

NIS+ Transition Guide

2550 Garcia Avenue
Mountain View, CA 94043
U.S.A.



SunSoft
A Sun Microsystems, Inc. Business

© 1994 Sun Microsystems, Inc.
2550 Garcia Avenue, Mountain View, California 94043-1100 U.S.A.

All rights reserved. This product and related documentation are protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or related documentation may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any.

Portions of this product may be derived from the UNIX[®] and Berkeley 4.3 BSD systems, licensed from UNIX System Laboratories, Inc., a wholly owned subsidiary of Novell, Inc., and the University of California, respectively. Third-party font software in this product is protected by copyright and licensed from Sun's font suppliers.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the United States Government is subject to the restrictions set forth in DFARS 252.227-7013 (c)(1)(ii) and FAR 52.227-19.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

TRADEMARKS

Sun, Sun Microsystems, the Sun logo, Sun Microsystems Computer Corporation, the Sun Microsystems Computer Corporation logo, SunSoft, the SunSoft logo, Solaris, SunOS, OpenWindows, DeskSet, ONC, ONC+, and NFS are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and certain other countries. UNIX and OPEN LOOK are registered trademarks of UNIX System Laboratories, Inc., a wholly owned subsidiary of Novell, Inc. All other product names mentioned herein are the trademarks of their respective owners.

All SPARC trademarks, including the SCD Compliant Logo, are trademarks or registered trademarks of SPARC International, Inc. SPARCstation, SPARCserver, SPARCengine, SPARCstorage, SPARCware, SPARCcenter, SPARCclassic, SPARCcluster, SPARCdesign, SPARC811, SPARCprinter, UltraSPARC, microSPARC, SPARCworks, and SPARCcompiler are licensed exclusively to Sun Microsystems, Inc. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK[®] and Sun[™] Graphical User Interfaces were developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

X Window System is a product of the Massachusetts Institute of Technology.

THIS PUBLICATION IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT.

THIS PUBLICATION COULD INCLUDE TECHNICAL INACCURACIES OR TYPOGRAPHICAL ERRORS. CHANGES ARE PERIODICALLY ADDED TO THE INFORMATION HEREIN; THESE CHANGES WILL BE INCORPORATED IN NEW EDITIONS OF THE PUBLICATION. SUN MICROSYSTEMS, INC. MAY MAKE IMPROVEMENTS AND/OR CHANGES IN THE PRODUCT(S) AND/OR THE PROGRAM(S) DESCRIBED IN THIS PUBLICATION AT ANY TIME.



Contents

Preface.....	xi
1. Introduction.....	1
Differences Between NIS and NIS+.....	1
Domain Structure.....	2
Interoperability.....	3
Server Configuration.....	4
Information Management.....	5
Security.....	6
Suggested Transition Process.....	6
Identify Major Transition Goals.....	7
Become Familiar With NIS+.....	8
Design Your Final NIS+ Namespace.....	9
Select Security Measures.....	9
Decide How to Use NIS-Compatibility Mode.....	10
Complete Prerequisites to Transition.....	10

Implement the Transition	10
2. Designing the NIS+ Namespace	11
Identify the Goals of Your Administrative Model.	11
Design the Namespace Structure	12
The Domain Hierarchy	12
Designing a Domain Hierarchy	13
Domain Names	18
The Email Environment	19
Select the Namespace Servers	19
Supported Domains.	20
Server Load.	22
Disk Space and Memory Requirements	23
Determine Table Configurations	24
Differences Between NIS+ Tables and NIS Maps	24
Use of Custom NIS+ Tables	28
Connections Between Tables	29
3. Selecting NIS+ Security Measures.	33
Determine the Impact of Implementing NIS+ Security	33
How NIS+ Security Affects Users	34
How NIS+ Security Affects Administrators.	34
How NIS+ Security Affects Transition Planning.	35
Select Credentials.	35
Selecting a Security Level	37
Forming NIS+ Groups.	37

Determine Access Rights to NIS+ Groups and Directories . . .	38
Determine Access Rights to NIS+ Tables	39
4. Using NIS-Compatibility Mode	45
Select the Domains That Will Be NIS Compatible	48
Determine the Configuration of NIS-Compatible Servers	48
Decide How to Transfer Information Between Services	49
A Note About Changing Information in the Password Table	50
Decide How to Implement DNS Forwarding.	51
NIS and NIS+ Command Equivalents in the Solaris 1.x and 2.x Releases.	53
NIS Commands Supported in the Solaris 2.x Release	53
Client and Server Command Equivalents	54
NIS and NIS+ API Function Equivalents	56
5. Prerequisites to Transition	59
Gauge the Impact of NIS+ on Other Systems.	59
Train Administrators.	60
Write a Communications Plan	60
Identify Required Conversion Tools and Processes.	61
Identify Administrative Groups Used for Transition	61
Determine Who Will Own the Domains	62
Determine Resource Availability	63
Resolve Conflicts Between Login Names and Hostnames	63
Examine All Information Source Files.	64
Remove the “.” From NIS Map Names	64

Document Your Existing NIS Namespace	65
Create a Conversion Plan for Your NIS Servers.....	65
6. Implementing the Transition	67
Phase I—Set Up the NIS+ Namespace	68
Phase II—Connect the NIS+ Namespace to Other Namespaces	69
Phase III—Make the NIS+ Namespace Fully Operational	70
Phase IV—Upgrade NIS-Compatible Domains.....	71
Index.....	73

Figures

Figure 1-1	NIS+ Domains	2
Figure 1-2	NIS+ Standard Tables	5
Figure 2-1	Sample NIS+ Hierarchy by Logical Organization	14
Figure 2-2	Sample NIS+ Hierarchy by Physical Location	14
Figure 2-3	Domain Names Syntax	18
Figure 2-4	Server Relationship to Domain.....	20
Figure 2-5	Assigning Servers to Domains	21
Figure 2-6	Adding Replicas to a Domain	22
Figure 2-7	Sample Name Service Switch File	28
Figure 2-8	Establishing a Path to Tables in a Higher Domain.....	29
Figure 2-9	Information Distribution Across an NIS+ Hierarchy	31
Figure 3-1	NIS+ Principals	36
Figure 3-2	Default Passwd Table Access Rights for NIS-compatible Domains	41
Figure 3-3	Default Passwd Table Access Rights for Standard NIS+ Domains	42
Figure 4-1	Transition to NIS-Compatibility Mode	47

Tables

Table P-1	Typographic Conventions	xiii
Table 2-1	NIS+ Tables	25
Table 2-2	Correspondence Between NIS Maps and NIS+ Tables	26
Table 3-1	Default Access Rights for NIS+ Objects	38
Table 3-2	Default Access Rights for NIS+ Tables and Columns	40
Table 4-1	Support for NIS Protocols by NIS+ Servers in NIS Compatibility- Mode	45
Table 4-2	NIS+ Data Transfer Commands	50
Table 4-3	Commands for Changing Passwords	51
Table 4-4	NIS Command Support in the Solaris 2.x Releases	54
Table 4-5	NIS Client Commands and Equivalent NIS+ Commands . . .	55
Table 4-6	NIS Server Commands and Equivalent NIS+ Commands . . .	56
Table 4-7	NIS API and NIS+ API Equivalent Functions	57
Table 5-1	NIS+ Commands for Groups	62

Preface

NIS+ Transition Guide describes how to convert a site running the network information service (NIS) name service to one running the network information service plus (NIS+) name service. This manual is part of the Solaris™ 2.4 System and Network Administration manual set.

Who Should Use This Book

This manual is for experienced system and network administrators who want to convert their site from NIS to NIS+. For information on configuring NIS+, see *Name Services Configuration Guide*. For NIS+ customizing information and detailed administration instructions, see *Name Services Administration Guide*.

Although this manual introduces some concepts relevant to NIS+, it makes no attempt to explain networking fundamentals or describe the administration tools offered by the Solaris environment. If you administer networks, this manual assumes you already know how the administration tools work and have already chosen your favorite tools.

How This Book Is Organized

This book contains six chapters:

Chapter 1, “Introduction,” describes the differences between NIS and NIS+ features and an overview of the suggested transition process.

Chapter 2, “Designing the NIS+ Namespace,” discusses how to design your NIS+ namespace.

Chapter 3, “Selecting NIS+ Security Measures,” describes the NIS+ security features and the effects they have on administration and transition planning.

Chapter 4, “Using NIS-Compatibility Mode,” describes how to run NIS and NIS+ clients concurrently and NIS+ servers in NIS-compatibility mode.

Chapter 5, “Prerequisites to Transition,” presents the steps that you need to take before beginning the actual transition.

Chapter 6, “Implementing the Transition,” lists the steps required to implement an NIS to NIS+ transition.

Related Books

You can consult the following for more information on NIS+ and DNS:

- *Name Services Configuration Guide*—Describes how to plan for and configure an NIS+ namespace
- *Name Services Administration Guide*—Describes how to administer a running NIS+ namespace and modify its security level
- *Network Interfaces Programmer’s Guide*—Describes the application programming interfaces for networks including NIS+
- *man Pages(1M): System Administration Commands*—Contains the man pages for the NIS+ commands
- *Administration Application Reference Manual*—Describes the administration tool window interface for modifying the data in NIS+ tables

These books are also part of the Solaris 2.4 System and Network Administration manual set.

What Typographic Changes and Symbols Mean

Table P-1 describes the type changes and symbols used in this book.

Table P-1 Typographic Conventions

Typeface or Symbol	Meaning	Example
AaBbCc123	The names of commands, files, and directories; on-screen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. system% You have mail.
AaBbCc123	What you type, contrasted with on-screen computer output	<div style="border: 1px solid black; padding: 2px;">system% su Password:</div>
<i>AaBbCc123</i>	Command-line placeholder: replace with a real name or value	To delete a file, type <code>rm filename</code> .
<i>AaBbCc123</i>	Book titles, new words or terms, or words to be emphasized	Read Chapter 6 in <i>User's Guide</i> . These are called <i>class</i> options. You <i>must</i> be root to do this.

Code samples are included in boxes and may display the following:

%	UNIX C shell prompt	system%
\$	UNIX Bourne and Korn shell prompt	system\$
#	Superuser prompt, all shells	system#

In the AnswerBook[®] on-line documentation tool, double-clicking on a cross-reference takes you to the page on which it begins.

Introduction

This chapter introduces the issues involved in converting from NIS to NIS+. It describes the differences between the two name services and outlines a suggested transition process.

<i>Differences Between NIS and NIS+</i>	<i>page 1</i>
<i>Suggested Transition Process</i>	<i>page 6</i>

Differences Between NIS and NIS+

NIS and NIS+ have several differences that have an impact on a transition. For instance, NIS uses a flat, non-hierarchical namespace with only one domain (or several disconnected domains), while NIS+ provides a domain hierarchy similar to that of DNS. This means that before you can convert to NIS+, you must design the NIS+ namespace. NIS+ also provides security, which limits access not only to the information in the namespace but also to the structural components of the namespace.

These and other differences demonstrate that NIS+ is not simply an upgrade to NIS, but an entirely new product. Therefore, the transition from NIS to NIS+ is largely directed by the differences between the products.

These differences are described in broad terms in the remainder of this chapter. Understanding them is critical to a successful transition to NIS+. They are:

- Domain structure
- Interoperability

- Server configuration
- Information management
- Security


Domain Structure

NIS+ is not simply an upgrade to NIS; it is designed to *replace* NIS. This becomes evident when you examine its domain structure. NIS domains are flat and lack the ability to have a hierarchy. NIS+ domains may be flat, but you can also construct NIS+ hierarchical domains. Such hierarchies consist of a root domain with an infinite number of subdomains under them.

The NIS domain structure addressed the administration requirements of client-server computing networks prevalent in the 1980s; in other words, client-server networks with a few hundred clients and a few multipurpose servers.

NIS+ is designed to support networks with 100 to 10,000 multivendor clients supported by 10 to 100 specialized servers located in sites throughout the world, connected to several “untrusted” public networks. The size and complexity of these networks requires new, autonomous administration practices. The NIS+ domain structure was designed to address these requirements. It consists of hierarchical domains similar to those of DNS, as shown in the following diagram:

Figure 1-1 NIS+ Domains



Hierarchical domains allow NIS+ to be used in a range of networks, from small to very large. They also allow the NIS+ service to adapt to the growth of an organization. The NIS+ domain structure is thoroughly described in the *Name Services Configuration Guide*.

Interoperability

NIS+ provides interoperability features designed for upgrading from NIS and for continuing the interaction with DNS originally provided by the NIS service.

To help convert from NIS, NIS+ provides an NIS-compatibility mode and an information transfer utility. The NIS-compatibility mode enables an NIS+ server running Solaris 2.x software to answer requests from NIS clients while continuing to answer requests from NIS+ clients. The information transfer utility helps administrators keep NIS maps and NIS+ tables synchronized.

NIS-compatibility mode requires slightly different setup procedures than those used for a standard NIS+ server. Also, NIS-compatibility mode has security implications for tables in the NIS+ namespace. These differences and implications are described in *Name Services Configuration Guide*.

NIS client machines interact with the NIS+ namespace differently from NIS+ client machines when NIS+ servers are running in NIS-compatibility mode. The differences are:

- NIS client machines cannot follow NIS+ table paths or links, or do read operations in other domains.
- NIS client machines can have their unsatisfied host requests forwarded to DNS if you run `rpc.nisd` with the `-Y -B` options, but the NIS+ server will not forward these requests for an NIS+ client. The forwarding of NIS+ client machines' DNS requests is controlled by the `resolv.conf` and `nsswitch.conf` files' configurations. See *Name Services Configuration Guide* for more information.
- Authorized NIS+ client users can update entries in the namespace and perform operations such as updating the password with the `nispasswd` command, and changing their secure-RPC credentials with the `chkey` command. NIS client users cannot make any updates directly to NIS+ servers running in NIS-compatibility mode. For example, you cannot use `yppasswd` to update a password in the NIS+ password table. You have to use `nispasswd` to update a password in the NIS+ password table. Since you cannot run `nispasswd` on an NIS client, the command must be run on the NIS+ server.

