White Paper:

A Scalable and Distributed Solution

Identity Management and XML Directory Services Solution

Credentials

ViewDS Identity Solutions is the owner and developer of the ViewDS Directory Suite that is deployed in Government, Defense, and private sector companies throughout the world. This technology has some unique features that deliver superior performance, reliability and operability in large distributed environments. This document discusses the ViewDS Directory Server capabilities to operate in a distributed and replicated environment.

Hierarchical Directory Information Tree (DIT)

By definition, a directory is a hierarchical database, as opposed to a relational database. Unlike many other vendors who decided to place a directory front end over an underlying relational database, ViewDS Directory development was based on a purpose-built hierarchical database.

Support for a hierarchical DIT allows information to be organized in a manner that represents a real-world structure. For example, in a PKI based environment, certificates are named using hierarchical X.500 distinguished names. Consequently, ViewDS will allow certificates to be stored in a location that matches their name.

While other directory vendors would prefer flat hierarchies, ViewDS places no restrictions on the depth of the DIT. Due to underlying support for X.500 and a hierarchical database, ViewDS does not suffer any performance degradation with deep DITs. Since many other directory vendors cannot perform the move or rename operation on non-leaf entries (due to the use of relational databases) ViewDS is completely functional with deep DIT environments.

Distributed DIT

ViewDS support for hierarchical directory trees extends into the X.500 Distributed DIT model. ViewDS supports the ability for different segments of the DIT to be mastered by different directory server instances.

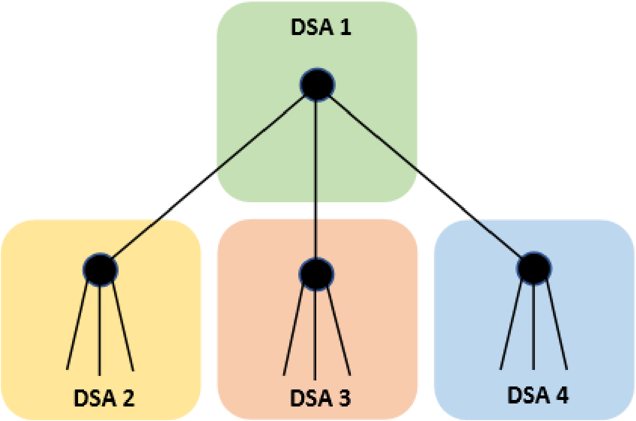

There is no limit on the number of different servers that are present within the distributed environment and each server can work cooperatively to service a client's request. DSA to DSA communications is achieved through the support of DSP – The Directory System Protocol, which allows a DSA to interact with other DSA's in order to service a request on behalf of the client.

DSP is the only standardized method that provides an interoperable method for multiple DSA's to cooperate in a distributed environment.

ViewDS has been participating in such interoperability trials as far back as the 997 Directory Challenge3 and more recently in the JWID and CWID defense interoperability demonstrations.

With the exception of hardware restrictions, the DSP protocol and support for hierarchical DITs allow ViewDS to scale to a virtually unlimited number of servers and entries.

Distributed Schema and Access Control Management

ViewDS support for large distributed environments is not limited to scalability and operability. The management of the directory data is also provided in a scalable manner. Schema and Access Controls are defined in ViewDS according to the rules specified in X.500. They are not stored in files or bundled in a single high-level entry. Rather, they are stored within the DIT at applicable locations. Schema and Access Controls are defined at points in the DIT and are applicable to subordinate entries until new schema or access controls are defined.

