#15

Monitoring Quality of Service in a Global Information Service

Colin Robbins

This paper was written by Colin Robbins (NEXOR), in his role as consultant to DANTE on NameFLOW-Paradise (August 1995).

DANTE IN PRINT is a track record of papers and articles published by, or on behalf of DANTE. HTML and Postscript versions are available from: http://www.dante.net/pubs/dip.html

For more information about DANTE or *DANTE IN PRINT* please contact:

DANTE
Francis House
112 Hills Road
Cambridge CB2 1PQ
United Kingdom

Tel: +44 1223 302 992
Fax: +44 1223 303 005
E-mail: dante@dante.org.uk

Monitoring Quality of Service in a Global Information Service

Colin Robbins

Version: 0.6
Status: DRAFT for comment

Abstract

This paper looks at the Quality of Service issues that affect the Internet OSI Directory Service, and identifies the management infrastructure required to monitor and improve the existing service.

1 Introduction

The PARADISE project [1] has been running a research based OSI Directory Service, based upon the X.500 standards, since 1990. The pilot now extends to 35 countries, holding in excess of 1.5 million entries. This has been built up largely on an ad-hoc, best effort basis. This has lead to Directory where the accessibility of nodes is variable: some are always available; others never and some intermittent. Once accessed, the accuracy and completeness of the data is also variable between nodes. Problems of this nature are not limited to the OSI Directory, but are a general problem in distributed informations systems. Whilst the paper focuses on X.500, a number of the concepts could be applied to other information services, such as the World Wide Web[4].

NameFLOW-Paradise is now being run on a commercial basis by DANTE(1), with a goal of migrating from a pilot system to a service. A key part of this transition is being able to manage and consequently improve the quality of the service offered.

This paper starts in Section 2, which looks at the service requirements for a Global Directory, and presents the importance of being able to measure the system against these requirements in Section 3. The paper then moves on to look at existing Directory Quality of Service (QOS), and explores what we already know about the QOS of the current system in Section 4. From this a new set of measurable quality parameters are discussed and defined in Section 5, together with a proposed infrastructure to be used to collect this information on a global basis. Finally, Section 6 draws some conclusions and identifies future work required. The paper assumes the reader has background familiarity with the concepts of X.500.

This paper has been funded by DANTE, as part of the NameFLOW-Paradise service. It is a discussion document, it does not imply any of the procedures will be implemented or enforced within the NameFLOW-Paradise service.

2 Requirements

Before presenting specific QOS issues and thus looking at solutions, it is important to clearly understand the requirements of the Global Directory system. Then the solutions can be judged to see if they impact the requirements. Firstly there is the user perspective of the role of the Directory, and secondly a technology requirement.

2.1 User Perspective

From the user perspective, one role of the Directory, and I suggest the key role, is as an on-line source of contact information. The information will typically be accessed at the time the information is required, for example just before a phone call is made, or during the composition of an electronic mail messages. This mode of usage dictates that the Directory is available for access at all times.

We should also be aware that the Directory is a global system, so there is no room for concepts such as daytime availability only, or one hour

Colin Robbins is senior technology consultant at NEXOR and has worked for DANTE as consultant on NameFLOW-Paradise. His e-mail address is C.Robbins@nexor.co.uk

downtime at midnight for backups. One user's midnight is another's midday. This leads to the user requirement for 24 hour data availability.

A second consequence of this mode of user interaction, is that a user will not wait forever for the information. If the system is slow to respond to a data access query, the user may look for the information elsewhere, and if successful may not use the service in the future. This leads to the user requirement for a real time response from the system.

Putting these two concepts together, gives us the system requirement:

> Requirement 1. The Directory service has to be reliable.

One of the key points this paper looks at is how to define a reliable directory, and consequently identify if the service we are providing meets this requirement.

If the service is to be of assistance to the user, the data they are looking for has to be contained within the system. This can be split into two requirements:

> Requirement 2. Data held about an organisation or individual is as complete and accurate as possible.

and most importantly:

> Requirement 3. Data is held about all organisations and individuals with which the user may wish to communicate.

