RSA ACE/Server 5.1
Scalability and Performance Guide

Contact Information

See our Web sites for regional Customer Support telephone and fax numbers.

RSA Security Inc.
www.rsasecurity.com

RSA Security Ireland Limited
www.rsasecurity.ie

Trademarks

License agreement

Note on encryption technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when exporting this product.

Distribution

Limit distribution of this document to trusted personnel.

RSA Security Notice

Protected by U.S. Patent #4,720,860, #4,885,778, #4,856,062, and other foreign patents.

The RC5™ Block Encryption Algorithm With Data-Dependent Rotations is protected by U.S. Patent #5,724,428 and #5,835,600.

First printing: January 2003
Second printing: October 2003

Contents

1 Introduction

This book presents a framework for planning the setup and use of RSA ACE/Server 5.1 in a network infrastructure. It is intended for information technology professionals responsible for an organization's network security.

The information provided here is designed to help you plan for an RSA ACE/Server 5.1 installation servicing a large number of users. It assumes a customer configuration with an RSA ACE/Server 5.1 Advanced license and multiple Replica Servers.

When planning installation and setup of RSA ACE/Server 5.1, the issues of scalability, network load, system administration and maintenance are important. This book discusses these issues, and includes the results of performance tests run on RSA ACE/Server 5.1.

This chapter describes the RSA ACE/Server 5.1 environment, and recommended configurations for supported server hardware. It also discusses scalability and performance issues, network loads, and administration and maintenance issues.

Chapter 2 (beginning on page 17) presents and discusses the results of a number of performance test profiles run on RSA ACE/Server 5.1, and the equipment used on these tests. It also discusses administration capacity, which is dependent on the size of the installation.

Frequently Asked Questions

Frequently asked questions about RSA ACE/Server 5.1 performance and scalability, which this document helps to answer, include:

- What is the maximum number of users that RSA ACE/Server can support? Per realm? With multiple realms?
- What is the maximum number of realms an installation can have?
- What maximum and sustained authentication rates are possible?
- How does cross-realm authentication affect the overall authentication rate?
- What network load is introduced by authentication? By cross-realm authentication? By remote authentication?
- How does replication affect performance? What are some of the other issues that affect performance?
- What are the hardware factors that affect performance?
- Can I increase the maximum number of remote administration sessions?
- Can I increase the maximum number of Quick Administration (Web-based) sessions?
- What are expectations regarding system maintenance?

RSA ACE/Server Environment

This section provides a brief overview of the RSA ACE/Server 5.1 environment. For details about system installation, administration, and use, refer to the *RSA ACE/Server 5.1 Installation Guide* (separate Windows and UNIX versions) and the *RSA ACE/Server 5.1 Administrator's Guide*.

Primary/Replica Servers

A typical RSA ACE/Server 5.1 installation runs on a specified array of computers that are part of a TCP/IP network. A computer designated as Primary Server and as many as 10 additional computers designated as Replica Servers make up a *realm*.

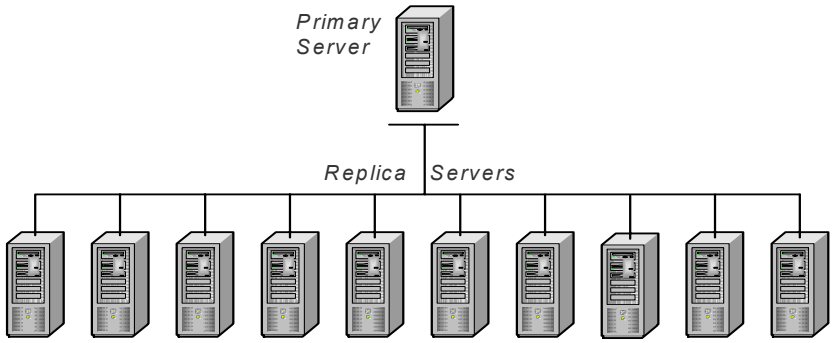

Note: The RSA ACE/Server Base license allows one Primary and one Replica server in one realm. If you want to deploy more than one Replica server or more than one Primary server (multiple realms), you must purchase an Advanced license. Contact your authorized RSA Security sales representative, or navigate to **www.rsasecurity.com/contact/upgrades.html** on the World Wide Web.

The Primary Server functions as the administration server. It replicates database changes to each Replica, and consolidates log messages from all of the Replicas into the Primary database. Replica Servers function as authenticating servers, with read-only database administration capabilities.

Note: RSA Security recommends that you enable authentication only on Replica Servers. This ensures that the Primary will have adequate cycles to perform replication and other administrative tasks. To disable authentication on the Primary, simply shut down the authentication service running on that machine. For information, see "Database Maintenance" for your platform in the *RSA ACE/Server 5.1 Administrator's Guide*.

For small to medium-sized organizations connected through a LAN, a single Primary/Replica array can provide enough bandwidth for authentication loads and to enable replication and administration. As your organization grows, you can add more Replicas to handle the increased authentication load and provide more flexibility in your network.

Realms

For larger organizations that have offices in distant remote sites, multiple instances of a Primary and its Replicas may be preferable.

Each Primary/Replica array in an installation is called a realm. With the RSA ACE/Server Advanced license, a single installation can include up to six realms, and you can configure a realm to authenticate and allow access to users from other realms. This is called *cross-realm authentication*.

The following figure illustrates the additional network traffic that is required to support cross-realm authentication. As shown in the figure, a user from Realm B is attempting to gain access to a protected server in Realm A. The protected server with RSA ACE/Server Agent Host (client) software asks for authentication of the user in Realm A.

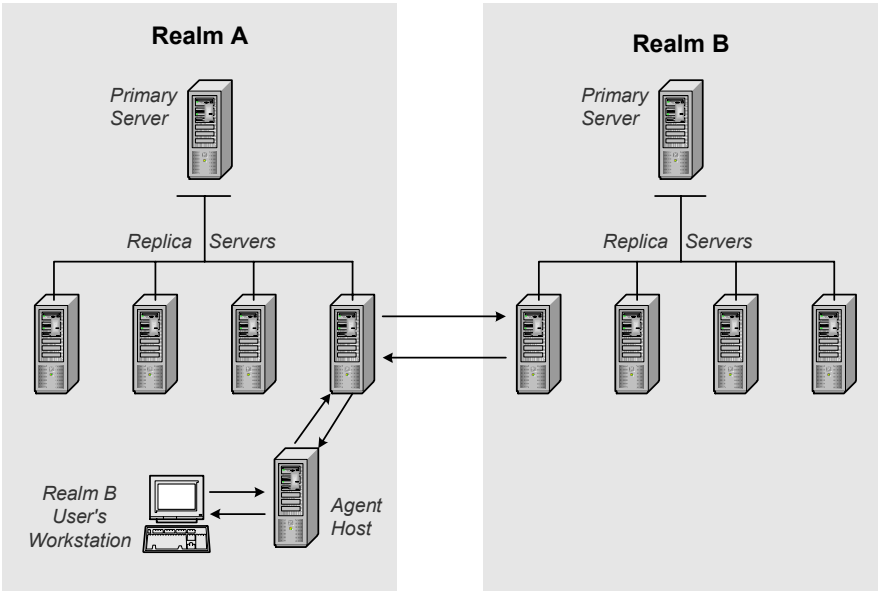

When the authenticating Replica in Realm A checks the database and does not find the user, it polls each realm until it finds the Server in Realm B that has a record for the visiting user. The authenticating Replica in Realm A then sends the authentication request to the Server in Realm B, which authenticates the user and passes this information back.

The authenticating Replica in Realm A then informs the Agent Host that the user is authenticated, and the user is admitted to the network.

If you have a very large number of users, or if you want to install Primary Servers at widely separated sites, you can set up multiple realms. This would enable you to distribute administration workload across the organization, and to improve network performance by enabling users to authenticate on nearby computers.

Important: Multiple realms apply only to customers with an Advanced license for RSA ACE/Server 5.1. An Advanced license allows up to six realms. If your site requires more than six realms, you can purchase multiple Advanced licenses, enabling RSA ACE/Server to support up to 20 realms. For a complete description of licensing options, refer to the *RSA ACE/Server 5.1 Administrator's Guide*.

Recommendations for Supported Server Hardware

RSA ACE/Server 5.1 runs on a variety of Windows- and UNIX-based hardware platforms. For complete information about system requirements, refer to the *RSA ACE/Server 5.1 Installation Guide* for your platform.

This section discusses additional considerations pertaining to server hardware. To maximize performance, RSA Security recommends that you designate dual-processor (or better) servers with sufficient memory and hard drive capacity for authentication purposes. RSA ACE/Server 5.1 detects the number of processors and starts multiple authentication "engines" for each processor.

The following table summarizes the *minimum* recommended specifications for dual-processor systems. Note that the recommended memory and storage type depend on the number of users in your RSA ACE/Server user database.

Platform	CPU/Clock Speed	User DB	Memory	Disk Drive
Windows	Dual Pentium III/500MHz or faster	10K	512MB	EIDE or SCSI-2
		100K	640MB	Ultra SCSI
		500K	1024MB	Ultra SCSI
Sun/Solaris	Ultra SPARC II/Dual 300MHz or faster	10K	256MB	EIDE or SCSI-2
		100K	300MB	Ultra SCSI
		500K	700MB	Ultra SCSI
HP/HP-UX	HP J2240/Dual 236 MHz PA-8200 CPU or faster	10K	256MB	SCSI-2
		100K	300MB	Ultra SCSI
		500K	700MB	Ultra SCSI
IBM/AIX	RS/6000/Dual 233 MHz processors or faster	10K	256MB	SCSI-2
		100K	300MB	Ultra SCSI
		500K	700MB	Ultra SCSI

In general, more and faster processors, faster hard drives, and more memory will result in better RSA ACE/Server 5.1 performance.

Note: RSA ACE/Server 5.1 requires that the Primary and Replicas in a realm all be running on the same platform and operating system. For both performance and security purposes, RSA Security strongly recommends that the computers designated as Primary and Replica servers be used exclusively for RSA ACE/Server purposes. You should avoid using these computers as file servers, firewalls, or for any other application.

Scalability and Performance Considerations

When planning your RSA ACE/Server 5.1 installation, there are a number of factors to consider regarding your user population and network traffic. The following subsections discuss these items.