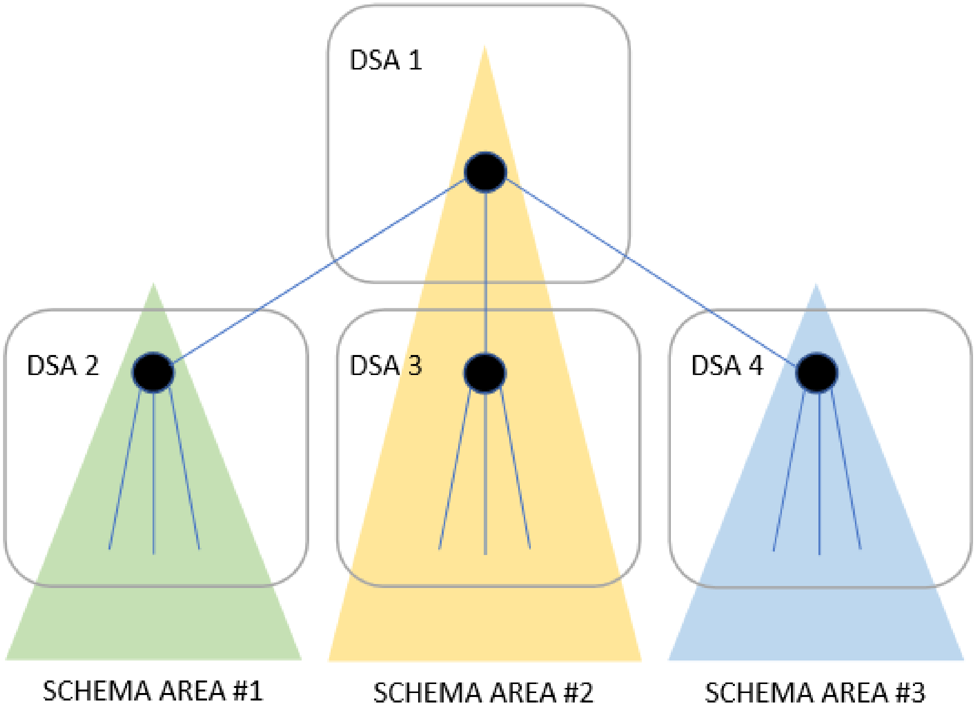

The ability to maintain schema in the DIT results in the efficient management of the schema in a distributed environment. In the image above, DSA #3 doesn't need to define new schema, since the distributed data model allows it to automatically retrieve it from its parent, DSA #1.

Alternatively, in the case of DSA's #2 and #4, it is possible for them to define new schema that is specific to their portions of the tree. Within a distributed schema environment, a top level DSA, such as DSA #1, need not be burdened with the unmaintainable task of managing schema for the entire DIT. Access Controls work independently of the schema but are defined in a similar fashion.

Replication in a distributed Environment

Within a distributed environment, a replication strategy can be employed allowing a copy of the data that is mastered within one ViewDS server to be replicated to another ViewDS server.

The replication is completed using the X.500 DISP protocol, which is the only standardized method for performing data replication amongst directory servers. Replication can be used to provide failover, to ensure that data is still available for other servers to query in the event that a directory server becomes unavailable. Replication can also be used to enhance performance by making data available locally within a DSA, allowing it to perform operations locally.

ViewDS servers can hold both mastered data and replica from an unlimited number of supplier DSA's.

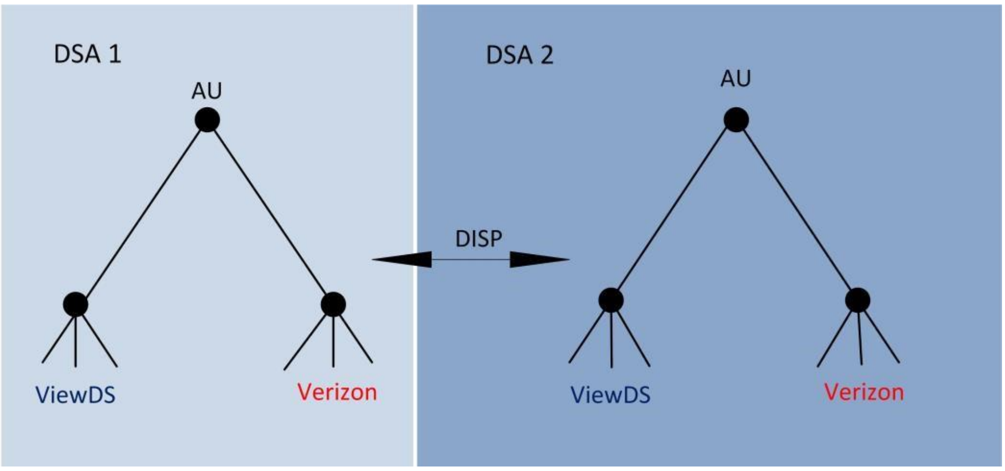

In the image above, both DSA 1 and 2 master data, and provide a replica to each other, thus resulting in each server having a copy of the other's information.