Requirement 3 is perhaps the hardest issue to address, and there has been numerous discussion within NameFLOW-Paradise and the various National pilots as to how to increase the database population. One key issue faced is "Why should organisations put data in the Directory - what is the cost benefit?". The simple answer at the moment is: there is no cost benefit.

For there to be a benefit, the quality of the system has to improve to the level where users rely upon it. Then, by not being in the Directory and organisation is at a competitive dis-advantage, as contacting them becomes harder for users and clients. In this respect, addressing requirements 1 and 2 are critical. This paper will focus upon the first two issues, in the belief that this will contribute to a resolution of the third. Some analysts have questioned this assertion, suggesting that a "critical mass" data set is required before this situation arises, leading to the suggestion that "organisation name" data is bulk loaded from external sources to provide some limited information. This leads to a set of questions like which data set, does it provide accurate data, and do we have the right to load it anyway. The author believes this will not help, as the quality of the data is unlikely to be good.

2.2 Transition

The existing pilot is based around the QUIPU [2] software - both public and commercial versions, some 90% of the 600 DSAs are QUIPU. QUIPU is an X.500(88) system [3] designed by Steve Kille, and implemented by the author whilst we were members of the Department of Computer Science at University College London. The main thrust of the work was to provide a research tool to demonstrate that distributed X.500 could be made to work in a wide area network. The success of this can be measured by the current NameFLOW-Paradise service. However, there are some problems.

QUIPU made some simplifications and also uses some extensions to the X.500 protocol, which has lead to subsequent interworking issues (see [6] for a detailed description of these issues). A key part of the next phase of NameFLOW-Paradise is to migrate from QUIPU to a system based upon the 1993 edition of the X.500 standards [4]. This includes the introduction of standards based replication and a migration away from the requirement for a root DSA.

Whilst planing this transition is not within the scope of this document, ensuring the quality of the system during and after the transition is. This leads to

> Requirement 4. The quality of the service must be maintained during the transition phase.

This in turn means we need to know what the quality of the current system is, so it can be compared with the new system. Looking more closely, this is essentially a manifestation of the first requirement, the DSA must be reliable - this reliability has to be maintained during the transition phase.

An example of this is in the plan to retire the non-standard root DSA proposed by Chadwick in [13].

During this transition we have to ensure the QOS of the system is not affected, so we need to be able to clearly identify the current QOS, so we can monitor the impact of the transition, and take corrective action if there is a problem.

2.3 Service Agreements

One of the key ways of managing a service, is the introduction of Service Level Agreements (SLA) to provide a management framework. Some recommendations for such an agreement are given an Internet draft (ID) [12], and there are other examples within other national Directory service operations, including [25] from the NADF. The ID recommendation includes percentage data and DSA uptime requirements for connection to the service. As I hope to show here, more detailed information on how these required service levels are going to be monitored is needed. In this paper I introduce some parameters that ought to be in the SLA.

In the companion document [26], a template SLA is given, incorporating the proposals in this paper.

3 Measurement

Before we can authoritatively talk about providing a Directory that meets these requirements, we need to have some metrics to base the discussions on. Quoting "Kelvin's Principle" from [14]:

> "...when you can measure what you are speaking about, and express it in numbers, you know something about it; but when you cannot measure it, when you cannot express it in numbers, your knowledge is of a meagre and unsatisfactory kind."

If we apply this principle to the QOS of the Directory, we need to start looking at exactly what we can measure! Ask any user of the current NameFLOW-Paradise service, and they will report on accessibility problems. But if you look at the probe results (see Section 4.1), you will see the average weekly DSA accessibility figures suggest an availability in the high ninety percents. There does appear to be a contradiction here - perhaps we are measuring the wrong thing.

In the following Section we will look at exactly what we can measure now, then asses how well the measurements help in identifying if the system we have matches our requirements.