- Even if all the servers on a local subnet no longer respond, the NIS+ client machines can still have their name service calls answered if they can contact any of the replicas of that domain. NIS client machines do not have access to information on the network outside their subnet unless the server names have been set with `ypset`, or for Solaris 2.x NIS clients only, with `ypinit`.
- NIS client machines cannot be sure that the data they are receiving comes from an authorized NIS server, while authorized NIS+ clients are certain that the data is coming from an authorized NIS+ server.
- Under NIS, if the server is no longer responding, the NIS `yp_match()` call continues to retry this call until the server starts responding and answers the request. The NIS+ API (application program interface) will return an error message to the application when this situation occurs.

In the Solaris 2.3 and later releases, the NIS-compatibility mode *supports* DNS forwarding. In the Solaris 2.2 release, support for DNS forwarding is available as a *patch*. The DNS forwarding patch is *not* available in the Solaris 2.0 and 2.1 releases.

Although an NIS+ domain cannot be connected to the Internet directly, the NIS+ client machines can be connected to the Internet with the name service switch. The client can set up its switch configuration file to search for information in either DNS zone files or NIS maps—in addition to NIS+ tables.

Server Configuration

The NIS+ client-server arrangement is similar to those of NIS and DNS in that each domain is supported by a set of servers. The main server is called the *master* server, and the backup servers are called *replicas*. Both master and replica servers run NIS+ server software and both maintain copies of NIS+ tables.

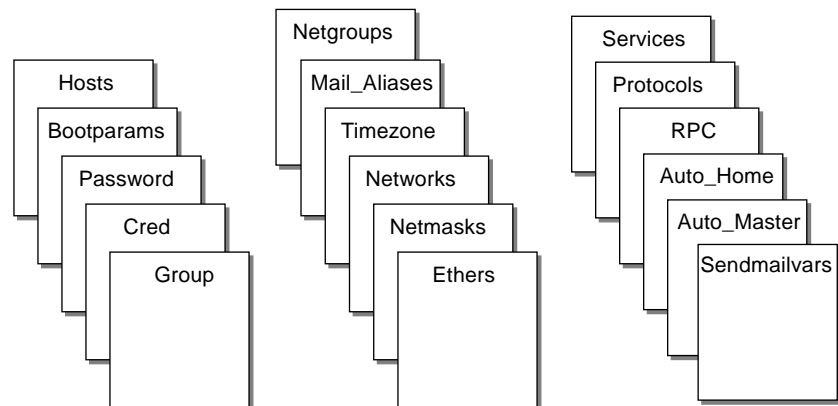
However, NIS+ uses an update model that is completely different from the one used by NIS. At the time NIS was developed, it was assumed that most of the information NIS would store would be static. NIS updates are handled manually and its maps have to be remade and propagated in full every time any information in the map changes.

NIS+, however, accepts *incremental* updates to the replicas. Changes must still be made to the master database on the master server, but once made they are automatically propagated to the replica servers. You don't have to "make" any maps or wait hours for propagation. Propagation is now a matter of minutes, perhaps as much as 200 seconds.

Information Management

NIS+ stores information in *tables* instead of maps or zone files. NIS+ provides 17 types of predefined, or *system*, table as is shown in Figure 1-2:

Figure 1-2 NIS+ Standard Tables



NIS+ tables are not ASCII files, but are tables in the NIS+ relational database. You can only view and edit their contents by using the NIS+ commands.

NIS+ tables provide two major improvements over the maps used by NIS. First, an NIS+ table can be searched by any searchable column, not just the first column (sometimes referred to as the "key"). To know whether a particular column is searchable, run the `niscat -o` command on a table. The command returns a list of the table's columns and their attributes, one of the which is whether a column is searchable. This searching ability eliminates the need for duplicate maps, such as the `hosts.byname` and `hosts.byaddr` maps used by NIS. Second, the information in NIS+ tables has access controls at three levels: the table level, the entry (row) level, and the column level.

NIS maps are located on the server in `/var/yp/domainname`, whereas NIS+ directories are located in `/var/nis/machinename`. The NIS+ tables are contained in the database. The tables' information is loaded into memory as requests are made to the database. Keeping data in memory in the order requested minimizes calls to the disk, thereby improving request response time.

Security

The security features of NIS+ protect the information in the namespace and the structure of the namespace itself from unauthorized access. NIS+ security is provided by two means: *authentication* and *authorization*. Authentication is the process by which an NIS+ server identifies the NIS+ *principal* (a client user or client workstation) that sent a particular request. Authorization is the process by which a server identifies the access rights granted to that principal, be it a client machine or client user.

In other words, before someone can access anything in the namespace, they must be an authenticated NIS+ client and they must have the proper permission to access that information if the highest level of NIS+ security has been implemented at a site. Furthermore, requests for access to the namespace are only honored if they are made either through NIS+ client library routines or NIS+ administration commands. The NIS+ tables and structures cannot be directly edited.

Suggested Transition Process

Following is a suggested transition process:

1. Identify major transition goals.
2. Become familiar with NIS+.
3. Design your final NIS+ namespace.
4. Select security measures.
5. Decide how to use NIS-compatibility mode.
6. Complete prerequisites to transition.
7. Implement the transition.

Identify Major Transition Goals

Before you begin the transition, you might find it helpful to state the goals of your site's transition before you begin. Here are some typical goals:

Consider the Alternatives to Making the Transition Immediately

You can defer the upgrade to NIS+ until after your site has completed its transition to the Solaris 2.x release. This would allow you to focus your resources on one transition effort at a time. Use the binary-compatibility mode Name Service kit to continue running NIS after you upgrade to the Solaris 2.x release. It is a much better solution than running NIS+ with the Solaris 1.x release.

Upgrading to NIS+ before upgrading to the Solaris 2.x release is possible but definitely *not* recommended. There are Solaris 1.x NIS+ binaries. These binaries were developed to enable sites to upgrade their network information service when they were unable to upgrade their servers to the Solaris 2.x release. There is no client-side NIS+ support for the Solaris 1.x release.

Keep Things Simple

You can take several steps to simplify the transition. While these steps will diminish the effectiveness of NIS+, they will consume fewer servers and less administrative time. Once the transition is complete, you can change the NIS+ setup to completely suit your needs. Here are some suggestions:

- Don't change domain names.
- Don't use any hierarchies, keep a flat NIS+ namespace.
- Use the NIS-compatibility features.
- Use default tables and directory structures.
- Don't establish credentials for clients.

Use a Single Release of Software

Decide which version of the Solaris 2.x software and NIS+ you will use for the transition. Since there are slight differences between versions, using multiple versions would complicate the transition process needlessly. Choose one version of the Solaris product and use its corresponding version of NIS+.

The current release has the most features (such as setup scripts). Make sure you compile a list of the SunOS patches that are required for normal operation, and make sure that all servers and clients have the same patches loaded.

Minimize Impact on Client Users

This goal implies two major considerations. First, users should not notice any change in service. Second, the transition phase itself should cause minimal disruption to client users. To ensure the second consideration, be sure the administrators responsible for each domain migrate their client machines to NIS+, rather than ask the users to implement the migration. This ensures the proper procedures are implemented, the procedures are consistent across client machines, and irregularities can be dealt with immediately by the administrator.

Miscellaneous “Don’ts”

- Don’t change the name services currently provided by NIS or the way NIS functions.
- Don’t change the structure of DNS.
- Don’t change the IP network topology.
- Don’t upgrade applications that use NIS to NIS+; leave the migration to NIS+ APIs for the future.
- Don’t consider additional uses for NIS+ during the implementation phase; add them later.

Become Familiar With NIS+

Familiarize yourself with NIS+, particularly with the concepts summarized earlier in this chapter and discussed in the remainder of this book. For details, see the following sources of information:

- *Name Services Configuration Guide* – Describes the structure of NIS+ objects, provides planning guidelines, complete setup instructions, task-oriented descriptions of NIS+ administrative commands, and how to use the NIS+ scripts, which help you set up an NIS+ namespace more easily and quickly than using the individual NIS+ administrative commands.
- *Name Services Administration Guide* – Describes how to administer a running NIS+ namespace and how to customize your namespace.
- *Network Interfaces Programmer’s Guide* – Describes the functions of the NIS+ API.
- NIS+ man pages in *man Pages(1M): System Administration Commands* – Describes all the NIS+ commands and their syntax.

One of the best ways to become familiar with NIS+ is to build a prototype namespace. There is no substitute for hands-on experience with the product; administrators need the opportunity to practice in a forgiving test environment.

Note – Do not use your prototype domain as the basis for your actual running NIS+ namespace. Deleting your prototype when you have learned all you can from it will avoid namespace configuration problems. Start anew to create the real namespace after following all the planning steps.

When you create the test domain(s), make small, manageable domains. For guidance, you can use *Name Services Configuration Guide* and *Name Services Administration Guide*. *Name Services Configuration Guide* shows you how to create a simple test domain and subdomain with or without NIS-compatibility mode set using the NIS+ scripts. The NIS+ scripts, which are only available starting with the Solaris 2.3 release, spare you the trouble of typing each individual NIS+ command when you are performing such tasks as setting up an NIS+ server. It also includes all the planning information you need to know before you make the transition to your real NIS+ namespace.

Design Your Final NIS+ Namespace

Design the final NIS+ namespace, following the guidelines in Chapter 2, “Designing the NIS+ Namespace.” While designing the namespace, do not worry about limitations imposed by the transition from NIS. You can add those later, once you know what your final NIS+ goal is.

Select Security Measures

NIS+ security measures provide a great benefit to users and administrators, but they require additional knowledge and setup steps on the part of both users and administrators. They also require several planning decisions. Chapter 3, “Selecting NIS+ Security Measures” describes the implications of NIS+ security and the decisions you need to make for using it in your NIS+ namespace.

Decide How to Use NIS-Compatibility Mode

The use of parallel NIS and NIS+ namespaces is virtually unavoidable during a transition. Because of the additional resources required for parallel namespaces, try to develop a transition sequence that reduces the amount of time your site uses dual services or the extent of dual services within the namespace (for example, convert as many domains as possible to NIS+ only).

Chapter 4, “Using NIS-Compatibility Mode” explains the transition issues associated with the NIS-compatibility mode and suggests a way to make the transition from NIS, through NIS compatibility, to NIS+ alone.


Complete Prerequisites to Transition

In addition to all these planning decisions mentioned above, you must complete several miscellaneous prerequisites, as described in Chapter 5, “Prerequisites to Transition.”

Implement the Transition

Chapter 6, “Implementing the Transition” provides a suggested series of steps to implement the transition you have planned in the previous steps.

Designing the NIS+ Namespace



This chapter provides general guidelines and recommendations for designing the final NIS+ namespace your site will have.

<i>Identify the Goals of Your Administrative Model</i>	<i>page 11</i>
<i>Design the Namespace Structure</i>	<i>page 12</i>
<i>Select the Namespace Servers</i>	<i>page 19</i>
<i>Determine Table Configurations</i>	<i>page 24</i>

When designing the namespace, do not worry about limitations imposed by the transition from NIS. You can modify your NIS+ domain later, once you know what your final NIS+ configuration will look like.

Identify the Goals of Your Administrative Model

Select the model of information administration, such as the domain structure, that your site will use. Without a clear idea of how information at your site will be created, stored, used, and administered, it is difficult to make the design decisions suggested in this section. You could end up with a design that is more expensive to operate than necessary. You also run the risk of designing a namespace that does not suit your needs. Changing the namespace design after it has been set up is costly.

Design the Namespace Structure

Designing the NIS+ namespace is one of the most important tasks you can perform, since changing the domain structure after NIS+ has been set up is a time-consuming, complex job. It is complex because information, security, and administration policies are woven into the domain structure of the namespace. Rearranging domains would require rearranging information, reestablishing security, and recreating administration policies.

When designing the structure of an NIS+ namespace, consider the following factors:

- The domain hierarchy
- Domain names
- The mail environment

The Domain Hierarchy

The main benefit of an NIS+ domain hierarchy is that it allows the namespace to be divided into more easily managed components. Each component can have its own security, information management, and administration policies. It is advisable to have a hierarchy if the number of clients you have exceeds 500, if you want to set up different security policies for a set of users, or if you have geographically distributed sites.

Unless there is a need for a domain hierarchy, not having a hierarchy will simplify your transition to NIS+ because the NIS+ servers for each subdomain are not part of the subdomain that they serve, with the exception of the root domain. The NIS+ servers are within the parent domain of the subdomain they serve. This relationship of server to subdomain creates problems for those applications that expect the servers to be able to get their name service data from the subdomain. For example, if a subdomain NIS+ server is also an NFS server, then the server would not get its netgroups information from the subdomain, but instead retrieve the information from its domain, which is the domain above the subdomain; this can be confusing. Another example of when a hierarchy could cause problems would be where the NIS+ server was also used by users to log in remotely and to execute certain commands that they could not execute from their own workstations. If you have only a single root domain, you will not have these problems because NIS+ root servers live in the domain that they serve.

When all users were in the same NIS domain, they were directly visible to each other without using fully qualified names. Creating an NIS+ hierarchy, however, puts users in separate domains, which means that the users in one domain will not be directly visible to users in another domain unless you use fully qualified names or paths. For example, if there are two subdomains, `sales.wiz.com.` and `eng.wiz.com.`, created out of the earlier `wiz.com.` domain, then for user “joe” in the `sales.wiz.com.` domain to be able to send mail to user “jane” in `eng.wiz.com.`, he would have to specify her name as `jane@hostname.eng.wiz.com.` (or `jane@hostname.eng`) instead of just “jane” as was sufficient when they were in the same domain. Remote logins also require fully qualified names between domains.

You could use the table path to set up connections between tables in one domain and another domain, but to do so would negate the advantages of having a domain hierarchy. You would also be reducing the reliability of the NIS+ service because now clients would have to depend upon the availability of not only their own home domains, but also of other domains to which their tables are pathed. Using table paths may also slow request response time.

Designing a Domain Hierarchy

To design a domain hierarchy, first read Part I of *Name Services Configuration Guide*. It describes NIS+ domain structure, information storage, and security.


