Samba and SMB/CIFS

The CIFS/SMB protocol has been notorious for each implementation doing their own little extensions to add features or work around problems in other implementations. Samba is no exception.

In this talk I'll describe some of the more interesting extensions that we have put into Samba. Although other implementors don't need to know about these extensions to inter-operate with Samba I'm hoping that airing these in public will give some useful ideas for other implementors and promote discussion.

WINS and *1B

Cross-subnet browsing has always been difficult with SMB. It can be made to work with the protocols as they exist in the spec but not without a lot of trouble and only if your network layout is "just right".

After looking at this we found that the core piece of missing technology is some way for a host to ask the WINS server for a complete list of Domain Master Browsers that it knows about. Given that information it is easy to construct enterprise wide browse lists.

We solved this by extending the Samba WINS server to accept name queries queries for the name *1B. The query will return the current list of known Domain Master Browser IP addresses.

Browse Synchronization

To use the *1B feature we modified our DMB implementation to periodically send *1B queries to the WINS server (if present). If the query is successful then a node status request is sent to any of the resulting IPs which are unknown to the DMB. The node status request is used to determine the workgroup name. Resulting workgroups are scheduled for browse synchronization via the normal mechanisms.

Inter-DMB synchronization

Our *1B browse synchronization solves the cross-subnet browsing issue when all DMBs are using the same WINS server, but does nothing to help the case where some DMBs use a different WINS server or use broadcast registration.

To solve this problem we added inter-DMB synchronization. Each Samba DMB will perform a workgroup browse sync with all currently known DMBs by choosing a random DMB to sync with at regular intervals. The frequency of these sync operations is fixed, regardless of the number of DMBs, thus preventing a N^2 explosion of network traffic.

These sync operations allow a Samba DMB to build a much more complete list of workgroups. The disadvantage is that "dead"

workgroups can survive for long after they have no members. We are thinking about a solution to that problem, which is caused by lack of TTL information if the NetServerEnum response.

Remote announce

When the above mechanisms fail, such as in very disjoint networks without a WINS server, it is sometimes useful to tell Samba to send host announcements to a remote subnet where they will be picked up by the local master browser. You can do this using a "remote announce" option.

This option allows you to force the appearance of a server in any workgroup listing you like as long as you know the broadcast address of the remote network. Even over an international link.

Netbios aliases

Samba can announce and register itself as any number of simultaneous netbios names on the network. This can be useful when the server has multiple roles which you may want to split onto separate machines at a later date.

Associated with each name can be a Samba config file, so you can configure each of the names to behave quite differently. For example one may be a user level security file server and another might be a share level security print server.

Encryption conversion

A constant headache for Unix SMB server implementations is the handling of the encrypted SMB password database and its synchronization with the standard unix password systems.

To try to reduce the pain a little Samba can auto-add users to the encrypted database as they login. This allows you to convert from non-encrypted to encrypted with a minimum of fuss.

sync options

Windows programmers don't seem to know the difference between sync() and flush(). We see quite a few programs (Explorer being one of them!) using O_SYNC when opening files in a quite inappropriate manner. This has an enormous performance penalty.

For this reason we give the administrator the option of ignoring sync requests (this can be set on a per share basis).

smbwrapper

For Samba 2.0 we have implemented a library of posix-like filesystem functions that can be preloaded by a unix loader, allowing existing unix binaries to see a virtual SMB filesystem and network neighborhood.

This provides a SMB client filesystem on a wide range of Unix flavors which finally allows unix systems to become true members of a SMB workgroup.

The system (called smbwrapper) also shows printer shares as directories, with printing via cp, print queue listing via ls and print queue deletion via rm.

SSL

Samba 2.0 has SSL support in both the client and server (contributed by Christian Starkjohann). This allows SMB sessions to be established over an underlying secure transport, giving a fully encrypted network filesystem.

This is particularly useful in conjunction with smbwrapper.

summary

Samba tries to be as compliant as possible with the SMB and CIFS specifications as they are currently known, as well as the implementation quirks of WinNT and Win95. Beyond that we have also implemented a number of extensions that make Samba more useful without breaking access from clients that don't have the extensions.

I hope that the description of some of these extensions will give ideas to other SMB implementors and promote discussion of the protocol.

Questions?