Microsoft Consulting Services

Iowa State University

Windows 2000 Active Directory Functional Specification

June 18, 2001

Table Of Contents

Iowa State University

Windows 2000 Active Directory – Functional Specification

1. Executive Summary

Iowa State University (ISU) is a large state regents institution located in Ames, Iowa. ISU is comprised of 9 colleges, 59 academic departments, and 20 support units. This includes over 25,000 students and approximately 6000 employees serving as faculty, professionals, and staff. Two of the largest IT providers on campus are the Office of Academic Information Technologies (AIT) and Administrative Data Processing (ADP).

As part of a larger project focused on providing campus students, faculty, and staff with a ubiquitous computing environment, in conjunction with several departmental IT providers, AIT and ADP have begun design of a model campus-wide Windows 2000 Active Directory architecture. AIT and ADP have engaged Microsoft Consulting Services to aid in revisiting the Active Directory architecture to ensure that the architecture is optimally deigned to meet stated goals and objectives. Iowa State University and Microsoft Consulting Services will focus on these efforts over a 12-week period beginning in April 2001.

This document describes the Active Directory architecture developed by the Iowa State Active Directory design team and Microsoft Consulting Services.

2. Windows 2000 Architecture

2.1 Active Directory Namespace Background

Active Directory plays many roles, from being the backbone of distributed security to providing a service-publishing framework. Active Directory provides a central service for administrators to organize network resources, to manage users, computers, and applications; and to secure Intranet and Internet network access.

Designing the Windows 2000 Active Directory namespace is a crucial step in the overall Windows 2000 design process. The Active Directory namespace is the building block for all design decisions. In the current release of Windows 2000, once the root of the namespace is selected and created it cannot be changed without a complete reinstallation.

2.2 Forest Architecture

2.2.1 Background Information

A Windows 2000 forest is a logical set of domain trees that together may form either a contiguous or disjoint namespace. A forest is named the same as the first domain that is installed in the forest. Two-way trusts exist between all parent and child domains in a forest tree and between peer top-level domains in the forest. Trust relationships within the forest are transitive; hence all domains trust all other domains in the same forest. Trusts between domains in separate forests are one-way and non-transitive.

All domains in a forest have the following in common:

- Schema – The formal definition of objects, properties, and their relationships
- Global Catalog – Partial replica of all objects in the forest

- Sites & Services Configuration – Enterprise configuration information including physical sites and enterprise network services

When designing a forest architecture, the recommended approach is to start with a single forest and add additional forests only if absolutely necessary, providing very specific justification for the creation of each forest. In almost all cases, it is envisioned that a single forest will work for even the largest organizations. However, for completeness, possible reasons for creating additional forests are itemized below:

- Multiple Schemas - Some organizations may require multiple directory schemas. For example, an organization may have a quasi-official relationship with a joint venture. Although the level of trust between the two organizations is high, thereby warranting some type of relationship between their Windows 2000 deployments, each organization may have unique schema requirements. In this case, multiple Windows 2000 forests can be deployed, each with their own schema, and perhaps with manual trusts established between the forests.
- Grass Roots Implementations, Mergers & Acquisitions - The initial release of Windows 2000 does not support merging multiple forests into a single forest. If an organization has multiple forests due to grass roots implementations, mergers, or acquisitions, it would not currently be possible to merge these into a single forest. They can, however, be logically linked with one-way non-transitive trusts.
- Politics – The “Enterprise Admins” and “Schema Admins” security groups exist in the forest root domain and have special permissions throughout the forest. Lack of trust between IT administrative groups may prevent agreement on membership of these groups. This issue can be mitigated by the implementation of multiple forests.

Some of the advantages of a multiple forest implementation are:

- Each organization can maintain its own directory schema, meaning that schema modifications are more localized and have less of an impact on the greater organization
- Divestitures can be handled more easily. A spin-off of an organization contained in a separate forest is a trivial task. Currently there are no tools or methods available to “split” a forest into multiple forests.

Some of the disadvantages of a multiple forest implementation are:

- No single enterprise directory exists in a multi-forest environment. Each forest will have its own independent directory. Directory synchronization between Forests is not possible in the initial release of Windows 2000 Active Directory.
- No cross-business security context exist by default, although Windows NT 4.0 style trusts can be explicitly added
- Transitive trust relationships do not exists between forests
- Mergers of organizations will require forest reorganization, which is a very significant operation in the initial release of Windows 2000
- User and resource transfers between forests is a significant operation
- No single enterprise-wide Global Address List (GAL) for Exchange 2000 exists in a multi-forest environment. For MAPI clients that have an Exchange 2000 mailbox, the GAL will be the Active Directory global catalog. Since the global catalog is specific to a forest, there will be only one global catalog, and hence one GAL, for each forest.

- Message flow between Exchange 2000 servers in the same forest is automatic and requires no additional configuration or connectors.
- The forest hides the structure of the network and the directory from end-users. The implementation of multiple forests requires end-users to have a greater understanding of the underlying directory structure.

2.2.2 Iowa State University Forest Design

The Active Directory design team recommends that Iowa State University implement a single Active Directory forest. The campus forest will be architected and policies developed in a manner that permits and encourages future campus adopters of Windows 2000 and Active Directory to participate in this forest.

Prior to recommending a single campus-wide forest, the Active Directory design team evaluated multi-forest and single forest options. The design team determined that a single forest model will best meet the needs of the university based on known business requirements and that potential benefits of this model outweigh concerns. Key factors in this decision include the reliance of Exchange 2000 on the forest architecture, the transitive trust relationships that exist between domains in the forest, and the ability to make the underlying network and directory structure transparent to end-users through the use of a single forest. Additionally, only the single forest architecture complies with AIT's desire to implement single identify for faculty, staff, and students. A more complete listing of benefits and drawbacks of a single campus forest can be found in Appendix A.

A more thorough review of the forest planning information used as a basis for discussion between the Active Directory design team and Microsoft Consulting Services (MCS) is included in the document titled "Iowa State University, Windows 2000 Active Directory Forest Planning".

2.3 Domain Architecture

2.3.1 Background Information

A Windows 2000 domain is identified by a boundary of security and replication, and is distinguished by a unique namespace.

Multiple domains can be joined hierarchically in a contiguous namespace to form a tree. Multiple domains and/or multiple trees can be joined to form a forest. All domain controllers in a forest share a common global catalog, configuration container, and schema.

