Advanced Server for UNIX

Concepts and Planning Guide

September 2002

Product Version: Advanced Server for UNIX Version 5.1B

Operating System and Version: Tru64 UNIX Version 5.1A or higher

This guide describes concepts related to planning and administering the Advanced Server for UNIX (ASU) software.

Hewlett-Packard Company
Palo Alto, California

Contents

3 Domain User Accounts and Groups

4 Disk Shares

Index

Figures

Tables

About This Guide

Concepts and Planning explains the concepts related to planning and administering the Advanced Server for UNIX (ASU) software.

Audience

This guide is intended for anyone who is responsible for planning, installing, configuring, and administering the ASU software.

Organization

The guide is organized as follows:

Chapter 1	Describes the ASU software.
Chapter 2	Describes the ASU domain environment.
Chapter 3	Describes domain user accounts and groups.
Chapter 4	Describes how the ASU server shares Tru64 UNIX based file systems.
Chapter 5	Describes how the ASU server shares Tru64 UNIX based printers.
Chapter 6	Describes how to monitor the ASU server.

Related Documentation

The following documents provide more information about the ASU software:

- ASU *Installation and Administration Guide* - Describes how to install, configure, and administer the ASU software.
- ASU *Release Notes* - Describes the latest information about the ASU software that might not be documented elsewhere.

Reader's Comments

HP welcomes any comments and suggestions you have on this and other Tru64 UNIX manuals.

You can send your comments in the following ways:

- Fax: 603-884-0120 Attn: UBPG Publications, ZKO3-3/Y32

- Internet electronic mail: `readers_comment@zk3.dec.com`

  A Reader's Comment form is located on your system in the following location:

  `/usr/doc/readers_comment.txt`

Please include the following information along with your comments:

- The full title of the manual and the order number. (The order number appears on the title page of printed and PDF versions of a manual.)
- The section numbers and page numbers of the information on which you are commenting.
- The version of Tru64 UNIX that you are using.
- If known, the type of processor that is running the Tru64 UNIX software.

The Tru64 UNIX Publications group cannot respond to system problems or technical support inquiries. Please address technical questions to your local system vendor or to the appropriate HP technical support office. Information provided with the software media explains how to send problem reports to HP.

Conventions

The following conventions are used in this guide:

%	
$	A percent sign represents the C shell system prompt. A dollar sign represents the system prompt for the Bourne, Korn, and POSIX shells.
#	A number sign represents the superuser prompt.
file	Italic (slanted) type indicates variable values, placeholders, and function argument names.
[|]	
{ | }	In syntax definitions, brackets indicate items that are optional and braces indicate items that are required. Vertical bars separating items inside brackets or braces indicate that you choose one item from among those listed.
. . .	In syntax definitions, a horizontal ellipsis indicates that the preceding item can be repeated one or more times.

`cat`(1)	A cross-reference to a reference page includes the appropriate section number in parentheses. For example, `cat`(1) indicates that you can find information on the `cat` command in Section 1 of the reference pages.
Return	In an example, a key name enclosed in a box indicates that you press that key.
Ctrl/*x*	This symbol indicates that you hold down the first named key while pressing the key or mouse button that follows the slash. In examples, this key combination is enclosed in a box (for example, Ctrl/C).

1

ASU Overview

The Advanced Server for UNIX (ASU) software is a Tru64 UNIX layered application that allows you to share UNIX based file systems and printers with Windows users as shares. Windows users connect to shares without modification to their software. Once connected, the file system or printer associated with a share appears as a transparent extension to a Windows user's local computing environment.

This chapter describes:

- The ASU server
- The ASU server process model
- The ASU server architecture
- ASU server administrative interfaces

1.1 ASU Server Overview

Each system on which you install, configure, and run the ASU software becomes an ASU server. An ASU server provides Windows users with access to UNIX file systems as disk shares and printers as printer shares.

The ASU server provides flexible disk and printer share security by enforcing either the Windows NT security model exclusively, or Windows NT combined with Tru64 UNIX security models. By default, the ASU server implements the Windows NT combined with Tru64 UNIX security model. This security model requires that a Window user have the following user accounts:

- A domain user account that you create. The ASU server uses this account to enforce Windows NT security.
- A Tru64 UNIX user account that is automatically created when you create the domain user account. The Tru64 UNIX operating system software uses this account to enforce Tru64 UNIX security policies.

An ASU server interoperates with and uses features of the Tru64 UNIX operating system, including preemptive multitasking, symmetric multiprocessing, and time-sharing. The way in which the ASU server interoperates with the Tru64 UNIX operating system software depends on the values assigned to value entries stored in a database called the ASU registry.

An ASU server participates in a Windows domain as a Primary Domain Controller (PDC), Backup Domain Controller (BDC), or member server to provide Windows NT Advanced Server Version 4.0 services described in Table 1–1 to Windows users, clients, and servers.

Table 1–1: ASU Services

Service	Description
`Alerter`	Used by the ASU server and other ASU services to notify selected users and computers of administrative alerts that occur on this computer.
`Browser`	Maintains an up-to-date list of controllers in a domain, and provides the list when requested.
`EventLog`	Records system, security, and application events in the event logs, and enables remote access to those logs.
`NetLogon`	Verifies the domain user name and password of each person who attempts to log in to the domain or to the ASU server.
`Netrun`	Lets users run UNIX system applications on a server from their workstation.
`Replicator`	Replicates directories and the files in those directories to other workstations.
`Server`	Provides file, print, named pipe sharing, and support for remote procedure calls.
`TimeSource`	Identifies a controller as the time source for a domain. Other controllers synchronize their clocks with the time server.

1.2 The ASU Server Process Model

The ASU server and client workstations communicate by using the Server Message Block (SMB) protocol, the native file-sharing protocol in the Microsoft Windows and OS/2 operating systems.

The client sends an SMB request to the ASU server. The ASU server receives the SMBs, maps the requests to equivalent Tru64 UNIX system semantics, interprets the client's intent, and performs the Tru64 UNIX system function to satisfy the client's request. The network architecture layers, such as transport and hardware, ensure reliable SMB exchange.

The ASU server is a combination of the system processes that work in close cooperation. The `lmx.ctrl` process is the ASU master control process and must be running. The UNIX network listener service passes new ASU client connection requests or sessions to the `lmx.ctrl` process. The `lmx.ctrl` process accepts a new client session and distributes it to an `lmx.srv` process. The `lmx.srv` process actually services the needs of the client.

Figure 1–1 shows the relationship of the network listener service, the `lmx.ctrl` process, the `lmx.srv` process, and clients.

Figure 1–1: ASU Process Model

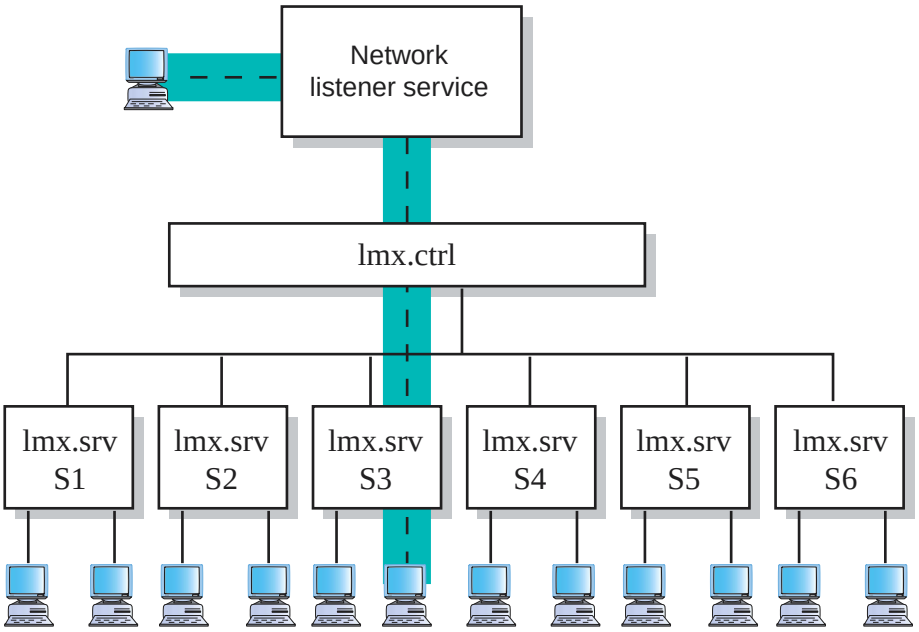

The number of `lmx.srv` processes that are created varies according to the number of client sessions. The ASU process model allows multiple client sessions to be serviced by a single `lmx.srv` process. The `lmx.ctrl` process determines if a new client session is serviced by an existing `lmx.srv` process or by a new `lmx.srv` process. This distribution of client sessions to `lmx.srv` processes enables more than one client request to be processed at a time.

1.3 The ASU Server Architecture

The ASU server sends SMBs on the network using the NetBIOS protocol on an Ethernet, FDDI, or Token Ring network adapter as supported by the Tru64 UNIX operating system software. The NetBIOS protocol is responsible for establishing logical names for workstations on the network, establishing sessions between workstations' logical names on the network, and supporting reliable data transfer between them.

The ASU software provides and uses the NetBIOS protocol over the following transports:

- NetBIOS over TCP/IP (`knbtcp`) is used over the system's installed TCP/IP transport software for local and wide area networking.
- NetBEUI transport is used exclusively for local area networking.

Figure 1–2 shows the ASU architecture. The ASU software provides the shaded components.

Figure 1–2: ASU Network Architecture

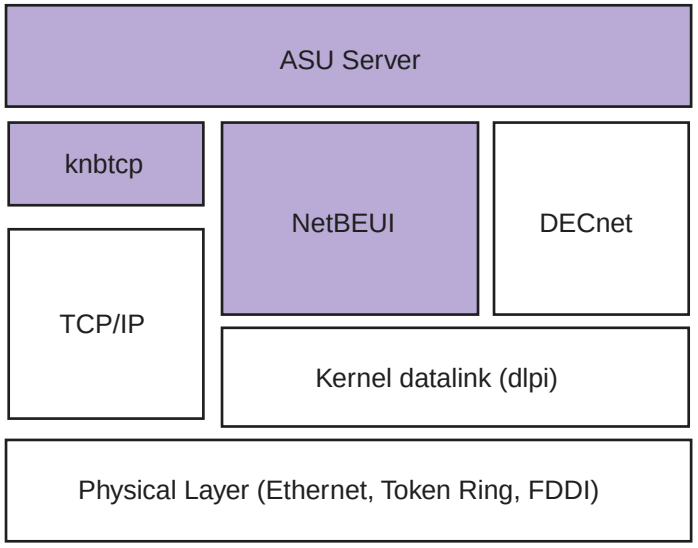

ZK-1785U-AI

1.4 ASU Server Administrative Interfaces

To administer the ASU software you can use the following:

- ASU commands
- `net` commands
- Tru64 UNIX commands and graphical user interfaces (GUIs)
- Windows GUIs

Note

If you plan to configure the ASU Server in a Windows 2000 domain, then you must administer the ASU server by using Windows 2000 interfaces.

1.4.1 ASU Commands

The ASU commands are Tru64 UNIX style commands that you can use to display information about, administer, and troubleshoot the ASU server and domain. You enter ASU commands in lowercase at the Tru64 UNIX command prompt on a system running the ASU software. See the ASU *Installation and Administration Guide* for more information on ASU commands.

1.4.2 net Commands

The `net` commands are Windows style commands that you can use to create shares, domain user accounts, and groups and to display information about and administer the ASU server, domain, shares, domain user accounts, and groups.

A `net` command begins with the word `net` followed by a keyword and options. You enter `net` commands in lowercase at the command prompt on a system running the ASU software in the following form:

```
# net keyword [/option]
```

See the ASU *Installation and Administration Guide* for more information on `net` commands.

1.4.3 Tru64 UNIX Commands and GUIs

The Tru64 UNIX user and file system commands and GUIs provide additional ASU related options that you can use to create and administer shares and domain user accounts. See *System Administration* for information on administering the ASU server using Tru64 UNIX commands and GUIs.

1.4.4 Windows GUIs

You can use the following Windows based GUIs to administer the ASU server and domain:

- The Server Manager creates, displays information about, and administers shares.
- The User Manager for Domains creates, displays information about, and administers domain user accounts and groups.
- The Policy Editor displays information about and administers the ASU registry.
- The Event Viewer displays ASU related application, security, and system events.

You can administer the ASU server by using the version of the Windows GUIs that are provided with a Windows NT Server Version 4.0 and a Windows 2000 Server. For a system running another type of Windows operating system software, you must install the version of the Windows GUIs that are provided with the ASU software. See the ASU *Installation and Administration Guide* for more information on installing the Windows based GUIs.

2

Advanced Server Domains

Each ASU server must belong to a Windows domain and have a role in that domain. A Windows domain is a grouping of Advanced Servers that share a common directory database in which all domain security, user account, and group information is stored. An Advanced Server can be an ASU server, Windows NT server, or a Windows 2000 server. Other Advanced Server and Windows NT documents may refer to the directory database as the Security Accounts Manager (SAM) database.

This chapter describes:

- Domain roles
- Common domain models
- Computer accounts
- Logging into a domain
- Managing domains

See the ASU *Installation and Administration Guide* for information on installing the ASU software and configuring the ASU server in a domain.

2.1 Domain Roles

The domain role of an Advanced Server determines which one maintains the directory database and which ones receive a copy of the directory database. Each Advanced Server must have one of the following roles in the domain:

- Primary domain controller (PDC). A PDC stores and maintains the directory database and authenticates domain user logon requests. A domain is created by configuring an Advanced Server as a PDC. There is only one PDC per domain.

  You cannot configure the ASU server as a PDC in a Windows 2000 domain.

- Backup domain controller (BDC). A BDC receives a copy of the directory database to authenticate domain user logon requests. This copy is synchronized periodically and automatically with the PDC. There can be many BDCs in a domain.

  You can configure the ASU server as a BDC in a Windows 2000 Server domain only if the Windows 2000 Server is configured for mixed mode.

