Best Practice Active Directory Design for Managing Windows Networks

A structured approach to Active Directory design makes enterprise-scale directory service deployment straightforward and easy to understand. This guide combines business and technical guidance to minimize the time and effort required to implement the Active Directory directory service.

This guide provides a step-by-step methodology based on best practices learned from customers that have already deployed Active Directory in their organizations. It provides all the tasks and decisions you need to develop an Active Directory design to manage Windows networks. The intended audience for this guide is the IT professional responsible for testing, piloting, and rolling out an Active Directory design.

On This Page

Introduction

With the Active Directory service of Windows® 2000, organizations can simplify user and resource management while creating a scalable, secure, and manageable infrastructure for deploying additional important and emerging technologies.

To help shorten planning cycles and ensure successful deployments Microsoft is publishing a series of scenario-based guides that provide prescriptive, task-based, and solution-oriented guidance.

The Best Practice Active Directory Design for Managing Windows Networks and its companion guide, *Best Practice Active Directory Deployment for Managing Windows Networks,* are part of this series. These guides provide a structured approach to designing and deploying Active Directory. Without this structured approach, implementing Active Directory in your organization can take longer than expected.

These guides encapsulate planning and deployment expertise from Microsoft's product team with lessons learned from customers who have already designed and deployed Active Directory in their organizations.

Active Directory Deployment Scenarios

Unlike special-purpose directories, Active Directory can play a variety of roles within an organization. These roles range from managing Windows networks to supporting directory-enabled e-commerce applications. However, the way you intend to use Active Directory will affect the way that you make important design and deployment decisions.

Active Directory for Windows Network Management

This guide focuses on providing best practice–based guidance for deploying Active Directory for the purpose of managing networks comprised of Windows clients, Windows servers and Windows-compatible applications and devices. This guide will refer to this as the network operating system (NOS) management role. Benefits of deploying Active Directory in a NOS management role include:

- Centralized management of very large Windows networks (Active Directory is designed to support millions of objects).
- The ability to eliminate resource domains, including the hardware and administration they entail.
- Policy-based desktop lockdown and software distribution.
- The ability to delegate administrative control over resources where appropriate.
- Simplified location and use of shared resources.
- For additional information about the business value of deploying Active Directory visit http://www.microsoft.com/windows2000.
- This guide only covers deploying Active Directory and DNS core services as part of managing a Windows network. Other services that are layered on Active Directory can be added later and do not affect the initial design. For example, Group Policy can simplify management by providing policy-based administration for users, groups, workstations, and servers. Some services that can be layered on Active directory are:
- Group Policy

- Exchange 2000
- Integrated public key infrastructure (PKI) services
- Domain-based DFS

Special Considerations for Branch Office Deployments

Microsoft has identified a number of special considerations for deploying Active Directory in branch office environments. The characteristics of a branch office environment include:

- A large number of physical locations that need to contain replicas of Active Directory data.
- A small number of users per location.
- A hub and spoke network topology where many branch offices rely on connectivity to a centralized hub site for communications to other parts of the organization.
- Slow network connectivity between the branch office locations and the hub site.

Because of the ramifications of these requirements, Microsoft has developed additional content focused on deploying Active Directory in branch office environments. The *Active Directory Branch Office Planning Guide* is available on-line at http://www.microsoft.com/windows2000/techinfo/planning/activedirectory/branchoffice/default.asp. This content is designed to be used together with the *Best Practice Active Directory Design for Managing Windows Networks guide* as needed.

Special Considerations for Exchange 2000 Deployments

This guide will help you to design an Active Directory deployment that could host Exchange 2000. However, the information needed to successfully deploy Exchange 2000 as part of your Active Directory is not presented here.

For details see *Microsoft Exchange 2000 Server Upgrade Series* at http://www.microsoft.com/technet/prodtechnol/exchange/exchange2000/deploy/upgrdmigrate/ex2kupgr/default.asp.

Top of page

About this Guide

This guide was written for IT professionals who are going to participate in an Active Directory planning and deployment project. It provides a best practice approach to designing Active Directory that combines business and technical guidance to minimize the time and effort required to implement Active Directory in your organization. It contains worksheets throughout this document that will assist you in recording your design.

The Best Practice Approach

The best practice approach is based on the experience gained by the Active Directory development team from organizations that have successfully implemented Active Directory. This approach shortens planning cycles by:

- Promoting standard, tested Active Directory designs.
- Focusing on design choices that are proven to work well in the Windows NOS management role.
- Clarifying task milestones and the order of tasks with flowcharts and worksheets.
- Providing scenarios to reinforce design concepts.

The Scope of this Document

While the guidelines presented in this document are appropriate for almost all NOS management deployments, they have been tested and validated specifically for environments containing:

- Less than 100,000 users.
- Less than 200 physical locations.
- For deployments beyond these ranges, Microsoft suggests that you consider engaging the services of a consulting firm that has experience with deploying Active Directory in more complex environments.

How to Use This Guide

Each phase of the design process, such as creating a forest design, will include its own process flowchart and list of tasks that must be performed. Proceed through your Active Directory design in this order. Note that each step in the process may involve new team members who will be responsible for making decisions. Each member should record his or her design decisions in the worksheets provided. It may also be helpful to cut and paste from this guide into your design document so that new project members can understand previous design decisions.

You should not proceed with the next step or task until you have completed the worksheets provided and had all responsible individuals approve the design choices. Each worksheet builds on the information and decisions recorded in the previous worksheet. If you plan to use the best practices guide for Active Directory deployment, these worksheets are referred to during the deployment phase.

Before you proceed with the design, ensure that you have:

- Set and understand business goals for this Active Directory deployment.
- Executive-level sponsorship to implement an Active Directory–managed Windows network as designed.

Active Directory infrastructure deployment can span both technology and business areas. Therefore, your ability to progress on the design will depend on your ability to articulate the value of Active Directory to IT and business decision makers.

Who Should Read This Guide

This guide was written for IT professionals who are going to participate in an Active Directory planning and deployment project. This includes architects, project managers, system integrators, and consultants.

Because Active Directory is best deployed as a corporate-wide infrastructure, the design team will likely involve many people in your organization. This guide will make it clear what types of representatives are needed at various stages of the project. Project teams must gain the buy-in of these representatives for the design decisions that affect their part of the organization. For example, deploying Active Directory in most companies requires integration with an existing DNS infrastructure. The people who manage these systems will be critical to the success of the project. At the same time, it is important to keep teams as small as possible to make decisions easier to reach.

It is very important to note that deploying Active Directory in a Windows network management role should be driven at the corporate level — not at the departmental level. If you are a departmental administrator and want to deploy Active Directory, you should contact your corporate IT administrator for assistance. Failure to do so may make it difficult to join your departmental deployment to a corporate-level Active Directory deployment in the future.

Project Roles

While there will be many individuals involved in a typical Active Directory deployment, it is especially important to staff two roles early on: a project architect and a project manager. In a large organization, several individuals might share these roles. Table 1 summarizes the Active Directory design roles.

Table 1 Active Directory Design Roles

Design role	Responsibility	Description
Project architect	Technical design	Specialist or consultant responsible for technical decision-making and for ensuring that the design fulfills the organization's business goals.
Project manager	Process planning	Acts as the single point of contact to drive progress on the design by involving the appropriate people and garnering consensus. Responsible for all planning and scheduling to support the design.

Project Architect

Each forest requires an Active Directory architect to oversee the Active Directory design and migration process. An information technology (IT) architect or IT systems planner who has prior directory management experience would be a likely candidate. Otherwise, consider hiring a consultant who has experience with Active Directory design and deployment. Hiring a consultant brings important experience and perspective to the design team and may be particularly helpful in working through cross-organizational issues.

The Active Directory architect's responsibilities include:

- Owning the Active Directory design.
- Understanding the rationale for key design decisions.
- Determining if the organization's business goals are being met.
- Suggesting other solutions that might better reflect business needs, if necessary.

The final Active Directory design must reflect a combination of business goals and technical decisions. Therefore the project architect will review design decisions, compare them to business goals, and make sure that the two remain in alignment.

Project Manager

To promote an effective design process, management should nominate a person, or a small committee, to the role of project manager. The project manager facilitates cooperation across business units as well as with groups that manage technologies such as DNS, networks, and Windows NT. A project manager's responsibilities include:

- Providing basic project planning, such as scheduling and budgeting.
- Driving progress on the Active Directory design.
- Involving the right people during each part of the design process.
- Serving as single point of contact for the directory project.
- Garnering consensus among teams.

Top of page

Active Directory Design: Key Concepts

As you approach your design, it is important to note that you will be designing both a logical model and a physical model.

Logical Models

Active Directory allows administrators to organize elements of a network (such as users, computers, devices, and so on) into a hierarchical, tree-like structure based on the concept of containership. The top level Active Directory container is called a *forest*. Within forests, there are *domains*. Within domains there are organizational units (OUs). This is called the logical model because it is designed independently from most physical aspects of the deployment, such as the number of replicas required within each domain and network topology.

To facilitate the management of large numbers of objects, Active Directory also supports the concept of *administrative delegation* at the container level. Through delegation, owners can transfer authority over objects to other users or groups. Delegation is important because it helps to distribute the management of large numbers of objects across a number of people trusted to perform management of specific groups and types of objects.

For example, Figure 1 illustrates the distribution of users in a fictitious North America–based organization. In this example, a single Active Directory forest contains all of the users, computers, devices, and other entities within the organization. To support geographically based administration, the organization created five domains (West, Midwest, Northeast, Southeast, and Latin America) as first-level divisions of the forest. To support further delegation, the organization subdivided the West division into OUs, represented by the dashed lines.

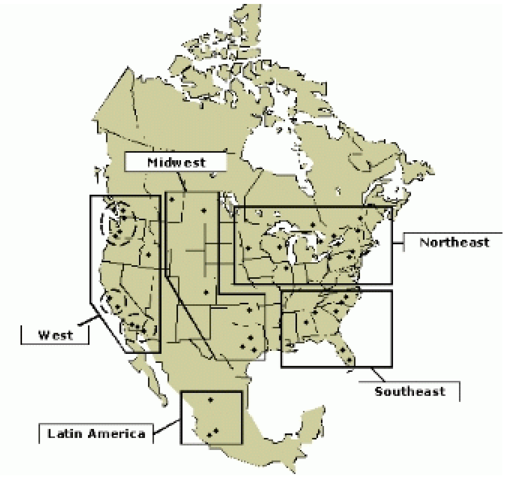

Figure 1: Delegation of administration within an organization

In this example, the organization has delegated some aspects of administration to the division manager of the West domain. The division manager of the West domain has in turn delegated some aspects of administration to its sub-division managers. In the same fashion, Active Directory supports a hierarchical structure that creates levels of administrative delegation for supporting the directory service and all forest objects.

As you design your logical model (following the step-by-step procedures later in the guide) you will essentially be deciding where to place forest, domain, and OU boundaries.

Physical Models

Once you have designed the logical model, the physical nature of the network will determine what additional tasks you need to perform. These tasks might include deciding where to place replicas of domain and global catalog data. You will also need to describe your network topology at the subnet level to Active Directory so that it can set up an optimized path for inter-local area network (LAN) communications, such as replication traffic.

You will want to pay particular attention to replication decisions because they impact both network traffic and scope of data visibility. For example, domain controllers do not replicate directory data between forests. Domain controllers hosting the global catalog contain a partial description of every object in the forest, and share this information forest-wide, but only with other domain controllers containing the global catalog. Within each domain, all data updates to objects within the domain replicate automatically to each of the domain's domain controllers, but not to domain controllers in other domains.

Again, this guide provides step-by-step guidance on all physical model decisions and procedures.

Additional Reading

This guide provides only a limited amount of background information on the concepts, technology, and terminology behind Active Directory. If you are not familiar with the Active Directory concepts presented in this guide, you should begin by reading and understanding the information contained in the following.

- Active Directory Overview, available on-line at:

  http://www.microsoft.com/windows2000/server/evaluation/features/dirlist.asp

- Active Directory Architecture, available on-line at:

  http://www.microsoft.com/windows

For further background information, we also recommend:

- The Active Directory Glossary, available on-line at:

  http://www.microsoft.com/windows2000/techinfo/howitworks/activedirectory/glossary.asp

- Active Directory Users, Computers, and Groups, is available on-line at:

  http://www.microsoft.com/technet/prodtechnol/windows2000serv/technologies/activedirectory/maintain/adusers.mspx

- The Directory Services section of the Windows 2000 Technical Library, is available on-line at:

  http://www.microsoft.com/windows2000/technologies/directory/default.asp

- Chapter 1, "Active Directory Logical Structure" in the Distributed Systems Guide of the Windows 2000 Server Resource Kit

- Chapter 9, "Designing the Active Directory Structure" in the Deployment Planning Guide of the Windows 2000 Server Resource Kit, is available on-line at:

  (http://www.microsoft.com/technet/prodtechnol/windows2000serv/reskit/deploy/dgbd_ads_heqs.asp)

Top of page

Part I: Determining the Number of Forests in Your Organization

Top of page

Determining the Number of Forests for Your Organization

The highest-level container in Active Directory is the forest. As a first step in the design process, the Active Directory architect and project manager must determine how many forests an organization requires. Because it is the simplest model to administer, you should strive for a single forest design for your organization. However, a single forest deployment has several constraints and not every organization can chose this model.

For example, individuals who currently manage the IT infrastructure for autonomous divisions within the organization may wish to assume the role of forest owner and proceed with their own forest design. However, in other situations, potential forest owners may choose to merge their autonomous divisions into a single forest to reduce the cost of designing and operating their own Active Directory or to facilitate resource sharing.

Part I will guide you through the issues and tradeoffs of deciding how many forests your deployment will need.

The Role of Forests in Windows Network Designs

A forest is a single instance of an Active Directory deployment and by definition is administratively autonomous from any other Active Directory deployment within the organization. In other words, as the highest level of ownership and control, the forest represents a complete Active Directory security and administrative boundary. Within this boundary there are also a number of shared elements. They include:

- Schema.

  The Active Directory schema contains information about every object class and all the attributes of every object class that can be stored in the forest. For example, every user account is an instantiation of the *user* class as defined in the schema. The schema description is stored in a portion of the directory database that replicates to every domain controller in the forest.

- Configuration data.

  An Active Directory deployment is described by a set of configuration objects that contain data defining infrastructure and topology elements such as domains, sites, and site links. The configuration data is also stored in a portion of the directory database that replicates to every domain controller in the forest.

- Global catalog of searchable directory objects.

  Administrators can designate domain controllers to hold, in addition to replicas of domain data, a partial copy of each object in the forest. The data is kept in a portion of the directory database called the *global catalog*. The global catalog concept allows fast, efficient object searches that span the entire forest. Every global catalog domain controller in the forest, regardless of the domain to which it belongs, will hold an identical copy of the global catalog.

- Trust relationships between all domains in the forest.

  Active Directory automatically creates transitive, two-way trust relationships between all domains in a forest. Any member computer can recognize and authorize access to any user or group from any domain in the forest.

Because forests can contain millions of objects, there are few technical reasons why the majority of organizations cannot deploy a single forest to meet their needs. However, depending on your organization's administrative model, you may need to deploy more than one forest. During the first phase of Active Directory design, the Active Directory architect will determine the number of forests your organization requires to satisfy the business goals of the organization and designate an owner for each new forest.

Elements of Determining the Number of Forests

The Active Directory architect and project manager are responsible for completing a forest plan for their organization. The plan should contain a list of forests to be designed for the organization with:

- Name of each forest owner.
- Scope of each forest.

Comparing Service and Data Ownership

Before you can understand the ownership roles presented in this guide, you should be aware of the differences between service and data ownership. Previous to Active Directory the Windows NT administrative model defined a domain as a group of objects managed by administrators. This guide introduces the concept of split IT administration. Two basic types of administrators exist for Windows 2000, data (or object) administrators and service (or directory) administrators.

The biggest advantage to the split administrative style is that individual business units do not need to staff and train personnel to support the directory service. Instead, individual business units can concentrate on their core business and only concern themselves with managing the data in the directory.

To manage the objects in the directory, data administrators must perform similar functions and have similar rights and permissions as your current Windows NT administrators. In Windows 2000, data administrators support the users and computers in the forest by performing functions such as:

- Adding and removing computers, users, groups, and OUs.
- Creating Group Policy settings.

Service administrators do not closely correspond to any distinct administrative role holder in the Windows NT domain model. Service administrators deliver the directory service, administer domains, own the domain controllers, and manage the configuration of the directory. For Windows 2000 Active Directory, service administrators support the forest infrastructure by performing functions such as:

- Creating or removing domains.
- Modifying the schema.
- Installing and removing domain controllers.
- Managing domain controller configuration.
- Monitoring domain controller health.
- Managing the site topology.

By separating data administration from service administration, an administrator can retain complete autonomy over all user and computers with a division while delegating service administration to another group. This distinction allows larger organizations to gain the benefits of a single central directory while retaining current business practices.

Throughout the remainder of the document, this guide distinguishes between service owners and data owners.

Analyzing the Forest Owner Role

The forest owner is ultimately responsible fore delivering directory services to the Windows network. The forest owner's responsibilities include:

- Ownership of the schema and configuration containers.
- Ownership of the data contained in the forest root domain.

Table 2 lists these and other roles assigned the forest owner and the responsibilities associated with each. The forest owner controls the forest through three security groups:

- Domain Admins of the forest root domain
- Enterprise Admins
- Schema Admins

Table 2 Roles and Responsibilities of the Forest Owner

Role	Responsibilites
Service owner for domain controllers in the forest root domain	Controls domain controller configuration throughout the forest to manage replication issues.
Administrative owner for data in the forest root domain	Controls membership of Domains Admins, Enterprise Admins, and Schema Admins security groups in the forest root domain.
Administrative oversight of all domain data	Through the Enterprise Admins group, the forest owner can correct errors anywhere in the directory. Although you can block Enterprise Admins administrative access to non-root domains, this is not the best practice.
Administrative owner of the schema	Through Schema Admins: · Sets policy for schema extensions · Sets process for schema extensions · Controls who extends the schema
Administrative (and policy) owner of the configuration	Through Enterprise Admins: · Acts as gatekeeper for new domains in the forest · Administrative owner of site topology

Note: The forest owner should always be a highly trusted individual. Therefore no valid reason exists for blocking the forest owner's access to data in non-root domains. Should a serious error occur in a domain's data, due to malicious or benign actions, the forest owner will need access to the domain to fix the problem.

Relationship between Service Owner Roles

Managing the directory and delegating administration of data throughout the forest creates new administrative roles in the best practice model. Figure 2 illustrates the lines of responsibility between IT administrative roles recommended as a best practice. These lines refer to the delegation of directory service responsibility and do not necessarily relate to an organization's reporting structure. For example, the DNS owner provides and configures DNS services as specified by the forest owner and possibly by each domain owner in the forest.

The forest owner delegates the management of individual domains to a group of domain owners, of site topology to a site topology owner, and of DNS service to a DNS owner. In turn, each domain owner delegates management of each OU to a group of OU owners and the DNS service to the DNS owner.