User Factors

The number of Servers in your realm, and whether to establish multiple realms, depends on your user population. Factors to consider include:

- Number of users
- Locations of users
- Arrival times

Number of Users

What is the size of your current user population and the expected growth curve? RSA Security recommends a limit of *one million users per realm*. So, for example, if your current user population is 50,000 and you expect it to grow to 100,000, you might want to start with a single realm.

Note: The RSA ACE/Server user database (**sdserv.db**) is limited to two gigabytes in size. This is typically more than adequate to service the data associated with one million (or more) users. For more information, refer to "Database Maintenance" for your platform in the *RSA ACE/Server 5.1 Administrator's Guide*.

Your realm initially could include a Primary and four Replicas. As your user population grows, you can add one or more Replicas to accommodate the additional authentication load.

See "Peak Authentication Measurements" on page 18 for data about RSA ACE/Server 5.1 authentication rates and guidance on optimum Server configurations within a realm.

Authentication rates plateau at six Replicas in a realm. Beyond this, adding Replicas creates an increased replication load which, in turn, slightly affects authentication performance. You may still want to add Replicas, however, to assure *failover* at each remote site, to provide more optimal load balancing, or to service new locations.

For example, with two Replicas at each site, even if one Server goes down, users would still be able to authenticate locally, maintaining company-wide load balancing and preventing increased network latency. (Even if both Replicas at a location were to go down, users would still be able to authenticate to other Replicas in other locations on the WAN.)

Locations of Users

When planning your RSA ACE/Server 5.1 installation, the locations of your user population, and the time zones in which they work, are also important considerations.

In organizations where all users are on a local-area network (LAN), a single realm is usually sufficient. The Primary and Replica Servers can be located in the same room, distributed throughout a single building, or located in different buildings on a campus.

For organizations that have multiple, geographically-separated offices (multiple LANs) connected by a wide-area network (WAN), the Primary and Replica servers can be located anywhere within the organization.

Similarly, for Web-based deployments (for example, an ISP), where users are widely dispersed, you could locate Replicas at strategic data centers around the world.

In some cases, rather than deploying more Replicas, it may actually be preferable to establish multiple realms. With an Advanced license, RSA Security allows you to establish up to six realms.

Note: For larger authentication loads, it is possible to purchase “stacked” Advanced licenses to allow *up to 20 realms*. Contact your RSA Security sales representative for more information.

In choosing the physical location of Primary and Replica servers in a realm, you should consider a number of factors: performance, maintenance, troubleshooting and security.

While it may be easier to upgrade or troubleshoot servers by having them rack-mounted in the same room, overall authentication performance might suffer, particularly if many users were logging on from different remote offices on a WAN.

In contrast, you might want to improve performance by locating Replicas physically closer to the users who will be authenticating through them. For example, in a corporation that has multiple remote sites, one Replica could be located in the corporate headquarters in New York, another in the manufacturing facility in Mexico, and a third in the research laboratory in California.

To maximize performance, you might decide to set up individual realms for each satellite office so that local users can authenticate locally. However, consider that multiple realms can add more overhead to installation, administration and troubleshooting. In addition, if you expect a great number of cross-realm authentications, performance will be impacted. See "Peak Authentication Measurements" on page 18 for information about authentication performance within the same realm and "Cross-Realm Authentication" on page 25 for information about cross-realm authentication performance.

Important: Regardless of the physical location of Primary and Replica Servers in your organization, it is critical that they be secured in a locked room accessible only by authorized personnel.

Peak Arrival Times

In addition to number and locations of users, another factor is the peak times at which users arrive at work and log in to the network.

For example, if your organization ran three shifts at each office location, you might have three to six peak periods each work day during which users are authenticating to your network.

In a single shift, the peaks and valleys of the authentication load might look like the following graph.

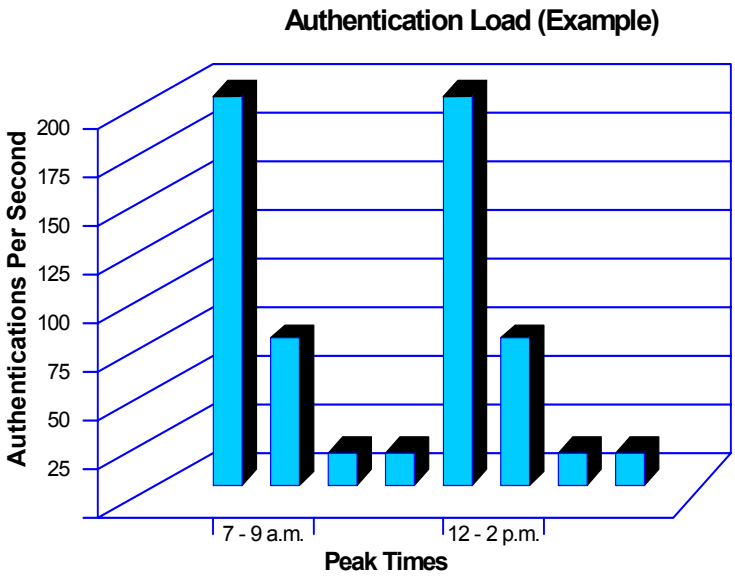

When planning an RSA ACE/Server 5.1 installation, it is important to allow for the peak authentication rates that your organization is likely to experience.

The questions to answer are: When are the peak periods during the day when the majority of users are attempting to log on to the network? What is the maximum number of users that might be logging in during those peak periods?

You should also allow ample room for expansion as your user population increases in size.

Network Traffic

RSA ACE/Server 5.1 authentication and replication processes involve the communication of multiple data packets across a local- or wide-area network.

The following subsections discuss the network traffic (number and size of packets) that you can expect from RSA ACE/Server 5.1 processes.

Authentications

During a single authentication, data packets are sent back and forth between the Agent and the Server until a user is authenticated or entry to the network is disallowed. The actual time to complete this communications loop depends on the speed of your network.

Each authentication involves 2,032 bytes of data (in four 508-byte packets) sent between the Agent and the Server. To begin an authentication, the Agent sends a "Name Lock" packet. The Server replies with a "Lock Response" packet. The Agent then sends an "Authentication Request" packet, which in turn generates an "Authentication Response" packet.

This per-authentication load will grow approximately another 600 bytes for every Replica in the realm. Therefore, on a fully-configured system, each authentication could generate up to 8000 bytes of network load. So, for example, a realm with a Primary and 10 Replica Servers experiencing 100 authentications per second (APS) would consume 800,000 bytes per second of network capacity:

Example Maximum Authentication Load (Primary/10 Replicas)	100 APS x 8000 bytes = 800,000 bytes per second

Replication

In RSA ACE/Server 5.1, replication is the process by which user databases on all servers in the realm are synchronized.

The Primary Server runs a separate instance of the replication service for each Replica Server in a realm. Each Replica Server runs a single instance of the replication service. The replication service enables the Primary and Replica to communicate and exchange information about changes to the database (called *delta records*) on a regular basis. Each exchange of delta records that occurs between the Primary and a Replica is called a *replication pass*.

After they start, the Primary and the Replicas exchange delta records at a specified frequency called the *replication interval*, a setting that you can specify for your system. The default replication interval is 100 seconds.

Each replication pass generates a minimum of 12 packets per Replica, each 16 bytes in size. So, in a realm with 10 Replicas, a replication pass would require approximately 1,920 bytes of network capacity. Given the default replication interval of 100 seconds (there are 86,400 seconds in a day), this would amount to approximately 1.7 megabytes of network usage per day.

Example Replication Network Load (Primary/10 Replicas)	1920 bytes x 864 replication passes = 1,658,880 bytes per day

Load Balancing with RSA ACE/Agent 5.0 Software

RSA ACE/Agent 5.0 software offers a *load balancing* feature that automatically distributes the authentication request load, helping to optimize authentication performance on your network.

If you deploy RSA ACE/Server 5.1 and RSA ACE/Agent 5.0 versions, you can take advantage of this load balancing capability.

Note: Although previous versions of RSA ACE/Agent software work with RSA ACE/Server 5.1, they do not take advantage of automatic load balancing. For a list of RSA ACE/Agent 5.0 software, refer to **http://www.rsasecurity.com/products/securid/specs/agentsupport.html**.

For load balancing, RSA ACE/Agent 5.0 software polls each Server in the realm and, based on the response time of each Server, determines a priority list. The Server with the fastest response will receive authentication requests from the Agent more frequently than other Servers until the Agent software sends another time request.

As an alternative to automatic load balancing, administrators have the option of balancing the load manually by specifying exactly which Servers each Agent Host should use to process requests.

For complete information about load balancing capabilities, refer to the *RSA ACE/Server 5.1 Administrator's Guide*.

Network Latency

Another item to consider when planning your RSA ACE/Server 5.1 installation is minimizing network latency. Network latency is the time that it takes for a data packet to travel through the network to its destination and a return packet to arrive.

Consider the example of a corporation with a headquarters and three remote sites. By placing a replica at each site, there would be less network latency, and better authentication performance. This is because users at each site would be authenticating to the nearest Replica.

Although the same setup would increase latency for replication, the replication load is significantly less than the authentication load, and replication can happen over a longer period of time. See "Database Replication" on page 28 for more information about replication load.

Another advantage of this approach is, if one of the Replicas goes down, users at that remote site would still be able to authenticate to one of the other Replicas in the realm.

Remote authentication

Another consideration when planning network capacity is the impact of remote authentication. People at branch offices, telecommuters, and people who are traveling may need to access your corporation's network. If a significant number of your users access your network remotely, either through dial-up connections or high-speed cable or DSL modems, this may add noticeable overhead to network throughput.

Remote access typically requires a RAS (Remote Access Server), a computer and associated software that is set up to handle remote users. RAS configurations usually include or are associated with a firewall server to ensure security and a router that can forward the remote access request to another part of the corporate network.

RAS devices are usually a component of a VPN (Virtual Private Network), which adds more overhead to your network. A VPN involves encrypting data before sending it through the public network, and decrypting the data at the receiving end. Some VPNs include an additional level of security that involves encrypting not only the data but also the originating and receiving network addresses, adding still more network overhead.

Administration Considerations

When planning an RSA ACE/Server 5.1 installation, you should also consider administration and database maintenance issues.