- Member server. A member server maintains its own local user account database and does not receive a copy of the directory database, and therefore, does not process domain logon requests. A member server can participate in a domain to offer shared resources; however, a user must have a user account on the member server or in a domain in which the member server trusts. Trusts are described in Section 2.2. There can be many member servers in a domain.

  You can configure the ASU server as a member server in a Windows 2000 Server domain if the Windows 2000 Server is configured for mixed or native mode.

2.2 Common Domain Models

By planning and organizing domains, you can simplify network administration and ensure that users can connect to shares throughout the network. There are two types of domain models:

- A trusted domain model in which Advanced Servers are configured in multiple domains and the domains establish a trust relationship. A trust relationship allows a trusting domain to trust that a trusted domain authenticated its users and, therefore, allows the users of the trusted domain to access to its shares.
- A single domain model in which each Advanced Server is configured in one domain.

2.2.1 Trusted Domain Models

There are two types of trust relationships:

- A one-way trust relationship is when one domain trusts another domain.
- A two-way trust relationship is when two or more domains trust each other.

There are two common types of trusted domain models:

- The single master domain model, which is best suited for environments with a small number of users.
- The multiple master domain model, which is best suited for environments with a large number of users.

With each of these domain models, you create two types of domains:

- A master domain in which you create domain user accounts.
- Resource domains in which you create disk and printer shares.

You create a one-way trust from the resource domain to a master domain. Users log into a master domain and, because of the one-way trust, can access

all of the disk and printers shares in the resource domains for which you have granted them access.

2.2.1.1 Single Master Domain Model

Figure 2–1 shows the single domain master model where all user accounts are created in a master domain called Personnel. All shares are created in the resource domains called Marketing, Sales, and Production.

Figure 2–1: Single Domain Master Model

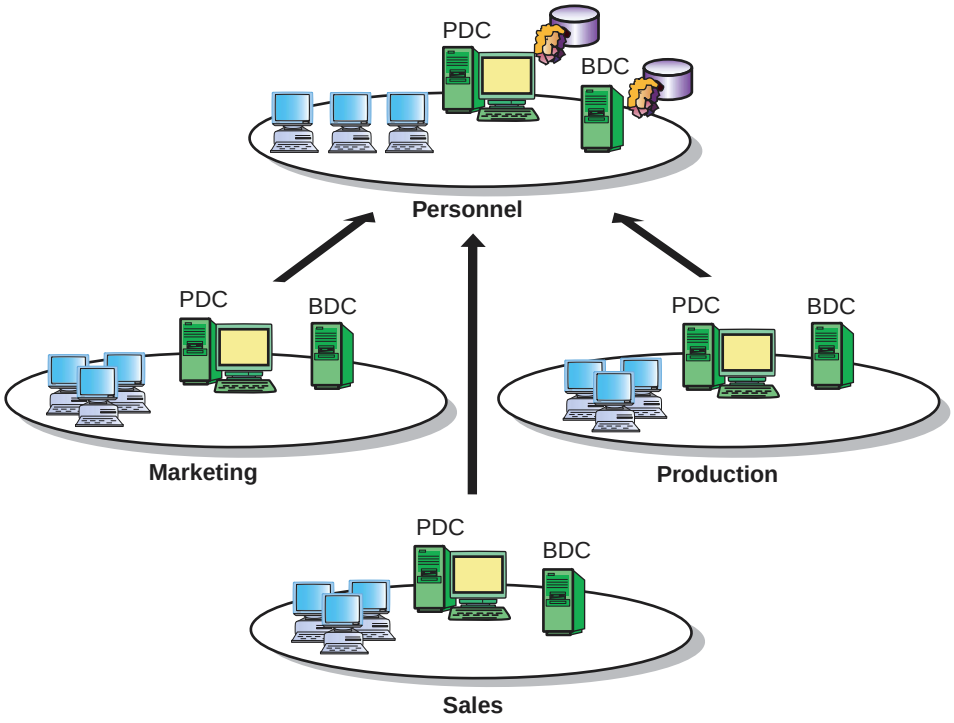

ZK-1782U-AI

Because every domain user account exists in the master domain and because each resource domain trusts the master domain, every user account can use shares in any of the resource domains for which you have granted them access.

Features of the single master domain model include:

- Centralized domain user account management. By creating all domain user accounts in a single domain and creating one-way-trust relationships between domains, you consolidate the administration of user accounts and avoid having to administer separate user account

databases. As a result, users need only one domain user account to use shares in any of the domains.

- Decentralized resource management or local system administration capability. Departmental domains can have their own administrators who manage the resources in the department.
- Resources can be grouped logically, corresponding to local domains.

2.2.1.2 Multiple Master Domain Model

The multiple master domain model is similar to the single domain model in that you create domain user accounts in a master domain and shares in a resource domain; however, as the name implies, you create more than one master domain.

Figure 2–2 shows the multiple domain master model where all domain user accounts are created in two master domains called A_Personnel and B_Personnel. All shares are created in the resource domains South_Sales, South_Marketing, North_Sales, and North_Marketing.

Figure 2–2: Multiple Domain Master Model

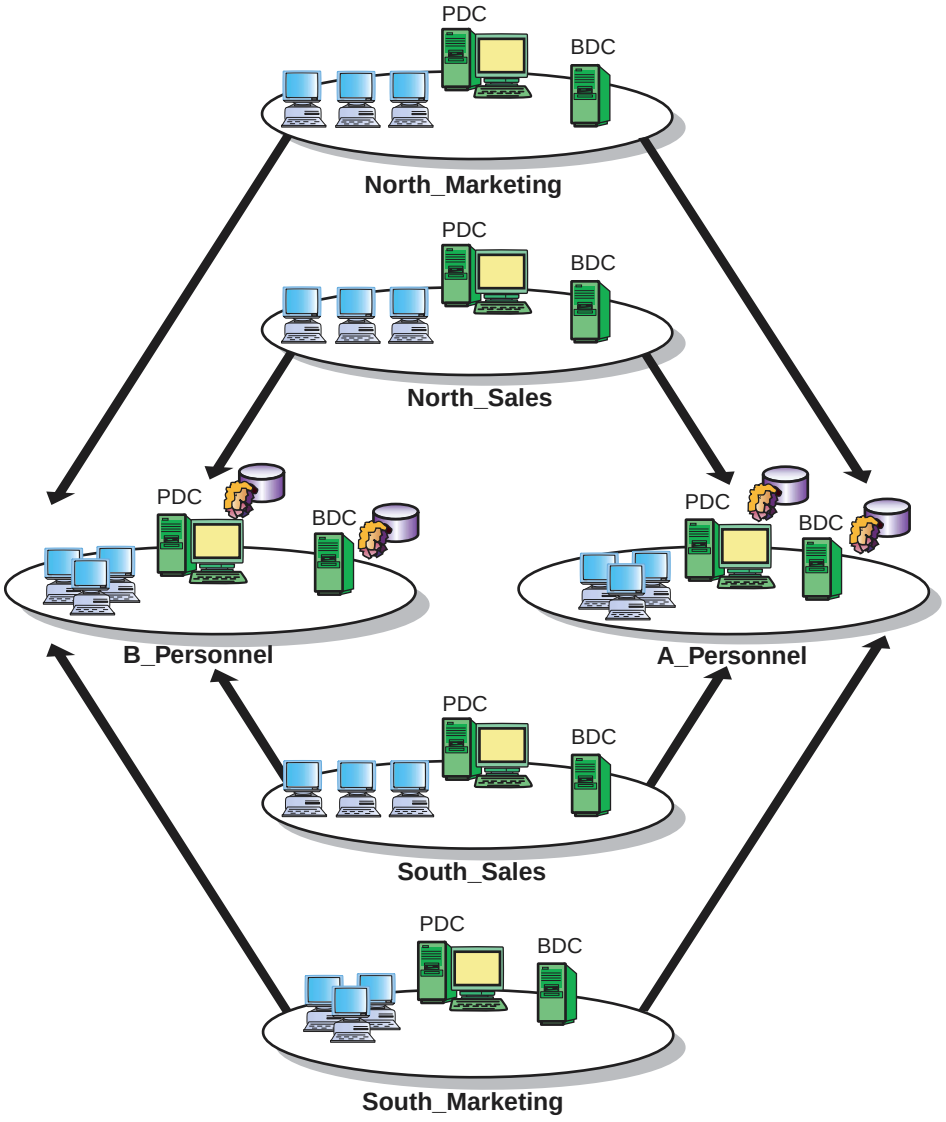

ZK-1783U-AI

Because every domain user account exists in one of the master domains and because each resource domain trusts each master domain, every user account can use shares in any of the resource domains for which you have granted them access.

The multiple master domain model incorporates all of the features of a single master domain model and has the following additional features:

- Is scaleable to networks with any number of users.
- Allows users to log on from anywhere in the network.
- Provides centralized or decentralized administration.

2.2.2 Single Domain Model

You can create a domain model in which domain user accounts and shares are created in the same domain. This means that you do not need to create any trust relationships because the domain user accounts and shares are in the same domain. If a domain user needs to access shares in other domains, you can create a domain user account with the same user name and password in those domains. If a user matches the user name and password, access is given regardless of which domain the user is logged on to.

The exception to this rule is when trust relationships are in operation. If a trust relationship is established between two domains, then the trusting domain is able to distinguish users in the trusted domain from users in the local domain, even if they have the same name and password.

Consider three domains: Athens, Berlin, and Cairo. Each domain has a user account, Peter, with the same password. The Athens domain trusts the Berlin domain but no other trust relationships are established. The following table describes the access that Peter has while logged on to each domain:

Logged On	Can Access	Cannot Access
Athens	Athens, Berlin, Cairo	
Berlin	Berlin, Cairo	Athens
Cairo	Athens, Berlin, Cairo	

It appears that the trust relationship between Athens and Berlin has restricted the access of Peter across the domains. In fact, the trust relationship has made access more controllable. The Athens domain now can distinguish the user Athens\Peter from the user Berlin\Peter and grant access accordingly.

2.3 Computer Accounts

A computer account is automatically created in the directory database for each BDC that joins a domain. Computer accounts are used to establish secure communications session.

A secure communications session is created when systems communicating in a connection are satisfied that the other system has identified itself correctly. Systems identify themselves by using their computer accounts.

When a secure communications session is established, communications can begin between the systems.

Examples of secure communication sessions include the following:

- A BDC creates a secure communications session with the PDC to receive a copy of the master directory database and updates.
- A PDC and BDCs in a trusting domain create a secure communications session with the PDC in trusted domains to pass user authentication information.
- A Windows NT Workstation or Windows NT Server system running as a member server in a domain creates a secure communications session to a PDC or BDC in a domain to pass user authentication information.

2.4 Logging Into a Domain

To log into a domain, a user must have a domain user account either in the domain or in a trusted domain. An administrator creates a domain user account by assigning a user name and password to an account, specifying the user's identification data, and defining the user's rights in the domain. The Advanced Server then assigns a unique security identifier (SID) to the new account. See Chapter 3 for more information on creating domain user accounts.

Users can logon in to a domain in the following ways:

- Interactive logon occurs when a user types information in the Logon Information dialog box displayed by the computer's operating system. In the Domain box, the user selects either the name of a domain or the name of the computer being used for logon, depending on where the user account is defined.
- Remote logon occurs when a user already is logged in to a domain and makes a network connection to another computer. For example, the user connects to another computer using the Map Network Drive dialog box.

When a user of a Windows computer attempts to log into a domain, the `NetLogon` service on their computer attempts to locate and establish a secure communications channel with the `NetLogon` service on a domain controller that can authenticate the user. The `NetLogon` service automatically starts when the ASU server starts. These secure communications channels are used to exchange user identification data between computers' `NetLogon` services. The `NetLogon` service on PDCs and BDCs likewise attempts discovery with all trusted domains. After a domain controller is discovered that can authenticate the user, it is used for subsequent user account authentication.

2.4.1 Interactive Logon Authentication

On the user's local system, the Logon Information dialog box prompts the user for a user name, password, and domain or computer name.

The user enters their user name and password in the User name and Password fields. The user selects a domain or computer name from a Domain list. The content of the Domain list depends on whether or not the computer participates in a domain. If the computer is configured as follows:

- As a member of a workgroup, the list contains the local computer name.
- To participate in a domain, the list contains the computer name and every domain (including trusted domains) in which user account can be authenticated.

To log on to the local computer, the user selects the local computer name. The computer checks its local directory database for the specified user name and password. If a match is found, the logon is approved. If a match is not found, the logon fails.

To log in to a domain, the user selects the domain name. The computer sends the domain name, user name, and password to a domain controller in the selected domain. The domain controller checks the domain name and then checks the user name and password against that domain's directory database and processes the request as follows:

- If the domain name is correct and the user name and password match an entry in the directory database, the domain controller notifies the computer that the logon is approved.
- If the domain name is recognized as a trusted domain, the domain controller passes the user name and password information to a trusted domain controller for authentication. If the trusted domain controller authenticates the account, the logon information is passed back to the initial domain controller and the user is logged on. If the account is not authenticated (that is, not defined in the trusted domain directory database), the logon fails.
- If the domain controller does not recognize the domain name or if there is no entry in the directory database for the user, the logon fails.

2.4.2 Remote Logon Authentication

A remote logon occurs when a user is logged on to a computer or domain and then makes a network connection to another computer. The credentials used at interactive logon are used for the remote logon unless the user overrides those credentials by typing a different domain or computer name and user name in the Connect As box in the Map Network Drive dialog box.

If the user makes a network connection to a computer in a domain that contains their domain user account, the logon proceeds as if the user were connecting using an account on the remote computer. The remote computer authenticates the logon credentials against its directory database. If the account is not in the directory database but the Guest account is enabled with no password set, the user is logged on with Guest privileges. If the Guest account is not enabled, the logon fails. See Chapter 3 for information about the Guest account.

On some administrative screens such as User Manager for Domains, a user might need to precede their user name with the name of the domain where their domain user account was created to distinguish it from another user. For example, user JohnL from the Sales domain might appear as Sales\JohnL. This name would distinguish him from a different JohnL in another domain, such as Engineering\JohnL.