Once you are familiar with the components of a domain hierarchy, make a diagram of how you expect the hierarchy to look when you are finished. The diagram will be a useful reference when you are in the midst of the setup procedure. At a minimum, you will need to consider the following issues:

- Organizational or geographical mapping
- Connection to higher domain
- Client support in the root domain
- Domain size compared to number of domains
- Number of levels
- Security levels
- Replicas and number of replicas
- Information management

Organizational or Geographical Mapping


One of the major benefits of NIS+ is its capability to divide the namespace into smaller, manageable parts. You could create a hierarchy of organizations, such as those of the hypothetical corporation, Wizard, Inc.:

Figure 2-1 Sample NIS+ Hierarchy by Logical Organization



You could also organize the hierarchy by buildings instead of organizations:

Figure 2-2 Sample NIS+ Hierarchy by Physical Location



The scheme you select depends primarily on how you prefer to administer the namespace and how clients will tend to use the namespace. For example, if clients of `eng.wiz.com.` will be distributed throughout the buildings of Wizard, Inc., you should not organize the namespace by building. Since the clients would need to constantly have access to other domains, you would need to add their credentials to the other domains and you would increase traffic flow through the root master server. A better scheme would be to arrange them by organization. On the other hand, building-sized domains are immune to the reorganizations that require organization-based domains to be restructured.

Do not be limited by the physical layout of the network, since an NIS+ namespace does not have to be congruent with the physical network, except for situations where it has to support NIS clients. The number of domains your namespace needs depends on the kind of hierarchy you select.

Consider future expansion plans. Will today's NIS+ root domain be beneath another NIS+ domain in the future? Changing this arrangement would entail a great deal of work. Try to estimate the need for future domains in the namespace and design a structure that can accommodate them without disruption.

Connection to Higher Domains?

Consider whether the NIS+ namespace will be connected to higher domains, such as those of the Internet or DNS. If you currently use NIS under a DNS hierarchy, do you want to replace the NIS domains only or do you want to replace the entire companywide DNS/NIS structure with an NIS+ namespace?

Client Support in the Root Domain

In the two Wizard, Inc., domain hierarchies illustrated in Figure 2-1 and Figure 2-2 on page 14, are all the clients placed in domains beneath the root domain? Or do some belong to the root domain? Is the purpose of the root domain only to act as the root for its subdomains or will it support its own group of clients? You could place all clients in the lowest layer of domains, and only those used for administration in the intermediate domains. For example, if you implemented this plan in the first illustration, all clients would belong to the `big.sales.wiz.com.`, `small.sales.wiz.com.`, and `eng.wiz.com.` domains, and only clients used for administration would belong to the `wiz.com.` and `sales.wiz.com.` domains.

Or you could place the clients of general-purpose departments in higher-level domains. For example, in the second illustration, Figure 2-2 on page 14, where the domain is organized by building, you could put the clients of the Facilities Department in the `wiz.com.` domain. It is not recommended that you do so, however, because the root domain should be kept simple and relatively unpopulated.

Domain Size Compared With Number of Domains

The current NIS+ implementation is optimized for up to 1000 NIS+ clients per domain and for up to 10 replicas per domain. Such a domain would typically have 10,000 table entries. The limitations come from the current server discovery protocol. If you have more than 1000 NIS+ clients, you should divide your namespace into different domains and create a hierarchy.

Creating a hierarchy, however, may introduce more complexity than you are prepared to handle. You may still prefer to create larger domains rather than a hierarchy because one large domain requires less administration than multiple smaller domains do. Larger domains need fewer skilled administrators to service them, since tasks can be automated more readily (with scripts you create), thus lowering the administrative expense. Smaller domains provide better performance, and you can customize their tables more easily. You also achieve greater administrative flexibility with smaller domains.

Number of Levels

NIS+ was designed to handle multiple levels of domains. Although the software can accommodate almost any number of levels, a hierarchy with too many levels would be difficult to administer. For example, the names of objects could become long and unwieldy. Consider 20 to be the limit for the number of subdomains for any one domain and limit the levels of the NIS+ hierarchy to 5.

Security Level

Typically, you will run the namespace at security level 2. However, if you plan to use different security levels for different domains, you should identify them now. Chapter 3, “Selecting NIS+ Security Measures,” provides more information about security levels.

Replicas and Number of Replicas

Any one domain should have no more than 10 replicas because of the increased network traffic and server load that occur when information updates are propagated to the replicas. Determining the number of replicas a domain requires depends on other factors as well. They are:

- The physical location of the servers
- The number of subnets in a domain
- Whether there are NIS clients in the NIS+ namespace

You should create a minimum of two servers (one master and one replica) for every domain and at least one replica for every physical location. You do not need a replica for every subnet, unless you have Solaris 1.x NIS clients and you want NIS+ servers in NIS-compatibility mode to support the NIS clients. NIS clients do not have access to servers that are not on the same subnet. The only exception are the Solaris 2.x NIS clients, which can use `ypinit (1M)` to specify a list of NIS servers. The netmask number in these cases would have to be set appropriately.

One way you can have a sufficient number of replicas per domain without using a multiplicity of machines is to create multi-homed servers. A multi-homed server is a machine with multiple ethernet or network interfaces. A multi-homed server can serve multiple subnets in a domain.

If the domain hierarchy that you design spans a WAN link, it would be prudent to replicate the domain on either side of the WAN link—with a master server on one side and a replica on the other. This could possibly enable clients on the other side of the link to continue with NIS+ service even if the WAN link were temporarily disabled. Putting servers on either side of a WAN, however, does change the structure of a namespace that is organized by group function rather than by physical layout, since the replica might possibly physically reside within the geographic perimeter of a different domain.

Domains Across Time Zones

Geographically dispersed organizations may determine that organizing their domain hierarchy by functional groups would cause a domain to span more than one time zone. It is *strongly* recommended that you do *not* have domains that span multiple time zones. If you do need to configure a domain across time zones, be aware that a replica's time will be taken from the master server, so the database updates will be synchronized properly, using universal (Greenwich Mean) time. This may cause problems if the replica machine is used for other services that are time critical. To make domains across time zones work, the replica's `/etc/TIMEZONE` file has to be locally set to the master server's time zone when you are installing NIS+. Once the replica is running, some time critical programs may run properly and some may not, depending on whether these programs use universal or local time.

Information Management


It is best to use a model of local administration within centralized constraints for managing the information in an NIS+ namespace. Information should be managed, as much as possible, from within its home domain, but according to guidelines or policies set at the global namespace level. This provides the greatest degree of domain independence while maintaining consistency across domains.

Domain Names

Consider name length and complexity. First, choose names that are descriptive. “Sales” is considerably more descriptive than “BW23A.” Second, choose short names. Avoid appending a long string such as “EmployeeAdministrationServices.WizardCorporation” to object names when you administer the namespace.

A domain name is formed from left to right, starting with the local domain and ending with the root domain. For example:

Figure 2-3 Domain Names Syntax



The first line above shows the name of the root domain. The root domain must always have at least two labels and must end in a dot. The second label can be an Internet domain name, such as “Com.” The second line above shows the name of a lower-level domain.

Also consider implications of particular names for email domains, both within the company and over the Internet.

Depending on the migration strategy chosen, it could be a viable alternative to change domain names on NIS to the desired structure, then migrate to NIS+ domain by domain.

The Email Environment

Because NIS+ can have a domain hierarchy while NIS has a flat domain space, changing to NIS+ can have effects on your mail environment. With NIS, only one mailhost is required. If you use a domain hierarchy for NIS+, you may need one mailhost for each domain in the namespace because names in separate domains may no longer be unique.

Therefore, the email addresses of clients who are not in the root domain may change. As a general rule, client email addresses can change when domain names change or when new levels are added to the hierarchy.

In Solaris 2.x, these changes required a great deal of work. Later releases provided several `sendmail` enhancements to make the task easier. In addition, NIS+ provides a `sendmailvars` table. The `sendmail` program first looks at the `sendmailvars` table (see Table 2-1 on page 25), then examines the local `sendmail.cf` file.

Note – Be sure that mail servers reside in the NIS+ domain whose clients they support. Also for performance reasons, do not use paths to direct mail servers to tables in other domains.


Consider the impact of the new mail addresses on DNS. You may need to adjust the DNS MX records.

Select the Namespace Servers

Each NIS+ domain is supported by a set of NIS+ servers. The servers store the domain's directories, groups, and tables, and answer requests for access from users, administrators, and applications. Each domain is supported by only one set of servers. However, a single set of servers can support more than one domain.

Remember that a domain is not an object, but a reference to a collection of objects. Therefore, a server that supports a domain is not actually associated with the domain, but with the domain's directories. A domain consists of three directories: `domain`, `org_dir.domain`, and `groups_dir.domain`:

Figure 2-4 Server Relationship to Domain



Any Solaris 2.x-based workstation can be an NIS+ server as long as it has its own hard disk of sufficient size. The software for both NIS+ servers and clients is included in the Solaris product. Therefore, any workstation that has a Solaris 2.x release installed can become a server or a client, or both.

When selecting the servers that will support the NIS+ namespace, consider the following factors:

- Supported domains
- Server load
- Disk space and memory requirements
- NIS clients

Supported Domains


When selecting servers, you must differentiate between the requirements imposed by the NIS+ service and those imposed by the traffic load of your namespace.

The NIS+ service requires you to assign at least one server, the master, to each NIS+ domain¹. How many other servers a domain requires is determined by the traffic load, the network configuration, and whether NIS clients are present.

The traffic loads you anticipate will determine the total number of servers used to support the namespace, how much storage and processing speed each will require, and whether a domain needs replicas to ensure its availability. How can you determine how many servers you need?

A good way to begin is by assigning one master server to each domain in the hierarchy:

Figure 2-5 Assigning Servers to Domains




If certain domains must always be available, add two or more replicas to them. Two replicas allow requests to still be answered even if one of the replicas is damaged. Requests may not be answered in a timely manner if a master has only one replica and it is being repaired or updated. A domain with only one

1. An NIS+ server is capable of supporting more than one domain, but use this configuration only in small namespaces or testing situations.

replica loses fifty percent of its load capacity when the replica is down. Always add at least one replica to a domain. In small to medium domains, configurations with two to four replicas are normal.

Figure 2-6 Adding Replicas to a Domain



In organizations with many distributed sites, each site often needs its own subdomain. Typically, the subdomain master is placed in a higher-level domain. As a result, there can be a great deal of traffic between point-to-point links. Creating local replicas can speed request response and minimize point-to-point traffic across the link. In this configuration, lookups may be handled locally.

Server Load

NIS+ master servers require fewer replicas than NIS servers did, since NIS+ does not depend on broadcasts on the local subnet.

Putting replicas on both sides of a weak network link (such as wide area links) is recommended. If the link breaks and the networks are decoupled, both sides of the network can still obtain service.

Do not put more than 10 replicas on one domain. If you can, put one on each subnet; otherwise, distribute the servers as best you can and optimize for the best performance. You don't need NIS+ servers on every subnet, unless they support NIS clients. In such cases, you may want to install NIS+ servers on multi-homed machines.

Try to keep fewer than 1000 clients in a domain. NIS+ clients present a higher load on servers than NIS clients. A large number of clients served by only a few servers may impact network performance.

Disk Space and Memory Requirements

How much disk space you need depends on four factors:

- Disk space consumed by the Solaris 2.x software
- Disk space for `/var/nis` (and `/var/yp`)
- Amount of memory
- Swap space required for NIS+ server processes

The Solaris 2.x software can require over 220 Mbytes of disk space, depending on how much of it you install. For exact numbers, see *SPARC: Installing Solaris Software* or *x86: Installing Solaris Software*. You should also count the disk space consumed by other software the server may use. The NIS+ software itself is part of the Solaris 2.4 distribution, so it does not consume additional disk space.

NIS+ directories, groups, tables, and client information are stored in `/var/nis`. The `/var/nis` directory uses about 5 Kbytes of disk space per client. For example purposes only, if a namespace has 1000 clients, `/var/nis` requires about 5 Mbytes of disk space. However, because transaction logs (also kept in `/var/nis`) can grow large, you may want additional space per client—an additional 10–15 Mbytes is recommended. In other words, for 1000 clients, allocate 15 to 20 Mbytes for `/var/nis`. You can reduce this if you checkpoint transaction logs regularly. You should create a separate partition for `/var/nis`. This separate partition will help during an operating system upgrade.

If you will use NIS+ concurrently with NIS, allocate an equal amount of space to the amount you are allocating to `/var/nis` for `/var/yp` to hold NIS maps that you may transfer from NIS.

Although 32 Mbytes is the minimum memory requirement for servers, it is better to equip servers of medium to large domains with at least 64 Mbytes.

You also need swap space equal to three times or more of the size of the NIS+ server process—in addition to the server's normal swap space requirements. The size of the `rpc.nisd` process as shown by the `ps -efl` command. Most of this space is used during callback operations or when directories are

checkpointed (with `nisping -C`) or replicated, because during such procedures an entire NIS+ server process is forked. In no case should you use less than 64 Mbytes of swap space.

Determine Table Configurations

NIS+ tables provide several features not found in simple text files or maps. They have a column-entry structure, accept search paths, can be linked together, and can be configured in several different ways. You can also create your own custom NIS+ tables. When selecting the table configurations for your domains, consider the following factors:

- Differences between NIS+ tables and NIS maps
- Use of custom NIS+ tables
- Table location
- Connections between tables

Differences Between NIS+ Tables and NIS Maps

NIS+ tables differ from NIS maps in many ways, but two of those differences are worth keeping in mind during your namespace design:

- NIS+ uses fewer standard tables than NIS
- NIS+ tables interoperate with `/etc` files differently than NIS maps did in the SunOS 4.x releases

NIS+ Standard Tables

Review the 17 standard NIS+ tables to make sure they suit the needs of your site. They are listed in Table 2-1 on page 25. Table 2-2 on page 26 lists the correspondence between NIS maps and NIS+ tables.

Do not worry about synchronizing related tables. The NIS+ tables store essentially the same information as NIS maps, but they consolidate similar information into a single table (for example, the NIS+ hosts table stores the same information as the `hosts.byaddr` and `hosts.byname` NIS maps). Instead of the key-value pairs used in NIS maps, NIS+ tables use columns and rows. (See *Name Services Configuration Guide*.) Key-value tables have two columns with the first column being the key and the second column being the

value. Therefore, when you update any information, such as host information, you need only update it in one place, such as the hosts table. You no longer have to worry about keeping that information consistent across related maps.

Note the new name of the automounter tables:

- `auto_home` (old name: `auto.home`)
- `auto_master` (old name: `auto.master`)

The dots were changed to underscores because NIS+ uses dots to separate directories. Dots in a table name will cause NIS+ to mistranslate names.

To make the transition from NIS to NIS+, you must change the dots in your NIS automounter maps to underscores. You may also need to do this on your clients' automounter configuration files.

