Global Information Assurance Certification Paper

Exploring Sybase Adaptive Server Security Features
Rochelle Mooney
April 4, 2003

Introduction

With today's security threats, it's vitally important to secure an organization's data. Systems should provide an array of features to protect most effectively. These security features should be user friendly, flexible, and easy to administer. Sybase Adaptive Server Enterprise (ASE) provides security features that aid in protecting an organization's sensitive data from inappropriate access and unauthorized disclosure. Sybase adaptive server enterprise provides several security features that can be used easily to protect the organization's valuable information assets. The purpose of this paper is to give you an overview of the major security features that are available in Sybase Adaptive Sever. I will discuss these security features and show you how they are used to enforce a security policy. These security features are easy to use and administer. As a user it is important that you understand these security mechanism so that you can apply them effectively. As with any system, some basic guidelines will need to take place, after ASE has been installed on your system. With the initial ASE installation a powerful user login "sa" is created. This single login is configured with the roles of System Administrator and System Security Officer. The "sa" login is configured with a "NULL" password, therefore, for security protection; a password must be assigned immediately. Once your ASE system has been secured, you can then take advantage of others security features, such as assigning logins, enabling auditing, etc.

User Identification and Authentication

The most basic method of user authentication is the user's password. The "user name" or "login" is the identity; with the password being the authentication of that identity. In order to access an adaptive server, a user is a given a unique login account. Access to the server is granted once you have entered a valid user login and password. All activities performed on the server are attributed to your login id – which is subject to audit. The SQL server password must be at least 6 bytes or longer in length. Passwords are stored in encrypted format in the syslogins table under the master database. A unique login account is created, via the system administrator (sa) for a user's system access. A default database can be created for the login account so that the user can start each adaptive server session in that database without issuing an additional command (i.e. use database-name). Identification and authentication ensures that only authorized users can log into the system. The followings requirements, and not limited to, should exist for the most effective system user identification and authentication:

- Each user should have only one user "login"

- User password should be unique and not shared with other users on the system
- User password should be at least six characters long
- System should require periodic password changes for each user (i.e. every 60 to 90 days)
- System should require that users are unable to use expired passwords when the system forces a password change
- If you allow guest logins, the number of days should not exceed five
- System should limit the number of unsuccessful login attempts (i.e. a maximum number of three attempts is recommended)
- When entering a user's password, the password should not echo back to the user's screen
- Access to the system should be suspended when user is setting idle for a period of time (i.e. five to ten minutes)
- If a user forgets his password, the "sa" should issue a temporary password, and the user should be required to change to password when he first logins in

Division of Roles

It is sometimes necessary to maintain individual accountability within your system. ASE offers roles such as System Administrator, System Security Officer, Operator and User-Defined. These roles and to whom they are assigned is vitally important for security protection. Division of roles is an important feature in the SQL server. Users in SQL server are granted special operational and administrative roles. These roles determine what system administrative and security-related actions a user can perform; which is subjected to audit. Roles are given to each individual login accounts created on the SQL server. The various roles are system administrator, system security officer, and operator. The system administrator will perform administrative tasks that are unrelated to specific applications. The administrative tasks included managing disk storage, sever configuration, login accounts, granting permissions, etc. The system administrative can be more than one person; in fact, in most cases the role is granted to several individual login accounts. The system security officer is held reliable for security-sensitive related task in SQL server; for example, create server login accounts, managing the audit system, granting and revoking system security officer and operator roles, etc. The system security officer only has access to databases but only have privileges on objects within sybsecurity

databases and its tables. Operators can only backup and load databases on a server-wide basis. A single user is granted the operator role. With this role, the user can use the dump and load command to backup and restore all databases on a server. System roles are provided by the adaptive server and the system security officer create user-defined roles. System administrator and system security officer are the system roles. The security officer can define both role hierarchies and mutual exclusivity of roles. Role hierarchies are roles that have other roles; these types of roles are user-defined or included in the system. Mutual exclusivity of roles is independent of one another. A login account on a server can be granted any role and can posses more than one role. (6) Establishing user-defined roles make it easy to grant permissions to access database for all users in the organization. With user-define roles you have more control over security-related features of adaptive server.