2.4.3 Cached Logon Information

The first time that a user logs on to a domain account from a given computer, a domain controller downloads validated logon information (from the directory database) to the computer. This downloaded information is cached on the computer. On subsequent logons, if a domain controller is not available, the user can log on to the domain account using the cached logon information.

Computers running Windows NT Server and Windows NT Workstation store the information used to authenticate the last several (the default is 10) users who logged on interactively. The credentials for users who log on to the local computer also are stored in that computer's local directory database.

2.5 Managing Domains

Managing domains includes the following:

- Synchronizing the directory database between the PDC and BDCs
- Promoting and demoting controllers
- Managing the domain account policy and audit policy
- Managing trust relationships

2.5.1 Synchronizing the Directory Database

The PDC automatically sends timed notices that signal the BDCs to request directory database changes from the PDC. The notices are staggered so that all BDCs do not request changes at the same time. When the BDC requests changes, it informs the PDC of the last change it received. Thus the PDC

always is aware of which BDC needs changes. If a BDC is up-to-date, then the BDC does not request changes.

2.5.1.1 Directory Database Changes

Changes to the directory database consist of any new or changed passwords, user accounts, groups, and any changes in group memberships and user rights.

Changes to the directory database are recorded in the change log. The log holds approximately 2000 changes. The size of the change log determines how long changes can be held. When the log is full and a new change is added, the oldest change is deleted. When a BDC requests changes, those changes which occurred since the last synchronization are copied to the BDC. If a BDC does not request changes in a timely manner, then the entire directory database must be copied to that BDC. For example, if a BDC is off-line for an extended period of time, the number of changes that can occur during that period may exceed the number that can be stored in the change log.

2.5.1.2 Full and Partial Synchronization

Replicating the entire directory database to a BDC is called full synchronization. Full synchronization is performed automatically when a new BDC is added to a domain or when changes have been deleted from the change log before replication takes place. The `NetLogon` service default settings for the timing of updates (every five minutes) and the size of the change log ensure that full synchronization will not be required under most operating conditions.

Replicating only directory database changes that have occurred since the last synchronization to a BDC is called partial synchronization. You can force a partial synchronization to all BDCs or to a particular BDC. For example, if a new user is added to the domain or if a user's password changes, you can perform a partial synchronization to quickly replicate the new user's account or password change to the directory database on all BDCs.

You can use the Computer menu in Server Manager GUI to synchronize the directory database. The options that are available depend on the type of computer that is selected as follows:

- When the PDC is selected, the Synchronize Entire Domain option is available. This command copies the latest directory database changes from the PDC to all of the BDCs in the domain. Synchronize Entire Domain initiates synchronization of all BDCs without waiting for completion of any synchronization in progress.

- When a BDC is selected, the Synchronize With Primary Domain Controller option is available. This command copies the latest directory database changes to the selected BDC only.

See Synchronizing a Backup Domain Controller with the Primary Domain Controller and Synchronizing All Servers of the Domain in the Server Manager Help for information on how to synchronize domain controllers.

2.5.2 Promoting and Demoting Controllers

There might be circumstances in which the PDC might become unavailable. For example, you might need to upgrade software or perform other maintenance operations.

If your domain consists of only a PDC and it becomes unavailable, users cannot log in. If your domain consists of a PDC and BDCs and the PDC becomes unavailable, users can log in to the domain; however, they cannot access shares that are on the PDC and you cannot create or modify user accounts because the master directory database is only on the PDC.

If the PDC becomes unavailable and you have a BDC in the domain, you can promote the BDC to be the PDC. When a BDC is promoted to PDC, an up-to-date copy of the domain's directory database is replicated from the old PDC to the new one, and the old PDC is demoted to a BDC.

If a BDC is promoted to PDC and the former PDC returns to service, you must demote the former PDC to BDC. Until the former PDC is demoted to a BDC, it will not run the `NetLogon` service nor participate in authentication of user logons, and its icon in the Server Manager window will be dimmed.

You cannot move a PDC to another domain until you promote a BDC. To move an ASU server to another domain, use the `joindomain` command. See `joindomain`(8) for more information.

Note

Usually, when a BDC is promoted to a PDC, the system automatically demotes the former PDC to a BDC. However, if Server Manager cannot locate the PDC, the PDC is not demoted and the user receives a message indicating this condition. You can choose to proceed without demoting the PDC or wait until the PDC can be demoted.

See Promoting a Backup Domain Controller to Primary Domain Controller and Demoting a Primary Domain Controller to Backup Domain Controller in Server Manager Help for more information.

2.5.3 Managing Domain Security Policies

ASU security policy settings can provide different levels of security for domain user actions on domain controllers, workstations, and member servers. You should establish a domain security policy as part of planning your domain.

When administering domains, security policy applies to the PDC and BDCs in the domain (they share the same security policy). When administering a computer running as a member server, security policy applies only to that computer.

You can define the following security policies:

- The domain user account policy defines how passwords are used by domain user accounts.
- The audit policy defines the types of events that are recorded in an event log.

2.5.3.1 Domain User Account Policy

The domain user account policy controls how domain user account passwords are used by all user accounts for a computer or domain and also determines the account lockout policy. Changes to account policy affect every domain user account on the computer or domain at the next logon. Table 2–1 describe the domain user account policy options that you can set.

Table 2–1: Domain User Account Policy Options

Option	Description
Maximum Password Age	The period of time that a password can be used before the system requires the user to change it.
Minimum Password Age	The period of time that a password must be used before the user is allowed to change it.
Minimum Password Length	The fewest characters that a password can contain.
Password Uniqueness	The number of new passwords that must be used by a user account before an old password can be reused.
Lockout After	The number of incorrect logon attempts that will cause the account to be locked. The range is from 1 to 999. A locked account remains locked until an administrator unlocks it or until the time specified by the Lockout Duration option passes.

Table 2–1: Domain User Account Policy Options (cont.)

Option	Description
Reset Count After	The maximum number of minutes that can occur between any two bad logon attempts. The range is from 1 to 99999. For example, if the Lockout After option is set to 5, and the Reset Count After option is 30 minutes, then 5 bad logon attempts, each 29 minutes apart, would cause lockout.
Lockout Duration	Whether or not user accounts are locked until an administrator unlocks them (Forever) or for the specified number of minutes (Duration).
Forcibly Disconnect Remote Users From Server When Logon Hours Expire	Whether or not a user account that exceeds the time set by the Logon Hours option for that user is disconnected from all controllers in the domain. The user receives a warning message a few minutes prior to expiration of the logon hours. If you disable this option, the user is disconnected when Logon Hours has been reached, but no new connections are allowed and a warning message is sent every five minutes. See Section 3.1.1 for information on setting logon hours for a user.
Users Must Log On In Order To Change Password	Whether or not users can change their own passwords when they expire. If you disable this option, users can change their own passwords when they expire without help from an administrator.

You can set domain user account policy options by using the following:

- The `net account` command. For more information on the `net account` command, enter:

  # **`net help accounts /options | more`**

- The User Manager for Domains and selecting the Account option from the Policies menu.

  See Managing the Account Policy in User Manager for Domains Help for more information.

2.5.3.2 Audit Policy

Auditing allows you to record in an event log file selected activities of domain user accounts. In a domain, the audit policy applies to the PDC and all BDCs in the domain.

The ASU server can audit and record a range of events including system-wide events such as a user logging on and attempts by a particular user to read specific files. Both successful and unsuccessful attempts to perform an event can be audited and recorded. Table 2–2 describes the events that you can audit and record.

Table 2–2: Audit Events

Event	Description
Logon and Logoff	A user logged on or off or made a network connection.
File and Object Access	A user opened a directory or a file that is set for auditing in the File Manager, or a user sent a print job to a printer that is set for auditing in the Print Manager.
Use of User Rights	A user used a user right (except those rights related to logon and logoff).
User and Group Management	A user or group was created, changed, or deleted. A user account was renamed, disabled, or enabled; or a password was set or changed.
Security Policy Changes	A change was made to the User Rights, Audit, or Trust Relationships policies.

Because an event log size is limited, select the events to be audited carefully and consider the amount of disk space you are willing to devote to the log. The maximum size of the log is defined in the Event Viewer.

You can set the audit policy by using the following:

- The `net auditing` command. For more information on the `net auditing` command, enter:

  # **`net help auditing /options | more`**

- The User Manager for Domains and selecting the Audit option from the Policies menu.

  See Managing the Audit Policy in User Manager for Domains Help for more information.

You can view the logs by using:

- The `elfread` command on the system on which the ASU server is running. See `elfread`(8) for more information.
- The Event Viewer on a Windows system. See Chapter 6 for more information.

2.5.4 Managing Trust Relationships

Trust relationships move the convenience of centralized administration from the domain level to the network level. By establishing trust relationships between the domains on your network, you enable domain user accounts to be used in domains other than the domain in which these accounts are created. You need to create each domain user account only once and, through trust relationships, the account can be given access to any computer on your network, not just the computers in one domain.

2.5.4.1 Creating a Trust Relationship

Creating a trust relationship between two domains creates a computer account to be used by computers in the trusting domain to establish secure communications to the trusted domain.

Both the trusting and trusted domains must be configured to establish a trust relationship as follows:

1. On the PDC in the trusted domain, add the name of the trusting domain to the list of domains that trusts it and assign a password for this relationship.
2. On the PDC in the trusting domain, add the name of the domain to be trusted to its list of trusted domains. This requires the password from Step 1.

Note

A trust relationship will not work if the password assigned to a trust relationship changes on the trusted or trusting domain. If the password changes on the trusted or trusting domain, both sides of the trust relationship must be removed and recreated. When you recreate the trust relationship, you again must assign matching passwords for the trusting and trusted domains.

You can create and remove a trust relationship by using the following:

- The `net trust` command. For more information on the `net trust` command, enter the following command:

  # **`net help trust /options | more`**

- The User Manager for Domains and selecting the Trust Relationship option from the Policies menu.

  See Trust Relationship Policy in User Manager for Domains Help for more information.

2.5.4.2 User Manager for Domains Limitations

The User Manager for Domains that is included in the Windows NT Server Tools program group for Windows 95 computers has limitations that affect the administration of trusted domains.

To use the Trust Relationships dialog box to trust another domain, and to use Users and Groups dialog box to grant privileges to users and groups in a trusted domain, at least one of the following conditions must exist:

- The other domain already must trust your domain.
- The domain user account that you are logged on to has the same name and password as a domain user account in the other domain.
- The other domain has enabled its Guest account, and the domain user account that you are logged on to does not have the same name as any domain user account in the other domain.

Additionally, the User Manager for Domains that is included in the Windows NT Server Tools program group cannot verify trust relationships between domains. Be sure to enter correct passwords for the trust relationship. If you are performing this procedure from Windows NT Server Tools, you will receive a message indicating the trust relationship could not be verified.

3

Domain User Accounts and Groups

By default, a user must be authenticated by Windows NT and Tru64 UNIX security policies when they request access to shares. Therefore, users who access shares must have the following:

- A domain user account that you create. The ASU server uses this account to authenticate the user to enforce the Windows NT security policy set for the share.
- A Tru64 UNIX user account that is automatically created when you create the domain user account. The Tru64 UNIX operating system software uses this account to authenticate the user to enforce the Tru64 UNIX security policy set for the file system or printer that is associated with a share.

A domain user account and Tru64 UNIX user account contain information about the user, including name, password, and other information that a security system uses to authenticate a user.

This chapter describes:

- Domain user accounts
- Tru64 UNIX user accounts
- Grouping domain user accounts

See the ASU *Installation and Administration Guide* for information on creating domain user accounts and groups.

3.1 Domain User Accounts

A domain user account contains information that defines a user in an Advanced Server domain.

Table 3–1 describes the elements of a domain user account.

Table 3–1: Domain User Account Elements

Account Element	Description
User name	The unique name that the user types when logging on.
Full name	The user's full name.

Table 3–1: Domain User Account Elements (cont.)

Account Element	Description
Description	Text describing the user or user account.
Password	The user's password. See Table 3–2 for setting password options.
Logon hours	The hours during which the user can log in to the domain. See Section 3.1.1 for more information about setting Logon Hours. Whether or not users are forced to log off when their logon hours expire is defined by the Forcibly Disconnect Remote Users From Server When Logon Hours Expire option in the domain's account security policy. See Section 2.5.3.1 for more information about setting domain account security policy.
Logon workstations	The computer names of workstations from which a user can log in to the domain. By default, the user can log in from any workstation.
Expiration date	A future date when the account automatically becomes disabled; it is useful to ensure that accounts for temporary employees or students are not kept active unnecessarily.
Logon script	A batch file or executable file that runs automatically when the user logs on. See Section 3.1.2 for more information.
Home directory	A directory that is private to the user. See Section 3.1.3 for more information.
Profile	The path to a folder containing information that is retained to create the user's desktop environment, such as program groups, network connections, and screen colors. See Section 3.1.4 for more information.
Account type	Account type is either global or local. Most accounts you create will be global accounts.

Table 3–2 describes the options that you can use to control a user's password.

Table 3–2: Domain User Account Password Options

Option	Description
User Must Change Password at Next Logon	Whether or not the user will be forced to change their password the next time they log in to the domain. This option is defined by the Maximum Password Age option in the domain's account security policy. See Section 2.5.3.1 for more information about setting domain account security policy.

Table 3–2: Domain User Account Password Options (cont.)

Option	Description
User Cannot Change Password	Whether or not the user can change their password. This restriction is useful for shared accounts. It does not apply to administrators.
Password Never Expires	Whether or not the user's passwords expires. This is used for accounts that represent services such as the `Replicator` service. It also is useful for accounts for which you never want the password to change, such as guest accounts.
Account Disabled	Whether or not the account is disabled and cannot be used to log on. Although the account is not removed from the directory database, the user cannot log in to the domain until the account is enabled.

A domain user account includes a security identifier (SID), a unique number that identifies the account. Every account on your network is issued a unique SID when the account is created. Internal processes refer to an account's SID rather than the account's user name. If you create an account, delete it, and then create an account with the same name, the new account will not have the rights or permissions that previously were granted to the old account because the accounts have different SID numbers. Renaming a user account retains the SID.

3.1.1 Specifying Logon Hours

By default, users can connect to the ASU server 24 hours a day, seven days a week; however, you can restrict access.