Table 2-1 NIS+ Tables

NIS+ Table	Information in the Table
hosts	Network address and hostname of every workstation in the domain
bootparams	Location of the root, swap, and dump partition of every diskless client in the domain
passwd	Password information about every user in the domain
cred	Credentials for principals who belong to the domain
group	The group password, group id, and members of every UNIX group in the domain
netgroup	The netgroups to which workstations and users in the domain may belong
mail_aliases	Information about the mail aliases of users in the domain.
timezone	The time zone of the domain
networks	The networks in the domain and their canonical names
netmasks	The networks in the domain and their associated netmasks
ethers	The ethernet address of every workstation in the domain
services	The names of IP services used in the domain and their port numbers
protocols	The list of IP protocols used in the domain
rpc	The RPC program numbers for RPC services available in the domain

Table 2-1 NIS+ Tables (Continued)

NIS+ Table	Information in the Table
auto_home	The location of all user's home directories in the domain
auto_master	Automounter map information
sendmailvars	Stores the mail domain

Table 2-2 Correspondence Between NIS Maps and NIS+ Tables

NIS Map	NIS+ Table	Notes
auto.home	auto_home	
auto.master	auto_master	
bootparams	bootparams	
ethers.byaddr	ethers	
ethers.byname	ethers	
group.bygid	group	Not the same as NIS+ groups.
group.byname	group	Not the same as NIS+ groups.
hosts.byaddr	hosts	
hosts.byname	hosts	
mail.aliases	mail_aliases	
mail.byaddr	mail_aliases	
netgroup	netgroup	
netgroup.byhost	netgroup	
netgroup.byuser	netgroup	
netid.byname	cred	
netmasks.byaddr	netmasks	
networks.byaddr	networks	
networks.byname	networks	
passwd.byname	passwd	
passwd.byuid	passwd	
protocols.byname	protocols	

Table 2-2 Correspondence Between NIS Maps and NIS+ Tables (Continued)

NIS Map	NIS+ Table	Notes
<code>protocols.bynumber</code>	<code>protocols</code>	
<code>publickey.byname</code>	<code>cred</code>	
<code>rpc.bynumber</code>	<code>rpc</code>	
<code>services.byname</code>	<code>services</code>	
<code>ypservers</code>		Not needed.

NIS+ has one new table for which there is no corresponding NIS table: `sendmailvars`. The `sendmailvars` table stores the mail domain used by `sendmail`.

NIS+ Tables Interoperate Differently With /etc Files

The way NIS interacted with `/etc` files in the SunOS 4.x environment—or other network information services—was controlled by the `/etc` files, using the `+/-` syntax. How NIS+, NIS, or DNS interacts with `/etc` files in the Solaris 2.x release—or other network information services—is determined by the name service switch. A configuration file, `/etc/nsswitch.conf`, located on every Solaris 2.x client, specifies the sources of information for that client: `/etc` files, DNS zone files (hosts only), NIS maps, or NIS+ tables. The `nsswitch.conf` configuration file of NIS+ clients resembles the example in Figure 2-7 on page 28.

Figure 2-7 Sample Name Service Switch File

```
passwd:      files nisplus
group:       files nisplus

hosts:       nisplus [NOTFOUND=return] files
services:    nisplus [NOTFOUND=return] files
networks:    nisplus [NOTFOUND=return] files
protocols:   nisplus [NOTFOUND=return] files
rpc:         nisplus [NOTFOUND=return] files
ethers:      nisplus [NOTFOUND=return] files
netmasks:   nisplus [NOTFOUND=return] files
bootparams:  nisplus [NOTFOUND=return] files

publickey:   nisplus

netgroup:    nisplus

automount:   files nisplus
aliases:     files nisplus
```

In other words, for most types of information, the source is first an NIS+ table, then an `/etc` file.

You can select from three versions of the switch configuration file or you can create your own. For instructions, see *Name Services Administration Guide*.

Use of Custom NIS+ Tables

Determine which nonstandard NIS maps you use and their purpose. Can they be converted to NIS+ or replaced with NIS+ standard maps?

Some applications may rely on NIS maps. Will they still function the same way with NIS+, and can they function correctly in a mixed environment?

To build a custom table in NIS+, use `nistbladm`. Remember that you cannot use dots in the table names.

If you want to be able to use NIS+ to support your custom NIS maps, you should create a key-value table, a table with two columns. The first column is the key and the second column is the value. If you then run the NIS+ servers in NIS-compatibility mode, the NIS clients will not notice any change in functionality.

Connections Between Tables

NIS+ tables contain information only about the resources and services in their home domain. If a client tries to find information that is stored in another domain, the client has to provide the other domain name. You can make this “forwarding” automatic by connecting the local table to the remote table. NIS+ tables can be connected in two different ways:


- Through *paths*
- Through *links*

Paths and links should not be used if you are going to have NIS clients in the NIS+ namespace, because NIS clients are unable to follow the paths or links to find the appropriate information.

Paths

If information in a particular NIS+ table is often requested by clients in other domains, consider establishing a path from the local NIS+ table to the one in the other domain. For example:

Figure 2-8 Establishing a Path to Tables in a Higher Domain



Such a path would have two main benefits. First, it would save clients in lower domains the trouble of explicitly searching through a second table. Second, it would allow the administrator in the higher-level domain to make changes in one table and render that change visible to clients in other domains. However, such a path would also hurt performance. Performance is especially affected when searches are unsuccessful, because the NIS+ service must search through two tables instead of one. When you use paths, a table lookup now also depends upon the availability of other domains. This dependence can reduce the net availability of your domain. For these reasons, use paths only if you do not have any other solution to your problem.

You should also be aware that since “mailhost” is often used as an alias, when trying to find information about a specific mailhost, you should use its fully qualified name in the search path (for example, mailhost.sales.wiz.com.); otherwise NIS+ will return *all* the mailhosts it finds in all the domains it searches through.

The path is established in the local table, with the `-p` option to the `nistbladm` command. To change a table’s path, you must have modify access to the table object. To find out what a table’s search path is, use the `niscat -o` command (you must have read access to the table).

Links

Links between tables produce an effect similar to paths, except that the link involves a search through only one table: the remote table. With a search path, NIS+ first searches the local table, and only if it is unsuccessful does it search the remote table. With a link, the search moves directly to the remote table. In fact, the remote table virtually replaces the local table.

The benefit of a link is that it allows a lower domain to access the information in a higher domain without the need to administer its own table.

To create a link, use the `nisl` command. You must have modify rights to the table object.


Deciding whether to use a path or to link NIS+ tables in a domain is a complex decision, but here are some basic principles to keep in mind:

- Every domain must have access to every standard table.
- Volatile, frequently accessed data should be located lower in the hierarchy. Such data should be located closer to where it is used most often.


- Data that is accessed by several domains should be located higher in the hierarchy, unless the domains need to be independent.
- The lower in the hierarchy you place data, the easier it will be to administer autonomously.
- Only NIS+ clients can see tables connected by paths and links. They cannot be seen by NIS clients.

Figure 2-9 summarizes this principle.

Figure 2-9 Information Distribution Across an NIS+ Hierarchy



Selecting NIS+ Security Measures



This chapter provides general guidelines and recommendations for making choices about security in your namespace.

<i>Determine the Impact of Implementing NIS+ Security</i>	<i>page 33</i>
<i>Select Credentials</i>	<i>page 35</i>
<i>Selecting a Security Level</i>	<i>page 37</i>
<i>Forming NIS+ Groups</i>	<i>page 37</i>
<i>Determine Access Rights to NIS+ Groups and Directories</i>	<i>page 38</i>
<i>Determine Access Rights to NIS+ Tables</i>	<i>page 39</i>

Determine the Impact of Implementing NIS+ Security

Because NIS+ provides security that NIS did not, it requires more administrative work. It may also require more work from users who are not used to performing `chkey`, `keylogin` or `keylogout` procedures. Furthermore, the protection provided by NIS+ is not entirely secure. Given enough computing power and the right knowledge, the Diffie-Hellman public key cryptography system can be broken. In addition, the secret key stored with the key server process is not automatically removed when a credentialed non-root user logs out unless that user logs out with `keylogout(1)`. Security may still be compromised even if the user logs out with `keylogout(1)` because the session keys may remain valid until they expire or are refreshed. (See `keylogout(1)` for more information.) Root's key, created by `keylogin -r`

and stored in `/etc/.rootkey`, remains until the `.rootkey` file is explicitly removed. The superuser cannot use `keylogout(1)`. Nevertheless, NIS+ is much more secure than NIS.

Keeping these trade-offs in mind, consider the impact NIS+ security will have on both users and administrators before you decide what security level to use.

How NIS+ Security Affects Users

NIS+ security benefits users because it improves the reliability of the information they obtain from NIS+ and it protects their information from unauthorized access. However, NIS+ security requires users to learn a bit about security and requires them to perform a few extra administrative steps.

Although NIS+ requires a network login, users are not required to perform an additional key login because the `login` command automatically gets the network keys for the client when the client has been correctly configured. A client is correctly configured when their login password and their network password are the same. The secret key for the user `root` is normally made available in the `/etc/.rootkey` file (a possible security problem as noted earlier). When the NIS+ user password and credential are changed with the `nispasswd` command, the credential information is automatically changed for the user.

- To change the NIS+ machine's local root password, run the `passwd` command.
- To change the root credential, run the `chkey` command.

However, if your site allows users to maintain passwords in their local `/etc/passwd` files in addition to their network passwords, and if these passwords are different from the network passwords, then users must run `keylogin` each time they run `login`. The reasons for this are explained in the security chapter in the *Name Services Administration Guide*.

How NIS+ Security Affects Administrators

Because Solaris 2.x bundles the DES encryption mechanism for authentication, administrators who need secure operation do not need to purchase a separate encryption kit. However, administrators must train users how to use the `nispasswd` and the `chkey` commands, and when to use them.

Furthermore, setting up a secure NIS+ namespace is more complex than setting up one without any security. The complexity comes not only from the extra steps required to set up the namespace, but from the job of creating and maintaining user and machine credentials for all NIS+ principals. Administrators have to remove obsolete credentials just as they remove dead account information from the `passwd` and `hosts` tables. Also, when servers' public keys change, administrators have to update the keys throughout the namespace (using `nissupdkeys`). Administrators also have to add LOCAL credentials for users from other domains who want to remote login to this domain and to have authenticated access to NIS+.

How NIS+ Security Affects Transition Planning

After you become familiar with the benefits and the administrative requirements of NIS+ security, you must decide whether to implement NIS+ security during or after the transition. It is recommended that you should use full NIS+ security even if you operate some or all servers in a domain in NIS-compatibility mode. (It is preferable though that all servers in a domain have the same NIS-compatibility status.) However, this entails a heavy administrative burden. If you prefer a simpler approach, you could set up the NIS+ servers and namespace with NIS-compatible security, but decline to create credentials for NIS+ clients. Administrators and servers would still require credentials. The NIS+ clients would be relegated to the nobody category, along with the NIS clients. This reduces training and setup requirements, but it has the following drawbacks:

- Users lose the ability to update any NIS+ tables and, in particular, their network passwords.
- Users will not be able to verify that the name service information is coming from an authenticated NIS+ server.

Select Credentials

NIS+ provides two types of credential: LOCAL and DES. All NIS+ principals need at least one of these credentials. If the namespace is running at security level 2 (the default), all NIS+ principals (clients) must have DES credentials in their home domain. In addition, all users (not workstations) must have LOCAL credentials in their home domains and in every other domain for which they need login access.


To determine the credential needs of your namespace, consider the

- Type of principal
- Type of credential

Type of Principal

NIS+ principals can be users or the superuser identity on the client workstation.

Figure 3-1 NIS+ Principals



When you determine the credentials you need to create, make sure you know which type of principal the credential is for. For instance, when you set up an NIS+ client (use `nisclient`), you will create credentials for both the workstation and for the user. Unless credentials for the user are also created, the user would only have the access rights granted to the nobody class. This can work perfectly well if that's how you choose to set up your namespace. But if you don't give any access rights to the nobody class, the namespace won't be available to the users.

NIS+ cannot distinguish between a human principal and a workstation principal when requests are made. Therefore, all user names must be different from machine names in the same namespace. For example, under NIS, it is acceptable to have a user with the login name of `jane` whose local machine is also named `jane`. So that user's network address is `jane@jane`. This user will have to change either her login or her machine name when her site converts to NIS+. Identical user and machine names are a problem even when the machine with the duplicate name does not belong to the user with the same name. The following examples illustrate duplicate name combinations not valid with NIS+:

- `jane@jane` in the same namespace
- `tree@jane` and `jane@tree` in the same namespace
- `joe@tree` and `tree@mike` in the same namespace

The best solution to this problem is to check all `/etc` files and NIS maps before you use the data to populate NIS+ tables. If you find duplicate names, change the machine names rather than the login names, and later create an alias for the machine's old name.

Selecting a Security Level

Even though NIS+ servers can operate at three security levels, you should not run your domain at lower than level 2, the highest level of NIS+ security unless you have a very good reason. Security level 2 is also the default level set when you use the NIS+ scripts to create your NIS+ namespace. If you run your domain at a lower level of security than level 2, then any user will have the access rights to modify the name service information and the namespace's structure. Security level 0 is only useful for creating the namespace. At this level, any user can become root on any machine in the network. Security level 1 uses only the user id for authorization and does no authentication of the user.

Security level 2 sends only requests that use DES credentials to appropriate NIS+ principals. Requests that only use LOCAL credentials or none at all are assigned the access rights granted to the nobody class. Requests that use invalid DES credentials are denied. See *Name Services Administration Guide* for more information on the NIS+ security levels.

Forming NIS+ Groups

NIS+ introduces a new type of group to name service administration that NIS does not have. An NIS+ group is only used as a means to provide NIS+ access rights to several NIS+ principals at one time. It is used only for NIS+ authorization.

The default name of the group created by NIS+ scripts for such purposes is the admin group. Member users of the admin group have special privileges, such as permission to make certain changes to the namespace information. For example, you could add several junior administrators to the admin group so that they can only modify the `passwd` and `hosts` tables, but they would be unable to modify any other tables. By using an admin group, you can distribute administration tasks across many users, and not just reserve them for the superuser of the entire hierarchy. The NIS+ admin group must have

credentials created for its members even if you are running the domain in NIS-compatibility mode because only authenticated users have permission to modify NIS+ tables.

