placed between the client and the WTE proxy servers, to load balance the traffic between the proxy servers. Although this is not a mandatory requirement, it is a very good idea to distribute the workload between multiple WTE proxy servers using RCA by using IBM SecureWay Network Dispatcher (ND), the Load Balancing component of IBM WebSphere Performance Pack. In particular, ND offers the WTE Advisor to perform peak load management, as we will see in Chapter 11, “Peak Load Management” on page 401.
2. Using the hash function, P3 knows that this URL may be cached only in P2.
3. P3 makes an RCA request to P2 asking if it has already cached the requested object.
4. At this point, two situations can happen:
   - P2 responds that it has the URL in its cache:
     1. P3 directly accesses P2’s cache to pick up the file corresponding to the requested URL. P2 will not pass all the content of the cached object to P3. In fact, since P3 knows that P2 owns the URL, and since P2’s cache directory is accessible to P3 via the shared file system, P3 accesses P2’s cache directory directly.
     2. P3 returns the cached object to the requesting client
   - P2 responds that it does not have the requested object in its cache:

1. P3 will grab the URL in its original source, the Web or FTP server.
2. P3 will create the cache entry of the URL in P2's cache directory and use RCA to tell P2 about the modification.

Using this procedure, RCA is guaranteed to be an extremely fast protocol, and provides a way to build a scalable array.

10.2 Planning for RCA

When planning to use RCA, you should remember the following:

- The participating proxy servers should be as close as possible, with high-bandwidth connections. It is recommended that members are on the same LAN segment, or communicate over an SP switch, if at all possible. In addition, consider dedicating a network segment just to the inter-node communication (the RCA messages and the shared file system).
- The same file system must be used by all members of an array. You cannot use, for example, AFS on some nodes, NFS on others and the Windows NT file system on others. The recommended file system is AFS.
- Membership in the RCA array should be long-term. The configuration should be as stable as possible.
- Proxy servers should have similar capabilities (for example, CPU, memory, size, cache size).
- Network outages between members of the array must be infrequent.
- There should be less than 100 members in any array.
- All members of the array must run WTE Version 2. You cannot share a cache between WTE Version 2 and Distributed Web Traffic Express (DWTE) Version 1.1[1].
- The members of the array should be load-balanced using ND.

Conversely, RCA is not appropriate when:

- The participating proxy servers are not in close proximity.
- Frequent network outages are expected.
- Servers differ widely in capacity. For example, a PC server with a 500 MB cache and an RS/6000 server with a 20 GB cache should not be in the same array.
- Membership in the RCA array is short term.

10.3 RCA Scenario

Installing and configuring an RCA environment is a typical case where you can install, configure and integrate all of the WebSphere Performance Pack components. In fact, RCA nodes need to share the same file system, and by using AFS you can take advantage of the powerful file sharing features of WebSphere Performance Pack. Moreover, as we have already said, a scenario involving multiple proxy servers sharing the same cache makes more sense if the

[1] DWTE is another IBM caching and filtering proxy server. DWTE also implements RCA. Originally, IBM implemented RCA only in DWTE. Then, the RCA feature was ported also to WTE V2.0.

workload is distributed between the servers. For this purpose, you can use the ND component of IBM WebSphere Performance Pack.

In this section, we show you how we implemented an RCA scenario similar to the one represented below:

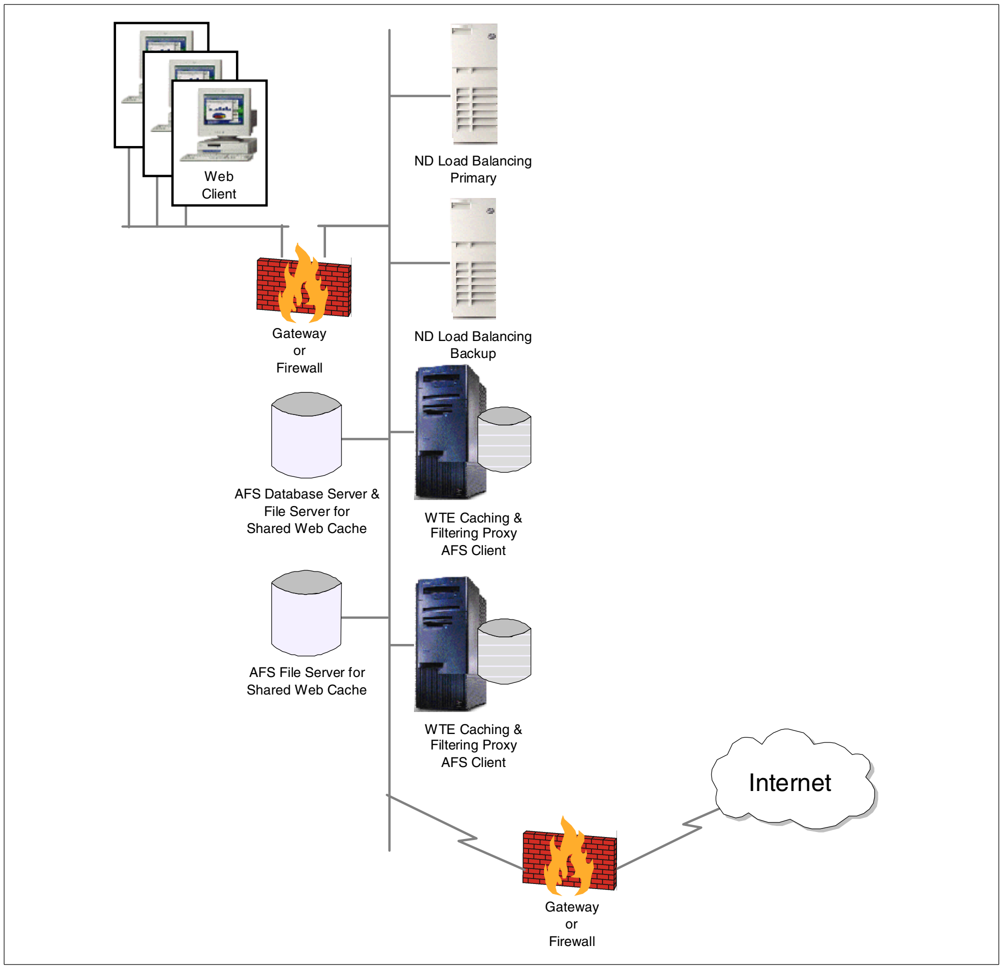

Figure 335. Graphical Representation of the RCA Scenario

10.3.1 Scenario Implementation

1. We constructed an AFS cell called cuzcuz.itso.ibm.com following the steps described in Chapter 3, “AFS Installation and Configuration on AIX” on page 53.

   One of the most useful characteristics of AFS is server ReadOnly volume replication, which means that we can have several copies of the same volume replicated across a number of file servers. This way, clients can access the closest server. Another advantage is that the workload will be balanced through all the servers. But in the case of RCA, the cache directory needs to be updated continuously, so the volumes must be accessed through the ReadWrite path. Even for a remote read access from another proxy, the ReadWrite volume must be available, because a ReadOnly volume is created and replicated manually and statically.

2. We created a separate volume for each proxy, but we didn't replicate those volumes.
3. We configured all the proxy servers as clients of the AFS cell.
4. We granted read and write permissions for all proxies when accessing any other proxy cache directory.

We set up two proxy AIX machines, pluto.itso.ral.ibm.com and rs600030.itso.ral.ibm.com. These machines had essentially the same configurations.

We also configured RCA. This can be done by either accessing the Configuration and Administration Forms or by including the following lines in the ibmproxy.conf configuration file of both the WTE machines:

```
Version    RCA/1.0
ArrayName Cake
Member rs600030 {
        RCAAddr         rs600030.itso.ral.ibm.com
        RCAPort         6969
        CacheSize       100M
        CacheRoot       /afs/.cuzcuz.itso.ibm.com/proxy-itso/cache-rs600030
        Timeout         1000 milliseconds
        BindSpecific    On
        ReuseAddr       Off
}

Member pluto {
        RCAAddr         pluto.itso.ral.ibm.com
        RCAPort         6969
        CacheSize       100M
        CacheRoot       /afs/.cuzcuz.itso.ibm.com/proxy-itso/cache-pluto
        Timeout         1000 milliseconds
        ReuseAddr       Off
}
```

Figure 336. RCA Configuration in ibmproxy.conf

Figure 336 on page 399 shows the configuration for an array called Cake. We can see that the cache directory for each proxy is inside the same AFS cell, and is under the ReadWrite path, .cuzcuz.itso.ibm.com. Through these lines in the WTE configuration file, both proxies know what the cache directories are, and can access them directly. The `RCAPort` directive tells each proxy which port it should connect to, and how to communicate with the other members.

It is very important that the RCA block of the configuration be identical in all the array members. Even the order in which members appear in the RCA block should not differ. One way to make this easier is to store the RCA configuration in a separate file, located somewhere in the shared file system (even on a replicated volume, if using AFS), and then including this configuration using the `RCAConfigFile` directive in the WTE ibmproxy.conf configuration file.

After restarting the WTE servers, we accessed the WTE Configuration and Administration Forms of both the WTE servers, and we verified that the forms reflected the configuration we had issued:

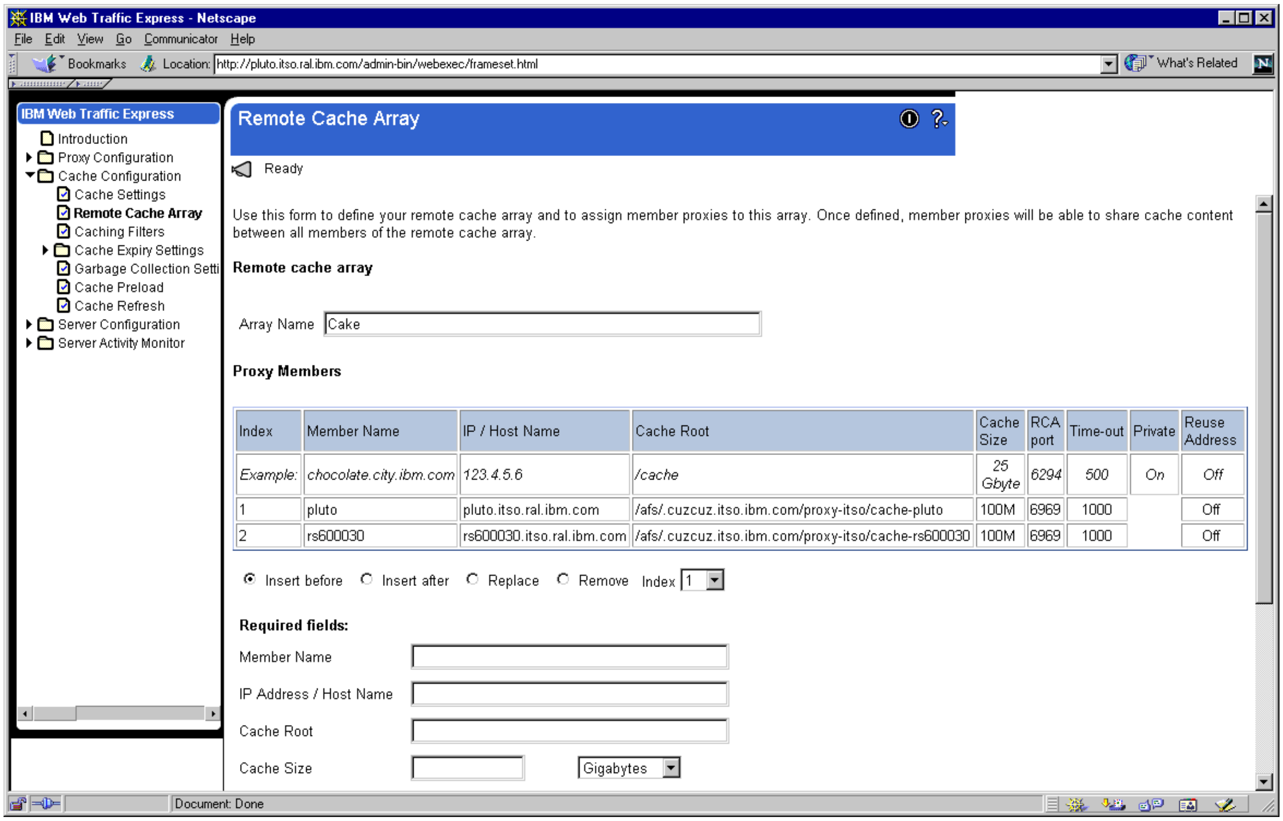

Remote Cache Array

Ready

Use this form to define your remote cache array and to assign member proxies to this array. Once defined, member proxies will be able to share cache content between all members of the remote cache array.

Remote cache array

Array Name: Cake

Proxy Members

Index	Member Name	IP / Host Name	Cache Root	Cache Size	RCA port	Time-out	Private	Reuse Address
Example:	*chocolate.city.ibm.com*	*123.4.5.6*	*/cache*	*25 Gbyte*	*6294*	*500*	*On*	*Off*
1	pluto	pluto.itso.ral.ibm.com	/afs/.cuzcuz.itso.ibm.com/proxy-itso/cache-pluto	100M	6969	1000		Off
2	rs600030	rs600030.itso.ral.ibm.com	/afs/.cuzcuz.itso.ibm.com/proxy-itso/cache-rs600030	100M	6969	1000		Off

● Insert before ○ Insert after ○ Replace ○ Remove Index 1

Required fields:

Member Name

IP Address / Host Name

Cache Root

Cache Size Gigabytes

Figure 337. RCA Configuration in the Web Interface

A simple test permitted us to demonstrate that RCA was functioning correctly.

Chapter 11. Peak Load Management

In WebSphere Performance Pack Version 2, a new Advisor, known as the WTE Advisor for SecureWay Network Dispatcher (ND), enhances load balancing management between multiple Web Traffic Express (WTE) nodes by preventing the ND server from sending new requests to a WTE node that is engaged in garbage collection or cache refresh. This is a new feature implemented for the first time in WebSphere Performance Pack Version 2.

