RSA ACE/Server 6.0
Scalability and Performance Guide

Contact Information

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

RSA Security Inc.	RSA Security Ireland Limited
www.rsasecurity.com	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 2005

Contents

1 Introduction

This book assists you in planning the setup and use of RSA ACE/Server 6.0 in your network infrastructure. It is intended for the information technology professionals responsible for your organization's network security.

The information is designed to help you plan for an RSA ACE/Server 6.0 installation servicing a large number of users. It assumes a you have an RSA ACE/Server 6.0 Advanced license and multiple Replica Servers.

This chapter describes the RSA ACE/Server 6.0 environment, and recommended configurations for supported server hardware. It also discusses scalability and performance issues, network loads, and administration and maintenance issues.

Chapter 2, "Performance Tests," presents and discusses the results of a number of performance test profiles run on RSA ACE/Server 6.0, and the equipment used in these tests. It also discusses administration capacity, which is dependent on the size of the installation.

Frequently Asked Questions

This document helps answer the following frequently asked questions about RSA ACE/Server 6.0 performance and scalability:

- 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 other issues 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 6.0 environment. For details about system installation, administration, and use, see the *RSA ACE/Server 6.0 Installation Guide* for your platform and the *Administrator's Guide*.

Primary/Replica Servers

A typical RSA ACE/Server 6.0 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*.

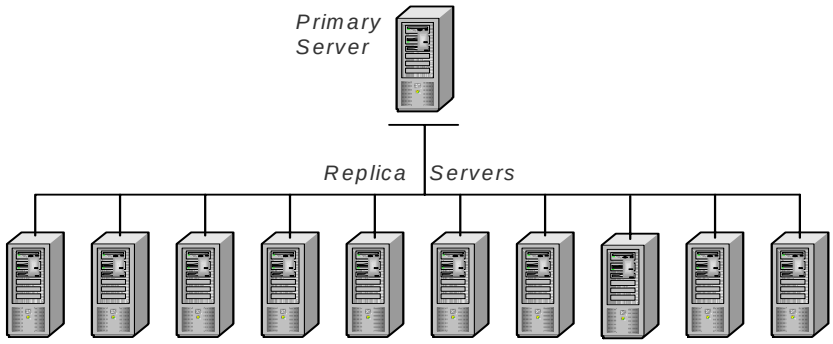

Note: The RSA ACE/Server Base license allows one Primary and one Replica Server in one realm. If you want to deploy more than one Replica Server or more than one Replica Server (multiple realms), you must purchase an Advanced license. Contact your authorized RSA Security sales representative, or go to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

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, shut down the authentication service running on that machine. For information, see "Database Maintenance" for your platform in the *Administrator's Guide*.

For small to medium-sized organizations connected through a local area network (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

Each Primary/Replica array in an installation is called a realm. For larger organizations that have offices in distant remote sites, multiple realms may be preferable.

With the RSA ACE/Server Advanced license, you can have up to six realms, and you can configure each 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.

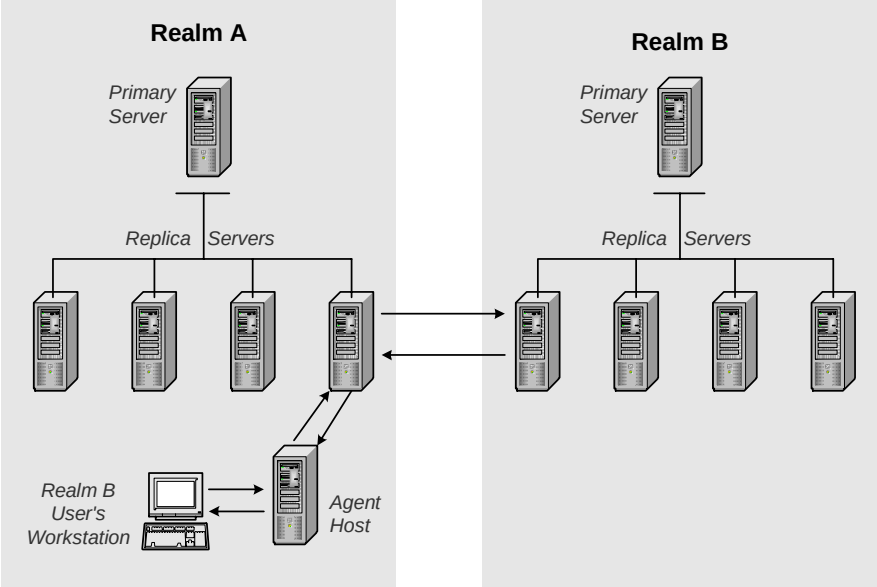

When the authenticating Replica in Realm A checks the database and does not find the user, it polls each realm until it finds the Server in Realm B that has a record for the visiting user. The authenticating Replica in Realm A then sends the authentication request to the Server in Realm B, which authenticates 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 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 6.0. 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, see the *Administrator's Guide*.

Offline Authentication

Offline authentication extends RSA SecurID authentication to users when the connection to the RSA ACE/Server is not available (for example, when users work away from the office, or when network conditions make the connection temporarily unavailable). You enable, disable, and configure offline authentication through the RSA ACE/Server Database Administration Application.

You can use offline authentication in the traditional Primary-Replica-Agent configuration, where offline data is downloaded to the desktop on which the RSA ACE/Agent software is installed. This is known as a Local Authentication Client. You can also use offline authentication in a Domain environment, where part of the Agent software is installed on a domain controller, and part is installed on a client of the domain controller. This is known as the Domain Agent Client-Domain Agent Host. In this case, offline authentication data is downloaded to the domain controller and to the client.

Note: The Domain Authentication Host is also known as the Domain Authentication Server.

The following graphic illustrates the Domain Agent Client-Domain Agent Host environment.

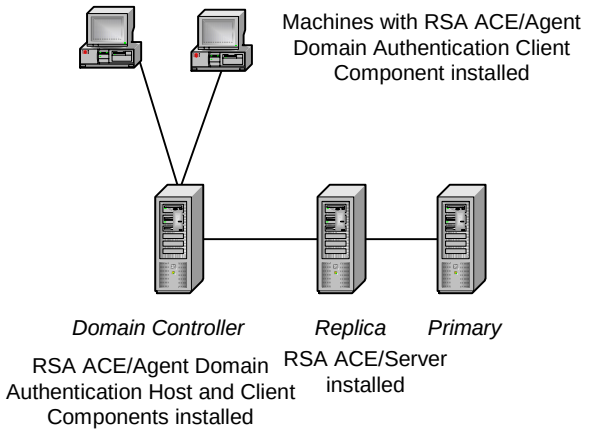

Recommendations for Supported Server Hardware

RSA ACE/Server 6.0 runs on a variety of Windows- and UNIX-based hardware platforms. For complete information about system requirements, refer to the *RSA ACE/Server 6.0 Installation Guide* for your platform.

This section discusses additional considerations for 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 6.0 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	10 K	512 MB	EIDE or SCSI-2
		100 K	640 MB	Ultra SCSI
		500 K	1024 MB	Ultra SCSI
Sun/Solaris	Ultra SPARC II/Dual 300MHz or faster	10 K	256 MB	EIDE or SCSI-2
		100 K	300 MB	Ultra SCSI
		500 K	700 MB	Ultra SCSI
Domain Controller	Dual processor 850 MHz or faster	10 K	2 GB	Ultra SCSI

In general, faster processors, faster hard drives, and more memory will result in better RSA ACE/Server 6.0 performance. Preliminary testing on higher-end equipment shows performance rates that are approximately twice those documented in this guide.

Note: RSA ACE/Server 6.0 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. Avoid using these computers as file servers, firewalls, or for any other application.

Scalability and Performance Considerations

When planning your RSA ACE/Server 6.0 installation, there are a number of factors to consider regarding your user population and network traffic. The following subsections discuss these factors.