After identifying the type of credentials you need, you should select the access rights that are required in the namespace. To make that task easier, you should first decide how many administrative groups you will need. Using separate groups is useful when you want to assign them different rights. Usually, you create groups by domain. Each domain should have only one admin group.

Determine Access Rights to NIS+ Groups and Directories

After arranging your principals into groups, determine the kind of access rights granted by the objects in the namespace to those groups, as well as to the other categories of principal (nobody, owner, group, and world). Planning these assignments ahead of time will help you establish a coherent security policy as shown in Table 3-1. NIS+ provides different default access rights for different namespace objects.

Table 3-1 Default Access Rights for NIS+ Objects

Object	Nobody	Owner	Group	World
Root directory object	r---	rmcd	rmcd	r---
Non-root directory object	r---	rmcd	rmcd	r---
groups_dir directory objects	r---	rmcd	rmcd	r---
org_dir directory objects	r---	rmcd	rmcd	r---
NIS+ groups	----	rmcd	r---	r---
NIS+ tables (see Table 3-2 on page 40)	varies	varies	varies	varies

You can use the default rights or assign your own. If you assign your own, you will have to consider how the objects in your namespace will be accessed. Keep in mind that the nobody class comprises all requests from NIS+ clients, whether authenticated or not. The world class comprises all authenticated requests from NIS+ clients. Therefore, if you don't want to provide namespace access to unauthenticated requests, don't assign any access rights to the nobody class; reserve them only for the world class. On the other hand, if you expect some clients—through applications, for instance—to make

unauthenticated read requests, you should assign read rights to the nobody class. If you want to support NIS clients in NIS-compatibility mode, you will have to assign read rights to the nobody class.

Also consider the rights each type of namespace object will assign to the NIS+ groups you specified earlier. Depending on how you plan to administer the namespace, you can assign all or some of the available access rights to the group. A good solution is to have the user root on the master server be the owner of the admin group. The admin group should have create and destroy rights on the objects in the root domain. If you only want one administrator to create and modify the root domain, then put just that administrator in the admin group. You can always add additional members to the group. If several administrators may be involved in the setup process, put them all in the group and assign full rights to it. That is easier than switching ownership back and forth.

Finally, the owner of an object should have full rights, although this is not as important if the group does. A namespace is more secure if you give only the owner full rights, but it is easier to administer if you give the administrative group full rights.

Determine Access Rights to NIS+ Tables

NIS+ objects other than NIS+ tables are primarily structural. NIS+ tables, however, are a different type of object: they are informational. Access to NIS+ tables is required by all NIS+ principals and applications running on behalf of those principals. Therefore, their access requirements are a bit different.

Table 3-2 on page 40 lists the default access rights assigned to NIS+ tables. If any columns provide rights in addition to those of the table, they are also listed. You can change these rights at the table and entry level with the `nischmod` command, and at the column level with the `nistbladm -u` command. "*Protecting the Encrypted Passwd Field*" on page 41 provides just one example of how to change table rights to accommodate different needs.

Table 3-2 Default Access Rights for NIS+ Tables and Columns

Table/Column	Nobody	Owner	Group	World
hosts table	r---	rmcd	rmcd	r---
bootparams table	r---	rmcd	rmcd	r---
passwd table	----	rmcd	rmcd	r---
name column	r---	----	----	----
passwd column ¹	----	-m--	----	----
uid column	r---	----	----	----
gid column	r---	----	----	----
gcos column	r---	-m--	----	----
home column	r---	----	----	----
shell column	r---	----	----	----
shadow column	----	----	----	----
group table	----	rmcd	rmcd	r---
name column	r---	----	----	----
passwd column	----	-m--	----	----
gid column	r---	----	----	----
members column	r---	-m--	----	----
cred table	r---	rmcd	rmcd	r---
cname column	----	----	----	----
auth_type column	----	----	----	----
auth_name column	----	----	----	----
public_data column	----	-m--	----	----
private_data column	----	-m--	----	----
networks table	r---	rmcd	rmcd	r---
netmasks table	r---	rmcd	rmcd	r---
ethers table	r---	rmcd	rmcd	r---
services table	r---	rmcd	rmcd	r---
protocols table	r---	rmcd	rmcd	r---

Table 3-2 Default Access Rights for NIS+ Tables and Columns (Continued)

Table/Column	Nobody	Owner	Group	World
rpc table	r---	rmd	rmd	r---
auto_home table	r---	rmd	rmd	r---
auto_master table	r---	rmd	rmd	r---

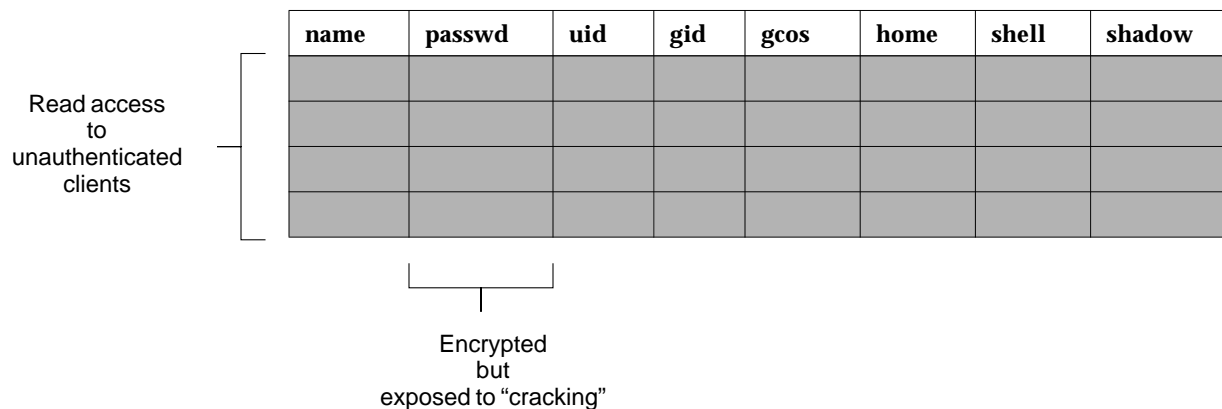
1. NIS-compatible domains give the nobody class read rights to the passwd table at the table object level.

Protecting the Encrypted Passwd Field

As you can see in Table 3-2, read access is provided to the nobody class by all tables except the passwd table. NIS+ tables give the nobody class read access because many applications that need to access NIS+ tables run as unauthenticated clients. However, if this were also done for the passwd table, it would expose the encrypted passwd column to unauthenticated clients.

The configuration shown in Figure 3-2 is the default set of access rights for NIS-compatible domains. NIS-compatible domains must give the nobody class read access to the passwd column because NIS clients are unauthenticated and would otherwise be unable to access their passwd column. Therefore, in an NIS-compatible domain, even though passwords are encrypted, they are vulnerable to decoding. They would be much more secure if they were not readable by anyone except their owner.

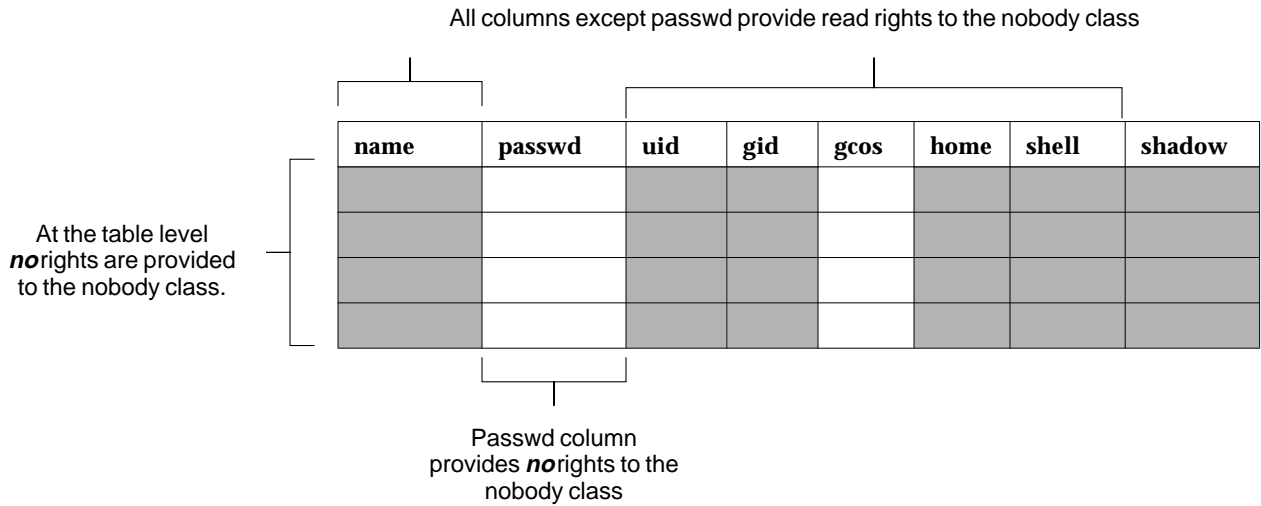
Figure 3-2 Default Passwd Table Access Rights for NIS-compatible Domains



Standard NIS+ domains (not NIS-compatible) provide that extra level of security. The default configuration (provided by `nissetup`) uses a column-based scheme to hide the `passwd` column from unauthenticated users while still providing access to the rest of the `passwd` table.

At the table level, no unauthenticated principals have read access. At the column level, they have read access to every column except the `passwd` column:

Figure 3-3 Default Passwd Table Access Rights for Standard NIS+ Domains



How does an entry owner get access to the `passwd` column? Entry owners have both read and modify access to their own entries. They obtain read access by being a member of the world class. (Remember that at the table level, the world class has read rights.) They obtain modify access by explicit assignment at the column level:

		Nobody	Owner	Group	World
Passwd Table Rights	:	----	rmd	rmd	r---
Passwd Column Rights	:	----	-m--	----	----

Keep in mind that table owners and entry owners are rarely and not necessarily the same NIS+ principals. Thus, table-level read access for the owner does not imply read access for the owner of any particular entry.

As mentioned earlier, this is the default setup from the Solaris 2.3 release forward. To use this scheme in earlier versions of the Solaris software, follow the instructions in the task titled, “How to Limit Access to the Pswd Column,” in *Name Services Administration Guide*.

Using NIS-Compatibility Mode

4

Deciding whether and how to run NIS+ in parallel to NIS—and when to stop—is probably one of the most difficult transition issues you will face. NIS+ provides several features that allow it to operate in parallel with NIS; notably, the NIS-compatibility mode.

The essential benefit provided by NIS-compatibility mode is that it does not require you to make any changes to NIS clients. The essential drawback is that you will not be able to take advantage of full NIS+ security and hierarchy. You may have to change those clients' domain names, and they will not be able to update their network passwords with the `yppasswd` command.

Table 4-1 shows which NIS protocols are supported by NIS+ servers in NIS-compatibility mode.

Table 4-1 Support for NIS Protocols by NIS+ Servers in NIS Compatibility-Mode

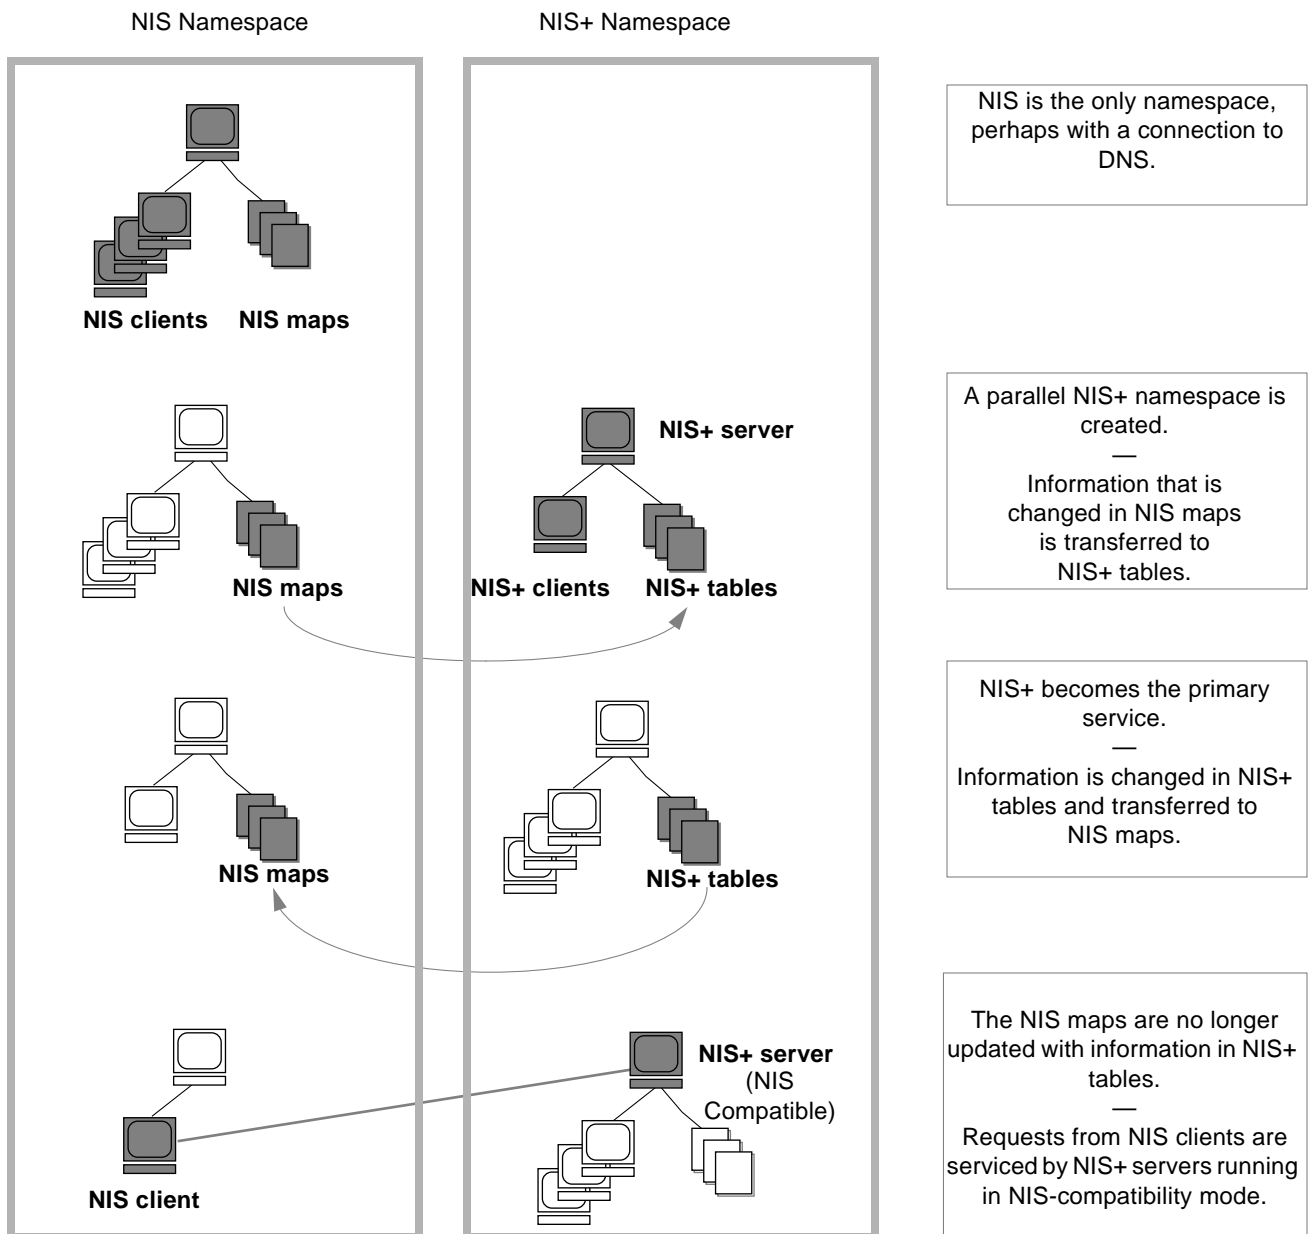
NIS Protocols	Compatibility Description
NIS client V2 protocol	Supported
NIS server-to-server protocol	Unsupported
NIS client update protocol	Unsupported
NIS client V1 protocol	Not supported except for <code>YPPROC_NULL</code> , <code>YPPROC_DOMAIN</code> and <code>YPPROC_DOMAIN_NONACK</code>

