

Stress Relief for Large Scale Directory Infrastructures - and Their Owners

By Phil Hunt, Director of Product Management and co-Founder, OctetString

When you strip away all the business-speak and technological drivel, there is one fundamental reason for developing enterprise applications: to make it easier for one interested party to transact business with another. Whether you're an employee accessing information on your 401(k) through a company portal or a corporation installing an ERP system, being able to launch an application, authenticate your identity, and do what you need to do with a few minutes worth of pointing and clicking provides a huge time savings in a world where there's more to do than time to do it. Theoretically.

Why the caveat? Because as organizations become successful and their network infrastructures grow to meet the demand, systems that used to work perfectly suddenly become overwhelmed. Often the blame is placed on system capacity. Organizations believe that adding more servers, load balancers, etc., will solve the problem. Yet, those solutions may be treating the symptom rather than the cause. They may provide temporary relief, but soon the systems are overburdened again, and the cycle continues to spiral downward.

Rather than taking two servers and calling the Network Admin in the morning, organizations need to look beyond the hardware to the way large scale directories are structured. Checking for the following two directory-related problems may save a lot of time and aggravation, and help provide a cure rather than a mask for the ills of the enterprise applications.

Connection Hog Syndrome

In many large scale directory environments, there may be several applications that "hog" connections. At application bootstrap time, those applications are assigned a directory server. They then establish a permanent connection with the server, whether they're actually processing transactions or not. Once enough of those connections are in place, two problems can occur:

Fault tolerance and load balancing is effectively bypassed. Since most traditional approaches use connection-based load balancing and failover, connection-hungry applications cannot be easily moved from an overloaded directory server to one with more capacity. They're stuck where they are until they release the connection, which may not be until well after it's needed by the system.

Since they're only using one directory, connection hogs may place too great a load on the node to which they are assigned. By never relinquishing the connection, there is no opportunity for the management systems to readjust the load.

Everybody's Talking At Me

This is the opposite scenario to the connection hog. In this case, too many applications are overloading the directory server by making too many TCP/IP connections and LDAP bind requests at the same time. When this happens, as often occurs during peak purchasing times,

many directory servers (both LDAP and X.500) can become unstable as they run out of system resources or processing threads. Many directory server benchmarks place the peak performance level at 7 to 10 simultaneous connection threads. While they may be rated to answer many more requests, the fact is peak throughput is diminished because the server spends its time establishing connections rather than servicing requests. Think of calling for a pizza at dinner time on Friday night. So many people are calling at the same time that the pizza joint quickly run out of phone lines and you get a busy signal. You keep spending your time redialing, and eventually get through to someone on the other end, who shouts "please hold" and puts you on hold before you can say a word. On the other end of the phone, the pizza guy is busy answering lines and putting them on hold - which means he's not taking orders, and he's not making pizzas either.

Everyone into the Pool

The problem directories (and databases) face is one of limited resources. Here's where virtual directories, which are already known for federation, integration and mapping, come to the rescue. When virtual directories are implemented as proxies they already know a lot about the applications they serve and the infrastructure they integrate with. They can provide a solution to both of these scenarios by changing the way connection loads are distributed through connection pooling. Connections are pooled at the virtual directory rather than going straight to the servers, reducing the number of connections going to the source directory servers while changing the path with which they connect.

In the case of the connection hog syndrome, virtual directories allow those applications to establish the permanent connections with the virtual directory rather than the source servers. The virtual directory then connects to the source servers as required, spreading the load between multiple servers. Most important, the virtual directory relinquishes the connection to the source server when the transaction is finished, freeing that server for other tasks.

Pooling connections also helps address the problems of too many connection requests at once. In this case, the connection requests from multiple clients go to the virtual directory, where they are pooled into a smaller, denser, more manageable number of connections before being passed on to the source server. Going back to the example of the pizza parlor, the effect is the same as the pizza guy suddenly becoming The Flash. He is able to organize the customers into fast-moving lines and handle the calls quicker, thus becoming super-efficient. Because existing connections are re-used, the amount of consumed resources on the proxy server is greatly reduced, allowing it to focus on LDAP transaction processing.

Return to Efficiency

Like water pressing against the foundation of a house, overstressing directory servers can cause cracks to form in the enterprise application infrastructure. And, as those applications get more use, these cracks can slow transaction processing to a crawl, thus defeating the purpose of having those large systems in the first place.

Virtual directories help make the number of directory server connections more manageable regardless of the cause, returning systems to peak efficiency. By helping it work better they not only relieve the stress on your infrastructure. They also relieve it on the people responsible for it.