4 Current Measurements Techniques.

Using a combination of measurement techniques we can currently measure

- DSA availability
- Data availability
- DSA operation throughput
- DSA performance

In the following sections we will evaluate the techniques to asses whether the parameters really give us what we expect, and in a form they can be used in a management environment.

4.1 DSA Availability

One obvious way of measuring the quality of the system is to use tools to measure it. Within NameFLOW-Paradise, two main techniques have been used to do this:

- Active DSA probing
- Passive probing

With active probing [19], a special purpose DUA periodically connects to a DSA and records the success or failure. This information is then correlated with figures on all DSAs within the Directory Management Domain (DMD) and from other probe sites to provide statistical information on DSA availability. Two tools for active probing have been developed. One by Steve Titcombe of UCL, and an improved successor by Alan Shepherd of NEXOR.

DSA availability tells us a lot about the availability of components of the Directory service, but not the availability of the service itself. Quoting Glib[14]:

> "Bear in mind that single components of a system may have a much higher availability that the total system of which they are part."

This is a very important point, and looking back at the requirements presented in Section 2, it is the system availability that is important. This helps to explain why the high DSA availability figures we currently see is in contradiction with the low availability user perception of the service. This aside, there are other problems with measuring DSA availability.

For a probe to contact a DSA it needs a presentation address. This contains a set of NSAP addresses, typically these are used to provide access points to different network services e.g, X.25, CLNS and TCP/IP (using [15] and [16]). This can only be obtained using implementation

specific insight into how knowledge of remote DSAs is stored. What is more, this information may not be available to user applications when this knowledge information is stored in operational 'specificKnowledge' attributes as defined by the X.500(1993) knowledge model.

Secondly there is a trend towards using the NSAP address to refer to different physical DSAs. The legality of this within the base standards is questionable, but it undoubtedly increases the system availability, so is likely to be a practice that will continue. In this light, we cannot be sure our DSA availability figures do actually reflect DSAs. Having provided a critique of the method, DSA probing has certainly been successful. Publishing a league table of DSA availability in the UK pilot, and at the root level, coincided with a general improvement in DSA availability - it does allow DSA managers to easily identify problems. Whether this needs to be done centrally is unclear. It is just as easy to set up the probe to monitor a set of DSAs within an organisational DMD as it is to probe from the national and root levels.

Finally it should be noted that probing from a remote site does in fact measure three parameters, without being able to clearly identify which parameter has failed in the case of DSA unavailability:

- Software availability
- Host availability
- Network availability

Separating network availability from host and software availability is critical to our understanding of the system. More importantly, resolving the network availability problems is likely to fall into the domain of a different set of managers and engineers than that of resolving software and host failures. Consequently network failures are not in our management domain; it is our responsibility to design a system that is robust in the event of such failures.

To summarise, DSA probing has historically provided a useful measure of the system, but may no longer be an appropriate mechanism for central monitoring.

4.2 Data Availability

In a paper on QOS of the Directory [7], Barker identifies two key problems with the active DSA probing technique:

> "The probe measures the availability of DSAs rather than of information."

and

> "The probe information does not provide a DUI(2) specific view of the network."

Barker then presents how these problems can be solved using passive probing - the technique of building an accessibility database into the DUI. As a DUI collects results of operations, timing and success / failure information is recorded in a database. This database is then used to warn future DUI users of operations that may take a long time, or even fail based upon previous attempts.

Whilst this unquestionably enables the DUI to provide valuable information to the user about pending doom, its in not clear if this helps in progressing any of the requirements identified in Section 2. Although it warns user of unavailable data, the real issue is how do we prevent it being unavailable in the first instance.

Data availability is a key measurement of the system. As shown in [7] passive probing can be used to provide a direct measure of:

> percentage of time the data is available from the overall system.

and

> length of time taken to access(3) the data.

These two measurements directly maps onto out reliability requirement. However, there are problems as pointed out by Barker:

> "The timings produced are specific to DE's(4) abstraction of the Directory, and thus several DUIs cannot readily share a QOS database."

