

Mesh of Multiple DAG servers - Results from TISDAG

Status of this Memo

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

Copyright Notice

Copyright (C) The Internet Society (2000). All Rights Reserved.

Abstract

The Common Indexing Protocol ([CIP1]) is designed to facilitate the creation not only of query referral indexes, but also of meshes of (loosely) affiliated referral indexes. The purpose of such a mesh of servers is to implement some kind of distributed sharing of indexing and/or searching tasks across different servers. So far, the TISDAG (Technical Infrastructure for Swedish Directory Access Gateways) project ([TISDAG], [DAGEXP]) has focused on creating a single referral index; the obvious next step is to integrate that into a larger set of interoperating services.

1. Introduction

1.1 Overview of mesh possibilities

Two different possibilities are possible for extending the TISDAG service to a mesh model (or some combination of both). First, it should be possible to create a mesh of DAG-based services. Or, it might be interesting to use the mesh architecture to incorporate access to other types of services (e.g., the Norwegian Directory of Directories). In either case, the basic principle for establishing a mesh is that interoperating services should exchange index objects, according to the architecture of the mesh (e.g., hierarchical, or graph-like, preferably without loops!).

As is outlined in the CIP documentation ([CIP1]), many possibilities exist for mechanisms for creating indexes over multiple referral servers -- for example, WDSP index objects could be passed along

untouched, or a referral index server's contents could be aggregated into a new index object, generating referrals back to that server.

The proposal is that the mesh should be constructed using index objects aggregated over participating services' servers. That is, referrals will be generated to other recognized services, not their individual participants. This can be done as a hierarchy or a level mesh one-layer deep, but the important reason for not simply passing forward index objects (unaggregated) is that individual services may support different ranges of access protocols, have particular security requirements, etc. Referrals should be directed to a CAP or CAPs -- either the standard ones used by the DAG system, or new ones established to support particular semantics of remote systems (e.g., other query types, etc). Within a given DAG system, referrals to these remote servers will look just like any other referral, although a particular SAP or SAPs may be established to provide query fulfillment (again, to enable translations between variations of service, to allow secure access if the relationship between the services is restricted, etc).

In the following scenarios of mesh traversal, the assumption is that the primary service in discussion (Country A in Scenario 1, Country B in Scenario 2) is a DAG-based service. The scenarios are presented in the light of interoperating DAG services, but in most cases it would be equally applicable if the remote service was provided by some other service architecture. Again, the key element for establishing a mesh of any sort is the exchange of the CIP index object, not internal system architecture.


1.1.1 Scenario 1: Top Down

Suppose 2 countries tie their services together. A user makes a query in Country A. A certain number of hits are made against the index objects of A's WDSs. There is also a hit in the aggregate index of Country B. There are 3 possible cases under which this must be handled:

Case 1:

Country A and Country B are running services that are essentially the same -- in terms of protocols, queries, and schema that are supported. In this case, one referral should be generated per protocol supported by Country B's service. The referral can be passed back as far as the client, if its protocol supports referrals. Alternatively, the CAP may chain the referral through an appropriate SAP, in the usual fashion. In other words, the CAPs of Country B's service act as WDSs to Country A's service.

Consider the following illustration (only relevant CAPs, SAPs, etc, are shown; others suppressed for lack of room):



where

- Prot[i] is some particular query protocol
- RI-A has an index over all A-WDSP[i] and RI-B
- RI-B has an index over all B-WDSP[i]
- (1) is the query to the Country A DAG system, which yields a referral based on the index object from RI-B
- (2) is that referral
- (3) is the resolution of that referral, which the client takes to the Country B DAG system directly (to find out which, if any, B-WDSP[i] have relevant information)

Case 2:

Country A and Country B are running services that address the same service type (e.g., whitepages), but are not using an identical collection of protocols, allowed queries, or schema. The index object that Country B sent to Country A's DAG service must be constructed in terms of Country A's service, in order for appropriate hits to be generated against the index object (i.e. for referrals to Country B's service). However, to resolve the referral, it will be necessary to do some further protocol/schema/query mapping. This can be done by a special SAP established within Country A's service, that maps Country A's service into the published service of Country B. Country A may then elect to support only one of Country B's access protocols, and the designated SAP will always contact one type of CAP at Country B.