Discretionary Access Controls

Discretionary access controls (DAC) are a very important security feature within ASE that offers access flexibility based on the users identity. Discretionary access controls are controls that are used at the discretion of object owners. They are "discretionary" because an object owner can choose to allow you to access an object or can disallow such access. (6) Discretionary access controls are enforced by using grant and revoke commands. Permissions can be granted to users, groups, and roles. To take away the privileges, a revoke command is used. Permissions are usually granted to create databases, create objects within a database, access tables, views, and columns, and execute stored procedures. There are two kinds of database permissions that can be granted or revoke: object access permission and object creation permission. Object access permissions are used to regulate the access to certain database objects. For example, in order for a user to have access to a table in a database, the user must be granted object access permission. A user can't select, update, insert, delete, reference, or execute an object unless the user has been granted object access permission. Object create permissions are used to regulate the ability to create objects. The system administrator or database owner grants this type of permission. A user can't create a database, default, procedure, rule, table or view in a database unless the user has object create permission. ASE's discretionary access control system recognizes certain types of users that include; System Administrators, System Security Officers, Operators, Database owners, Database object owners, Public users.

Auditing

Auditing provides information that describes the basic transaction cycles of the system, thus ensuring accurate and reliable information processing. Auditing also identifies and describes how the data is captured, processed, stored and reported. Among other things, computer fraud, concepts, principles, and methods used to commit fraud are also identified with auditing. The Sybase Adaptive Server auditing function provides the ability to record the occurrence of many events, such as logins,

logouts, database access, and administrative actions. This auditing feature is powerful in determining who performed specific actions or detecting malicious or unauthorized activities. Auditing can be installed manually by using a sequence of steps, but it's highly recommended that you use the Sybase provided installation program, "auditinit" for installation. (7) Auditing must be properly installed, turned on and configured to audit the appropriate events before it can work effectively. Auditing is used to ensure the accountability of events on the system. Events that occur on a SQL server can be audited. An audit log will contain the date, time, the user login, and the success or failure of the event. The following events can be audited; logins and logouts, server boots, use of data access commands, attempts to access particular objects, and a particular user's actions. The system security officer uses the audit trail to evaluate the impact of events. Having the auditing system in place and operational can deter computer system fraud.

Policy-Based Access Control

Policy-based access control framework provides a powerful and flexible way to protect data, all the way down to the row level. With this capability, administrators have the ability to define security policies that are strictly relevant to the value of individual data elements. The server is used to transparently enforce the security policies. Once the policies are defined, they are automatically invoked when the affected data is queried, regardless of how the data was queried. Data can be queried via an application, ad hoc query, store procedure, or view. Policy –based access control simplifies the security administration of an ASE installation as well as the application development process. This is possible because the server is used to enforce security. Developers can now focus on implementing business functionality because the administrators will define a security policy that can be enforced consistently across the entire server. Which is achieved through the combined capabilities of access rules, login triggers, and domain integrity rules.

Access Rules

Access rules are essential building blocks of policy-based access control because an access rule is bound to a specific column and is invoke on any select, update, or delete operation on a corresponding table. Access rules can be used to define user-defined data types. Applying an access rule to a user-defined data type will activate tables that have columns with that data type. The outcome of an access rule is either true or false which is used to determine which rows to return to the user. When the condition is met, the row is returned to the user. If the condition is not met, a row is not returned to the user. An access rule can also be bound to many different columns in a table and is combined with either "AND"or "OR". Each access rule is designated as either an AND rule or an OR rule. All of the AND rules will need to equal TRUE in order for a row to be returned to a user. But only one of the OR rules will need to equal TRUE for a row to be returned to a user. AND is the

default designation for an access rule and provides greater flexibility when defining a set of access rules that address the business needs; even the complicated ones. Access rules can be used in java functions as well. Access Rules can embody user-defined functions written in Java that can look up information in JDBC-accessible databases and LDAP=compliant directories, and then use that information as the basis for making security policy decisions. (1) Which means that access rules can be based solely on dynamic and existing user account information. Database owner (DBO) or table owner (TBO) controls the rows users can access by granting/revoking privileges at the row level. Control is accomplished through "Access rules". Access privileges can be given to individuals, groups, or roles.