Using the WTE Advisor, an ND server has the ability to detect and react to sudden increases in activity on a WTE node belonging to a cluster. The new WTE Advisor on ND opens a connection to the WTE node, sends a *WTE-specific* HTTP GET request, and interprets the response as a WTE load.

Notice that the WTE Advisor integrates the load balancing component with the caching and filtering component of WebSphere Performance Pack. This new feature is called *peak load management* and is part of the quality-of-service enhancements offered by IBM WebSphere Performance Pack Version 2.

11.1 How Peak Load Management Works

Peak load management is a new feature implemented for the first time in IBM WebSphere Performance Pack Version 2. You can perform peak load management by activating the WTE Advisor on a Dispatcher machine that is load balancing the workload between two or more WTE proxy servers.

The WTE Advisor is a process running on the Dispatcher machine, but is specific for WTE. This process will not work if the load balanced proxy server is not WTE. When the WTE Advisor runs on a Dispatcher machine, this machine sends a specific HTTP request to the load balanced WTE servers, and these reply back with four pieces of information:

- A boolean saying if the cache is reloading
- A boolean saying if garbage collection is running
- The number of active threads
- The maximum number of threads

The WTE Advisor is able to understand these pieces of information passed to it by the WTE machines, and feeds them back to the Dispatcher Manager. The Manager uses the data retrieved by the WTE Advisor to modify the weights associated with the WTE proxy servers, if necessary, in proportion to the importance given by the Dispatcher administrator to the Advisor’s input.

This way, the Dispatcher can detect when a WTE proxy server is performing some CPU-intense activities better than it could if the WTE Advisor were not running. The Dispatcher takes into account all these resource-intensive activities and makes its load balancing decision appropriately, as we demonstrate in the following section.

11.2 Peak Load Management Scenario

In this section, we show you a basic configuration for an ND machine load balancing multiple WTE proxy server nodes. We built the network with five workstations: one client, one Web server, two WTE proxy servers and one ND machine load balancing the traffic between the two WTE servers. We installed the components as shown in the following diagram:

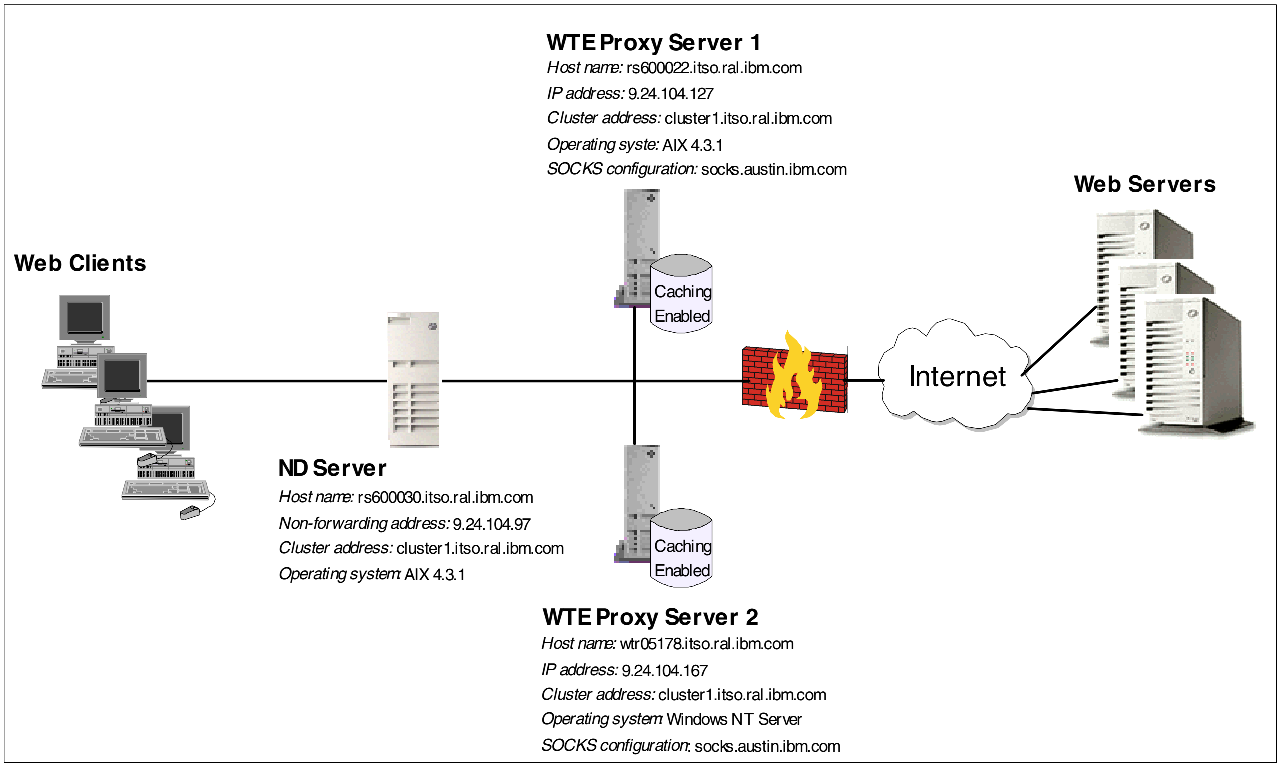
WTE Proxy Server 1
Host name: rs600022.itso.ral.ibm.com
IP address: 9.24.104.127
Cluster address: cluster1.itso.ral.ibm.com
Operating syste: AIX 4.3.1
SOCKS configuration: socks.austin.ibm.com
Web Servers
Web Clients
Caching Enabled
Internet
ND Server
Host name: rs600030.itso.ral.ibm.com
Non-forwarding address: 9.24.104.97
Cluster address: cluster1.itso.ral.ibm.com
Operating system: AIX 4.3.1
Caching Enabled
WTE Proxy Server 2
Host name: wtr05178.itso.ral.ibm.com
IP address: 9.24.104.167
Cluster address: cluster1.itso.ral.ibm.com
Operating system: Windows NT Server
SOCKS configuration: socks.austin.ibm.com

Figure 338. Peak Load Management Scenario

Notice that the diagram above tries to reflect a real-life situation, where multiple clients and multiple Web servers come into the scenario.

The following table provides more details on the WTE cluster configuration:

Table 9. WTE Cluster Configuration

Machine Role	Host Name	IP Address	Operating System
Dispatcher	rs600030	9.24.104.97	AIX 4.3.1
	cluster1	9.24.104.156	
WTE proxy server 1	rs600022	9.24.104.127	AIX 4.3.1
WTE proxy server 2	wtr05178	9.24.104.167	Windows NT Server 4.0

All the machines above belonged to the domain itso.ral.ibm.com. Notice that 9.24.104.156 was the cluster address in this configuration. The host name associated with this cluster address was cluster1.itso.ral.ibm.com.

All the machines shown in Figure 338 on page 402 were provided with a token-ring interface and connected to the same LAN.

The software we used in our scenario is described in the following list:

- Netscape Navigator 4.5 was the Web browser running on the Web client machine.

- The load balancing function was provided by the Dispatcher component of ND Version 2.1.
- The proxy server function was provided by WTE Version 2.0.
- The Web server function was provided by Lotus Domino Go Webserver V4.6.2.5.

11.2.1 Dispatcher Configuration

For the installation and configuration of the ND component of IBM WebSphere Performance Pack, refer to *IBM WebSphere Performance Pack - Load Balancing with IBM SecureWay Network Dispatcher,* SG24-5858. We configured the Dispatcher using the ND configuration graphical user interface (GUI).

The following diagram shows the Dispatcher GUI after the configuration was completed:

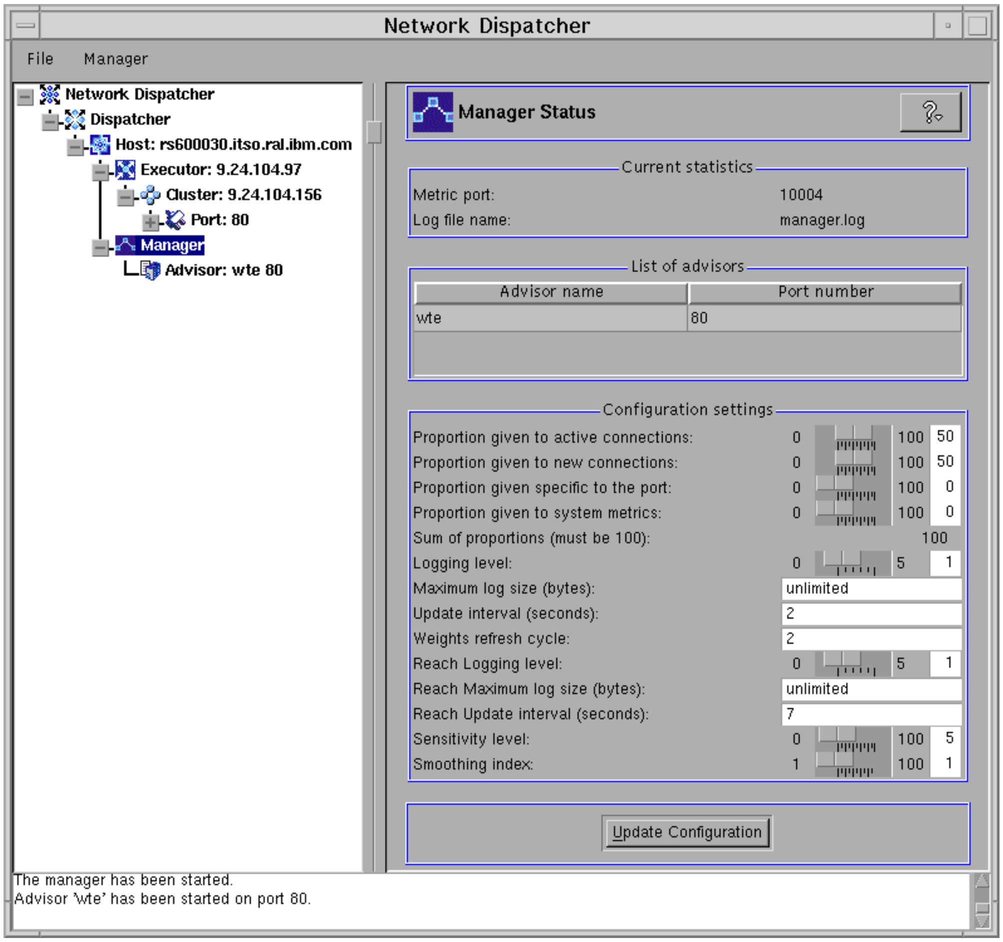

Network Dispatcher
Dispatcher
Host: rs600030.itso.ral.ibm.com
Executor: 9.24.104.97
Cluster: 9.24.104.156
Port: 80
Manager
Advisor: wte 80

Manager Status

Current statistics

Metric port: 10004
Log file name: manager.log

List of advisors

Advisor name	Port number
wte	80

Configuration settings

Proportion given to active connections:	0	100	50
Proportion given to new connections:	0	100	50
Proportion given specific to the port:	0	100	0
Proportion given to system metrics:	0	100	0
Sum of proportions (must be 100):			100
Logging level:	0	5	1
Maximum log size (bytes):			unlimited
Update interval (seconds):			2
Weights refresh cycle:			2
Reach Logging level:	0	5	1
Reach Maximum log size (bytes):			unlimited
Reach Update interval (seconds):			7
Sensitivity level:	0	100	5
Smoothing index:	1	100	1

The manager has been started.
Advisor 'wte' has been started on port 80.

Figure 339. Network Dispatcher GUI - WTE Cluster with WTE Advisors

As you can see from the figure above, we defined a cluster of two WTE proxy servers, load balanced by an ND machine. In this cluster, we also activated the WTE Advisor on port 80. Advisors can be activated as explained in *IBM WebSphere Performance Pack - Load Balancing with IBM SecureWay Network Dispatcher,* SG24-5858. In our configuration, you can see that the Dispatcher was configured to load balance the traffic to the two WTE servers giving 40% of importance to the active connections, 40% of importance to the new connections, and 20% of importance to the feedback coming from Advisors (in this case, from the WTE Advisor). In this scenario, no importance was given to the input coming from system monitoring tools, such as ISS, because no system monitoring tool was part of the configuration.

It would not be appropriate to give a high importance to the Advisors' inputs, but in this case we preferred to set the importance to 20% as we wanted to see the real effects of the WTE Advisors in the load balancing decisions made by the Dispatcher.

11.2.2 Configuration of the WTE Proxy Servers

No special configurations other than a basic one are required on the WTE servers when the WTE Advisor is running on the load balancing Dispatcher machine.

11.2.3 Testing the Peak Load Management Scenario

With the configuration described above, we performed some tests to verify how the Dispatcher reacts to a sudden activity increase on one of the WTE machines.