Remote Administration

With the RSA ACE/Server Remote Administration tool, authorized administrators can perform all necessary database administration tasks from any workstation on the LAN or WAN. This includes adding and deleting users, editing user information, assigning tokens, editing system parameters and system extension data, and so on.

For a complete list of functions that can be performed remotely, refer to the *RSA ACE/Server 5.1 Administrator's Guide*.

Remote Administration runs from any Windows-based PC to administer Windows or UNIX databases in a local realm and in registered remote realms. There are limitations to the number of remote administration sessions that can be open at a given time within a single realm. In addition, there is network traffic associated with Remote Administration. This should be taken into account when considering how much of the administration load will be performed remotely. See "Remote and Web-based Administration Sessions" on page 31 for more information, including the number of sessions RSA ACE/Server default settings can support, and also how you might adjust those settings to increase the number of sessions your system will support.

Web-based Administration (Quick Admin)

The Web-based administration tool, Quick Admin, enables a system or help desk administrator to use a Web browser to view and modify user, token, and extension record data in the Primary RSA ACE/Server database.

Using Quick Admin, administrators can perform such common tasks as assigning a temporary password to a user, marking a token as lost, resetting a token, generating a token report, or editing user information.

There are limitations to the number of Web-based administration sessions that can be open at a given time in a single realm.

In addition, there is network traffic associated with Quick Admin. This should be taken into account when considering how much of the administration load will be performed via the Web. See "Remote and Web-based Administration Sessions" on page 31 for more information, including the number of sessions RSA ACE/Server default settings can support, and also how you might adjust those settings to increase the number of sessions your system can handle.

LDAP Import and Synchronization

In RSA ACE/Server 5.1, enhanced LDAP import and synchronization capabilities supporting Microsoft Active Directory, Netscape iPlanet, and Novell Netware are included. If your organization uses one of these LDAP services to manage your user data, you can import this user data to RSA ACE/Server's database, and schedule automatic updates to keep the two databases in sync.

Although the LDAP import and update functions will use network bandwidth, you can use the built-in scheduler to run these functions at off-peak times. Performance test data for LDAP import and compare operations can be found on page 29.

Log Maintenance

RSA ACE/Server 5.1 stores a variety of data in a commercial database developed by Progress Software Corporation and integrated into the RSA ACE/Server software.

One type of data, called log data, is an audit trail of all authentication and administrative activity. Left unattended, this log database (**sdlog.db**) would continue to grow until it exceeded its allowable size limit, or the Server ran out of disk space.

Important: When **sdlog.db** reaches 2,147,155,968 bytes in size on the Primary Server, all the replication engines shut down, and sdadmin, log monitor, replication management, and any other process that needs to access the log database, become unusable. In this event, users would be denied network access because RSA ACE/Server does not authenticate a user unless it can log the event.

Consequently, it is important to maintain the log database. One way to prevent log overflow is to periodically purge older log records from the database (for example, deleting all the log records created before a certain date). You can perform log maintenance manually, or by using the RSA ACE/Server Automated Log Maintenance (ALM) tool.

When log records are deleted, disk space is made available for new log records, but **not** automatically freed for other uses. RSA Security provides a database compression utility that enables you to reclaim disk space used by the RSA ACE/Server databases. It is always a good practice to compress the database after deleting a large number of log records. See "Log Maintenance" on page 34 for log maintenance test results.

Note: For log maintenance, you can also use the *log filtering* tool in RSA ACE/Server 5.1. Log filtering provides a way to select the log messages that go into the RSA ACE/Server log database. By filtering out certain messages, you can slow the growth of the log database and increase replication, authentication, and administration performance.

Disaster Recovery

The Primary/Replica model of RSA ACE/Server 5.1 enables failover potential and quick recovery in a variety of emergency situations. By employing a Primary and two or more Replicas in your configuration, you can quickly react to a Server machine going down.

Should your Primary Server machine fail, the Replicas will continue authenticating users. If you expect your Primary to be down for a day or more, RSA ACE/Server 5.1 provides a *Nominate* feature that enables you to select and configure a Replica Server to become the Primary. See the *RSA ACE/Server 5.1 Administrator's Guide* for complete information.

For the case in which an authenticating Replica goes down, RSA ACE/Server 5.1 provides a disaster-recovery mechanism, called *DBPush*. After the Replica hardware is brought back online, you can use the DBPush tool to restore necessary data to the Replica through the network connection. See "Database Push" on page 29 for DBPush test results.

Note: For organizations in which high availability of network resources is a requirement, RSA ACE/Server 5.1 supports the Hewlett Packard ServiceGuard and Veritas Cluster Server (on Solaris 9) high availability hardware systems.

2 Performance Tests

This chapter examines the results of performance tests developed and run by RSA Security on RSA ACE/Server 5.1. An analysis of these test results should help network administrators plan a more optimal installation and deployment of RSA ACE/Server 5.1 in their own network environments.

Test Systems and Network Environment

This section describes the test environments that RSA Security used for RSA ACE/Server 5.1 performance tests.

Note: The test systems are typical configurations, and are not intended as recommendations by RSA Security. For best performance, deploy the fastest, most powerful computers possible. RSA ACE/Server 5.1 is a multi-threaded application that takes advantage of multi-processor systems. Dual- and quad-processor systems will help optimize performance, as will faster networks.

The tests described here were conducted on both Intel/Windows and Sun/Solaris server equipment connected to Fast Ethernet (100BASE-T) local area networks.

Each configuration included a Primary Server, with up to 10 Replica Servers, and various Agent Hosts and remote administration machines.

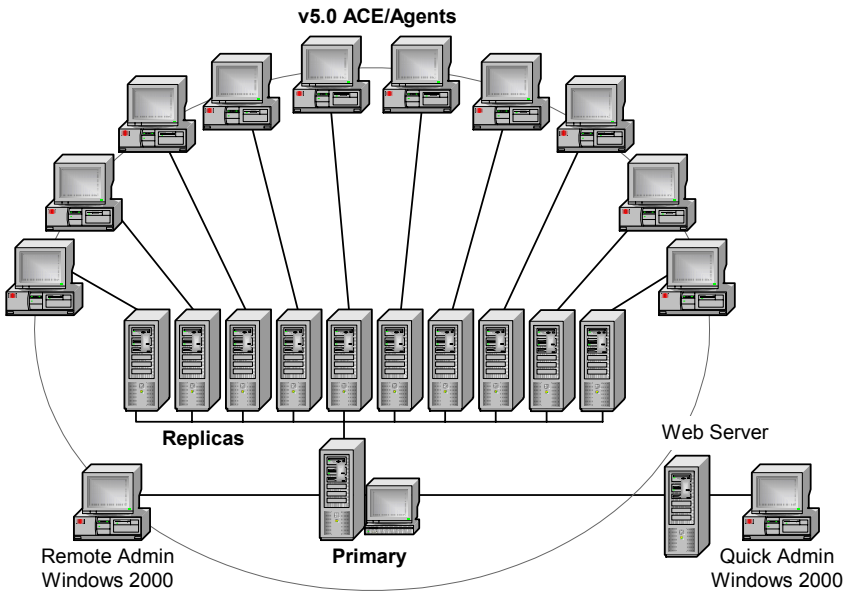

The following table describes the specific Server hardware used in these performance tests. Where appropriate, additional information about hardware, specific to a performance test, appears in the test-related sections later in this chapter.

Platform	Server Hardware	Operating System
Windows	Dual-processor Pentium III (933MHz), 1024MB RAM	Windows 2000 Advanced Server
Sun	Sun Fire V100 workstations, single UltraSPARC IIe 500MHz CPU, 1024MB RAM	Solaris 8

Peak Authentication Measurements

Peak authentication is the highest number of authentications per second (APS) that RSA ACE/Server 5.1 can achieve for short, concentrated periods. This section describes and analyzes peak authentication performance in single-realm Windows 2000 and Solaris 8 environments, as well as through the RSA RADIUS Server.

Data was gathered from tests conducted both in isolated and "real world" network configurations.

The "real world" tests, while not as comprehensive as the isolated network tests, are useful in gaining some insight to performance when administrative loads are added during authentication. These administrative loads were created by starting regular, periodic Quick Admin, Remote Administration, and Administration Toolkit (ATK) sessions while authentication proceeded.

For data and analysis covering cross-realm authentication, see "Cross-Realm Authentication" on page 25.

Peak Authentication on the Windows 2000 Test System

The following two graphs show the measured peak authentication rates for the Windows 2000 test configuration.

The graphs show peak authentication rates in both an isolated network setup, and under simulated “real world” conditions. The tests used version 5.0 Agents exclusively.

