

Spontaneous Windows File Sharing via Virtual Groups*

Stefan Götz, Klaus Wehrle
Protocol-Engineering & Distributed Systems Group
University of Tübingen
{Stefan.Goetz,Klaus.Wehrle}@uni-tuebingen.de

Abstract: Although file and printer sharing services have been deployed almost ubiquitously for a long time as part of Microsoft Windows, only recent peer-to-peer applications popularized file-sharing on a global scale. As the Windows CIFS protocol was designed for local area networks, its use has been confined to relatively small environments. We propose a mechanism to set up spontaneous virtual groups that allow to use legacy Windows file and printer sharing globally in virtual LANs.

1 Introduction

Microsoft Windows[©] file and printer sharing is one of the most popular and widely available services in local area networks. Since it is an integral part of the Windows operating system, it often forms a center-piece for user collaboration in companies and home networks. Its tight integration into the operating system makes it intuitively and easily useable and probably similarly common-place as the web and e-mail.

Since the underlying CIFS protocol (also, and incorrectly, known as SMB) [Her03] relies on broadcast communication, it can only provide its services within a LAN. Also, peers behind NAT gateways and firewalls cannot connect to each other. This is often the case for peers in different administrative domains, e.g. a dial-up user trying to access his desktop machine in a company network. Today's standard solution to this problems are Virtual Private Networks (VPNs). However, they are too cumbersome or even impossible to set up and administer for regular users who just want to share files with each other.

Our alternative approach connects peers to a virtual LAN as a common broadcast medium. This preserves the semantics expected by CIFS and thus does not require modifications of existing CIFS implementations. At the same time, we leverage the existing infrastructure and connectivity support of the i3 Internet Indirection Infrastructure [SAZ⁺02]. Our approach can also benefit from other i3 features such as firewall and NAT tunneling or mobility support [SLW04, LSW⁺04].

2 Background

This section provides an overview of the Common Internet File System *CIFS* protocol as well as a brief introduction to the Internet Indirection Infrastructure *i3*.

*This work has been supported by Landesstiftung Baden-Württemberg under grant xxx.yyy

2.1 Common Internet File System – CIFS

The primary functionality of the CIFS protocol suite [Her03] is to support file and printing services in LANs. The participating peers can act both as servers and clients when offering or accessing these services, respectively. Services are identified through a simple naming scheme: `\\ <server> \ <share>` names the machine hosting a service in a non-hierarchical manner and the service itself (termed a *share*).

CIFS mainly consists of two protocol levels. At the lower level, the *NBT* protocol provides the NetBIOS API on top of TCP/IP to access peers on a LAN [RFC 1001,RFC 1002]. NBT implements a name, session, and datagram service. The name service is available through a decentralized multicast or a centralized unicast mode or a combination of the two. The session and datagram services are transport services similar to TCP and UDP.

When a peer wants to connect to a file service, it resolves the name of the server to an IP address using the NBT name service and then establishes an NBT datagram or session connection to the server. The higher-level Server Message Block (*SMB*) protocol allows the peer to authenticate with and access the services offered by the server.

2.2 Internet Indirection Infrastructure – i3

The core idea of the Internet Indirection Infrastructure (i3) [SAZ⁺02] is to communicate across one or more points of indirection in contrast to end-to-end communication. Indirection points are identified by unique IDs. Peers can register *triggers* with IDs, which instruct i3 to forward data for an ID to the receiver set in the trigger. This scheme decouples the act of sending from the act of receiving and can thus provide additional features like multicast, anycast, mobility support, or service composition [SLW04, ZLS⁺03, LRSS02].

3 Design

Our design leverages the concepts and infrastructure of i3 in order to connect CIFS peers on separate LANs to form a distributed virtual LAN. Via i3, peers on the same virtual LAN can communicate over a virtual broadcast facility as well as directly with each other as shown (cf. Fig. 1). These i3 services are transparently interposed between the involved peers avoiding prohibitive source code level modifications of the peer systems.

3.1 Virtual LANs

Our approach connects peers via i3 so they form non-local broadcast realms which we call virtual LANs. At the i3 level, routing CIFS broadcast traffic to form a virtual LAN is straight-forward: all VLAN members associate with a common i3 multicast point. I.e., they all register triggers for a common ID – e.g. the SHA-1 hash value of the VLAN-ID – and associate these triggers with their host addresses. Consequently, i3 forwards all traffic for the VLAN-ID to the registered peers.

Assuming no access control mechanism, peers only need to know the VLAN-ID to join a specific virtual LAN. Peers can learn VLAN-IDs through a variety of mechanisms. For simplicity, we identify virtual LANs by textual names and derive the VLAN-ID by hashing the name which needs to be known by all potential VLAN participants.