11.2.3.1 Load Balancing Basic WTE Activities

On the Web client machine, we configured the Web browser to access the Internet through the proxy server 9.24.104.156. This was the cluster address of our WTE proxy server cluster (see Table 9 on page 402). Then, we requested multiple URLs from the clients simultaneously. From the Dispatcher Server monitor GUI, we could see that the two WTE proxy servers were in use simultaneously as shown in the following figure:

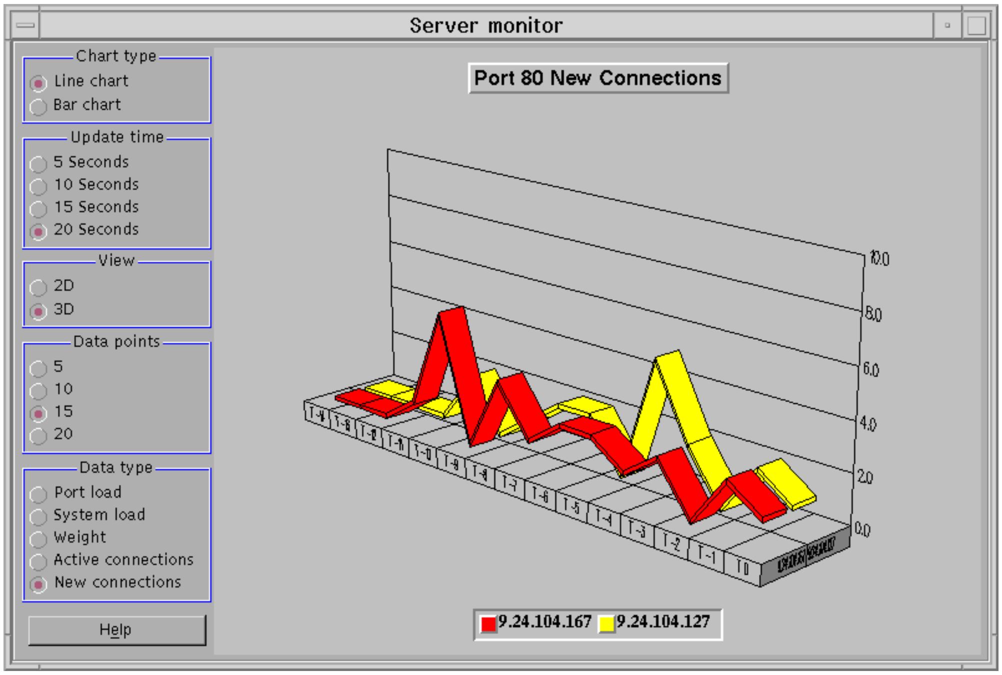

Server monitor

Chart type
● Line chart
○ Bar chart

Update time
○ 5 Seconds
○ 10 Seconds
○ 15 Seconds
● 20 Seconds

View
○ 2D
● 3D

Data points
○ 5
○ 10
● 15
○ 20

Data type
○ Port load
○ System load
○ Weight
○ Active connections
● New connections

Figure 340. Network Dispatcher GUI - WTE Proxy Servers Running

11.2.3.2 Consequences of the WTE Advisor Activation

As mentioned earlier, the WTE Advisor for the ND enhances load management on multiple WTE servers by preventing the Dispatcher from sending new requests to a WTE node that is currently engaged in garbage collection or cache refresh. This function is very useful, since garbage collection and cache refresh are very CPU-intensive activities for a WTE server, and it is very important that the Dispatcher have a way to detect when a WTE server is using most of its resources for these resource-consuming activities.

The following figure shows that, at the beginning of our tests, when there was no particular activity on either of the WTE proxy servers, both WTE servers had the same weights, as shown:

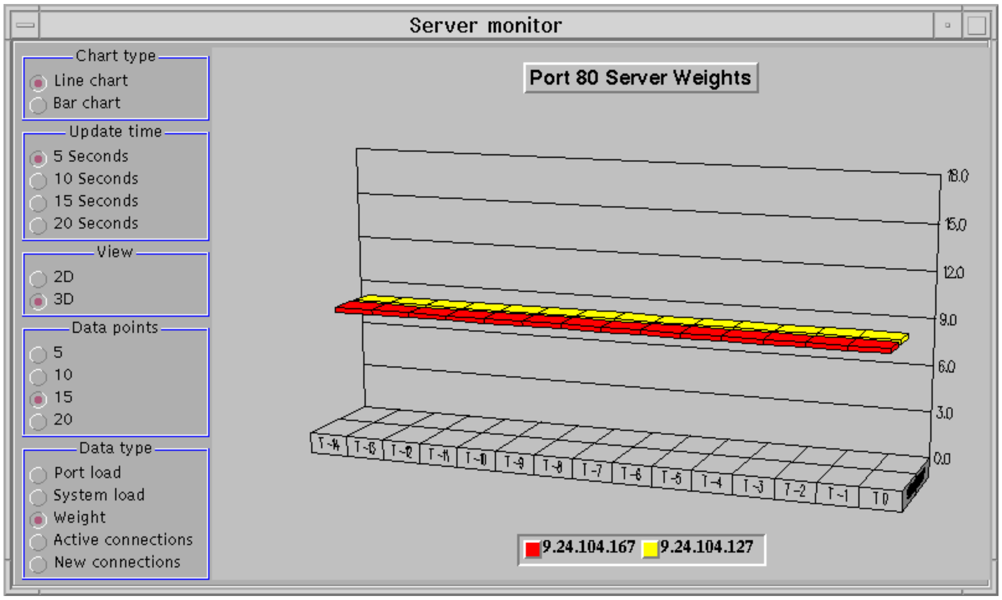

Server monitor

Chart type
● Line chart
○ Bar chart

Update time
● 5 Seconds
○ 10 Seconds
○ 15 Seconds
○ 20 Seconds

View
○ 2D
● 3D

Data points
○ 5
○ 10
● 15
○ 20

Data type
○ Port load
○ System load
● Weight
○ Active connections
○ New connections

Figure 341. Network Dispatcher GUI - WTE Proxy Servers - No Load

We then activated the cache refresh process on the WTE proxy server 1 (see Table 9 on page 402), whose hostname was rs600022. Although cache refresh can start automatically at specific times, for this test we manually forced the WTE proxy server 1 to run the cache refresh process. We did this with the command:

```
cacheagt -r /etc/ibmproxy.conf
```

Through the WTE Advisor, the Dispatcher immediately registered that cache refresh was running on the WTE proxy server 1, and the Dispatcher Manager decreased the weight of this WTE server to reduce the number of connections that were going to be forwarded to it. The Dispatcher Server monitor immediately registered this weight variation:

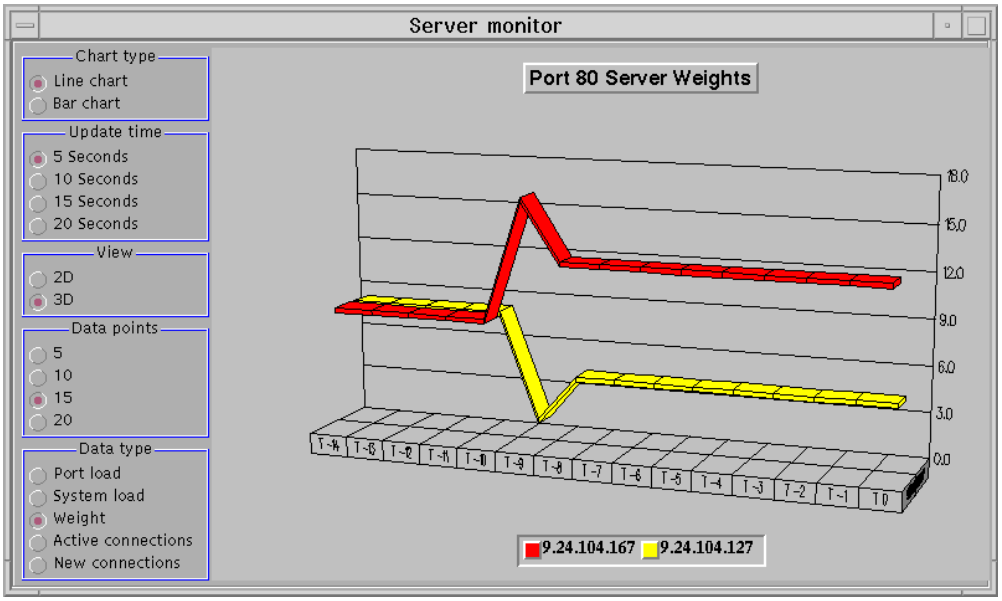

Server monitor

Chart type
● Line chart
○ Bar chart

Update time
● 5 Seconds
○ 10 Seconds
○ 15 Seconds
○ 20 Seconds

View
○ 2D
● 3D

Data points
○ 5
○ 10
● 15
○ 20

Data type
○ Port load
○ System load
● Weight
○ Active connections
○ New connections

Figure 342. Network Dispatcher GUI - WTE Proxy Servers - Cache Refresh in Progress

We can see that the weight of the WTE proxy server 1, having IP address 9.24.104.127, is now much lower than the weight of the WTE proxy server 2, having IP address 9.24.104.167. The following diagram shows the Dispatcher configuration GUI and the weight of each of the WTE proxy servers when the cache refresh process was running:

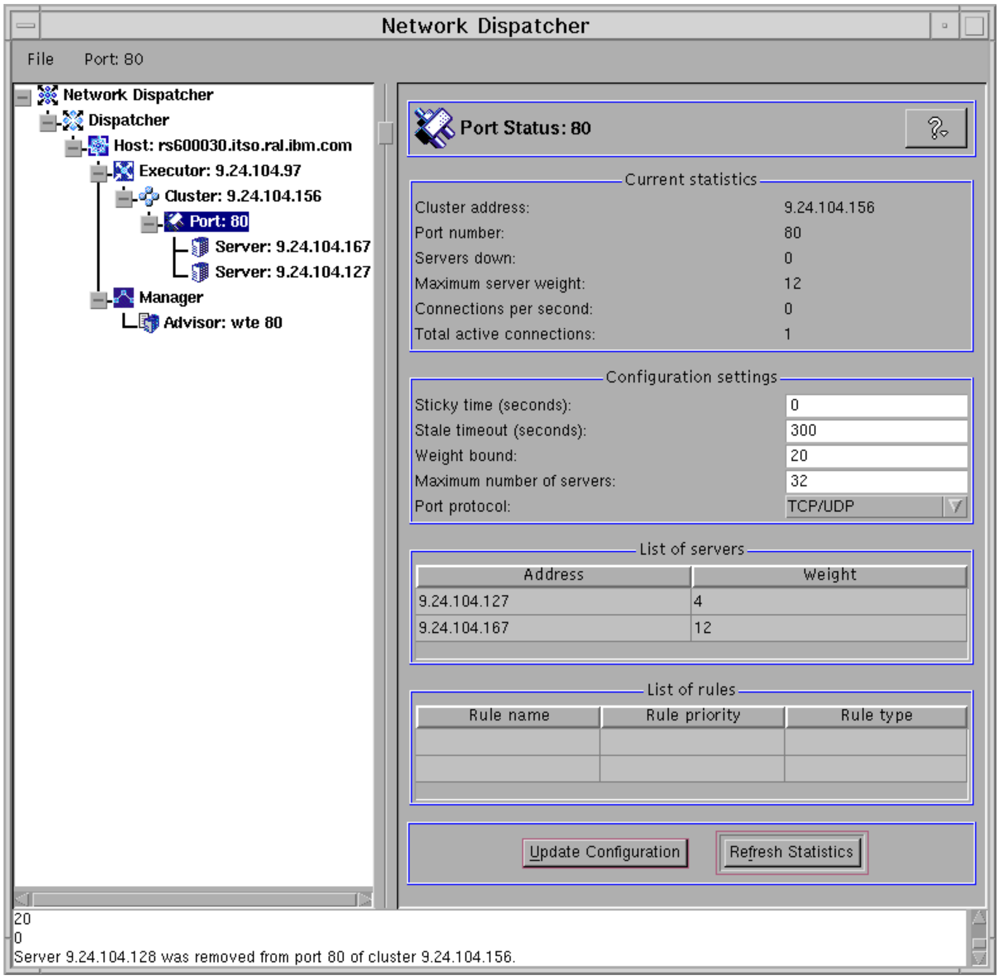

- Network Dispatcher
  - Dispatcher
    - Host: rs600030.itso.ral.ibm.com
      - Executor: 9.24.104.97
        - Cluster: 9.24.104.156
          - Port: 80
            - Server: 9.24.104.167
            - Server: 9.24.104.127
      - Manager
        - Advisor: wte 80

Port Status: 80

Current statistics

Cluster address:	9.24.104.156
Port number:	80
Servers down:	0
Maximum server weight:	12
Connections per second:	0
Total active connections:	1

Configuration settings

Sticky time (seconds):	0
Stale timeout (seconds):	300
Weight bound:	20
Maximum number of servers:	32
Port protocol:	TCP/UDP

List of servers

Address	Weight
9.24.104.127	4
9.24.104.167	12

List of rules

Rule name	Rule priority	Rule type

```
20
0
Server 9.24.104.128 was removed from port 80 of cluster 9.24.104.156.
```

Figure 343. Network Dispatcher GUI - Cache Refresh in Progress

This demonstration completes our scenario on the peak load management feature of IBM WebSphere Performance Pack Version 2.

Chapter 12. Common Configuration

The Common Configuration utility is a new feature of IBM WebSphere Performance Pack Version 2. In Version 2, it provides access to seven browser-based wizards that can be used to configure specific aspects of the three components of the product:

- IBM AFS Enterprise File System (AFS):
  - Getting Started with AFS
  - Add an AFS Web volume
  - Replicate an AFS Web volume
  - Move an AFS Web volume
  - Delete an AFS Web volume
- IBM Web Traffic Express (WTE):
  - Platform for Internet Content Selection (PICS) filtering wizard
- IBM SecureWay Network Dispatcher (ND):
  - Add a Network Dispatcher cluster

The wizards do not provide access to all the functions of the individual component configuration programs. Use the GUIs provided with each of the components to perform additional product-specific configuration.

12.1 Installation

The Common Configuration utility must be installed and configured as a separate component. The installation and configuration are supported on the three platforms where IBM WebSphere Performance Pack itself is supported.

12.1.1 Planning for the Installation

The first step is to choose a machine that will perform the function of *configuration server* for your environment. The Common Configuration utility allows for remote configuration of the WebSphere Performance Pack components. In other words, it is not necessary to install the Common Configuration utility on the same ND, AFS, or WTE machine you will use the Common Configuration utility for.

Access to the Common Configuration utility machine is Web based.