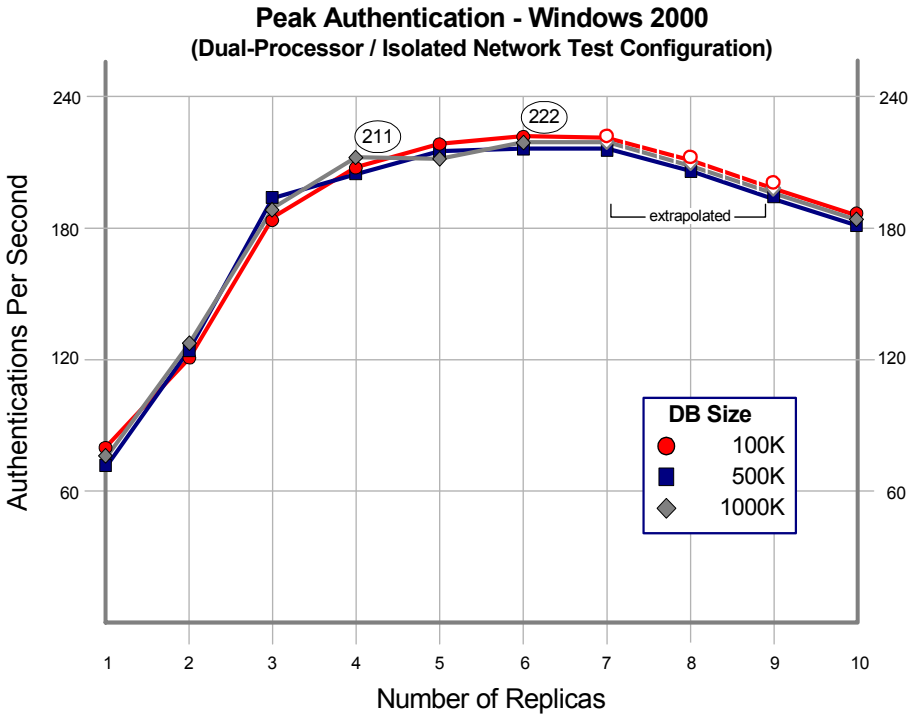

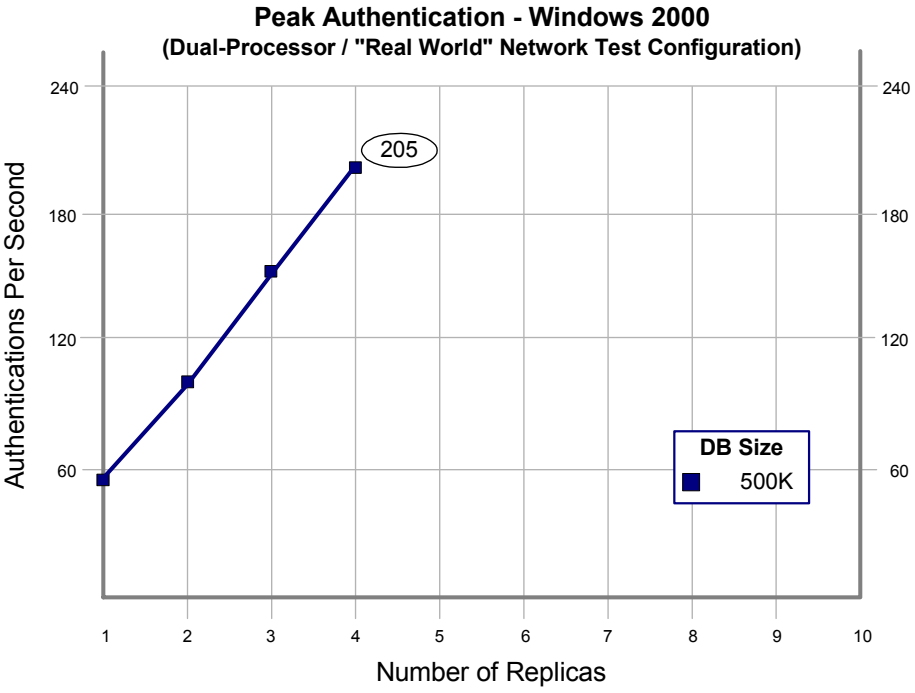

Peak Authentication on the Sun Solaris Test System

The following two graphs show the measured peak authentication rates for the Sun Solaris test configuration. The graphs show peak authentication rates in both an isolated network setup, and under simulated “real world” conditions. The tests used version 5.0 Agents exclusively.

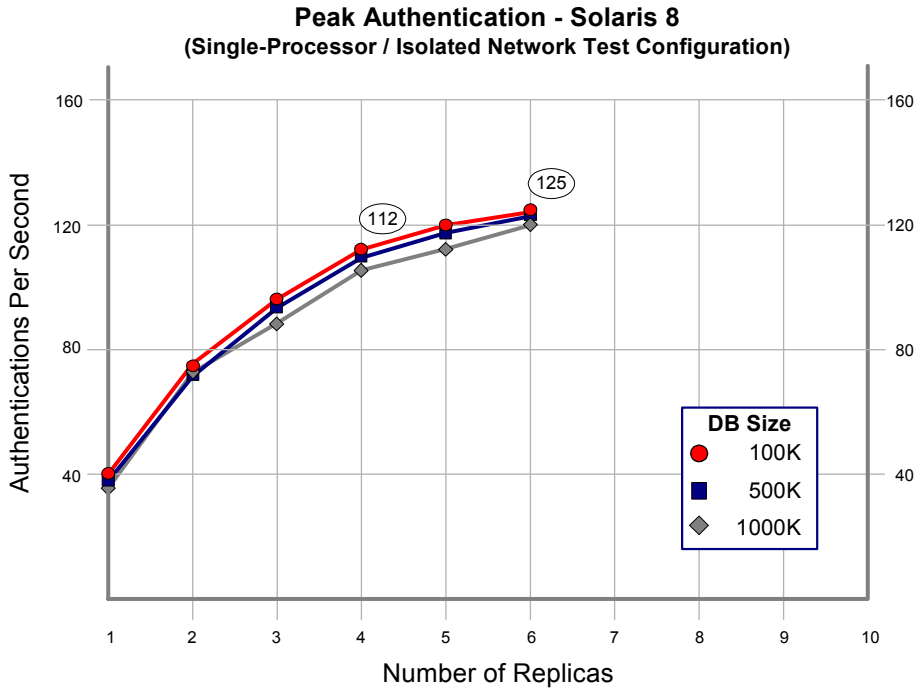

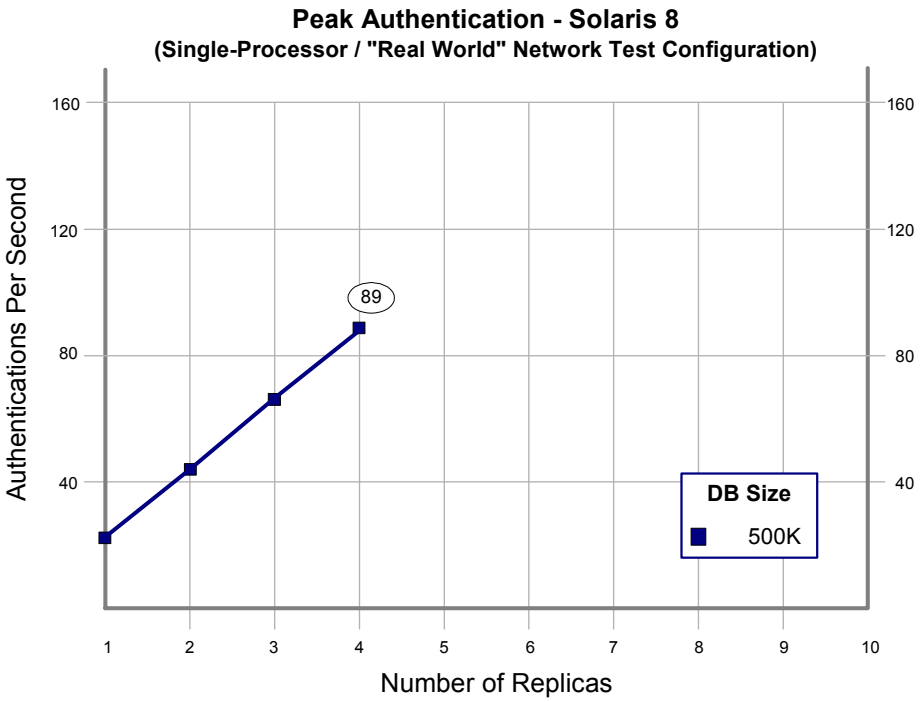

Peak Remote Authentication Rate through RADIUS

The following two graphs show the measured peak remote authentication rate through the RSA Security RADIUS server, the first in an isolated network environment, the second under simulated "real world" conditions. In these tests, the RADIUS server was set up as an agent on the same machine as the RSA ACE/Server.

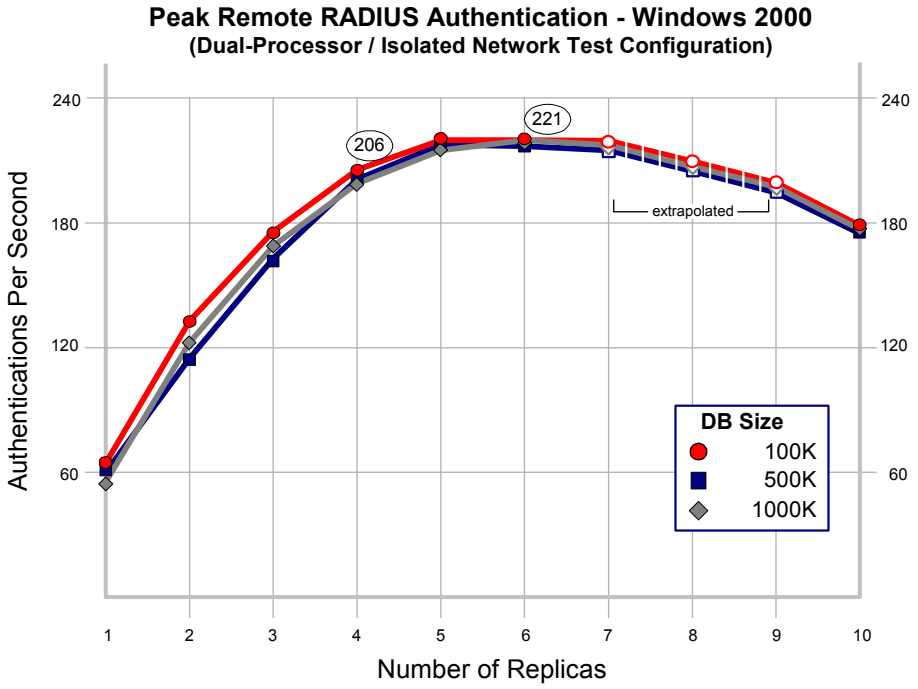

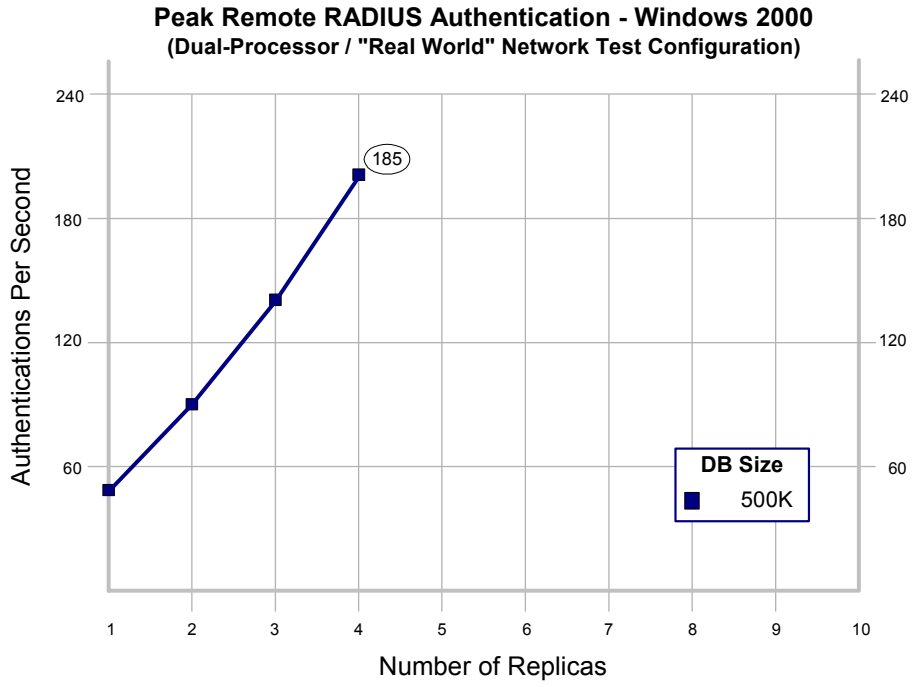

Analysis of Peak Authentication Meaurements

It is important to understand that RSA ACE/Server 5.1 is not designed to achieve peak authentication rates indefinitely, as replication and logging would be compromised. Use the peak authentication rates published here as guidelines for your organization's peak requirements. Take into account your employee peak arrival times, time zones, and so on. The peak authentication tests indicate the following:

- Peak performance was **222 authentications per second** on a representative dual-processor configuration. This compares to 175 APS for RSA ACE/Server 5.0, a 26 percent improvement, and 35 APS for version 4.1 of RSA ACE/Server, a 534 percent improvement.
- The highest authentication rates were achieved with a configuration of **one Primary and six Replicas**.
- Remote authentications through the RADIUS server were only marginally lower than those measured on a local area network. RADIUS performance was similar to local authentication, peaking at **221 APS** with five or six Replicas.
- The "real world" tests, which added database administration loads during authentication, showed some reduction in APS compared to the isolated network tests. The reduction was less pronounced on the dual-processor Windows 2000 configuration (**205 APS real-world versus 207 APS isolated** with four Replicas).
- On the single-processor Sun Solaris 8 configuration, the administrative load applied during the "real world" tests had a more significant impact (**89 APS real-world versus 110 APS isolated** with four Replicas). This is attributed to the single-processor machines not taking full advantage of RSA ACE/Server's multi-threaded architecture.
- For RADIUS authentications, which were conducted on the dual-processor Windows 2000 configuration, the "real world" differences (**185 APS real-world versus 200 APS isolated** with four Replicas), were attributed to the additional load created by the RADIUS agents, which were installed on the Replicas.
- Because of increased replication and logging traffic, more than six Replicas per realm begins to show a small decrease in performance for all the database sizes used in the tests. While adding more than six Replicas in a realm does not improve peak authentication rates, it does provide you with increased failover, as well as more options for locating Server hardware in your organization.With automatic load balancing enabled, additional strategically-located Replicas can help to minimize latency (delay) on your network.
- Your actual peak authentication rate depends on several factors: your hardware configuration, the types of RSA ACE/Agents you are using, whether your users are authenticating locally or remotely, whether you have established multiple realms and the amount of cross-realm authentication in your organization.
- Database size has a statistically insignificant effect on authentication rate. Authentication rates for 100K, 500K, and 1000K databases were within percentage points of each other at each Replica count in the test configuration.
- If you are upgrading from version 4.1, even without adding server hardware, you can still expect a **128 percent improvement** in peak authentication performance per realm.

Sustained Authentication Rate

This section describes a test measuring **sustained authentication rate**, which is the authentication rate that RSA ACE/Server 5.1 can perform on an ongoing basis while maintaining acceptable replication and administrative loads.

Why Replication Is Important

Since multiple Servers running RSA ACE/Server software can authenticate users simultaneously, it is essential that all of these Servers' databases are frequently updated so they are identical and current. Among other things, this prevents an attacker from stealing a legitimate user's identity and attempting to gain access from another computer.

During replication, administrative changes made to the Primary Server's database (add a new user, delete a user, and so on) propagate to all of the Replica Servers' databases. Likewise, changes made to one Replica Server's database (a user's last authentication date, a change in a token's status, a new agent registration, and so on) are made to the database of the Primary Server and, in turn, the databases of all other Replica Servers.

As this information is propagated (*replicated*), if a collision — more than one change to the same database record — occurs, it is automatically detected and resolved. To maintain the ability to detect and resolve collisions, it is important that an RSA ACE/Server installation have the capacity to fully reconcile the Server databases on a regular basis, while still handling the peaks and valleys of anticipated authentication loads. The faster full reconciliation can take place, the more secure an RSA ACE/Server installation will be.

Sustained Authentication Test Results

Several tests were run over 24-hour periods to find a balance between reasonable reconciliation rates and acceptable sustained authentication numbers.

The tests were run on the Windows 2000 platform, employing a Primary and six Replica Servers. Six additional machines were set up as Agent Hosts running RSA ACE/Agent version 5.0 software. (Refer to "Test Systems and Network Environment" on page 17 for the hardware specifications of these machines.) The user database contained 500,000 users.

Only the Replicas were set up to authenticate users. One Agent Host was assigned to each Replica, and generated a steady stream of authentication requests over a 24-hour period. Load balancing was disabled.

During the 24-hour period, each Agent Host ran eight concurrent daemons, which each continuously sent an authentication request to the target Replica, waited for a response, then sent another request.

The following table shows the results of three sustained authentication tests.

	Reconciliation Time (minutes)	APS per Replica	Total APS
Test 1	81	3.5	21
Test 2	107	6	36
Test 3	147	8.5	51

Analysis

With a configuration of a Primary and six Replicas, RSA ACE/Server 5.1 was able to sustain **21 to 51 authentications per second** (APS). With the default replication interval of 100 seconds, full reconciliation between the Primary and all Replicas occurred between 81 and 147 minutes.

Based on this, an optimally-configured RSA ACE/Server 5.1 system would be capable of **75,000 to 183,000 authentications per hour** per realm, with full reconciliation of the Server databases happening at reasonable intervals.

Actual sustained authentication performance in an organization can be affected by other network factors, such as network speed, latency, and load caused by other traffic.

Also, it is important to note that an RSA ACE/Server installation does not have the type of continuous authentication load reflected in these tests. When there are fluctuations in the authentication load, the Server uses any “idle time” to allow the replication process to catch up on the distribution of database changes.

Example

Consider, for example, a corporation with 75,000 employees based at three locations in three time zones. The corporation has a wide area network connecting multiple LANs through leased high-speed telecommunications lines.

The employees are evenly distributed within the three time zones (25,000 per site). Also, the corporation operates in two eight-hour shifts at all of its locations (12,500 employees per shift per location).

For this example, assume that each employee requires authentication to enter the corporate network twice a day, when arriving to work at the beginning of the shift, and in a one-hour window after meal time. So, in each time zone, there are four peak periods, each lasting approximately an hour, during which 25,000 employees require authentication.

For this example, assuming a single realm, the distribution of servers might look like this:

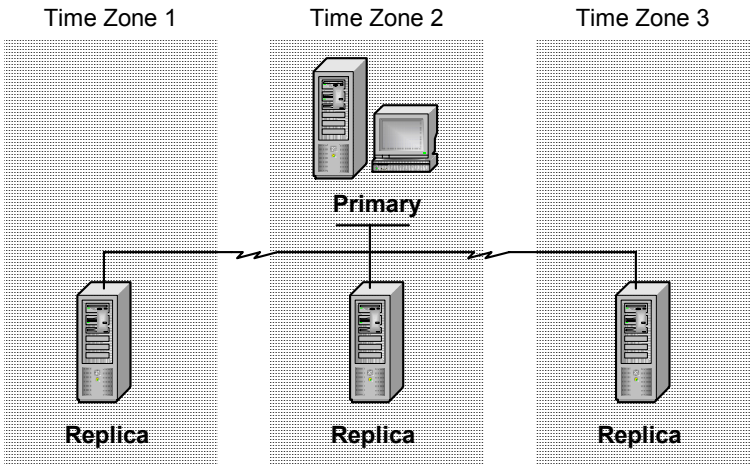

Each site would have at least one authenticating Server. Note that this example shows three authenticating Replicas, half the number used in the RSA Security test. With this type of configuration, RSA ACE/Server 5.1 could sustain approximately 37,000 to 61,000 authentications per hour indefinitely, and even higher rates for short durations during peak authentication periods. This would be more than adequate for the organization in this example. Even if one Replica went down for some reason, the other Replicas in the realm would provide failover protection.

Cross-Realm Authentication

If you have office locations that are widely dispersed geographically, you can establish multiple realms, each site having its own RSA ACE/Server 5.1 Primary/Replica Server array. You then can configure each realm to enable users from other realms to gain access throughout the organization. This is known as *cross-realm authentication*. (Refer to "Realms" on page 7 for an introduction to cross-realm authentication.)

Although a single cross-realm authentication produces more network traffic than an authentication within a single realm, the number of cross-realm authentications is typically a small percentage of an organization's total authentication traffic. Additionally, having multiple realms typically reduces long-distance replication and logging traffic.

Cross-Realm Authentication Test Results

In cross-realm authentication, all RSA ACE/Agents, including version 5.0 Agents, use pre-5.0 protocol to communicate with the home realm. This means that each Agent will only authenticate to one Server in the home realm. For version 4.4 (and earlier legacy) Agents, this is the Server designated as the acting Master Server. For 5.0 Agents, the designated Server is referred to as the Preferred Server.

For that reason, and because the transmission of additional data packets is necessary, each cross-realm authentication is slower than an authentication within a realm. Refer to "Network Traffic" on page 12 for more information about network data packets transmitted by various RSA ACE/Server 5.1 processes.

The following graph shows the results of the cross-realm authentication tests. Both the home realm and the remote realm were established on the Windows 2000 platform. The remote realm was set up with a Primary and one Replica server, and had a database of 80K users. The home realm had from one to four Replicas, and had a database of 100K users. Version 5.0 Agents were used exclusively.