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 *two and a half million users per realm*. 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. This is typically more than adequate to service the data associated with one million (or more) users. For more information, see "Database Maintenance" for your platform in the *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.

For data about RSA ACE/Server 6.0 authentication rates and guidance on optimum Server configurations within a realm, see "Peak Authentication Measurements" on page 22.

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 fails, 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 fail, users can still authenticate to other Replicas in other locations on the WAN.)

If you are using domain controllers in a Domain Agent Client-Domain Agent Host environment, you need to consider the number of users per domain controller. The following table is based on recommendations by Microsoft on the number of domain controllers and CPU speed, based on the number of users.

Users per Domain in a Site	Minimum Number of Domain Controllers Required per Domain in a Site	Minimum CPU Speed Required per Domain Controller
1 - 499	One	Uniprocessor 850 megahertz (MHz) or faster
500 - 999	One	Dual processor 850 MHz or faster
1,000 - 2,999	Two	Dual processor 850 MHz or faster
3,000 - 10,000	Two	Dual processor 850 MHz or faster
10,000 users	One for every 5,000 users	Dual processor 850 MHz or faster

Locations of Users

When planning your RSA ACE/Server 6.0 installation, the locations of your user population, and the time zones in which they work, are 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 internet service provider), 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 enables you to establish up to six realms.

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

In choosing the physical location of Primary and Replica Servers in a realm, 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. For information about authentication performance within the same realm, see "Peak Authentication Measurements" on page 22. For information about cross-realm authentication performance, see "Cross-Realm Authentication" on page 32.

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, or return from lunch, and log in to the network.

For example, if your organization runs 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.

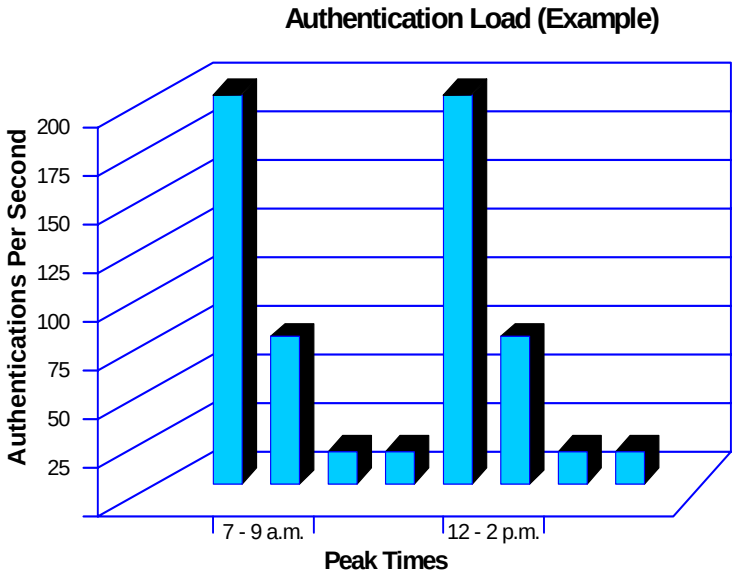

When planning an RSA ACE/Server 6.0 installation, it is important to allow for the peak authentication rates that your organization is likely to experience.

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

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

Network Traffic

RSA ACE/Server 6.0 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 6.0 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 denied. The actual time to complete this communications loop depends on the speed of your network.

Configuration	Authentication		With Offline Data		Each Additional Day of Offline Data	
	Packets (UDP)	Bytes	Packets (TCP)	Bytes	Packets (TCP)	Bytes
Local Authentication Client-RSA ACE\Server	6	2532	61	51787	33	48464
Domain Agent Client-Domain Agent Host	342 (71from RSA SecurID authentication + 271 from Microsoft domain authentication)	86527 (9265 + 77262)	52	50644	57	53814
Domain Authentication Host-RSA ACE\Server	6	2532	64	51907	40	48824

Local Authentication Client to ACE/Server Traffic

Each authentication involves 2,532 bytes of data (in six 422-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 849 bytes (with no downloading of offline data) or 1402 bytes (with downloading of offline data) for every Replica in the realm. Therefore, on a fully-configured system, each authentication could generate up to 10,500 bytes (with no downloading of offline data) or 16,500 bytes (with downloading of offline data) of network load. So, for example, a realm with a Primary and 10 Replica Servers experiencing 100 authentications per second (APS) would consume 1,050,000 bytes (with no downloading of offline data) or 1,650,00 bytes (with downloading of offline data) per second of network capacity.

Domain Agent Client-Domain Agent Host Traffic

Each authentication involves 86,527 bytes of data (9,265 from RSA SecurID authentication and 77,262 from Microsoft domain authentication) sent between the Agent and the Server.

This per-authentication load will grow approximately another 849 bytes (with no downloading of offline data) or 1402 bytes (with downloading of offline data) for every Replica in the realm. Therefore, on a fully-configured system, each authentication could generate up to 17,600 bytes (with no downloading of offline data) or 23,000 bytes (with downloading of offline data) of network load. So, for example, a realm with a Primary and 10 Replica Servers experiencing 100 authentications per second (APS) would consume 1,760,000 bytes (with no downloading of offline data) or 2,300,00 bytes (with downloading of offline data) per second of network capacity.

Replication

In RSA ACE/Server 6.0, 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 complete 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 approximately 5.5 packets per Replica totaling 617 bytes. The following table shows the additional packets added to a replication pass for specific actions resulting from authentication and administration.

Action	Additional Packets (Bytes)
Adding an Agent Host	325
Logging a message	163
Adding a user	476
Adding a token	518
Single authentication (no downloading of offline data)	849
Single authentication (with downloading of offline data)	1402

Load Balancing with RSA ACE/Agent Software

RSA ACE/Server 6.0 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 6.0 and RSA ACE/Agent 5.0 or later, you can take advantage of this load balancing capability.

Note: Although previous versions of RSA ACE/Agent software work with RSA ACE/Server 6.0, they do not take advantage of automatic load balancing. For a list of RSA ACE/Agent software, go to **http://www.rsasecurity.com/node.asp?id=1174**.

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 must use to process requests.

For complete information about load balancing capabilities, see the *Administrator's Guide*.

Network Latency

Another item to consider when planning your RSA ACE/Server 6.0 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, because users at each site would be authenticating to the nearest Replica.

Although the same setup increases latency for replication, the replication load is significantly less than the authentication load, and replication can be distributed over a longer period of time. For more information about replication load, see "Database Replication" on page 34.

Another advantage of this approach is if one of the Replicas fails, users at that remote site can still 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 Remote Access Server (RAS), 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.

Offline Authentication Log Consolidation

When an Agent Host (in either Local Authentication Client or Domain Agent Client-Domain Agent Host environments) reconnects to the RSA ACE/Server, accumulated log entries are sent to the Server audit log, which is consolidated on the Primary. Uploading the accumulated log entries increases network traffic. In the Local Authentication Client environment, each entry generates 26 packets (3,513 bytes). In the Domain Agent Client-Domain Agent Host environment, each entry generates 26 packets (3,481 bytes).

You can disable consolidation of the log entries through the Database Administration application. Click **System** > **System Administration** > **Edit Offline Auth Config**, clear **Upload offline authentication log entries when user reconnects**, and click **OK**.

Administration Considerations

When planning an RSA ACE/Server 6.0 installation, you must 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, see the *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. 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 supports, see "Remote and Web-Based Administration Sessions" on page 37.

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. Take this into account when considering how much of the administration load will be performed via the web. 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, see "Remote and Web-Based Administration Sessions" on page 37.

LDAP Import and Synchronization

In RSA ACE/Server 6.0, 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 the RSA ACE/Server database, and schedule automatic updates to keep the two databases in sync.

Although the LDAP import and update functions use network bandwidth, you can use the built-in scheduler to run these functions at off-peak times. For information on performance test data for LDAP import, see "LDAP Import" on page 36.