When a user is connected to a controller and the logon hours are exceeded, the user either will be disconnected from all server connections or will be allowed to remain connected but denied any new connections, depending on if you enabled the Forcibly Disconnect Remote Users From Server When Logon Hours Expire option. See Section 2.5.3.1 for more information on setting this option.

You specify the logon hours for a user by using the following:

- The `net user /TIMES:{times | ALL}` command. For more information on the `net user` command, enter the following command:

  # **`net help user /options | more`**

- The User Manager for Domains and viewing properties for a selected user, then selecting Hours in the User Properties dialog box.

  A Logon Hours dialog box is displayed. The Logon Hours dialog box displays a one-week calendar, with logon hours displayed in one-hour

increments across seven days. A box represents each hour. For example, the first box in each row represents the hour from midnight through 12:59 A.M., and the last box in each row represents the hour from 11:00 P.M. through 11:59 P.M. The filled boxes indicate when the user is allowed to connect to controllers; the empty boxes indicate when a user is prohibited from connecting.

Note

The logon hours are in the time zone of the PDC, not in the time zone of the workstation or server to which the user is logging on or connecting.

3.1.2 Logon Script

A logon script is an executable or batch file that contains commands that automatically run when a user logs on to a system on which the ASU server is running. Logon scripts are used to automatically configure users' working environments, such as making network connections and starting applications.

The best way to ensure that logon scripts always are available is to use the `Replicator` service. The `Replicator` service maintains identical copies of a directory tree on multiple controllers. When you make a change to a file in the master copy of the tree (located on the export server), the `Replicator` service automatically copies the change to the import controllers.

3.1.2.1 Creating a Logon Script

You create logon scripts by using a text editor, then using the `Replicator` service to distribute the logon scripts to domain controllers.

When using the `Replicator` service, the ASU server searches for logon scripts in the `/usr/net/servers/lanman/shares/asu/repl/import/scripts` directory on import servers and in the `/usr/net/servers/lanman/shares/asu/repl/export/scripts` directory on export servers. If you do not use the `Replicator` service, you must make sure that the logon scripts are in the same directory on each controller because any controller can validate the logon request, then execute the logon script. See Section 4.2 for more information on directory replication.

Table 3–3 describes special parameters that you can use when creating logon scripts.

Table 3–3: Logon Script Parameters

Parameter	Description
%HOMEDRIVE%	The user's local workstation drive letter connected to the user's home directory
%HOMEPATH%	The full path of the user's home directory
%HOMESHARE%	The name of the share containing the user's home directory
%OS%	The operating system of the user's workstation
%PROCESSOR_ARCHITECTURE%	The processor type of the user's workstation
%PROCESSOR_LEVEL%	The processor level of the user's workstation
%USERDOMAIN%	The domain containing the user's account
%USERNAME%	The domain user name

3.1.2.2 Assigning a Logon Script

You assign a logon script to a domain user account by using the following:

- The `net user` *`username`* `/scriptpath:[`*`pathname`*`]` command. For more information on the `net user` command, enter the following command:

  # **`net help user /options | more`**

- The User Manager for Domains and selecting the properties for a user, clicking on the Profile button, and entering the path of the logon script in the Logon Script Name box.

3.1.3 Home Directory

A user's home directory is a directory that is accessible to the user and contains files and programs for that user. When a user logs on at a workstation, a connection is made automatically to that user's home directory; this becomes that user's default directory for the File Open and Save As dialog boxes, for the command prompt, and for all applications that do not have a working directory defined. By default, when you create a domain user account, the ASU server creates a home directory for the user in the `/usr/users` directory using the user's Tru64 UNIX account name.

If you do not specify a home directory, the default home directory is the `\USERS\DEFAULT` directory on the user's local drive.

Because home directories collect user files in one location, they make it easy for an administrator to back up user files and delete user accounts.

3.1.3.1 Assigning Home Directories

You assign a home directory to a user account by using the following:

- The `net user` *`username`* `/homedir:`*`pathname`* command. For more information on the `net user` command, enter the following command:

  # **`net help user /options | more`**

- The User Manager for Domains and selecting the properties for a user, clicking on the Profile button, and entering the path of the home directory in the Home Directory box.

If you specify a home directory on a controller and it does not exist, it is created. If you specify a directory on the user's local system and it does not exist, it is not created.

In the Home Directory box, `%USERNAME%` can be substituted for the last entry in the path. The system later substitutes the user name of the domain user account. This substitution is useful when multiple domain user accounts are selected. For example, you have selected eight domain user accounts. In the Home Directory box, you might select Connect, specify a drive letter of K, select the To box, and type `\\SALES\home\%username%`. When you choose OK to save the User Environment Profile, the actual user name will be substituted for each `%USERNAME%` entry.

When a domain user account is copied, the home directory is copied in one of the following two ways:

- If the last subdirectory of the home directory path of the user account being copied is the user's name, the new account substitutes the new user's name in the home directory path. For example, if the original domain user account has a user name of PETER and a home directory of \\SETTER\USERS\PETER, a copied (new) domain user account with the user name of EVAN would have a home directory of \\SETTER\USERS\EVAN.
- If the last subdirectory of the home directory path of the user account being copied is not the user's name, then the home directory path is copied exactly. For example, if the original domain user account has the user name of PETER and a home directory of \\HOUND\USERS\HOME, then a copied (new) domain user account with the user name of EVAN would have the same home directory of \\HOUND\USERS\HOME.

3.1.4 User Profile

A user profile defines the work environment settings that are set when a user logs into a domain. These settings include all the user-specific settings of a user's Windows environment, such as which options in Control Panel

they can use, screen colors, share connections, mouse settings, shortcuts, window size, and position.

Types of user profiles include the following:

- Local user profiles that are created automatically on the computer the first time a user logs on to a computer running the Windows operating system software. Each user's individual user profile is available to that user on successive logons at that computer.
- Roaming user profiles that are available on computers running the ASU server and on computers running the Windows operating system software. To enable roaming user profiles, an administrator enters a user profile path into the user account. The first time the user logs off, the local user profile is copied to that location. Thereafter, the copy of the user profile is downloaded each time the user logs on (if it is more current than the local copy). Both the local and server copies are updated each time the user logs off.
- Mandatory user profiles that are roaming profiles that are created for the user and cannot be changed by the user. When the user logs off, the local user profile is not saved and a copy of the local user profile is not copied to the server.

User profiles are available on computers running Windows 95; however, a user profile created on Windows 95 is not available to the user on a computer running the ASU server or Windows NT, even if the user profile is stored on these systems.

3.1.4.1 Creating a User Profile

You use the System Policy Editor to create a user profile. The System Policy Editor is a Windows based interface that you can use to view and manage policies that define the environment for specific Windows computers, users, or groups when logged in to a domain.

Using the System Policy Editor you can set the following:

- A user-specific policy that applies to each domain user account or group. Most policies are user-specific.
- A machine-specific policy that applies to all users on a Windows system and does not change according to user since it does not follow users as they move between different systems.

The System Policy Editor saves settings in a policy (`.POL`) file. When a user logs in, a program called the policy downloader starts. The policy downloader is installed on every Windows client. The policy downloader looks on the network for the policy file, opens the policy file, looks for an entry using the local computer name or user name, and merges the administrator's

registry settings as defined in the policy file into the local registry. If the downloader does not find an entry with the local computer name or user name in the policy file, then it looks for the DEFAULT USER or DEFAULT COMPUTER entry and uses those registry settings for the merge. If there are no entries for a specific user or computer and default entries do not exist, then no merge takes place.

See the ASU *Installation and Administration Guide* for more information about the System Policy Editor.

3.1.4.2 Assigning Profiles

You assign a profile to a domain user account by using the following:

- The `net user` *`username`* `/profilepath:[`*`pathname`*`]` command. For more information on the `net user` command, enter the following command:

  # **`net help user /options | more`**

- The User Manager for Domains and selecting the properties for a user, clicking on the Profile button, and entering the path of the profile in the User Profile Path box.

3.1.5 Managing the User Rights Policy

A right authorizes a user to perform certain actions on a computer system, such as backing up files and directories, logging on to a computer interactively, or shutting down a system. Rights exist as capabilities for using either domain controllers at the domain level or workstations or member servers at the local level. Rights can be granted to domain user accounts or groups.

A user who logs on to a domain user account or belongs to a group to which the appropriate rights have been granted to perform actions can carry out the actions. When a user does not have appropriate rights to perform an action, an attempt to carry out that action is blocked by the ASU server.

Rights apply to the system as a whole and are different from permissions, which apply to specific objects. A permission is a rule associated with an object (usually a directory, file, or printer), and it regulates which users can have access to the object and in what manner. Most often the creator or owner of the object sets the permissions for the object. However, because all rights are not associated with a specific object and are applied at the domain (domain controllers) or local (workstation or member server) level, they sometimes can override permissions set on an object. For example, a user logged on to a domain account that is a member of the Backup Operators group has the right to perform backup tasks for all servers of the domain. Doing so requires the ability to read all files on those servers, even files

on which their owners have set permissions that explicitly deny access to all users, including members of the Backup Operators group. A right, in this case, the right to perform a backup, takes precedence over all file and directory permissions.

Table 3–4 describes user rights. See Section 4.1 for more information about permissions.

Table 3–4: User Rights

User Right	Allows
Access this computer from network	Connecting over the network to a controller. This right cannot be revoked from the Administrator's local group.
Add workstations to domain	Adding a workstation to the domain, allowing the workstation to recognize the domain's user accounts and global groups and those of trusted domains.
Back up files and directories	Backing up files and directories and reading of all files. This right supersedes file and directory permissions, and also applies to the Registry.
Change the system time	Setting the time for the internal clock of a controller.
Force shutdown from a remote system	This right is not currently implemented. It is reserved for future use.
Load and unload device drivers	Installing and removing device drivers.
Log on locally	Logging on locally at the controller.
Manage auditing and security log	Specifying what types of share access (such as file access) are to be audited. Viewing and clearing the security log. This right does not allow a user to set system auditing. This ability is held only by the Administrators group.
Restore files and directories	Restoring (writing) files and directories. This right supersedes file and directory permissions, and also applies to the Registry.
Shut down the system	Shutting down the controller.
Take ownership of files or other objects	Taking ownership of files, directories, and other objects on a controller.

3.1.5.1 Assigning User Rights

You assign rights to a domain user account by using the User Manager for Domains and selecting Policies, then User Rights.

When you assign user rights for a domain, the rights apply to all the controllers. When you administer user rights on a workstation or member server, the rights apply only to the workstation or member server.

3.1.6 Built-in Domain User Accounts

Two built-in user accounts are created automatically when you install the ASU software:

- The Administrator account
- The Guest account

3.1.6.1 Built-in Administrator User Account

The built-in Administrator account is allowed to perform domain management tasks on a controller, workstation, or member server that belongs to that domain.

When you install the ASU software you are prompted for a password to the built-in Administrator account. This password should be guarded carefully because if the password is forgotten or the person who knows the password becomes unavailable, the built-in Administrator account is unusable. The password can be changed and it does not expire.

The built-in Administrator account can never be deleted or disabled. This feature distinguishes the Administrator account from other members of the Administrators local group.

Following installation of the ASU software, you should create additional administrative accounts with administrative-level capabilities, while reserving the built-in Administrator account for emergency purposes. When administrative users have separate accounts, their actions can be audited on an individual user account basis rather than trying to audit the built-in Administrator account.

See the ASU *Installation and Administration Guide* for information about installing the ASU software. See Chapter 6 for information about auditing.

3.1.6.2 Built-in Guest User Account

The built-in Guest account can be used by a user who needs to log in to the domain, but does not have a domain user account in the domain or in a trusted domain. A user whose account is disabled (but not deleted) also can use the Guest account.

The Guest account is disabled by default, does not require a password, and has no predefined rights or permissions in the domain however, you can

set a password and rights and permissions for the Guest account like any other domain user account.

There are two types of guest logons as follows:

- A local guest logon. A local guest logon occurs when a user logs on interactively at a computer running Windows software and specifies Guest as the user name in the Logon Information dialog box. Because the Guest account on these computers (but not on domain controllers) has the built-in right to log on locally, the guest user can then work at that computer (subject to the rights and permissions you have granted the Guest account) and use it to access the network.
- A network guest logon. A network guest logon occurs when a user makes a network connection to a controller and that controller does not recognize the user's user name, domain name, and password. A network guest logon is approved only if the Guest account of the destination controller is enabled and has no password set. The guest user then has all rights, permissions, and group memberships on the controller that are granted to the Guest account, even though the user did not specify Guest as their user name.

3.1.6.2.1 Enabling the Guest Account

You enable the Guest account on each controller by using the following:

- The `net user Guest /active:yes` command. For more information on the `net user` command, enter the following command:

  # **`net help user /options | more`**

- The User Manager for Domains and selecting the Guest user, then clicking in the box next to Account Disabled to remove the check mark.

3.2 Tru64 UNIX User Accounts

By default, the Tru64 UNIX operating system software authenticates a user before they can access a file system or printer that is assocated with a share.

By default, when you create a domain user account, a Tru64 UNIX user account is automatically created in the local `/etc/passwd` file if an account with the same name does not exist. The Tru64 UNIX operating system software uses this account to authenticate the user. However, you can configure the Tru64 UNIX operating system software to send authentication requests to a network information service (NIS), a Windows 2000 Server, or to a Windows NT Server Version 4.0. The NIS, Windows 2000 Server, or Windows NT Server Version 4.0 uses the information in its user account database to authenticate users on behalf of the Tru64 UNIX operating system software and sends the results back. This is useful if you have a user

account database on a NIS, Windows 2000 Server, or on a Windows NT Server Version 4.0 and you do not want to create a user account database on the Tru64 UNIX system.

See the ASU *Installation and Administration Guide* for more information on configuring user account authentication.

3.2.1 Associating Domain User Accounts to Tru64 UNIX User Accounts

By default, a domain user account is associated with a Tru64 UNIX user account. This association allows your Tru64 UNIX files to be owned by your Tru64 UNIX system user account and to be accessed by the domain user account. Domain user accounts that are not mapped to a specific Tru64 UNIX user account are mapped by default to the `lmworld` Tru64 UNIX user account.