When designing domain architecture, the recommended approach is to start with a single domain and add additional domains as necessary, while providing justification for each additional domain. Possible reasons for creating additional domains are itemized below:

- Unique Policies - Different locations or organizations may have unique security and administration policies, which cannot be reconciled. In these instances, separate domains are required.
- Network Traffic - By creating multiple domains, network traffic may be reduced because only changes to the global catalog, configuration container, and schema are replicated between domains. Within a domain the entire domain database, as well as the contents of the SYSVOL partition, which includes logon scripts and group policy objects, are replicated to all domain controllers in the domain. Although the entire domain database is replicated to each domain controller in a domain, the impact on link traffic is mitigated. After initial replication, only changed properties are replicated. However, the primary method for controlling

replication traffic in Windows 2000 is with the use of “sites”, between which traffic can be compressed and controlled.

- Network Connectivity - If the network does not support RPC communication between locations, the SMTP Inter-Site Connector must be used. However, the SMTP Inter-Site Connector cannot be used to replicate the domain-naming context. The result is that multiple domains would be required.
- Capacity - If the number of objects in a domain is likely to significantly exceed one to two million objects, a separate domain should be considered.
- In-place Upgrade of Existing Domains - One migration strategy is to upgrade existing domains in-place with a plan of eventually consolidating domains.
- Administrative Considerations (Politics) - Decentralized organizations often have different groups of administrators responsible for managing users and resources throughout the organization. These administrators may not want to share control of a domain and its resources. For example, a domain administrator can always gain access to and take ownership of objects within the domain. However, a delegated organizational unit (OU) administrator does not have this ability.

2.3.2 Iowa State University Domain Design

The Active Directory design team recommends that Iowa State implement an architecture based on a single-domain, as shown in the following figure. Additionally, requests for “Unique Requirement” domains will be evaluated on a case-by-case basis, and only be granted when a technical requirement cannot be met with the single domain.

Note: *Domain names listed in this figure are only representative, and not necessarily the actual domain names that will be configured. Additionally, no particular namespace is identified in this diagram.*

Prior to recommending the single-domain model with potential support for “Unique Requirement” domains shown in the previous figure, the Active Directory design team evaluated several domain models ranging in size from a single domain for all university users and resources to separate domains for any college or campus group requesting a domain.

The design team found no technical limitations that would prevent implementation of the single domain model. Additionally the single domain model can provide independent IT providers nearly the same level of autonomy that they experience today with multiple Windows NT 4.0 domains, while eliminating several areas of duplicated effort. Potentially, one department or another may have a unique technical requirement that cannot be met with the single campus domain, yet the department still wishes to realize the many benefits of participating in the campus Active Directory forest. Requests for these “Unique Requirement” domains in the campus forest will be evaluated on a case-by-case basis and only be granted if a unique requirement cannot be met with the single domain. Implementation of a single domain model requires participating departments to agree on a limited number of domain wide configuration settings, such as password complexity requirements. While specific policies have not yet been determined, it appears that agreement on configuration of these settings can be reached.

The Active Directory design team also evaluated a unit domain model that provides a separate domain in the campus forest for each qualified requesting organization. While technically, the

unit domain model does provide the absolute greatest level of unit autonomy in a single forest, it also presents its own unique challenges while providing few benefits. As the number of domains in the forest grows, the difficulty and complexity of managing and troubleshooting the forest grows. The use of multiple domains would also make it more difficult to achieve the goal of single identity for faculty, staff, and students. Additionally, administrators of any domain in the forest must be equally trusted as poor security management and domain controller maintenance may affect services and workload for the entire enterprise.

After evaluating several models, the Active Directory design team recommended a single-domain with support for "Unique Requirement" domains. The design team feels that this model will provide the autonomy required by various units, while still providing the security and stability required by the larger enterprise. It is also recommended that ISU form an Enterprise Services Team with the responsibility of developing security policies and delegating top-level organizational units. Microsoft Consulting Services recommends that additional domains only be created when a specific justifiable need exists. A more complete listing of benefits and drawbacks of a single campus domain can be found in Appendix A.

A more thorough review of the domain planning information used as a basis for discussion between the Active Directory design team and Microsoft Consulting Services (MCS) is included in the document titled "Iowa State University, Active Directory Domain Planning".

The following sections outline the purpose of domains that will initially be included in the campus Active Directory forest.

2.3.2.1 IASTATE.EDU Domain

In this domain design, all user accounts and Windows based resources for Iowa State University will reside in a single Windows 2000 domain. This domain is also the root domain (first domain installed in the forest), and will be named IASTATE.EDU *(W2K)* or IASTATE *(Legacy).* The organizational unit structure defined in section 2.4.2 will be developed to facilitate delegated administration of user accounts and resources within the IASTATE.EDU domain. Enterprise Admin, Schema Admins, and Domain Admins for the IASTATE.EDU domain will contain a limited number of enterprise administrators from AIT and ADP. Policies and usage of Enterprise & Schema Admin privileges will be developed, documented, and approved by the larger IT community.

This Windows 2000 domain, IASTATE.EDU had already been previously created through the installation of new Windows 2000 domain controllers. Additionally, a number of production resources already reside in the existing IASTATE.EDU domain. Prior to recommending leaving this domain at its present location in the DNS namespace, the Active Directory design team evaluated alternatives, such ADS.IASTATE.EDU. While some potential advantages may exist for moving Active Directory down a level in the DNS namespace, such as alternative methods of interoperability with the existing UNIX environment, the design team determined that these potential advantages were minimal in comparison to the effort and disruption of services necessary to facilitate this move.

Only default trust relationships will exist between the IASTATE.EDU domain and other domains. No trust relationships will be implemented between the IASTATE.EDU domain and any other domain, other than temporarily for migration purposes.

As noted previously, a number of policies have domain wide scope, and cannot be set individually at the organizational unit level. These include password policies, account lockout policies, and Kerberos policies. It is believed that default Kerberos polices will be acceptable to all participant of the Active Directory forest. The Active Directory design team has proposed the following settings for domain wide password polices, which are currently under review:

Password Length	8 character minimum
Password Complexity	Must contain at least 1 character form two separate character classes (lowercase, uppercase, numeric, special)

New Password	Cannot be the same as the previous password
Password Change	Password changes must occur at least once every 90 days
Account is Locked Out	After 10 consecutive attempts in X amount of time
Account Lockout Duration	24 hours

2.3.2.2 Unique Requirement Domains

In this domain design, a limited set of "Unique Requirement" domains may be implemented. These will exist only in cases where a specific technical requirement cannot be met in the single domain. Requests for "Unique Requirement" domains will be evaluated on a case-by-case basis. Microsoft Consulting Services recommends that an Enterprise Services Team document an objective set of principals to guide the evaluation of "Unique Requirement" domains. In these instances user credentials for any particular user will still only reside in EITHER the IASTATE.EDU domain OR the "Unique Requirement" domain.