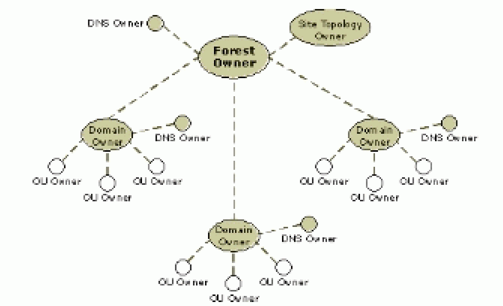

Figure 2: Hierarchy of directory service administrative roles for a single forest with four domains

Table 3 separates the major Active Directory Administrative roles into service ownership and data ownership roles.

Table 3 Active Directory Service and Data Ownership Roles

Service owners	Data owners
Forest owner Domain owners DNS owners Site topology owners	Forest owner (for root domain) OU owners

In a large forest, each owner role may be shared by a group of administrators, each with the necessary administrative authority to perform the requisite duties. In a smaller organization, one individual may perform several roles.

Note: The forest owner alone determines who becomes a domain owner. Domain owners have administrative control over the domain controllers in the forest and therefore must be highly trusted.

Best Practice Forest Models

The best forest design depends upon your organization's IT business practices. The three general forest models recommended for an organization are:

- Single, global forest model.

- Subscription model.
- Multiple forests corresponding to business units.

Single Forest Model

In the single forest model an organization chooses to contain and manage all of the directory objects within the organization in a single forest. A single forest is ideal when possible, because it represents the least possible administrative overhead. Figure 3 represents a single forest design for an organization with three divisions. In this Active Directory design, all three divisions become parts of a single forest.

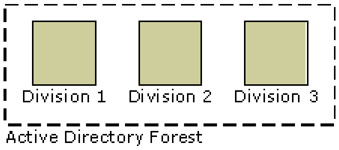

Figure 3: A single forest design for an organization with three divisions
Subscription Forest Design

In a subscription forest model there are some autonomous, division-level forests but most divisions are consolidated within a larger forest operated by a centralized IT organization. Subscription forests have the advantage of providing rapid deployment for divisions that are early adopters of Active Directory while allowing other business units that are concerned about security or sharing responsibilities to design and manage their own forests.

Figure 4 represents a subscription forest design for an organization with three business units. In this example, two business units chose to participate in a single forest while a third division created a second forest.

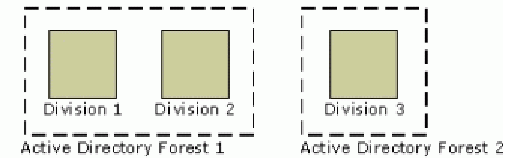

Figure 4: A hybrid forest model for an organization with three divisions
Multiple Forest Design

In a multiple forest model the majority of business units within an organization chose to deploy their own Active Directory forest. The multiple forest model works best for organizations with business units that want to (or need to) retain administrative autonomy. This model also has the highest possible administrative overhead.

Figure 5 represents a multiple forest design for an organization with three divisions. Each business unit historically and presently manages its own IT infrastructure. As a result, each business unit chose to design and deploy an independent Active Directory forest.

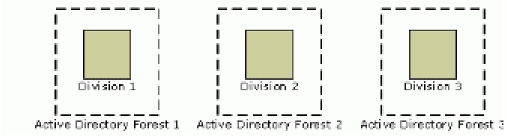
Division 1
Division 2
Division 3
Active Directory Forest 1
Active Directory Forest 2
Active Directory Forest 3

Figure 5: A multiple forest model resulting in three Active Directory forests

Forest Design Process

Figure 6 illustrates the forest design process. During the forest design process, the project manager and Active Directory architect will determine the number of forests to design for your organization. Because it represents the smallest possible management overhead, a single forest is ideal. However, due to several constraints, a single forest may not be possible and the goal becomes minimizing the total number of forests. To determine the number of forests needed:

1. Identify all candidate directory service providers.
2. Assign divisions to a forest based on technical and organizational factors.

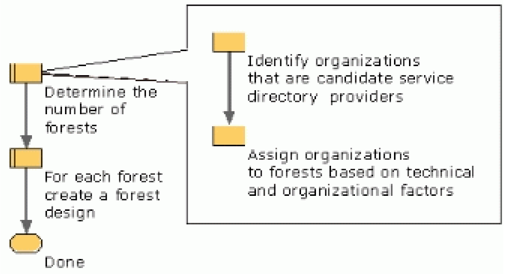
Determine the number of forests
For each forest create a forest design
Done
Identify organizations that are candidate service directory providers
Assign organizations to forests based on technical and organizational factors

Figure 6: Task flow for determining the number of forests in your organization

Identifying Candidate Directory Service Providers

Determine the highest level of IT administrators within your organization that have a mandate to deliver directory services for a Windows network. You may discover you have more than one entity in your organization capable of delivering directory services.

Your organization might have a centralized IT model with a chief information officer (CIO) who delegates IT administration to divisional administrators as needed. This situation indicates a single candidate directory service provider and therefore a single candidate forest.

In contrast, your organization might have a decentralized IT model, where division-level IT administrators perform network administration and delegation within their autonomous organizations. Peer CIOs commonly occur in organizations with distributed authority. This situation suggests a multiple candidate directory service providers and therefore multiple candidate forests.

Your organization may have already deployed Active Directory for a different purpose. Consider any current forest owners as candidate directory service providers. The existing directory may support messaging, Windows network services for limited number of business units, or some other purpose. Whether your existing Active Directory might become the foundation for expanded directory services depends upon its purpose and how it was deployed.

Assigning Forests

Determine the *minimum* number of forests that your organization requires to satisfy its business needs. To assign the correct number of forests to your organization:

1. If you have identified only one candidate directory service provider, then this individual, who will become your forest owner, should proceed to Part II: Creating a Forest Design.
2. If you identified multiple candidate directory service providers, then determine if you can designate one candidate as the forest owner with all other candidates as forest participants. The following section, "Participating in a Single or Subscription Forest, will help you make these decisions.
3. If not all candidate directory service providers can agree on a single forest, determine if a subscription forest will work for most of your candidate directory service providers.

   Note: Set a time by which you, the candidate directory service provider, will make your decision whether to join a forest. If creating a single forest requires that you change your current business practices and it is taking far longer than you had anticipated, you should probably create your own forest.
4. If your attempt to implement a subscription forest failed, then each candidate directory service provider should create a new, separate forest.

   The ramifications of making such a decision are detailed in the following section, "Creating Your Own Forest."

One candidate directory service provider may lack the interest, budget, or mandate to proceed with the deployment. In this case, this candidate directory service provider can opt out of deploying Active Directory at this time and can revisit deployment at a later date without blocking the remaining candidate directory service providers from proceeding with their designs.

Participating in a Single or Subscription Forest

A forest might already exist or a group within your organization may have offered to manage the directory for other divisions. Candidate directory service providers should become a forest participant if possible.

Choosing forest ownership or forest participation gives rise to different set of responsibilities for the candidate. The forest participant, as data owner, is responsible for managing the data, whereas the forest owner, as service owner, is responsible for maintaining the directory as a service. For example, the data owner ensures that a user's password is properly set, whereas the service owner ensures that the password replicates to all domain controllers. Table 4 provides a comparison of the responsibilities incurred when owning and participating in a forest.

Table 4 A Comparison of Owning and Participating in a Forest

If you own a forest	If you participate in a forest
Control service delivery. · Accountable for quality of service. · Delegate data management to individual data owners.	Outsource service delivery. · Manage only your data.

When subscribing to directory services provided by another group, you reduce your administrative costs by:

- Not initiating your own design process.
- Not duplicating directory management, such as:
  - Expert planning teams for directory deployment and operations.
  - Domain controller hardware management.
  - Directory structure management.

- Not isolating your users from users, resources, and services in other forests.

A shared forest works in many situations and greatly simplifies Active Directory management. Sharing a single, consistent configuration across all domains in the forest has specific advantages that include the following:

- Schema changes only need to be applied once to affect all domains.
- Configuration changes only need to be applied once to affect all domains.
- Single common global catalog for all users.
- Automatic trusts between all domains.

Your current business practices determine the value and feasibility of sharing a forest.

Some candidate directory service providers may opt to remove themselves from a shared forest design because they have distinct security requirements or have network and administrative constraints that are not addressed in the shared forest. Table 5 summarizes the characteristics present in most shared forests.

Table 5 Characteristics of a Shared Forest

Factor	Shared forest characteristics
Required	
Security	· You can trust the forest owner and all domain owners in the forest. · All domain controllers will be in secure locations or you can accept the risks involved in having domain controllers in unsecure locations.
Administration	· You can abide by common schema and configuration policies set out by the forest owner.
Name resolution	· DNS service can resolve names throughout the forest.
Recommended	
Networking	· No firewalls separating domain controllers. · No Network Address Translation (NAT) devices separating domain controllers.
Inter-organizational collaboration	· Domain trusts between organizations exist today. · Wide range of resources is currently shared among divisions.
IT organization	· Common IT infrastructure group. · If directory service management is outsourced, it is outsourced to a single provider only.

Important: Every domain controller in the forest contains a writable copy of the schema and configuration containers. If an individual can install software on or gain physical access to a domain controller the individual may be able to circumvent security features built into Windows 2000 and make changes to these containers. This represents a security risk to the forest. Best practice guidelines recommend that:

- All administrators such as domain administrators, who can install software on a domain controller, should be highly trusted.
- Domain controllers be kept in secure locations.

Lacking any of the required characteristics prohibits forest sharing. A potential forest participant can deviate from the recommended characteristics and still share a forest. However, if a potential forest participant lacks too many of these characteristics, then forest participation becomes technically difficult and less beneficial. Deviating from the recommended characteristics in Table 5 is not a best practice. If your design deviates, you should seek expert assistance for your design and deployment.

If you choose to join a forest, verify that the forest has a clear ownership. Sharing the ownership of a forest with another IT group or splitting the management across multiple outsourcers might present organizational challenges that you may wish to avoid.

Creating Your Own Forest

Although sharing a forest has a number of benefits, these benefits are proportional to the level of cooperation and collaboration between its divisions. You probably should not plan a single forest for your organization if:

- You require autonomy.

  You must have full control over service delivery or you cannot accept the terms offered by another forest owner.

- You do not collaborate with other divisions as a rule.

  If collaboration occurs, try to determine the nature of the collaboration. Collaboration can mean exchanging emails or something closer, such as accessing each other's intranet, accessing shares, and sharing applications. If divisions collaborate closely, then they should reside in the same forest.

  If collaboration is the exception rather than the rule, you can configure trusts on a case-by-case basis with domains in other forests. For example, you may currently restrict the level of collaboration between divisions with firewalls.

- Your previous infrastructure deployments have failed due to the diverse cultures or missions within your greater organization.

When putting together your forest design remember that administrative autonomy has a cost associated with it. The trade-off between autonomy and management efficiency should be carefully considered.

Note: Divestitures can create uncertainty about how to proceed with forest planning. Should you include the division that you plan to spin off in your forest or put the division in its own forest? The recommendation is that you create a separate forest for a division if you are absolutely sure that you are spinning off that division. Otherwise, include the division. Migrating the division to a new forest should the divestiture succeed requires the same level of effort as merging the division into the forest should the divestiture fail.

Implications of Interforest Collaboration

Because each forest is administered separately, adding additional forests to your design increases your organization's management overhead. When debating whether to create your own forest, weigh the ramifications listed in Table 6 against any benefits. If your division rarely collaborates with other business units, these factors may be relatively unimportant to you in making your decision.

Table 6 Implications of Interforest Collaboration

Change in functionality	Ramification
No automatic transitive trusts	To access resources in another forest, you must establish explicit trusts between the two domains involved.
No Kerberos authentication	Authentication between forests lacks delegation of authentication and mutual authentication.
No single global catalog of objects	In order to create a single view across multiple forests, you will need to set up directory synchronization software, such as Microsoft Metadirectory Services.Having multiple global catalogs may impact certain applications such as Exchange 2000.
Cannot log on with email-style user principal name (UPN)	If the machine account and user account are in different forests, must use old-style names to identify users.

Delegating the Next Steps to the Forest Owners

At this point in the design process, a list of forests and forest owners has been identified. Each forest owner can now proceed independently with the next phase in the design process "Part II: Creating the Forest Design." Each forest owner should proceed by designating a project manager and an Active Directory architect to oversee the forest design phase.

Figure 7 illustrates the tasks required to create a forest design. Each forest owner designated for the organization creates a separate Active Directory design.

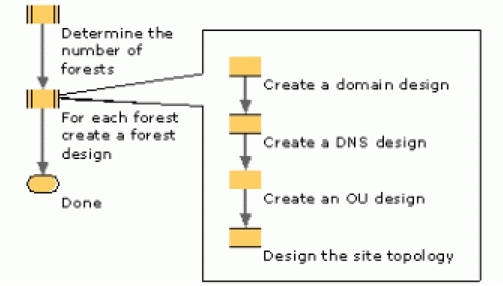

Figure 7: Task flow performed by each forest owner to design the forest

Top of page

Part II: Creating the Forest Design

Top of page

Creating a Domain Design

At this point in the Active Directory design process, a forest owner has been identified and is ready to proceed with the design. As the next step in the design process you will design the forest root domain and identify all the child domains of this forest root.

The Role of Domains in Windows Network Designs

A domain is a single partition of the Active Directory forest. Organizations partition the Active Directory forests into multiple domains to avoid replicating data to places where it is not needed. This allows the directory to scale globally over a network with limited available bandwidth.

For Active Directory, the domain continues to function as the administrative boundary for managing objects, such as users, groups, and computers. In addition, the domain supports a number of core functions related to administration. In the context of a network operating system, a domain's important functions include:

- Authentication.

  Each domain contains security principal objects, such as users, groups, and computers that can be granted or denied access to resources on the network. A user can only be authenticated by a domain controller in the domain that hosts the user's account.

- Policy-based Administration.

  Active Directory addresses the problem of standardizing the management of potentially hundreds of thousands of objects in a domain with policy-based administration. For example, you can use Group Policy to standardize user and computer configurations. Group Policy can be created and applied as part of a Windows 2000 migration or at any time after deploying Active Directory.

- Security policies for user accounts.

  A small set of security policies that apply to unique domain user accounts can only be set on a per-domain basis. These include:

  - Password policy
  - Account lockout policy
  - Kerberos ticket policy

- Serving as a directory for publishing shared resources.

  Active Directory provides a place for services to publish connection information about shared resources. For example, printer resources might be published in a domain to facilitate searches by users.

Contrasting Windows NT and Active Directory Domains

A typical organization running Windows NT has multiple domains. These domains exist to accommodate different business practices, geographical divisions, and Windows NT limitations. Three Windows NT product limitations have resulted in a large number of domains:

- Limited accounts database size.
- Reliance on the primary domain controller (PDC) alone for database updates.
- Inability to delegate administration within a domain.

To avoid these limitations, larger organizations typically deployed Windows NT as several master user domains (MUDs) for accounts and a large number of resource domains (RDs). Administrators within these domains created trusts to link users and resources throughout the organization. Figure 8 shows a typical organization's Windows NT master domain model structure.

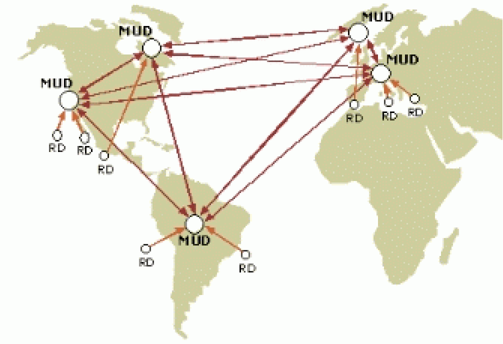

Figure 8: Windows NT master domain model with typical trust relationships
Windows 2000 Active Directory has none of these limitations. As a result, many domains created in Windows NT can be consolidated into a few Active Directory domains.

Elements of a Domain Design

The forest owner is responsible for completing a domain design for their organization. The design should contain the following:

- Name of the forest root domain.
- List of all domains with:

    - Name of the domain.
    - Scope of the domain.
    - Number of users in the domain.
- A list of existing Windows NT domains slated for upgrade or consolidation.

The Domain Owner Role

The domain owner provides directory service administration at the domain level. The domain owner controls the domain through the Domain Admins security group and other built-in security groups. Table 7 summarizes the domain owner's responsibilities and assigned tasks.

Table 7 Domain Owner's Roles and Responsibilities

Area of responsibility	**Associated tasks**
Managing domain controller operations	· Creating and removing domain controllers. · Domain controller health monitoring. · Managing services running on domain controllers. · Backing up and restoring.
Configuring domain-wide settings	· Creating domain and domain user account policies such as password, Kerberos, and account lockout.
Delegating data-level administration	· Creating OUs and delegating administration. · Repairing problems in the OU structure that OU owners do not have sufficient access rights to fix.
Managing External trusts	· Creating trust links with domains outside of the forest.

Caution: The level of trust required of domain owners demands that these individuals be carefully screened. In turn, the domain owner should tightly control the membership of all built-in and well-known administrative users and groups.

Best Practice Domain Design

The best practice domain design will vary with the number of users in your forest and the bandwidth available on your network to support domain controller replication. A best practice domain design includes:

- A dedicated root domain managed by the forest owner.
- A single domain or multiple geographically based domains as children of the root domain managed by domain owners designated by the forest owner.
- All domains are in native mode.

Note: During a domain's in-place upgrade from Windows NT to Windows 2000, some domain controllers will still be running Windows NT. To accommodate both Windows 2000 and Windows NT domain controllers on the same Windows network, Active Directory can function in mixed mode. In mixed mode certain features of Active Directory are disabled to allow interoperability with Windows NT backup domain controllers (BDCs). When all BDCs have been removed from the domain, then the domain can be switched to native mode, enabling the following Windows 2000 features:

- Universal groups
- Group nesting
- Domain local groups
- SID History

Of these features, domain local groups and SID History are necessary to support user and computer migrations from Windows NT domains.

To reap the maximum benefits from Active Directory, try to minimize the number of domains in the finished forest. For most organizations an ideal design is a forest root domain and one global domain.

The components of a best practice domain design are discussed in the following section.

Dedicated Forest Root Domain

The first domain created in a forest is automatically assigned the role of forest root domain. All other domains build upon the forest root domain to define the directory hierarchy. The forest root hosts the two special administrative groups tasked with supporting the Active Directory infrastructure: the Enterprise Admins and the Schema Admins.

The best practices approach to domain design dictates that the forest root domain be dedicated exclusively to administering the forest infrastructure. A dedicated forest root is recommended for the reasons explained in Table 8.