The Tru64 UNIX user account name that is assigned to the domain user account will be the same as or similar to the domain user account name. Differences can arise in cases of long, duplicate, or special character user account names. If a domain user account user is associated to a non-existent Tru64 UNIX user account, or if the Tru64 UNIX user account for the user is deleted, the user will not have access to any shares.

You use the `mapuname` command to control the association of domain user accounts to Tru64 UNIX user accounts. The default relationship between domain user accounts and Tru64 UNIX user accounts is controlled by values assigned to user-related registry value entries in the ASU registry.

See the ASU *Installation and Administration Guide* for more information about the ASU registry.

3.3 Grouping Domain User Accounts

To ease administration, you can group domain user accounts and administer the group as one unit. Domain user accounts added to a group become members of the group and immediately acquire the rights and permissions granted to the group. Changes made to the group affect each member. Group membership provides an easy way to grant common capabilities to sets of users.

There are the following three types of groups:

- Local groups
- Global groups
- Special groups

3.3.1 Local Groups

A local group contains domain user accounts and global groups from the local domain and from domains that it trusts. A local group cannot contain other local groups. Table 3–5 describes the built-in local groups.

Table 3–5: Local Groups

Group	Description
Administrators	Members are automatically granted every built-in right and ability to manage the domain and controllers. By default, the Domain Admins global group is a member of the Administrators local group.
Users	Members have normal user access to and capabilities for the domain. By default, the Domain Users global group is a member of the Users local group.
Guests	Members have no rights on controllers. However, they do have certain rights at their individual workstations. By default, the Domain Guests global group is a member of the Guests local group.
Account Operators	Members can create, modify, and delete most domain user accounts and groups, log on to and shut down domain controllers, and add controllers or member servers to a domain. Members cannot administer security policies, modify or delete the Domain Admins global group, the Administrators, Account Operators, Backup Operators, Print Operators, or Server Operators local groups, any global groups belonging to these local groups, or accounts of members of any of these groups.
Backup Operators	Members can back up and restore files on a PDC and BDCs.
Print Operators	Members can create, delete, and manage printer shares on the domain's controllers and can shut down these controllers.
Server Operators	Members can create, delete, and manage shares and change the system time.
Replicator	Members can mange the `Replicator` service of the PDC and the BDCs in the domain. See Chapter 4 for information about directory replication.

Not all built-in local groups exist on ASU servers, Windows NT servers, Windows NT workstations, and member servers. Table 3–6 identifies which

built-in local groups exist on ASU servers, Windows NT servers, Windows NT workstations, and member servers.

Table 3–6: Built-in Local Groups

ASU Servers and Windows NT Servers	Windows NT Workstations and Member Servers
Administrators	Administrators
Users	Users
Guests	Guests
Account Operators	Backup Operators
Backup Operators	Replicator
Print Operators	Power Users
Server Operators	
Replicator	

3.3.2 Global Groups

A global group contains domain user accounts from the domain in which it is created. A global group cannot contain local groups, other global groups, or domain user accounts from other domains. You can add a global group to local groups in the same domain, to domains that trust that domain, or to member servers or computers running Windows NT Workstation in the same or a trusting domain. You can grant a global group permissions and rights in its own domain, on workstations or member servers, or in trusting domains. You cannot create a global group on a computer running Windows NT Workstation or on computers running as a member server.

Table 3–7 describes the built-in global groups. These groups cannot be deleted.

Table 3–7: Global Groups

Group	Description
Domain Admins	Members can administer the domain, the controllers and workstations of the domain, and a trusting domain that has added the Domain Admins global group from this domain to the Administrators local group in the trusting domain. Only Administrators can modify the Domain Admins global group. The built-in Administrator user account is a member of the Domain Admins global group. The Domain Admins global group is a member of the Administrators local group. To provide a domain user account with administrative-level capabilities, add the account to the Domain Admins global group.
Domain Users	Members have normal user access to and capabilities for the domain and the computers in the domain running Windows NT Workstation and Windows NT Server as a member server. Only Administrators and Account Operators can modify the Domain Users global group. The built-in Administrator user account is a member of the Domain Users global group. By default, all new domain user accounts created in the domain are added to the Domain Users group, unless you specifically remove them. The Domain Users global group is a member of the Users local group for the domain and of the Users local group for every computer in the domain running Windows NT Workstation or member servers running Windows NT Server.
Domain Guests	Members have no rights on controllers. The built-in Guest user account is a member of the Domain Guests global group. The Domain Guests global group is a member of the Guests local group. Add user accounts that are intended to have more limited rights and permissions than typical domain user accounts to the Domain Guests group and remove them from the Domain Users group. Only Administrators and Account Operators can modify the Domain Guests global group.

3.3.3 Special Groups

Table 3–8 describes the other groups that used for special purposes. Because the memberships of these groups cannot be altered, the groups are not listed in User Manager for Domains.

When you administer the ASU server, these special groups sometimes appear in the list. For example, they can appear when assigning permissions to shares and files.

Table 3–8: Special Groups

Group	Description
Everyone	Members include all local and remote users (that is, the Interactive and Network groups combined). In a domain, members can access the network and connect to disk and printer shares.
Interactive	Members include anyone using the computer locally.
Network	Members include users connected over the network to the computer.
System	Members include the operating system.
Creator Owner	Members include users who create a subdirectory, file, or print job. For a directory, if permissions are granted to the Creator Owner group, the creator of a subdirectory or file will be granted those permissions for that subdirectory or file. For a printer, if permissions are granted to the Creator Owner group, the creator of a print job will be granted those permissions for that print job.

3.3.4 Strategies for Using Groups

Some rights are only provided to users who are members of a specific group. For example, the only way to allow a person to create domain user accounts is to add that person's domain user account to either the Administrators or Account Operators local group on the domain.

You can add a user to more than one built-in group. For example, a user in both the Print Operators and Backup Operators groups has all the rights granted to Print Operators and all the rights granted to Backup Operators.

In a domain, rights are granted and restricted on the domain level; if a group has a right in a domain, its members have that right on all the controllers in the domain. On member servers, rights granted apply only to that computer.

In a multiple-domain setting, you can think of global groups as a means of adding users to the local groups of trusting domains. To extend users' rights and permissions to resources on other domains, add their accounts to a global group in your domain and then add the global group to a local group in a trusting domain.

Even if you maintain a single domain, keep in mind that additional domains may be added in the future. You can use global groups added to local groups for granting all rights and permissions. Later, if another domain is created, the rights and permissions assigned to your local groups can be extended to a new domain's users by creating a trust relationship and adding global groups from the new domain to your local groups. Likewise, if the new

domain trusts your domain, your global groups can be added to the new domain local groups.

Domain global groups also can be used for administrative purpose on computers running Windows NT Workstation or on member servers running Windows NT Server. For example, the Domain Admins global group is added by default to the Administrators built-in local group on each workstation or member server that joins the existing domain. Membership in the workstation or member server local Administrators group enables the network administrator to manage the computer remotely by creating program groups, installing software, and troubleshooting computer problems.

Most network administrators have dual roles. They are both administrators and users of the network. A network is more secure if an administrator uses two domain user accounts. One of these accounts should be in the Domain Admins global group and should be used to perform network management tasks and the other account should be in the Domain Users global group and should be used at all other times. While logged on as a regular user, an administrator cannot inadvertently change aspects of the network that only administrators can change. And, if an administrator were to accidentally introduce a virus, the program would not have the rights of an administrator and would not modify the operating system.

Table 3–9 provides some guidelines for using global and local groups:

Table 3–9: Group Guidelines

Purpose of Group	Use	Comments
To group users of a domain into a single unit for use in other domains or user workstations	Global	The global group can be put into local groups or given permissions and rights directly in other domains.
To group users who need permissions and rights only in one domain	Local	The local group can contain users and global groups from this and other domains.
To group users who need permissions on computers running Windows NT Workstation or on member servers	Global	A domain's global groups can be given permissions on these computers, but a domain's local groups cannot.

Table 3–9: Group Guidelines (cont.)

Purpose of Group	Use	Comments
To contain other groups	Local	The local group can contain global groups and users; however, no group can contain other local groups.
To group users from multiple domains	Local	The local group can be used in only the domain in which it is created. If you need to be able to grant this local group permissions in multiple domains, you must create the local group in every domain in which you need it.

3.3.5 Managing Groups

A group name must be unique to the domain or to the computer being administered. A global group name can contain up to 20 characters. A local group name can contain up to 256 characters. A group name can contain any uppercase or lowercase alphanumeric characters.

When a group name is displayed and when the distinction is necessary, the ASU server identifies the domain or workstation the group is from by displaying the group name in the form *DOMAINNAME**groupname* or *COMPUTERNAME**groupname*. For example, a group named Managers from a domain named Engineering would be displayed as ENGINEERING\Managers.

You can copy a group or create one. By copying, you ensure that the new group has the same members as the original group. However, the permissions and rights of the original group are not copied to the new group.

A group includes a security identifier (SID), a unique number that identifies the group. Every group on your network is issued a unique SID when the account is first created. Internal processes refer to a group SID rather than the group name. If you create a group, delete it, then create a group with the same name, the new group will not have the rights or permissions that previously were granted to the old group because the new group will have a different SID. Renaming a group retains the SID. Groups that you create can be deleted, but the built-in groups cannot. Deleting a group removes only that group; it does not delete the domain user accounts or global groups that are members of the deleted group.

4

Disk Shares

The ASU server enables you to share UNIX directories as disk shares, then make those shares available to domain user accounts and groups. For example, when a directory is shared, authorized Windows users can make connections from their workstations to the share and access its files.

The only way to make a file accessible over the network is to share one of its parent directories as disk share. When you share a directory, Windows users can gain access to that directory, the files in it, and all of the subdirectories and their contents. Every point on the directory tree below the shared directory is available to users. You use Windows NT permissions to manage user access to the share and Windows NT file system (NTFS) and Tru64 UNIX permissions to manage user access to files and directories in shares.

This chapter discusses the following topics:

- Disk share permissions
- Directory replication
- Managing disk share usage

See the ASU *Installation and Administration Guide* for information on creating disk shares.

4.1 Disk Share Permissions

Every file and directory has an owner. The owner controls how permissions are set on the file or directory and can grant permissions to others. When a file or directory is created, the person creating the file or directory automatically becomes its owner. By default, a user must pass the following levels of security before they can access a file or directory in a disk share:

- Windows NT share level security
- Windows NT File System (NTFS) security
- Tru64 UNIX file and directory security

The following steps describe how permissions are checked when a user maps a drive to a disk share and requests access to a file in the disk share:

1. From a system running the Windows operating system software, a user connects to a disk share. By default, all users have permission to connect to a share.

   The user's Windows system provides the ASU server with authentication information about the user, including the user's name, password, and security ID.

2. The ASU server checks the user's name and password in the directory database.

   If the ASU server authenticates the user's information, a unique ID is assigned to the user's Windows system. The Windows system must present this ID when the user makes subsequent requests to shares.

3. The user attempts to open a file in the share.

   The ASU `lmx.srv` process services the user's request. Normally, the `lmx.srv` process runs as root, the highest Tru64 UNIX privilege level.

4. The `lmx.srv` process determines if the user has the correct Windows NT share permissions to access the share.

   If the permissions are not correct, the `lmx.srv` returns an access denied error to the Windows system.

5. The `lmx.srv` process determines if the user has the correct NTFS permissions to access the file in the share.

   If the permissions are not correct, the `lmx.srv` process returns an access denied error to the Windows system.

6. The `lmx.srv` process determines Tru64 UNIX access based on the mapping of the domain user account to a Tru64 UNIX user account.

7. The `lmx.srv` process changes its effective user ID from root to the ID of the corresponding Tru64 UNIX account and tries to open the file.

8. The Tru64 UNIX operating system determines if the user has the correct Tru64 UNIX permissions.

   If the permissions are correct, the file is opened. If the permissions are not correct, the `lmx.srv` process returns an access denied error to the Windows system.

4.1.1 Windows NT Permissions

Table 4–1 describes the Windows NT permissions that you can set for a disk share.

Table 4–1: Windows NT Permissions

Permission	Purpose
No Access	Prevents a user from accessing the disk share
Read	Allows: • View file and subdirectory names • Move to subdirectories • View data in files • Run application files
Change	Allows everything Read allows, plus: • Add files and subdirectories • Change data in files • Delete subdirectories and files
Full Control	Allows everything Read and Change allows, plus: • Change Windows NT and NTFS permissions • Set Windows NT and NTFS permission to take ownership of files and subdirectories The default is to grant the Everyone group Full Control permission to new disk shares.

Note

Permissions are cumulative except that the No Access permission overrides all other permissions. For example, if the Coworkers group has Write permission for a file while the Finance group has only Read permission and John is a member of both groups, John will be granted Read and Write permissions. However, if you change the Finance group's permission for the file to No Access, John will not be able to use the file even though he is a member of a group that has write access to the file.

See the ASU *Installation and Administration Guide* for information on setting Windows NT permissions.

4.1.2 Windows NTFS Permissions

There is a standard set of NTFS permissions that you can set or you can customize NTFS permissions to meet you needs. Table 4–2 describes the standard Windows NTFS permissions that you can set. Table 4–3 describes the custom Windows NTFS permissions that you can set.

