Network Working Group
Requests for Comments 1275

S.E. Hardcastle-Kille
University College London
November 1991

Replication Requirements to provide an Internet Directory using X.500

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard. Distribution of this memo is unlimited.

Abstract

This RFCconsiders certain deficiencies of the 1988 X.500 standard, which need to be addressed before an effective open Internet Directory can be established using these protocols and services [CCI88]. The only areas considered are primary problems, to which solutions must be found before a pilot can be deployed. This RFCconcerns itself with deficiencies which can only be addressed by use of additional protocol or procedures for distributed operation.

1 Distributed Operation Extensions

The Internet Directory will operate DSAs over TCP/IP using RFC 1006 [RC87], and DSAs over the an ISO Network Service. Distributed operation procedures should not require full connectivity.

2 Knowledge Replication

Knowledge information is critical to resolution of names, and performing searches. Knowledge information high up the tree needs to be widely available. Consider resolving a name below "Country=US". To do this, a DSA needs to have full knowledge at this point. Many DSAs need to be able to do this, in order to give reasonable response and availability. It would be an unacceptable bottleneck to force such resolution to a single or small number of DSAs. To replicate this knowledge widely, a systematic approach to replication is needed.

3 Data Replication

Searches are often made at the root and country level, and this is a vital service (e.g., an approximate match of an organisation name). Data needs to be collected in such a way that this sort of searching is reasonably efficient. The usual X.500 approach of subordinate references militates against this. At a node in the DIT, subordinate references to the entries below are held. These entries will be in many DSAs, each of which needs to be accessed in order to perform the single level search. It is suggested that replication of data is necessary to achieve this.

The major requirement for this replication is high up the DIT, where information must be replicated between different implementations. At lower levels of the DIT, it is reasonable for DSAs to be of the same implementation and to use implementation specific techniques in order to achieve performance and availability.

4 Alternate DSAs

When a DSA Referral is returned, only the master DSA is indicated. This will lead to a single point of failure. It seems important to allow for additional references to slave copies, in order to get better availability. This needs to be solved in conjunction with the problem described in the previous section.

5 Guidelines for use of Replication

To be effective, the replication specification needs to provide guidelines for deployment in the pilot, in order to meet the desired service criteria.

6 Some scaling targets

Most techniques for replication have scaling limits. It is important that mechanisms used do not stress the limits of the mechanism. The order of magnitude envisioned in the pilot is 100 000 non-leaf entries and several million leaf entries.

References

[CCI88] The Directory — overview of concepts, models and services, December 1988. CCITT X.500 Series Recommendations.

[RC87] Marshall T. Rose and Dwight E. Cass. ISO Transport Services on top of the TCP. Request for Comments 1006, Northrop Corporation Technology Center, May 1987.

7 Security Considerations

Security considerations are not discussed in this memo.

8 Author's Address

Steve Hardcastle-Kille
Department of Computer Science
University College London
Gower Street
WC1E 6BT
England

Phone: +44-71-380-7294

EMail: S.Kille@CS.UCL.AC.UK