Figure 4-1 on page 47 illustrates how you can convert from an NIS-only namespace to a namespace that responds to both NIS and NIS+ requests. Remember, the principal benefit of the NIS-compatibility mode is that it provides NIS service to NIS clients. Its principal drawback is that it allows any unauthenticated client to obtain information from NIS+ tables.

If you plan to use the NIS-compatibility mode, you have to make the following decisions:

<i>Select the Domains That Will Be NIS Compatible</i>	<i>page 48</i>
<i>Determine the Configuration of NIS-Compatible Servers</i>	<i>page 48</i>
<i>Decide How to Transfer Information Between Services</i>	<i>page 49</i>
<i>Decide How to Implement DNS Forwarding</i>	<i>page 51</i>

Figure 4-1 Transition to NIS-Compatibility Mode



Select the Domains That Will Be NIS Compatible

Make a list of your NIS clients and group them in their eventual NIS+ domains. If the NIS+ domain running in NIS-compatibility mode does not have the same name as its NIS clients' original NIS domain, you must change the NIS clients' domain name to the NIS+ domain's name that is being supported by the NIS-compatible NIS+ server.

In NIS-compatibility mode, NIS client users can no longer use the `yppasswd` command to change their own passwords. Instead, system administrators with authorization or members of the NIS+ admin group have to run the `nispasswd` command for them.

Changing passwords raises the subject of synchronization of information between the NIS+ and NIS domains. You must devise a way to keep the NIS+ domains in synchronization with the NIS domains. At first, you should update the NIS+ tables nightly with the latest NIS information. If you are using NIS domains as the master source of name service information at your site, all changes made to the NIS+ tables between synchronization updates will be erased every time you propagate an NIS update. For example, passwords that were changed with the `nispasswd` command will be overwritten by the previous NIS password, unless you also run `yppasswd` whenever you run the `nispasswd` command. When you convert to an all NIS+ namespace, you will no longer need to run the `yppasswd` command.

At first, NIS will no doubt be the primary service. As you become familiar with the intricacies of sharing information, you will be able to plan a transition to make NIS+ the primary service.

Some NIS+ users may want the capability to switch back and forth between the main NIS domain and the new NIS+ domain. The `nisclient` script can help them do this easily when backup files are made.

Determine the Configuration of NIS-Compatible Servers

Take stock of your NIS servers, keeping in mind the requirements for your NIS+ servers. If you plan to eventually use them for the NIS+ service, upgrade them to the NIS+ recommendations. Identify which NIS servers will be used to support which NIS+ domains, and in what capacity, (either master or replica). Remember that NIS+ servers belong to the domain *above* the one they support

(except for the root domain servers). Since NIS+ servers do not belong to the domain they serve, you will not be able to use the same machines for other services that require domain dependent information.

If possible, plan to use your NIS+ server machines only for NIS+. This arrangement could require you to transfer other network services, such as DNS name services, boot server, home directories, NFS servers, and so on, to non-NIS+ server machines.

At many sites, the NIS server plays multiple roles, such as NFS server, compute server, `rlogin` server, and `mailhost` server. Because the NIS server uses the same information to resolve its names as do its clients, the NIS server can provide other services as well. As discussed in “The Domain Hierarchy” on page 12, except for root domains, all NIS+ servers live in the domain above the ones that they serve. So either do not run those services on an NIS+ server that require access to the name service, or use other means, such as files in `nsswitch.conf`, to acquire this same information. This problem would be solved if there is no hierarchy; in which case, the NIS+ root servers will live in the domain that they serve. The resource requirements of an NIS+ server are greater than those of an NIS server; hence it is advisable to not run other services along with NIS+.

If you have non-Solaris machines on your network, then either you can continue to run your NIS+ servers in NIS-compatibility mode, or you can move all such machines into their own domain. If you move all non-Solaris machines to one subnet, you can remove the restriction of having NIS+ servers on the same subnet as their NIS-compatible clients require. This will reduce the number of replicas required for any domain.

Decide How to Transfer Information Between Services

To keep the information synchronized, be sure to make one namespace subordinate to the other. At first, the NIS namespace may be the dominant one, in which case you would make changes to the NIS maps and load them into the NIS+ tables. In effect, the NIS namespace would be the master database.

An NIS+ server in NIS-compatibility mode supports standard NIS maps. An exhaustive list of these maps is in the NOTES section of the `ypfiles(4)` man page. But there are some limitations on map support. The NIS+ server serves

`yptest` requests only on the `netgroup` map, and not on the reverse maps. It does not support enumeration requests. (For example, `yptest`.) The `passwd.adjunct` map is not supported, either.

Eventually, the NIS+ namespace would be dominant. When that is the case, you would make changes in the NIS+ tables and copy them to the NIS maps.

The NIS+ command `nisaddent` and the NIS+ script `nispopulate` transfer information between NIS maps and NIS+ tables, as summarized in Table 4-2.

Table 4-2 NIS+ Data Transfer Commands

NIS+ Command	Description
<code>/usr/lib/nis/nisaddent -y</code>	Transfers information from an NIS map to an NIS+ table after you run <code>yptest</code> to transfer maps from an NIS server to the local disk. Nonstandard NIS maps can be transferred to NIS+ tables if the information is in key/value pairs. Multi-columned maps will be not be transferred.
<code>/usr/lib/nis/nisaddent -d</code>	Copies information from an NIS+ table into a file, which can then be transferred into an NIS map with standard NIS utilities.
<code>/usr/lib/nis/nispopulate -Y</code>	Transfers information from NIS maps into NIS+ tables.

A Note About Changing Information in the Password Table

While using parallel NIS and NIS+ namespaces, be sure to keep the information in the NIS+ `passwd` table synchronized with the information in the NIS `passwd` map. The `nispasswd` and `yppasswd` commands affect only their respective databases.

NIS+ uses the `nispasswd` command to change information stored in the NIS+ `passwd` table. It is the equivalent of the `yppasswd` command for NIS. The `nispasswd` command provides capabilities that are similar to those offered by other commands. Table 4-3 on page 51 summarizes their differences.

Table 4-3 Commands for Changing Passwords

Command	Description
<code>passwd</code>	Changes information in the workstation's <code>/etc/passwd</code> and <code>/etc/shadow</code> file.
<code>yppasswd</code>	Changes information in the NIS <code>passwd</code> map. Has no effect on the NIS+ <code>passwd</code> table or local files.
<code>nispasswd</code>	Changes and displays information in the NIS+ <code>passwd</code> table. When a principal's password is changed, <code>nispasswd</code> tries to update the principal's secret key in the <code>cred</code> table. Although <code>nistbladm</code> could be used to edit the NIS+ <code>passwd</code> table, <code>nispasswd</code> is better because its options are customized for the <code>passwd</code> table, including password encryption, which makes the command easier to use.
<code>nistbladm</code>	Creates, changes, and displays information about any NIS+ table, including the <code>passwd</code> table. Although the <code>nispasswd</code> command is easier to use for the <code>passwd</code> table, <code>nistbladm</code> allows you to: <ul style="list-style-type: none"> - Create new entries - Delete an existing entry - Change the UID and GID fields in the <code>passwd</code> table - Change the <code>LastChanged</code>, <code>Inactive</code>, and <code>Expired</code> fields in the <code>shadow</code> table - Change access rights and other security-related attributes of the <code>passwd</code> table. NOTE: (Do <i>not</i> use it to change the passwords themselves, though.)

If you are the owner of the `passwd` table, you can change password information without constraints. However, if you are not the owner, you must comply with construction constraints.

In domains created with NIS-compatibility mode, the permissions are slightly different: permissions at the column level provide read access to the `nobody` category.

Decide How to Implement DNS Forwarding

NIS servers can forward DNS requests made from Solaris 1.x NIS clients. NIS+ servers running in NIS-compatibility mode also provide DNS forwarding, but only starting with the Solaris 2.3 or later releases. (This feature is available in

the Solaris 2.2 release through a patch.) As a result, Solaris 2.x NIS clients must have appropriate `/etc/nsswitch.conf` and `/etc/resolv.conf` files installed locally.

Solaris 1.x NIS clients being supported by Solaris 2.0 or 2.1 servers running in NIS-compatibility mode will not be able to take advantage of DNS forwarding. You must upgrade those servers to Solaris 2.3 or later releases.

If the DNS domains are repartitioned, you must redefine new DNS zone files. Clients, however, may require updates to their `/etc/resolv.conf` file. A client, if it is also a DNS client, can set up its Name Service Switch configuration file to search for host information in either DNS zone files or NIS maps—in addition to NIS+ tables.

DNS Forwarding for NIS+ Clients

NIS+ clients do *not* have implicit DNS forwarding capabilities like NIS clients do. Instead, they take advantage of the Name Service Switch. To provide DNS capabilities to an NIS+ client, change its hosts entry to:

```
hosts: nisplus dns [NOTFOUND=return] files
```

DNS Forwarding for Solaris 2.x NIS Clients

If an NIS client is using the DNS forwarding capability of an NIS-compatible NIS+ server, the client's `nsswitch.conf` file should *not* have the following syntax in the hosts file:

```
hosts: nis dns files # not for NIS clients
```

Since DNS forwarding automatically forwards host requests to DNS, the syntax shown above would cause both the NIS+ server and the name service switch to forward unsuccessful requests to the DNS servers, impacting performance.

In addition, NIS clients running the Solaris 1.x release must have a resolver library. Be sure to install the `shlib_custom` software category. In this case, Solaris 1.x NIS clients would only use DNS for host lookup.

NIS and NIS+ Command Equivalents in the Solaris 1.x and 2.x Releases

The tables in this section give a quick overview of the differences between Solaris 1.x NIS commands, Solaris 2.x NIS commands, and their NIS+ equivalents. Table 4-4 on page 54 describes which NIS commands are supported in the Solaris 2.x release. Table 4-5 on page 55 and Table 4-6 on page 56 describe the NIS+ equivalents to NIS client and server commands in the Solaris 2.x release. Table 4-7 on page 57 contains a list of the NIS application programmer's interface functions and their NIS+ API equivalents, if they exist. See the appropriate man pages for details.

NIS Commands Supported in the Solaris 2.x Release

Only some NIS commands are supported in the Solaris 2.x release. NIS server commands are not shipped with the Solaris 2.x release. Just the NIS client commands are included. Whether these NIS commands run also depends on whether a Solaris 2.x NIS client is making requests of an NIS server or of an NIS+ server in NIS-compatibility mode. NIS clients cannot make updates to NIS+ servers that are running in NIS-compatibility mode. For example, such clients cannot run the commands `chkey`, `newkey`, and `yppasswd`. Table 4-4 on page 54 lists the NIS commands supported in the Solaris 2.x release.

Table 4-4 NIS Command Support in the Solaris 2.x Releases

Command Type	NIS Commands Supported in the Solaris 2.x Release	NIS Commands Not Supported in the Solaris 2.x Release
Utilities	ypinit ypxfr ypcat ypmatch yppasswd ypset ypwhich	yppush yppoll ypchsh ypchfn ypmake
Daemons	ypbind	ypserv ypxfrd rpc.yppupdated rpc.yppasswdd
NIS API	yp_get_default_domain() yp_bind() yp_unbind() yp_match() yp_first() yp_next() yp_all() yp_master() yperr_string() ypprot_err()	yp_order() yp_update()

Client and Server Command Equivalents

The two tables in this section contain NIS commands and their approximate NIS+ equivalents. The commands have been divided into two categories; Table 4-5 on page 55 contains those commands that are name service client to name service server and Table 4-6 on page 56 contains those commands that are name service server to name service server.

The commands shown in Table 4-5 on page 55 are client to name server commands. These commands are typed on name service client machines and request information of name service servers. The commands in column one

will run on Solaris 1.x or 2.x NIS clients connected to Solaris 1.x NIS servers. The commands in column two will run on Solaris 1.x or 2.x NIS clients connected to Solaris 2.x NIS+ servers running in NIS-compatibility mode.

Note – The command `ypinit -c` is only available, however, on Solaris 2.x NIS clients.

The commands in column three will only run on Solaris 2.x NIS+ clients connected to Solaris 2.x NIS+ servers. Commands are approximately equivalent across rows. N/A indicates that an equivalent command does not exist for that case.