Log Maintenance

RSA ACE/Server 6.0 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 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, are unusable. In this event, users are 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. RSA Security recommends compressing the database after deleting a large number of log records. For log maintenance test results, see "Log Maintenance" on page 41.

For log maintenance, you can also use the *log filtering* tool in RSA ACE/Server 6.0. 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.

For information about consolidating offline authentication logs, see "Offline Authentication Log Consolidation" on page 16.

Disaster Recovery

The Primary/Replica model of RSA ACE/Server 6.0 enables failover protection 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 failing.

If your Primary Server machine fails, the Replicas will continue authenticating users. If you expect your Primary to be out of service for a day or more, RSA ACE/Server 6.0 provides a *Nominate* feature that enables you to select and configure a Replica Server to become the Primary. For complete information, see the *Administrator's Guide*.

If an authenticating Replica fails, RSA ACE/Server 6.0 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. For DBPush test results, see "Database Push" on page 36.

Note: For organizations in which high availability of network resources is a requirement, RSA ACE/Server 6.0 supports the 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 6.0. An analysis of these test results helps network administrators plan a more optimal installation and deployment of RSA ACE/Server 6.0 in their own network environments.

Test Systems and Network Environment

This section describes the test environments that RSA Security used for RSA ACE/Server 6.0 performance tests.

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 6.0 is a multi-threaded application that takes advantage of multi-processor systems. Dual-processor systems help optimize performance, as do faster networks.

Note: Preliminary testing on higher-end equipment shows performance rates that are approximately twice those documented in this guide.

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

Local Authentication and Domain Authentication

RSA ACE/Server 6.0 includes new features that enable authentication through domain controllers (Domain Authentication Client-Domain Authentication Host). Domain authentication requires users to provide RSA SecurID passcodes to authenticate to their Microsoft domain. In addition, you can use the traditional method and authenticate to the RSA ACE/Server (Local Authentication Client). Both configurations allow users to authenticate to machines that are not connected to the network (offline authentication).

Normally, a user's name and passcode are authenticated when the RSA ACE/Agent sends them to the RSA ACE/Server. Offline authentication extends the RSA SecurID solution when a user is disconnected from the corporate network or a connection to the RSA ACE/Server is temporarily unavailable.

For domain authentication, domain controllers can act as proxies in case a network outage causes the RSA ACE/Server to become temporarily unavailable. Using offline authentication data on their local machines, users are still able to authenticate and gain access to the domain. The RSA ACE/Agent uses strongly encrypted offline authentication data on the local machine or domain controller.

New offline authentication data is downloaded each time users reconnect their computers to the network, or the domain controller reconnects to the RSA ACE/Server. Offline authentication data expires when the protected resource has been offline beyond a certain amount of time (as specified by the RSA ACE/Server administrator).

For the Local Authentication Client tests, each configuration included a Primary Server, with up to six Replica Servers, and various Agent Hosts and Remote Administration machines.

Local Authentication Client Model

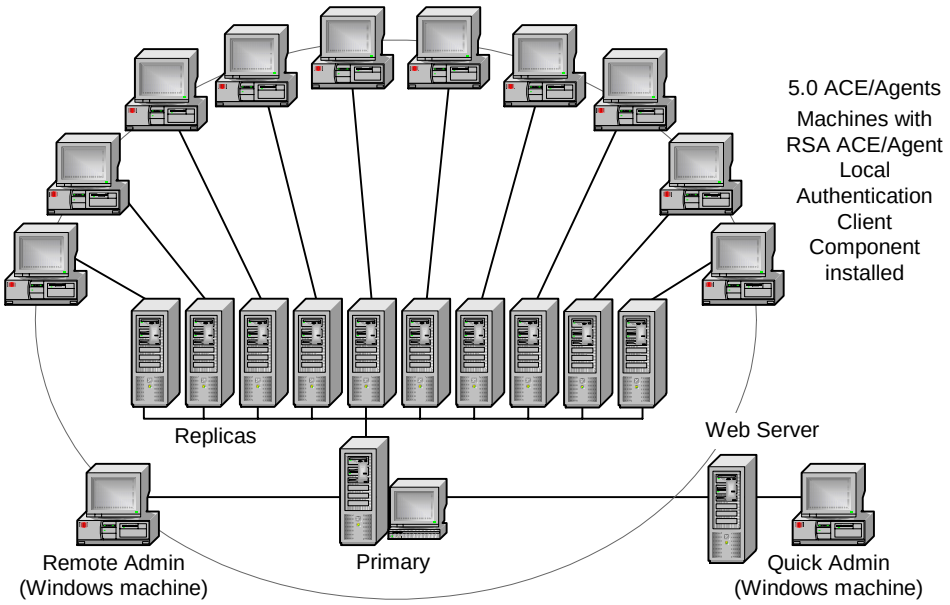

For the Domain Agent Client-Domain Agent Host tests, each configuration included a Primary Server, with up to six Replica Servers, and two domain controllers per Replica, as well as various Agent Hosts and Remote Administration machines.

Domain Authentication Client-Domain Authentication Host Model

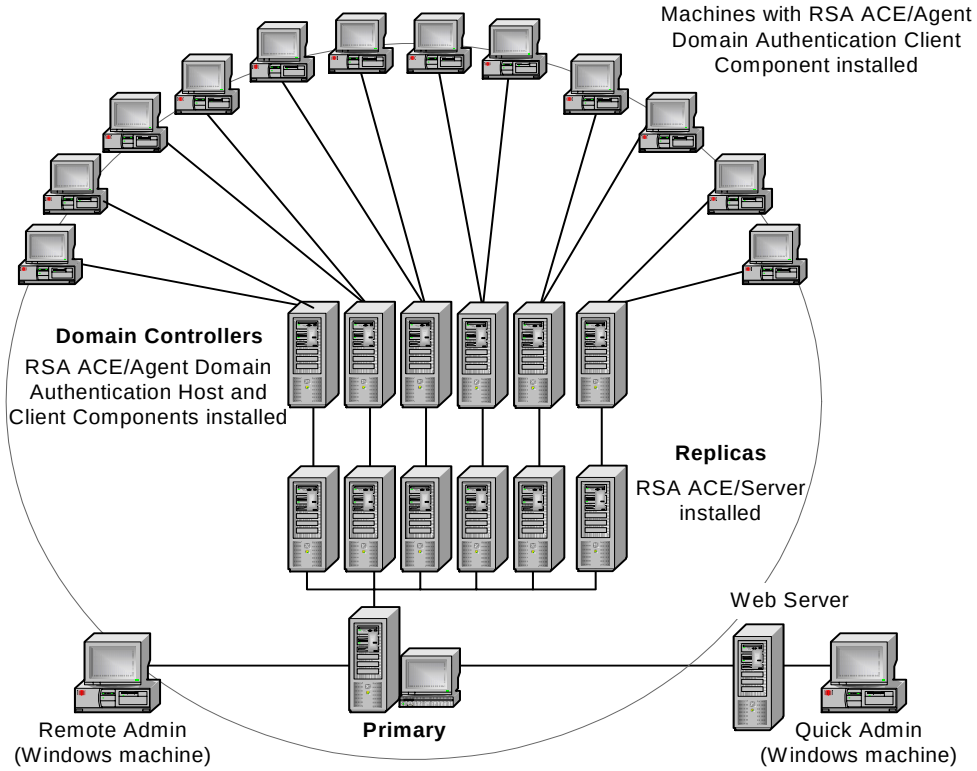

The following table describes the specific server hardware used in these performance tests. Any hardware differences in specific tests are indicated in the test-related sections later in this chapter.

Platform	Server Hardware	Operating System
Windows	Dual-processor Pentium III (750 MHz or 933 MHz), 1024 MB RAM	Windows 2000 Advanced Server
Windows Domain Controllers	Dual-processor Pentium 4 Xeon (3.2 GHz), 2048 MB RAM	Windows 2000 Advanced Server
Sun	Sun Fire V100 workstations, single UltraSPARC IIe 500 MHz CPU, 1024 MB RAM	Solaris 9

Note: Windows tests ran on dual-processor systems. Solaris tests ran on single processor systems.