Table 8 Reasons for Including a Dedicated Forest Root in Your Design

Reason	Explanation
Fewer administrators can make forest-wide changes	Limiting the forest root domain administrative membership reduces the likelihood that an administrative error will impact the entire forest.
Easily replicated for forest backup	A small root domain can be easily replicated anywhere on your network to provide protection against geographically centered catastrophes.
Never becomes obsolete	You can never retire the root domain, even if your organization changes. A dedicated root domain never becomes obsolete because it functions solely as the forest root.
Ownership easily transferred	Transferring ownership of the root domain to transfer forest ownership does not involve migrating production data or resources.

The role of the forest root domain centers on defining and managing the infrastructure. Managing the directory infrastructure requires new administrative roles and responsibilities. Plan to reserve the dedicated root domain for forest administration exclusively. Avoid including any users or resources not dedicated to forest administration in the forest root domain.

Single, Global Child Domain

A single global child domain of the forest root domain has all users, computers, and group accounts in a single child domain, except those of directory administrators residing in the forest root. Figure 9 illustrates a forest with a single global child domain.

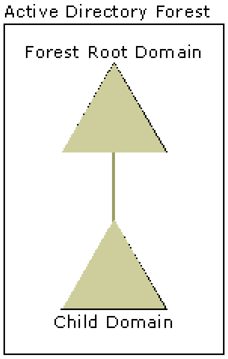

Figure 9: Forest with a root domain and a single global child domain
Geographically Based Domains

All object data corresponding to a domain is fully replicated to all domain controllers within the domain. For this reason, forests containing a large number of users distributed among widely-spaced locations with poor wide-area network (WAN) connections may require a geographically based set of child domains. If your organization cannot support a single, global child domain design, then geographically based domains are a best practice because they:

- Map well to network WAN connectivity.
- Map well to the global IT management groups in the typical organization.
- Are based on a stable structure.

Note: In some organizations divisions developed as autonomous units that currently provide their own IT infrastructure and do not coordinate services with any other divisions.

A forest owner, who is ultimately responsible for delivering services to the forest, is generally not willing to delegate responsibility for directory services to an IT administrator from an autonomous division. Furthermore, security considerations demand that all domain owners in a shared forest trust one another. The IT administrator from an autonomous division probably wants an independent domain to achieve some level of autonomy. Maintaining this level of autonomy can only be accomplished by creating a separate forest. For these reasons, organizationally based domains are not a best practice.

Figure 10 illustrates a forest containing geographically based child domains.

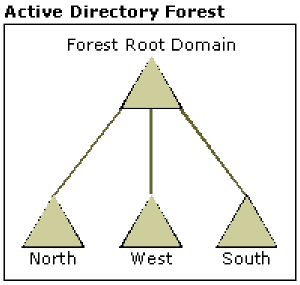

Figure 10: Forest with a dedicated forest root and three child domains

Domain Design Process

During the domain design phase, you will evaluate the existing IT environment and current administrative design to determine how best to combine current Windows NT domains into a minimum number of Active Directory domains. Deliverables include specifying:

1. The number of Active Directory domains in the design.
2. The scope of each new Active Directory domain.
3. A name for each domain.
4. Whether each Active Directory domain is created or upgraded from Windows NT.

The domain design process, which should help you to achieve one of the best practice domain models, includes the following steps:

1. Design the dedicated forest root domain.
2. Select a single global or regional domain model.
3. If a single, global domain is insufficient, divide the global organization into regional domains.
4. For each domain, specify whether you will upgrade an existing Windows NT domain or create a new domain.
5. Map all remaining domains to one of the regional domains for consolidation.

Figure 11 illustrates the tasks involved in completing the domain design.

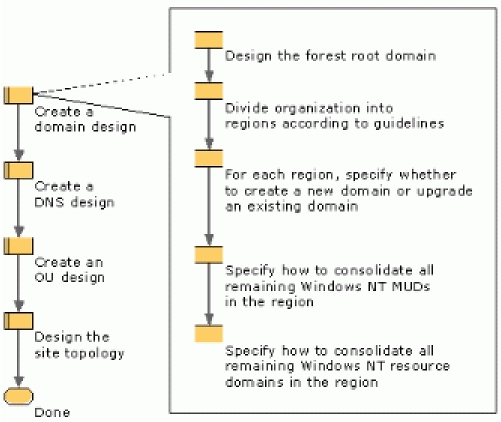
Create a domain design
Create a DNS design
Create an OU design
Design the site topology
Done
Design the forest root domain
Divide organization into regions according to guidelines
For each region, specify whether to create a new domain or upgrade an existing domain
Specify how to consolidate all remaining Windows NT MUDs in the region
Specify how to consolidate all remaining Windows NT resource domains in the region

Figure 11: Tasks performed for the domain design process

Designing the Forest Root Domain

To design the forest root domain you need to assign it a domain name. Active Directory domains have two types of names:

- Domain Name System (DNS) names used by Windows 2000 clients.
- NetBIOS names used by clients running earlier versions of Windows.

The forest root domain name is also the name of the forest. To name the forest root domain:

1. Identify your organization's DNS owner and determine what registered DNS names you have available on the network that will host Active Directory.

   Keep in mind that the names available on this network may be distinct from the names that your company exposes on the Internet. For example, the name your organization uses on the Internet might be contosopharma.com and the name used on your internal corporate network might be contoso.com. In this case the name that you select is contoso.com

If you do not have a registered domain name, you should register a name with an Internet DNS registration authority.

Note: As a best practice use DNS names registered with an Internet authority in the Active Directory namespace. Only registered names are guaranteed to be globally unique. If another organization later registers the same DNS domain name, or if your organization merges with, acquires, or is acquired by other company that uses the same DNS names then the two infrastructures can never interact with one another.

Add a prefix that is not currently in use to the registered DNS name to create a new subordinate name. For example, if your DNS root name were contoso.com then you should create an Active Directory forest root domain name such as concorp.contoso.com, where the namespace concorp.contoso.com is not already in use on the network. This new branch of the namespace will be dedicated to Active Directory and Windows 2000 and can easily be integrated with the existing DNS implementation. The rules for selecting a prefix are listed in Table 9.

Table 9 Rules for Creating a Prefix for an Active Directory Name

Prefix rule	Explanation
○ Not likely to become outdated	○ Domains cannot be renamed, so avoid names such as a business line or operating system that could become obsolete. Geographical names are recommended.
○ Internet standard characters only	○ A-Z, a-z, 0-9, and (-), but not entirely numeric.
○ 15 characters or less	○ If you choose a prefix length of 15 characters or less, then the NetBIOS name used by down-level clients remains the same as the prefix.

2. Arrange with the DNS owner to delegate ownership of that name to you.

   The DNS design section of this document discusses DNS names and DNS deployment planning in greater detail.

Selecting the Single Global or Regional Domain Model

Determine if your design will be a single, global domain or multiple, geographically based model. A single, global domain model is preferred for the following reasons:

- Users never need to be moved between domains.
- Duplicate Group Policy settings that span domains are not necessary.
- Any domain controller can process authentication for any user.

For some large organizations, the slowest network link connecting a domain controller may not be capable of handling the replication traffic of a single, global domain. In such a case, select the regional domain model. Refer to Table 10 to determine the maximum number of users that a single, global domain can accommodate. Consider growth in your organization when using these guidelines.

Table 10 Replication Sizing Guidelines for a Single, Global Domain

Slowest link connecting a domain controller (kbps)	Create a single, global domain no larger than (users)
9.6	20,000
14.4	30,000
19.2	40,000
28.8	50,000
38.4	75,000
56.0 (and higher)	100,000

Note: These estimates are for forests with up to 100,000 users. Larger forests are possible but not covered under these best practice guidelines. These values provide conservative estimates. Assumptions are:

- 10% of the minimum bandwidth is available to handle replication.

- All domain controllers (DCs) are global catalogs
- Incoming new user rate is 20% per year
- Outgoing user rate is 15% per year
- Group membership changes are a dominant factor in replication payload.
- 1:1 ratio of users to computers
- Directory service–integrated DNS is in use.
- DNS scavenging is in use

Use this list only as an approximation. Confirm that your network can handle your replication traffic by performing lab testing before implementing your domain plan. For additional information about domain sizing see *Building Enterprise Active Directory Services: Notes From the Field* by Microsoft Consulting Services, 2000, Redmond: Microsoft Press.

If you have already deployed multiple Windows NT domains, you will migrate those domains into the new global Active Directory domain. However, if you have already consolidated your Windows NT MUDs into less than ten geographically based domains, and then you may wish to upgrade those domains in-place rather than migrating your users to a single, global domain. Moving accounts between domains can impact end users. Before making your decision, evaluate the long-term administrative benefits of a single, global domain against the cost of migrating those users. This tradeoff is described in the next section, "Dividing the Global Organization into Regional Domains."

Dividing the Global Organization into Regional Domains

If size constraints require that you create regional domains for your forest, then you must:

1. Divide your organization into regions based on a regional feature that makes sense to your organization, such as:
    - Continent
    - Network connectivity
    - Administrative region

   An Active Directory domain will be created for each region. Your goal is to minimize the number of domains that you create. As a best practice, plan for no more than ten domains. When you consider adding domains in your forest design, remember that increasing the number of domains is a tradeoff between optimizing your replication bandwidth and minimizing your administrative costs. Table 11 lists these costs.

   Table 11 Management Effort Associated With Adding Domains to a Forest

Effort involved	Implications
○ Management of multiple Domain Admins groups	○ Because domain administrators have full control over a domain, the membership of the domain administrators group for a domain must be carefully controlled. Each added domain in a forest incurs this management overhead.
○ Additional domain controller hardware	○ A domain controller can host only a single domain. Each new domain that you create will require at least two computers to meet reliability and availability requirements. Because all domain controllers can accept and originate changes, you must physically guard them.
○ Trusts	○ For a user in one domain to access a resource in a different domain, domain controllers in both domains must be contacted. This communication represents an added possible point of failure. The fewer domains, the less your users will rely on communications with multiple domain controllers to maintain service.
○ Group Policy and access control management that is common to multiple domains	○ Group Policy and access control applied within a domain do not flow automatically into other domains. If you have Group Policy settings or delegated administration through access control that is uniform across many domains, they must be applied separately to each domain.
○ Increased likelihood of having to move objects between domains	○ The ease of restructuring within a domain favors a minimum number of large, stable domains in your domain design. As you add domains to the design, the likelihood increases that you will need to move an object, such as a user, or a group of objects from one domain to another. Moving a user within a domain is trivial. Moving a user between domains can impact the user.

2. For each region, use Table 12 to determine if the region meets the sizing guidelines then:
    a. Find the slowest network link connecting domain controllers, including links that connect this region with other regions.

b. Finding the row on the chart corresponding to the slowest link (rounding down).

c. If the number of users in the region exceeds the guideline, split the region into smaller regions

Consider growth in your organization when using these guidelines.

Table 12 Replication Sizing Guidelines for Regional Domains

Minimum bandwidth (kbps)	Create a forest no larger than (users)	Create a regional domain no larger than (users)
9.6	25,000	15,000
14.4	50,000	15,000
19.2	50,000	25,000
28.8	75,000	40,000
38.4	100,000	45,000
56.0 (and higher)	100,000	100,000

3. **Note:** These estimates are for forests with up to 100,000 users. Larger forests are possible but are not covered under theses best practice guidelines. These values provide conservative estimates. Assumptions are:

    - 10% of the minimum bandwidth is available to handle replication.
    - All DCs are global catalogs
    - Incoming new user rate is 20% per year
    - Outgoing user rate is 15% per year
    - Group membership changes are a dominant factor in replication payload.
    - 1:1 ratio of users to computers
    - Directory services–integrated DNS is in use.
    - DNS scavenging is in use

      Use this chart only as an approximation. Confirm that your network can handle your replication traffic by performing lab testing before implementing your domain plan. For additional information about domain sizing see: *Building Enterprise Active Directory Services: Notes From the Field* by Microsoft Consulting Services, 2000, Redmond: Microsoft Press.

4. If you have a domain with more than 50,000 users, specify a *dedicated PDC*.

    You will configure the dedicated PDC to handle PDC operations only. *Best Practices Active Directory Deployment for Managing Windows Networks* discusses configuring a dedicated PDC in further detail.

Figure 12 illustrates a regional domain design. In this example, there are a total of 60,000 users in the forest. With a minimum bandwidth of 32 kbps, these guidelines recommend no more than 50,000 users in one domain. This organization would like to implement a three-domain design reflecting geographic divisions. This design would create 40,000, 15,000, and 5,000 user domains, all of which fall within the domain replication guidelines.

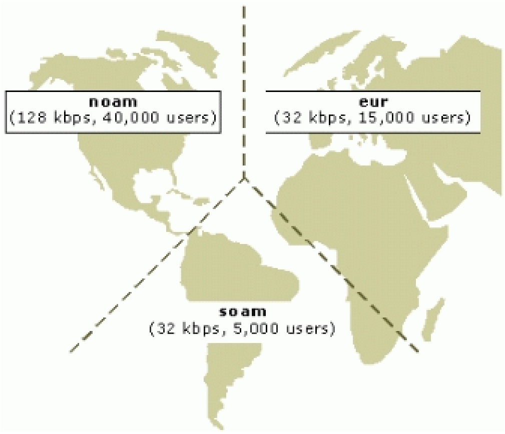

Figure 12: An organization's regional mapping of its users and computers

Specifying a New or Upgraded Domain for Each Region

Each region that you specified represents an Active Directory domain. These domains can arise from the creation of a new empty domain into which users and computers will migrate or by upgrading an existing Windows NT domain. If you are upgrading from Windows NT, several master user domains (MUDs) will likely already exist.

Use the following steps to specify a new or upgraded domain for each region

1. If you are migrating from Windows NT, evaluate each MUD in the region and determine if any of the MUDs are upgrade candidates to become the regional domain.

   More information about this topic is provided in the following section, "Evaluating Current Master User Domains in a Region.

2. If you are migrating from an operating system other than Windows NT or if no Windows NT MUD upgrade candidates exist, plan to create a new domain for the region.

   More information about this topic is provided in the section, "Designing New Regional Domains" later in this part.

Evaluating Current Master User Domains in a Region

Important: This topic was created for organizations migrating from Windows NT. If your organization uses another operating system, you can skip this topic.

If your organization is migrating from Windows NT, your users and resources currently reside in the Windows NT domains. In this phase of the design the forest owner will determine how best to migrate and consolidate these Windows NT domains into a finished Active Directory design.

Before beginning, create a master list of MUDs in this region that will join the forest. For each of these MUDs, complete a master user domain worksheet. For each MUD, you will ultimately decide if you want to:

- Upgrade the MUD in-place, meaning upgrading all of its domain controllers to Windows 2000 Active Directory.
- Retire the MUD and migrate its user accounts to a new or upgraded domain.

Table 13 will help you to determine whether to upgrade a MUD or migrate its contents.

Table 13 Reasons For and Against Upgrading a MUD In-place

Upgrade in-place because	Create a new domain and migrate because
· Current domain is a good fit for the regional model. · Faster than migrating users to a new domain.	· Domain spans regions and its contents need to be split among other regional domains. · It will take too long to upgrade the domain to native mode (all DCs upgraded).

Important: Keep in mind that if you choose to upgrade a Windows NT domain as your regional domain, this domain cannot act as a target domain for domain consolidation until it is in native mode. Getting to Active Directory native mode might take longer than expected if there are many BDCs in remote locations or needing hardware upgrades. In this case you might wish to choose another Windows NT domain to upgrade or build a new domain as the consolidation target.

Ideally, you will choose one MUD from each region to upgrade to Windows 2000. However, you may not have a good upgrade candidate in a region or may identify more than one upgrade candidate. If you do not have a good upgrade candidate, you can create a new domain. This is covered in the next section.

If you identified more than one MUD that is an upgrade candidate, you can either:

- Upgrade the largest MUD and consolidate the remaining MUDs into it.
- Define additional regions and upgrade each MUD in-place.

  If you decide to define additional regions, remember to consider the additional management effort associated with adding domains to the forest.

For each domain that the forest owner chooses to in-place upgrade as part of the Active Directory design:

1. Assign a domain a DNS name.

   Assign <domain_prefix>.<forest name> as the domain name. As a default, Windows 2000 will use the current NetBIOS name as the domain prefix. You can change the domain prefix during the upgrade, but as a best practice keep the default name. This ensures that when you are viewing the domain name in either Windows 2000 or earlier tools, you will see the same name.

2. Assign someone to the domain owner role in the new domain.
3. Ensure that the current domain owner agrees to the plan.

Designing New Regional Domains

The forest owner will need to create new domains as part of the domain design if:

- The organization does not currently have Windows NT domains.
- None of the existing domains are suitable for upgrade.

Once new regional domains have been specified, distribute accounts and resources among these domains.

For each new domain that the forest owner chooses to add to the Windows 2000 design:

1. Assign the domain a DNS name.

   Use <*new name*>.<*forest name*> as the domain name. Review the rules for creating a naming prefix listed in Table 9.

2. Assign someone to the domain owner role in the new domain.

Example of Specifying New or Upgraded Domains

In Figure 13, the domain design calls for a MUD from both the North American and South American regions to be in-place upgraded to the position of Windows 2000 regional domain. In Europe, both MUDs have numerous domain controllers in remote locations. The forest owner determined that in either case an in-place upgrade would take too long to reach native mode. Therefore, a new domain was specified in the domain design. The three remaining MUDs will be consolidated into the regional domain in their areas.

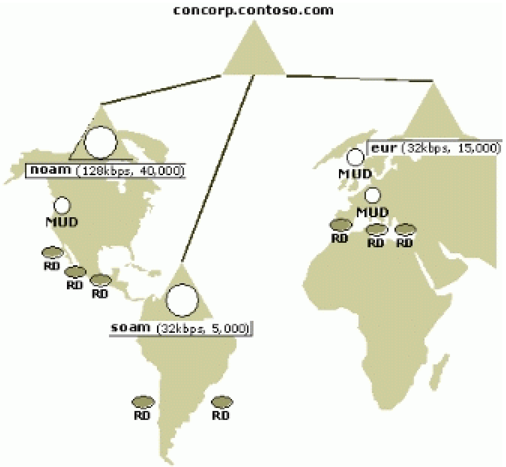

Figure 13: Creating Active Directory domains by either upgrading Windows NT MUDS (triangle with a circle) or by creating a new domain (triangle only) Consolidating Remaining Windows NT Account Domains

You should consolidate all remaining MUDs with their respective regional domains. Keep in mind that when you consolidate a MUD into another domain that:

- You can configure OUs to preserve the original domains organization and administration.

  OU design is covered in a later section of this guide.

- The consolidation process has a one-time impact on those users being migrated.

  The user migration process is covered in detail in the companion guide, *Best Practice Active Directory Deployment for Managing Windows Networks*.

Note: If you plan to split the contents of a Windows NT domain across multiple regional domains, as a best practice you should migrate the contents into other target domains and not upgrade this domain in-place.

Create a mapping of all remaining MUDs to regional domains. Figure 14 illustrates how all the remaining Windows NT MUDs are consolidated into the Active Directory regional domains for concorp.contoso.com.