The following diagram shows how the Common Configuration utility works in relation to all of the other WebSphere Performance Pack components:

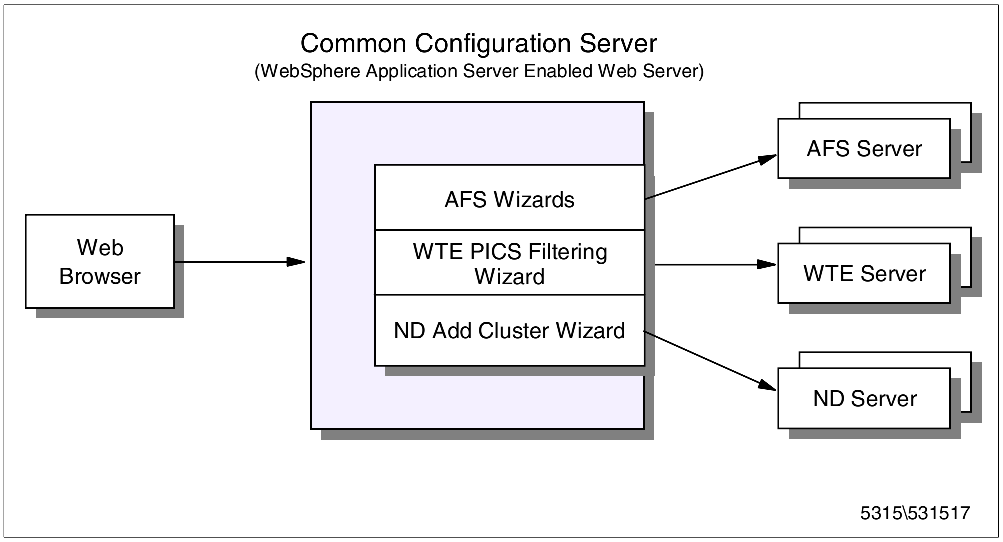
Common Configuration Server
(WebSphere Application Server Enabled Web Server)
Web Browser
AFS Wizards
WTE PICS Filtering Wizard
ND Add Cluster Wizard
AFS Server
WTE Server
ND Server
5315\531517

Figure 344. Common Configuration Server Environment

Because the Common Configuration utility is implemented by using IBM WebSphere Application Server and associated Java servlets (see 1.7.1, “IBM WebSphere Application Server” on page 23), it is necessary that your chosen configuration server machine meet these two prerequisites:

1. A Web server must be installed on the Common Configuration machine. This Web server must be supported by WebSphere Application Server on your platform. See 12.1.2, “Installing the Common Configuration Utility” on page 410 for a list of the supported Web servers on your platform.

   You need to know the document directory and the configuration file directory for whichever Web server you have installed.

2. The Web server must have the ability to serve files that are contained on one of the machine’s local file systems. It cannot be set up to serve files only from AFS.

Notice that it is not necessary to install WebSphere Application Server on the Common Configuration server machine. If WebSphere Application Server is not installed yet, the Common Configuration installation process installs the WebSphere Application Server filesets transparently during the installation routine.

Once these prerequisites are met, you can proceed to install the Common Configuration on your Common Configuration server machine.

12.1.2 Installing the Common Configuration Utility

The process of installing the Common Configuration utility with the Java InstallShield differs only slightly on the three platforms that it can be installed on. We will demonstrate the installation and note where the differences are.

We performed our installation of the Common Configuration utility on a uniprocessor IBM RS/6000 43P having 192 MB of RAM, 2.2 GB of hard disk and one token-ring interface. This RS/6000 had AIX Version 4.3.1, Java Runtime Environment (JRE) 1.1.6 and IBM HTTP Server Version 1.3.3 installed on it.

We also installed the Common Configuration utility on an IBM PC 365 running Windows NT Server4.0. This machine had Java Development Kit (JDK) 1.1.6 and IBM HTTP Server Version 1.3.3 installed on it.

One of the Web servers listed below that is supported by the WebSphere Application Server must be installed on your machine:

- On the AIX platform:
  - IBM HTTP Server V1.3.3
  - Lotus Domino Go Webserver V4.6.2
  - Netscape Enterprise V3.01
  - Netscape Enterprise V3.51
  - Apache HTTP Server V1.3.2
- On the Windows NT platform:
  - IBM HTTP Server V1.3.3
  - Lotus Domino Go Webserver V4.6.2
  - Netscape Enterprise v3.01
  - Netscape Enterprise v3.5
  - Apache v1.3.2
  - IIS 2/3.0
  - IIS 4.0
- On the Solaris platform:
  - IBM HTTP Server V1.3.3
  - Apache Server V1.3.2
  - Lotus Domino Go Webserver V4.6.2.5
  - Netscape Enterprise Server V3.01 and V3.51 (recommend V3.5.1)
  - Netscape FastTrack Server V3.01

The installation program for the Common Configuration utility makes use of Java InstallShield’s setup class. For this reason you are required to pre-install the Java Virtual Machine (JVM). On AIX, only the Java Runtime Environment (JRE) is sufficient as it contains the JVM, the Java platform core classes, and supporting files. On Windows NT we found it necessary to install the entire JDK to perform the installation. We used JDK Version 1.1.6.

If Java is not installed on your machine or installed at a level lower than 1.1.6, you can find the install image for JDK 1.1.6:

- For AIX at `http://www.ibm.com/java/jdk/download/index.html` or on the IBM WebSphere Performance Pack Version 2 CD-ROM
- For Windows NT and Sun Solaris at `http://www.javasoft.com` or on the IBM WebSphere Performance Pack Version 2 CD-ROM

The installation of the JDK is described in the IBM redbook *Network Computing Framework Component Guide,* SG24-2119.

To prepare for the installation of the Common Configuration utility, follow the steps listed below:

1. Insert the IBM WebSphere Performance Pack Version 2 CD-ROM in the CD-ROM drive.
2. This step will differ on each platform:
   - On AIX, from a command line, enter the following commands:

```
mkdir /cdrom
mount -rv cdrfs /dev/cd0 /cdrom
cd /cdrom/aix
```

- On Windows NT, enter:

  ```
  E:
  cd nt
  ```

  where E is assigned to the CD-ROM drive.

- On Solaris, insert the IBM Websphere Performance Pack CD-ROM in the CD-ROM drive. The system will automatically mount the WebSphere Performance Pack CD-ROM as /cdrom/websphere. Then, from a command prompt, enter the following command:

  ```
  cd /cdrom/websphere/sun
  ```

3. To start the Java InstallShield installation program, enter:

   ```
   java setup
   ```

The first screen you will see is the Welcome window. After clicking the **Next** button, you are prompted to enter the destination location. Note that this is a working directory for Websphere Performance Pack. The default location on Windows NT is C:\WSPP; on Solaris /opt/WSPP; and on AIX /usr/lpp/WSPP, as shown in the following window:

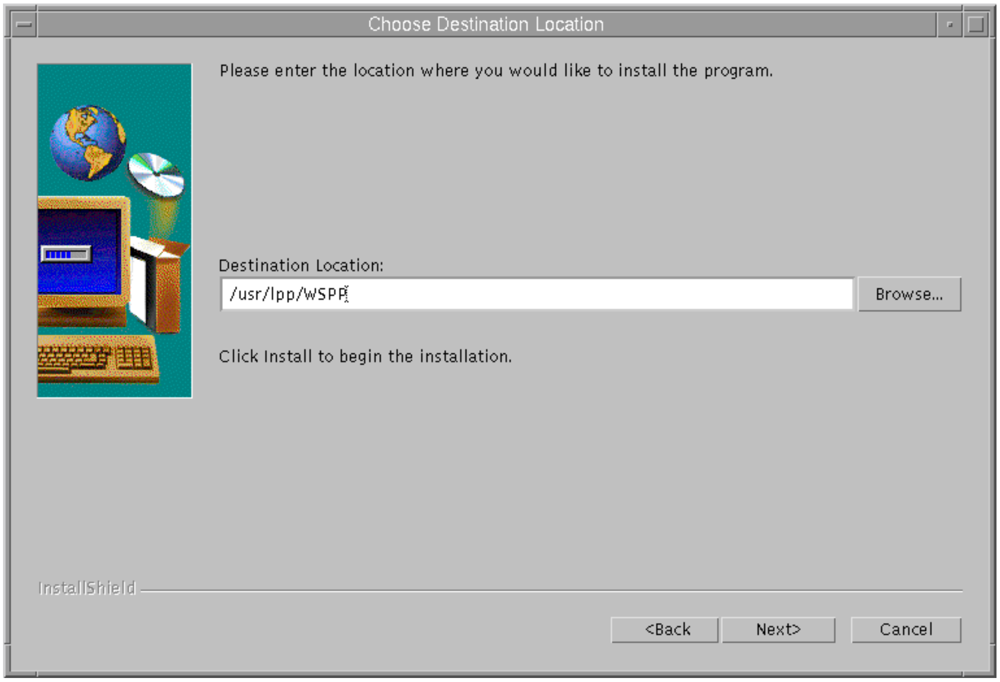

Please enter the location where you would like to install the program.

Destination Location:

/usr/lpp/WSPP

Click Install to begin the installation.

Figure 345. Choose Destination Location

In the above window, you are prompted to click **Install** to proceed. The button does not exist, due to a known InstallShield for Java problem. Click the button labeled **Next** to continue instead. If the destination directory does not exist on your system, you will be presented with a question dialog asking you if you would like the location to be created. Then, you will be presented with a window allowing you to choose which IBM WebSphere Performance Pack V2 components you want to install:

- File Sharing
- Load Balancing
- Caching and Filtering
- Common Configuration

This is shown in the following figure:

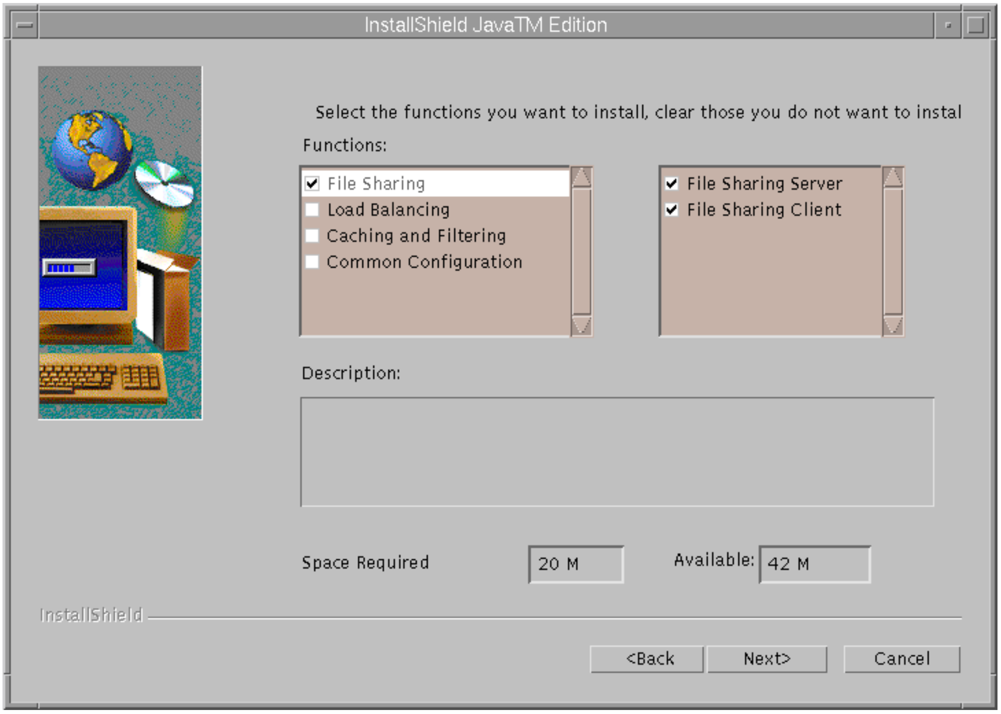

InstallShield JavaTM Edition

Select the functions you want to install, clear those you do not want to instal

Functions:

☑ File Sharing
□ Load Balancing
□ Caching and Filtering
□ Common Configuration

☑ File Sharing Server
☑ File Sharing Client

Description:

Space Required 20 M

Available: 42 M

Figure 346. Choose the Component to Install (accinst01)

When we selected **Common Configuration**, a description of that component was displayed as well as the space required to install it. This window also shows you how much space is available in the destination location file system.

If you do not have enough free space in your destination location file system and you click **Next**, you will see a warning window letting you know that there is not enough space to perform the installation. You will then have to use the operating system utilities to increase the size of the file system so that there is enough space for the installation to be performed. After we had successfully increased the size of the /usr file system, we were not able to continue with the Java InstallShield process. In order for the available space field (as seen in Figure 346 on page 413) to be updated to reflect the new amount of available space in the target directory, it was necessary to backup one screen by clicking the **Back** button followed by the **Next** button. Then, with the Common Configuration component selected, we were able to click the **Next** button.

In the next step, we had to provide some information to the install process. The type of information that the user is prompted for next depends on whether or not IBM WebSphere Application Server is already installed on the machine. Furthermore, in both cases the information that the user is prompted for is different for each of the supported platforms and so in the next sections we will show all of the possibilities.

12.1.2.1 Providing Information on AIX and Solaris

The information you provide during the installation depends on whether or not WebSphere Application Server is already installed on your system.