Peak Authentication Measurements

Peak authentication is the highest number of authentications per second (APS) that RSA ACE/Server 6.0 can achieve for short, concentrated periods. This section describes and analyzes peak authentication performance in single-realm Windows 2000 and Solaris 9 environments.

Data was gathered from tests conducted in isolated configurations.

For data and analysis covering cross-realm authentication, see "Cross-Realm Authentication" on page 32.

Peak Authentication on the Windows 2000 Test System (Local Authentication Client)

The following graph shows the measured peak authentication rates for the Windows 2000 Local Authentication Client test configuration and the impact on authentication rates of downloading 0, 1, 3, and 30 days of offline data. Note that while the impact of downloading 30 days of offline data is significant, users would seldom download offline data for that many days at one time, and then only during the initial download of offline data, if you set the parameter to 30 days. After the initial download, the number of days of offline data is kept up-to-date. For example, a user who takes his or her laptop home on the weekend would download two days of offline data after returning to work on Monday.

The tests were performed in an isolated network setup with a 100,000-user database, a Primary, from one to six Replicas (with one Domain Authentication Host per Replica), and two 6.0 Agents per Replica, each running 250 Agent threads. Test results from RSA ACE/Server 5.1 are included for comparison.

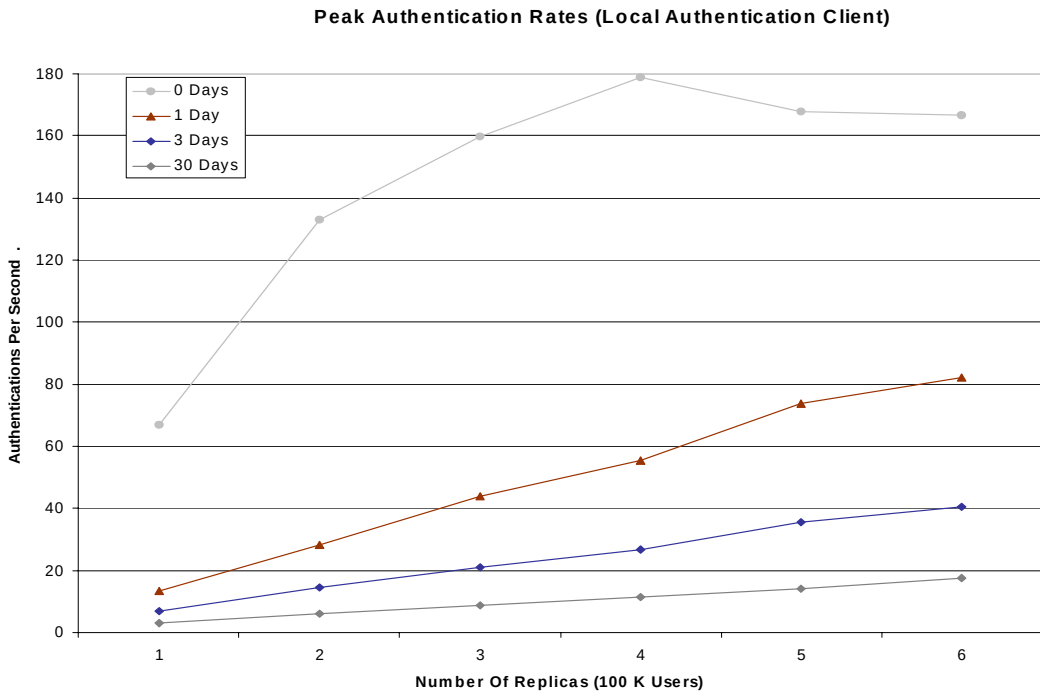

Peak Authentication on the Windows 2000 Test System (Domain Agent Client-Domain Agent Host)

The following graph shows the measured peak authentication rates for the Windows 2000 Domain Agent Client-Domain Agent Host test configuration (one Domain Agent Host per Replica), and the impact on authentication rates of downloading 0, 1, and 3 days of offline data.

The tests were performed in an isolated network setup with a 50,000-user database (100,000 users for the five and six Replica configurations), a Primary, from one to six Replicas (with one Domain Authentication Host per Replica), and two 6.0 Agents per Replica, each running 250 threads.

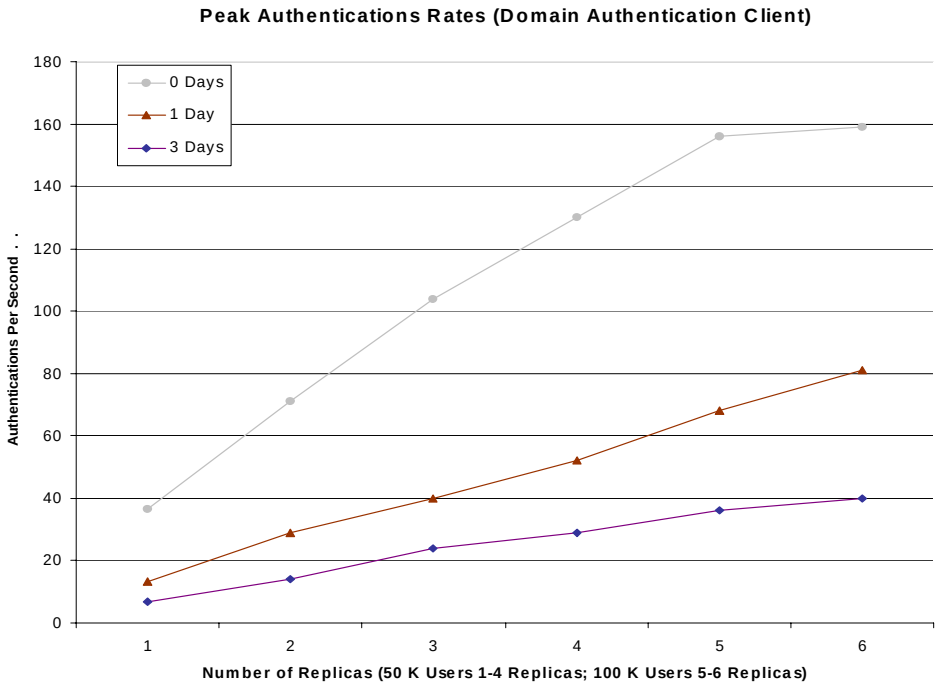

Peak Authentication on the Sun Solaris Test System (Local Authentication Client)

The following graph shows the measured peak authentication rates for the Sun Solaris test configuration in the Local Authentication Client and Domain Authentication Client-Domain Agent Host environments. The tests were performed in an isolated network setup with a 100,000-user database, a Primary, from one to six Replicas, and two 6.0 Agents per Replica, each running 250 threads. Test results from RSA ACE/Server 5.1 are included for comparison.