2.4 Organizational Unit Design

2.4.1 Background Information

Organizational units are containers used to organize objects within a Windows 2000 domain. A best practice is to implement organizational units for the purposes of delegated administration and application of group policy and not simply to create structure. It is not necessary to create an organizational unit structure that is aesthetically pleasing or that mirrors the organizational structure of the institution. If security doesn't need to be delegated and if group policy objects are not applied to a collection of objects, there is no need to create an organizational unit. Important points to consider about the use of organizational units include the following:

- Organizational units can be nested within other organizational units
- Group policies can be applied to organizational units and filtered by access control lists
- Organizational units can be used to restrict access to Active Directory objects
- Directory and file level permissions are granted to security groups and not to organizational units
- Unlike forests and domains, organizational units can be easily added, removed, renamed, and moved
- Active Directory objects can be easily moved between organizational units
- Active Directory object friendly names must be unique within an organizational unit but not across organizational units

2.4.2 Iowa State University Organizational Unit Design

The primary objective behind the Iowa State University Organizational Unit base architecture is delegated administration. The base organizational unit structure for ISU will include 4 tiers as shown in the following diagram:

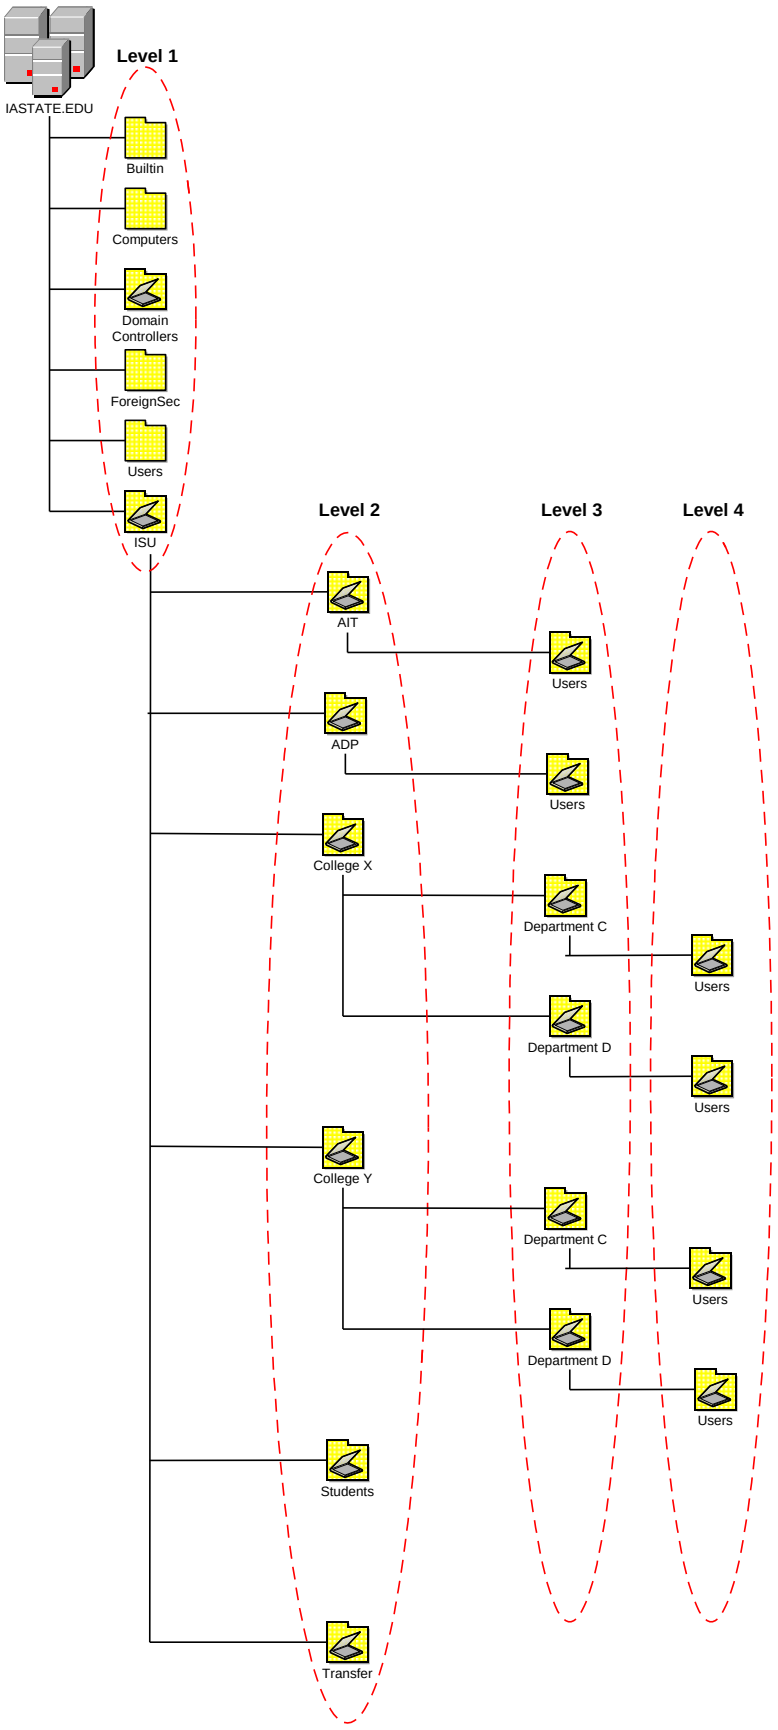
Level 1
IASTATE.EDU
Builtin
Computers
Domain Controllers
ForeignSec
Users
ISU
Level 2
Level 3
Level 4
AIT
Users
ADP
Users
College X
Department C
Users
Department D
Users
College Y
Department C
Users
Department D
Users
Students
Transfer

2.4.2.1 First Level Containers

The first level of the structure contains both containers and organizational units. A standard set of containers is created during the installation of a new Windows 2000 domain. Containers are less functional than organizational units because containers cannot be nested and group policy objects cannot be applied to containers. The following table contains the names of the standard Windows 2000 containers and their function.

Containers	Function
Users	Contains domain administrator accounts
Computers	Unused
Built-in	Contains built-in groups
ForeignSecurityPrincipals	Used by Active Directory
LostAndFound	Used by Active Directory
System	Used by Active Directory

Windows 2000 Standard Containers

2.4.2.2 First Level Organizational Units

As noted above, the first level also includes organizational units. The following table contains the names and purpose of first level organizational units in the ISU environment.