1. **WebSphere Application Server is not already installed**

   When WebSphere Application Server was not preinstalled on our AIX machine, we saw the following window:

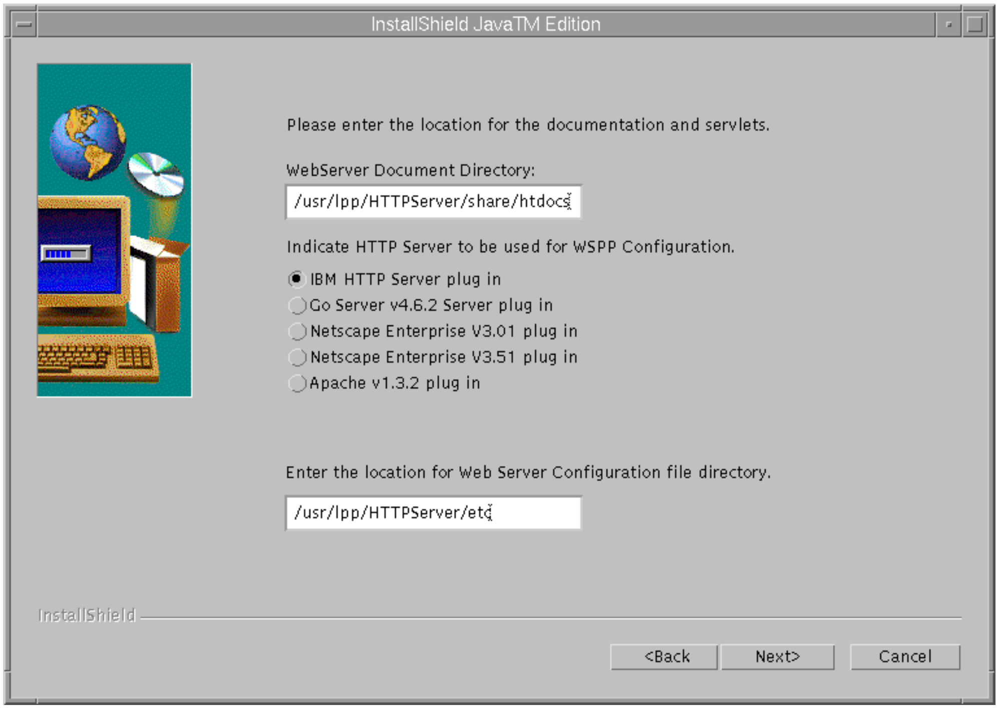

InstallShield JavaTM Edition

Please enter the location for the documentation and servlets.

WebServer Document Directory:

/usr/lpp/HTTPServer/share/htdocs

Indicate HTTP Server to be used for WSPP Configuration.

● IBM HTTP Server plug in
○ Go Server v4.6.2 Server plug in
○ Netscape Enterprise V3.01 plug in
○ Netscape Enterprise V3.51 plug in
○ Apache v1.3.2 plug in

Enter the location for Web Server Configuration file directory.

/usr/lpp/HTTPServer/etc

Figure 347. Request to Supply Web Server Configuration Details

The above window will be displayed if you do not have WebSphere Application Server already installed. In order for the correct Web server plug-in to be installed, you will need to select one of the radio buttons to inform the installation process which Web server is installed on the machine.

In addition, you are asked to provide two directories:

   1. The first piece of information the installation process asks you for is the Web server document directory. The installation process requires this information as it places some files in this directory that are used when the Common Configuration utility is launched.
   2. The second piece of information the installation process asks you for is the Web server configuration file directory. Because the installation in this case is going to install WebSphere Application Server, it must make modifications to the Web server configuration file so that the Web server will be able to work together with WebSphere Application Server.

You will receive warning messages if you try to continue without filling in either of the two fields or if you set the location for the Web server configuration file incorrectly. As we were using the IBM HTTP Server, we selected the radio button for it and filled in the associated directory information.

2. **WebSphere Application Server is already installed**

When WebSphere Application Server was already installed on our AIX machine, we saw the following window:

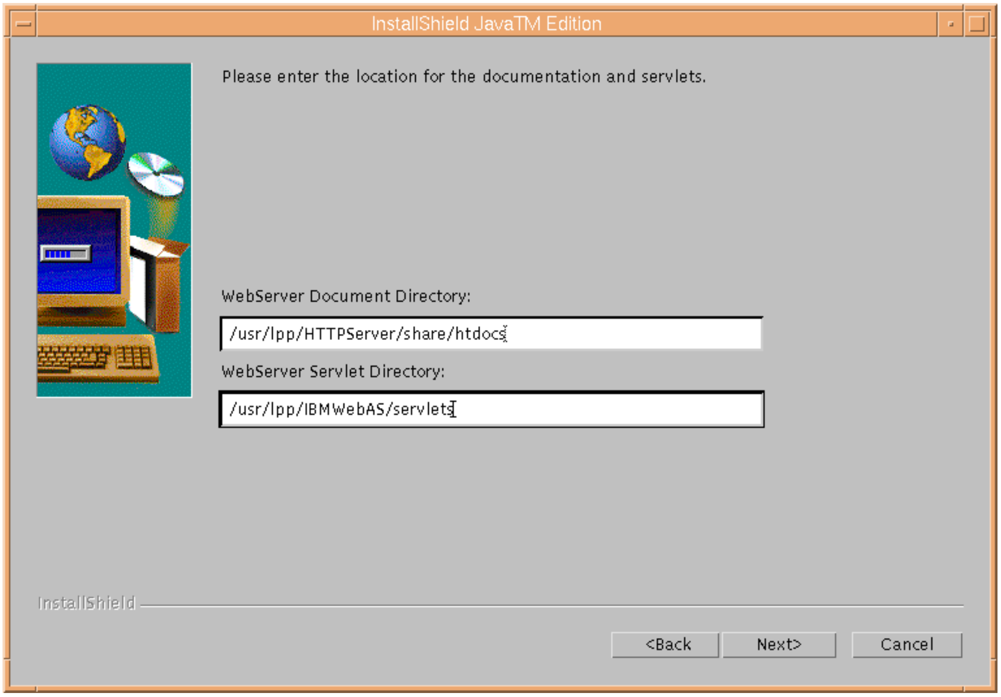

InstallShield JavaTM Edition

Please enter the location for the documentation and servlets.

WebServer Document Directory: /usr/lpp/HTTPServer/share/htdocs

WebServer Servlet Directory: /usr/lpp/IBMWebAS/servlets

Figure 348. Request to Supply Web Server and Servlet Configuration Details

In this case, you are again asked to supply two directory locations to the install process:

1. The first piece of information it asks you for is the same as in Point 1 on page 414: that is, the Web server document directory. The installation process requires this information as it places some files in this directory that are used when the Common Configuration utility is launched. We supplied the default document directory for IBM HTTP Server as that was the Web server we had preinstalled.
2. The second piece of information is different from Point 2 on page 414. In this case, it asks where the Web server servlet directory is. This directory is where the Common Configuration servlets will be placed during the installation process. We supplied the default servlet directory for WebSphere Application Server on AIX in the second field.

What we have just seen for AIX is very similar in the Solaris platform.

12.1.2.2 Providing Information on Windows NT

The information you provide during the installation depends on whether or not WebSphere Application Server is already installed on your system.

1. **WebSphere Application Server is not already installed**

   When WebSphere Application Server was not preinstalled on our Windows NT machine, we saw the following window:

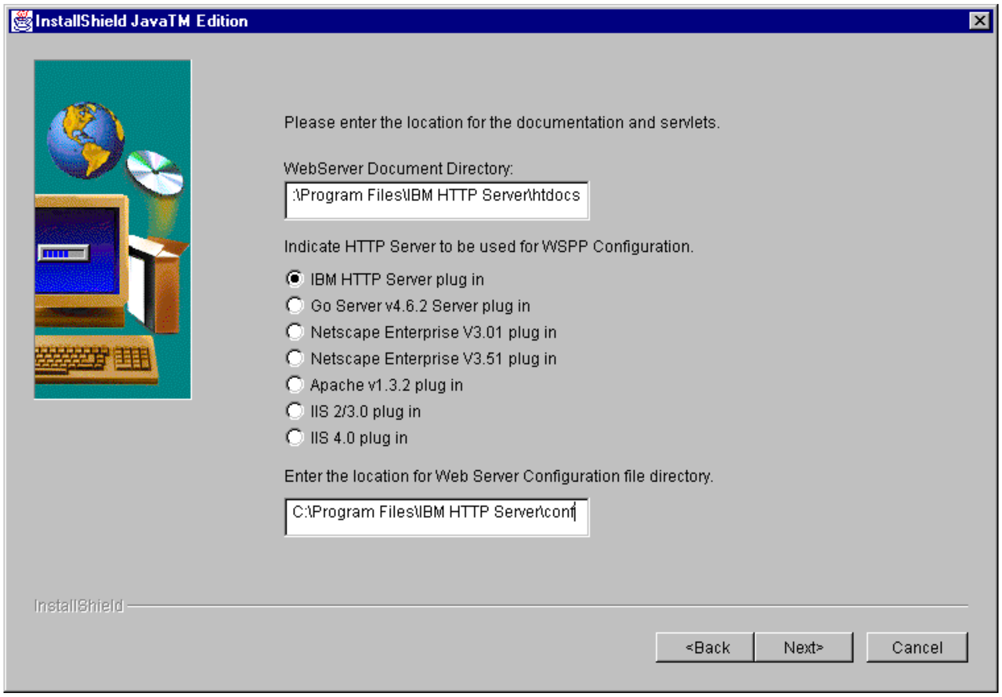

Please enter the location for the documentation and servlets.

WebServer Document Directory:

:\Program Files\IBM HTTP Server\htdocs

Indicate HTTP Server to be used for WSPP Configuration.

- ● IBM HTTP Server plug in
- ○ Go Server v4.6.2 Server plug in
- ○ Netscape Enterprise V3.01 plug in
- ○ Netscape Enterprise V3.51 plug in
- ○ Apache v1.3.2 plug in
- ○ IIS 2/3.0 plug in
- ○ IIS 4.0 plug in

Enter the location for Web Server Configuration file directory.

C:\Program Files\IBM HTTP Server\conf

Figure 349. Request to Supply Web Server Configuration Details

The above window will be displayed if you do not have WebSphere Application Server already installed. In order for the correct Web server plug-in to be installed, you will need to select one of the radio buttons to inform the installation process which Web server is installed on the machine.

In addition, you are asked to provide two directories:

1. The first piece of information the installation process asks you for is the Web server document directory. The installation process requires this information as it places some files in this directory that are used when the Common Configuration utility is launched.
2. The second piece of information the installation process asks you for is the Web server configuration file directory. Because the installation in this case is going to install WebSphere Application Server, it must make modifications to the Web server configuration file so that the Web server will work together with WebSphere Application Server.

You will receive warning messages if you try to continue without filling in either of the two fields or if you set the location for the Web server configuration file incorrectly. As we were using the IBM HTTP Server, we selected the radio button for it and filled in the associated directory information.

2. WebSphere Application Server is already installed

When WebSphere Application Server was already installed on our Windows NT machine, we saw the following window:

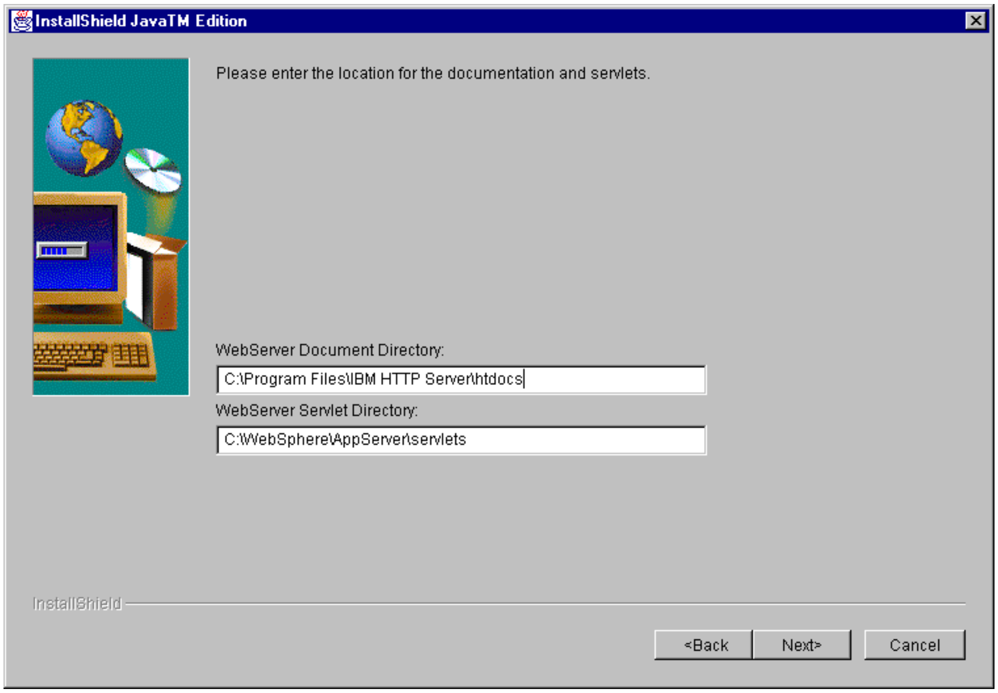

InstallShield JavaTM Edition

Please enter the location for the documentation and servlets.

WebServer Document Directory:

C:\Program Files\IBM HTTP Server\htdocs

WebServer Servlet Directory:

C:\WebSphere\AppServer\servlets

Figure 350. Request to Supply Web Server and Servlet Configuration Details

In this case, you are again asked to supply two directory locations to the install process:

1. The first piece of information the installation process asks you for is the same as in Point 1 on page 415: the Web server document directory. The installation process requires this information as it places in this directory some files that are used when the Common Configuration utility is launched. We supplied the default document directory for IBM HTTP Server as that was the Web server we had preinstalled.
2. The second piece of information is different from Point 2 on page 416. Now, in fact, the installation process asks you where the Web server servlet directory is. This directory is where the Common Configuration servlets will be placed during the installation process. We supplied the default servlet directory for WebSphere Application Server on Windows NT in the second field.