There are a number of potential solutions to this. Firstly define the format of the database, or an API into the database, so that it can be agreed between vendors. This would allow a wide range of DUIs to supply data to the QOS database. The draw back of this approach is it requires software changes that vendors may be unwilling to make, and will certainly take time to propagate.

An alternative approach is to feed the database from the DSA rather than DUI. It is probably true to say there are less DSA implementations than DUI implementations, so the impact of the

requirement will be less. Secondly DSAs almost certainly keep some of this information internally anyway, so may only require minor modifications. This is particularly true if the SNMP MIBs discussed in Section 4.3 are implemented, as some DSA availability information will already be maintained by the system.

Where the data is to be stored also depends upon how the data is to be used. DE uses the data directly to inform users of progress. In this case it may be more appropriate to have the DUI maintain the database. One interesting concept Barker is currently researching is storing the QOS database within the DIT, in this way the DSA could maintain the database, but it would be available to the DUI for presentation to the user. One of the problems with this whole approach is as service providers, we don't really have control of the software used, or control of the extended features the software vendors provide. We could make software requirements, but this will only limit the software available. Consequently this will be detrimental to the overall cause of running a wide scale open Directory. The concepts of DE are important, but we need to see how the information may alternatively be used.

4.3 Operation Throughput

An SNMP MIB for monitoring X.500 has been defined in [8]. The MIB provides information is three categories of information:

- Summary of DSA accessed, operations and errors.
- Summary of the entries held by the DSA and cache performance.
- Summary of the interactions with other DSAs.

This information can be used by the local DSA manager to identify the performance of a particular DSA.

The SNMP MIBs can be used to provide a measure of operation throughput. Whilst this is a useful measure of the DSAs spare capacity, it does not help towards providing information relative to the initial set of requirements.

The information provided is of more use to a DSA administrator so they can establish when the DSA capacity is about reached, so can consider remedies before there is a drop in the required QOS. For this reason use of SNMP will not be considered further here.

4.4 DSA Performance

When looking at the QOS of an overall system it is important to understand the capabilities of the software components that build the system. Two RFCs [10] and [11] define a set of metrics that can be used to asses DUAs and DSAs respectively. Whilst DUAs are an important part of the system, they are not considered within the scope of this paper.

The DSA metrics RFC allows DSA vendors to publish some bench mark tests for DSA performance. Clearly this needs to be taken into account, as the performance of the overall system will not be able to exceed the capabilities of the chosen software.

Metrics defining the performance of a DSA are of questionable use in looking at the QOS of the overall system. Each DSA installation will perform differently to the published figures due to differing data sets and local environments. However, they are a useful tool for the DSA administrator when choosing software to see if the software is sufficiently powerful for the task in hand.

As an aside, I recommend the RFC on DSA metrics [11] needs updating, as the current measures are based on single operations so provide results in the range of a second or two. More of the "query rate" style tests would be useful.

Secondly the current tests can all performed when the DSA is unloaded -i.e., just responding to the tests. It would be more interesting for example to know who it performs when subject to multiple simultaneous queries - a situation that is more likely to be found in practice.

4.5 QOS attributes

In the ID [9] Kille describes "a mechanism for specifying the Quality of Service for DSA Operations and Data in the Internet Pilot Directory Service". The mechanism involves storing attributes that described the data completeness and DSA status within the Directory.

Although implemented within QUIPU, this has not been widely adopted by DSA managers, so the information is not readily available. One reason for this is that it is not mandatory, so has been overlooked by many DSA managers.

Secondly, the definitions used are open to mis-

interpretation, so have been treated cautiously. For example, data completeness is defined by one of the values:

NONE,
SAMPLE,
SELECTED,
SUBSTANTIAL,
FULL.

These are very loose definitions, consequently its not obvious which value to set, and similarly it is not obvious what to do with this information if it is presented. In Section 5.3 on page 9 a more precise way of representing this information is presented.