The standard SMB user authentication mechanisms continue to work as expected in such an environment. However, communication at the i3-level may need to be secured. As a

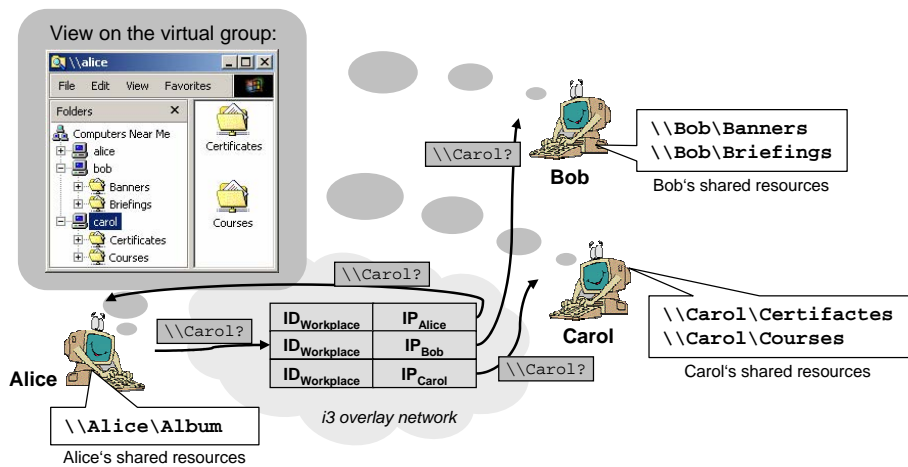


Figure 1: Example Scenario: Virtual workgroup formed by Alice, Bob and Carol on top of the i3 overlay. Alice is issuing a lookup for Carol, which is distributed via the i3-multicast feature.

simplistic measure, VLAN members could e.g. symmetrically encrypt all piggy-backed NBT traffic such that it can only be decyphered by peers knowing the VLAN-specific key. Alternatively, i3 service composition allows interposing centralized security schemes between the VLAN members and the multicast point, such as authentication and encryption services. These security considerations are orthogonal to our approach and thus not discussed in further detail.

3.2 Virtual Workgroups

Windows supports only a single workgroup which is tedious to configure. Furthermore, users might prefer remaining in their physical workgroup while also joining one or more virtual LANs. Therefore, we decided to slightly extend the idea of a standard workgroup.

In contrast to a single physical workgroup, peers can join several virtual workgroups which represent virtual LANs. The user may freely configure the virtual workgroups to join in the i3 client software instead of Windows. The client software also modifies browsing information such that VLAN peers appear in the corresponding virtual workgroup. Thus, a high degree of consistency is maintained from the user's point of view.

3.3 Routing CIFS Traffic via i3

Existing CIFS implementations do not support the i3 protocol natively. The effort of adding and maintaining i3 support in their implementations is prohibitive. Proprietary systems like Windows would not allow such changes at all.

Thus, we bridge this gap with a more generic solution by running the i3 proxy [LSW⁺04] on every peer in a virtual LAN. It already provides the basic infrastructure required such as packet filtering, re-writing, address translation, and access to the i3 service. The CIFS-related features are encapsulated in a protocol-specific extension module of the proxy. The main responsibilities of this module are to re-route CIFS packets to and from i3 and

to translate the addresses of other peers so they become meaningful for the local peer.

3.4 Routing Unicast CIFS Traffic via i3

The NBT name service can work as expected in a virtual LAN so peer names get resolved to the real IP addresses of the respective peers. However, these addresses may not be reachable if the peers are located behind a NAT gateway or firewall. Only by re-routing connection attempts to such IP addresses via i3, connectivity can be ensured.

In order for receiving peers to be accessible via i3, the proxy on each peer inserts a peer-specific trigger into i3. Each proxy maps the IP addresses of remote peers to the corresponding trigger ID and re-routes CIFS traffic accordingly.

3.5 Address Translation

Network address translation gateways can break the routing scheme presented above because IP addresses might not be unique. In such a case, the i3 proxy is not able to maintain unambiguous associations of remote peer addresses and trigger IDs. Furthermore, the IP addresses of remote peers can be in the address range of the local area network, inhibiting a distinction of local and remote CIFS traffic.

The proxy is free to present the local peer with fake IP addresses of remote peers as long as the CIFS protocol is not violated. As the proxy intercepts all packets from remote peers, it rewrites those that advertise IP addresses of remote peers and passes them on to the local system. Through this indirection, the proxy can associate each remote peer with an IP address that does not conflict with addresses of other remote or local peers in a fashion transparent to the local system.

4 Outlook

We currently investigate mechanisms to let peers access any services in remote LANs in which only a single gateway peer runs the i3 proxy. Such use-cases also require authentication and authorization at the i3 level. Furthermore, an elegant approach to resolving clashes of peer names is still a matter of research.

References

- [Her03] Christopher R. Hertel. *Implementing CIFS: The Common Internet File System*, volume 1. Prentice Hall, August 2003.
- [LRSS02] K. Lakshminarayanan, A. Rao, I. Stoica, and S. Shenker. Flexible and Robust Large Scale Multicast using i3. Technical Report CS-02-1187, UC Berkeley, 2002.
- [LSW⁺04] K. Lakshminarayanan, Ion Stoica, Klaus Wehrle, et al. Supporting Legacy Applications over i3. Technical Report UCB/CSD-04-134, UC Berkeley, May 2004.
- [SAZ⁺02] I. Stoica, D. Adkins, S. Zhaung, et al. Internet Indirection Infrastructure. In *Proceedings of ACM SIGCOMM'02*, August 2002. Pittsburgh, PA.
- [SLW04] I. Stoica, K. Lakshminarayanan, and K. Wehrle. Support for Service Composition in i3. In *Proceedings of ACM Multimedia*, October 2004. New York.
- [ZLS⁺03] S. Zhuang, K. Lai, I. Stoica, et al. Host Mobility Using an Internet Indirection Infrastructure. In *Proceedings of ACM MobiSys*, 2003.