12.1.3 Remaining Installation Steps

After clicking **Next** we saw a window asking if we wanted the installation program to replace other programs already installed. We selected **No** in this case, since we had not installed any programs on our system yet:

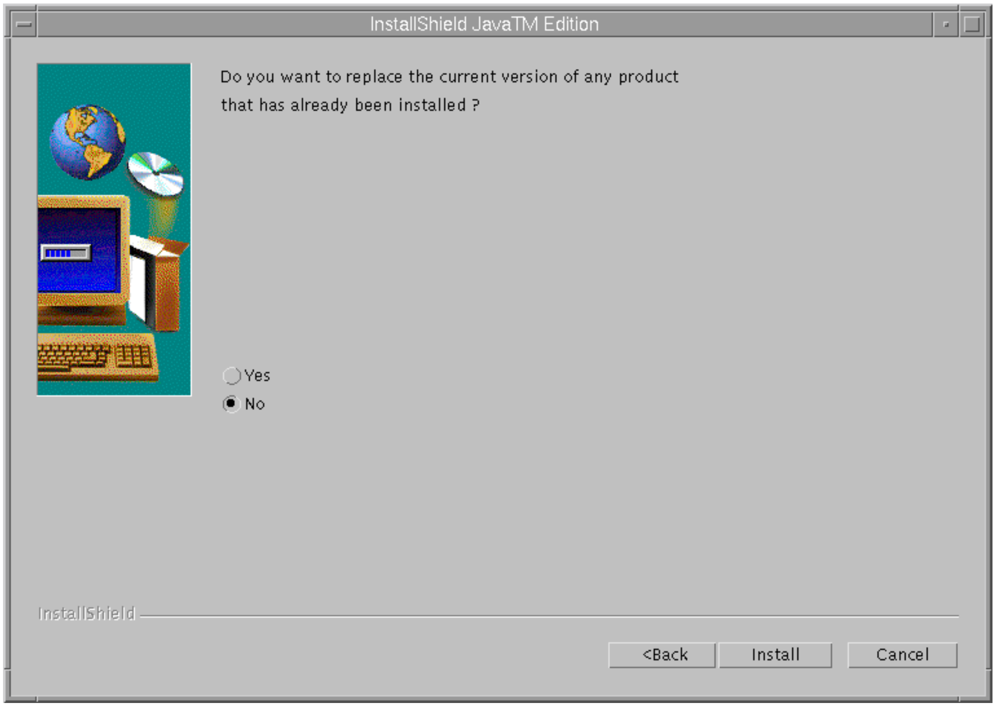

Do you want to replace the current version of any product that has already been installed ?

○ Yes

● No

Figure 351. Choose to Replace Version

After clicking **Install**, the Common Configuration servlets were installed. When WebSphere Application Server was not already installed on our system, we first saw the following screen displayed, showing the installation progress for WebSphere Application Server:

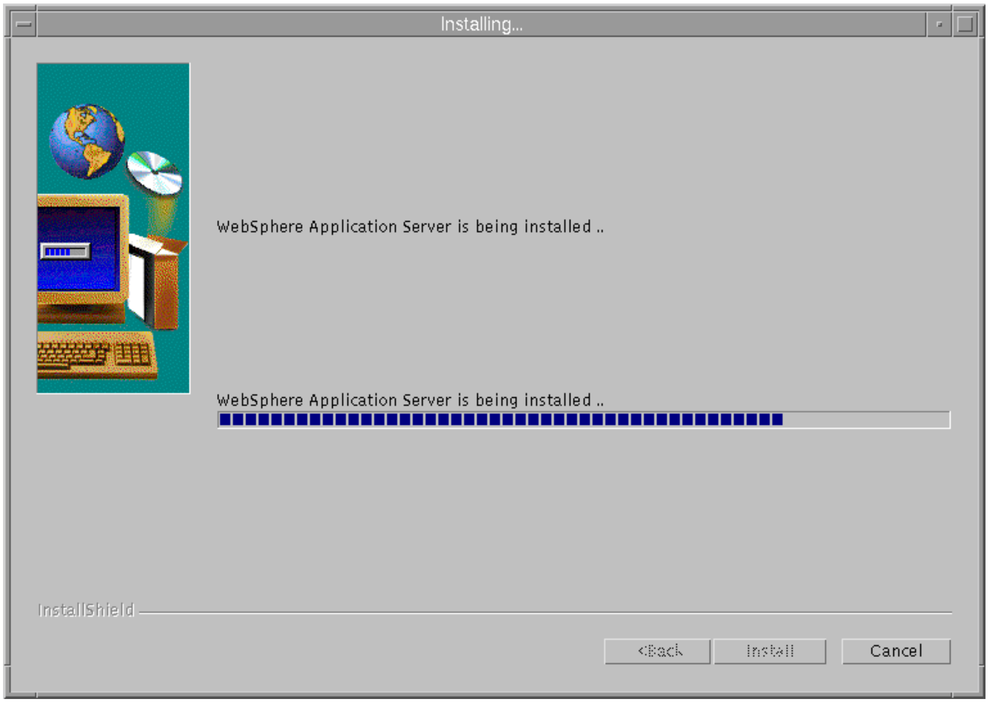

WebSphere Application Server is being installed ..

WebSphere Application Server is being installed ..

Figure 352. WebSphere Application Server Installation Progress Indicator

The following screen is always displayed during the installation of the Common Configuration utility, showing the installation progress for the Common Configuration utility itself:

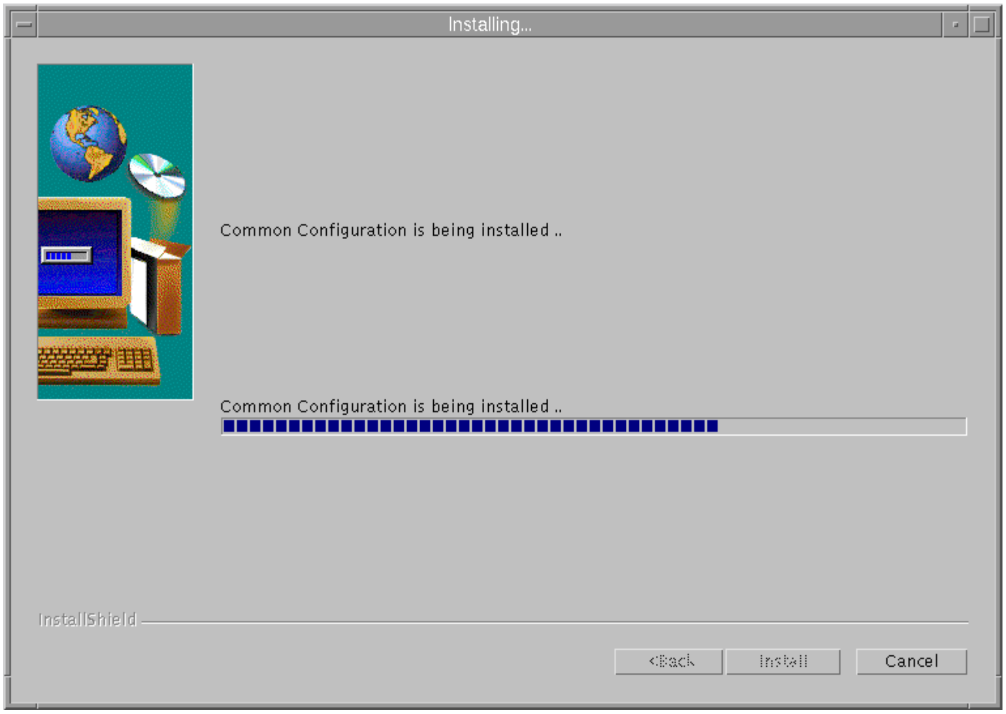

Common Configuration is being installed ..

Common Configuration is being installed ..

Figure 353. Common Configuration Installation Progress Indicator

Then, you will be informed that the installation completed successfully:

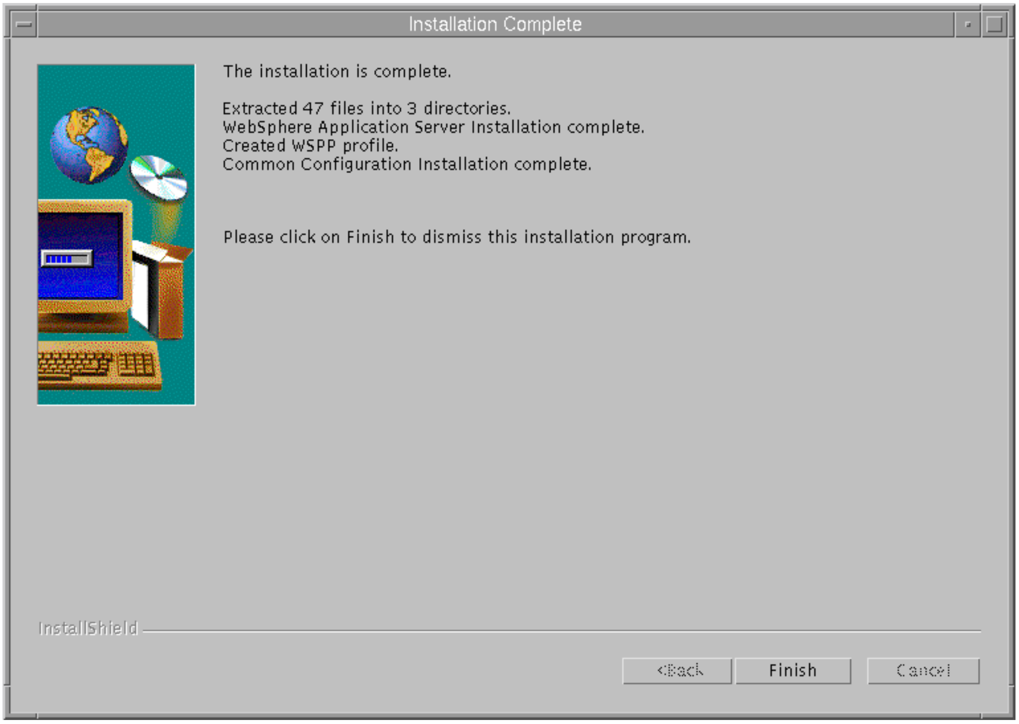

The installation is complete.

Extracted 47 files into 3 directories.
WebSphere Application Server Installation complete.
Created WSPP profile.
Common Configuration Installation complete.

Please click on Finish to dismiss this installation program.

Figure 354. Successful Completion Indication

We clicked on **Finish** to dismiss the installation process.

At this point, we could verify that the installation of the Common Configuration utility transparently installs WebSphere Application Server on the Common Configuration server. For example, on our AIX system, we saw that the filesets shown in the following figure were the only ones that were installed on our system as a result of the Common Configuration installation:

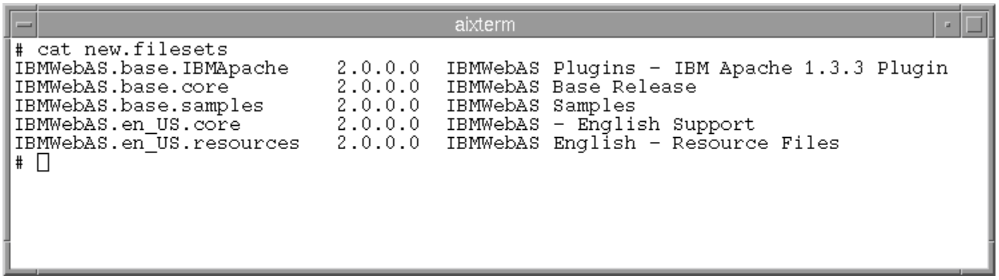

```
# cat new.filesets
IBMWebAS.base.IBMApache    2.0.0.0  IBMWebAS Plugins - IBM Apache 1.3.3 Plugin
IBMWebAS.base.core         2.0.0.0  IBMWebAS Base Release
IBMWebAS.base.samples      2.0.0.0  IBMWebAS Samples
IBMWebAS.en_US.core        2.0.0.0  IBMWebAS - English Support
IBMWebAS.en_US.resources   2.0.0.0  IBMWebAS English - Resource Files
# 
```

Figure 355. Installing the Common Configuration Utility Installs WebSphere Application Server

12.2 Preparing for the Common Configuration Utility

Communication between the Common Configuration server machine and the WebSphere Performance Pack Version 2 machine that the Common Configuration server is configuring is done via Remote Method Invocation (RMI).

A machine where WebSphere Performance Pack has been installed, and that needs to be configured through the Common Configuration utility, is also known as *WebSphere Performance Pack server machine*. On a WebSphere Performance Pack server machine, the installation program creates a file that must be executed before you can use the Common Configuration to configure WebSphere Performance Pack on this machine. The file is called wsppserver on AIX and Solaris, and wsppserver.cmd on Windows NT. It is located in the rmi subdirectory of the installation directory that you instructed the installation program to use at installation time.

Notice that the wsppserver file must not be executed on the Common Configuration server machine, but on the WebSphere Performance Pack server machine that is going to be the target of the Common Configuration utility. In other words, the wsppserver script must be executed on the machine where one of the three components (AFS, ND, and WTE) of WebSphere Performance Pack has been installed, and needs to be configured (typically remotely) through the Common Configuration utility.

On AIX, we followed these steps to execute the wsppserver file:

```
cd /usr/lpp/WSPP/rmi
chmod +x wsppserver
./wsppserver
```

We confirmed the success of the wsppserver execution by looking at the wsppcc.log file that is also present in the same directory as wsppserver:

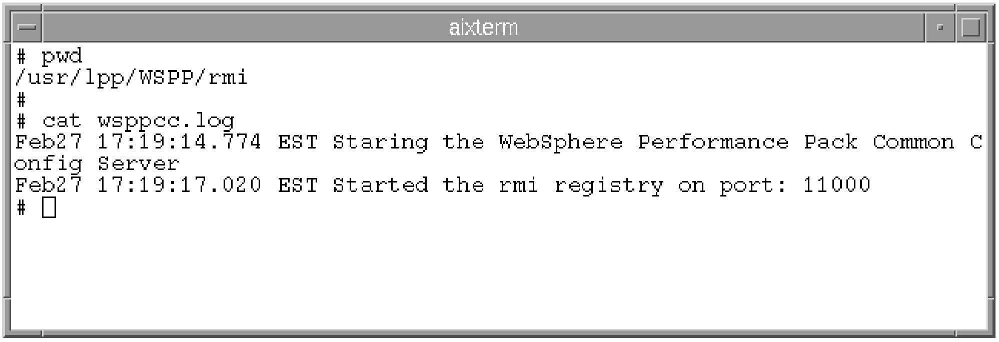

```
# pwd
/usr/lpp/WSPP/rmi
#
# cat wsppcc.log
Feb27 17:19:14.774 EST Staring the WebSphere Performance Pack Common C
onfig Server
Feb27 17:19:17.020 EST Started the rmi registry on port: 11000
# 
```

Figure 356. Confirming the Success of the wsppserver Execution

If you would like wsppserver to be run each time the WebSphere Performance Pack machine is rebooted, you could add it to an rc file on AIX or Solaris; on a Windows NT machine, add wsppserver to the Startup folder or registry.

