iDAR FAQs, Features, and Troubleshooting

This document contains useful information on iPlanet Directory Access Router (iDAR). It contains answers for frequently asked questions (FAQs), clarifications on certain iDAR features, and troubleshooting information.

The document contains the following sections:

iDAR FAQ

What is iDAR?

iDAR is an LDAP proxy for LDAP clients and LDAP servers. Requests from LDAP clients are forwarded to LDAP servers based on rules defined in iDAR's configuration. Results from the server are passed back to clients, again based on rules defined in the configuration. This process is totally transparent to the clients, which connect to iDAR just as they would to any LDAP server.

Why do I need an LDAP Proxy Server?

Many enterprises want to make some part of their directory information externally visible, while keeping other parts internally private. With iDAR you can accomplish this goal easily, and without assigning directory passwords to external clients. iDAR can also be used as a high availability solution for enterprise directory service with load balancing and failover features.

Additional security features such as protection from denial of service attacks and search limits are also provided.

What version of the LDAP protocol does iDAR support?

iDAR supports LDAP clients or chaining LDAP servers that use either the LDAPv2 or the LDAPv3 protocol.

Does iDAR support secure authentication and encryption?

iDAR supports SSLv3 services for public-key based data encryption using certificates. Secure authentication and encryption available to LDAP clients can either use the secure LDAP port or the Internet Transport Layer Security (TLS) model, which uses the Diffie-Hellman, Digital Signature Standard (DSA), and Triple-DES algorithms.

Does iDAR work with any LDAP-enabled directory server?

iDAR will work with any LDAP-conformant directory server. Some directory product vendors claim to implement LDAP in their marketing literature, but the reality is often a different story. iDAR has been the most thoroughly tested with the iPlanet Directory Server.

If the iDAR 5.0 Console is used, the iPlanet Directory Server 5.0 is the supported configuration depository.

Is there a configuration utility available to configure iDAR?

iDAR 5.0 includes a Java-based GUI (console) that can be used to configure iDAR. The console uses the iPlanet Directory Server to store the configuration it generates.

iDAR Features

Can iDAR be used to prevent denial-of-service attacks?

Yes. You can limit the number of simultaneous operations processed per connection, number of operations allowed per connection, total number of concurrent connections, maximum concurrent connections per defined group (network, subnetwork or based on bind DN), and maximum concurrent connections for a single IP address.

Does iDAR support “reverse” proxying?

In a strict sense, iDAR is a reverse proxy; however, the LDAP protocol does not support the concept of reverse proxying.

Can iDAR be used to prevent “trawling” of an LDAP directory?

Yes. Trawling refers to very broad queries designed to download large portions of your directory, a practice many sites wish to prohibit. iDAR can prohibit or limit trawling in a number of ways:

- The scope of searches can be limited to a single level of the directory tree, entire subtrees can be hidden, and a hard limit on the number of entries returned in response to a query can be set.
- Inequality searches can be forbidden, thus disallowing searches that return many results based on exclusion and substring searches can be restricted by length; for example, prohibiting searches for all entries with a surname starting with the letters A-Z.
- iDAR can also be configured to deny un-indexed searches. Un-indexed searches are inefficient and can possibly have a negative impact in performance.

Does iDAR do automatic load balancing of queries?

iDAR supports automatic server load balancing among a set of back-end LDAP servers. iDAR also supports automatic fall-over to a secondary LDAP server if the primary LDAP server is down.

How many LDAP servers can one iDAR server load balance?

The performance needs of the directory server and the complexity of work being done by iDAR determines the optimal number of directory servers that iDAR should load balance. For example, if iDAR is doing complicated work, such as attribute renaming, the number of directory servers iDAR is configured to load balance should be reduced. Consider adding more iDAR units to compensate for possible performance impacts of complex iDAR configurations.

Can search requests be filtered?

Yes. You can configure iDAR to refuse searches that attempt to search on a particular attribute. In addition, you can configure iDAR to modify incoming search requests to conform to a designated minimum search base, search scope, and time limit.

Can search results be filtered?

Yes. Results can be filtered both in terms of number of entries returned and the attributes that are included in the result set. Search result entries can also be filtered based on the entry DN or content.

How are access groups defined?

Varying levels of access to the directory are provided to clients based on the network address of the client. Thus, different levels of access can be granted to clients outside the corporate firewall, inside the firewall, on the executive subnetwork, and even to individual machines. Further, access level can be changed upon a successful completion of a LDAP Bind operation by the client or when a SSL session is established.

Does iDAR support protected password authentication?

Yes. Through the use of the SASL mechanism a variety of protected password authentication schemes can be implemented. These mechanisms must be supported by the back-end directory server. iDAR does not support SASL mechanisms with connection protection and SASL EXTERNAL mechanism.