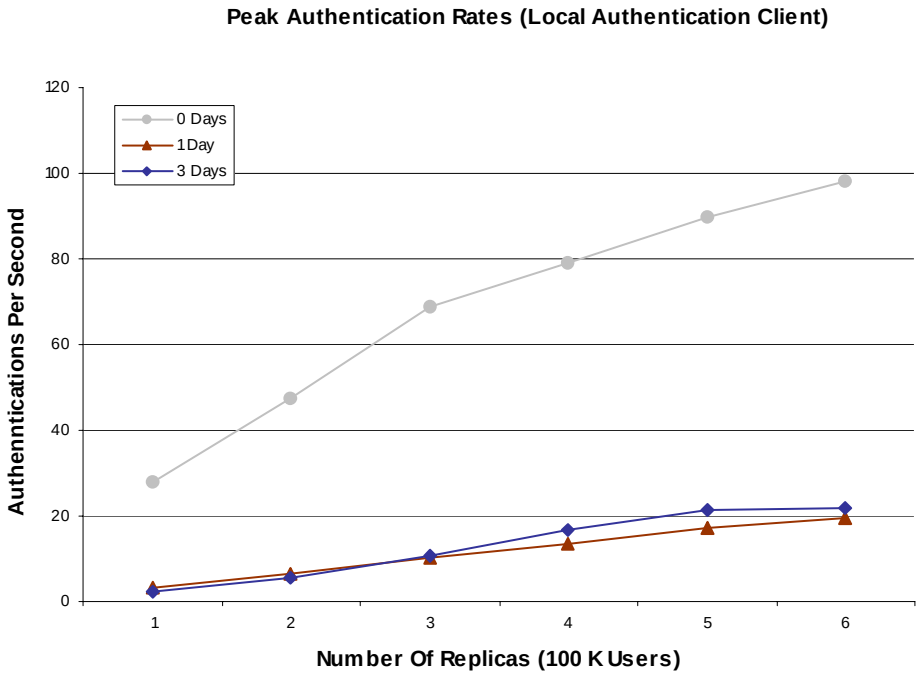

Peak Authentication on the Sun Solaris Test System (Domain Agent Client-Domain Agent Host)

The following graph shows the measured peak authentication rates for the Sun Solaris test configuration in the Domain Authentication Client-Domain Agent Host environment.The tests were performed in an isolated network setup with a 100,000-user database, a Primary, from one to six Replicas (with one Domain Authentication Host per Replica), and two 6.0 Agents per Replica, each running 250 threads.

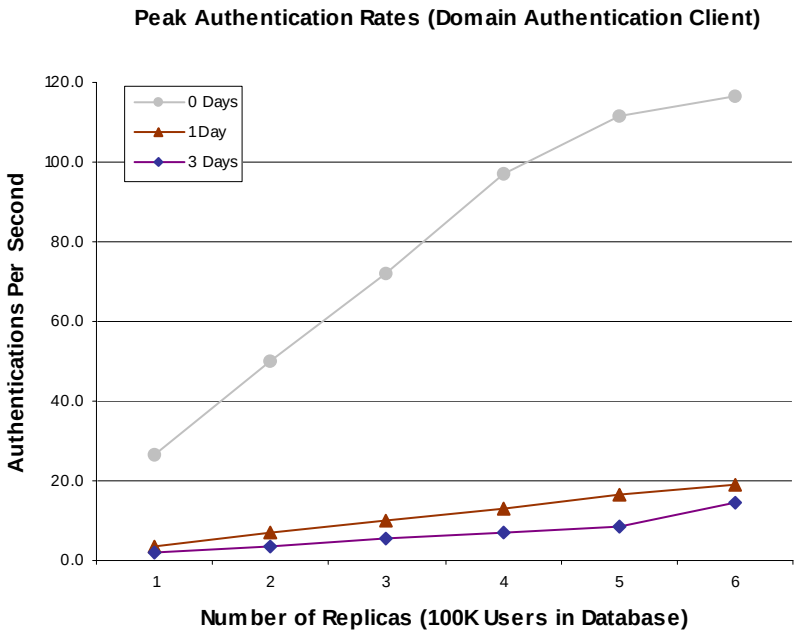

Analysis of Peak Authentication Measurements

RSA ACE/Server 6.0 is not designed to achieve peak authentication rates indefinitely, because replication, logging, and downloads of offline data would be compromised. Use the peak authentication rates in this book 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 **180 authentications per second** on a representative dual-processor configuration on the Windows 2000 platform.
- The highest authentication rates were achieved with a configuration of **one Primary and four Replicas**, with no offline data downloaded.
- Increasing the number of Replicas per realm increases performance for all the database sizes used in the tests. Adding more Replicas provides 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 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, the amount of cross-realm authentication in your organization, and the amount of offline data being downloaded. Preliminary testing on higher-end equipment shows performance rates that are approximately twice those documented in this guide.
- 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.

Effect of Authentication Load

This section contains charts based on the peak authentication test results for the Domain Agent Client-Domain Agent Host configuration. The charts extrapolate performance under various authentication loads. The charts project the length of time it takes for various Primary-Replica configurations to completely process different authentication loads, and the length of time it takes to deliver offline data (with one or three days of offline data downloaded) for each user in a 50,000 user database.

The tests were run using the same configurations used in the peak authentication tests for "Peak Authentication on the Windows 2000 Test System (Domain Agent Client-Domain Agent Host)" on page 24. This configuration enabled a total of 3,000 authentication requests to be sent simultaneously. The projected results estimate the effect of heavy authentication loads on authentication and on delivery of offline data.