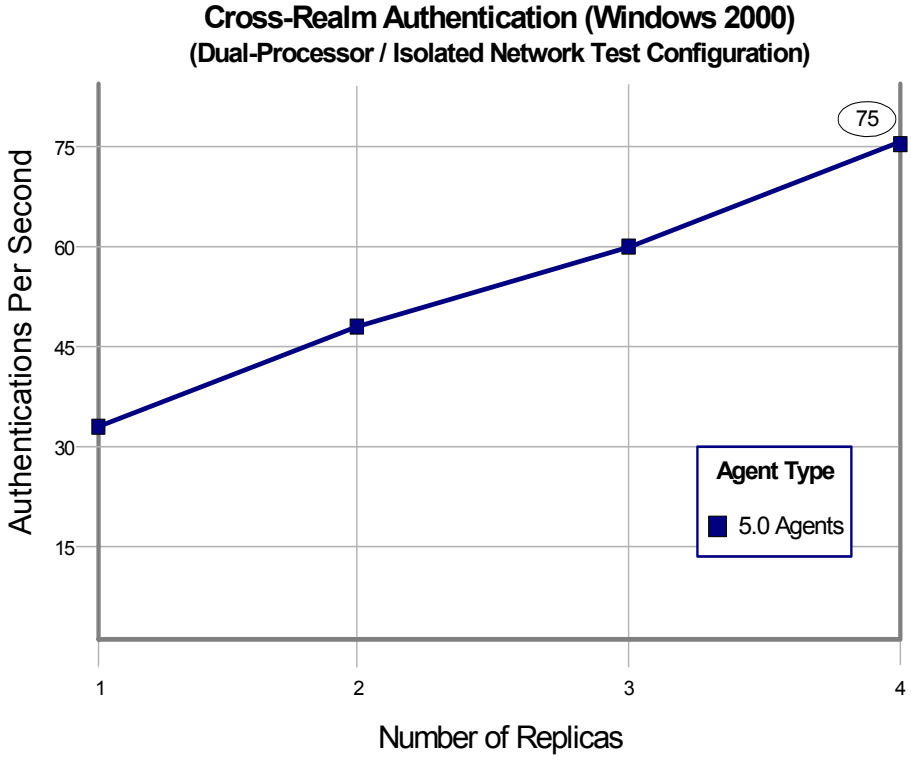

Analysis

The cross-realm authentication tests revealed the following:

- Using 5.0 Agents, peak cross-realm authentication was approximately 35 APS when the home realm was configured with one Replica and one Agent. With each additional Replica (and Agent), the cross-realm APS grew approximately 15 APS, peaking at **75 APS** when the home realm was configured with four Replicas and four Agents.

  Note: One of the design goals of RSA ACE/Server 5.x architecture was to enable organizations using older versions to upgrade and consolidate multiple Master/Slave realms into a single 5.x realm. This results in dramatic gains in authentication rates, and provides other advantages.

- Very large organizations, and/or those with multiple, widely dispersed offices, can establish multiple 5.1 realms. This provides improved in-realm authentication rates (over version 4.1), as well as equivalent cross-realm authentication performance.

Single Realm versus Multiple Realms

RSA Security recommends that organizations consolidate multiple realms into one realm. However, sometimes it makes sense to establish more than one realm. The following table summarizes the key points when choosing a single-realm or a multiple-realm environment.

Choose	Reasons	Example
Single Realm	• Expected peak rate is less than 222 APS. • Total number of users is less than one million. • Fewer than five remote offices are on the WAN. • Centralized administration is preferred. • Faster authentication is preferred when all users must be on the same network.	A company with offices in Boston, New York, and Chicago connected through a WAN serviced by leased high-speed telecommunications lines. (Configuration is a Primary and five Replicas. Each site has the recommended two Servers to maintain failover.)
Multiple Realms	• Expected peak rate is greater than 222 APS. • Total number of users is greater than one million. • Distributed administration is preferred. • More efficient network use (less latency) when amount of cross-realm authentication is low (less than 10 percent).	A large multinational manufacturing organization with offices in San Francisco, New York, London, Paris, Hong Kong, and Melbourne. (Each site is set up and administered as a separate realm with a Primary and four Replicas.)

Database Replication

In the RSA ACE/Server 5.1 Primary/Replica model, the Primary functions as the administration Server, and periodically *replicates* database changes (delta records) to each Replica.

As part of administering RSA ACE/Server, you can specify a *replication interval*, the frequency at which a replication pass is initiated. The default is 60 seconds.

Most changes to the Primary database result from administrator actions — for example, adding a user to the database and assigning an RSA SecurID token to that user. Adding one user would be a small change that would be quickly replicated to the other Servers. A larger change, for example, importing 100 new users from an LDAP database, would take longer to replicate. Replication is a background process and can also be slowed by higher-priority processes (for example, authentication).

Server Hardware

Operating System	CPU	RAM
Windows 2000 Advanced Server	Single-processor 800 MHz Pentium IIIs	512MB

Replication Test Results

The statistics provided in this section gauge the impact of the replication process in a single-realm configuration with different database sizes. The reported time is how long it took to complete the replication pass once it started.

Configuration: Primary and Single Replica

Task	User DB	Time
Replicate 100 New Users to DB	50,000	19 seconds
	100,000	73 seconds

Analysis

After initial installation and deployment of RSA ACE/Server 5.1, the replication process typically places a very small additional load on system resources. This varies to some degree based on the replication interval (the default is 60 seconds). A shorter replication interval would increase the load. A longer interval would reduce it.

Replication rate is also affected by peaks in authentication requests. As replication runs at a lower priority than authentication, it is temporarily suspended, or slowed down, as the authentication load increases.

Replication can also be affected by network speed and the geographical distance between the Primary and Replica Servers. The more distance, and the slower the network connection, the slower replication will be.

Database Push

In RSA ACE/Server 5.1, *DBpush* (database push) is a function that enables the administrator to copy the entire latest database files from the Primary to one or more Replicas over a LAN or WAN. The administrator can use DBpush during installation, or as part of a failure-recovery process after a Replica Server goes down.

Server Hardware

Operating System	CPU	RAM
Windows 2000 Advanced Server	• Primary: Single-processor 500 MHz Pentium III • Replica: Single-processor 400 MHz Pentium III	128MB

Database Push Test Results

The statistics provided in the following table show the length of time required for the database push process in a Primary/One Replica configuration with different database sizes.

Configuration: Primary/One Replica	
Push 50K User Database	4 Minutes, 57 Seconds
Push 100K User Database	21 Minutes, 4 Seconds
Push 500K User Database	52 Minutes, 56 Seconds

Analysis

Restoring the entire user database to a Replica is a straightforward process. Larger databases take longer to push. DBpush should be done during off-peak hours.

LDAP Import and Synchronization

RSA ACE/Server 5.1 provides tools to import an LDAP database into the Server database, and to schedule automatic updates to keep the databases in agreement. With these tools, your LDAP database can be used to populate and maintain the RSA ACE/Server database.

This section provides statistics on the initial import of users into the RSA ACE/Server database, as well as compare/update processes.

Server Hardware

Operating Systems	CPU	RAM
RSA ACE/Server: Windows 2000 Advanced Server	Single-processor 800 MHz Pentium III	512MB
LDAP Server (Active Directory SP2): Windows 2000 Advanced Server	Single-processor 800 MHz Pentium III	256MB

LDAP Import and Synchronization Results

The following table provides the results of the LDAP import and synchronization tests.

Configuration: Primary and Single Replica		
Task	**User DB**	**Time**
Import 10K new users from LDAP	100K	5 minutes, 10 seconds
Import 50K new users from LDAP	100K	34 minutes, 22 seconds
Update DB with 5K (5%) new users	100K	36 seconds
Update DB with 10K (10%) new users	100K	3 minutes, 56 seconds

Analysis

Populating the RSA ACE/Server database by importing users from a supported LDAP database, and keeping the databases in sync, are straightforward administrative processes. Larger databases take longer to import and keep in sync. With the scheduling tools built into RSA ACE/Server, however, LDAP import and sync can be automated and done during off-peak hours.

Remote and Web-based Administration Sessions

RSA ACE/Server provides tools to administer its databases from remote machines, either through a desktop application or a Web browser.

The **Remote Administration** tool is a desktop application that enables authorized personnel to administer RSA ACE/Server databases from a Windows-based computer on your organization's local or wide area network. Any administrative function can be performed through Remote Administration.

The **Quick Admin** tool enables Help Desk and other authorized personnel to perform a subset of administrative tasks through their Web browsers.

This section discusses support of multiple concurrent host-mode, remote, and Web-based administrative sessions on an RSA ACE/Server installation. It provides some examples of typical installations for comparison, and describes system settings that you can adjust, if necessary, to improve your system's administrative capacity. Finally, it shows the results of tests to illustrate remote and web-based session limits under simulated system-wide administration loads.

System Settings that Affect Administrative Capacity

The maximum number of concurrent administrative sessions your installation will need to support depends on the size of your user population and the number of administrative personnel involved in supporting your users.

On all platforms, RSA ACE/Server installs default parameter file (.PF) settings for the built-in Progress relational database management system. In addition, on supported UNIX platforms (Solaris, HP-UX, IBM AIX), you may need to adjust some kernel settings to achieve optimum performance.

Refer to the appendix "System Capacity and Resource Utilization" in this book for further discussion of these settings. For procedures to adjust these settings, refer to the *RSA ACE/Server 5.1 for UNIX Installation Guide*.

Scenarios

For most installations, the default parameter settings built into RSA ACE/Server 5.1 should be sufficient to support multiple local and remote administration sessions. However, larger installations can modify certain settings to increase the total number of administration sessions.

The following three subsections provide examples of typical RSA ACE/Server installations, and discuss the required system settings and, in the case of UNIX-based systems, the kernel settings.

Small Installation

A typical small RSA ACE/Server installation might include a Primary Server with one Replica Server, supporting no more than 5000 users. Both Servers would be running on single-CPU machines, and no extra Administration Toolkit (ATK) sessions would be required.

With default settings, this type of installation could support 56 simultaneous local and/or Web-based administration sessions, and 30 Remote administration sessions, for a total of 86 simultaneous administration sessions. This would typically be more than adequate to service a user population of 5000, thus no changes to the default settings would be necessary.