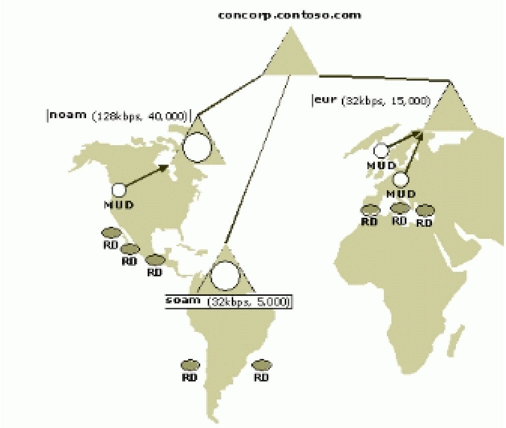

Figure 14: Consolidating Windows NT MUDs into Active Directory domains
Consolidating Windows NT Resource Domains

As a best practice you should plan to consolidate all Windows NT resource domains into the Active Directory domain in its region.

Determining to which domain the contents of the resource domain will move is often straightforward. Because users want their data in close proximity, most resource domains are geographically organized to correspond to users in a particular location. This means that most resource domains can readily be consolidated into the Active Directory domains that contain user accounts from that location. In cases where the contents of the Windows NT resource domain span geographies, specify which resources should be migrated to which target domains. Keep in mind that when you consolidate this MUD into a target Windows 2000 domain that you can configure OUs to preserve the original domain's organization and administration.

In other cases, you might preserve the domain as a Windows NT resource domain until a later date when you consolidate the domain or eventually eliminate it through attrition. This can occur if there are applications running on domain controllers in the domain that cannot run on Windows 2000.

Note: If you have a resource domain dedicated to hosting Exchange 5.5, leave that domain running Windows NT for the time being. To plan your Exchange migration, seethe *Microsoft Exchange 2000 Server Upgrade Series,* available on-line at http://www.microsoft.com/technet/prodtechnol/exchange/exchange2000/deploy/upgrdmigrate/ex2kupgr/default.asp.

Figure 15 illustrates how the forest owner might choose to consolidate existing resource domains into Active Directory domains in the same region.

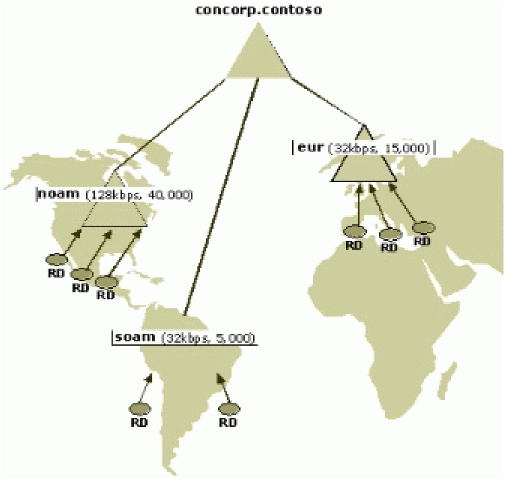

Figure 15: Consolidating Windows NT resource domains into Windows 2000 domains
In this scenario, all resource domains were organized regionally. This made it easy to determine where in the regional domain hierarchy to consolidate each resource domain.

Figure 16 illustrates the final Active Directory forest structure for concorp.contoso.com. The final design consists of a forest root domain plus three regional domains.

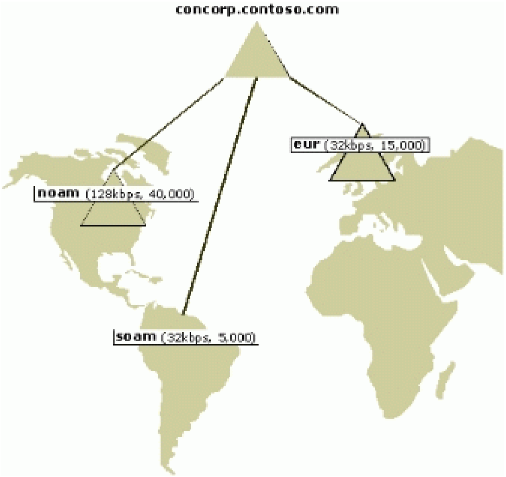

Figure 16: Final Active Directory domain design for concorp.contoso.com

Taking the Next Step

During the domain design process, each forest owner specified a forest root domain name and the logical structure for Active Directory. Since Active Directory relies on DNS for its name resolution, the forest owner appoints a DNS owner for Active Directory and specifies the DNS design as the next step.

If the organization has already implemented a DNS service, the forest owner and DNS owner for Active Directory work with the current DNS owner to delegate the forest root DNS name to Active Directory. If no DNS service has been implemented, the forest owner and the DNS owner for Active Directory must plan to implement one.

The next topic covers the details of this process.

Top of page

Creating a DNS Design for Active Directory

For the forest owner and Windows NT administrators DNS might be a new concept. DNS provides a translation of names to IP addresses. Network devices use IP addresses to locate and connect to hosts. However, users find remembering IP addresses difficult and typically prefer friendly names such as ftp.contoso.com. DNS enables users to enter hierarchical, friendly names to connect to computers and other resources on IP networks.

For Windows 2000, DNS represents a better-scaling, long-term replacement for NetBIOS name resolution, which is performed by the Windows Internet Name Service (WINS) in Windows NT–based networks. However, you still need to run WINS to support clients running versions of Windows earlier than Windows 2000, because such clients do not understand how to use DNS for domain location.

This design document explains how to create a design based on the choices you made in the "Domain Planning" section earlier in this document. This guide shows best practices for designing a DNS service for Active Directory on a network with:

- No existing DNS service
- Existing DNS service

The Role of DNS in Windows Network Designs

To understand how Windows 2000 and Active Directory use DNS, it is important to understand what DNS is and how it works.

DNS is a distributed database that contains all the information needed for any client to look up any name. The database is a hierarchical structure represented as an inverted tree. This tree is also known as a *namespace*. The Internet represents one kind of namespace. Any client in the Internet namespace can look up any name in the Internet namespace. Some companies also have a private or internal namespace.

DNS is designed so that any DNS server can answer queries about any name within its namespace. A DNS server has three options for answering queries:

- If the answer is in its cache, it answers the query from the cache.
- If the answer is in a zone hosted by the DNS server, it answers the query from its zone.

  A *zone* is a portion of the DNS tree containing all the records needed to answer queries for names within a zone. When a DNS server hosts a zone, it is *authoritative* for that zone and can answer queries for any name in the zone. For example, a server hosting the zone contoso.com can answer queries for any name in that zone.
- If the server cannot answer the query from its cache or zones, it queries other servers.

For a DNS root server to be able to successfully answer queries about any name, it must have a direct or indirect path to every server and zone in the namespace. Figure 17 shows how the DNS root server knows how to contact other servers in the namespace.

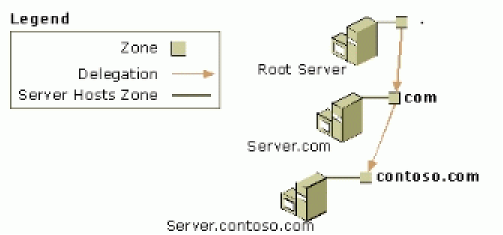

Figure 17: Delegation example.
The DNS root server hosts the root zone, represented as a dot (.). The root zone contains a delegation to a zone in the next level of the hierarchy, the com zone. A *delegation* is a record in the parent zone that lists the name server authoritative for the delegated zone. In this example, the authoritative name server is Server.com. The delegation in the root zone tells the DNS root server that, to find the com zone, it must contact Server.com. Likewise, the delegation in the com zone tells Server.com that to find the contoso.com zone, it must contact Server.contoso.com.

Note: A delegation uses two types of records. The name server (NS) resource record gives the name of an authoritative server. The host (A) resource record gives the IP address of an authoritative server.

In this manner, the DNS root server can find any name in the namespace. The root zone has delegations that lead directly or indirectly to all other zones. Any server that can query the DNS root server can use the information in the delegations to find any name in the namespace. DNS servers have *root hints,* a list of names and IP addresses, which tell them how to query the DNS root servers.

In some configurations, some servers do not have root hints. Instead, they forward all queries they cannot answer to another server.

Recursive Name Resolution

Clients depend on DNS servers to answer all name resolution queries. In a process called *recursive name resolution*, clients use a recursive query to ask servers to resolve names and to specify that the server must give a definitive answer, whether or not the name exists. When a server cannot answer authoritatively, it sends the query to another DNS server. The DNS server determines to which server the query goes in one of two ways: using root hints or by forwarding.

Resolving Names with Root Hints

Figure 18 illustrates how DNS resolves a name using root hints.

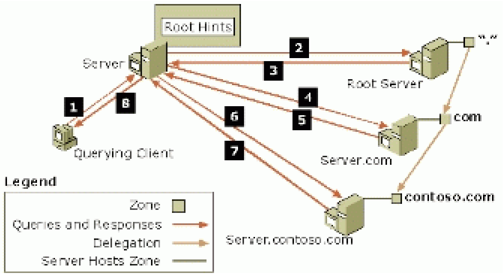

Figure 18: Recursive name resolution with root hints example
In this example, (1) the querying client sends a recursive query to a DNS server for the name ftp.contoso.com. Since the DNS service is not authoritative for the name and does not have the answer in its cache, the DNS server checks its root hints to find the IP address of the DNS root server, then (2) uses an iterative query to ask the DNS root server to resolve the name ftp.contoso.com. With an *iterative query*, the server is not necessarily asking for a definitive answer. Instead, the server accepts a pointer to another server that might have a better answer. Because the name ftp.contoso.com ends with the label com, (3) the DNS root server returns a delegation to the com zone. (4,5)The DNS server then continues making iterative queries until (6) it reaches Server.contoso.com, which finds the answer in its zone files and (7) responds with a definitive answer. (8) The server then returns the result to the client.

Resolving Names by Forwarding

Forwarding enables you to route name resolution through certain servers instead of using root hints. It is an administrative choice and not necessary for support of Active Directory. If you have an existing DNS service, someone has already made this choice. Figure 19 shows an example of a forwarded query.

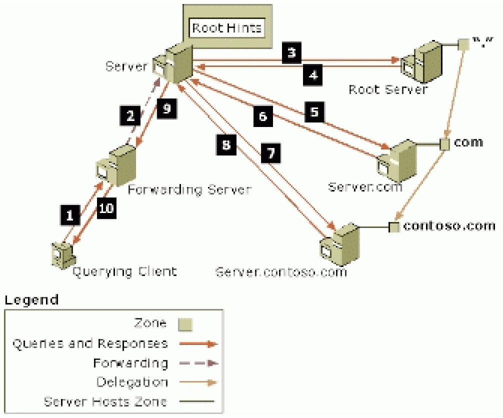

Figure 19: Forwarding example.
This example is the same as the previous example, except that the client queries a server that is configured to forward queries. When (1) the client queries for the name client.contoso.com, (2) the forwarding server forwards that query to another server, known as a *forwarder*. The remaining steps are the same as in the previous figure.

Forward and Reverse Lookups

If a client asks a server to give it an IP address that corresponds to a given name, it is asking the server to perform a *forward lookup*. The answer is located in a *forward lookup zone*.

On the other hand, if a client asks a server to give it a name that corresponds to a given IP address, it is asking the server to perform a *reverse lookup*. The answer is located in a *reverse lookup zone*. Computer IP addresses are stored in records called *pointer* (PTR) records.

Reverse lookup capability is not required for the proper operation of Active Directory. If you have an existing DNS namespace with existing reverse lookup zones, the existing DNS administrator can continue maintaining those zones. If you do not have an existing DNS namespace, you do not need to create reverse lookup zones in order to deploy Active Directory.

Locating Active Directory Domain Controllers

Clients communicate with domain controllers for such operations as processing logon requests and searching the directory for published resources, like printers. Figure 20 shows how clients use DNS to locate Active Directory domain controllers. In this example, (1) the domain controller registers in DNS. (2) The client queries a DNS server for the names of the domain controllers in the domain. The DNS server performs recursive name resolution to find the names of domain controllers, and then returns the names and IP addresses to the client. (3) The client uses the IP address to contact the domain controller.

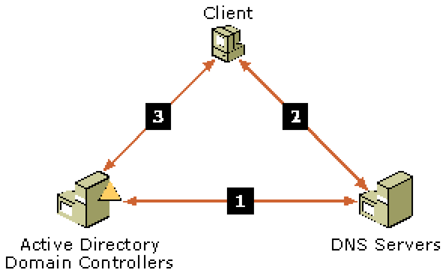

Figure 20: How clients locate domain controllers
Clients need the IP address of a DNS server in order to locate a domain controller. Domain controllers register a variety of records in DNS to help clients and other computers locate them. These records are called *locator records*.

Windows 2000 Active Directory Integrated DNS

The Windows 2000 DNS server can store its zones in Active Directory. Active Directory integration has the following advantages:

- Creates multiple masters for DNS replication, meaning:
    - Any DNS server can accept updates for that zone.
    - Requires no separate DNS replication topology.
- Supports secure dynamic updates.

  Secure dynamic update allows an administrator to precisely control which computers update which names, and prevents unauthorized computers from overwriting existing names from DNS.
- Supports out-of-date record scavenging.

For more information about DNS, consult one of the sources listed in Table 14.

Table 14 More Information about DNS

Resource	General DNS Information	Windows 2000 DNS Information
"Introduction to DNS" in the *TCP/IP Core Networking Guide* of the *Windows 2000 Server Resource Kit*		
"Windows 2000 DNS" in the *Networking Guide* of the *Windows 2000 Server Resource Kit*		
DNS and BIND, 3d ed., by Paul Albitz and Cricket Liu, 1998, Sebastopol, CA: O'Reilly & Associates		

RFCs 1034 and 1035.		
Windows 2000 Server Help.		

For more information about RFCs 1034 and 1035, see the Request for Comments link on the Web Resources page at http://www.microsoft.com/windows/reskits/webresources/default.asp.

Computer Naming

Windows 2000 introduces the concept of a primary DNS suffix to a computer name. When a Windows 2000 computer joins a domain, the computer by default registers itself with a name comprised of the host name of the computer and the DNS name of the domain the computer has joined (<computer_name>.<primary_suffix>). This fully qualified DNS name for the computer is known as the *primary* name.

A computer may already have a DNS name that was statically entered into a DNS server or registered by an integrated DNS/DHCP service. The primary name of the computer is distinct from these previously registered names.

Elements of the DNS Design

The forest owner and DNS owner are responsible for completing a DNS design for their organization. The plan should contain:

- DNS server configuration including:
    - DNS server placement.
    - Zone placement and type.
    - Recursive name resolution method.
- DNS client configuration including:
    - Computer naming scheme.
    - Resolver configuration.

The DNS Owner Role

Each domain should have a DNS owner who reports to the domain owner. The DNS owner is responsible for completing the DNS design for Windows 2000.

If your organization is migrating from Windows NT, then you currently have an individual who supports WINS name resolution. The DNS owner supports DNS in a manner analogous to your current WINS support.

The DNS owner is responsible for maintaining contact with the organization's DHCP and DNS groups. Involve these groups in the Active Directory DNS design process so that each group is aware of what your team is planning and can input early in your design.

Best Practice DNS Design with No Existing DNS Service

The best practice model for DNS depends on whether DNS has already been implemented on the network. When no DNS service has been implemented, then the forest owner and DNS owner for Active Directory should plan to implement one as specified in best practices.

The companion guide, *Best Practice Active Directory Deployment for Managing Windows Networks*, includes detailed instructions for deploying DNS. The information provided here is for reference only.

Server Configuration

Create DNS domains that correspond to the Active Directory domains you have already specified. Figure 21 illustrates one possible DNS design for an organization that has not previously implemented DNS.

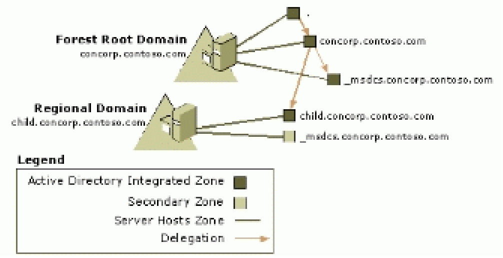
.
Forest Root Domain
concorp.contoso.com
concorp.contoso.com
_msdcs.concorp.contoso.com
Regional Domain
child.concorp.contoso.com
child.concorp.contoso.com
_msdcs.concorp.contoso.com
Legend
Active Directory Integrated Zone
Secondary Zone
Server Hosts Zone
Delegation

Figure 21: New DNS structure

In this design, the DNS server authoritative for the name of the forest root domain hosts a private root zone. All other servers list the DNS root servers in their root hints.

Table 15 provides an overview of the best practice DNS configuration for a network with no existing DNS service.

Table 15 Best Practice DNS Design for an Organization with No DNS Service

Design element	Configuration
DNS server placement	Every domain controller should run DNS.
Recursive name resolution	Root hints on all servers except forest root domain controllers.
Zone placement	The forest root contains zones for: · DNS root · Forest root domain · _msdcs.<forest-root> Each child domain contains a zone for: · Its own child domain · _msdcs.<forest-root>

Note: Active Directory uses a special set of locator records, the forest-wide locator records, to help replication partners find each other and to help clients find global catalog servers. Active Directory stores all the forest-wide locator records in the zone _msdcs.<forest_name>. Because the information in the zone must be widely available, this zone is replicated to all DNS servers in the forest.

Zone Configuration for Forest Root DNS Servers

Design zone configuration for the DNS root zone as specified in Table 16.

Table 16 DNS Root Zone

Design Element	Design
Zone type	Active Directory-integrated
Dynamic updates	Disabled
Scavenging	Disabled
Delegations from the root zone to other zones	A delegation is created for the forest root zone. The delegation includes an NS record for each forest root domain controller.

Design zone configuration for the forest root zone as specified in Table 17.

Table 17 DNS Forest Root Zone

Design Element	Design
Zone type	Active Directory–integrated.
Dynamic updates	Secure dynamic updates only.
Scavenging	Enabled, with default settings.
Delegations from the forest root zone to other zones	· A delegation is created for the zone _msdcs.<forest-root>. The delegation will include an NS record for each forest root domain controller. · A delegation is created for the appropriate regional domains. Each delegation includes an NS record for all regional domain DCs in the appropriate regional domain.

Design zone configuration for the _msdcs.<forest-root> zone as specified in Table 18.

Table 18 Configuration for the _msdcs.<forest-root> Zone

Design Element	Design
Zone type	Active Directory_integrated
Dynamic updates	Secure dynamic updates only

Scavenging	Enabled

Zone Configuration for Regional Domain DNS Servers

The zone for the regional domain is configured as specified in Tables 19 and 20.

Table 19 Zone for Child Domain.