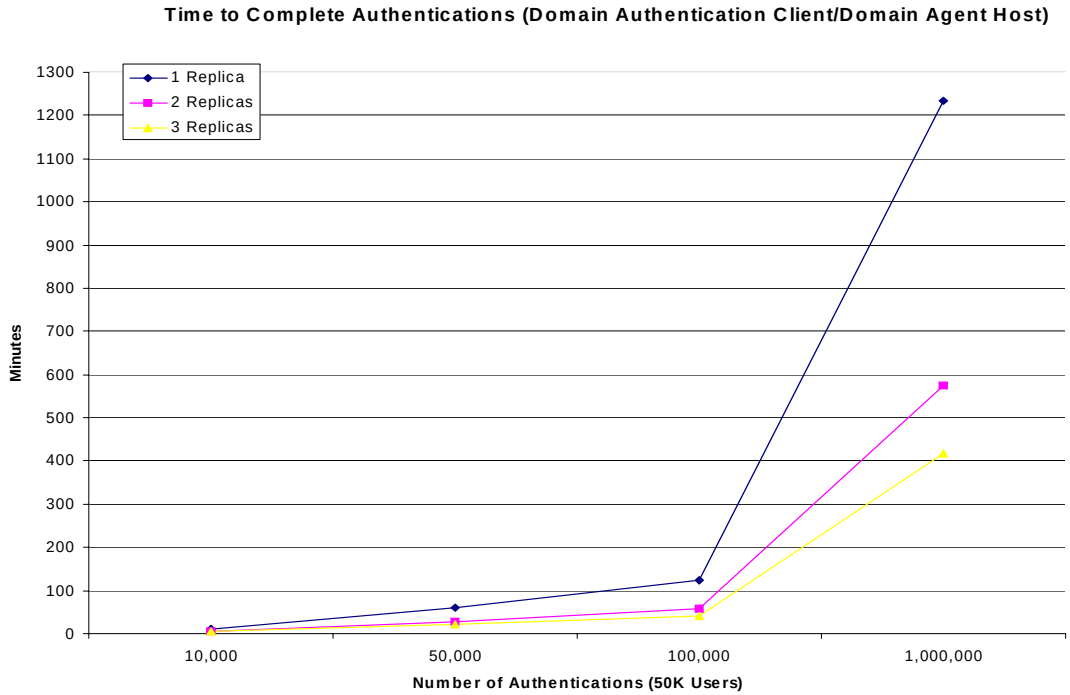

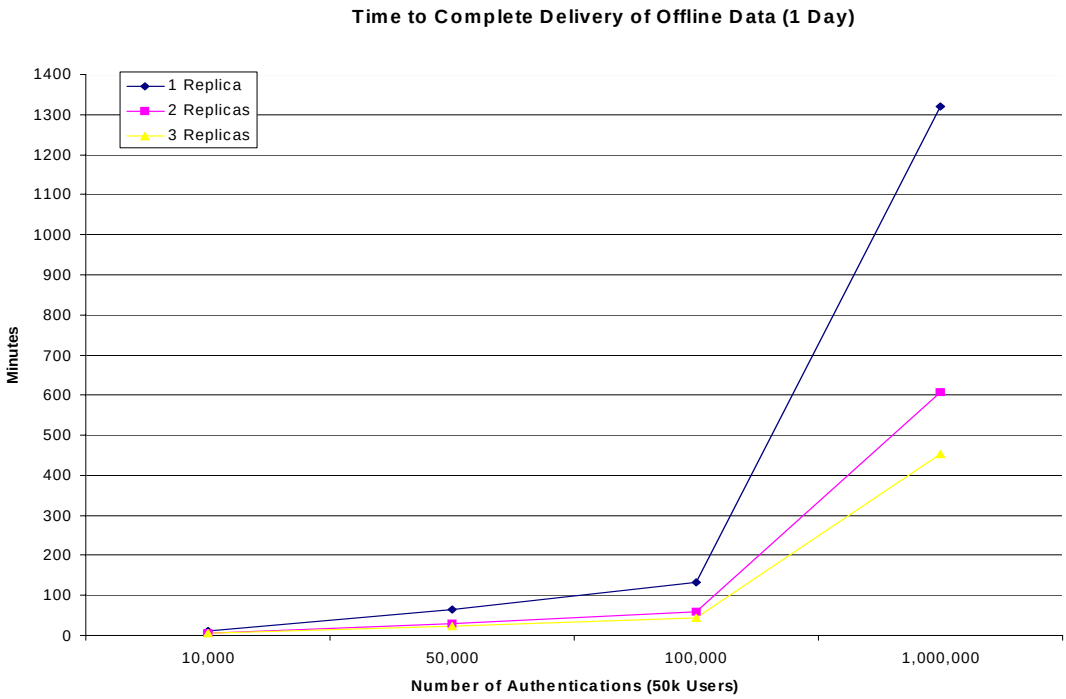
Time to Complete Delivery of Offline Data (1 Day)
1 Replica
2 Replicas
3 Replicas
Minutes
1400
1300
1200
1100
1000
900
800
700
600
500
400
300
200
100
0
10,000
50,000
100,000
1,000,000
Number of Authentications (50k Users)

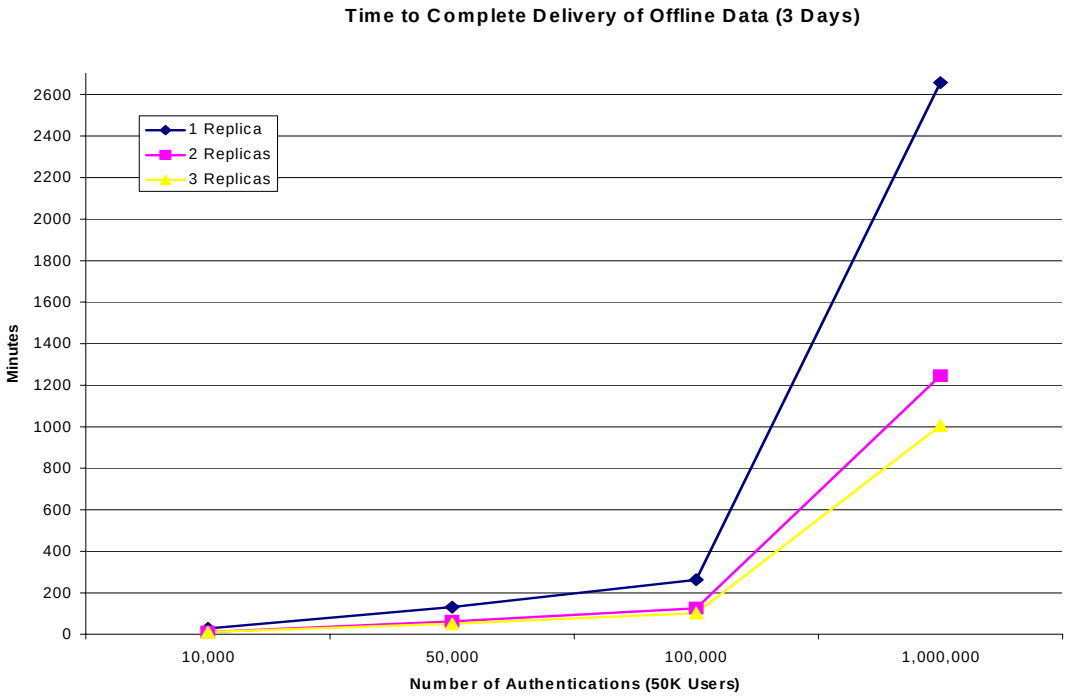
Time to Complete Delivery of Offline Data (3 Days)
1 Replica
2 Replicas
3 Replicas
Minutes
2600
2400
2200
2000
1800
1600
1400
1200
1000
800
600
400
200
0
10,000
50,000
100,000
1,000,000
Number of Authentications (50K Users)

Sustained Authentication Rate

This section describes a test measuring **sustained authentication rate**, which is the authentication rate that RSA ACE/Server 6.0 performs on an ongoing basis while maintaining acceptable replication and administrative loads.

Why Replication Is Important

Since multiple Servers running RSA ACE/Server software authenticate users simultaneously, it is essential that all of these Servers' databases are frequently updated so they are identical and current. 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 at the same time—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 has 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 takes place, the more secure an RSA ACE/Server installation is.

Sustained Authentication Test Results

Tests were run every 15 minutes over a 24-hour period 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 6.0 software. (For the hardware specifications of these machines, see "Test Systems and Network Environment" on page 19.) The user database contained 100,000 users. A Local Authentication Client configuration was used with no offline data downloads.

Only the Replicas were set up to authenticate users. A dedicated Agent Host for each Replica generated a stream of authentication requests over a 24-hour period. Load balancing was disabled.

During the 24-hour period, each Agent Host ran 60 concurrent daemons. Each daemon sent an authentication request to the target Replica, waited for a response, then sent another request.

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

Average Reconciliation Time (minutes)	Authentications per Second per Replica	Total Authentications per Second
78	8.65	51.9

Analysis

With a configuration of a Primary and six Replicas, RSA ACE/Server 6.0 was able to sustain **51.9 authentications per second** (APS). With the default replication interval of 100 seconds, full reconciliation between the Primary and all Replicas occurred in 78 minutes.

Based on this, an optimally-configured RSA ACE/Server 6.0 system is capable of **186,800 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 keep up-to-date 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 (WAN) connecting multiple LANs through leased high-speed telecommunication 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 12,500 employees require authentication.

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

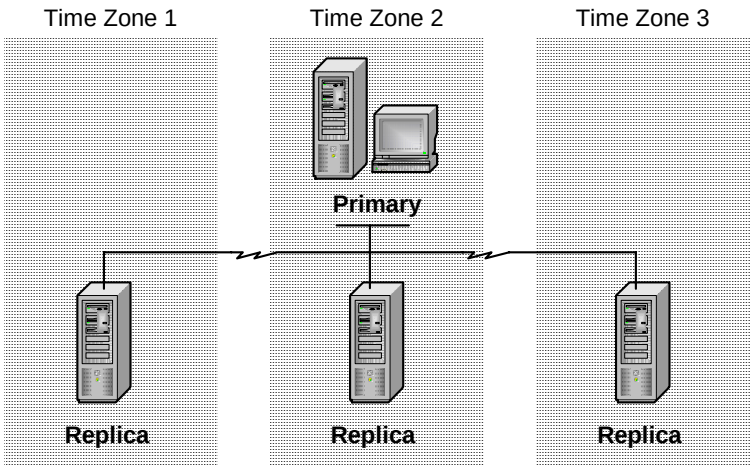

Each site has at least one authenticating Server. Note that this example shows three authenticating Replicas, half the number used in the RSA Security test. With this type of configuration, RSA ACE/Server 6.0 sustains approximately 90,000 authentications per hour indefinitely, and even higher rates for short durations during peak authentication periods. This would be more than adequate for the organization in this example. Even if one Replica failed, the other Replicas in the realm would provide sufficient 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 6.0 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*. (For an introduction to cross-realm authentication, see "Realms" on page 7.)