Login Triggers

Login triggers are ASE store procedures that execute automatically in the background during a login process. Login triggers can be used for anything that you normally would perform in a store procedure. Login triggers can be used to enforce account usage policies. For example, if you want to restrict users from logging in the system outside of regular business hours, you can create a login trigger. Login triggers are also useful in the area of employee turnover since there is a delay between an employee departure and when the account is removed. Login triggers do not exist until you set them up using sp_modifylogin. The command must be executed in the login's default database. Keep in mind that the stored procedure must exist in this database as well. After configuring a store procedure as a login trigger, you cannot drop the store procedure until it has been reconfigured. You will have to drop the login trigger first and then change the login trigger to another stored procedure. The command to drop the login trigger is sp_modifylogin my_login, "login script", NULL. The command to change the login trigger is sp_modifylogin, "login script", new procedure. The login trigger object ID is stored in the column syslogins.procid. You must have sso_role to set, change or drop a login trigger. The command to display the current login trigger is sp_displaylogin my_login. Once a login trigger is configured, it is automatically executed in the background. The output of a login trigger is written to the ASE error log file. It is not a good idea to use a login trigger on the "sa" login, and to do extensive processing. Creating a login trigger on the "sa" login can lock you out of the ASE server because if it is effectively locked by a failing login trigger. If a process takes longer than a few seconds, it has the risk of creating a block or can cause a deadlock. Avoid using extensive processing in a login trigger.

Domain Integrity Rules

Domain integrity rules are existing ASE server-enforced integrity mechanisms that can be used in conjunction with access rules to provide security policy control over the flow of information into and through the server. (1) Domain integrity rules are bound to columns and is invoked during an update and insert operation. Data

entered into a table must be appropriate to the columns it is entered into. The validity of a domain may be as broad as specifying only a data type (text, numeric, etc.) or as narrow as specifying just a few available values. Different database engines will have varying means of enforcing the validity of the domain of a column, but most will permit the entry of at least a simple Boolean expression that must be true for the value in the column to be accepted. The domain of an entry may be dependent on two or more columns in a table. (4) Domain integrity rules can be written with the use of JavaTM user-defined function. This type of function is used to access other databases and directories that support information regarding security policy decisions. However, when working with domain integrity rules it is necessary to comprehend the design of the database, the type of expected data, and the rules that applies to it.

Application Context Facility

Application context facility provides a secured environment for accessing data. Application on a database server must limit access to the data. Applications are carefully coded to consider the profile of the user. For example, a Human Resource Application is coded to know which users are allowed to update salary information. (5) Application allows greater access to sensitive data than when accessed via ad hoc query. Application context facility is composed of four built-in functions used to provide a secure environment for accessing data. The application context facility consists of context name, attribute name, and attribute value. "Context" is the name of a context that is user defined and can be used by one or more applications. (1) A separate security context can be used by each application or you can use the same security context for a set of applications as long as they have similar access control requirements. Attributes are variables used in access rules and domain integrity rules. Attributes are assigned a specific value when setting up a context. Contexts are set-up on a session-by-session basis and so allow security policy to be based on properties of both application and the user invoking the application. (1) Users are responsible for defining the context name, attributes, and the values for each context. Users are able to define various contexts and various correlating attribute/value pairs for each context. With application context facility, users can define, store, and retrieve users profiles and application profiles currently being used. Context is not persistent across session – it is only specific to a session. ACF is defined in a table created by the system administrator. There are several drawbacks to this approach. The first is ensuring that a user can't bypass the application and access the data directly. The second is giving the application developer the additional burden of incorporating security checks into the application itself. (1) ASE's application context facility supports application-specific policies which are enforced in the sever itself. Therefore, they cannot be changed and updated without making changes to the application.

Secure Socket Layer Encryption

ASE provides SSL encryption, which is used to protect the confidentiality and integrity of the entire client-server session. A SSL can be invoked on a session-by-session basis. This will allow the level of protection to be matched with the requirements of the application and environment. No administrative work is needed; only the purchase of a digital certificate is required. The digital certificate is used to authenticate the ASE server to its clients. With the use of both password-based user authentication and sever certificate authentication, mutual authentication is established for all client-sever and sever-server connections. Sybase's SSL support is the strongest that can be found in the world, using encryption implementations provided by Certicom. (1) There are several levels of security involved in the implementation of secure socket layer encryption on adaptive server. First, the server authenticates itself, which is a process used to determine if it is the server you intended to contact. An encrypted SSL session will begin before data is transmitted. Second, once a SSL session has been established, the client can send his user name and password over the secure, encrypted connection. Third, comparing the digital signature on the server certificate can be used to determine if the data received by the client has been modified before it reaching the intended recipient. The SSL protocol is implemented as a filter appended to the master and query lines of the interface file on the adaptive server. The address and port numbers are configured to accept connections specified by multiple networks, different protocols, and alternate ports.