Design Element	Design
Zone type	Active Directory–integrated
Dynamic updates	Secure dynamic updates only
Scavenging	Enabled

Table 20 Configuration for the Secondary Copy of the _msdcs.<forest-root> Zone

Design Element	Design
Zone type	Secondary
Dynamic updates	Not applicable
Scavenging	Not applicable
Other configuration	The zone transfer source is a DNS server on a nearby forest root DC

Client Configuration

To configure the DNS client, the DNS owner for Active Directory must specify the computer naming scheme and how the client will locate DNS servers. Table 21 summarizes these specifications.

Table 21 Client Configuration for DNS When There is No Existing DNS Service

Design element	Configuration
Computer naming	Use default naming. When a Windows 2000 computer joins a domain, the computer assigns itself a primary DNS name comprised of the host name of the computer and the name of the domain.
Client Resolver Configuration	Configure clients with the addresses of at least two DNS servers running on Active Directory domain controllers.

Figure 22 shows the client configuration model. This model assumes you are already running a DHCP server, but not a DNS server or an integrated DNS/DHCP solution.

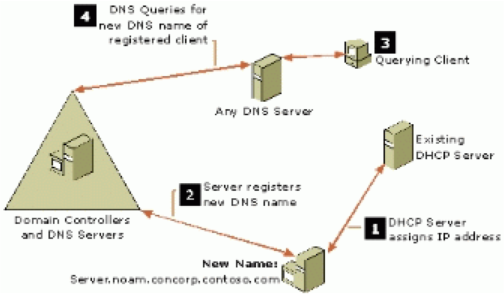
4 DNS Queries for new DNS name of registered client
3 Querying Client
Any DNS Server
Existing DHCP Server
2 Server registers new DNS name
Domain Controllers and DNS Servers
1 DHCP Server assigns IP address
New Name:
Server.noam.concorp.contoso.com

Figure 22: Computer naming without DNS but with an existing DHCP service
In this example, a DHCP server assigns the computer Server.noam.corp.contoso.com an IP address. The computer registers its DNS name in your new DNS infrastructure. When a client queries DNS for the name of the computer, the DNS server finds the name and resolves the query.

With this model, you do not need to make any changes to the DHCP service you already have.

Best Practice DNS Design with Existing DNS Service

The best practice model for DNS depends on whether DNS has already been implemented on the network. When a DNS service has been implemented, the forest owner and DNS owner for Active Directory should work with the existing DNS owner to delegate the forest root domain DNS name to Active Directory and configure an Active Directory–integrated DNS service as specified.

The companion guide, *Best Practice Active Directory Deployment for Managing Windows Networks*, includes detailed instructions for deploying DNS. The information provided here is for reference only.

Server Configuration

Create DNS domains corresponding to the Active Directory domains you have already specified. Table 22 provides an overview of the best practice DNS configuration for a network with an existing DNS service.

Table 22 Best Practice DNS Design for an Organization with Existing DNS Service

Design element	Configuration
DNS server placement	Every domain controller should run DNS.
Recursive name resolution	Determine whether your existing DNS service uses forwarding or root hints. Configure DNS running on domain controllers to match.
Zone placement	The forest root contains zones for: · the forest root domain · _msdcs.<forest-root> Each regional domain contains a zone for: · Its own child domain · _msdcs.<forest-root>

Note: Active Directory uses a special set of locator records, the forest-wide locator records, to help replication partners find each other and to help clients find global catalog servers. Active Directory stores all the forest-wide locator records in the zone _msdcs.<forest_name>. Because the information in the zone must be widely available, this zone is replicated to all DNS servers in the forest.

This design has the advantage that the existing DNS structure remains intact. You do not need to migrate any servers or zones. You are simply adding domains to an existing DNS hierarchy in exactly the same way that the DNS designers intended when they wrote the RFCs.

Note: Although there are other DNS server implementations that can support Active Directory, using Windows 2000 Active Directory-integrated DNS is a best practice because it supports secure dynamic update.

The server configuration will vary depending on whether your DNS service uses root hints or forwarding for recursive name resolution.

Root hints

Figure 23 illustrates the server configuration model for a network that uses root hints. A DNS server somewhere hosts a root zone. This zone could either be the Internet root or a private root. For your configuration, this distinction is irrelevant.

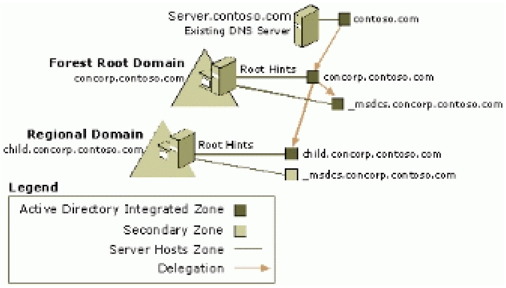

Figure 23: Integration with existing service using root hints

In this example, the DNS servers have root hints that indicate the location of the DNS root servers. When a DNS server needs to resolve a query and cannot resolve it from its cache or zone files, the DNS server queries the DNS root server.

Forwarding

Figure 24 illustrates the server configuration model for a network that uses forwarding.

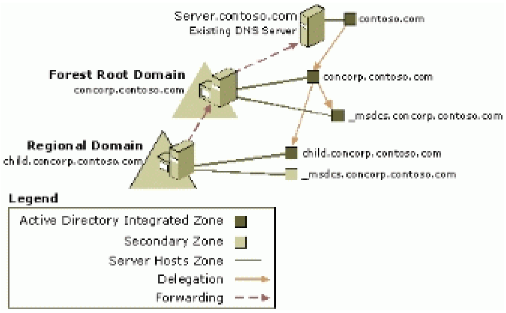

Figure 24: Integration with existing service using forwarding
In this configuration, child DNS servers forward queries to the nearest forest root DNS server, which in turn forwards queries to the same forwarder used by existing DNS servers today.

Zone Configuration for Forest Root DNS Servers

Design zone configuration for the DNS forest root zone as specified in Table 23.

Table 23 DNS Forest Root Zone

Design zone configuration for the DNS _msdcs.<forest-root> zone as specified in Table 24.

Design Element	Design
Zone type	Active Directory–integrated.
Dynamic updates	Secured dynamic updates only.
Scavenging	Enabled.
Delegations from the root zone to other zones	· A delegation is created for the zone _msdcs.<forest-root>. The delegation includes NS records for all forest root domain controllers. · Delegations are created for the appropriate regional domains. Each delegation includes an NS record for all regional domain DCs in the appropriate regional domain.
Other configuration	The parent zone is updated to include a delegation for this zone.

Table 24 DNS _msdcs.<forest-root> Zone

Design Element	Design
Zone type	Active Directory–integrated
Dynamic updates	Secure dynamic updates only
Scavenging	Enabled

Zone Configuration for Regional Domain DNS Servers

The zone for each regional domain is configured as specified in Tables 25 and 26.

Table 25 Zone for the Regional Domain.

Design Element	Design
Zone type	Active Directory–integrated
Dynamic updates	Secure dynamic updates only
Scavenging	Enabled

Table 26 Configuration for the Secondary Copy of the _msdcs.<forest-root> Zone

Design Element	Design
Zone type	Secondary
Dynamic updates	Not applicable
Scavenging	Not applicable
Other configuration	The zone transfer source is a DNS server on a nearby forest root DC

Client Configuration

To configure the DNS client, the DNS owner for Active Directory must specify the computer naming scheme and how the client will locate DNS servers. Table 27 summarizes these specifications.

Table 27 Client Configuration for DNS When There is an Existing DNS Service

Design element	Configuration
Computer naming	Use defaults naming. When a Windows 2000 computer joins a domain, the computer assigns itself a primary DNS name comprised of the host name of the computer and the name of the domain.
Client Resolver Configuration	For existing clients, do not change the resolver configuration. Windows 2000 clients do not need to point directly to a Windows 2000 DNS server and can use any DNS server on the network.

If your clients already have a name registered in DNS, your clients now have two names: the existing name and the new primary name. Clients can still be located by either name. The primary names are automatically created and updated through dynamic update and automatically cleaned up through scavenging.

A computer may have an existing DNS name if the organization has previously implemented:

- A DNS service.
- An integrated DHCP/DNS solution

Note: Using default naming for your Windows 2000 clients does not affect any authentication mechanism you use in the existing network

If you want to take advantage of Kerberos authentication when connecting to a server running Windows 2000, you must make sure that the client connects to the server using the primary name.

Each of these designs leaves any existing DNS, DHCP, or integrated DNS/DHCP solution intact. You could use an arbitrary DNS suffix instead of the Active Directory domain name if the DNS and Active Directory names were different. However, due to the additional management overhead of such a configuration, this is not a best practice.

Note: If you already have reverse lookup zones configured, leave the reverse lookup zones and all the PTR records in place. All existing clients can continue to be registered by using their existing names.

Clients with Names Statically Registered in an Existing DNS Server

In this model, your existing network already has a DNS server for which your administrator manually enters names. Figure 25 shows an example of this configuration.

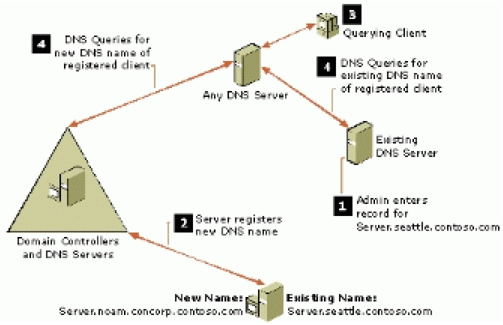
4 DNS Queries for new DNS name of registered client
3 Querying Client
Any DNS Server
4 DNS Queries for existing DNS name of registered client
Existing DNS Server
1 Admin enters record for Server.seattle.contoso.com
Domain Controllers and DNS Servers
2 Server registers new DNS name
New Name: Server.noam.concorp.contoso.com
Existing Name: Server.seattle.contoso.com

Figure 25: Computer naming with existing DNS service
In this example, the existing DNS server hosts a zone containing the name Server.seattle.contoso.com. After joining a domain, the server also registers a new name, Server.noam.concorp.contoso.com, with a DNS server running on a domain controller.

Any client can find the server by either DNS name. When a client queries any DNS server, the DNS server performs name resolution to find the IP address for the name. Depending on how DNS is configured, the DNS server might forward the query, or it might use recursive name resolution. Either way, it queries the appropriate server for the appropriate name and then responds to the client with the appropriate answer.

This model does not affect any DNS naming system you might already have in place. If, before your deployment, your DNS server hosts a zone containing names for your computers, it can continue doing so after your deployment.

Clients with Names Registered by an Integrated DHCP/DNS Solution

In this model, the existing network uses an integrated DHCP/DNS solution such as Lucent QIP. Figure 26 shows an example of this configuration.

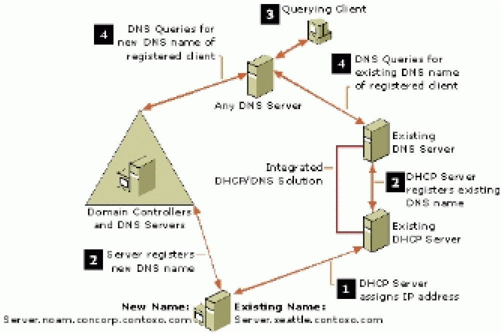
3 Querying Client
4 DNS Queries for new DNS name of registered client
Any DNS Server
4 DNS Queries for existing DNS name of registered client
Existing DNS Server
Integrated DHCP/DNS Solution
2 DHCP Server registers existing DNS name
Domain Controllers and DNS Servers
Existing DHCP Server
2 Server registers new DNS name
1 DHCP Server assigns IP address
New Name: Server.noam.concorp.contoso.com
Existing Name: Server.seattle.contoso.com

Figure 26: Computer naming with integrated DHCP/DNS solution
In this example, the server, Server.seattle.contoso.com, relies on the DHCP server to register its name with the DNS server and to provide an IP address. After joining a domain, the server also registers a new name, Server.noam.corp.contoso.com, with a DNS server running on a domain controller.

Any client can find the server by either DNS name. When a client queries any DNS server, the DNS server performs name resolution to find the IP address for the name. Depending on how DNS is configured, the DNS server might forward the query, or it might use recursive name resolution. Either way, it queries the appropriate server for the appropriate name, and then responds to the client with the appropriate answer.

This model does not affect any integrated DHCP/DNS solution you might already have in place. If, before your deployment, your DHCP server registers DNS names for your computers, it can continue doing so after your deployment.

DNS Design Process

The DNS planning process should help you achieve one of the best practice DNS designs. The steps in the DNS design process include:

1. Plan the DNS server placement.
2. Plan the recursive name resolution.
3. Configure the DNS zones.
4. Plan the computer naming.
5. Plan the client resolver configuration.

Figure 27 illustrates the DNS design process.

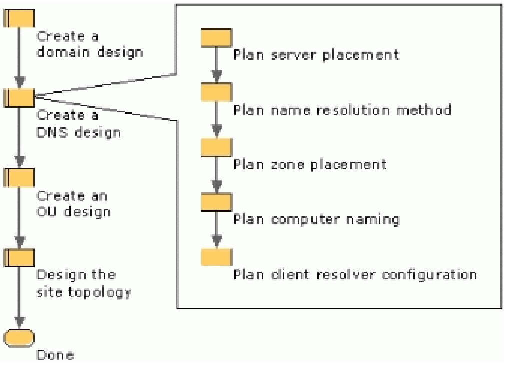

Figure 27: The task flow for the DNS design for Active Directory
The decisions involved in DNS planning are simple and mechanical. If you have an existing DNS structure, your existing structure will dictate many of your choices. As you design each component, fill in the blanks of the corresponding DNS worksheet. After you have finished, review the design with your external DNS and DHCP groups.

Taking the Next Step

As part of the domain design, the forest owner assigned domain ownership. A domain owner exists for every domain identified during the domain design process. Each domain owner can now proceed independently with the next step in the design process, which is discussed in the following section.

Top of page

Creating an Organizational Unit Design

After domain planning is complete, an OU structure can be designed. In the best practices OU model, departments within the domain manage their internal operations, while the domain's IT staff manages the overall infrastructure. In other words, each department manages its objects in the directory, while the domain IT staff manages the configuration of the directory service itself.

Best practices for creating an OU design introduces the role of "OU owner." The Active Directory OU owner is comparable to most Windows NT domain administrators. This means that domain administrators who manage users and resources in a Windows NT domain will manage the same resources in an Active Directory domain, but will be owners of OUs.

Expect to make periodic changes to your OU structure to reflect changes in your administrative structure and to support policy-based administration. OUs are designed to be easily changed.

The Role of OUs in Windows Network Designs

OUs are containers within domains that can contain other OUs, users, groups, computers, and other objects. These OUs and sub-OUs form a hierarchical structure within a domain, and are primarily used to group objects for management purposes.

Note: There are no practical limitations on how many levels OUs can be nested. When designing sub-tiers of OUs, you should compare the value of additional granularity of control with the added complexity of managing the structure. As a best practice, create OU structures no more than ten levels deep.

When designing an OU structure, keep in mind that the OU hierarchy does not need to mirror your organization's departmental hierarchy. Every OU you create should have a defined purpose (such as delegation or policy) and should add value to your system; otherwise, you will spend additional time maintaining the structure without gaining a corresponding benefit.

The initial goal in designing an OU structure is delegation of administration. After this structure is in place, you can further refine it by creating any sub-tiers of OUs you require for other purposes, such as applying Group Policy or placing objects in separate OUs to restrict their visibility.

Delegation of Administration

Delegation of administration allows you to designate groups of users who have control over the users, computers, or other objects in an OU. To accomplish this, place the user who has control into a group, place the set of objects to be controlled into an OU, and then delegate administration of the OU to that group. In Windows 2000, you have very fine control over the administrative tasks that can be delegated; for example, you could create one group that has full control of all objects in an OU, another group that can only create, delete, and manage user accounts in the OU, and another that can only reset the passwords of user accounts. These permissions can be made inheritable so that they cover not only a single container, but also sub-containers.

OUs replace Windows NT domains as the target for delegation of administration to:

- Minimize the number of administrators who must have high levels of access.
- Make individual groups in your organization responsible for local administration.

Group Policy

Although you will first design your OU structure for delegation of administration, some organizations will later create additional tiers of OUs for Group Policy purposes.

Group Policy can use security groups to filter its scope, that is, to apply a Group Policy setting to a subset of the objects in an OU without creating a sub-OU. For example, you could create a Group Policy setting to apply only to managers, even if your OU contains all users. To accomplish this, you apply the Group Policy setting to the OU, then create a group of managers and change permissions on the policy to apply only to that group. In contrast, to delegate administration for a subset of an OU you must create a sub-OU. Because you can achieve precise control of Group Policy scope with security groups, delegation of administration takes priority in the OU design process.

Note: You cannot apply Group Policy settings to the default Users and Computers containers. These are not OUs, and you cannot place OUs in them, so they cannot be used for Group Policy.

It is recommended that you accept the defaults that are set within the Default Domain Policy and the Default Domain Controllers Policy, with the exception of the nodes listed below. These nodes contain settings that you should customize to meet your security requirements.

- Default Domain Policy
  - Password Policy
  - Account Lockout Policy
  - Kerberos Policy
- Default Domain Controllers Policy
  - User Rights Assignment

Other Group Policy settings (those not related to security) should not be set within these default Group Policy settings; instead you should create new Group Policy objects (GPOs). For more information about planning and implementing Group Policy, read the following topics:

- Group Policy planning, see the *Windows 2000 Change and Configuration Management Deployment Guide*, available at:

  http://www.microsoft.com/windows2000/techinfo/reskit/deploy/CCM/default.asp

- Template GPOs for building six different managed scenarios, see *Implementing Common Desktop Management Scenarios*, available at:

  http://www.microsoft.com/downloads/details.aspx?familyid=354B9F45-8AA6-4775-9208-C681A7043292&displaylang=en

- Group Policy:

  http://www.microsoft.com/windows2000/techinfo/howitworks/management/grouppolwp.asp

Elements of an OU Plan

The domain owner is responsible for completing an OU design for the domain. The design should contain:

- A diagram of the OU hierarchy.
- A list of OUs. For each OU:
  - Its purpose.
  - A list of users or groups that have control over the OU or the objects in it.
  - The type of control they have over the class of objects in the OU.

Record your answers in the OU design worksheets provided later in this guide.

The Organizational Unit Owner Role

The domain owner designates an OU owner for each OU. Each OU owner is a data manager with control over a sub-tree of objects in Active Directory. OU owners do not own or control operation of the service, it remains under the control of the domain owner. This allows you to separate ownership and administration of the service from that of the objects contained within the service, and to reduce the number of service administrators who have high levels of access.

Note: Even though the OU owner is delegated control over a sub-tree of objects, the domain owner retains full control over all sub-trees. This should not be changed so that the domain owner can correct mistakes such as an improper access control list (ACL) or recovery of lost sub-trees when administrators are terminated.

Within their scope of management, OU owners can control how administration is delegated, and how policy is applied to objects. They may also create new sub-trees and delegate administration of OUs.

The Best Practices OU Model

In the best practice OU model, each domain contains a standard set of OUs regardless of whether you upgraded or newly created the domain. Each of these OUs delegates administrative control to a data owner.