First Level OUs	Contains
Domain Controllers	Contains all domain controller computer accounts. Allows a common group policy to be applied to all domain controllers.
ISU	Contains additional organizational units, allows domain wide policies to be applied to all non-domain admin accounts, and functions as a placeholder for additional organizational units.

First Level Organizational Units

2.4.2.3 Second Level Organizational Units

Second level organizational units provide for delegated administration based on organization or college, as show in the following figure.

Second Level OUs	Contains
AIT	Contains additional organizational units that will contain AIT user accounts and Windows based resources.
ADP	Contains additional organizational units that will contain ADP user accounts and Windows based resources. May also contain user accounts and Windows based resources for ADP customers.
College X	Separate organizational unit for each college. Contains additional organizational units for departmental delegation.
Students	Contains all user accounts for students.
Transfer	Organizational unit used to facilitate inter-collegiate/department account transfers.

Second Level Organizational Units

2.4.2.4 Third Level Organizational Units

Third level organizational units provide for delegated administration based on collegiate departments and default user account containers for AIT and ADP, as show in the following figure.

Containers	Contains
AIT/ADP Users	Default container for AIT and ADP user accounts.
Department A…	Separate organizational unit for each department. Contains additional organizational units for departmental user accounts and Windows based resources.

2.4.2.5 Fourth Level Organizational Units

Fourth level organizational units provide for delegated administration based on collegiate departments and default user account containers for department, as show in the following figure.

Containers	Contains
Department A … Users	Default container for departmental faculty and staff user accounts.

The organizational unit structure outlined above is the default structure that will be implemented by the enterprise services group. Additional organizational unit may be developed at the collegiate and department level at the discretion of the individual colleges and departments.

User accounts for real university recognized personnel will be automatically populated into one of the “Users” or “Students” containers via an AIT developed process. If an individual college or department would like their faculty/staff to be automatically created in a different OU they may create a new OU and notify AIT of the requested change, so that the automatic process can be modified to reflect the change.

2.5 Site Design

2.5.1 Background Information

Active Directory utilizes the concept of “sites”, as does Systems Management Server 2.0 and Exchange 2000. A site is a collection of IP subnets with fast and reliable connectivity. Sites are used to control the direction, schedule, and frequency of replication for the Active Directory, Active Directory integrated DNS, and SYSVOL information. In addition to controlling replication traffic, sites are used to find resources that are closest to the requestor. For example, when a logon request is looking for a domain controller for authentication or when a Distributed File System replica is being chosen, site affinity will be checked to find a resource that exists in or near a requestor’s site.

2.5.1.1 Replication Contexts Protocols

Active Directory Replication is divided into three naming contexts (NC).

NC=Domain Replication (AD)

- RPC over IP

NC=Global Catalog Replication

- RPC over IP
- SMTP (Inter-Site only)

NC=Configuration/Schema Replication

- RPC over IP
- SMTP (Inter-Site only)

In addition to the Active Directory information, System Volume (SYSVOL) information is also replicated. The SYSVOL contains:

- SYSVOL share
- NETLOGON share
- Windows 9x & Windows NT system policies
- Windows 2000 group policy objects
- User logon and logoff scripts

2.5.1.2 Intra-Site Replication

Replication within a site (Intra-Site) is defined by the following characteristics

- RPC based
- No data compression
- Notification replication process - 5 minute pause interval (can be set in the registry)

2.5.1.3 Inter-Site Replication

Replication between sites (Inter-Site) is defined by the following characteristics

- RPC communication or SMTP communication (Inter-Domain only)
- Data compression - RPC over IP offers approximately 85% - 90% compression
- Scheduled replication - Replication schedules can be set to any frequency, depending on the data consistency required and the need to manage peak bandwidth usage.

2.5.2 Iowa State University Site Design

Iowa State University primarily consists of a LAN environment with the exception of the Campus Book Store, Extension offices, and a few buildings located in Ames, but not on the main campus. Network bandwidth throughout the main campus is no less than 10Mb.

The ISU Active Directory design will initially consist of two sites. The first site will include all subnets other than the Campus Book Store Subnets. Due to the abundance of network bandwidth throughout the main campus environment there is little or no benefit to replication compression and site affinity. The second site will include only Campus Book Store Subnets. This site has been created due to limited bandwidth between the book store and the main campus and the desire for DFS site affinity by a department occupying the book store.

Initially, Extension offices located throughout the state of Iowa will be part of the main campus Active Directory site and will not contain domain controllers. If in the future it is determined that domain controllers will be necessary at Extension offices, it is recommended to create separate sites for these locations.

2.6 Services Locations

The following sections discuss recommendations for the placement of Windows 2000 domain controller servers, global catalog servers, and DNS servers in the Iowa State University environment.

2.6.1 Domain Controller Server Placement

Based on estimated potential Active Directory usage in the next 2 years, it is recommended that the IASTATE.EDU domain include five domain controllers for fault tolerance and performance. It is also recommended that four of these domain controllers reside in separate key campus hub locations in physically secure rooms and that network access to the security groups within this domain be tightly controlled. Additionally, it is recommended that one domain controller reside in a physically secure room in the Campus Book Store.

Recommendations for domain controller hardware and configuration are included in Appendix B.

2.6.2 Global Catalog Server Placement

In a single domain Active Directory forest no additional replication traffic is incurred from global catalog replication. Because the Iowa State University Active Directory forest will consist of a single domain, it is recommended that all domain controllers in the IASTATE.EDU domain be configured as global catalog servers. This will allow greater service availability while incurring no additional replication traffic and require only minimally more domain controller disk space.

If there is a period of time in which multiple domains will be present in the ISU Active Directory forest, it is recommended that all domains in the forest include at least one domain controller not configured as a global catalog server and that the Infrastructure Master FSMO role reside on these servers.

2.7 Operations Masters

Active Directory defines five operations master roles: schema master (SM), domain naming (DM) master, relative identifier (RID) master, primary domain controller emulator (PDCE), and infrastructure master (IM). The schema master and domain naming master are per-forest roles, meaning that there is only one schema master and one domain naming master in the entire forest.

2.7.1 Schema Master (Forest Level FSMO)

The domain controller that holds the schema master role is the only domain controller that can perform write operations to the directory schema. Those schema updates are replicated from the schema master to all other domain controllers in the forest.

In the ISU environment, it is recommended to place the schema master on the same domain controller as the domain naming master, and that this server be configured as a global catalog server.

2.7.2 Domain Naming Master (Forest Level FSMO)

The domain controller that holds the domain naming master role is the only domain controller that can add new domains to the forest and remove existing domains from the forest. In Windows 2000 Active Directory, the domain naming master must reside on a global catalog server.