Table 4-5 NIS Client Commands and Equivalent NIS+ Commands

SunOS 4.x NIS Server	NIS+ Server in NIS-Compatibility Mode	NIS+ Server
<code>ypwhich -m</code>	<code>ypwhich -m</code>	<code>niscat -o</code>
<code>ypcat</code>	<code>ypcat**</code>	<code>niscat</code>
<code>ypwhich</code>	<code>ypwhich</code>	N/A
<code>ypmatch</code>	<code>ypmatch</code>	<code>nismatch/nisgrep</code>
<code>yppasswd</code>	N/A	<code>nispasswd</code>
<code>ypbind</code>	<code>ypbind</code>	N/A
<code>yppoll</code>	N/A	N/A
<code>ypset</code>	<code>ypset</code>	N/A
N/A	<code>ypinit -c</code>	<code>nisinit</code>

Note – ** The `ypcat` command is not supported for queries directed to the `netgroup` table. The NIS client's request will time-out before an answer is received because this table's format is so different from the `netgroup` NIS map's format.

The commands shown in Table 4-6 on page 56 are name server to name server commands. The NIS server commands are not included in the Solaris 2.x release, so they are not available to either NIS+ servers nor NIS+ servers in NIS-compatibility mode. In addition, an NIS server cannot make updates to an

NIS+ server, nor is the reverse situation possible. The third column of this table lists the NIS+ server commands that are equivalent to the NIS server commands listed in the first column. There are no exact equivalents for servers in NIS-compatibility mode because NIS-compatibility mode only refers to client commands.

Table 4-6 NIS Server Commands and Equivalent NIS+ Commands

SunOS 4.x NIS Server	NIS+ Server in NIS-Compatibility Mode	NIS+ Server
ypxfr	N/A	N/A
makedbm	N/A	nisaddent
ypinit -m ypinit -s	N/A	nissserver
ypserv	rpc.nisd -Y	rpc.nisd
ypserv -d	rpc.nisd -Y -B	no DNS forwarding required, use /etc/nsswitch.conf
ypxfrd	N/A	N/A
rpc.yupdated	N/A	N/A
rpc.yppasswd	N/A	N/A
yppush	N/A	nisping
ypmake	N/A	nissetup,nisaddent

NIS and NIS+ API Function Equivalents

To completely convert your site to NIS+, all applications have to be ported to NIS+ in addition to changing the name service. Any internally created applications that make NIS calls have to be modified to use NIS+ calls. Otherwise, you will always have to run your NIS+ servers in NIS-compatibility mode with all the drawbacks that this mode entails. External applications may force you to run your namespace in NIS-compatibility mode until they are updated as well.

Table 4-7 contains a list of the NIS application programmer's interface functions and their NIS+ API equivalents, if they exist.

Table 4-7 NIS API and NIS+ API Equivalent Functions

NIS API Functions	NIS+ API Functions
<code>yp_get_default_domain()</code>	<code>nis_local_directory()</code>
<code>ypbind()</code>	N/A
<code>ypunbind()</code>	N/A
<code>ypmatch()</code>	<code>nis_list()</code>
<code>yp_first()</code>	<code>nis_first_entry()</code>
<code>yp_next()</code>	<code>nis_next_entry()</code>
<code>yp_all()</code>	<code>nis_list()</code>
<code>yp_master()</code>	<code>nis_lookup()</code>
<code>yperr_string()</code>	<code>nis_perror()</code> , <code>nis_sperrno()</code>
<code>ypprot_err()</code>	<code>nis_perror()</code> , <code>nis_sperrno()</code>
<code>yp_order()</code>	N/A
<code>yp_update()</code>	<code>nis_add_entry()</code> , <code>nis_remove_entry()</code> , <code>nis_modify_entry()</code>

Prerequisites to Transition

5 

This chapter describes several miscellaneous tasks that must be carried out before beginning the transition:

<i>Gauge the Impact of NIS+ on Other Systems</i>	<i>page 59</i>
<i>Train Administrators</i>	<i>page 60</i>
<i>Write a Communications Plan</i>	<i>page 60</i>
<i>Identify Required Conversion Tools and Processes</i>	<i>page 61</i>
<i>Identify Administrative Groups Used for Transition</i>	<i>page 61</i>
<i>Determine Who Will Own the Domains</i>	<i>page 62</i>
<i>Determine Resource Availability</i>	<i>page 63</i>
<i>Resolve Conflicts Between Login Names and Hostnames</i>	<i>page 63</i>
<i>Examine All Information Source Files</i>	<i>page 64</i>
<i>Remove the "." From NIS Map Names</i>	<i>page 64</i>
<i>Document Your Existing NIS Namespace</i>	<i>page 65</i>
<i>Create a Conversion Plan for Your NIS Servers</i>	<i>page 65</i>

Gauge the Impact of NIS+ on Other Systems

Perform a formal introduction, testing, and familiarization program at your site, not only to train administrators, but also to uncover dependencies of other systems or applications on NIS that will be affected by a transition to NIS+.

For instance, some applications may rely on some of the NIS maps. Will they function with standard or custom NIS+ tables? How will their need for access affect your overall security plan?

What nonstandard NIS maps are being used at your site? Can you convert them to NIS+ tables or create nonstandard NIS+ tables to store their information? Be sure to check their access rights.

Does your site use locally-built applications that depend on NIS? Do you have commands or applications that make direct NIS calls such as embedded `yp_match()` function calls? (See "*NIS and NIS+ API Function Equivalents*" on page 56 for more information.)

Do you have any duplicate user and host names in your namespace? (See "*Type of Principal*" on page 36 for more information.)

How will the network installation procedures be affected by the transition to NIS+? Analyze the changes required, if any.

Gauging the impact of NIS+ on your site administrative practices will help uncover potential roadblocks.

Train Administrators

Another goal of the introduction and familiarization program discussed in "*Become Familiar With NIS+*" on page 8 is to give the administrators at your site an opportunity to become familiar with NIS+ concepts and procedures. Classroom instruction alone is insufficient. Administrators need a chance to work in a harmless test environment. The training should consist of:

- A formal NIS+ concepts and administration course
- Basic NIS+ troubleshooting information and practice
- Information about your site's implementation strategy and plans.

Write a Communications Plan

Prepare a plan to communicate your intentions to users long before you actually begin converting clients to NIS+. Tell them about the implementation plan and give them a way to obtain more information. As mentioned in Chapter 1, "Introduction," a typical transition goal is to keep the impact of the

transition on clients to a minimum, but users might become concerned about the upcoming change. Send out email notices, conduct informative seminars, and designate email aliases or individuals to which users can send questions.

Identify Required Conversion Tools and Processes

Consider creating or obtaining transition tools to help with the implementation. If your site already uses automated tools to administer individual systems or network services, consider porting them to operate under the versions of Solaris software and NIS+ that you will be using for the transition (see “Use a Single Release of Software” on page 7). Here are some suggestions for scripts you might want to write:

- A script to convert users to NIS+—make additions to the `nisclient` shell script
- A check script to verify the correctness of a user’s NIS+ environment
- Backup and recovery scripts
- `crontab` entries for routine NIS+ maintenance
- Procedures for notification of outages

Scripts such as these will ensure that the transition is carried out uniformly across domains, speed the transition, and reduce complaints.

You should also prepare a set of standard configuration files and options, such as `nsswitch.conf`, that all clients across the namespace can use.

Identify Administrative Groups Used for Transition

Be sure that the NIS+ groups created as part of your namespace design (see “Forming NIS+ Groups” on page 37) correspond to the administrative resources you have identified for the transition. You could require a different set of NIS+ groups for the transition than for routine operation of an NIS+ namespace. Consider adding remote administrators to your groups in case you need their help in an emergency.

Make sure that group members have the proper credentials, that namespace objects grant the proper access rights to groups, and that the right group is identified as the group owner of the right namespace objects.

You might find this summary of commands that operate on NIS+ groups and group permissions useful:

Table 5-1 NIS+ Commands for Groups

Command	Description
<code>nisgrpadm</code>	Create or delete groups, add, change, list, or delete members
<code>niscat -o</code>	Display the object properties of an NIS+ group
<code>nissetup</code>	Create the basic structure of the directory in which a domain's groups are stored: <code>groups_dir</code>
<code>nisls</code>	List the contents of a directory
<code>NIS_GROUP</code>	An environment variable that overrides the value of <code>nisdefaults</code> for the shell in which it is set
<code>nischmod</code>	Changes an object's access rights
<code>nischown</code>	Changes the owner of an NIS+ object
<code>nischgrp</code>	Changes the group owner of an NIS+ object
<code>nistbladm -u</code>	Changes access rights to NIS+ table columns
<code>nisdefaults</code>	Displays or changes the current NIS+ defaults

Determine Who Will Own the Domains

To take complete advantage of the features inherent in a domain hierarchy, distribute the ownership of domains to the organizations they are dedicated to supporting. This will free the administrators of the root domain from performing rudimentary tasks at the local level.

Once you know who owns what, you can provide guidelines for creating administrative groups and setting their access rights to objects.

Consider how to coordinate the ownership of NIS+ domains with the ownership of DNS domains. Here are some guidelines:

- The administration of the DNS domain structure should remain the responsibility of the highest-level administrative group at the site.
- This same administrative group also owns the top-level NIS+ domain.

- Responsibility for the administration of lower-level DNS and NIS+ domains is delegated to individual sites by the top-level administrative group. If the NIS+ domains will be created along the same principles as the DNS domains (for instance, organized geographically), this delegation will be simple to explain.

Determine Resource Availability

Determine what administrative resources will be required for the implementation. These will be above and beyond the resources required for normal operation of NIS+. If your transition will involve a long period of NIS+ and NIS compatibility, additional resources may be required.

Consider not only the task of implementing the namespace design, but also of converting the numerous clients and dealing with special requests or problems. Keep in mind that NIS+ has a steep learning curve. Administrators may be less efficient for a while at performing support functions with NIS+ than they were with NIS. Consider not only formal training, but extensive lab sessions with hands-on experience.

Finally, even after the transition is complete, administrators will require extra time to become familiar with the everyday work flow of supporting NIS+.

Also consider the hardware resources. NIS servers are often used to support other network services such as routing, printing, and file management. Because of the potential load on an NIS+ server, you should use dedicated NIS+ servers. This load-balancing simplifies the transition because it makes troubleshooting and performance monitoring simpler. Of course, you incur the cost of additional systems. How many servers you will need and how they should be configured is described in Chapter 2, “Designing the NIS+ Namespace.”

Remember, these servers are required in *addition* to the NIS servers. Although the NIS servers might be decommissioned or recycled after the transition is complete, the NIS+ servers will continue to be used.

Resolve Conflicts Between Login Names and Hostnames

The NIS+ authentication scheme does not allow workstations and users to use the same names within a domain; for example, joe@joe is not permitted. Since NIS+ does not distinguish between credentials for hosts and for login names,

you can only use one credential type per name. If you have duplicate names in your namespace and you must keep the duplicate host name for some other reason make the following change: Retain the user login name and alias the duplicate host names. Create a new name for the host and use the old name as an alias for the new name. See "*Type of Principal*" on page 36 for examples of illegal name combinations.

You will need to resolve name conflicts before the implementation can begin, but you should also plan on permanently checking new workstations and user names during routine NIS+ operation. The `nisclient` script does name comparisons when you use it to create a client credential.

Examine All Information Source Files

Check all `/etc` files and NIS maps for empty fields or corrupted data before configuring NIS+. The NIS+ table populating scripts and commands may not succeed if the data source files contain empty fields or extraneous characters. Fill blank fields or fix the data before you start. It is better to delete questionable users or machines from the `/etc` files or NIS maps before running NIS+ scripts, and then add them back later after NIS+ is installed, than to proceed with the scripts and possibly corrupt data.

Remove the "." From NIS Map Names

As described in Chapter 2, "Designing the NIS+ Namespace", NIS+ automounter tables have replaced the "." in their names and file contents with underscores. You also need to make this change to the names of NIS maps that you will use during the transition. If you do not, NIS+ will confuse the "." in the name with the periods that distinguish domain levels in object names.

Note – Be sure to convert the "." to underscores for *all* NIS maps, not just those of the automounter. Be aware, however, that changing the names of nonstandard NIS maps from dots to underscores may cause applications that use those nonstandard maps to fail unless you also modify the applications to recognize NIS+ syntax.

Document Your Existing NIS Namespace

Documenting your current configuration will give you a clear point of departure for the transition. Make a list of the following items:

- Name and location of all current NIS domains and networks
- Hostname and location of all current NIS servers, both master and slave
- Configuration of all current NIS servers, including:
 - Hostname
 - CPU type
 - Memory size
 - Disk space available
 - Name of administrators with root access
- Nonstandard NIS maps

Correlate the list of your NIS clients with their eventual NIS+ domains. They will need to be upgraded to the Solaris 2.x release.

Create a Conversion Plan for Your NIS Servers

Take stock of your NIS servers. Although you can recycle them after the transition is complete, keep in mind that you will go through a stage in which you will need servers for *both* services. Therefore, you cannot simply plan to satisfy all your NIS+ server needs with your existing NIS servers.

You will find it helpful to create a detailed conversion plan for NIS servers, identifying which NIS servers will be used for NIS+ and when they will be converted. Do not use the NIS servers as NIS+ servers during the first stages of NIS to NIS+ transition. As described in Chapter 6, “Implementing the Transition”, the implementation is most stable when you check the operation of the entire namespace as a whole before you convert any clients to NIS+.

Assign NIS servers to NIS+ domains and identify each server’s role (master or replica). Once you have identified the servers you plan to convert to NIS+ service, upgrade them to NIS+ requirements (see “Disk Space and Memory Requirements” on page 23).

Implementing the Transition

6

Once you have performed the tasks described in the previous chapters, most of the hard work is done. Now all you have to do is set up the namespace you designed and add the clients to it. This chapter describes how to do that. Before performing these steps, verify that all your pre-transition tasks have been completed and that users at your site are aware of your plans.

If you will be running NIS+ domains alongside DNS domains, you will set up one NIS+ sub domain with each DNS domain. After you have set up a complete NIS+ namespace along with the first DNS domain and have verified that everything is working right, then you can set up the other NIS+ namespaces in parallel.