5 What Should we be Measuring?

Returning to the key requirements defined at the start of this paper:

- Requirement 1. The Directory service has to be reliable.
- Requirement 2. Data held about an organisation or individual is as complete and accurate as possible.

The measurements we can make now only really address parts of requirement 1 from a specific DUI perspective. For the second requirement, our measurements fall short of providing reliable data that can be used to present the completeness of the service. The other techniques described are of use to the local administrator, but not the overall service provider.

Looking at the requirements, and the measurements we have available, it appears we may have been looking at the wrong sort of data in the past. Quoting Glib [14],

> [The measurement parameters] "should be specified in terms of the final end-user result demanded."

How can we directly measure the system performance against the requirements?

5.1 Distributed Monitoring

Whatever measurements are taken, them measurement process cannot be performed centrally. The DIB itself is distributed for reasons of scale and localisation, for very similar reasons the monitoring of the service has to be distributed. From the NameFLOW-Paradise perspective of running the service at the first level, the distribution of monitoring can follow a simple model. Each country or organisation running a first level DSA will provide data as defined in the SLA. The NameFLOW-Paradise service provider will then combine the figures to provide overall statistics on the service, for example it should be possible to say:

> For week x the NameFLOW-Paradise service was n% available, with a data completeness of m%.

This could then be put into a graphical form as the weeks go by, so that the effects of changes within the system can be monitored.

It should be noted that in the following sections the "national service provider" is referenced, but as demonstrated by the NADF, more than one service provider will generally be responsible for providing the service. In this context, it is expected each service provider will have SLAs, and part of this will involve the exchanging of the appropriate measurement parameters, to provide overall country level statistics.

Secondly, in many countries a second level of distribution of monitoring could be appropriate. For example within the pilot in the US, each state could set up a monitoring system and feed information back to the national providers.

5.2 Reliability.

In order to produce this overall figure, how do we determine the parameters locally? What is a reliable directory? How can we define it? In software engineering terms, Shooman [18] offers the definition:

> "Software reliability is the probability that the program performs successfully according to specification, for a given period of time."

Applying this to a distributed Directory, I offer the definition:

> "Directory reliability is the probability that a server agent will, in reasonable time, successfully respond to a user query, for a given period of monitoring as seen from distributed nodes in the network."

The addition of the phrase "as seen from distributed nodes in the network" in the definition is a key point, as the reliability may well be different from different nodes within the network. This implies there will be a set of such data, from different locations that need to be merged to provide the overall reliability figure.

The implication of the definition is that we ought to be able to perform the appropriate measurements to authoritatively say

"NameFLOW-Paradise was n% reliable last week". How do we measure this reliability probability directly?

Using the techniques of passive probing we could simply could the number of queries made, and the number of operations that failed. This would give reliability figure for a specific DUI view of the network. There is a problem here that some organisation will be accessed more than others, so will bias the figures.

An alternative would be for the passive probe to keep track of the number of successes and failures for each organisation accessed, and average the results. To ensure each organisation was accessed, a "power search" could be frequently issued to update the database. If this probing was performed from a number of key sites and merged it could give an overall country level reliability figure.

Rather than using this DUI, a different flavour of active probe could be developed. This would maintain a list of organisations to be monitored, then instead of trying to access the DSA as now, simply try an X.500 operation upon the organisation's data set. A database could then be kept of the success or failure of each organisation. As with the passive probe, use of the data probe could be distributed to key sites within the DMD. In summary, any of the techniques could be used, the important point is to make the measurement so the parameters can be monitored.

5.2.1 Distributed Probing

Above, I have presented some options on how reliability can be measured at the National level, the exact choice has to be defined by the organisation running the service. There will be many such organisations combining to provide the overall Global NameFLOW-Paradise service. Each Country will have a different view of the reliability of the overall service, and will be in the best position to assess the reliability of the service as seen by their key access points.