In the ISU environment, it is recommended to place the domain naming master on the same domain controller as the schema master, and that this server is configured as a global catalog server.

2.7.3 RID Pool Master (per Domain FSMO)

A new security principal object (User, Group, or Computer) can be created on any domain controller. However, after creating several hundred security principal objects, a domain controller must communicate with the domain controller holding the domain's RID master role before creating the next security principal object. Then, another several hundred security principal objects can be created, and when this set of objects has been created, the process of contacting the RID master repeats. If a domain controller's RID pool is empty, and the RID master is unavailable, you cannot create new security principal objects on that domain controller.

In the ISU environment, it is recommended to place the RID master on the same domain controller as the PDC emulator.

2.7.4 PDC Emulator (per Domain FSMO)

The domain controller holding the PDC Emulator provides backward compatibility to down-level backup domain controllers (when running in Mixed Mode). The PDC emulator also serves other roles including time synchronization and password latency control.

In the ISU environment, it is recommended to place the PDC emulator on the same domain controller as the RID master.

2.7.5 Infrastructure Master (per Domain FSMO)

The domain controller holding the infrastructure master role for a security group's domain is responsible for updating the cross-domain group-to-user reference to reflect any modifications to the user's name. The infrastructure master updates these references locally and uses replication to bring all other replicas of the domain up to date. If the infrastructure master is unavailable, these updates are delayed.

Unless all domain controllers in the forest are configured as global catalog servers, it is recommended to place the infrastructure master on a non-global catalog domain controller.

2.8 Time Services

Time synchronization is important to the Windows 2000 Active Directory. Windows 2000 Active Directory relies on time stamps to perform certain conflict resolution while performing multi-master replication. Synchronized time is also critical in Window 2000 because the default authentication protocol (Kerberos V5) uses timestamps as part of the authentication ticket generation process.

2.8.1 Time Services Design

Microsoft provides an NTP compliant time service on all versions of Windows 2000. The time service is called W32TIME.

The Windows Time Synchronization service (W32Time) is a fully compliant implementation of the Simple Network Time Protocol (SNTP) as detailed in IETF RFC 1769.

2.8.2 W32Time Basic Operation

Windows 2000 Professional workstations and Windows 2000 Member Servers, referred to as the *client*, synchronize their local time to the authenticating domain controller. By default the client performs a check of local time every eight hours. If the client’s local time is off from the target time

by more than two seconds, the interval check period is divided in half. The checks continue and the interval is reduced until the client is within two seconds of the target time. Once the client is within two seconds, the interval period starts to increase by double until a steady state period is reached and the client is within two seconds of the target. The minimum interval period is 45 minutes and the maximum is eight hours.

2.8.3 Time Convergence Hierarchy

By default, time synchronization follows the hierarchy of a Windows 2000 Active Directory Forest. Within a domain, the PDC emulator (PDCE) is the master time source. All domain controllers in the domain synchronize their time with the domain PDCE. Each child domain synchronizes time with its parent domain. This is accomplished by the PDCE for the child domain synchronizing with the PDCE for the parent domain. The following figure illustrates the time synchronization hierarchy for the IASTATE.EDU domain and any potential child domains.