Row-Level Access Control

Row-level access control provides several features to automate the control of accessing data at the row level. Row-level access control is often referred to as "data security policies". There are three features of Row-level access control: access rule, application context facility, and login trigger. Database Owners and table owners can restrict access to a table's data rows by defining access rules and binding those rules to the table. Access to data can be further controlled by setting application contexts and creating login triggers. With row-level access control, database owner or table owner controls the rows in a table that users can access based on their identification or profile and the privileges the user has from the application level. The benefits of row-level access control includes:

- The ability to set permissions for individual rows.
- Automatic data filtering according to group, role, and application
- Data-level security encoded in the server
- Increase flexibility: for example, if a user responsibility changes, an administrator can change the user's group, role, or application context to affect data access

- More flexibility than views.

Adaptive Server enforces row-level access control for all data manipulation languages (DMLs), preventing users from bypassing the access control to get to the data. (8) The following command is used to configure your system for row-level access control:
sp_configure "enable row level access", 1. When row-level access control is used, adaptive sever memory usage increases. Row-level access control is a dynamic option; therefore, there is no need to reboot the SQL Server after enabling row level access.

Conclusion

As information technology increases, the importance of security continues to become even much more critical. Sybase Adaptive Server Enterprise meets the need for a wide-range of security features for any organization, from the least to the most critical applications. ASE's security features provide a most safe, secure and flexible environment to ensure that vital data is protected. Each of the security features discussed in this paper can be utilized within ASE independently or in combination. The first feature addressed "User Identification and Authentication" is not optional and must be implemented to gain access to ASE. The next feature "Division of Roles" is optional and will allow for more control over the security related operations to individuals. "Discretionary Access Controls" feature is also optional and at the system administrator's discretion, other system users can be given permissions to perform system type functions. The "Auditing" feature, which is optional, is an important feature and will monitor the integrity of system operations among users, thus deterring fraudulent activity. The "Policy-Based Access Control" feature is optional and will allow the security control to operate at yet a more granular level for protecting data. "Access Rules" are optional and work in conjunction with the previous feature "Policy-Based Access Control", therefore allowing more robust to security controls. "Login Triggers" another security feature; it's seldom used but is convenient when performing system maintenance. "Domain Integrity Rules" is an optional feature and can be used in conjunctions with of integrity rules, such as domain. "Application Context Facility" is optional, and is very convenient when tying to control the security of various users within the organization. "Secure Socket Layer Encryption" is optional and is a very powerful feature, which encrypts, therefore protecting system confidentiality and integrity to the highest. The last feature addressed in this paper "Row-Level Access Control" is optional, works in conjunction with the three previously mentioned features (access rule, application context facility, and login trigger) and often used when implementing data security policies. As you can see, ASE offers a variety of security features to aid in protecting an organization's data. These features will allow any organization to provide protection to the utmost. With the ease of use, the features will allow for one to implement security for its system/data in a minimum amount of time with little training, thus providing required security in a reasonable time.

References

1. “New Security features in Sybase Adaptive Server Enterprise” A Sybase Technical White Paper, 2001, URL: http://www.sybase.com/detail?id=1013009

2. “Login triggers in ASE 12.5” URL: http://www.sypron.nl/logtrig.html

3. “Data Integrity” URL: http://admin.nj.devry.edu/~kjudge/DataIntegrity.htm

4. “Data integrity” URL: http://216.239.39.100/search?q=cache:3eulyAXTLYYC:www.cf.ac.uk/chemy/msc/access2.doc+%22domain+integrity+rules%22&hl=en&ie=UTF-8

5. “Using the Application Context Facility” URL: http://manuals.sybase.com/onlinebooks/group-as/asg1250e/sag/@Generic__BookTextView/38861

6. “Sybase Adaptive Server Enterprise Security Features User Guide”, Sybase, Release 11.5.x.

7. “Sybase Adaptive Server auditing configured incorrectly” URL: http://www.iss.net/security_center/static/3644.php

8. “Row-level access control”. URL: http://manuals.sybase.com/onlinebooks/group-as/asg1250e/sag/@Generic__BookTextView/38430