It is recommend that DANTE set up an international framework to coordinate these measurements. Each first level DSA in Global NameFLOW-Paradise, as part of the SLA, will be required to supply DANTE with:

- regular, measured reliability figure for their countries service,
- overall figure for the reliability of the other countries top level, obtained by probing a given 'typical' node.

Both of these figures can be obtained using any of the techniques outlined in the previous section. DANTE will then obtain these figures from all first level DSA providers within the service, and combine them to provide an overall Global NameFLOW-Paradise figure.

To provide a more accurate figures, each DMD providing the data may wish to merge statistics from a number of different key monitoring sites, in a similar way to the gathering of the current probing statistics.

5.3 Data Completeness

There are several parameters that can be used to assess the completeness and accuracy of data. A key problem is being able to identify the correctness of the data supplied. In the following sections I propose some parameters that could be collected and statistically verified.

These parameters could be systematically checked by a special DUA robot, but this is almost certainly not going to be acceptable to the DSA manager community. However, as identified earlier, measurement is critical if we are to really understand the system. To reconcile this, I propose that the parameters are defined by the systems administrator, and checked on a statistical basis, that is occasionally a small sample of data is looked at by a DUA, and the parameter verified. If the sample confirms the administrator declared figures, all is well. If the sample is not, the administrator is asked to re-check the declared parameters, or some other action, as declared in the SLA, is taken.

The following sub sections, describe some possible ways forward. More work will be required is significant progress is to be made.

5.3.1 Data Coverage

Data coverage falls into two categories. At the top level of the Directory hierarchy, it would be useful to know the chances of a given organisation being represented in the Directory. Then at the organisational level we need to know the probability of finding the data we required about a specific user.

Nationally, we could quote the percentage of all organisations that have data in the Directory, but

being realistic it will be many years before this will become a meaningful measure. The only thing you can really do is provide forecasts of how many organisations are expected in the Directory, within certain time frames, then measure the actual number of organisations represented. This is relatively easy to do, but gives a very real measure of service performance. If the forecast is not met, it is a definite indication that the service is not meeting the requirements of the community, and it is time to re-visit the service definition.

At the organisation level, we need to have an indication of the number of employees or organisational roles covered, and the data available for these entries. It is unlikely that organisations will present figures on how many entries they have within their databases. In [23] Kille proposes a mechanism for counting the DIT, one of the reasons this did not progress was commercial objection to making this data available.

This problem can be solved by using percentages. An organisation is less likely to have a problem with publishing "n% of employees are represented in the Directory". This is a parameter that would have to be supplied by an administrator, with no easy way of being verified externally. For verification a mechanism like that proposed by Kille would be needed. Without a direct verifiable measurement, it is not clear if this parameter will be of use in assessing the data completeness.

What could be verified is the completeness of the entries available. The SLA could declare that each user entry Directory should contain a specific set of attributes, for example: phone, and email addresses. The administrator could then provide information on the percentage of entries in the database that contains the required attributes. This can be measured on a statistical basis. (I believe there is an ITU work group specifying a minimal attribute set, as part of the F.500 standard.)

This information could be published within the X.500 Directory, by updating the ID [9] and issuing as an RFC.

If these parameters were collected from each organisation within the service, an overall data completeness figure for the national and consequently international Directory could be obtained.

5.3.2 Data Accuracy

As noted in ID on publishing information on the Internet [21], it is important to record the time a data item was last verified. If a user is aware that a phone number in the Directory has not been checked for two years, and when using the number is unable to get through, they might consider trying an alternative mechanism to locate the number. Similarly, it could be part of the SLA that data not verified every 12 months should be removed.

X.500(93) already gives us an operational attributes for "created time" and "modified time". Similar parameters are also defined in COSINE / Internet schema [22] for X.500(88) systems, but neither give us last verified time.

An extension to this, would be to add a lasted verified attribute into the X.500 schema definitions. This will allow some measure of the accuracy of data to be assessed.