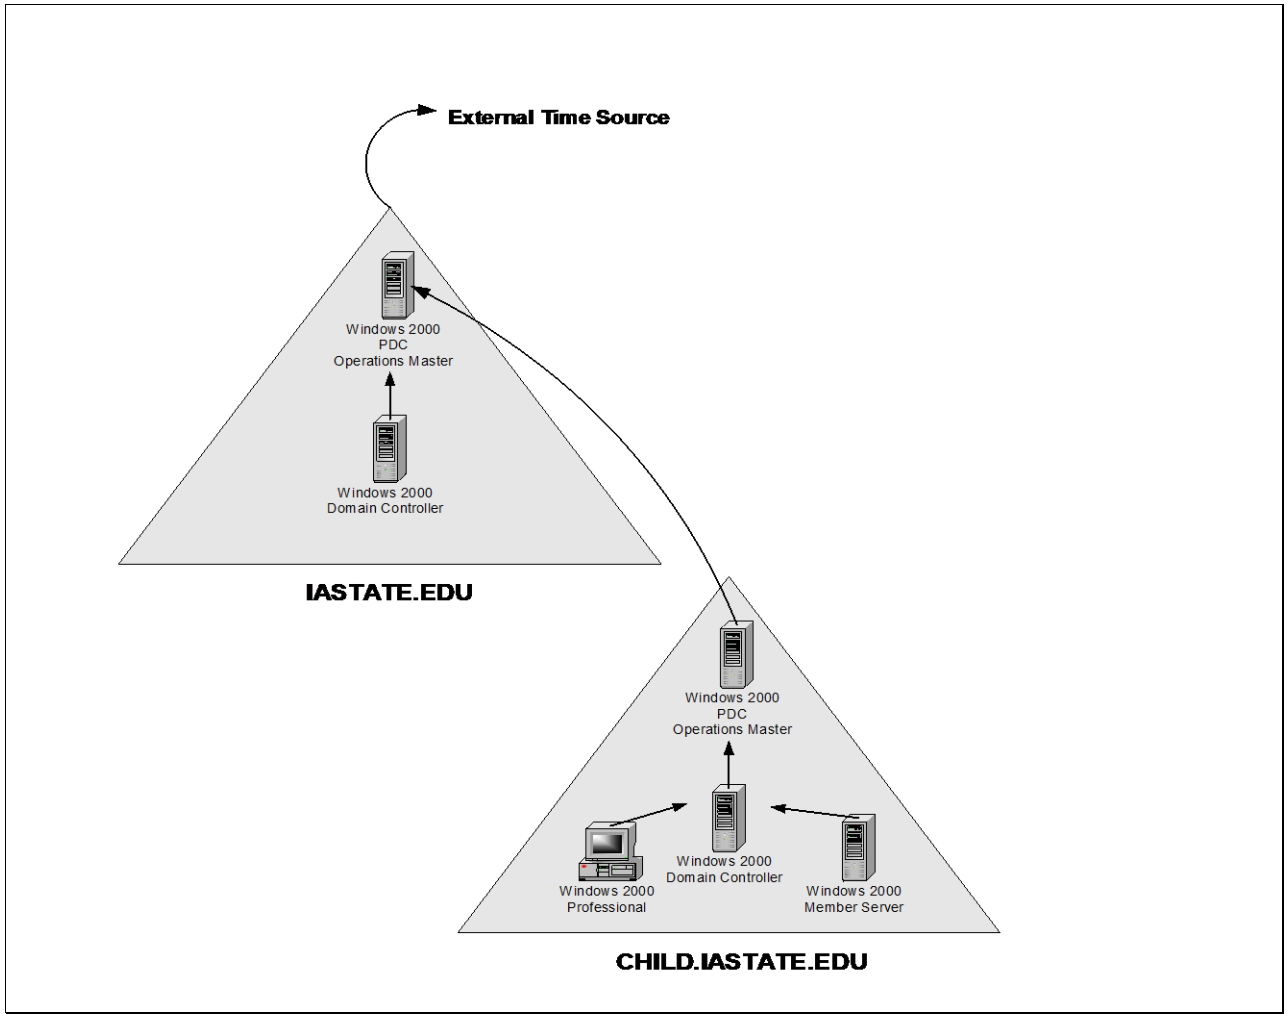

W32Time Time Synchronisation Hierarchy

In the campus forest, the PDCE in the IASTATE.EDU domain is authoritative for the forest, and can be manually set to synchronize with an outside time source. All time synchronization within

the forest is automatically configured and enabled, except for the PDC Operations Master in IASTATE.EDU.

Time for the PDC Operations Master in IASTATE.EDU should be configured manually using the NET TIME command:

```
NET TIME /SETSNTP:<ntp server list>
```

3. DNS Architecture

The Domain Name System (DNS) is a set of protocols and services on a TCP/IP network that allows users of the network to utilize a hierarchical user-friendly name when looking for other hosts (that is, computers) instead of having to remember and use their IP addresses.

Windows 2000 Active directory places specific requirements on the DNS platform, namely Dynamic DNS support and SRV record support. In addition, the Windows 2000 operating system places a greater emphasis on DNS for name resolution than did previous Windows operating systems. DNS has replaced WINS as the primary name resolution service for Windows 2000. DNS name resolution and registrations are critical to the successful operation and user experience in a Windows 2000 environment.

3.1 Iowa State University DNS Architecture

Iowa State University currently hosts the majority of Iowa State DNS domains on a BIND platform. Iowa State will augment the existing BIND DNS environment with Windows 2000 DNS to provide support for the deployment of Windows 2000 and Active Directory. Deployment of Active Directory will leverage Windows 2000 DNS and will require minimal modification to the existing DNS environment. Four new DNS domains; _tcp.iastate.edu, _msdcs.iastate.edu, _sites.iastate.edu, and _udp.iastate.edu, will be created and delegated to Windows 2000 DNS to support the Active Directory IASTATE.EDU domain. The following figure outlines relevant portions of the current DNS environment and resulting environment following the deployment of Active Directory.

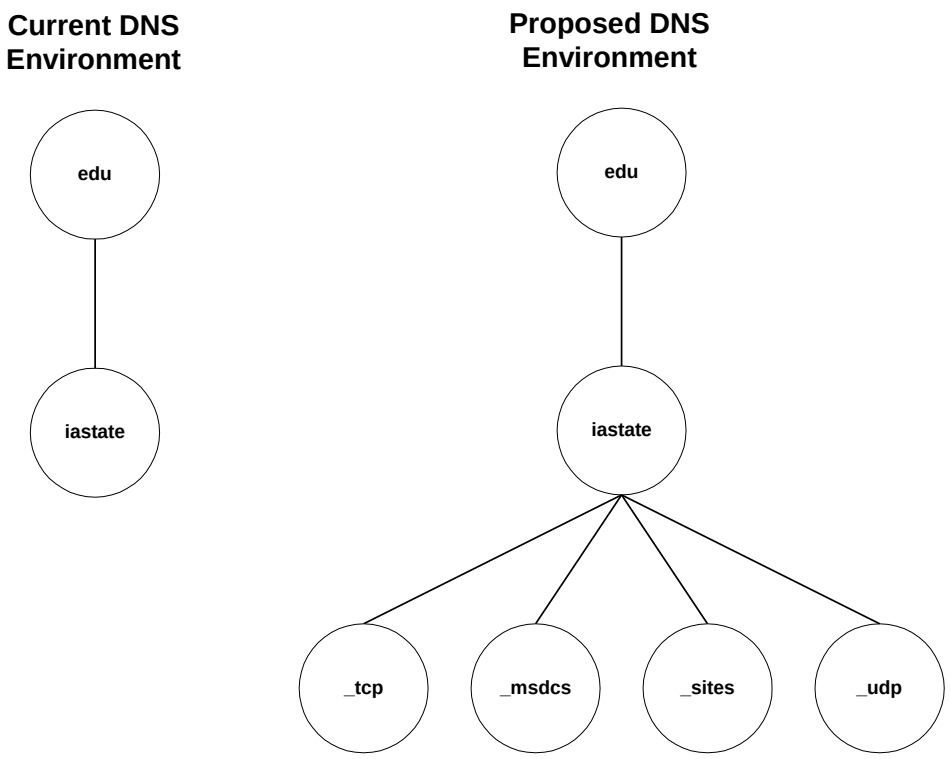

3.1.1 IASTATE.EDU DNS Domains

Primary DNS for the iastate.edu DNS domain will continue to be managed by AIT and hosted on existing DEC Alpha servers running BIND version 8.2.3. Dynamic updates will not be enabled on these servers. These servers will be configured to delegate authority for the _tcp.iastate.edu, _msdcs.iastate.edu, _sites.iastate.edu, and _udp.iastate.edu DNS domains to Windows 2000 DNS servers.

3.1.2 _* DNS Domains

Primary DNS for the _tcp.iastate.edu, _msdcs.iastate.edu, _sites.iastate.edu, and _udp.iastate.edu DNS domains will be managed by the Active Directory enterprise services team and hosted on Windows 2000 domain controllers/DNS servers in the IASTATE.EDU Active Directory domain. These domain controllers/DNS servers will be configure for Active Directory Integrated DNS, will permit secure dynamic updates, and will forward unresolved requests to DNS servers hosting the iastate.edu DNS domain.

3.1.3 Client Hostname Registration

In a Windows 2000 environment, the ability to manage and access resources on desktop clients will become increasingly reliant on the ability to locate clients via DNS rather than WINS. Therefore, in the future it may become desirable to dynamically update client host name records in DNS rather than statically maintaining this information as it is done today throughout much of campus. ISU currently has a process for automatically registering client host names in DNS for client systems residing in residence halls. This system provides automatic registration for various client operating systems including Windows based systems. In the future, this automatic registration process will be expanded campus wide.