To summarize, the steps that must be done before using the Common Configuration utility are:

1. Identify a Websphere Performance Pack server machine to be the target of the Common Configuration utility and on this machine execute the wsppserver script file. This file carries a .cmd extension on Windows NT.
2. Identify a configuration server machine and on this machine:

- Install the Common Configuration utility.
- Start the WebSphere Application Server-compliant Web server if it is not already running.

The next step that you must perform is dependent upon which WebSphere Performance Pack component you plan to configure with the Common Configuration utility.

12.3 Launching the Common Configuration Utility

To launch the Common Configuration utility from a Web browser, go to `http://common.configuration.server.name/wspp/startServer.html`, where common.configuration.server.name is the host name of the Common Configuration server machine. Following is the Common Configuration home page:

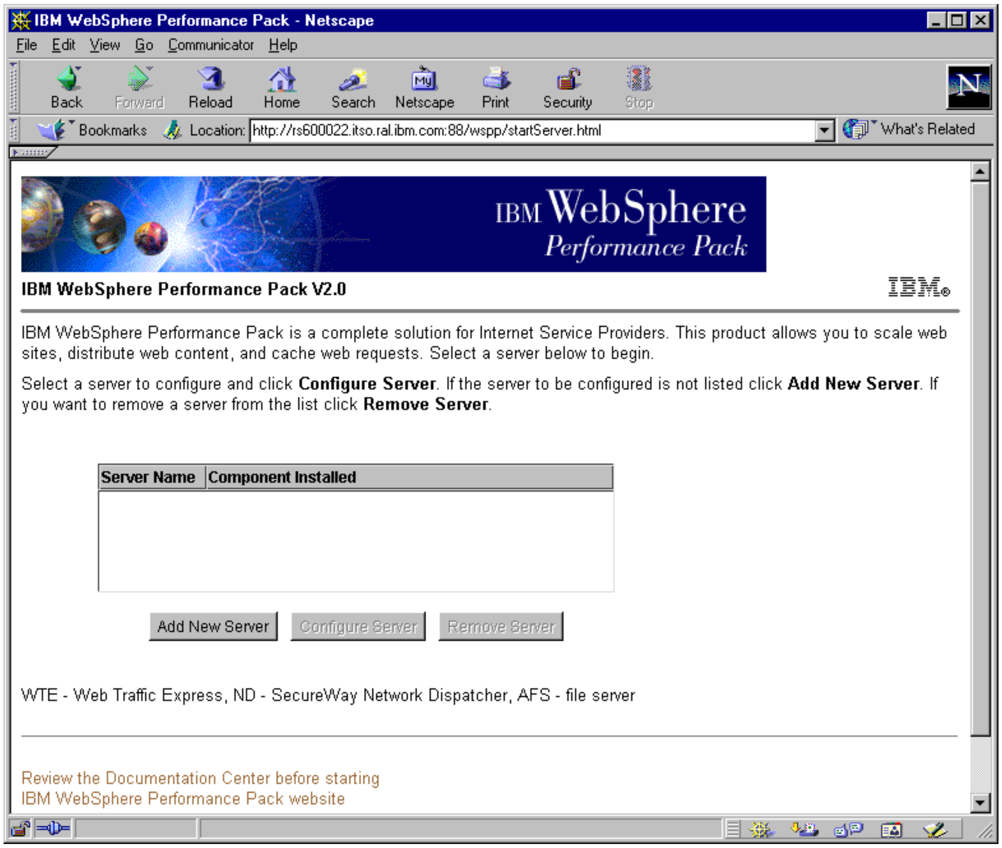
IBM WebSphere Performance Pack - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://rs600022.itso.ral.ibm.com:88/wspp/startServer.html
What's Related
IBM WebSphere Performance Pack

IBM WebSphere Performance Pack V2.0

IBM WebSphere Performance Pack is a complete solution for Internet Service Providers. This product allows you to scale web sites, distribute web content, and cache web requests. Select a server below to begin.

Select a server to configure and click **Configure Server**. If the server to be configured is not listed click **Add New Server**. If you want to remove a server from the list click **Remove Server**.

Server Name	Component Installed

WTE - Web Traffic Express, ND - SecureWay Network Dispatcher, AFS - file server

Review the Documentation Center before starting
IBM WebSphere Performance Pack website

Figure 357. Common Configuration Main Page

From the server selection list on the page, you can either click **Add New Server** or select a server from the list (if the list is not empty) and then click **Configure Server** or **Remove Server**.

The first time we accessed the Common Configuration utility, we selected **Add New Server** and in the subsequent Server Addition Request flield, entered the name of our WebSphere Performance Pack server machine:

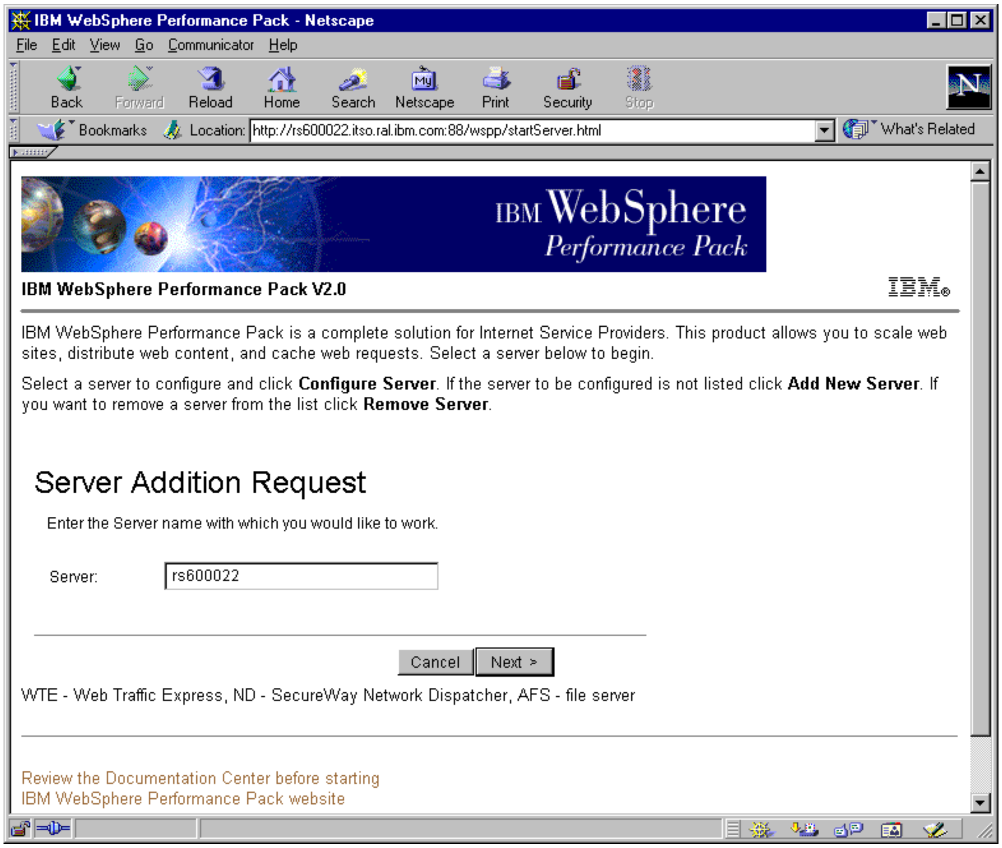
IBM WebSphere Performance Pack - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://rs600022.itso.ral.ibm.com:88/wspp/startServer.html
What's Related
IBM WebSphere Performance Pack
IBM

IBM WebSphere Performance Pack V2.0

IBM WebSphere Performance Pack is a complete solution for Internet Service Providers. This product allows you to scale web sites, distribute web content, and cache web requests. Select a server below to begin.

Select a server to configure and click **Configure Server**. If the server to be configured is not listed click **Add New Server**. If you want to remove a server from the list click **Remove Server**.

Server Addition Request

Enter the Server name with which you would like to work.

Server: rs600022

Cancel Next >

WTE - Web Traffic Express, ND - SecureWay Network Dispatcher, AFS - file server

Review the Documentation Center before starting
IBM WebSphere Performance Pack website

Figure 358. Server Addition Request

As you can see from the figure above, in this example, the Common Configuration server machine (see the host name in the URL) is the same as the WebSphere Performance Pack server machine (see the host name in the Server Addition Request field). However, there is no need for this to happen. The WebSphere Performance Pack server machines that a Common Configuration server can administer can be either remote or local.

After clicking the **Next** button, we saw a refreshed version of the Common Configuration main page, showing the name of the WebSphere Performance Pack server machine. The Common Configuration utility can immediately detect the WebSphere Performance Pack components that are installed on a WebSphere Performance Pack server and that, therefore, are subject to be configured through the Common Configuration utility. The following figure shows how the Common Configuration utility immediately detected the presence of WTE in the WebSphere Performance Pack server machine whose host name has been added in the Server Addition Request field:

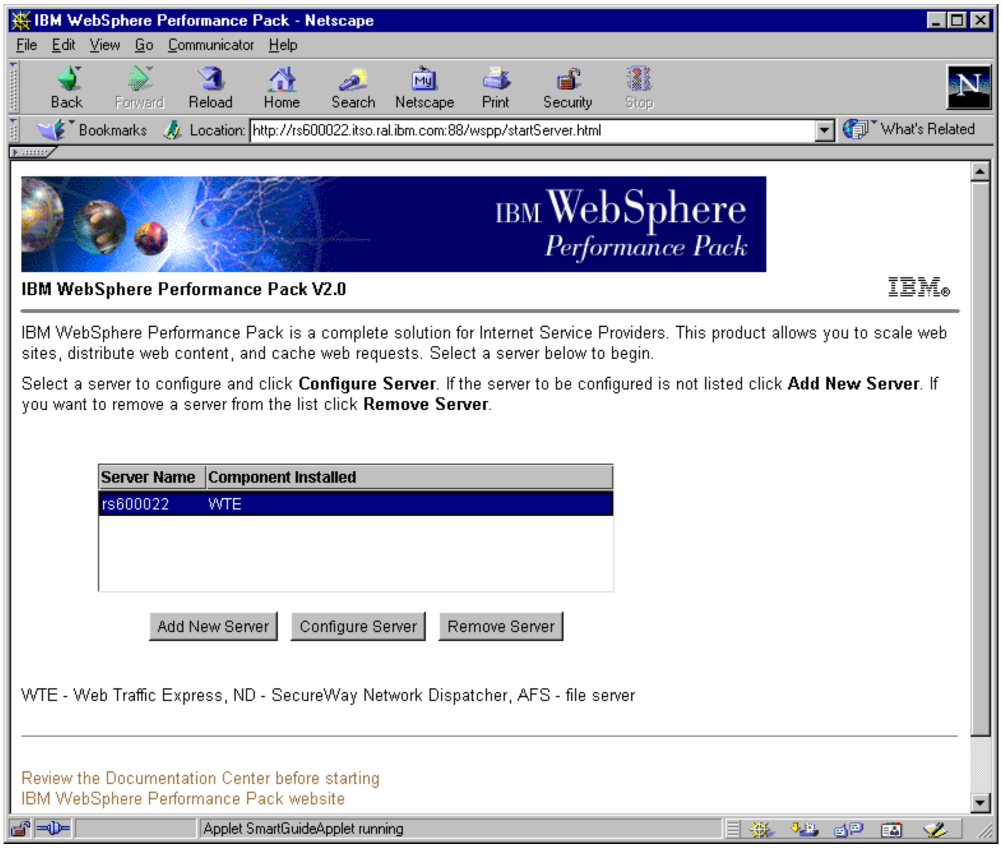
IBM WebSphere Performance Pack - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://rs600022.itso.ral.ibm.com:88/wspp/startServer.html
What's Related
IBM WebSphere Performance Pack

IBM WebSphere Performance Pack V2.0

IBM WebSphere Performance Pack is a complete solution for Internet Service Providers. This product allows you to scale web sites, distribute web content, and cache web requests. Select a server below to begin.

Select a server to configure and click **Configure Server**. If the server to be configured is not listed click **Add New Server**. If you want to remove a server from the list click **Remove Server**.