Does iDAR automatically follow referrals?

The following of referrals is configurable based on access group. Various access groups can be configured to automatically follow referrals, return referrals, or discard referrals.

Does iDAR cache search result information?

iDAR 5.0 does not support search result caching.

Can iDAR do attribute renaming?

iDAR can transparently rename attribute names between clients and servers.

Troubleshooting

How can I analyze logs of connection attempts?

iDAR can be configured to either use `syslog` or write to a specified log file. A popular UNIX utility known as swatch is freely available from Stanford University (`ftp://ftp.stanford.edu/general/security-tools/swatch`). Swatch can be used to monitor the log files generated by iDAR and to notify the administrator when defined events occur.

I have configured iDAR to follow referrals. However, when I perform a search with a LDAPv2 client I get error 32 (No such object) or some other error.

In order for iDAR to receive referrals from the back-end servers, it must use LDAPv3. Make sure you have selected “LDAP version 3 only” on LDAP version to use selection for each of your LDAP server properties.

iDAR is load balancing across ‘n’ servers. When one of these servers go down, some of my clients seem to hang.

Make sure that the server in question is reachable from the host iDAR is running on. An unreachable server (caused by bad network or a hung server) can cause iDAR to wait on a long timeout. This may cause the clients that are virtually connected to the unreachable server to appear to be hung. However, a server that is not running should be detected quickly via “Connection refused.” If clients still appear to be hanging under these conditions, then you should verify that ICMP messages can be exchanged between the host on which the server is supposed to be running and the host on which the iDAR is running. A firewall may block these messages.

I notice in the log files that some idle client connections are routinely failed over even though all my back-end servers are up.

Your back-end directory server is timing out idle connections and closing them. iDAR fails over these closed connections. You must set an idle connection timeout for iDAR as well. This will clean up idle and leaked client connections and also guard against one form denial of service attack.

Is there a way to restrict search requests containing the presence filter?

iDAR 5.0 does not have any direct mechanisms to restrict clients from using the presence filter. There are two indirect ways to address this issue.

You can set the `ids-proxy-con-forbidden-compare` to the name of the attribute that you do not want to be compared. This method is over restrictive, as it will reject searches containing both (mail=*) and (mail=Andy*) filters.

On the other hand, since presence filters (attrName=*) always generate the same result (assuming the data did not change), we can use the ids-proxy-con-size-limit attribute and the `ids-proxy-sch-SizeLimitProperty` to limit the damage. Although, LDAP does not require entries to be returned in a given order, under most (all) implementations, the set of result will either be returned in sorted order or unsorted order, and these will be the same every time. Hence, if iDAR is configured with a size limit, (using the size-limit attribute or the SizeLimitProperty) only the first 'n' of these sets will be returned every time. Because there can only be two sets of these 'n' entries, the risk of trawling the directory is greatly reduced.

Note that iDAR tries to set this size limit in the request itself when possible, and therefore the directory server will not be burdened with sending all the entries.

The size limit property gives you the option of applying exceptions to size limits imposed when necessary. Suppose, for example, that you have an entry of o=A, under which there are 400 organization units. Under each of those OU's there are people. If you want clients to see all the OU's but only see 5 people at a time, you can set up the `SizeLimitProperty` such that no limit is applied for a search with base o=A and one level scope. For all other searches a limit of 5 applies.

I am using JNDI. iDAR rejects some (not all) of my search queries saying it could not decode search request. If I send the query directly to my directory server, it works fine.

iDAR is more strict about attribute type as defined in LDAP standards. If you are requesting specific attributes to be returned as part of that search result, make sure that the attribute type names conform the LDAP standards stated in RFC 2251 section 4.1.4, paragraph 3.

"A specification may also assign one or more textual names for an attribute type. These names MUST begin with a letter, and only contain ASCII letters, digit characters and hyphens. They are case insensitive. (These ASCII characters are identical to ISO 10646 characters whose UTF-8 encoding is a single byte between 0x00 and 0x7F.)"

JNDI is known to send "" strings attribute names as part of attribute list to be returned for a search. As stated above empty strings are illegal LDAP attribute type names.

When I try to execute a task or perform some console function, I get an error message saying I need to make sure the Administration Server is running properly and that this host is permitted to connect to the Administration Server.

Log in to the Administration Server that is managing the iDAR whose console produced the error messages. It may be necessary to start the iPlanet Console on the host machine of the Administration Server. Open the server console for the Administration Server that is managing the iDAR on which you are unsuccessfully trying to invoke tasks. Click the Configuration tab and then the Network tab. Under Connection Restrictions, make sure that the host machine of the iPlanet Console that is unsuccessfully trying to manage iDAR is not restricted from accessing the Administration Server. See the iPlanet Console *Server Management Guide* for more information.