Show Changes
Edit
Print
Recent Changes
Subscriptions
Lost and Found
Find References
Rename

Search

History

8/1/2004 8:17:42 AM

List all versions

What Is A Service Principal Name SPN

Many daemons are configured to run with domain credentials. For instance, consider a Windows service that runs as **Network Service** on a machine in a domain. This service has the domain credentials of the machine, and clients can use Kerberos to authenticate with it. But if you look at how Kerberos works (What Is Kerberos), it seems as though the client needs to know exactly what account the server is running under, because she must ask the KDC for a ticket for that account (the ticket is normally encrypted using a key derived from the service account's password). If the client can't provide this information, how is she ever supposed to authenticate with the server?

Another problem we have is mutual authentication. When a client uses Kerberos to authenticate with, say, **mydaemon** on machine **FOO** via port 4761, the client should receive some assurance that **mydaemon** isn't being spoofed. For example, what if the real machine **FOO** has been disabled and a Trojan machine with the same name has been started on the network? Kerberos can help **mydaemon** assure the client of its authenticity, but how, when the client has no idea what account the daemon is supposed to be using?

Both of these problems are solved by using a Service Principal Name (SPN). If we want the client to obtain tickets and authenticate with a daemon called **mydaemon** running on machine **FOO** and listening on port 4761, we let the client ask for a ticket using a name constructed from that information: **mydaemon/foo:4761**.[1] To make this work, we need to configure Active Directory with a mapping from this name to the account that **mydaemon** is *supposed* to be running under. Now the KDC can use this information to issue the correct ticket. Remember, the directory is the trusted oracle where we store security policy, so we should use it to indicate under which account we want a particular service to run. The name **"mydaemon/foo:4761"** is called an SPN.

With this infrastructure in place, imagine the difficulty an attacker would have trying to spoof a daemon on the network. First he must knock out the real daemon, perhaps by disconnecting it from the network or flooding it with fake connection requests. Then he must start up a new machine on the network with the same name and expose his Trojan service. But since the client is using Kerberos with mutual authentication, she won't talk to the Trojan until he proves that he knows the password for the account under which the real service is configured to run. So the attacker now needs to either discover this password or compromise Active Directory to reconfigure the SPN mapping.

See HowToUseServicePrincipalNames to learn how to use service principal names in practice.

---

[1] Note that the text **mydaemon** is completely arbitrary: Each service just needs to decide on a unique string that it will use to construct its SPN, and any clients that want to talk to that service will need to form the SPN using that string. This is called the "service class," and for real NT services a reasonable convention is to simply use your service's short name.

PortedBy RubenBartelink