Figure 28 shows the elements of the model. Note that this may not be identical to the OU hierarchy you design because it shows both possible locations for Resource OUs — under the domain root and under an Account OU. Your design may have both types, or it may have one or the other. See "Resource OUs" later in this section for more information.

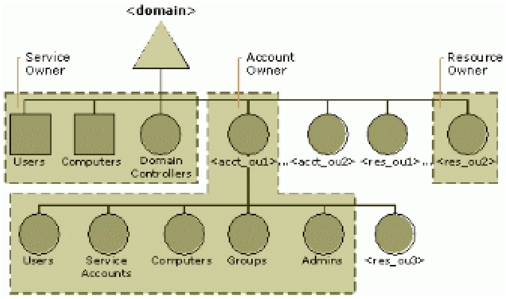

Figure 28: Best practices OU model — overview

Default Containers and OUs

There are several Active Directory default containers, including:

- Users and Computers containers
- Domain Controllers OU

These containers should remain under the control of the domain owner.

Users and Computers Containers

When you perform an in-place upgrade, existing users and computers are automatically placed into the Users and Computers containers. However, these containers are not used. Instead, you will move users and computers into OUs under the control of individual data-level owners.

Note: If you manage your Windows NT domain with earlier tools, such as User Manager, or third-party tools, then new users and groups are created in the Users container in the Active Directory domain. Until you upgrade your tools, you will need to manually move these users to their correct OUs.

When a computer joins an Active Directory domain, if no account exists for that computer, one will be created in the Computers container. To prevent this, pre-create the computer account in the appropriate OU.

Well-known and Built-in Accounts

By default, several well-known and built-in users and groups are created in a new domain. These objects should remain in the default containers and under the control of the domain owner. Table 28 lists these users and groups.

Table 28 List of Well-Known Users and Groups and Built-In Accounts

Well-known users/groups	Built-in accounts
Administrator Guest KRBTGT Domain Admins Domain Users Domain Guests Domain Computers Cert Publishers Schema Admins (root domain only) Enterprise Admins (root domain only).	Administrator Guest Internet Guest Account Launch IIS Process Account

Domain Controller OU

Objects created in the domain and domain controllers OUs have default Group Policy settings applied. A best practices recommendation is that the contents of these containers and their Group Policy settings be left untouched. The domain owner can create new Group Policy settings where necessary.

Account OUs

Account OUs manage identity, that is, they define the set of users and groups that can potentially be granted access to resources. They are analogous to Windows NT Master User Domains (MUDs). The Account OU model is shown in Figure 29.

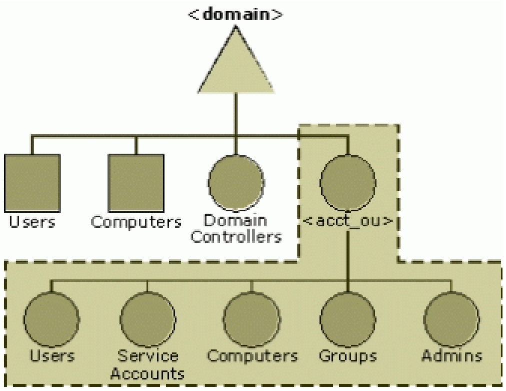

Figure 29: Account OU model
Table 29 lists and describes the OUs created under the account OU structure.

Table 29 OUs Created by Users and Computer Migrated from Windows NT MUDs

OU	Purpose
Users	User accounts for non-administrative personnel.
Service Accounts	Some services that require access to network resources are run under user accounts. This OU is created to separate and distinguish service user accounts from the human user accounts contained in the Users OU. Also, by placing the different types of user accounts in separate OUs, you can manage them according to their different administrative requirements.

Computers	Computer accounts other than Domain Controllers.
Groups	Groups of all types, except administrative groups, which are managed separately.
Admins	User and group accounts for administrative personnel, to allow them to be managed separately from regular users. Auditing should be enabled for this OU so you can track changes to administrative users and groups.

Account OU Administration

Figure 30 defines the administrative group design for an account OU. These groups are all local groups. Create an owner for the Account OU. In this example the group is named <acct_ou>_OU_admins. This group has full control of the OU, and will create the standard set of sub-OUs and the groups to manage them.

Groups that manage the sub-OUs are granted full control only over the specific class of objects they are to manage. For example, <acct_ou>_group_admins has control only over group objects.

Note that there is no separate administrative group to manage the Admins OU; rather, it inherits ownership from its parent OU, so it is managed by <acct_ou>_OU_admins.

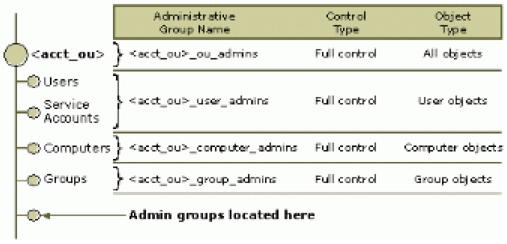

Administrative Group Name	Control Type	Object Type
<acct_ou>_ou_admins	Full control	All objects
<acct_ou>_user_admins	Full control	User objects
<acct_ou>_computer_admins	Full control	Computer objects
<acct_ou>_group_admins	Full control	Group objects

Figure 30: Account OU administrative groups
The OU owner may choose to create additional administrative groups. For example, the optional group <acct_ou>_helpdesk_admins might be created in the Admins OU with control of password resets.

Resource OUs

Resource OUs are used to manage access to resources. They are analogous to Windows NT resource domains. The OU owner creates computer accounts so they can attach servers with resources such as file shares, databases, and printers to the domain. The owner also creates groups to control access to the resources. Figure 31 illustrates two possible locations for the Resource OU.

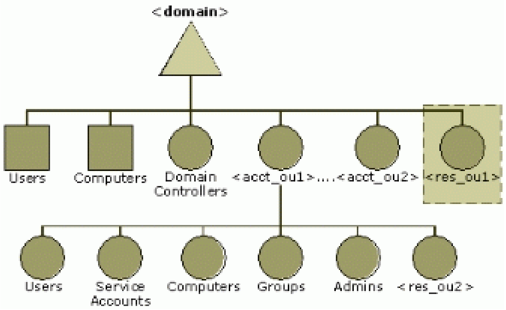

Figure 31: The resource OU model
The Resource OU model differs from the Account OU model in three ways:

- Location in the OU administrative hierarchy.

  If you have Windows NT resource domains that are independently managed by separate IT groups, you can place the new OUs on the top tier, under the domain root. If a Windows NT MUD owner owned and managed the contents of the former resource domain, then the resource OU can be created as a sub-OU of the corresponding Account OU.

  Note: Resource domain owners may prefer to be subordinate to the Windows NT MUD owner rather than the domain owner, to gain better access to support. For example, if an administrative error causes objects in the Resource OU to be inaccessible, the resource administrator needs help from an upper-tier administrator. In large organizations, the Windows NT MUD owner may be nearby and have a relatively smaller scope of management. In contrast, the domain owner could be in a different region or country, and be responsible for a large number of OUs.

- Single OU design.

  Resource OUs have no standard sub-OUs. Computers and groups are placed directly in the OU.

- The Resource OU owner owns only the objects within the OU.

  The resource OU owner does not own the OU container itself. Resource OU owners manage only computer and group objects; they cannot create other classes of objects within the OU. This is done to explicitly limit Resource OU owners to managing computer and group objects, and to prohibit them from creating sub-OUs.

Note: The creator or owner of an object has the ability to set the ACL on the object no matter what permissions are inherited from the parent container. If a Resource OU owner can reset the ACL on an OU, they can essentially create any class of object in the OU, such as users. For this reason, resource OU owners are not permitted to create OUs.

Figure 32 defines the administrative group design for a Resource OU. For each Resource OU, create a group to be the resource owner. This group has full control over the group and computer objects in the OU, but not over the OU container itself. The <res_ou>_OU_admin group manages its own membership and resides in the Resource OU.

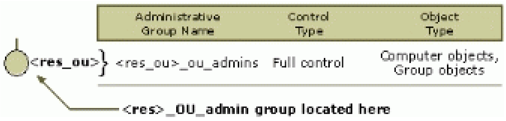

	Administrative Group Name	Control Type	Object Type
<res_ou>	<res_ou>_ou_admins	Full control	Computer objects, Group objects

<res>_OU_admin group located here

Figure 32: Resource OU administrative group design

Note: In Windows NT, resource domain owners typically had full control over computers that joined their domains because they were automatically made a member of the local Admins group on all machines in their domain. In a Resource OU, it is the domain administrator who is the local administrator on all machines. To grant full control over computers in an OU to Resource OU administrators, use restricted group Group Policy. For more information, see the additional reading about Group Policy.

OU Design Process

The domain owner designs the domain's OU structure by creating Account OUs and Resource OUs. Figure 33 illustrates this process.

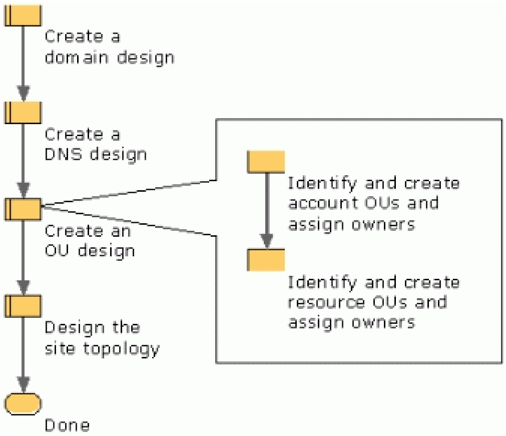
Create a domain design
Create a DNS design
Create an OU design
Identify and create account OUs and assign owners
Identify and create resource OUs and assign owners
Design the site topology
Done

Figure 33: Task flow for the OU design process

Creating Account OUs

Consult your domain design to determine if this domain is in-place upgraded or newly created, and to determine if this domain is the target of MUD consolidation. Use the following steps to create account OUs:

1. If the domain will be upgraded in-place, create an OU for the upgraded accounts.

   During deployment, you distinguish users and groups who belong to the domain owner from regular users and groups. Domain owner users and groups will stay in the default Users container. Regular users and groups will be moved to the appropriate Account OU.

2. If the domain is the target of MUD consolidation, create an OU for each independently managed source MUD.

   If multiple MUDs managed by the same IT group are being consolidated into the domain, you can migrate accounts from those domains into a single Account OU instead of using multiple Account OUs.

3. If the domain is newly created and is not the target of MUD consolidation, create an Account OU for the domain itself.

   This Account OU is necessary to separate the OU (data) owner from the domain (service) owner.

Creating Resource OUs

Consult your domain design to determine if this domain is the target of resource domain consolidation.

If the domain is the target of resource domain consolidation:

1. Create an OU for each independently managed source resource domain.

   If multiple resource domains managed by the same IT group are being consolidated into the domain, you can migrate objects from those domains into a single Resource OU instead of using multiple Resource OUs.

For resource domains managed by Windows NT MUD owners, you can migrate objects into the corresponding Account OU sub-tree (into the Computers and Groups sub-OUs) instead of using a separate Resource OU under the Account OU.

2. Place the OU under the domain root or under an Account OU. The resource domain owner becomes the OU owner.

If the domain is not the target of a resource domain consolidation, create Resource OUs as necessary based on your administrative requirements.

Top of page

Creating a Site Topology

The definition of a site is a set of well-connected (LAN speeds or greater) IP subnets. To create the site topology, identify areas of high connectivity as sites and the WAN connections between them as site links. Once you create sites and site links, Active Directory automatically generates a replication topology between domain controllers. By defining sites according to your LAN/WAN topology, you can ensure a replication topology that avoids WAN connections unless intersite communication is required.

The Role of Site Topology in Windows Network Designs

Active directory sites are a collection of IP subnets constituting a LAN and connected by site links. Active Directory uses sites to:

- Optimizes replication between domain controllers.
- Locate the closest domain controller for client logon and directory searches.

Control Replication

Site topology controls Active Directory replication to achieve the best balance between replication speed and replication cost by distinguishing between replication that occurs within a site and replication that must span sites.

Within sites, replication is optimized for speed — data updates trigger replication and the data is sent without the overhead required by data compression. Conversely, replication between sites is compressed to minimize the cost of transmission over WAN links.

Route Replication

Active Directory uses a multimaster, store-and-forward method of replication. A domain controller communicates directory changes to a second domain controller, which then communicates to a third, and so on, until all domain controllers have received the change. When replication occurs between sites, a single domain controller per domain at each site collects directory changes, stores these change, and communicates them at a scheduled time to a domain controller in another site.

Client Affinity

Active Directory clients locate domain controllers according to their site affiliation. A client locates a domain controller within the same site whenever possible. By finding a domain controller in the same site, the client avoids communications over WAN links.

Network Topologies

An organization's network topology should reflect its business needs. In some organizations, business needs result in users being situated in a few large locations that are well connected to one another. Alternatively, in other organizations users are situated in many small satellite locations connected to one of a few, well-connected hub sites. Such network topologies fall into one of three general types: ring, hub and spoke, and complex. These are illustrated in Figure 34.

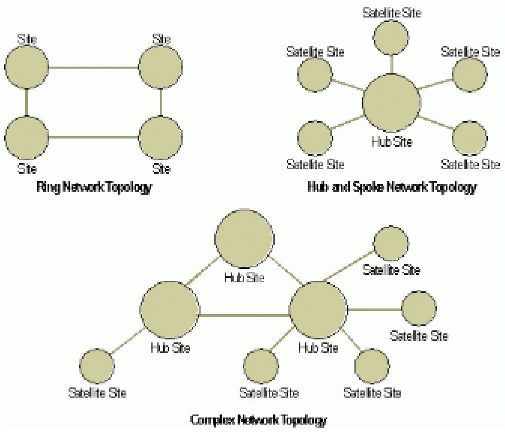

Figure 34: Network topologies
Elements of a Site Topology

The forest owner and site topology owner are responsible for completing a site topology design for the forest. The plan should contain:

- A location map indicating:
  - The number of users and computers at each location.

    - The speed and utilization of WAN links.
    - The IP addresses used at each site.
- A list of sites, stating for each:
    - The name of the site.
    - The number of users and computers.
    - The domains relevant to that site.
    - For each domain, the number of domain controllers and its specified hardware.
    - Domain controllers that are also global catalogs.
- A list of site links, stating for each:
    - The sites connected by the link.
    - The cost associated with each.
    - The scheduled replication time on each link.
    - The replication interval for each link.
- Replication latency calculations.

The Site Topology Owner Role

The site topology owner understands the conditions of the network between sites and has the authority to change settings in Active Directory to implement changes to the site topology. Changes to the site topology affect changes in the replication topology. The site topology owner's responsibilities include those summarized in Table 30.

Table 30 Site topology owner responsibilities

Role	Responsibility
Control changes to site topology	When network connectivity changes, the owner makes the respective changes to the site topology.
Move domain controllers between sites	If a domain controller's IP address changes, which puts it on a subnet in a different site, or if the subnet moves to a different site, you must move the domain controller to the new site manually.
Liaison with network group	The owner obtains and maintains information about connections and routers that affect site topology. · To effectively set costs on site links, the owner understands the issues of network speed and capacity that affect site topology. · Maintains a list of subnet addresses, subnet masks, and the location to which each belongs.

Best Practices Design for Site Topology

This document offers best practice guidelines for creating a site topology rather than a model.

- Delegate authority for site topology management to the site topology owner.
- Use the default configuration for intersite replication where possible. These include:
    - Do not disable the Knowledge Consistency Checker (KCC).
    - Do not disable site link transitivity.
    - Do not specify bridgehead servers.
    - Keep the replication schedule open as long as is practical.
- Calculate the expected replication latency between sites in your topology.

For advanced topics about replication, see *The Active Directory Branch Office Planning Guide*, available on-line at http://www.microsoft.com/windows2000/techinfo/planning/activedirectory/branchoffice/default.asp.

The following are best practice guidelines for designing site links:

- Map site links to WAN links.

- Make sure no single site is directly connected to more than 20 other sites.

  This condition can occur in large hub-and-spoke deployments where most sites are branch sites that communicate with a centralized hub site. If this condition exists and there are more than 20 site links from the hub site to branch sites, the hub site can be divided into multiple sites to provide additional bridgehead servers to handle the replication volume. In a site, a single bridgehead server is active per domain. If the site has more than 20 site links, the bridgehead servers can become overloaded.

Site Topology Design Process

To design a best practice topology, perform the following tasks in the order specified:

1. Create a location map.
2. Place domain controllers in locations.
3. Create sites based on locations.
4. Connect sites with site links.
5. Perform domain controller capacity planning.

Figure 35 illustrates the site topology design process.

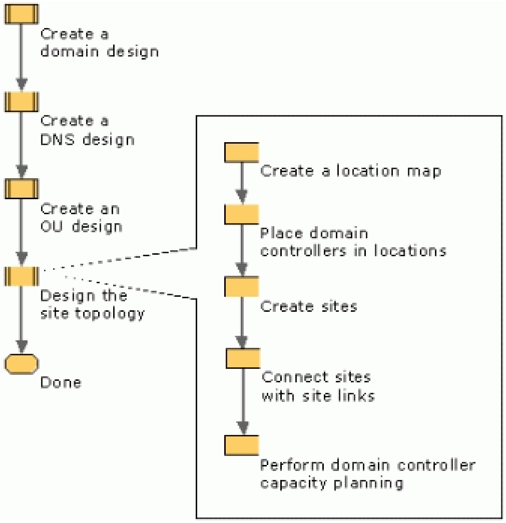

Figure 35: Task flow for the site topology design process

Creating a Location Map

A location map determines where users are located and which locations are best suited for replication with which other locations. As a best practice, Active Directory sites should map to the locations of LAN segments of your network. To create a location map:

1. Create a location map covering all physical locations in your forest connected by LAN-speed or better links.

   Consider a physical location to be a group of users with internal connectivity of 10 Mbps or better.

2. For each location, determine:
   - The number of users.
   - The number of workstations.
   - The list of local IP subnets.
3. Add network WAN connections to the map indicating:
   - WAN speed between all locations.
   - Current use of the WAN link.

Figure 36 shows a sample location map. Each location includes the number of users, the number of workstations, and the IP addresses of network subnets present. Each network connection includes its WAN speed and percentage or bandwidth used.

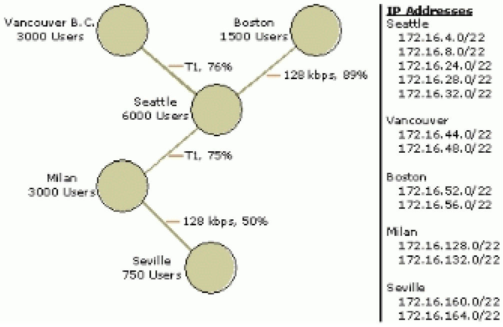

Figure 36: A location map with LAN segments and WAN links

Placing Domain Controllers in Locations

In conjunction with domain owners, the site topology owner determines the location of the domain controllers for the forest root domain and regional domains. Place domain controllers in locations so, should the WAN link fail, there is local authentication for users.