4. Object Management

Within Active Directory, users and resources are represented as user objects and resource objects such as computers, printers, and file shares. One of the key decisions that must be made during the design process is to determine who will manage these objects within Active Directory.

4.1 Resource Objects

Each Windows based resource may be represented as an object. Each object will have various attributes associated with it that define configuration settings for the resource. Common resource objects that exist in Active Directory include computer objects which are associated with a particular Windows NT or Windows 2000 server or workstation, printer objects which are associated with a particular Windows based print server, and file share objects which are associated with a particular Windows based file share.

In the case of resource objects, a single IT provider normally owns supports, manages, and provides security for the particular resource. Therefore, in the Iowa State campus Active Directory, which ever IT provider owns the resource will be fully responsible for management of the associated object in Active Directory, all attributes associated with the object, and access to the resource. No other campus IT provider will have access to modify the object unless granted permissions by the IT provider that owns the object. Resource objects will be created by the IT provider in an Active Directory organizational unit supported and managed by the respective IT provider. Creation, deletion, and modification of resources objects will normally be done by the object owner via native Windows 2000 tools such as Active Directory Users and Computers.

4.2 User Objects

Within Active Directory, each user of network resources is represented as a user object. Each user object will have various attributes associated with it such as username, password, department, phone number, mailbox location and home directory. It is this user object (username and password) that a user will utilize when accessing network resources such as file shares, printers, and Exchange based collaboration resources. In addition to “real” people, user objects will also be necessary for application service accounts and Exchange schedulable resources such as conference rooms.

One of the primary objectives when developing the campus Active Directory architecture is to ensure that one and only one user object exists for each user of network resources. By creating only one account for each “real” person, the user can be granted seamless access to Windows based resources throughout the campus. If multiple accounts exist for a person, the user must be aware of the underlying network structure and know which account has access to which resources.

In the campus environment it is often difficult to determine which IT provider is responsible for management of any individual Active Directory user object which represents a “real” person, as the person is likely to access resources and services from multiple providers. However, it is

important to understand that regardless of which IT provider manages any individual Active Directory user object, all IT providers may grant and restrict the user's access to resources owned and managed by the respective IT provider.

4.2.1 Faculty and Staff User Objects

In the case of user objects for faculty and staff, which IT provider should manage the Active Directory user object is normally apparent, based on which department or organization the faculty or staff is a member. Based on university records, each faculty and staff member user object will be assigned to a particular IT provider for management. For instances in which the faculty or staff member has a dual appointment, the Active Directory user object will be defaulted to one of the two IT providers based on university records. If the faculty or staff member determines a need to have his/her user object managed by a different IT provider, it is up to the faculty/staff member to have this changed. This may be done by changing the central meta-directory information or via a secure software facility the user may use to select the "preferred" IT provider from the list available within the central meta-directory. The exact mechanism has not been determined yet.

Each IT provider will have full control of user objects for "their" faculty and staff, with the exception of university owned attributes defined in the following table. The IT provider may manage faculty and staff user objects using native Windows 2000 tools such as Active Directory Users and Computers, or may develop and utilize custom tools and scripts. Population and removal of faculty and staff user objects from departmental organizational units will be automated based on university records. There will be no need for organizational unit administrators to manually create or delete user objects for faculty and staff.

A number of attributes associated with faculty and staff user objects contain general information and have no direct bearing on user functionality in the Windows environment. This information would likely be set no differently no matter who directly manages any individual user object. In the Iowa State environment this information will be kept in synch with Iowa State's meta-directory to ensure that this information is consistent across all Iowa State directory platforms. Therefore, university owned attributes for faculty and staff user objects will not be available for modification by OU administrators. Attributes that will be supplied by the meta-directory are listed in the following table.

University Owned Attributes
Proper Name
On/Off campus addresses
On/Off campus phone numbers
Logon name (including downlevel)
Account options (expiration, password complexity, etc)
Title
Department
Company

4.2.2 Resource User Objects and Groups

As noted previously, Active Directory user objects are also used to represent some application services, Exchange based schedulable resources such as conference rooms, and temporary personnel that are not university recognized. In these cases, the individual IT provider responsible for the application service, schedulable resource, or non-university recognized personnel may create representative Active Directory user objects in an organizational unit managed by the IT provider. The IT provider may create, delete, and manage resource user objects using native Windows 2000 tools such as Active Directory Users and Computers, or may develop and utilize custom tools and scripts.

To ensure naming conflicts do not occur between IT provider created resource user objects and university generated user objects for "real" personnel, naming conventions will be developed and enforced. At the time of this writing, the exact convention has not been determined. However, the convention will dictate that IT provider created resources user objects MUST end with a special character, such as the ! symbol and university generated user objects may NEVER end with the special character. If the IT provider creates a resource user object that does not adhere to the convention, the user object will be automatically renamed by a nightly process to adhere to the convention.

Similar to user objects, the university will automatically generate a number of groups in Active Directory via university records. Examples of automatically generated groups include class lists. To ensure that naming conflicts do no occur between IT provider created groups and university-generated groups, a convention similar to the user naming convention will be adopted and enforced for group naming.

4.2.3 Student User Objects

Management of student user objects presents unique challenges as ownership of students can often not be easily determined as many students have dual majors and most students will spend much of their time accessing resources provided by multiple IT providers. Additionally, a small subset of the student population is also staff in one or more departments. Therefore, Iowa State will utilize a concept of centralized management of student user objects, with a provision for departmental sponsored accounts.

Under this model, all student Active Directory user objects will be automatically generated and placed into an unmanaged location. No modifications will be permitted to these student Active Directory user objects. Attributes such as "home directory", "logon script", and "profile path", will be populated with variables, so that departmental IT providers can provide home directories, logon scripts, and roaming profiles to students without directly accessing or modifying the student user objects. Even though student user objects will reside in an unmanaged organization unit, any IT provider may grant or deny the student access to resources owned by that IT provider.

For instances in which an IT provider desires to provide some functionality to a student that cannot be achieved without modification to the unmanaged student user object, the provider may create a second user object for the student in the IT provider organizational unit. An example of a function that could not be achieved without modification to the student user object would be the desire to specify and Exchange mailbox for the student. This "sponsored" user account must adhere to the special character naming convention previously outlined. The IT provider may utilize native Windows 2000 tools to create, delete, and manage sponsored user accounts.

5. Appendix A – Benefits and Drawbacks

The following table outlines benefits and drawbacks of participation in the campus Active Directory forest for end-users, departmental administrators, enterprise administrators, and Iowa State University.