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 6.0 Agents, use pre-5.0 protocol to communicate with the home realm. This means that each Agent authenticates to only 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 6.0 Agents, the designated Server is referred to as the Preferred Server.

For that reason, and because it is necessary to transmit additional data packets, each cross-realm authentication is slower than an authentication within a realm. For more information about network data packets transmitted by various RSA ACE/Server 6.0 processes, see "Network Traffic" on page 13.

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 to four Replica Servers. The home realm had a Primary and one Replica. Version 6.0 Agents were used exclusively.

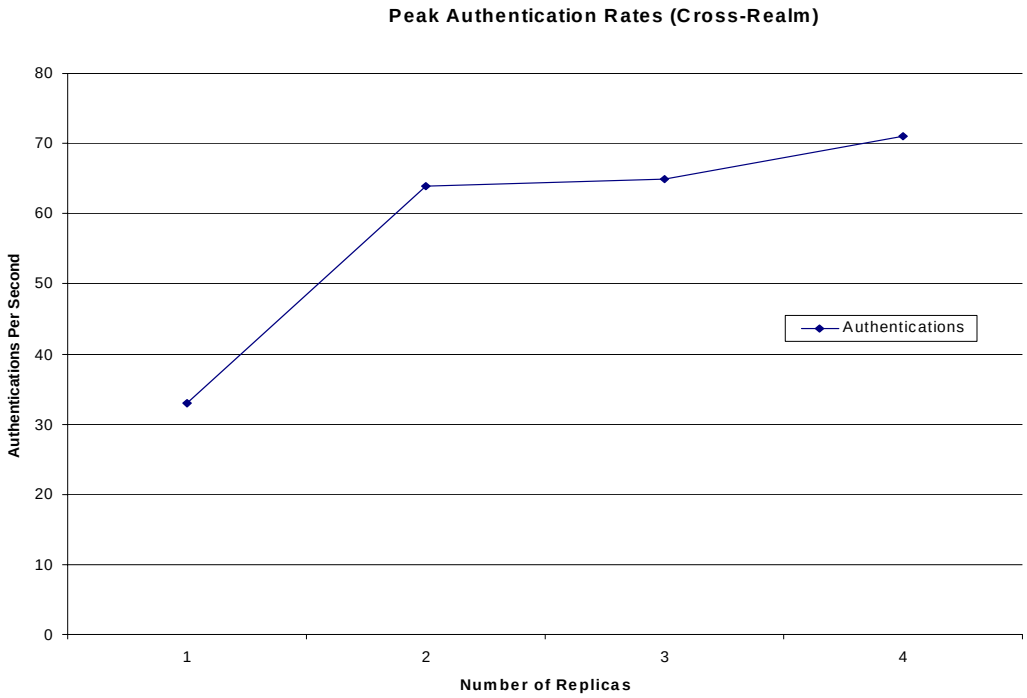

Analysis

The cross-realm authentication tests reveal the following:

- Using version 6.0 Agents, peak cross-realm authentication is approximately 33 APS when the remote realm is configured with one Replica and one Agent. With the addition of one Replica (and Agent), the cross-realm APS grows significantly, peaking at **71 APS** when the home realm is configured with four Replicas and four Agents.
- Very large organizations, and/or those with multiple, widely dispersed offices, can establish multiple realms. This provides improved in-realm authentication rates, 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 180 APS. • Total number of users is less than 250,000. • 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 connects through a WAN serviced by leased high-speed telecommunication 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 180 APS. • Total number of users is greater than 250,000. • 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 Primary/Replica model, the Primary functions as the administration Server, and periodically *replicates* database changes to each Replica.

As part of administering RSA ACE/Server, you can specify a *replication interval*, which is 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 is a small change that is quickly replicated to the other Servers. A larger change, for example, importing 100 new users from an LDAP database, takes longer to replicate. Replication is a background process and can be slowed by higher-priority processes (for example, authentication).

Server Hardware

Operating System	CPU	RAM
Windows 2000 Advanced Server	Dual-processor Pentium III (933 MHz)	1024 MB

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 the time needed to complete the replication pass once it has started.

Configuration: Primary and Single Replica

Task	User DB	Time
Replicate 100 New Users to DB	50 K	10 seconds
	100 K	10 seconds

Analysis

After initial installation and deployment of RSA ACE/Server 6.0, 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 increases the load. A longer interval reduces 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 is.

Database Push

In RSA ACE/Server 6.0, *DBpush* (database push) is a function that enables the administrator to copy the 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 fails.

Server Hardware

Operating System	CPU	RAM
Windows 2000 Advanced Server	• Primary: Single-processor 500 MHz Pentium III • Replica: Single-processor 400 MHz Pentium III	256 MB

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 50 K User Database	18 seconds
Push 100 K User Database	37 seconds
Push 500 K User Database	5 minutes, 9 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

RSA ACE/Server 6.0 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.

Server Hardware

Operating Systems	CPU	RAM
RSA ACE/Server: Windows 2000 Advanced Server	Single-processor 750 MHz Pentium III	512 MB
LDAP Server (Active Directory SP2): Windows 2000 Advanced Server	Single-processor 550 MHz Pentium III	1024 MB

LDAP Import Results

The following table provides the results of the LDAP import test.

Configuration: Primary and Single Replica		
Task	**User DB**	**Time**
Import 5K new users from LDAP	100 K	42 seconds
Import 50K new users from LDAP	100 K	23 minutes, 48 seconds

Analysis

Populating the RSA ACE/Server database by importing users from a supported LDAP database is a straightforward administrative process. 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, 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 needs to support depends on the size of your user population and the number of administrative personnel involved in supporting your users at any given time.

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

For further discussion of these settings, see Appendix A, "System Capacity and Resource Utilization." For procedures to adjust these settings, see the *RSA ACE/Server 6.0 for UNIX Installation Guide*.

Scenarios

For most installations, the default parameter settings built into RSA ACE/Server 6.0 are 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 Administration Toolkit (ATK) sessions would be in use.

With default settings, this type of installation can support 56 simultaneous local and/or web-based administration sessions, and 30 Remote Administration sessions, for a total of 86 simultaneous administration sessions. This is typically more than adequate to service a user population of 5000, therefore no changes to the default settings are necessary.

In addition, on a UNIX-based system, you can increase the amount of kernel memory available to the operating system by reducing kernel settings without affecting RSA ACE/Server operation.

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 comfortably supports 5,000 to 20,000 users.

With default settings, this type of installation supports 41 simultaneous local and/or web-based administration sessions, and 30 Remote Administration sessions, for a total of 71 simultaneous administration sessions. This is typically adequate to service a user population of up to 20,000, therefore no changes to the default settings are necessary.

In addition, on a UNIX-based system, you can increase the amount of kernel memory available to the operating system by slightly reducing kernel settings without affecting RSA ACE/Server operation.

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 needs to run as many as 100 concurrent Remote Administration sessions and 50 concurrent Quick Admin sessions.

Consequently, the system administrator has to make changes to the Progress Software 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, see Appendix A, "System Capacity and Resource Utilization."

Remote Administration Session Limits

To verify Remote Administration session limits, a 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 500.

The test setup included a Primary and four Replicas running on the dual-processor Windows 2000 machines described in the table on page 21. The user database was approximately 500K.

With multiple Server processes already running, new Remote Administration sessions were opened until a limit of **432 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 Remote Administration 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**, the maximum number of remote clients (**-n**) was set to different values. The results are described in the following table.

Value of -n	Number of Sessions
100 (default)	89
200	188
300	218
400	218

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 21. The user database was approximately 500K.

With multiple Server processes running, new Quick Admin sessions were opened until a limit of **218 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 Remote Administration 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 Remote Administration 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 is more than adequate to service a user population of 100,000 or more.

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

If an installation requires 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 following data is 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 750 MHz Pentium III	256 MB

Log Maintenance Test Results

The following table shows data related to maintenance (editing and compression) of large log databases in RSA ACE/Server 6.0.