Begin evaluating locations by considering the placement of domain controllers in your hub sites. Once you have completed added domain controllers to all hub sites, consider all satellite locations. You may choose not to add domain controllers to your smallest locations if:

- The location contains too few users and workstations to warrant a domain controller.
- Domain controller physical security cannot be guaranteed.

Figure 37 illustrates a map of all sites showing domains that are relevant to each site.

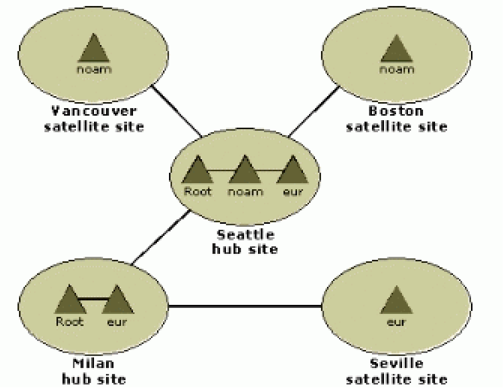

Figure 37: Domains present in a set of sites

Placing Forest Root Domain Controllers

For each of your hub sites, place one domain controller from the forest root domain. For each satellite location, evaluate the need for a forest root domain controller. Keep in mind that for a client to access a resource in a domain other than its own, communication must occur with a forest root domain controller. For this reason, any location that has domain controllers from two or more regional domains should also have a forest root domain controller.

Placing Regional Domain Controllers

Place a domain controller for regional domains into a location if the users in that location require the ability to log on when the WAN link is down. If the local domain controller fails, the users at that location will log on over the WAN. To improve logon performance in the case of a single domain controller failure, consider adding a second domain controller.

Placing Primary Domain Controllers

Each domain will have a PDC emulator operations master. Place the PDC in a location that is:

- Well-connected to other sites.
- Contains a large number of users from that domain.

The site where you place the PDC is designated the primary site for that domain.

Placing Global Catalog Servers

A global catalog server is required for logon to native-mode Active Directory domains. To eliminate the need to contact a global catalog server in a distant site for logons and for forest-wide searches, designate at least one domain controller per site as a global catalog.

Best practice design dictates that half of all domain controllers in the site be global catalogs, with at least two global catalogs if your site has multiple domain controllers. If you use the single, global domain model, plan for all domain controllers for the global domain to also be global catalog servers. Because the forest root domain is very small, making all domain controllers global catalogs requires very little additional resources.

Creating Sites Based on Location

Each location that contains a domain controller constitutes a site. To create a site:

1. Name the site
2. Specify the IP addresses associated with the site.

   Do not specify sites for locations that do not contain domain controllers. Instead, associate the subnets for these locations with the sites where you would like for these users to be authenticated.

Connecting Sites with Site Links

Intersite replication occurs according to the settings on site link objects that site administrators create in Active Directory. Each site link object represents the WAN connection between two or more sites.

Use the following steps to control how replication occurs between sites:

1. Connect sites with site links to model the way that your locations are connected.
2. Name the site links <name_of_site1>-<name_of_site2>.
3. Set site link parameters:
   - Cost
   - Schedule
   - Interval

Note: If multiple sites have the same connectivity and availability to each other, you can connect them with the same site link.

Figure 38 illustrates a site map with site links and WAN speeds indicated.

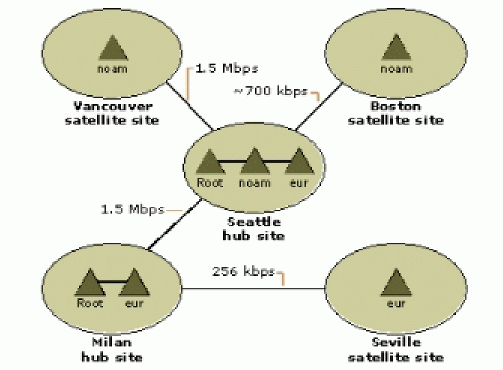

Figure 38: Network map of site links and WAN speeds

Setting the Cost

To determine the costs to place on site links, perform these tasks:

1. Assess the link speed.
2. Create a table containing the connection speed for each site link.
3. Use Table 31 to calculate the cost of each site link.

   To get a relative cost factor for site links that correlates to the available bandwidth, use the following formula:

$$Cost = \frac{1024}{\log(Available\ Bandwidth(Kbps))}$$

Table 31 Setting Costs on Site Links Based on WAN Speeds

Available bandwidth (kilobits/second)	**Cost**
9.6	1042
19.2	798
38.4	644
56	586
64	567
128	486
256	425
512	378
1024	340
2048	309
4096	283

These costs do not reflect differences in reliability between network links. If some of your links are less reliable than others, then you should not rely on these links for replication. Set higher costs on any failure-prone network links.

When a site link goes down, replication failover can be predicted and controlled by setting site link costs. Figure 38 illustrates how site link costs can be used to show the relative cost of linking any two sites in the topology.

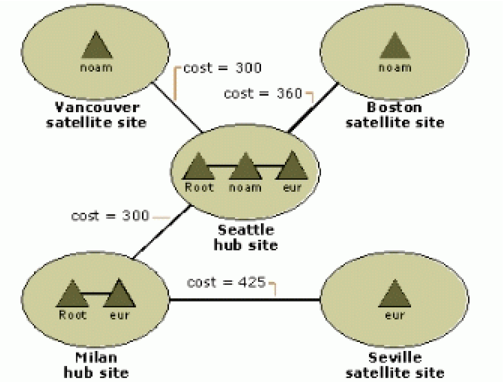

Figure 39: Site Map Containing Link Costs
Setting the Schedule

The site topology owner determines when site links are available for replication by setting a scheduled replication window. However, a replication cycle that starts before the end of the schedule will run until completion even if the scheduled replication window closes.

Control site link availability by setting a schedule on site links. You might use the default (100 percent available) schedule on most links, but block replication traffic during peak business hours on links to certain remote locations. By blocking replication, you give priority to other traffic, but you also increase replication latency.

When replication between two sites traverses multiple site links, the intersection of the replication schedules on all relevant links determines the connection schedule between the two sites. In addition, when schedules on each link never overlap, replication can never occur.

Setting the Interval

Within the scheduled replication window, interval determines how often domain controllers poll replication partners for changes.

During each window of the schedule that is open for replication, decide the interval of time that determines how often replication occurs within the schedule window. Consider the following criteria:

- A small interval decreases latency but increases the amount of WAN traffic.
- To keep domain directory partitions up-to-date, low latency is preferred.

Calculating Maximum Latency

With a store-and-forward replication strategy determining just how long a directory update might take to be replicated to every domain controller is not easily determined. To provide a conservative estimate of maximum latency perform these tasks:

1. Create a table of all the hub sites on your network. Your table should resemble the one shown in Table 32.

   Table 32 Latency of replication between hub sites

	Seattle	Boston	Vancouver	Milan
○ Seattle	○ 0.25			
○ Boston		○ 0.25		
○ Vancouver			○ 0.25	
○ Milan				○ 0.25

2. An average, worst-case latency within a site is estimated to be 15 minutes.

3. From the replication schedule, determine the maximum replication latency possible on any site link connecting two hub sites.

   For example, if replication occurs between Seattle and Boston for one hour a day, then the maximum delay for replication between these sites is 23 hours. If the replication delay between Boston and Seattle is the longest scheduled delay among all hub sites, then the maximum latency between all hubs is 23 hours.

4. For each hub site, create a table of the maximum latencies between the hub site and any of its satellite sites.

   For example, if replication occurs between Boston and Bar Harbor every four hours and this is the largest replication delay between Boston and any of its sites, then the maximum latency between the Boston hub and its satellites is four hours.

5. Combine these maximum latencies to determine the maximum latency for the entire network.

   For example, if the maximum latency between Milan and its satellite site in Seville is one day, then the maximum replication latency for this set of links is 52 hours (4+24+24).

   Review your maximum replication latency and revise your replication schedule if necessary.

Domain Controller Capacity Planning

The following operations create loads on domain controllers:

- User and workstation logon
- Directory lookup operations
- Replication
- Operations masters
- Other services running on the domain controller

User and workstation logons followed by directory lookups most impact domain controller performance.

The size of the domain is not usually an indication of the performance of its domain controllers. Instead, the usage patterns and the number of users logging on to a domain controller in a specific site most closely determine a domain controller's performance. Table 33 shows how services affect domain controller performance.

Table 33 Effect of services on domain controller performance

Operation	Service	Performance impact
User logon	Interactive user logon,User network logon, User batch logon	High
Workstation logon	Workstation boot process	Medium
Look-up operation	Search operations, LDAP searches	Depends on usage
Replication with 1 to 10 partners	Active Directory replication	Low - Medium

Replication with >10 partners	Active Directory replication	High
PDC operations master	Password change forwarding, logon forwarding	High
Infrastructure operations master	Validation of links to moved objects	Low
RID pool operations master	Distribute RID pools to domain controllers in this domain	Low
KCC with less than 200 sites	Compute replication topology	Low - Medium
KCC with more than 200 sites	Compute replication topology	High
Global catalog service	Universal group membership lookups,Forest-wide searches	Low - High (with Exchange 2000)
Infrastructure services	DNS, WINS, DHCP	Medium
Other services	File and print, Database operations	Depends on usage

Processor Requirements

Table 34 shows the size of domain controllers and services that can be loaded:

Table 34 Domain Controller Capacity Planning

Users in site	Domain controller	Global Catalog	Dedicated DC?	Server/Advanced Server
< 100	One uniprocessor PIII 500, 512 MB	Domain controller is a global catalog.	No	Server
101 – 500	One uniprocessor PIII 500, 512 MB	Domain controller is a global catalog.	Yes	Server
500 – 1,000	One Dual PIII 500, 1 GB.	Domain controller is a global catalog.	Yes	Server
1,000 – 10,000	Two Quad PIII XEON, 2 GB.	Both domain controllers are global catalogs.	Yes	Advanced Server
> 10,000 users	One Quad PIII XEON, 2 GB for every 5,000 users.	One global catalog for every two domain controllers with a minimum of two global catalogs.	Yes	Advanced Server

Note: The best practice guidelines for domain controller hardware are:

- For 100 to 200 sites use multiprocessor hardware.
- For >200 sites, consult the branch office documentation.
- For sites connected to 10 or more sites use multiprocessor hardware.
- For domains with > 50,000 users, use a Quad PIII XEON with 2 GB of memory for the dedicated PDC.

This table provides a starting point for domain controller sizing. Your actual hardware needs will depend on your usage pattern. If you require more accurate hardware size determinations or if you will be using Exchange 2000 in your deployment:

1. Download and run the ADSizer.
2. Test the hardware setup in the lab before deploying.

Disk Size Requirements

Domain controllers and global catalogs can use large amounts of data storage for the Active Directory database. Determine how much storage to provide for the Active Directory database from the information provided in Table 35.

Table 35 Storage Requirements for the Active Directory Database

Server	Active Directory database storage requirements

Domain controller	0.4 GB of storage for each 1,000 users.
Global catalog server	$= DC\ storage\ requirement + \frac{\sum DC\ storage\ requiement\ \ for\ other\ domains}{2}$

For example, in a forest with two 10,000-user domains, all domain controllers need 4 GB of storage. All global catalogs require 6 GB of storage.

Disk Layout Requirements

A domain controller requires space on storage devices for its operating system, log files, database, and SYSVOL. For domain controllers in sites with less than 10, 000 users, all four components can reside on a single RAID 1 array. A RAID 1 array provides protection against single disk failures.

Over time if you find that a domain controller becomes I/O bound, you can enhance the performance of the computer by adding additional memory (up to 2 GB). This alleviates the need to move the operating system, database, or logs files to separate RAID 1 systems.

For domain controllers in sites with more than 10,000 users, use separate RAID arrays as specified in Table 36.

Table 36 Domain controller disk capacity planning

Directory service component	Operations performed	RAID system
Operating system	Read and write operations	RAID 1
Log files	Mostly write operations	RAID 1
Database and SYSVOL	Mostly read operations	RAID 1 or RAID 0+1

Top of page

Implementing Your Design

After completing all of the tasks in the design process, implement your design by testing it in a lab environment and rolling it out in your production environment. The companion to this guide, *Best Practice Active Directory™ Deployment for NOS Environments*, will assist you with the implementation process.

Top of page

Part III: Worksheets

Top of page

Number of Forests in Your Organization Worksheets

After completing all of the tasks in the design process, implement your design by testing it in a lab and pilot environments and rolling it out in your production environment. The companion to this guide, *Best Practice Active Directory™ Deployment for Managing Windows Networks*, will assist you with the implementation.

The following worksheets correspond to the headings in this guide. The worksheets are designed to help you gather and organize the information you'll need to customize your organization's Active Directory deployment.

As you read this document and conduct planning sessions, you should complete the corresponding worksheets. Collectively, the worksheets become the core of your organization's Active Directory deployment project plan.

Worksheet: Number of Forests Design Team

Individuals participating in committee tasked with determining the number of Active Directory forests that your organization will design and deploy.

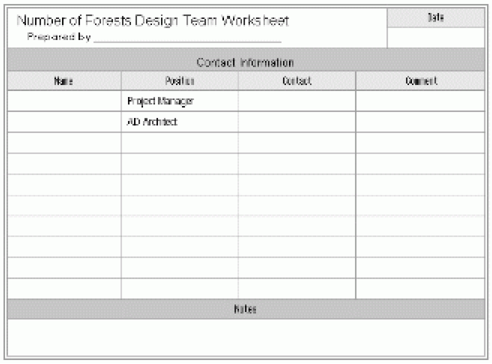

Number of Forests Design Team Worksheet

Prepared by ______________________

Date

Contact Information

Name	Position	Contact	Comment
	Project Manager		
	AD Architect		

Notes

Add or remove rows from the table to customize the form for your organization.

Worksheet: Candidate Directory Service Providers

The following divisions were identified as possible directory service providers and participated in the number of Active Directory forests design discussion.

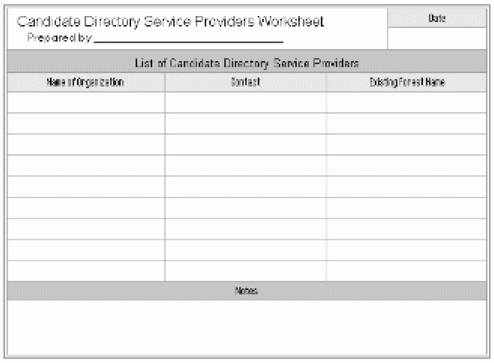

Candidate Directory Service Providers Worksheet

Prepared by: ____________________

Date

List of Candidate Directory Service Providers

Name of Organization	Contact	Existing Forest Name

Notes

List each organization that is a candidate directory service provider, and include the name of the contact person or representative. If the organization has already deployed Active Directory, specify the existing forest name. Add or remove rows from the table to customize the form for your organization.

Worksheet: Forest Design Justification

The following worksheet provides the design rationale behind each candidate forest identified in the previous worksheet. Create a worksheet for each candidate forest in your organization.

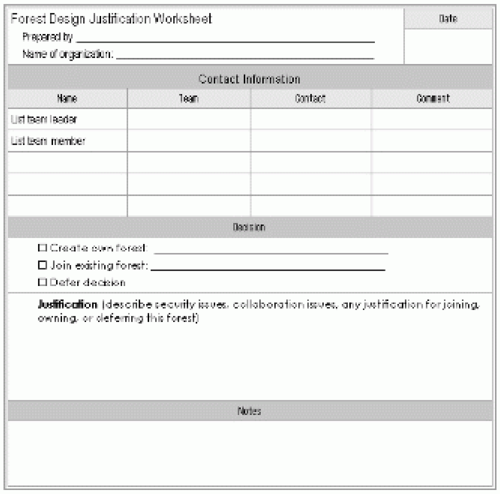

Forest Design Justification Worksheet

Prepared by ____________________

Name of organization: ____________________

Date

Contact Information

Name	Team	Contact	Comment
List team leader			
List team member			

Decision

□ Create own forest: ____________________

□ Join existing forest: ____________________

□ Defer decision

Justification (describe security issues, collaboration issues, any justification for joining, owning, or deferring this forest)

Notes

Add or remove rows from the table to customize the form for your organization.

Worksheet: Proposed Forests

The following worksheet provides a list of the forests that are proposed for your organization.

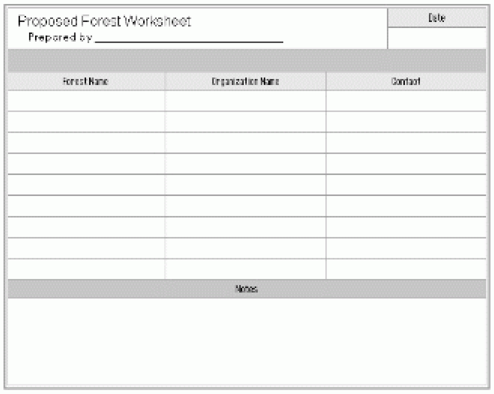

Proposed Forest Worksheet

Prepared by: ____________________

Date

Forest Name	Organization Name	Contact

Notes

Add or remove rows from the table to customize the form for your organization.

Top of page

Active Directory Forest Design Worksheets

You will complete one set of the remaining worksheets for each forest that you identified in the *Number of Forests* worksheets. For each forest listed, forward this document to the contact individual for that forest. The appropriate contact should then complete the remaining worksheets in this guide to finalize that forest's Active Directory design.

Worksheet: Forest Design Team

The following worksheets list the Planning Team and Technology Owners responsible for the Active Directory domain design for your forest.

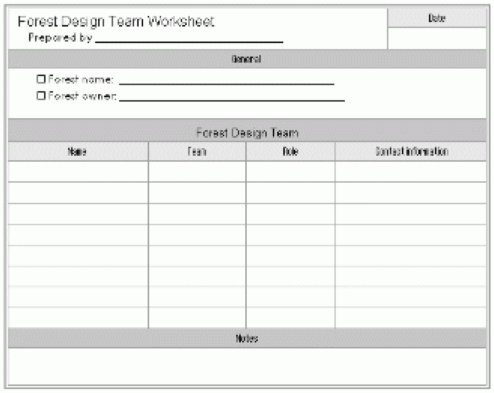

Forest Design Team Worksheet

Prepared by ______________________

Date

General

☐ Forest name: ______________________

☐ Forest owner: ______________________

Forest Design Team

Name	Team	Role	Contact Information

Notes

Complete the table by identifying the technology owners and team members for the Active Directory design of your forest. Add or remove rows from the table to customize the form for your organization.

Worksheet: Forest Root Domain Administration

The following worksheet identifies the forest root domain administrators for your forest.