In secure environments, DISP is used to replicate less sensitive information from an internal directory server, into a server that is outside of the organization's secure domain. The ViewDS directory server's DISP implementation can be configured to only replicate a partial set of the data.

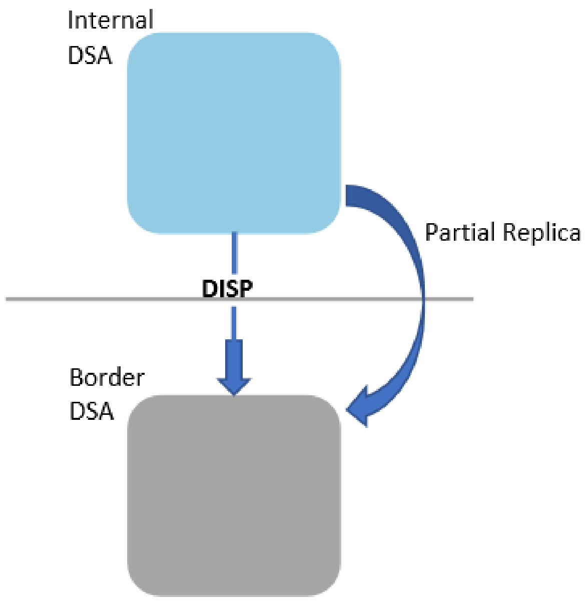

In fact, this scenario has been further tested and deployed in several defense projects such as CWID using the Clearswift EAL4 (Common Criteria evaluated) Directory Bastion solution.

Directory Bastion is a specialized standalone product, which is a new addition to the Clearswift Bastion product family. Its purpose is to allow X.500 Directory data to be synchronized between two X.500 Directory System Agents (DSAs) on two otherwise disjoint networks, while maintaining an assured separation between the two networks it interconnects.

The Directory Bastion ensures that the only communication that it allows to traverse between the two networks conforms to the Directory Information Shadowing Protocol (DISP, defined in ITU-T Rec. X.525) between explicitly identified DSAs.

Because DISP is the standard protocol to synchronize directory data between DSAs, a Directory Bastion can be inserted between two DSAs without requiring anything other than normal shadowing agreement configuration on the DSAs. Apart from network level addressing, the Directory Bastion is entirely transparent to the DSAs.

Scalable Certificate Storage

LDAP or X.500 directories are frequently utilized in X.509 Public Key Infrastructure (PKI) applications as repositories for certificates and certificate revocation lists (CRLs). There are three approaches PKI applications can take to locate certificates and CRLs of interest that have been stored within a directory.

1. The PKI application fetches the certificates and CRLs from all the likely candidate entries and filters the returned results itself to identify the data of interest.

2. When storing a certificate or CRL in the directory, the PKI application pulls out the values of certain component parts of the certificate or CRL and stores these values as separate attributes alongside the certificate or CRL. This is sometimes called the extracted attributes approach. For example, the serial number from a certificate is extracted and stored as a separate serial number attribute in the same directory entry as the certificate. The serial number attribute can then be used in a search filter to find an entry containing a certificate with a particular serial number. Since the serial number attribute has a primitive syntax (i.e., integer), all off-the-shelf directory implementations can handle indexing and searching of such an attribute. In a typical deployment, a number of components of certificates and CRLs would be extracted as separate attributes.