Table 4–2: NTFS Standard Permissions

Permission	For File	For Directory
Add	Cannot read the contents of current files, change them, or list the files	Can add files to the directory
AddRead	Can read and execute files but cannot change files	Can read, write, and execute files in the directory
Change	Can change the contents of current files	Can read and add files
Full control	Can read and change files, add new ones, change permissions for files, and take ownership of file	Change permissions for the directory and take ownership of the directory
NoAccess	Not applicable	Cannot access the directory in any way, even if the user is a member of a group that has been granted access to the directory
List	Cannot access files	List the files and subdirectories in this directory and change to a subdirectory of this directory
Read	Can read the contents of files and run applications	Allows viewing the names of files and subdirectories

Table 4–3: NTFS Custom Permissions

Permission	For File	For Directory
Change Permissions (p)	Allows changing the file's permissions	Allows changing the directory's permissions
Delete (d)	Allows deleting the file	Allows deleting the directory
Execute (x)	Allows running the file if it is a program	Allows changing to subdirectories
Read (r)	Allows viewing the file's data	Allows viewing the names of files and subdirectories
Take Ownership (o)	Allows taking ownership of the file	Allows taking ownership of the directory
Write (w)	Allows changing the file's data	Allows adding files and subdirectories

NTFS displays two sets of permissions: the permissions set on the directory and the permissions set on files in the directory. For example, the following output would display if you set AddRead permission on a share for a user name peter. The (RWX) means Read, Write, and Execute permissions on the directory, and (RX) means Read and Execute permission on its files.

```
Resource:    c:\usr\net\servers\lanman\shares\share1
Owner:       server1.dom\Administrators
Name:                              Permissions:
-------------------------------------------------------------------------------
*Administrators                    FullControl(All)(All)
*Everyone                          Read(RX)(RX)
peter                              AddRead(RWX)(RX)
```

The ASU server displays groups by preceding the group name with an asterisk (*).

NTFS Permissions on files in a directory can be set to NotSpecified. This means that no permissions will be set for that user or group on the files in the directory or that are created after setting this permission. A user or group cannot access files in the directory until permission is granted.

When you set permissions on a directory, you can use the Creator Owner special group to allow users to control only the subdirectories and files that they create within the directory. Permissions set on the Creator Owner group are transferred to the user who creates a directory or file within the directory. To change permissions on the directory, you must be the owner of the directory or have been granted permission to do so by the owner.

Note

By default, Windows NTFS permissions grant read and execute permission to the Everyone group, of which every domain user account is a member. You must grant Windows NTFS write permission to the domain user account or group that will write files to the disk share.

See the ASU *Installation and Administration Guide* for information on setting Windows NTFS permissions.

4.1.3 Tru64 UNIX Permissions

By default, subdirectories created in a disk share have the following Tru64 UNIX permissions:

- Owner has read and write permission
- Group has read permission
- Other has read permission

By default, files created in a disk share have the following Tru64 UNIX permissions:

- Owner has read and write permission
- Group has read permission
- Other has read permission

See *System Administration* for information on setting Tru64 UNIX permissions.

4.1.4 Disk Share Permission Considerations

The user who creates a file or directory is the owner of that file or directory; however, users who are members of the Administrators group always can take ownership of a file or directory. The owner can control access to the file or directory by changing the permissions set on it. When changing the permissions on a directory, you choose whether to apply the changes to all files and subdirectories in the directory. Users cannot use a directory or file unless they have been granted permission to do so or belong to a group that has permission to do so. Each permission that you set specifies the access that a group or user can have to the directory or file. For example, when you set Read permission for the group called Coworkers on the file `MY_IDEAS.DOC`, the users in the Coworkers group can display the file's data and attributes, but they cannot change the file or delete it.

The easiest way to administer security is by setting permissions for groups, not individual users. Typically, a user needs access to many files. If the user is a member of a group that has access to the files, you can terminate the user's access by removing the user from the group rather than changing the permissions on each of the files. Note that setting permission for an individual user does not override the access granted to the user through groups to which the user belongs.

When copying files or directories, the current security permissions, ownership, and auditing information is discarded. The copied files or directories inherit a new set of permissions from the directory into which they are copied and the person copying the file or directory is the owner.

4.2 Directory Replication

One of the helpful tasks that the ASU server performs is keeping shared resources current, which can be accomplished by using the `Replicator` service. If you have directories and files that you want to distribute to many users, you can use the `Replicator` service to set up and maintain identical directory trees on controllers and workstations to balance the work among them.

The Replicator service requires that you configure one computer as an export server, place the master copies of directories and files on the export server, and configure other computers as import computers. You only need to create one copy of a directory or file on the export server and import computers automatically receive an identical copy of the directory or file. When you update a directory or file in the directory tree on an export server, the updated directory or file is copied automatically to all the import computers.

Every export server maintains a list of computers to which directories and files are exported, and each import computer maintains a list of computers from which directories and files are imported. Export servers can export to domain names and import computers can import from those domain names. This is a convenient way to set up directory replication for many computers; each export server and import computer needs to specify only a few domain names for export or import rather than a long list of computer names.

You can configure an ASU server as an export server and as an import computer.

4.2.1 Directory Replication Overview

The Replicator service must be running on each export server and import computer. The Replicator service sends updated directories and files on each export server to import computers that participates in replication. The Replicator service on each computer logs on to the same user account, which you create for this purpose.

A domain can have multiple export servers; however, to ensure the integrity of replicated information, do not export duplicate subdirectories. The /usr/net/servers/lanman/shares/asu/repl/export directory is the default export path on the export server and should contain the directories and files to be replicated. When changes are saved to subdirectories and files in this directory, the subdirectories and files automatically replace the existing subdirectories and files on all of the import computers.

You also can specify whether the export server sends changes as they occur or wait until an export subdirectory has been stable for two minutes. This prevents exporting partially changed subdirectory trees. In addition, you can lock an export or import directory. Changes to a locked directory are not exported or imported until you unlock the directory.

On the export server, you specify the import computers that are to receive copies of the directories and files that the export server is exporting. An export server has only one list of import computers to which it replicates. All directories to be replicated are exported as subdirectories in the export path. Subdirectories created in the export path and files in those subdirectories

are exported automatically. Export servers can replicate any number of subdirectories (limited only by available memory) with each exported subdirectory having up to 32 subdirectory levels in its tree.

Imported subdirectories and their files are automatically placed in the `/usr/net/servers/lanman/shares/asu/repl/import` directory. You do not need to create import subdirectories. They are created automatically when directory replication occurs.

4.2.2 Configuring Directory Replication

Follow these steps to configure directory replication on export servers and import computers:

1. Create a domain user account for the `Replicator` service to use to log on. Be sure the user account has the Password Never Expires option selected and all logon hours allowed.

   See the ASU *Installation and Administration Guide* for information about creating domain user accounts.

2. For each computer that will be configured as an export server, use the Server Manager, select Properties from the Computer menu, then select Replication and configure:

   - The `Replicator` service to start automatically.
   - To log on using the domain user account that was created for the `Replicator` service.
   - The names of the import computers and domains by adding them to the To List.

   See Server Manager help for more information on configuring the `Replicator` service.

3. On a computer that will be configured as an export server, create the directories to be exported. They must be subdirectories of the replication export path, `/usr/net/servers/lanman/shares/asu/repl/export`.

4. For each computer that will be configured as an import computer, use the Server Manager, select Properties from the Computer menu, then select Replication and configure:

   - The `Replicator` service to start automatically.
   - To log on using the domain user account that was created for the `Replicator` service.
   - The names of the of the export servers and domains by adding them to the From List.

See Server Manager Help for more information on configuring the `Replicator` service.

You can set up an export server to replicate a directory tree to itself (from its export directory to its import directory). This replication can provide a local backup of the files, or you can use the import version of these files as another source for users to access, while preserving the export version of the files as a source master.

4.2.3 Managing Export Subdirectories

Use the Server Manager to manage exported subdirectories by clicking on Manage under Export Directories in the Directory Replication dialog box. Aspects of directory replication that you can manage include the following:

- Locking a subdirectory to prevent it from being exported to any import computers. For example, if you know that a directory will be receiving a series of changes that you do not want partially replicated, you can put one or more locks on the subdirectory in the export path. Until you remove the lock or locks, the subdirectory will not be replicated. The date and time the lock is placed is displayed so that you know how long a lock has been in effect.
- Stabilizing a subdirectory, which causes the export server to wait two minutes after changes before exporting the subdirectory. The waiting period allows time for subsequent changes to take place so that all intended changes are recorded before being replicated.
- Specifying to export all of the export subdirectories or only the first-level subdirectory in the export directory path.

4.2.4 Managing Import Subdirectories

You can use locks to prevent imports to subdirectories on an import computer. Locking a subdirectory on an import computer prevents the replication of subdirectories to that computer until the lock is removed. Locking a subdirectory on an import computer affects replication to only that computer, not to other import computers.

Use the Server Manager to manage locks on subdirectories and also view the status of each subdirectory by clicking on Manage under Import Directories in the Directory Replication dialog box.

The Status column can have one of the following four entries:

- OK indicates that the subdirectory is receiving regular updates from an export server and that the imported data is identical to the data that is being exported.

- No Master indicates that the subdirectory has received updates in the past but is not receiving updates currently. The export server might not be running or a lock may be in effect on the export server
- No Sync indicates that although the subdirectory has received updates, the data is not up-to-date. This could be a result of a communications failure, open files on the import computer or export server, the import computer not having access permissions at the export server, an export server malfunction, or a large subdirectory replication in progress.
- No entry (blank) indicates that replication never occurred for that subdirectory. Replication may not be configured correctly for this import computer, for the export server, or both.

The Last Update column shows the date and time of the latest change to the import subdirectory or to any of its subdirectories.

See To View a List of, or Manage Locks for, Import Subdirectories in Server Manager Help for more information on managing locks.

4.2.5 Replicating Logon Scripts

Logon scripts are files that can be assigned to domain user accounts. Every time a user logs on, the assigned logon script runs. Logon scripts allow administrators to affect domain users' environments without managing every aspect of it. When a controller processes a logon request, the system locates the logon script that you specified for the user. See Section 3.1.2 for more information on assigning logon scripts.

If you use logon scripts in a domain that has a PDC and at least one BDC, you should replicate logon scripts among the domain controllers. Master copies of every logon script for a domain should be stored in the replication export directory on one server, which can be the PDC but does not need to be. Copies of the logon scripts should be replicated to every controller that participates in authenticating logons for the domain. If this is done, only one copy of each logon script will need to be maintained and every controller that participates in authenticating domain logons will have identical copies of every user logon script.

By default, the ASU server exports logon scripts from the `/usr/net/servers/lanman/shares/asu/repl/export/scripts` directory to the `/usr/net/servers/lanman/shares/asu/repl/import/scripts` directory on import computers.

4.2.6 Using Directory Replication

Suppose you manage an environment in which you have a domain that contains two directory trees that you want to replicate: one for logon scripts

and one for other data. The computers that need to import the two directory trees are different. Four domain controllers need the logon scripts; however, only two domain controllers and two Windows NT Servers need to import the other data.

The best solution is to set up two export servers: one for the scripts directory tree and one for the data directory tree. Remember that a single export server has only one list of import computers to which it replicates. If you set up a single export server for the two directories, it exports both directory trees to all import computers even though not all import computers use both directory trees.

4.2.7 Replication Troubleshooting Tips

Directory replication problems can have a variety of causes. When the `Replicator` service generates an error, it is displayed in the Event Viewer. The Event Viewer contains information about the Status column in the Manage Import Directories dialog box and information about messages that appear while you are configuring directory replication.

See Chapter 6 for more information about the Event Viewer.

4.2.7.1 Access Denied

If the Event Viewer shows `access denied` errors for the `Replicator` service, be sure that:

- The `Replicator` service on the export server and import computer is configured to log on using the same domain user account.
- The domain user account used by the import computer's `Replicator` service has permission to read the files on the export computer.

  The default permissions for an export directory grant FullControl to the Replicator local group. If FullControl permission is removed from the directory, exported files are copied to the import computers but receive the wrong permissions, and an access denied error is written to the event log. If necessary, use the Server Manager, select properties for the share that is associated with the export directory, select Permissions, and grant FullControl permission to the Replicator local group

4.2.7.2 Exporting to Specific Computers

Be sure to specify export servers and import computers in the To List and From List in the Directory Replication dialog box. If you fail to do so, exporting will occur to all import computers in the local domain, and importing will occur from all export servers in the local domain.

An export server has only one list of import computers to which it replicates. All import computers in that list receive all the exported directories and files. You cannot be selective about which import computer receives which export file or directory. For this level of control, you must use multiple export servers and configure them accordingly.

4.2.7.3 Lost Permissions in Import Directory

Do not use the Explorer or File Manager to examine permissions in the `/usr/net/servers/lanman/shares/asu/repl/import` directory. If you do, special permissions initially set there might be lost. These initial permissions enable directory replication to work; you do not need to change them.

4.2.7.4 Replication to a Domain Name Over a WAN Link

Directory replication to a domain name does not always succeed when some or all import computers are located across a wide area network (WAN) bridge from an export server. When adding names to the export To List on an export server, and when adding names to the import From List on an import computer, specify the computer names (instead of or in addition to specifying the domain name) for those computers separated by a WAN bridge.

4.3 Managing Disk Share Usage

Table 4–4 describes the disk share usage information that you can display by using the Properties option in the Server Manager or `net` commands.

Table 4–4: Disk Share Usage

Option	Description	net Command
Sessions	The users who are remotely connected to the computer and the shares to which they are connected.	# **net session**
Open Files	The number of shared resources opened on the computer.	# **net file**
File Locks	The number of file locks on open resources of the computer.	# **net file**
Open Named Pipes	The number of named pipes open on the computer. In some cases, a print job is monitored as an open named pipe.	# **net file**

See the ASU *Installation and Administration Guide* and the `net help` command for more information on `net` commands.

4.3.1 Disconnecting Users From Shares

You can use the Server Manager or the `net` commands to disconnect one or all users from a share. To prevent data loss, always warn users before disconnecting them. See Section 4.3.2 for more information on sending messages.

To use the `net` command to disconnect all users, enter the following command:

```
# net session /delete
```

Follow these steps to use the Server Manager disconnect one or all users from shares:

1. If you want to disconnect users who are connected to shares on a remote controller, choose Select Domain to and enter the domain for the controller. While you are administering another computer remotely, your user account is listed as a user connected to the IPC$ resource. It cannot be disconnected.
2. From the Computer menu, select Properties.

   A Properties dialog box is displayed.
3. Select:
   - Users to disconnect a user or all users from all shares to which they are connected.

     A Users Sessions dialog box is displayed. To disconnect a user, select the user in the list and select Disconnect. To disconnect all users, select Disconnect All.
   - Shares to disconnect a user or all users from a particular share to which they are connected.

     A Shared Resource dialog box is displayed. Select the name of the share in the Sharename window. To disconnect a user, select the user in the Connected Users window and select Disconnect. To disconnect all users, select Disconnect All.