Database Before Maintenance		Maintenance			Database After Maintenance	
Total Entries	**Size**	**Log Entries Removed**	**Remove Time (min:sec)**	**Compress Time (min:sec)**	**Total Entries**	**Size (KB)**
500 K	74 MB	50 K (10%)	3:45	23:36	450,051	67264
500 K	74 MB	500 K (100%)	24:00	00:03	14	448
1000 K	149 MB	100 K (10%)	07:06	44:58	900,046	133824
1000 K	149 MB	1000 K (100%)	49:38	00:06	10	2499

Analysis

A log database can grow to a very large size 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 passcode incorrectly a few times can add a dozen or more entries to the log. These and other typical transactions add up quickly. RSA Security recommends regular log maintenance. Otherwise, a failure could occur during a peak authentication period.

As the data shows, removing log entries and compressing large databases is 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 6.0 provides a log filtering feature that you can set up to slow log growth. Another tool provided by RSA ACE/Server 6.0 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, such as the number of:

- CPUs
- Disk I/O speed
- 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 Software RDBMS (relational database management system) parameter settings
- Operating system kernel settings

All these factors 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 Software RDBMS require.

The following information is 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 6.0 for UNIX Installation Guide*, January 2005
- 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 obtaining their documentation.

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

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

Users

The RSA ACE/Server 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 page 43 that contribute to a system's capacity, the first four items (CPUs, Replicas, LDAP synchronization jobs, and 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 default 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 calculates 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 starts. 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 Replica Server is 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 processes
- Quick Admin daemon (**sdcommd**)
- Automated Audit Log Maintenance (AALM) daemon
- Offline Authentication daemon (**sdoad**)

Estimating, our settings are based on using a value of 10 system processes.

Local_Apps

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 Software server processes. Progress Software configuration parameters control how these servers manage the remote user connections.

Progress Software 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 Software parameter (-n) controls the number of user entries each database instance supports. 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 Software RDBMS system capacities. Consult Progress Software documentation for a more complete description of the settings available in the parameter file.

Warning: Do not change this file 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 database servers (processes that handle remote client connections)	6 (specified)
-Ma	Maximum number of remote clients per database server (connections per Server)	5 (Progress Software default)
-Mi	Minimum number of connections per Server	1 (Progress Software 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 Software default)
-bibufs	Before-Image Buffers	7 (UNIX-only, specified) 5 (Windows-only, Progress Software 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 connection is handled by an existing server. Parameter file settings control the maximum number of servers ('-Mn') and the maximum number of database connections one server supports ('-Ma'). The total available Remote Client connections are:

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

RSA ACE/Server specifies the number of servers as six. By default, Progress Software allows up to five 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 Software 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
	semmns	123	500
	semmnu	123	500
	semmsl	41	200

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 Software 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 Software recommends using the following equation 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 Software is 16 MB (16,777,216 bytes). Almost all operating systems currently in use come with a default shared segment size of 64 MB.

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

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

Semaphores

Maximum Number of Semaphore Identifiers - semmni

This parameter specifies the maximum number of semaphore identifiers allowed for the system. Progress Software 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.

Maximum Number of Semaphores Per Identifier - semmsl

This parameter specifies the maximum number of semaphores allowed per semaphore identifier. Progress Software 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 semaphores is recommended.

Total Number of Semaphores - semmns

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

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

For the RSA ACE/Server this 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 Software recommends that this value be equal to the total number of semaphores (**semmns**).

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

General UNIX 'ulimit' Command

This command can be used to specify process-specific resource limits. It is commonly used to alter the maximum number of file descriptors (concurrently open files) that the system allows 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 must 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 does not occur using an RSA Security application. If this occurs with internally developed applications, verify that files opened by the application are being closed under all conditions.

B Authentication Performance Test Data

The peak authentication graphs starting on page 22 are 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, 500 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 successfully handled by the RSA ACE/Server. Therefore, one TPS is equivalent to one authentication request per second, or APS.

Peak Authentication Data - Windows (Local Authentication Client)

The following table corresponds to the graph in "Peak Authentication on the Windows 2000 Test System (Local Authentication Client)" on page 23.

Number of Replicas	Number of Agents	TPS			
		Days of Offline Data: 0	1	3	30
1	500	66.7	13.4	6.9	3.2
2	1000	133	28.2	14.4	6
3	1500	159.7	43.9	20.9	8.8
4	2000	178.8	55.4	26.9	11.5
5	2500	167.9	73.7	35.4	14
6	3000	166.5	82.1	40.4	17.6

Peak Authentication Data - Windows (Domain Agent Client-Domain Agent Host)

The following table corresponds to the graph in "Peak Authentication on the Windows 2000 Test System (Domain Agent Client-Domain Agent Host)" on page 24.

Number of Replicas	Number of Agents	TPS		
		Days of Offline Data 0	1	3
1	500	37	14	7
2	1000	71	29	14
3	1500	104	40	24
4	2000	130	52	29
5	2500	156	68	36
6	3000	159	81	40

Extrapolated Data

The following table corresponds to the graphs in "Effect of Authentication Load" on page 27.

Time to complete authentications for all users.

Number of Replicas	Minutes			
	10 K Users	50 K Users	100 K Users	1,000 K Users
1	12.3	61.7	123.5	1234.6
2	5.7	28.7	57.5	574.7
3	4.2	20.8	41.7	416.7

Time to complete delivery of offline data (One Day)

Number of Replicas	Minutes			
	10 K Users	50 K Users	100 K Users	1,000 K Users
1	13.2	66.0	132.1	1321.0
2	6.1	30.3	60.7	606.6
3	4.5	22.6	45.1	451.4

Time to complete delivery of offline data (Three Days)

Number of Replicas	Minutes			
	10 K Users	50 K Users	100 K Users	1,000 K Users
1	26.5	132.7	265.4	2654.3
2	12.5	62.3	124.5	1245.2
3	10.1	50.3	100.7	1006.9

Peak Authentication Data - Solaris

The following table corresponds to the graph in “Peak Authentication on the Sun Solaris Test System (Local Authentication Client)” on page 25.

Number of Replicas	Number of Agents	TPS			
		Days of Offline Data	0	1	3
1	500	27.7		3.2	2.3
2	1000	47.7		6.4	5.7
3	1500	69		10.2	10.7
4	2000	79		13.6	16.7
5	2500	89.5		17.0	21.4
6	3000	98.2		19.7	21.9

The following table corresponds to the graph in "Peak Authentication on the Sun Solaris Test System (Domain Agent Client-Domain Agent Host)" on page 26.

Number of Replicas	Number of Agents	TPS		
		Days of Offline Data 0	1	3
1	500	26.7	3.5	1.9
2	1000	50.2	6.9	3.6
3	1500	72.1	10.1	5.6
4	2000	96.9	13.2	7.1
5	2500	111.6	16.3	8.6
6	3000	116.5	19.1	14.5

Days of Offline Data Download

The following table shows performance numbers for downloading offline data, with and without an authentication load running. The test systems are the same as those used for testing peak authentication rates in the Domain Authentication Client-Domain Authentication Host environment.

Number of Replicas	Number of Agents	Days of Offline Data Per Second Windows		Solaris	
		With Authentication	Without Authentication	With Authentication (One Day of Offline Data)	Without Authentication (Three Days of Offline Data)
1	500	12	25	3.5	3.9
2	1000	27	36	6.8	8.3
3	1500	35	60	10.0	14.9
4	2000	50	101	13.2	18.9
5	2500	67	126	16.2	23.8
6	3000	81	143	19.0	37.1

Cross-Realm Authentication

The following table corresponds to the graph "Cross-Realm Authentication Test Results" on page 32. 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 100 K).

Number of Replicas	Number of Agents	TPS (100 K Database)
1	500	33
2	1000	64
3	1500	65
4	2000	71

Index

Q

R