KEY:

EU = End-User

DA = Departmental Administrators

HD = Help Desk

EA = Enterprise Administrators

ISU = Iowa State University ($$)

EAD = Enterprise Application Developers

Single Forest Architecture with Single Identity - Benefits	Who
End-Users may access resources throughout the university with a single account ID and password. This results in easier access to resources for end-users and reduced help desk calls.	EA, HD
End-Users may access a consistent set of resources regardless of campus location. This results in greater efficiency and productivity for users of network resources.	EA
End-Users (faculty & staff) may collaborate via Exchange 2000 with all faculty and staff. This allows end-users to efficiently collaborate on projects, schedule meetings, schedule resources, participate in instant messaging, and video conferencing.	EU, ISU
A single forest schema facilitates easier development of enterprise applications based on Active Directory	EAD, ISU
Within a single forest, there is little need to create manual Windows NT 4.0 style trust relationships between domains. This reduces administrative overhead for departmental administrators and facilitates easier access to resources.	DA, EU
A single forest facilitates automated creation of single identities for faculty, staff, and students. This means departmental administrators no longer must manually create user accounts. This eliminates duplication of administrative effort across departments and frees departmental administrators to focus on other tasks.	DA, ISU
Single identity allows user accounts to easily and completely be suspended or deleted, which facilitates greater security across the enterprise.	DA, EA, ISU
A single forest leads to greater consistency of security settings, physical housing of domain controllers, and secured access to critical enterprise resources. This benefits all users by ensuring the availability and security of enterprise resources.	EU, DA, ISU
Every forest requires maintenance of Active Directory sites, IP subnets, site based group policies, creation of enterprise and schema admin policies, and creation/deletion of user credentials. All of these tasks are duplicated and must be separately maintained in a multi-forest environment, which saps administrative resources and drives up cost.	DA, ISU
Transfer of users between departments in a single forest is a trivial task and allows end user to retain the same network account. This is a benefit to departmental administrators, as they do not have to manually create and delete accounts when users transfer between departments/majors. This also benefits end-users, as they do not have to relearn accounts upon transfer and may retain home directories and profiles.	DA, EU
The forest abstracts the underlying network architecture from the end user, which allows them to access a consistent set of resources from anywhere on campus without regard for domain structure. This benefits the end user by making it easier to access resources and help desk staff by reducing support calls.	EU, HD

A single forest permits the use of smart cards for authentication, thus improving network security, which is a benefit to the entire university.	EU, ISU
Roaming profiles may be used when accessing resources from any workstation that resides in the same forest as the user's credentials. This provides the end user with a consistent environment from anywhere on campus.	EU, HD

Single Forest Architecture with Single Identity - Drawbacks	Who
Departments loose the ability to modify forest schema independently of one another. Because of this, more rigorous testing of schema modifications may need to occur, resulting in delay. This may be a drawback to those waiting on a schema change that will enable some new functionality.	EU
In a single forest with single identities, only one account will exist for each student. Because no one department owns a student account, some customization of student accounts may be lost or may require greater agreement between departments.	EU, DA
A minimal set of forest wide activities may only be performed by enterprise/schema administrators. This requires autonomous IT departments to agree on membership and usage of enterprise and schema administrator groups, resulting in some loss of independence.	DA

Single Domain with Unique Requirement Domains - Benefits	Who
The single domain architecture eliminates complexity and reduces enterprise troubleshooting effort by eliminating the need for universal groups, global catalog servers and inter-domain replication configuration.	EA
The single domain architecture eliminates the need for departmental IT units to maintain and support authentication services.	DA
The single domain architecture allows for a well maintained, highly secure, and widely distributed authentication source, thus providing greater availability of resources.	EU
The single domain architecture requires reduced DNS complexity and allows DNS services to be maintained by and experienced DNS management staff.	EU, DA, EA, EU
The single domain architecture requires that a consistent set of security policies be applied to all user accounts. This enforces greater resource security as a consistent set of policies must be followed by all faculty, students, and staff.	ISU
The single domain architecture is easier and less complex to support for account synchronization with other directories.	EA

Single Domain with Unique Requirement Domains - Drawbacks	Who
All IT departments must agree on a common set of password, account lockout, and Kerberos policies. This reduces some of the autonomy that departmental IT possesses today.	DA
Initially, a higher level of delegation effort is required in a single domain environment, thus resulting in greater enterprise administrator effort and potentially resulting in some delay for departmental admins.	DA, EA
The single domain environment requires that departmental IT agree on usage/modification to domain-wide resources, such as installation of application on domain controllers and usage of terminal server licensing servers.	DA

6. Appendix B – Capacity Planning

The following table outlines 2-year projected Active Directory objects, activities, and hardware recommendations for the IASTATE.EDU campus Active Directory forest.

	IASTATE.EDU
User Accounts	
Total number of users	50000
Percent of users concurrently active during peak hour	20
Additional user attributes	15
Password expiration frequency in days	90
Average Interactive logon rate	3
Average Batch logon rate	1
Average Network logon rate	5
Computers and other objects	
Windows 2000 computers	3000
Non windows 2000 computers	3000
Other Active Directory objects	5700
Desired average Domain Controller CPU utilization	20

Administration	
Administrative load interval	Daily
Administrative additions	100
Administrative deletions	25
Administrative modifications	175
Exchange 2000 (These numbers are based on 50,000 users)	
Using Exchange 2000	Yes
Average messages sent by each user per day	2
Average number of recipients per message	2
Exchange 2000 Servers	
Percent Outlook mail clients	
Percent POP mail clients	
Percent SMTP mail clients	
Percent other mail clients	
Services Using Active Directory	
Total number of dial-in connections per day	
DHCP lease expiration period (days)	8
DNS scavenging NoRefreshInterval	7
Custom application searches per second	1
Custom application adds per second	
Custom application deletes per second	
Custom application modifications per second	

The following table outlines minimum hardware recommendations for the IASTATE.EDU domain. These recommendations are based on input into the Active Directory Sizer tool from the previous table.

BH – Bridge Head

DC – Domain Controller

GC – Global Catalog Server

Domain	Role	Site	Servers	DB Size	Processor	CPUs	Memory
IASTATE	BH / DC / GC	Main	1	576 Mb	Pentium III Xeon 550 MHz	1	256 Mbytes
IASTATE	BH / DC / GC	CBS	1	576 Mb	Pentium III Xeon 550 MHz	1	256 Mbytes
IASTATE	DC / GC	Main	3	576 Mb	Pentium III Xeon 550 MHz	1	256 Mbytes

For optimal performance, the following disk configuration is recommended on all IASTATE.EDU domain controllers:

Role	Disks	Configuration
System	2	RAID 1
Log Files	2	RAID 1
Database	3	RAID 5