4.3.2 Sending a Message to Users

You can send a message to all users who are connected to shares by using the Send Message command on the Computer menu in Server Manager. For example, you can do this before you disconnect one or more users or before you stop the `Server` service on a controller.

For a message to be sent and received, the `Messenger` service must be running on the computer sending the message and on the computers receiving the message.

See Sending a Message to Connected Users in Server Manager Help for more information about sending messages.

5

ASU Printer Shares

You can make printers hosted by Tru64 UNIX servers on which the ASU software is installed available to Windows users as printer shares.

Windows users can browse the network for printer shares and configure their Windows systems to use a printer share by using the Add Printer Wizard. Once configured, the printer that is associated with the printer share appears as a transparent extension to a user's local computing environment. For example, using an application such as Microsoft Word, users can print directly to the printer share, which forwards the jobs to the printer.

This chapter discusses the following topics:

- Planning your printing operations
- Print share properties

See the ASU *Installation and Administration Guide* for information on creating ASU printer shares.

5.1 Planning Your Printing Operations

You should ensure that network print operations are efficient and cost-effective. Among the decisions that you need to make include the following:

- Which printers to use
- Which computers to provide ASU printer shares
- How to configure printer shares for maximum use

5.1.1 Choosing Printers

Choosing print devices includes selecting from among devices that are designed specifically for network use. These devices offer options such as automatic port and emulation switching, dual paper bins, and double-sided printing. Before deciding on your network printers, consider the following questions:

- Do you need a few high-volume print devices or several less expensive personal print devices?

High-volume printers generally have more features but affect many more users if they break down.

- How many pages do you expect to print?

  You probably will experience fewer maintenance problems if you match printing volume with a printer's duty cycle (the number of pages that it can print per month).

- Which types of graphics support do you need?

  The combination of Windows NT and TrueType technology makes it possible to print sophisticated graphics and fonts on most printers, even those that normally support only bitmaps and text. TrueType fonts are integrated with the operating environment so that every Windows NT application can use them without changes or upgrades. If you intend to print many graphs, charts, or halftone photographs, consider a printer that supports 600 dpi or greater.

- How important is printing speed?

  Print devices can be attached to the network through serial or parallel ports on computers or directly attached to the network using built-in local area network (LAN) cards. Printers directly attached to the network usually provide faster throughput than parallel and serial buses. However, print throughput rates also depend on network traffic, the network interface card (NIC), the protocol, and the type of print device.

5.1.2 Choosing Computers to Be Print Servers

On a network of any size, you probably will concentrate your printer installations at a few servers. A computer acting as a print server can act as a file server or database server at the same time. File operations have insignificant impact on printers that are attached directly to the server.

A dedicated print server may be desirable if a server is required to manage many frequently-used printers. The decision to combine print and file servers may depend on security concerns. While printers always should be available to those persons using them, you may want to restrict physical access to file servers by keeping them in secured rooms.

No special hardware requirements exist for print servers except that they have appropriate printers (output ports) for parallel or serial print devices. Managing a large number of printers or many large documents requires an adequate amount of memory. The ASU server can control many network-interface printers, depending on the server's processing capability, the amount of installed memory, and the size and number of documents typically sent to the print server. To maintain high server throughput levels, increase memory as you add print devices. Disk space requirements are

minimal except in cases where large or numerous documents are likely to accumulate.

5.1.3 Planning How Users Access Printer Shares

Before creating a printer share, you need to be aware of ASU configuration options that can improve the flexibility and efficiency of network printing. With ASU, it is not necessary to have a one-to-one relationship between a printer share and a print devices. By associating printer shares and print devices in different ways, you can offer users flexibility in their printing operations. For example, you can configure the following:

- A single printer share to single print device.
- Multiple printer shares to a single print device.

  The capability to assign more than one printer share to a print device gives users flexibility in printing documents. For example, two printer shares that are associated with a single print device can offer different print properties: one may print separator pages and the other may not. Or, one printer share might hold documents and print them at night, while another processes documents 24 hours per day.
- A single printer share to multiple identical print devices (printing pool).

  There are no limit on the number of printers in a pool. Whichever print device is idle receives the next document. This configuration maximizes use of print devices while minimizing the amount of time users must wait for documents.

  All devices in the pool are the same hardware model and act as a single unit. Print share property settings apply to the whole pool. Printer ports can be of the same type or mixed (parallel, serial, and network).

5.1.4 Printer Drivers

Different hardware platforms and operating systems require different printer drivers. For example, to use a printer share created on an ASU server, a client running Windows NT on an Alpha computer requires the appropriate Alpha printer driver for that printer.

Printer drivers can be installed locally or can be provided by the ASU server. The ASU server stores printer drivers for Windows 95, Alpha, Power PC, MIPS, and x86-based clients in a disk share called PRINT$. The printer drivers are then available for clients to automatically download.

The ASU server determines whether incoming print requests are Alpha, Power PC, MIPS, or x86-based and automatically sends the appropriate driver to the client.

You must install the printer drivers on clients that are running Windows NT Version 3.1, 3.5, 3.51, or 4.0.

See the ASU *Installation and Administration Guide* for information on installing printer drivers.

Table 5–1 identifies settings that you can try for an unsupported print device. If your print device is not in this table, contact the manufacturer to determine if drivers are available.

Table 5–1: Unsupported Print Devices

For	Use
HPPCL (LaserJet) compatible	Hewlett-Packard LaserJet Plus
Color PostScript	QMS-ColorScript
35-font Plus font set or superset	Apple LaserWriter® Plus
9-pin dot matrix IBM compatible	IBM Proprinter
9-pin dot matrix Epson® compatible	Epson FX-80 for narrow or FX-100 for wide carriage
24-pin dot matrix IBM 24-pin compatible	IBM Proprinter X24
24-pin dot matrix Epson LQ compatible	Epson LQ-1500

5.2 Print Share Properties

The following sections describes printer share properties.

5.2.1 Separator Page

The ASU server prints a separator page or banner automatically before each print job. You can alter the default banner page and create one more appropriate to your needs.

You use the `net print` command to set or change a separator page.

See the ASU *Installation and Administration Guide* and the `net help` command for more information on `net` commands.

5.2.2 Using a Print Processor Script

A print processor script can send print jobs directly to a file or terminal instead of to a printer, or to a remote Tru64 UNIX system using the `uucp` command, or to another Tru64 UNIX system process, such as `troff` or `nroff`.

When you create a print processor script, you must share a queue that uses it to allow users to access it. Users access this queue as they would any other print queue. To avoid affecting service to other users, execute scripts in the background. Table 5–2 describes the environment variables that you can use in print processor scripts.

Table 5–2: Print Processor Script Environment Variables

Variable	Description
$CLIENT	The computer name from which the job was sent.
$COPIES	The number of copies to be printed (1 and above).
$PRIO	The Tru64 UNIX system `lpr` priority of the print job (1 to 39).
$DEST	The Tru64 UNIX system `lpr` printer class (server queue) to which the job was sent.
$FILENAME	The full path name of the file to be processed.

Follow these steps to create a print processor script:

1. Use a text editor to create a shell script and save it in the `lanman/customs` directory.
2. Make the script executable by using the `chmod +x` command.
3. Use the `net print /PROCESSOR:`*`pathname`* command to configure the printer share to use the print processor script.

See the ASU *Installation and Administration Guide* and the `net help` command for more information on `net` commands.

5.2.3 Scheduling and Spooling Settings

One way to maximize use of print devices is to stagger printing times. For example, if printer traffic is heavy during the day, you can postpone printing of less important documents by routing them through a printer share that prints only during off-hours. When you specify printing times, the print spooler accepts documents at any time but it does not print to the destination print device until the designated start time. At the stop printing time, the spooler stops sending documents to the print device and saves any documents remaining until it is scheduled to start printing again.

Table 5–3 describes the scheduling and spooling options that you set by using the Scheduling tab of the Properties sheet for a printer or `net` commands.

Table 5–3: Scheduling Options

Option	Description	net Command
Available	Defines when the printer is available.	# **net print /AFTER:time** # **net print /UNTIL:time:time**
Priority	Sets up a varied priority print queue based on document priority.	# **net print /PRIORITY:#** (1 is the highest and 9 the lowest)
Start printing after last page is spooled	Prevents delays when the print server prints pages faster than clients can provide them. (Default cannot be changed.)	No equivalent `net` command

See the ASU *Installation and Administration Guide* and the `net help` command for more information on `net` commands.

5.2.4 Controlling Access to Printer Shares

To control printer usage, set permissions for each printer share. By default, all of the shared printers that you create are available to all network users. To restrict access to a shared printer queue you must change its permission settings for a domain user account or group. Table 5–4 describes the permissions that you can set on printer shares. To change permissions on a printer share, you must be the owner of the printer or have been granted FullControl permission.

Table 5–4: Printer Share Permissions

Permission	Allows
No Access	No printing to the printer share
Print	Printing to the printer share

Table 5–4: Printer Share Permissions (cont.)

Permission	Allows
Manage Documents	Setting controls for documents and pausing, resuming, restarting, and deleting documents
Full Control	Print and Manage Documents permissions and: • Changing the printing order of documents • Pausing, resuming, or purging the printer • Changing print properties • Deleting the printer share • Changing permissions • Taking ownership By default, Administrators, Print Operators, and Server Operators have Full Control permissions.

Although permissions are cumulative, the No Access permission overrides all other permissions.

5.2.5 Auditing Printer Shares

You can audit a printer share to track its usage. For a particular printer, you can specify the actions to audit for domain user accounts and groups. You can audit both successful and failed actions. The ASU server stores the information generated from auditing in a log file that you can view by using the Event Viewer. See Chapter 6 for more information about the Event Viewer.

To audit a printer share, you can enable the spooler event logging and use the User Manager for Domains to set the audit policy.

Table 5–5 describes the printer share events that you can audit.

Table 5–5: Printer Share Audit Options

Option	Audits
Print	Printing documents
Full Control	• Changing job settings for documents • Pausing, restarting, moving, and deleting documents • Sharing a printer • Changing printer share properties
Delete	Deleting a printer share

Table 5–5: Printer Share Audit Options (cont.)

Option	Audits
Change Permissions	Changing printer share permissions
Take Ownership	Taking ownership

5.2.6 Custom Forms

Any user with FullControl permission can define a new form by using the server's Properties Forms property sheet. For example, you could create a form called Customer Receipt Form that uses letter-size paper and nonstandard margins. You also can create multiple forms with the same paper size or margins (or both) to meet specific user needs. For example, you can create forms that have unique names but the same paper size and image area (margins) to identify different departmental letterheads.

New form definitions are added to the print server's database and are stored per controller, not per printer. You assign forms to a specific print device and tray by using the printer's Properties Settings property sheet.

5.2.7 Setting Device-Specific Properties

Device-specific printer properties describe the physical configuration of a print device, such as which paper trays are loaded and how much memory a device has. These properties vary from device to device. When you create a printer share, default settings are used. Although default settings work for many printing needs, some special printing options, such as those available with PostScript printer drivers, require specific settings.

The following sections describe device-specific properties. To view or change device-specific properties select the printer icon in the Printers folder, select Properties on the Printer menu, and select the printer's Properties Device Settings tab.

5.2.7.1 Setting Printer Memory

Because page printers must store an entire page in memory, they require relatively large amounts of memory. If you are using a page printer, such as a laser printer, make sure that the amount of memory available in the device matches the value shown in the Device Settings tab. If the print device has substantially more or less memory than what is shown in the Device Settings tab, print throughput may suffer.

5.2.7.2 Using Print Forms

The ASU server uses a form-based printing model rather than a tray-based printing model. Under a form-based model, the Print Server administrator configures the ASU server by defining the form loaded in each paper source (tray). The form is defined using the following criteria:

- Size
- Printer area margins
- Form name

Using Windows based applications running on a Windows NT based computer, each user can select a desired print form. This frees the user from having to know which tray contains which form. The ASU server interprets tray and form assignment data and sends instructions to the print device to select the correct tray.

Windows based applications can use different forms within a document. For example, you can use Envelope for the first page, Letterhead for the second page, and Letter for the third and following pages.

5.2.7.3 Choosing Font Types

Fonts are collections of characters and symbols that have a specific design and resolution. Print devices use three types of fonts:

- Device fonts reside in the hardware of your print device. They can be built into the print device or can be provided by a font cartridge or font card.
- Screen fonts are Windows NT fonts (including TrueType) that can be translated for output to the print device. To install screen fonts on your Windows NT computer, use the Fonts option in the Control Panel folder.
- Downloadable soft fonts are installed using the Device Settings tab of the printer's Properties sheet. Clients that use soft fonts and that print to ASU printer shares should install soft fonts locally.

The ASU server supports three types of screen fonts that can be reproduced on printers:

- TrueType fonts are device-independent fonts that can be reproduced on all print devices. TrueType fonts are stored as outlines and can be scaled and rotated. To be reproduced on a print device, fonts only need to be present on the computer originating the document. The greatest benefit of using TrueType fonts in a networking environment is their portability; documents with TrueType fonts are independent of any print device, application, or system.

- Raster fonts are stored as bitmaps and are device-dependent. If a print device does not support raster fonts, it will not print them. Raster fonts cannot be scaled or rotated.
- Vector fonts are useful for devices such as pen plotters that cannot reproduce bitmaps. They can be scaled to any size or aspect ratio.

For each document, the client computer downloads required screen and soft fonts to the ASU server, which then sends them to the print device. To improve printing times, use device fonts which already are present at the print device.

Not all devices can use all three types of printer fonts. Pen plotters, for example, normally cannot use downloaded soft fonts or print raster screen fonts.

5.2.8 Setting Document Defaults

It is easy to confuse device-specific settings with document properties. Document properties do not rely on a device's physical settings. When applications create a new document, they often ask the printer for the default document settings.

Table 5–6 lists typical device-specific and document properties.

Table 5–6: Typical Settings

Device-Specific Properties	Document Properties
Color	Number of copies
Resolution	Page orientation
Memory	Two-sided printing
Font cartridge name	Collate copies
Form location	Form
Plotter pen	

To view Document Properties for a printer share, open the Printers folder, select the printer, and then select Document Defaults on the File menu.