These are the major implementation phases:

<i>Phase I—Set Up the NIS+ Namespace</i>	<i>page 68</i>
<i>Phase II—Connect the NIS+ Namespace to Other Namespaces</i>	<i>page 69</i>
<i>Phase III—Make the NIS+ Namespace Fully Operational</i>	<i>page 70</i>
<i>Phase IV—Upgrade NIS-Compatible Domains</i>	<i>page 71</i>

Phase I—Set Up the NIS+ Namespace

Set up the namespace with full DES authentication, even if the domains will operate in NIS-compatibility mode. Use the NIS+ scripts in *Name Services Configuration Guide* to set up your namespace and see *Name Services Administration Guide* for more explanation of the basic steps. Then perform the following steps:

1. Set up the root domain.

If you are going to run the root domain in NIS-compatibility mode, use the `-Y` flag of `rpc.nisd` and `nissetup` (only necessary if not using scripts).

2. Populate the root domain tables.

You can transfer information from NIS maps or text files, using `nispopulate`. Of course, you can also create entries one at a time with `nistbladm` or `nisaddent`.

3. Set up clients of the root domain.

Set up a few clients in the root domain so that you can test its operation properly. Use full DES authentication. Some of these client machines will later be converted to root replica servers and some will serve as workstations for the administrators who support the root domain. NIS+ servers should never be an individual's workstation.

4. Create or convert site-specific NIS+ tables.

If the new NIS+ root domain will require custom, site-specific NIS+ tables, create them, using `nistbladm`, and transfer the NIS data into them, using `nisaddent`.

5. Add administrators to root domain groups.

Remember, the administrators must have LOCAL and DES credentials (use `nisaddcred`). Their workstations should be root domain clients and their root identities should also be NIS+ clients with DES credentials.

6. Update the sendmailvars table, if necessary.

If your email environment has changed as a result of the new domain structure, populate the root domain's `sendmailvars` table with the new entries.

7. Set up root domain replicas.

First convert the clients into servers (use `rpc.nisd` with `-Y` for NIS compatibility and use also `-B` if you want DNS forwarding), and then associate the servers with the root domain using `nissserver -R`.

8. Test the root domain's operation.

Develop a set of installation-specific test routines to verify a client's functioning after the switch to NIS+. This will speed the transition work and reduce complaints. You should operate this domain for about a week before you begin converting other users to NIS+.

9. Set up the remainder of the namespace.

Do not convert any more clients to NIS+, but go ahead and set up all the other domains beneath the root domain. This includes setting up their master and replica servers. Test each new domain as thoroughly as you tested the root domain until you are sure your configurations and scripts work properly.

10. Test the operation of the namespace.

Test all operational procedures for maintenance, backup, recovery, and other scenarios. Test the information-sharing process between all domains in the namespace. Do not proceed to Phase II until the entire NIS+ operational environment has been verified.

11. Customize the security configuration of the NIS+ domains.

This may not be necessary if everything is working well, but if you want to protect some information from unauthorized access, you can change the default permissions of NIS+ tables so that even NIS clients would be unable to access them. You could also rearrange the membership of NIS+ groups and the permissions of NIS+ structural objects to align with administrative responsibilities.

Phase II—Connect the NIS+ Namespace to Other Namespaces

12. [optional] Connect the root domain to the DNS namespace.

An NIS+ client can be connected to the Internet using the name service switch. Workstations, if they are also DNS clients, can have their name service switch configuration files set to search for information in DNS zone files—in addition to NIS+ tables or NIS maps.

Configure each client's `/etc/nsswitch.conf` and `/etc/resolv.conf` files properly. The `/etc/nsswitch.conf` file is the client's name service switch configuration file. The `/etc/resolv.conf` lists the IP addresses of the client's DNS servers, and is described in *Name Services Configuration Guide*.

13. Test the joint operation of NIS+ with DNS.

Verify that requests for information can pass between the namespaces without difficulty.

14. If operating NIS+ in parallel with NIS, test the transfer of information.

Use the `nispopulate` script to transfer information from NIS to NIS+. To transfer data from NIS+ to NIS, run `nisaddent -d` and then `ypmake`. (See the man pages for more information.) Use scripts to automate this process. Establish policies for keeping tables synchronized; in particular, the `hosts` and `passwd` tables. Test the tools used to maintain consistency between the NIS and NIS+ environments. Decide when to make the NIS+ tables the real source of information.

15. Test operation of NIS+ with both DNS and NIS.

Test all three namespaces together to make sure the added links do not create problems.

Phase III—Make the NIS+ Namespace Fully Operational

16. Convert clients to NIS+.

Convert clients one workgroup at a time, and convert all workgroups in a subnet before starting on those of another subnet. That way, when you convert all the clients in a subnet, you can eliminate the NIS service on that subnet. Run the verification script after converting each client to make sure that the conversion worked properly. That verification script should inform the user of the support structure that is in place to help with problems and how to report them. The actual steps required depend on the site.

Use the `nisclient` script to convert NIS clients to NIS+ clients. If you need to modify the clients' DNS configuration, you will have to write your own scripts to automate that process.

You can also save time if your site has a shared, mounted central directory similar to `/usr/local`. You could put the script in the central directory and, on the day of conversion, send email to clients asking them to run the script as superuser.

17. Monitor the status of the transition as clients are being converted.

Track progress against your plan and all serious complications not anticipated in the planning stages. Announce your status so that interested parties can track it.

18. Decommission NIS servers.

As all the clients on a subnet are converted to NIS+, decommission the NIS servers. If a particular subnet has some clients that require NIS service, use the NIS-compatibility feature of the NIS+ servers, but do not retain the NIS servers.

19. Evaluate NIS+ performance.

Once the implementation is complete, test to see that NIS+ is working correctly.

20. Optimize the NIS+ environment.

Based on the results of your performance evaluation, modify the NIS+ environment as needed. These improvements could be as simple as adding selected replicas in domains with high loads, or as involved as rearranging the storage of NIS+ information for a group of domains.

21. Clean up new domains.

If you did not change old domain names during the transition for the sake of simplicity, upgrade them now to the new NIS+ naming scheme. For instance, if you left some domains with geographic labels while you converted to an organizational hierarchy, you would change the geographic names to their organizational versions.

Phase IV—Upgrade NIS-Compatible Domains

22. Convert the last NIS clients to NIS+.

As soon as you can, eliminate the need for NIS-compatible NIS+ domains. Upgrading the last NIS clients to NIS+ will allow you to take advantage of NIS+ security features. You will not be able to eliminate the need for NIS-compatible NIS+ domains if you are running nonSolaris machines on your network.

23. Adjust your security configuration.

Once you have no more NIS clients, you can restart the NIS+ servers in standard mode and run `nischmod` on the NIS+ tables to change permission levels to eliminate the security hole caused by NIS compatibility. If you did not create credentials for NIS+ principals before, you must do that now. Restrict unauthenticated principals' access.

24. Establish miscellaneous evaluation and improvement programs.

Evaluate operational procedures to determine which ones can be improved; in particular, procedures used to recover from problems. Plan for new NIS+ releases and possible functional enhancements. Track the development of Solaris components that might require new NIS+ tables. Look for automated tools that enable you to perform NIS+ administration functions more efficiently. Finally, work with internal developers to help them take advantage of the NIS+ API.

Once you have completed the above steps, your transition to NIS+ is complete.

Index

A

- access rights, 38
 - default for NIS+ objects, 38
 - default for tables, 40
 - modifying, 42
 - NIS compatibility mode, 51
 - nobody class, 38
 - password table, 42
- automounter
 - converting maps to tables, 64
 - files, 25
 - tables, 25

C

- changing secure RPC credentials, 3
- chkey, 3, 33
- classes, 38
 - nobody, 38
 - world, 38
- clients
 - difference from a server, 20
 - DNS forwarding for, 51
 - duplicate login and machine names, 36
 - impact of transition on, 8
 - NIS compatibility mode capabilities, 3

- number per domain, 23
- configuration of servers, 4
- connecting tables, 29
- credentials, 35, 37

D

- data transfer commands, 50
- database
 - updating, 4
- designing NIS+ namespace, 11
- disk requirements, 23
- DNS
 - forwarding for NIS clients, 3, 4, 51
 - forwarding for NIS+ clients, 52
 - nsswitch.conf file, 52
 - server support, 49
 - zone files, 4, 52
- domain structure, 2
 - modifying, 11
- domains
 - across time zones, 17
 - choosing ownership, 62
 - crossing domains, 29
 - determining NIS-compatible, 48
 - determining number of servers, 20
 - number of clients per, 23
 - server support, 48

E

- /etc files, 27, 64
- /etc/.rootkey file, 34
- /etc/TIMEZONE file, 17

F

- forwarding DNS requests, 3

G

- groups, administrative, 37
 - access rights, 39
 - list of NIS+ commands, 62

H

- heterogeneous networks, 49

K

- keylogin, 33, 34
- keylogout, 33

L

- links between tables, 30

M

- mailhost alias, 30
- map and table synchronization, 3, 49
 - data transfer commands, 50
- maps
 - converting to tables, 64
 - verifying data, 64
- maps location, 6
- maps, *see also* tables
- master servers
 - definition, 4
 - requirements, 20
 - traffic loads, 20
- memory requirements, 23
- modifying access rights, 42
- modifying domains, 11

N

- name service switch, 4
 - defined, 27
 - sample configuration file, 28
- namespace
 - access rights, 38
 - choosing servers, 19
 - designing, 11
 - documenting, 65
 - duplicate names, 36, 63
 - parallel NIS and NIS+, 50
- network size, 2
- NIS
 - client capabilities, 3
 - commands supported under NIS+, 53
 - differences from NIS+, 1
 - domain structure, 2
 - list of equivalent NIS+ commands, 55
 - map and NIS+ table
 - synchronization, 3, 48
 - map conversion to tables, 64
 - map data transfer commands, 50
 - maps location, 6
 - NIS+ tables differences from NIS maps, 24
 - protocols, 45
- NIS compatibility mode, 3, 10, 45 to 57
 - access permissions, 51
 - client capabilities, 3
 - configuring servers, 48
 - DNS forwarding, 52
 - heterogeneous networks, 49
 - NIS+ security, 35
 - password table default access rights, 41
 - selecting domains, 48
 - servers not responding, 4
 - set up, 3
 - setting server names, 4
 - supported NIS commands, 53
 - supported NIS protocols, 45
 - transition process overview, 46
 - using custom tables, 29

-
- using links and paths, 29
 - NIS+
 - access rights, 38
 - additional information, 8
 - authorized servers, 4
 - changing secure RPC credentials, 3
 - credentials, 35
 - custom tables, 28
 - designing the namespace, 11
 - differences from NIS, 1
 - disk and memory requirements, 23
 - domain structure, 2
 - duplicate names, 36
 - editing tables, 6
 - for Solaris 1.x, 7
 - group commands, 62
 - groups, 37
 - heterogenous networks, 49
 - interoperability, 3
 - interoperation with `/etc` files, 27
 - linking tables, 29
 - links, 30
 - list of equivalent NIS commands, 55
 - NIS compatibility mode, 3, 10
 - NIS compatible domains, 48
 - NIS map and NIS+ table
 - synchronization, 3, 48, 49
 - NIS map data transfer commands, 50
 - nobody class, 38
 - password table access rights, 42
 - password table, parallel namespaces, 50
 - paths, 29
 - permissions, 35
 - principal, 6, 36
 - security, 6, 33
 - simplifying transition, 7
 - Solaris 2.x releases, 7
 - standard tables, 24
 - subdomains, 22
 - supported NIS commands, 53
 - swap space requirements, 23
 - table access rights, 39
 - table configurations, 24
 - table location, 6
 - tables, defined, 5
 - transition goals, 7
 - transition process overview, 6
 - updating database, 4
 - `nisaddent`, 50
 - `nispasswd`, 3, 34, 48
 - `nispopulate`, 50
 - `nistbladm`, 51
 - nobody access rights, 38, 41
 - `nsswitch.conf` file, 3
 - DNS forwarding, 52
- P**
- `passwd`, 51
 - password table
 - access rights, NIS-compatible domains, 41
 - access rights, standard NIS+ domains, 42
 - changing password commands, 51
 - encrypting `passwd` field, 41
 - parallel NIS and NIS+ namespaces, 50
 - passwords
 - changing user's, 34, 48
 - commands for changing, 51
 - updating, 3
 - paths, 29
 - permissions, 35
 - principal, 6
 - defined, 36
 - duplicate names, 36
- R**
- replicas
 - definition, 4
 - determining number required, 21
 - distribution, 22
 - `resolv.conf` file, 3
 - `.rootkey` file, 33
 - `rpc.nisd`, 3

S

security, 6

- access rights, 38
- changing passwords, 34
- choosing credentials, 35
- classes, 38
- credentials, 35, 37
- how it affects administrators, 34
- impact of implementation, 33
- keylogin, 33
- keylogout, 33
- NIS+ groups, 37
- nobody class, 38
- public keys, 35
- removing keys, 33
- selecting levels, 9, 33, 37

servers

- choosing a workstation for, 20
- configuration, 4
- configuring NIS compatible, 48
- connection to directory, 20
- determining number required, 20
- difference from a client, 20
- domain support, 48
- load distribution, 22
- master, definition, 4
- overview, 19
- replicas, defined, 4
- supporting multiple domains, 19

slaves, *see* replicas

Solaris 1.x NIS+, 7

subdomains, 22

subnets, 4, 22

swap space requirements, 23

T

tables

- defined, 5
- access rights, 39
- automounter, 25
- changing password commands, 51
- changing password information, 50
- configurations, 24

correspondence to NIS maps, 26, 49

custom tables, 28

default access rights for NIS-compatible domains, 41

default access rights for standard NIS+ domains, 42

differences from NIS maps, 24

editing, 6

file type, 5

files location, 6

linking tables, 29

links, 30

list of default access rights, 40

list of standard tables, 25

name syntax, 25

NIS map synchronization, 3, 49

NIS+ standard tables, 24

password table access rights, 41

paths, 29

populating, 64

time zones, 17

transaction logs, 23

transition process, 6

changing automounter map syntax, 25, 64

conversion model, 47

documenting NIS namespace, 65

goals, 7

impact on clients, 8

implementation, 67 to 72

prerequisites, 59 to 65

security impact, 35

simple version, 7

transition to Solaris 2.x, 7

V

/var/nis, 6

/var/yp, 6

Y

YP compatibility, *see* NIS compatibility

ypfiles, 49

yppasswd, 3, 48, 51

ypset, 4