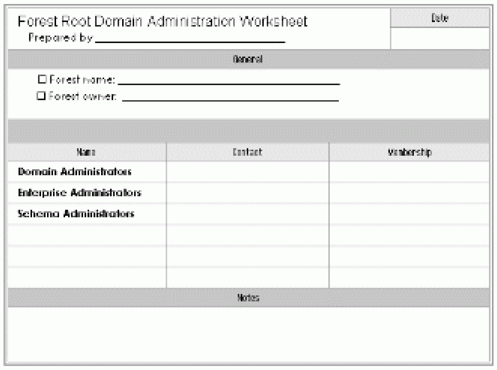

Forest Root Domain Administration Worksheet

Prepared by ______________________

Date

General

☐ Forest name: ______________________

☐ Forest owner: ______________________

Name	Contact	Membership
Domain Administrators		
Enterprise Administrators		
Schema Administrators		

Notes

Add or remove rows from the table to customize the form for your organization.

Worksheet: Forest Root Domain Design

The following worksheet identifies the forest root domain name design for the your forest.

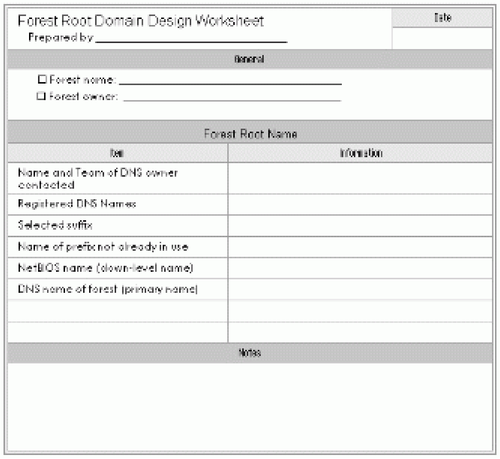

Forest Root Domain Design Worksheet

Prepared by ______________________

Date

General

☐ Forest name: ______________________

☐ Forest owner: ______________________

Forest Root Name

Item	Information
Name and Team of DNS owner contacted	
Registered DNS Names	
Selected suffix	
Name of prefix not already in use	
NetBIOS name (down-level name)	
DNS name of forest (primary name)	

Notes

Add or remove rows from the table to customize the form for your organization.

Worksheet: Identifying Regions

The following worksheet identifies the regions in the your forest.

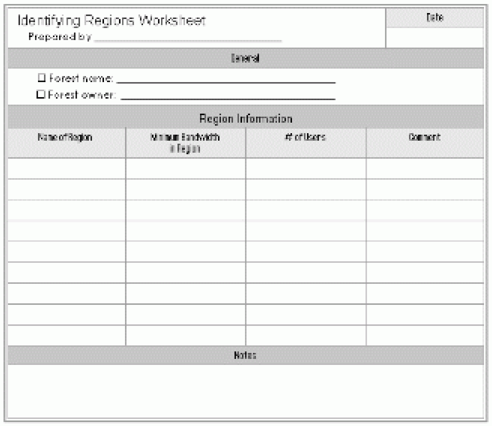

Identifying Regions Worksheet

Prepared by ______________________

Date

General

□ Forest name: ______________________

□ Forest owner: ______________________

Region Information

Name of Region	Minimum Bandwidth in Region	# of Users	Comment

Notes

Worksheet: Proposed Domain Design

The following worksheet identifies the Domain Model for the your forest.

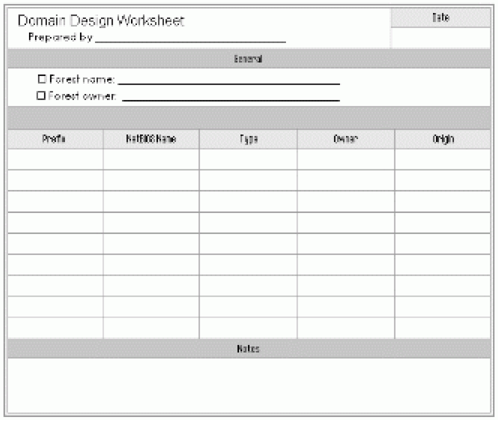

Domain Design Worksheet

Prepared by ______________________

Date

General

□ Forest name: ______________________

□ Forest owner: ______________________

Prefix	NetBIOS Name	Type	Owner	Origin

Notes

Add or remove rows from the table to customize the form for your organization.

Worksheet: Inventory of Windows NT Master User Domains

The following worksheet identifies existing Windows NT Master User domains in your forest.

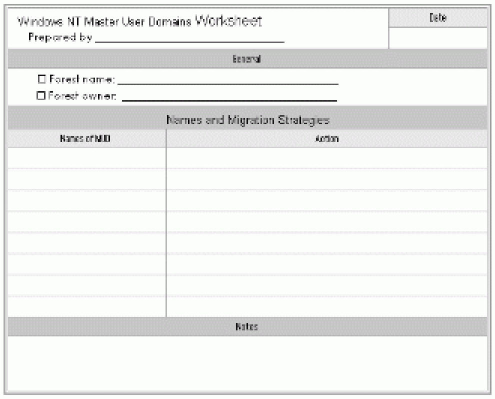

Windows NT Master User Domains Worksheet

Prepared by ______________________

Date

General

□ Forest name: ______________________

□ Forest owner: ______________________

Names and Migration Strategies

Names of MUD	Action

Notes

For each existing MUD in your forest copy and complete the "Quick Reference" worksheet. Add or remove rows from the table to customize the form for your organization.

Worksheet: Master User Domain Information

For each row in the previous worksheet, copy and complete the following detail worksheet explaining the details/justification of your migration strategy.

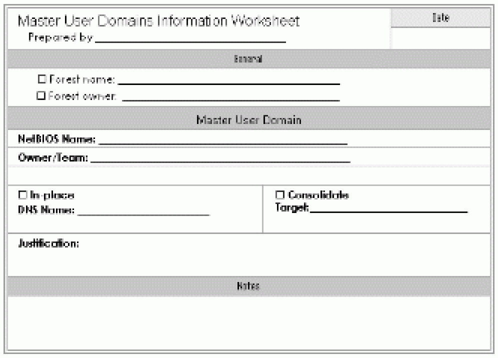

Master User Domains Information Worksheet

Prepared by ______________________

Date

General

☐ Forest name: ______________________

☐ Forest owner: ______________________

Master User Domain

NetBIOS Name: ______________________

Owner/Team: ______________________

☐ In-place DNS Name: ______________	☐ Consolidate Target: ______________

Justification:

Notes

Add or remove rows from the table to customize the form for your organization.

Worksheet: Windows NT Resource Domain Mapping

Identify each existing Resource Domain in your forest by completing the worksheet below.

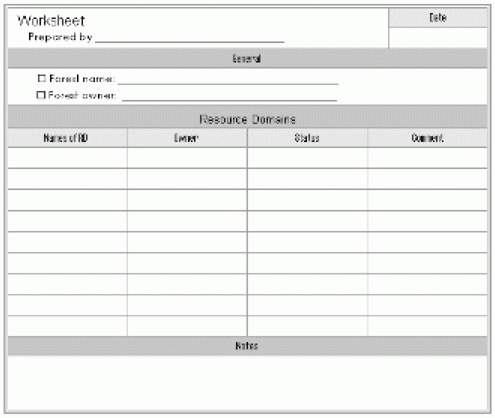

Worksheet

Prepared by ______________________

Date

General

☐ Forest name: ______________________

☐ Forest owner: ______________________

Resource Domains

Names of RD	Owner	Status	Comment

Notes

Add or remove rows from the table to customize the form for your organization.

Worksheet: Existing DNS Service Inventory

If you have existing DNS service in the your forest, complete all of the worksheets in this section. If you don't have existing DNS service, ignore and delete the inventory worksheets in this section.

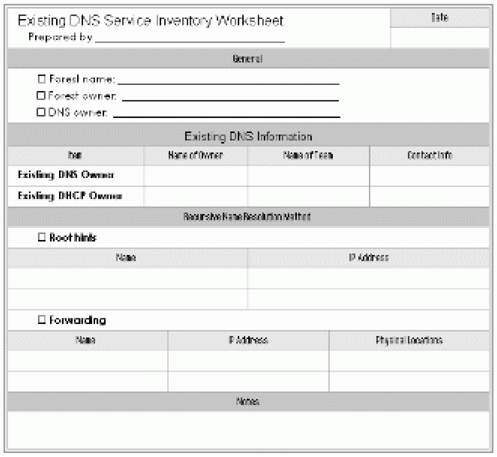

Existing DNS Service Inventory Worksheet

Prepared by ______________________

Date

General

☐ Forest name: ______________________

☐ Forest owner: ______________________

☐ DNS owner: ______________________

Existing DNS Information

Item	Name of Owner	Name of Team	Contact Info
Existing DNS Owner			
Existing DHCP Owner			

Recursive Name Resolution Method

☐ Root hints

Name	IP Address

☐ Forwarding

Name	IP Address	Physical Locations

Notes

Add or remove rows from the table to customize the form for your organization.

Worksheet: DNS Client Information

Answer the following questions to describe the DNS client computers in the your forest. The following describes the DNS client inventory in the your forest.

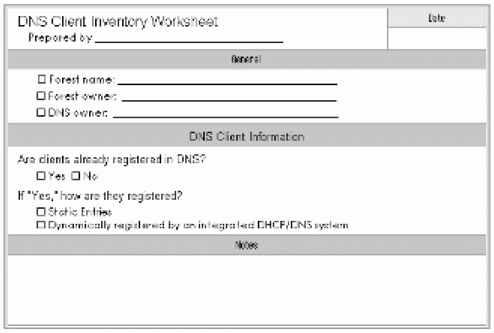

DNS Client Inventory Worksheet

Prepared by ______________________

Date

General

□ Forest name: ______________________

□ Forest owner: ______________________

□ DNS owner: ______________________

DNS Client Information

Are clients already registered in DNS?

□ Yes □ No

If "Yes," how are they registered?

□ Static Entries

□ Dynamically registered by an integrated DHCP/DNS system

Notes

Add or remove rows from the table to customize the form for your organization.

DNS Client Configuration

Create a figure that matches your client model.

Worksheet: DNS Server Design

Create a figure that matches your DNS for Active directory model (DNS with root hints, DNS with forwarding, and/or integrated DNS/DHCP solution).

Worksheet: Organizational Unit Design Team

The following worksheet identifies individuals on the OU design team for your domain.

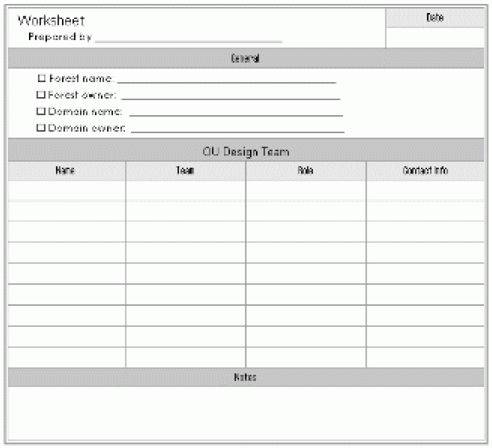

Worksheet

Prepared by ______________________

Date

General

□ Forest name: ______________________

□ Forest owner: ______________________

□ Domain name: ______________________

□ Domain owner: ______________________

OU Design Team

Name	Team	Role	Contact Info

Notes

Worksheet: Identifying OUs for Each Domain

The following worksheet identifies the OUs for each domain in the your forest. For each domain complete the following worksheet indicating the OUs in each domain.

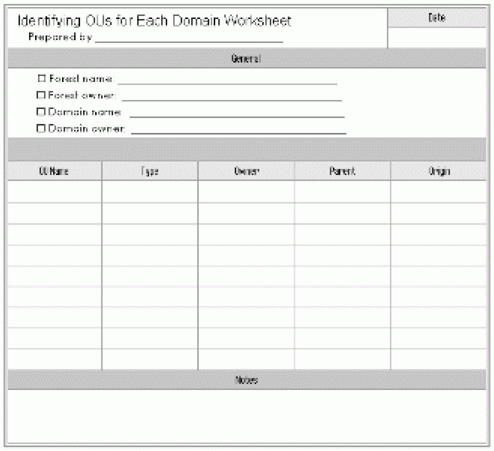

Identifying OUs for Each Domain Worksheet

Prepared by ______________________

Date

General

□ Forest name: ______________________

□ Forest owner: ______________________

□ Domain name: ______________________

□ Domain owner: ______________________

OU Name	Type	Owner	Parent	Origin

Notes

Worksheet: Site Topology Design Team

The following worksheet identifies members of the site topology design team.

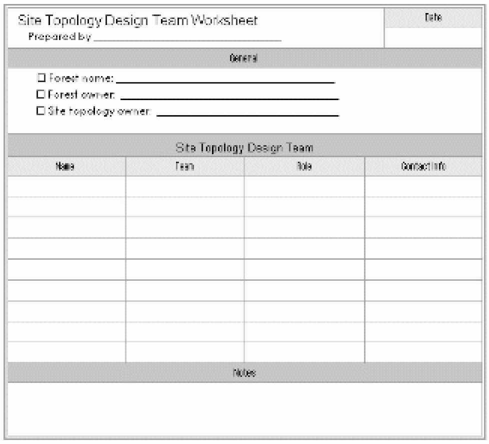

Site Topology Design Team Worksheet

Prepared by ______________________

Date

General

□ Forest name: ______________________

□ Forest owner: ______________________

□ Site topology owner: ______________________

Site Topology Design Team

Name	Team	Role	Contact Info

Notes

Add or remove rows from the table to customize the form for your organization.

Worksheet: Geographic Locations and Communication Links

On a geographic or organizational map summarize your organization's locations and links. Include the number of users at each location and the link speed.

Worksheet: Geographic Locations and Links

This worksheet identifies each location and the other locations directly connected to the location, the type of connection, and the available network capacity (remaining after other network traffic).

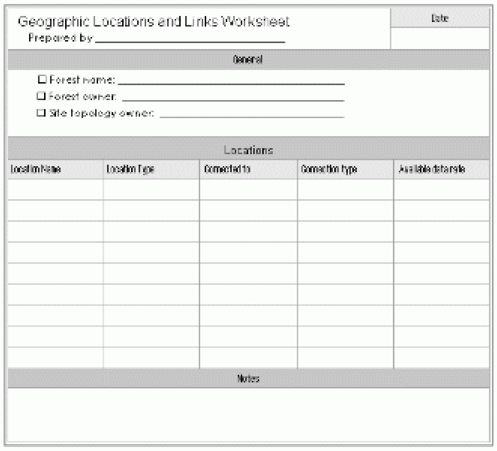

Geographic Locations and Links Worksheet

Prepared by ______________________

Date

General

□ Forest name: ______________________

□ Forest owner: ______________________

□ Site topology owner: ______________________

Locations

Location Name	Location Type	Connected to	Connection type	Available data rate

Notes

Add or remove rows from the table to customize the form for your organization.

Worksheet: Domain Controller Placement

The following worksheets describe existing and planned domain controller placement for the your forest. In the Location Name column, list each geographic location. For each Location, then list the names of the domains at that site, and the numbers of users, domain controllers, and global catalogs used as domain controllers at that site.

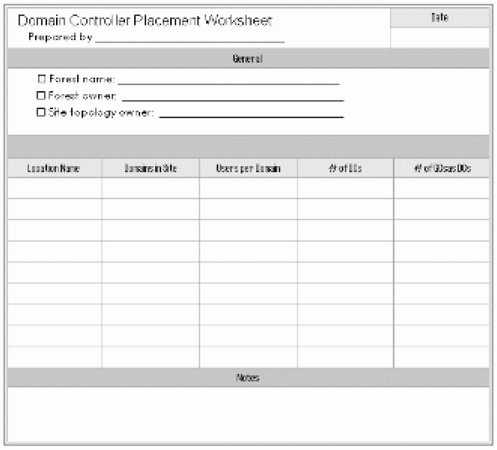

Domain Controller Placement Worksheet

Prepared by ______________________

Date

General

□ Forest name: ______________________

□ Forest owner: ______________________

□ Site topology owner: ______________________

Location Name	Domains in Site	Users per Domain	# of DCs	# of GCs as DCs

Notes

Add or remove rows from the table to customize the form for your organization.

Worksheet: Mapping Sites to Locations and Subnets

The following worksheet maps sites to their locations and subnets.

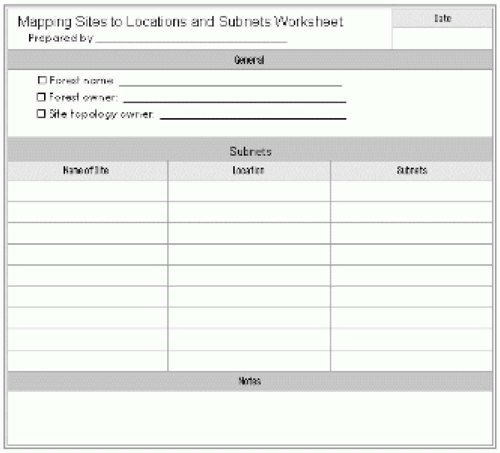

Mapping Sites to Locations and Subnets Worksheet

Prepared by ________________

Date

General

□ Forest name: ________________

□ Forest owner: ________________

□ Site topology owner: ________________

Subnets

Name of Site	Location	Subnets

Notes

Add or remove rows from the table to customize the form for your organization.

Worksheet: Site Links Parameters

The following worksheet describes the associated cost, schedule, and replication interval.

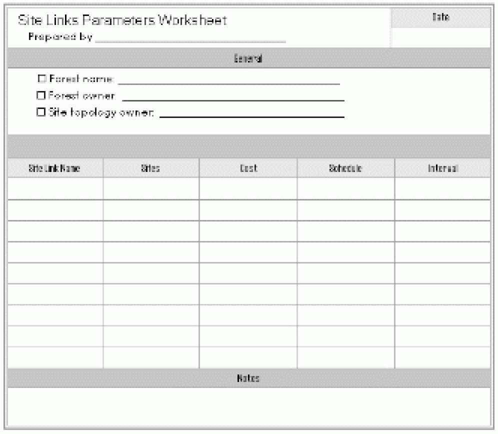

Site Links Parameters Worksheet

Prepared by ________________

Date

General

□ Forest name: ________________

□ Forest owner: ________________

□ Site topology owner: ________________

Site Link Name	Sites	Cost	Schedule	Interval

Notes

Add or remove rows from the table to customize the form for your organization.

Acknowledgements

We would like to thank the following people for reviewing the guides and providing valuable feedback. (Acknowledged individuals are employees of Microsoft Corporation unless otherwise noted.)

Micky Balladelli (Avanade Inc.), Dung Hoang Khac (Compaq Global Services), Vic Gupta, Shaun Hayes, Michael Kostersitz, Doug Lindsey, Alain Lissoir (Compaq Global Services), Tony Mosunmade, Tony Redmond (Compaq Global Services), Matthijs ten Seldam, Paul Thompson (NetIQ Corporation), Ruud van Velsen, Eddie Hanif, Frank Plawetzki, Connie Shih, Steve Towle, Christopher Westpoint.

Lab Staff: Robert Thingwold and David Meyer

Lab Partners: Compaq, Inc. and Cisco Systems

Top of page