Note

Document properties that are set from an application always override document defaults set in the printer's property sheets. However, if an application does not set a document property (such as page orientation or paper size), the print device defaults to the

document properties that were set in the printer's Document Properties sheets.

6

Monitoring Events

An event is any significant occurrence in the system (or in an application) that requires notification. Some critical events are noted in on-screen messages. An event that does not require immediate attention is noted in an audit entry in an event log file. An audit entry shows the activity that occurred, the user who performed the action, and the date and time of the activity. You can audit both successful and failed attempts. The audit trail can show who actually performed actions on the network and who tried to perform actions that are not permitted.

You can use the information in an event log to troubleshoot various hardware and software problems, and to monitor the ASU server for security events. You can view an event log by using the following:

- The Windows-based Event Viewer graphical user interface
- The Tru64 UNIX `elfread` command

This chapter describes how to monitor and view events by using the Event Viewer. See `elfread`(8) for more information about the `elfread` command.

This chapter discusses the following topics:

- Event Viewer overview
- Enabling auditing
- Logging events options
- Interpreting events
- Using Event Viewer
- Using event logs to troubleshoot problems

6.1 Event Viewer Overview

The ASU server records events and entries in the following types of logs:

- The system log contains events logged by ASU server system components. For example, the failure of an ASU service to start when the ASU starts is recorded in the system log. The types of events that are logged by system components are determined by the ASU server.
- The security log contains valid and invalid logon attempts and events related to resource use, such as creating, opening, or deleting files or

other objects. For example, if you use User Manager for Domains to enable logon and logoff auditing, attempts to log on to the system are recorded in the security log.

- The application log contains events logged by applications. For example, a database program might record a file error in the application log. Application developers decide which events to log.

System and application logs can be viewed by all users; security logs can be viewed only by system administrators.

6.2 Enabling Auditing

Event logging starts automatically when the ASU server starts; however, events are not audited by default. Administrators can use the User Manager for Domains to specify an Audit policy. The Audit policy determines the amount and type of events that are logged. Because the event logs are limited in size, carefully select the events to be audited and consider the amount of disk space you are willing to devote to the logs. The maximum size of the security log is defined in Event Viewer.

When you audit a file or folder, an entry is written to the Security log whenever the file or folder is accessed in a certain way. You determine which files and folders to audit, whose actions to audit, and exactly which types of actions are audited.

To audit a file or folder, use User Manager for Domains and enable auditing of File and Object Access, then use Explorer to specify which files to audit and which type of file access events to audit. Table 6–1 describes the directory and file actions that you can audit.

Table 6–1: Auditing Directories and Files

Auditing Directories	Auditing Files
Displaying names of files in the directory	Displaying file data
Displaying directory attributes	Displaying file attributes
Changing directory attributes	Displaying file owner and permissions
Creating subdirectories and files	Changing the file
Going to the directory's subdirectories	Changing file attributes
Displaying the directory's owner and permissions	Running the file
Deleting the directory	Deleting the file

Table 6–1: Auditing Directories and Files (cont.)

Auditing Directories	Auditing Files
Changing directory permissions	Changing file permissions
Changing directory ownership	Changing file ownership

6.3 Logging Events Options

Logging stops when an event log becomes full and cannot overwrite itself either because you set it for manual clearing or because the first event in the log is not old enough. When a log is full, you can free the log by clearing it.

Use the Log Settings command on the Log menu to define logging parameters for each type of log. You can set the maximum size of the log and specify whether the events are overwritten or stored for a certain period of time. Although you can increase (to the capacity of the disk and memory) or decrease the maximum log size, each log file has an initial maximum size of 512 KBytes. Before decreasing a log's size, you must clear the log.

The Event Log Wrapping option lets you define how events are retained in the log selected in the Change Settings For dialog box. (The default logging policy is to overwrite logs older than seven days.) You can customize this option for different logs. Table 6–2 describes event logging options.

Table 6–2: Event Logging Options

Use	To
Overwrite Events As Needed	Have new events continue to be written when the log is full. Each new event replaces the oldest event in the log. This option is a good choice for low-maintenance systems.
Overwrite Events Older Than [ ] days	Retain the log for the number of days you specify before overwriting events. This option is the best choice if you want to save log files weekly. This strategy minimizes the chance of losing important log entries and at the same time keeps log sizes reasonable.
Do Not Overwrite Events	Clear the log manually rather than automatically. Select this option only if you cannot afford to miss an event, for example, for the security log at a site where security is extremely important.

See To Manage the Audit Policy in User Manager for Domains Help for information on how to set the Audit policy.

6.4 Interpreting Events

Event logs consist of a header, a description of the event (based on the event type), and optionally additional data. Most security log entries consist of the header and a description.

The Event Viewer displays events from each log separately. Each line shows information about one event, including date, time, source, category, Event ID, user account, and computer name.

6.4.1 Event Header

Table 6–3 describes the contents of an event header.

Table 6–3: Event Header

Field	Displays
Date	The date the event occurred.
Time	The time the event occurred.
User	The username of the user on whose behalf the event occurred. If the event is not logged by a user, then the Security ID of the logging entity is displayed.
Computer	The name of the computer on which the event occurred.
Event ID	A number identifying the particular event type. The first line of the description usually contains the name of the event type. For example, 6005 is the ID of the event that occurs when the Event log service is started. The first line of the description of such an event is "The Event log service was started." The Event ID and the Source can be used by product support representatives to troubleshoot system problems.
Source	The software module that logged the event, which can be either an application name or a component of the system or of a large application, such as a service name.
Type	A classification of the event severity: Error, Information, or Warning in the system and application logs; Success Audit or Failure Audit in the security log. In the Event Viewer's normal list view, these are represented by symbols.
Category	A classification of the event by the event source. This information is used primarily in the security log. For example, for security audits, this corresponds to one of the event types for which success or failure auditing can be enabled in the User Manager for Domains Audit Policy dialog box.

6.4.2 Event Description

The format and contents of the event description vary, depending on the event type. The description is often the most useful information, indicating what happened or the significance of the event.

Table 6–4 describes types of events.

Table 6–4: Event Types

Event Type	Indicates
Error	Significant problems, such as a loss of data or loss of functions. For example, an Error event might be logged if an ASU service was not started when the ASU server started.
Warning	Events that are not necessarily significant but that indicate possible future problems. For example, a Warning event might be logged that the ASU server is low on key resources.
Information	Infrequent significant events that describe successful operations of major ASU server services. For example, when an ASU service starts successfully, it might log an Information event.
Success Audit	Audited security access attempts that were successful. For example, a user's successful attempt to log on to the system might be logged as a Success Audit event.
Failure Audit	Audited security access attempts that failed. For example, if a user tried to access a network drive and failed, the attempt might be logged as a Failure Audit event.

The optional data field, if used, contains binary data which can be displayed in bytes or words. This information is generated by the application that was the source of the event record. Because the data appears in hexadecimal format, its meaning can be interpreted only by someone who is familiar with the source application.

6.5 Using the Event Viewer

You determine which event log to view by switching between the system, security, and application logs. You also can use the Event Viewer to view logs on other computers.

6.5.1 Selecting a Log

Although the system log of the local computer appears the first time you start the Event Viewer, you can choose to view the security or application log. Use the Log menu to select a log for viewing.

6.5.2 Selecting a Computer

When you first start the Event Viewer, the events for the local computer appear.

To view events for another computer, click Select Computer on the Log menu. (It can be a Windows NT Workstation, an ASU server, a Windows NT server, or a LAN Manager 2.x server.)

If the computer you select is across a link with slow transmission rates, select Low Speed Connection. If this option is selected, the ASU server does not list all of the computers in the default domain, thereby minimizing network traffic across the link. (If slow transmission rates are normal, click Low Speed Connection on the Options menu.)

If you select a LAN Manager 2.x server for viewing, the Event Viewer can display its error (system) log and its audit (security) log.

See Select Computer in Event Viewer Help for information on how to select a computer for event viewing.

6.5.3 Refreshing the View

When you first open a log file, the Event Viewer displays the current information for that log. This information is not updated automatically. To see the latest events and to remove overwritten entries, choose the Refresh command.

See Refresh in Event Viewer Help for more information.

6.5.4 Changing the Font

You can change the font used in the Event Viewer. Changing this font affects only the display of the list of events in the main Event Viewer window.

See Changing the Font Selection in Event Viewer Help for more information.

6.5.5 Saving Log Files

You can save an event log in log-file format so that you can reopen it later in the Event Viewer. The log can also be saved in text format or comma-delimited text format so that you can use the information in other applications.

For example, you can save security logs so that you can monitor security events over a period of time. You can also save application logs so that you can track the Warning and Error events that occur for specific applications.

When you save a log file, the entire log is saved, regardless of any filtering options specified in the Event Viewer. If you changed the sort order in the Event Viewer, event records are saved exactly as displayed if you save the log in a text or comma-delimited text file.

A log file can be saved in the following formats:

- Log file format, which enables you to view the information again in the Event Viewer.
- Text file format, which enables you to use the information in an application, such as a word processor or electronic mail.
- Comma-delimited text file format, which enables you to use the information in an application, such as a spreadsheet or a flat-file database.

The binary event data is saved if you save a log in log file format but it is discarded if you save the log in text file format or in comma-delimited text file format. The event description is saved in all saved logs.

When you save a sorted log, the sort order affects the order in which event records are saved in a text file format or comma-delimited text file format. However, sort order does not affect the order of event records in a log saved in log file format. In either case, the sequence of data within each individual event record is record in the following order:

- Date (depends on the sort order specified on the View menu)
- Time
- Source
- Type
- Category
- Event
- User
- Computer
- Description

Saving a log file has no effect on the current contents of the active log. To clear the original log, you must select Clear All Events on the Log menu. To remove a saved log file, delete the file as you would other kinds of files.

You can view a saved file in the Event Viewer only if the log was saved in event log-file format. You cannot click the Refresh or Clear All Events commands to update the display or to clear a saved log.

Note

If you do not specify the correct log type (application, security, or system), the Description displayed for the saved log in the Event Detail dialog box will not be correct.

6.5.6 Viewing Specific Logged Events

After you select a log to view in the Event Viewer, you can do the following:

- View descriptions and additional details that the event source logs.
- Sort events from oldest to newest or from newest to oldest.
- Filter events so that only events with specific characteristics are displayed.
- Search for events based on specific characteristics or event descriptions.

6.5.6.1 Viewing Details About Events

For many events, you can view more information than is displayed in Event Viewer by double-clicking the event.

The Event Detail dialog box shows a text description of the selected event and any available binary data for the selected event. This information is generated by the application that was the source of the event record. Because the data appears in hexadecimal format, its meaning can be interpreted only by someone who is familiar with the source application. Not all events generate such data.

6.5.6.2 Sorting Events

By default, the Event Viewer lists events by date and time of occurrence from the newest to the oldest. To change the order from oldest to newest, click Oldest First on the View menu. If the Save Settings On Exit command on the Options menu is checked when you quit, the current sort order is used the next time you start the Event Viewer.

When a log is saved, the sort order affects the order in which event records are saved in a text format or comma-delimited text format file; sort order does not affect the order of event records saved in log file format.

6.5.6.3 Filtering Events

By default, the Event Viewer lists all events recorded in the selected log. To view a subset of events that have specific characteristics, click Filter Events on the View menu. When filtering is on, a check mark appears by the Filter command on the View menu and (Filtered) appears on the title bar. If Save Settings On Exit on the Options menu is checked when you quit the Event Viewer, the filters remain in effect the next time you start the Event Viewer.

Filtering has no effect on the actual contents of the log: it changes only the view. All events are logged continuously whether the filter is active or not. If you save a log from a filtered view, all records are saved even if you select a text format or comma-delimited text format file.

Table 6–5 describes the options available in the Filter dialog box.

Table 6–5: Event Filters

Filter	Filters
View From	Events after a specific date and time. By default, this is the date of the first event in the log file.
View Through	Events up to and including a specific date and time. By default, this is the date of the last event in the log file.
Information	Infrequent significant events that describe successful operations of major server services. For example, when a service starts successfully, it may log an Information event.
Warning	Events that are not necessarily significant but that indicate possible future problems. For example, a Warning event may be logged when the server is low on key resources.
Error	Significant problems, such as a loss of data or loss of functions. For example, an Error event may be logged if an ASU service was not started when the ASU server started.
Success Audit	Audited security access attempts that were successful. For example, a user's successful attempt to log on to the system may be logged as a Success Audit event.
Failure Audit	Audited security access attempts that failed. For example, if a user tried to access a network drive and failed, the attempt may be logged as a Failure Audit event.
Source	A source for logging events, such as an application, a system component, or a service.
Category	A classification of events defined by the source. For example, the security event categories are Logon and Logoff, Policy Change, Privilege Use, System Event, Object Access, Detailed Tracking, and Account Management.
User	A specific user that matches an actual user name. This field is not case-sensitive.

Table 6–5: Event Filters (cont.)

Filter	Filters
Computer	A specific computer that matches an actual computer name. This field is not case-sensitive.
Event ID	A specific number that corresponds to an actual event.

6.5.6.4 Searching for Events

To search for events that match a specific type, source, or category, click Find in the View menu. Searches can be useful when you are viewing large logs. For example, you can search for all Warning events related to a specific application or search for all Error events from all sources.

Your choices in the Find dialog box are in effect throughout the current session. If Save Settings On Exit on the Event Viewer Options menu is checked when you quit, the current filter settings are available the next time you start the Event Viewer.

6.6 Troubleshooting Using Event Logs

Careful monitoring of event logs can help you to predict and identify the sources of problems. Logs also can confirm problems with application software. If an application crashes, an application event log can provide a record of activity leading up to the event.

The following are guidelines for using event logs to identify problems:

- Save logs in log format. The binary data associated with an event is discarded if you save data in text or comma-delimited format.
- If a particular event seems related to system problems, try searching the event log to find other instances of the same event or to judge the frequency of an error.
- Note Event IDs. These numbers match a text description in a source message file. This number can be used by product-support representatives to understand what occurred in the system.

Index

E

G

H

L

M

P

R

S

T

U