Server Name	Component Installed
rs600022	WTE

Add New Server | Configure Server | Remove Server

WTE - Web Traffic Express, ND - SecureWay Network Dispatcher, AFS - file server

Review the Documentation Center before starting
IBM WebSphere Performance Pack website

Applet SmartGuideApplet running

Figure 359. One Performance Pack Server Added to the Common Configuration Page

Multiple WebSphere Performance Pack server machines (the local machine as well as remote machines) can be added to the list. The Common Configuration server can automatically detect the WebSphere Performance Pack components that are installed on all the WebSphere Performance Pack server machines, so that these can be configured through the Common Configuration utility. The following figure shows an example of this:

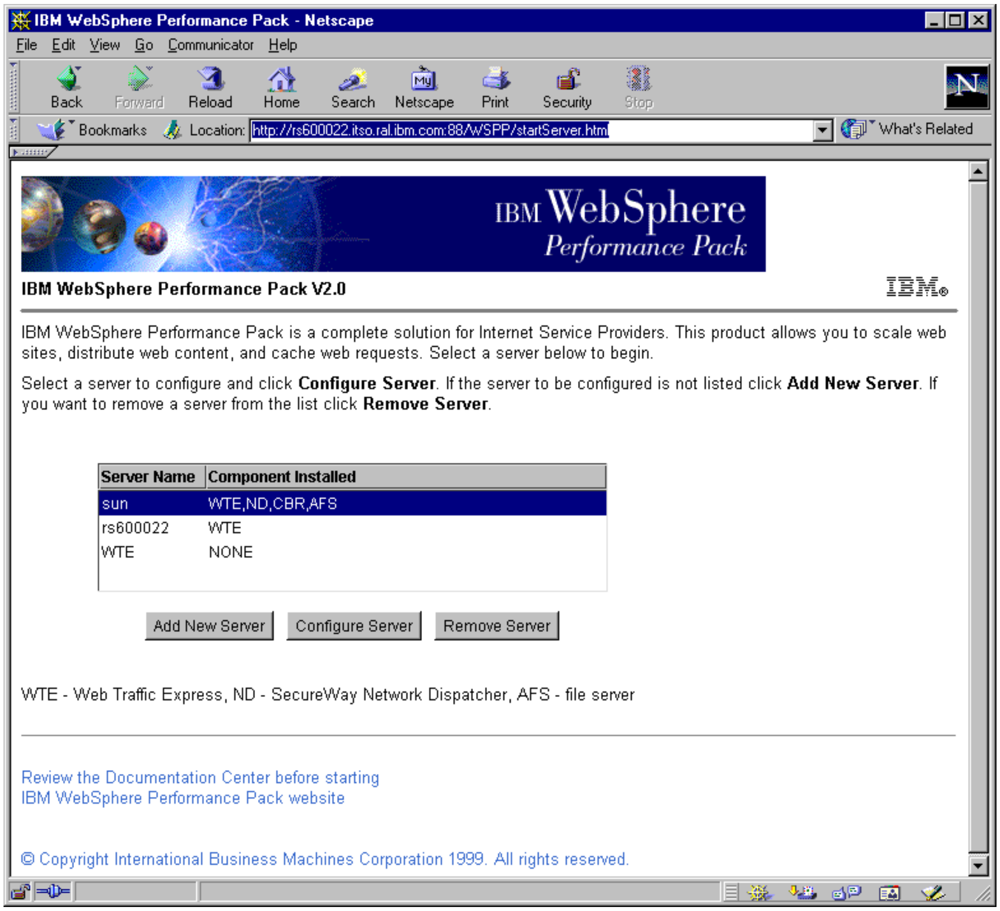
IBM WebSphere Performance Pack - Netscape
File Edit View Go Communicator Help
Back Forward Reload Home Search Netscape Print Security Stop
Bookmarks Location: http://rs600022.itso.ral.ibm.com:88/WSPP/startServer.html
What's Related
IBM WebSphere Performance Pack
IBM

IBM WebSphere Performance Pack V2.0

IBM WebSphere Performance Pack is a complete solution for Internet Service Providers. This product allows you to scale web sites, distribute web content, and cache web requests. Select a server below to begin.

Select a server to configure and click **Configure Server**. If the server to be configured is not listed click **Add New Server**. If you want to remove a server from the list click **Remove Server**.

Server Name	Component Installed
sun	WTE,ND,CBR,AFS
rs600022	WTE
WTE	NONE

WTE - Web Traffic Express, ND - SecureWay Network Dispatcher, AFS - file server

Review the Documentation Center before starting
IBM WebSphere Performance Pack website

© Copyright International Business Machines Corporation 1999. All rights reserved.

Figure 360. Common Configuration with Multiple WebSphere Performance Pack Servers

With a server selected, click **Configure Server** to begin the configuration. Once activated, the Common Configuration does not differ too much from the graphical user interface (GUI) tools that can be used to configure AFS, ND and WTE respectively. The advantage of the Common Configuration utility is that it allows centralized, remote, and secure administration of entire WebSphere Performance Pack sites. Only authenticated users can administer a WebSphere Performance Pack server through a Common Configuration server.

Appendix A. Special Notices

This publication is intended to help you to plan for, install, configure, use, tune and troubleshoot IBM AFS Enterprise File System, the File Sharing component of IBM WebSphere Performance Pack.The information in this publication is not intended as the specification of any programming interfaces that are provided by WebSphere Performance Pack Version 2. See the PUBLICATIONS section of the IBM Programming Announcement for WebSphere Performance Pack Version 2 for more information about what publications are considered to be product documentation.

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent program that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program or service.

Information in this book was developed in conjunction with use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to the IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact IBM Corporation, Dept. 600A, Mail Drop 1329, Somers, NY 10589 USA.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The information about non-IBM ("vendor") products in this manual has been supplied by the vendor and IBM assumes no responsibility for its accuracy or completeness. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Any pointers in this publication to external Web sites are provided for convenience only and do not in any manner serve as an endorsement of these Web sites.

Any performance data contained in this document was determined in a controlled environment, and therefore, the results that may be obtained in other operating environments may vary significantly. Users of this document should verify the applicable data for their specific environment.

Reference to PTF numbers that have not been released through the normal distribution process does not imply general availability. The purpose of including these reference numbers is to alert IBM customers to specific information relative to the implementation of the PTF when it becomes available to each customer according to the normal IBM PTF distribution process.

The following terms are trademarks of the International Business Machines Corporation in the United States and/or other countries:

ADSTAR
AS/400
eNetwork
Netfinity
OpenEdition
Parallel Sysplex
RISC System/6000
S/390
SP
ThinkPad
VisualAge
WebSphere

AIX
Deep Blue
IBM
NetView
OS/390
RACF
RS/6000
SecureWay
System/390
TXSeries
VTAM

The following terms are trademarks of other companies:

C-bus is a trademark of Corollary, Inc. in the United States and/or other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and/or other countries.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States and/or other countries.

PC Direct is a trademark of Ziff Communications Company in the United States and/or other countries and is used by IBM Corporation under license.

ActionMedia, LANDesk, MMX, Pentium and ProShare are trademarks of Intel Corporation in the United States and/or other countries.

UNIX is a registered trademark in the United States and/or other countries licensed exclusively through X/Open Company Limited.

SET and the SET logo are trademarks owned by SET Secure Electronic Transaction LLC.

Other company, product, and service names may be trademarks or service marks of others.

Appendix B. Related Publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

B.1 International Technical Support Organization Publications

For information on ordering these ITSO publications see "How to Get ITSO Redbooks" on page 431.

- *IBM WebSphere Performance Pack: Load Balancing with IBM SecureWay Network Dispatcher,* SG24-5858
- *IBM WebSphere Performance Pack: Caching and Filtering with IBM Web Traffic Express*, SG24-5859
- *IBM WebSphere Performance Pack Usage and Administration*, SG24-5233
- *Network Computing Framework Component Guide*, SG24-2119

B.2 Redbooks on CD-ROMs

Redbooks are also available on the following CD-ROMs. Click the CD-ROMs button at `http://www.redbooks.ibm.com/` for information about all the CD-ROMs offered, updates and formats.

CD-ROM Title	**Collection Kit Number**
System/390 Redbooks Collection	SK2T-2177
Networking and Systems Management Redbooks Collection	SK2T-6022
Transaction Processing and Data Management Redbooks Collection	SK2T-8038
Lotus Redbooks Collection	SK2T-8039
Tivoli Redbooks Collection	SK2T-8044
AS/400 Redbooks Collection	SK2T-2849
Netfinity Hardware and Software Redbooks Collection	SK2T-8046
RS/6000 Redbooks Collection (BkMgr Format)	SK2T-8040
RS/6000 Redbooks Collection (PDF Format)	SK2T-8043
Application Development Redbooks Collection	SK2T-8037
IBM Enterprise Storage and Systems Management Solutions	SK3T-3694

B.3 Other Publications

These publications are also relevant as further information sources:

- *AFS Installation Guide* (shipped with the product or downloadable from `http://www.transarc.com/Support`)
- *SecureWay Network Dispatcher User's Guide Version 2.1 for Solaris, Windows NT and AIX*, GC31-8496

B.4 Referenced Web Sites

- `http://www.redbooks.ibm.com`
- `http://www.javasoft.com`
- `http://www.software.ibm.com/webservers/appserv`

- `http://www.transarc.com`
- `http://www.transarc.com/Support`
- `http://www.ibm.com/java/jdk/download/index.html`
- `http://w3.itso.ibm.com`
- `http://w3.ibm.com`
- `http://www.elink.ibmlink.ibm.com/pbl/pbl`

How to Get ITSO Redbooks

This section explains how both customers and IBM employees can find out about ITSO redbooks, redpieces, and CD-ROMs. A form for ordering books and CD-ROMs by fax or e-mail is also provided.

- **Redbooks Web Site** `http://www.redbooks.ibm.com/`

  Search for, view, download, or order hardcopy/CD-ROM redbooks from the redbooks Web site. Also read redpieces and download additional materials (code samples or diskette/CD-ROM images) from this redbooks site.

  Redpieces are redbooks in progress; not all redbooks become redpieces and sometimes just a few chapters will be published this way. The intent is to get the information out much quicker than the formal publishing process allows.

- **E-mail Orders**

  Send orders by e-mail including information from the redbooks fax order form to:

	e-mail address
In United States	usib6fpl@ibmmail.com
Outside North America	Contact information is in the “How to Order” section at this site: `http://www.elink.ibmlink.ibm.com/pbl/pbl/`

- **Telephone Orders**

United States (toll free)	1-800-879-2755
Canada (toll free)	1-800-IBM-4YOU
Outside North America	Country coordinator phone number is in the “How to Order” section at this site: `http://www.elink.ibmlink.ibm.com/pbl/pbl/`

- **Fax Orders**

United States (toll free)	1-800-445-9269
Canada	1-403-267-4455
Outside North America	Fax phone number is in the “How to Order” section at this site: `http://www.elink.ibmlink.ibm.com/pbl/pbl/`

This information was current at the time of publication, but is continually subject to change. The latest information may be found at the redbooks Web site.

IBM Intranet for Employees

IBM employees may register for information on workshops, residencies, and redbooks by accessing the IBM Intranet Web site at `http://w3.itso.ibm.com/` and clicking the ITSO Mailing List button. Look in the Materials repository for workshops, presentations, papers, and Web pages developed and written by the ITSO technical professionals; click the Additional Materials button. Employees may access MyNews at `http://w3.ibm.com/` for redbook, residency, and workshop announcements.

IBM Redbook Fax Order Form

Please send me the following:

Title	Order Number	Quantity

First name Last name

Company

Address

City Postal code Country

Telephone number Telefax number VAT number

□ Invoice to customer number

□ Credit card number

Credit card expiration date Card issued to Signature

We accept American Express, Diners, Eurocard, Master Card, and Visa. Payment by credit card not available in all countries. Signature mandatory for credit card payment.

List of Abbreviations

ACL	access control list
AFS	Enterprise File System
AIX	advanced interactive executive
ATM	Asynchronous Transfer Mode
API	application programming interface
APPN	Advanced Peer-to-Peer Network
ARP	Address Protocol Request
BDM	Binary Distribution Machine
BOS	Basic OverSeer
CB	Component Broker
CBR	Content Based Routing
CORBA	Common Object Request Broker Architecture
DMZ	demilitarized zone
DNS	Domain Name System
DWTE	Distributed Web Traffic Express
EJB	Enterprise JavaBeans
EJS	Enterprise Java Services
FDDI	Fiber Distributed Data Interface
FTP	File Transfer Protocol
GC	garbage collector
GEM	Global Enterprise Manager
GUI	graphical user interface
HACMP	High Availability Cluster Multiprocessing
HP-UX	Hewlett-Packard UNIX
HTTP	Hypertext Transfer Protocol
IBM	International Business Machines Corporation
IP	Internet Protocol
ISP	Internet Service Provider
ISS	Interactive Session Support
ITSO	International Technical Support Organization
JDK	Java Development Kit
JFS	Journaled File System
JRE	Java Runtime Environment
JSP	JavaServer Pages
JVM	Java Virtual Machine
LAN	local area network
LANA	LAN adapter
LDAP	Lightweight Directory Access Protocol
MAC	Media Access Control
MB	megabytes
Mbps	megabits per second
MBps	megabytes per second
MHz	megahertz
MIB	management information base
MS-DOS	Microsoft disk operating system
MVS	multiple virtual storage
NAP	network access points
NCSA	National Center for Supercomputing Applications
ND	Network Dispatcher
NFS	Network File System
NNTP	NetNews transfer protocol
NTP	Network Time Protocol
NTFS	Windows NT File System
OS/2	Operating System/2
PICS	Platform for Internet Content Selection
POP	points of presence
POP3	Post Office Protocol 3
RCA	Remote Cache Access
RMI	Remote Method Invocation
RPC	remote procedure call
SCM	System Control Machine
SDA	Server Directed Affinity
SSL	Secure Sockets Layer
SMTP	Simple Mail Transfer Protocol
SRC	System Resource Control
TCP/IP	Transmission Control Protocol/Internet Protocol
TEC	Tivoli Enterprise Console
UDP	User Datagram Protocol

UFS	UNIX file system
URL	Universal Resource Locator, Uniform Resource Locator
VL	Volume Location
VLDB	Volume Location Database
VM	virtual memory
WAN	wide area network
WAND	Wide Area Network Dispatcher
WTE	Web Traffic Express
WLM	workload manager
WWW	World Wide Web
XML	Extensible Markup Language

Index

ITSO Redbook Evaluation

IBM WebSphere Performance Pack: Web Content Management with IBM AFS Enterprise File System
SG24-5857-00

Your feedback is very important to help us maintain the quality of ITSO redbooks. **Please complete this questionnaire and return it using one of the following methods:**

- Use the online evaluation form found at `http://www.redbooks.ibm.com/`
- Fax this form to: USA International Access Code + 1 914 432 8264
- Send your comments in an Internet note to `redbook@us.ibm.com`

Which of the following best describes you?
_ Customer _ Business Partner _ Solution Developer _ IBM employee
_ None of the above

Please rate your overall satisfaction with this book using the scale:
(1 = very good, 2 = good, 3 = average, 4 = poor, 5 = very poor)

Overall Satisfaction __________

Please answer the following questions:

Was this redbook published in time for your needs? Yes___ No___

If no, please explain:

What other redbooks would you like to see published?

Comments/Suggestions: (THANK YOU FOR YOUR FEEDBACK!)

SG24-5857-00

Printed in the U.S.A.

IBM WebSphere Performance Pack: Web Content Management with IBM AFS Enterprise File System

SG24-5857-00