Alternatively, Country B can establish a particular CAP that does the mapping from Country A's service into something that is most appropriate against the internal structure of its service. In this case, Country A's referral will be to a special CAP in Country B's service (which, again, will look like a WDSP to the Country A service); in fact, the referral may be handled directly by the client software. The difference between the two possible approaches lies in the responsibility of managing the relationship between the 2 service types. On the one hand, Country A could handle it if it knows its service as well as the published access to Country B. On the other, Country B could be responsible for establishing a CAP for every country that may want to connect to it. The latter can, in some cases, be justified by the amount of internal optimization that can be done, and because it reduces the overhead for Country A's service (can pass the referral directly back to the client software).

Consider the following illustration (only relevant CAPs, SAPs, etc, are shown; others suppressed for lack of room):

Case 3:

The third possibility is, in fact, a refinement of the first. If Country A and Country B are running services that are every way identical except for the data (WDSPPs covered), then it may make sense to NOT aggregate Country B's WDSPP index objects, but to copy them to Country A's server. Then, Country A's CAPs might be given access to the SAPs of Country B in order to carry out chaining directly at the remote service (instead of implicating Country A's SAPs and Country B's CAPs, as in the first example above). The answer does not come from technology -- it depends entirely on the nature of the relationship that can be established between Country A and Country B's services.

1.1.2 Scenario 2: Working Up

The above scenario implicitly assumes that Country A's server had received index objects from Country B's server. This will be the case if Country A's server is higher in the levels of a hierarchy of services (established by agreements between the service operators), or if the network is comprised of servers that share their index objects with all others, for example. In the latter case, searching at any one of the servers in the service yields the full range of results -- referrals will be made to any other server that might have data that fulfills the user's query. The sharing of the index objects is a mechanism to allow each server to manage local data, while enabling distributed load-sharing on the basic query handling.

However, if a hierarchical, or at least not-completely-connected model is used for the server network, queries carried out at a level other than the top of the hierarchy, or in one particular branch of the hierarchy, will not actually be matched against all index objects. Therefore, there may be other servers to which the query should be directed if the full space needs to be searched. Suppose, for example, that in the above example Country B is in fact lower in the hierarchy than Country A. A user sending a query to Country B's service may be content to limit the scope of the query to that country's information (this is true in enough real-life situations that this hierarchical relationship becomes an effective mechanism for scoping queries and avoiding having to flood the entire network with every single query or keep full copies of all data in every server).

Still in theoretical stages, the DAG/IP provides control constructs to allow DAG components to act according to the topology of the mesh. A CAP might use the "polled-by" system command to establish what other servers in the mesh exist in higher levels (and therefore would be worth contacting if the scope of the search is to be increased).

In the example above, a CAP in Country B's system could determine that Country A's service was polling Country B, and therefore make it a logical target for expanding the scope of the query. More experience (primarily with server mesh topologies) is necessary before it will be clear how to best make use of these capabilities:

- . should the CAP always broaden the scope? only if there are no local referrals? under user direction?
- . should the CAP use a local SAP to contact the remote service's CAP?
- . is it better to completely connect the mesh of servers, or produce some kind of hierarchy?
- . etc

2. Other considerations

Depending on the context in which a mesh is established (e.g., between national white pages services, or different units of a corporate organization, etc), it may be useful to allow individual WDSs to indicate whether they are willing to have their data included in a DAG system's aggregated index object (i.e., allowing the DAG system to receive referrals from other systems in the mesh).

3. Security Considerations

This document describes different configurations for sharing information between information services. It introduces no security considerations beyond those attendant in (and addressed by) particular directory service access protocols.

4. Acknowledgements

The work described in this document was carried out as part of an on-going project of Ericsson. For further information regarding that project, contact:

Bjorn Larsson
bjorn.x.larsson@era.ericsson.se

5. Authors' Addresses

Leslie L. Daigle
Thinking Cat Enterprises

EMail: leslie@thinkingcat.com

Thommy Eklof
Hotsip AB

EMail: thommy.eklof@hotsip.com

6. References

Request For Comments (RFC) and Internet Draft documents are available from numerous mirror sites.

[CIP1] Allen, J. and M. Mealling, "The Architecture of the Common Indexing Protocol (CIP)", RFC 2651, August 1999.

[TISDAG] Daigle, L. and R. Hedberg "Technical Infrastructure for Swedish Directory Access Gateways (TISDAG)," RFC 2967, October 2000.

[DAGEXP] Eklof, T. and L. Daigle, "Wide Area Directory Deployment Experiences", RFC 2969, October 2000.

[NDD] Hedberg, R. and H. Alvestrand, "Technical Specification, The Norwegian Directory of Directories (NDD)", Work in Progress.

7. Full Copyright Statement

Copyright (C) The Internet Society (2000). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