3. The directory client uses PKI matching rules to directly match certificates and CRLs having particular values for nominated components. For example, the certificateMatch matching rule can be used in a search filter to find an entry containing a certificate with a particular serial number. A directory server that supports these matching rules can handle in-situ indexing and searching of the various component parts of certificates and CRLs, without those components needing to be extracted as separate directory attributes. This means that the client is freed from the requirement to provide and maintain the extracted attributes.

The first approach can be very inefficient if large numbers of certificates or CRLs must be examined. The extracted attributes approach and the PKI matching rules can have comparable performance, however the extracted attributes approach has a number of disadvantages compared to using the PKI matching rules. The extracted attribute's is a much less scalable approach due to the fact that it requires much more information to be stored within the directory, as well as extra controls and safeguards to ensure that the extract information is kept synchronized with the actual signed certificate content.

ViewDS provides support for all three of these methods and uniquely provides support for Component Matching. While PKI Matching rules impose some restriction on the aspects of a certificate being matched, Component Matching support allows ViewDS to evaluate filters based on any arbitrary content of the certificate.

For example, using component matching, the LDAP filter to find an entry containing a user certificate that has, for example, bob@eb2bcom.com as a subject alternative name would be:

```
(userCertificate:componentFilterMatch:=item:{       component
"toBeSigned.extensions.*.extnValue.
content.(2.5.29.17).*.rfc822Name",
rule caseIgnoreIA5Match, value "bob@eb2bcom.com" })
```

Through the ViewDS ability to understand and search the underlying structure of PKI information it provides a fast, scalable and flexible environment for PKI data storage.

Conclusion

ViewDS provides a high quality directory service that leverages standardized methods for data storage, discovery, retrieval, management, security, distribution and replication. The robustness and integration capability of ViewDS is the result of decades of development effort based on a design goal of providing a highly available and scalable service.

ViewDS provides Scalable Data Storage by allowing data to be stored in either a flat or hierarchical manner without imposing any limitations on the number of entries, volume of data or the depth of the DIT. Utilizing a hierarchy allows schema, access control and data to be maintained and distributed in a scalable and modular way.

ViewDS facilitates Deep DITs and their inherent volatility by providing Scalable Data Management through its support of the core LDAP and X.500 move and rename operations on leaf and non-leaf entries.

Management of data held within the DIT is provided through Scalable Schema and Access Control Management, through the use of a hierarchy and conformance with the X.500 data models. Allowing multiple schema and access control administrative areas throughout the DIT (including distributed and replicated environments) provides a modular and scalable method for managing schema and access control in large deployments.

ViewDS can operate in a Scalable Distributed Environment through its support for the X.500 DSP protocol – the only standardized directory protocol for distributed operations, whose proven capability is highlighted by its usage in military, defense, aviation and government environments.

ViewDS also fully supports the storage, searching and retrieval of XML data within the directory, such as XACML security policies. The eXtensible Access Control Mark-up Language (XACML) is an XML dialect for the server-side representation of access control policy and access control decisions. These rules can be expressed in an application- independent manner, making it extremely versatile. XACML polices can reference other policies and can intelligently combine policies with competing or overlapping rule sets.

Therefore, because both Identity data and XACML based policies can be combined in the ViewDS Directory, they can also be replicated either partially or fully in a highly distributed and efficient manner. This is of critical importance in areas such as defense, intelligence & government environments which operate over a wide area, but which require consistent access control regimes. Combining both Identity data and XACML based access control policies together with scalable distributed operations makes the ViewDS capability unique in the market today.

ViewDS can also provide Failover and High Availability through support for the X.500 DISP protocol, which allows information to be replicated across multiple DSA servers. The specific types of information that is replicated is controllable, allowing Partial Replicas to exist, holding a subset of the data contained within an internal or secure server.

Scalable PKI Data Management is provided through the ViewDS support for searching capabilities such as PKI Matching rules and Component Matching. Such technologies remove the need to store extraneous information to allow the discovery of PKI data. ViewDS permits PKI data to be searched in its natural form. Such technology not only allows certificates to be discovered by search for DNs (or component within a DN) or Serial Numbers, but it also allows you to search by expiration date or if a given serial number appears within a CRL.