In addition, on a UNIX-based system, you could reduce kernel settings without impacting RSA ACE/Server operation, and thus increase the amount of kernel memory available to the operating system.

Medium-Sized Installation

A typical medium-sized installation might include a Primary with three Replicas, all running on dual-CPU machines. Such an installation might also have 10 custom ATK applications running. It could comfortably support 5,000 to 20,000 users.

With default settings, this type of installation could support 41 simultaneous local and/or Web-based administration sessions, and 30 Remote administration sessions, for a total of 71 simultaneous administration sessions. This would typically be adequate to service a user population of up to 20,000, thus no changes to the default settings would be necessary.

In addition, on a UNIX-based system, you could slightly reduce kernel settings without impacting RSA ACE/Server operation, and thus increase the amount of kernel memory available to the operating system.

Large Installation

A large installation might include a Primary with six Replicas, all running on quad-CPU machines. Such an installation might have as many as 25 custom ATK applications running, and could support 100,000 or more users.

To support 100,000 users, a large installation might need to run as many as 100 concurrent Remote Admin sessions and 50 concurrent Quick Admin sessions.

Consequently, the system administrator would have to make changes to the Progress database parameter file (**Startup.pf** on Windows, **sdserv.pf** and **sdlog.pf** on UNIX), as shown in the following table:

Parameter	Value	Description
-Mn	10	Maximum number of database servers (processes that handle remote client connections)
-Ma	10	Maximum number of remote clients per database server
-n	250	Maximum number of remote clients

The modified parameters provide sufficient capacity for 100 Remote Administration sessions and 99 Host Mode (local) or Quick Admin sessions. No changes to default kernel settings are required.

Note: For complete information about parameter and kernel settings, refer to the appendix "System Capacity and Resource Utilization" in this book.

Remote Administration Session Limits

To check Remote Administration session limits, a "real world" test, designed to simulate a large, typically busy system environment, was run on an RSA ACE/Server configuration. In **startup.pf**, parameter "n" (maximum number of remote clients) was set to 300.

The test setup included a Primary and four Replicas running on the dual-processor Windows 2000 machines described in the table on page 18. The user database was approximately 500K.

With multiple Server processes already running, new Remote Administration sessions were opened until a limit of **263 sessions** was reached. These other Server processes included:

Authentication—Each of the four Replicas had a load of 400 Agents generating authentication requests for a 10-minute period during the test.

Cross-realm—A connection to another realm was made, although left idle.

ATK—A customized ATK (administration toolkit) program ran on the Primary to add, delete, and list users, token groups and sites.

Log monitor—Log monitoring was continuously active on the Primary and four Replicas.

Web Express—A Web Express connection to the Primary was established, although left idle.

Remote Administration—One RA connection was started to provide load by running a Token Statistics report on the Primary.

Quick Admin Session Limits

To check Quick Admin session limits, a theoretical limit was calculated, then tested in a busy system environment. In **startup.pf**, parameter 'n' was set to 300.

As with the Remote Administration test, the system included a Primary and four Replicas running on the dual-processor Windows 2000 machines described in the table on page 18. The user database was approximately 500K.

With multiple Server processes running, new Quick Admin sessions were opened until a limit of **63 sessions** was reached. These other Server processes included:

Authentication—Each of the four Replicas had a load of 400 Agents generating authentication requests for a 10-minute period during the test.

Cross-realm—A connection to another realm was made, although left idle.

ATK—A customized ATK (administration toolkit) program ran on the Primary to add, delete, and list users, token groups and sites.

Log monitor—Log monitoring was continuously active on the Primary and four Replicas.

Web Express—A Web Express connection to the Primary was established, although left idle.

Remote Administration—One RA connection was started to provide load by running a Token Statistics report on the Primary. For additional load, a continuous loop of up to 30 RA processes was set up to connect and disconnect from the database during the test.

Analysis

On the test system, the actual limit of concurrent Remote Administration sessions was **263**, which should be more than adequate to service a user population of 100,000 or more.

The limit of concurrent Quick Admin sessions was **63**, which should be more than adequate to accommodate a large company's Help Desk.

Should an installation require a higher number of remote or Web-based connections to the RSA ACE/Server database, employing faster (quad-processor) machines with more memory is recommended.

Log Maintenance

In the RSA ACE/Server environment, log data is an audit trail of all authentication and administrative activity. It is important to maintain the log database so that it does not use up all available disk space and block authentication.

The data provided here are the result of several log database maintenance tests on the Primary of a single realm.

Server Hardware

Operating System	CPU	RAM
Windows 2000 Advanced Server	Single-processor 750MHz Pentium III	256MB

Log Maintenance Test Results

The following table shows data related to maintenance (editing and compression) of large log databases in RSA ACE/Server 5.1.

Database Before Maintenance		Maintenance			Database After Maintenance	
Total Entries	**Size**	**Log Entries Removed**	**Remove Time**	**Compress Time**	**Total Entries**	**Size**
500K	75MB	50K (10%)	9.5 min.	18.33 min.	450K	67MB
500K	75MB	500K (100%)	92 min.	4 sec.	21	448KB
1000K	150MB	500K (50%)	97 min.	20.66 min.	500K	75MB
1000K	150MB	1000K (100%)	185 min.	7 sec.	30	448MB

Analysis

A log database can grow to very large sizes over a relatively short period of time. For example, a single cross-realm authentication can add several entries to a log database. A user typing the wrong PASSCODE incorrectly a few times can add a dozen or more entries to the log. These and other typical transactions add up very quickly.

Hard drive sizes have gone up exponentially, while prices have come down. Despite this, regular log maintenance is recommended. Otherwise, a failure could occur during a peak authentication period.

As the data shows, removing log entries and compressing large databases can be a time-consuming task. Compression, in particular, places a processing load on the Primary Server, and should be done during off-peak hours.

RSA ACE/Server 5.1 provides a log filtering feature that you can set up to slow log growth. Another tool provided by RSA ACE/Server 5.1 is Automated Log Maintenance, which you can set up to perform log maintenance automatically at scheduled intervals.

A System Capacity and Resource Utilization

Many factors contribute to the capacity of an RSA ACE/Server installation. These are factors such as the number of:

- CPUs
- Replicas
- LDAP synchronization jobs
- Admin Toolkit (ATK) applications
- Local Administration Interface Sessions
- Concurrent Remote Administration Sessions
- Concurrent Quick Admin Sessions

These settings are related to, and in some cases are controlled by:

- Progress RDBMS (relational database management system) parameter settings
- Operating system kernel settings

All these combine to present a complicated set of interrelated values. The values combine the user capacity of the system versus the resources those users and the Progress RDBMS require.

These notes are intended to provide basic and consistent operation of the system. Other configurations may be required depending on specific business requirements.

Reference Documents

The following documents provide useful background information:

- RSA Security Inc., *RSA ACE/Server 5.1 for UNIX Installation Guide*, January 2003
- Progress Software Corporation, *System Administration Guide*, May 1997 (Chapter 14)
- Progress Software Corporation, *System Administration Reference*, November 1996 (Chapter 4)

Note: Contact Progress Software (**www.progress.com**) for information about acquiring publications from that company.

- Installation Guide and/or Administrator's Guide for your OS platform

Note: Always consult the latest documentation from your platform vendor for information pertaining to system setup and configuration.

Users

The Progress database system separates users into two categories: local and remote. The total number of users is obtained by adding these two categories together.

```
 Local Users
+Remote Users
 ============
 Total Users
```

Local Users

From the list of factors on the preceding page that contribute to a system's capacity, the first four items (CPUs, Replicas, LDAP synchronization jobs, Local Apps) contribute to the number of local users. Each local application (process) running on the RSA ACE/Server is considered a "local" user.

The following sum describes the method by which the number of local users is determined:

```
 Back-End processes
 Front-End Process (1)
 Replication Processes
 System Processes
 Local Apps
+QuickAdmin Sessions
 ======================
 Local Users
```

Each of these values is described in detail in the following sections.

Back-End Processes

The number of Back-End processes depends on the number of CPUs. If one CPU is present, two back-end processes are started. If more than one CPU is present, the following equation is used to calculate the number of Back-End processes to start:

```
Number of Back-End Processes = (Number-Of-CPUs x 2) + 1
```

For example, if there are two CPUs, a total of five back-end processes will be started. There is always a single Front-End process that also connects as a local user.

Replication Processes

The number of replication processes varies depending on whether the system is a Primary or Replica server. For the purposes of these calculations, a Primary server will be used. (A Primary server runs one Replication process for each Replica Server in the Realm.)

System Processes

There is generally a fixed number of system processes. These include processes such as:

- the Remote Administration daemon
- the Scheduled Job Executor (**jsed**)
- LDAP synchronization process(es)
- Quick Admin daemon (**sdcommd**)
- Automated Audit Log Maintenance (AALM) daemon.

For estimation purposes, our settings are based around using a value of 10 system processes.

Local_Apps

This is the number of local administrative interfaces (**sdadmin**) and Admin Toolkit (ATK) applications running on the system.

Quick Admin Sessions

Each Quick Admin session contributes the number of local users. Each session has an ATK server process (**apidemon**) started on its behalf.

Remote Users

Server processes are started for each Remote Admin session. Remote users are given access to the database through Progress server processes. Progress configuration parameters control how these servers manage the remote user connections.

Progress RDBMS

In general, RSA ACE/Server uses two types of database connections. The Server, its associated processes, and Quick Admin sessions connect to the database as a "Self-Service Client." Remote Admin sessions connect as a "Remote Client" through a "User Server."

A Progress parameter (-n) controls the number of user entries each database instance will support. The value of this parameter must be large enough to handle "Total Users." This parameter, among others, is specified in the Parameter File.

Parameter File (PF)

The parameter file controls the settings for many Progress RDBMS system capacities. Consult Progress documentation for a more complete description of the settings available in the Parameter File.

Warning: This file should not be changed without understanding the parameters being added or altered. Incorrect settings may prevent you from starting or administering your RSA ACE/Server. They may also decrease the performance (authentication rate) of the Server. By default, the Progress database is well-tuned. If you do make changes, be sure to make a backup copy of the original file beforehand.

On Windows, there is a single startup parameter file (**startup.pf**) located in the ACE **rdbms32** directory.

On UNIX, there are two separate parameter files for the Server and Log databases. These are located in the *ACEPROG* directory and are named **sdserv.pf** and **sdlog.pf** respectively.

The following table describes the capacity-related parameters and their values:

Parameter Flag	Description	Default or Value
-n	Number of Users (Total)	100 (specified)
-L	Lock Table Entries	64000 (specified)
-Mn	Maximum Number of Servers	6 (specified)
-Ma	Maximum Number of Connections-Per-Server	5 (Progress Default)
-Mi	Minimum Number of Connections-Per-Server	1 (Progress Default)

The following table describes the performance-related parameters and their values:

Parameter Flag	Description	Default or Value
-B	Blocks in Database Buffers	800 (Progress default)
-bibufs	Before-Image Buffers	7 (UNIX-only, specified) 5 (WINDOWS-only, Progress default)

Settings for the performance-related parameters are outside the scope of this document.

User Servers

When a Remote Client attempts to establish a connection, either a new server is started or the user is added to an existing server. Parameter file settings control the maximum number of servers ('-Mn') and the maximum number of database connections one server will support ('-Ma'). The total available Remote Client connections available is:

```
Remote Clients = Servers x Connections-Per-Server
```

RSA ACE/Server specifies the number of servers to be 6. By default, Progress allows up to 5 connections per server. Without any modifications, RSA ACE/Server supports:

```
6 Servers (-Mn) x 5 Connections-Per-Server (-Ma) =
            30 Remote Admin sessions
```

Larger installations may need to increase the **-Mn** parameter or add a **-Ma** parameter to increase the number and capacity of the Remote Client servers.

Active Databases

Some kernel parameters use the number of active databases in their calculations. In RSA ACE/Server, there are three databases:

- Server database
- Log database
- Report database

Kernel Parameters

The Progress database uses shared memory to exchange data between processes and semaphores to coordinate access to the shared memory. To handle the number of system and user processes in RSA ACE/Server, larger kernel settings for shared memory and semaphores are required.

Parameter Summary

The following table shows the calculated versus recommended values. These calculations are based on a dual-CPU system with 10 Replica systems:

Parameter		Calculated	Recommended
Shared	shmseg	16	16
	shmmni	48	64
	shmmax	16777216	16777216
Semaphores	semmni	3	4 - 16 (**Note**: On HP-UX systems, do not set this parameter higher than 4.)
	semmns	123	500
	semmnu	123	500
	semmsl	41	200

These parameters are set on HP-UX and Solaris systems. AIX does not require any kernel adjustments for Progress.

Shared Memory

Local user processes use shared memory to exchange data in the database.

Shared Memory Segments Per Process - shmseg

This parameter specifies the number of shared memory segments each user process can attach. Progress recommends a starting value of 16 segments per process.

Maximum Number of Shared Memory Segments - shmmni

This parameter specifies the maximum number of shared memory segments available on the system. Progress recommends the following equation be used to calculate this value:

```
shmmni = (shmseg * Active_Databases)
```

Based on this equation, the RSA ACE/Server value could be:

```
16 Segments Per Process x 3 Active Databases =
48 Shared-Memory Segments
```

To ensure sufficient system capacity, a value of 64 shared memory segments is recommended.

Maximum Size of A Shared Memory Segment - shmmax

This value is the maximum size of a single shared memory segment. The default recommended by Progress is 16Mb (16,777,216 bytes). Almost all operating systems currently in use come with a default shared segment size of 64Mb.

If Progress only uses 16Mb shared memory segment and can attach up to 16 segments to each process, a single process is conceptually capable of consuming 256Mb (268,435,456 bytes).

If the **shmmax** value is less than 16Mb, change it to 16Mb.

Semaphores

Maximum Number of Semaphore Identifiers - semmni

This parameter specifies the maximum number of semaphore identifiers allowed for the system. Progress recommends this be equal to the number of Active Databases (which is three for RSA ACE/Server).

To ensure sufficient system capacity, a value of 16 is recommended.

Note: For HP-UX systems, the **semmni** value should be no higher than 4.

Maximum Number of Semaphores Per Identifier - semmsl

This parameter specifies the maximum number of semaphores allowed per semaphore identifier. Progress provides the following equation for this value:

```
Semmsl =(Local-users + (Remote-users/Clients-per-DB-Server)
+ 4)
```

For RSA ACE/Server, the following value could be used:

```
31 Local-Users + (30 Remote Users/5 Clients-Per-DB) + 4 =
41 Semaphores-per-Identifier
```

To ensure sufficient system capacity, a value of 200 is recommended.

On HP-UX 11, this value is not available as a kernel parameter. On HP-UX 11i (11.11) this parameter is available.

Total Number of Semaphores - semmns

This parameter specifies the total number of semaphores allowed for the system. Progress provides the following equation for this value:

```
Semmns = (semmsl x Active_Databases)
```

For the RSA ACE/Server this is could be:

```
200 (semmsl) x 3 = 600 semaphores
```

Since the **semmsl** value is already considerably over the minimum value, a value of 500 provides more than enough system capacity.

Total Number of Semaphore Undo Structures - semmnu

This parameter specifies the maximum number of semaphore undo structures allowed. Progress recommends that this value be made equal to the total number of semaphores (**semmns**).

Like **semmns**, the currently recommended value is 500 Undo structures.

Other Special Requirements

Some operating systems have other parameters that need to be set for proper operation.

HP-UX 'nproc'

This parameter specifies the maximum number of processes that can run on the system. This value is increased from 276 to 640 on HP-UX systems.

AIX 'fsize'

This parameter specifies the maximum file size that the system can create. This is only kernel parameter altered on AIX. This value must be increased if you are planning for a large database. See the *RSA ACE/Server 5.1 for UNIX Installation Guide* for more information.

General UNIX 'ulimit' Command

This command can be used on any supported UNIX platform to specify process-specific resource limits. It is commonly used to alter the maximum number of file descriptors (concurrently open files) that the system will allow a single process to open (**ulimit -n**). If you experience an error such as “Cannot open file: Too many open files,” the current **ulimit -n** value should be increased. The current default value is obtained by running the **ulimit** command without a value (for example, **ulimit -n**). Increase the current value and repeat the process with which you originally encountered the error. In general, this type of error should never occur using an RSA Security application. If this occurs with internally developed applications, check that files opened by the application are being closed under all conditions.

B Authentication Performance Test Data

The peak authentication graphs starting on page 19 of this book were based on the data provided in this appendix.

Glossary

To interpret the test data, it is helpful to understand the following terminology.

Agents—In the peak authentication tests, *virtual load agents* were developed to simulate RSA ACE/Agent authentication requests. For each Replica, 400 virtual load agents were set up to send continuous authentication requests to the RSA ACE/Server.

TPS—Abbreviation for *transactions per second*, this is the load of authentication requests being imposed by virtual load agents that are successfully handled by the RSA ACE/Server. Therefore, one TPS is equivalent to one authentication request per second, or APS.

Local Authentication - Windows

The following table corresponds to the graph on page 19 entitled, “Peak Authentication - Windows 2000 (Dual-Processor/Isolated Network Configuration).”

Replica(s)	Agents	TPS		
		DB 100K	500K	1000K
1	400	80.006	69.810	72.732
2	800	121.910	124.878	127.565
3	1200	185.030	194.030	187.680
4	1600	209.380	206.560	211.830
5	2000	219.510	217.773	215.540
6	2400	222.390	220.030	222.230
7	2800	n/a	219.530	n/a
8	3200	209.860	209.720	n/a
9	3600	n/a	196.170	n/a
10	4000	186.025	185.495	185.620

The following table corresponds to the graph on page 19 entitled, “Peak Authentication - Windows 2000 (Dual-Processor/ “Real World” Network Test Configuration).”

Replica(s)	Agents	TPS	
		DB	500K
1	400	57.60	
2	800	103.30	
3	1200	156.08	
4	1600	205.57	

Local Authentication - Solaris

The following table corresponds to the graph on page 20 entitled, “Peak Authentication - Solaris 8 (Single-Processor/Isolated Network Configuration).”

Replica(s)	Agents	TPS			
		DB	100K	500K	1000K
1	400	40.30		38.55	35.78
2	800	75.60		73.14	68.51
3	1200	95.66		93.62	88.63
4	1600	112.69		110.40	102.40
5	2000	119.52		116.46	111.56
6	2400	125.69		124.87	119.97

The following table corresponds to the graph on page 20 entitled, “Peak Authentication - Solaris 8 (Single-Processor/ ‘Real World’ Network Test Configuration).”

Replica(s)	Agents	TPS	
		DB	500K
1	400	22.09	
2	800	44.31	
3	1200	66.68	
4	1600	88.97	

Remote Authentication through RADIUS

The following table corresponds to the graph on page 21 entitled, “Peak Remote RADIUS Authentication - Windows 2000 (Dual-Processor/Isolated Network Configuration).”

Replica(s)	Agents	TPS		
		DB 100K	500K	1000K
1	400	64.01	60.06	55.597
2	800	133.06	115.258	122.550
3	1200	175.46	156.485	169.263
4	1600	206.87	200.949	199.198
5	2000	221.41	219.795	215.758
6	2400	221.03	217.653	219.270
7	2800	n/a	n/a	n/a
8	3200	n/a	n/a	n/a
9	3600	n/a	n/a	n/a
10	4000	179.11	177.974	179.502

The following table corresponds to the graph on page 21 entitled, “Peak Remote RADIUS Authentication - Windows 2000 (Dual-Processor/ ‘Real World’ Network Test Configuration).”

Replica(s)	Agents	TPS	
		DB	500K
1	400	49.109	
2	800	93.952	
3	1200	144.557	
4	1600	185.304	

Cross-Realm Authentication

The following table corresponds to the graph on page 26 entitled, “Cross-Realm Authentication (Windows 2000) (Dual-Processor / Isolated Network Configuration).” The data refers to the **remote realm**, which is processing the authentication requests. A second realm, the home realm, is also involved in the process. The remote realm polls the home realm to confirm that each authenticating user resides in the home realm’s database (which was 80K in size).

Replica(s)	Agents	TPS	
		DB	100K
1	400	34.30	
2	800	49.40	
3	1200	63.51	
4	1600	75.54	

Index

Q

R

S

T

U

V

W