As an alternative to adding a last verified attribute, the semantics of the existing modified time attributes could be changed to mean verified. When the data is verified, an operation to similar to the UNIX 'touch' operation could be performed. Within QUIPU implementations this can be achieved by deleting the last modified time attribute, the DSA will then automatically put it back, with the current time. Similar mechanisms are probably available in other implementations. This could be published by each organisation in the form of the date when data within the Directory was last verified. Each national service could then publish a set of data verification figures, for example percentage of data verified in the last 6 month. This would give a useful accuracy figure, a figure that could in principle be checked by a special purpose DUA based upon a statistical sample.

6 Conclusions

The Directory service is distributed. The management and monitoring of the service also needs to be distributed. In the previous sections I have identified some mechanisms that could be used to provide this monitoring.

One of the dangers of implementing the procedures proposed in the paper, is it could deter some organisations from joining the service. This is a risk that has to be balanced against the potential increase in the quality of the service. As

I claimed in the opening paragraphs, if the service is of sufficient quality, organisations cannot afford to be omitted from the service.

6.1 Centrally we need:

Centrally means both at national level by the service provider, and overall at the root level by the NameFLOW-Paradise project administrators.)

- A revised probing mechanism to asses reliability of organisational data at the country level.
- A revised probing mechanism to assess the reliability of International access points (i.e., different country data).
- New mandatory QOS attributes to reflect completeness and accuracy of organisational data.
- Forecasting of service uptake, with verified milestones.
- Service level agreements to define minimal acceptable requirements.

6.2 Each DMD needs to consider:

- DSA metrics definitions of the software used.
- Local DSA monitoring to ensure the DSA is available, for example, using SNMP or probing techniques.

7 Future work

In this Section, I identify some areas where further work is required to define our quality requirements, and I put forward some suggestions as to where we may be able to improve the QOS of the current NameFLOW-Paradise Directory.

7.1 DUA quality

In Section 4.4 I analysed the DSA metrics paper, a similar paper for DUA exists [10]. There is a need to review this RFC and see if it needs extending. It may not.

7.2 DUA Performance

As well as DSA performance, the DUA performance will also affect the view of the service from the users perspective. Does using a different DUA technology inherently mean the DUA will preform better. For example, does the use of LDAP imply a faster DUA than the use of LDAP? Some experimental evidence suggests this does, despite the extra server in the loop suggesting it cannot be the case.

7.3 DSA relaying

The work presented in this paper has focused upon the components of the Directory system. However, within the overall scenario, the Directory Service runs over a variety of network technologies, including TCP / IP, X.25 and CLNS. As described in [20] a set of heuristics can be built to allow Directory operations to proceed in such an environment, including the technique of DSA relaying as described in [19]. These heuristics and methods are implementation dependant. How do we analyse their effectiveness within the context of the overall service?

If the relay mechanism fails, we could find we have a service consisting of many disjoint data islands. The reliability metrics described in this paper may not be able to detect this. By its very definition, the problem is different for each DSA. Different users may see a very different service level.

7.4 Replication

Replication is a key part of providing a reliable service, but it can introduce problems of it own. If the replication agreement expires, or replication does not take place for some reason, the shadow data is likely to become out of date, reducing the accuracy of the data presented. Some monitoring of replication would be useful to ensure the replica data is accurate.

Off site replication is a sensible way of protecting against local network unavailability, but there are privacy concerns. Before an SLA recommends off site replication, these need to be understood.

7.5 Pointer Attributes

Included within the X.500 schema definitions are attribute whose values are or contain distinguished names, or pointer to other entries within the DIT. However X.500 does not provide a management environment for these attributes, so they can become out of date reasonably quickly. A suite of pointer checking tools would be useful, a prototype was developed by Steve Titcombe of UCL and myself, but did not progress to a completed state. There are also some simple checks the COSINE bulk loading tools make.

7.6 Security

Security is another area that needs to be considered. As part of their privacy requirements an organisation may only allow their data to be accessed using higher levels of authentication, or only allow selected sites access, or have a guardian / fire wall DSA to control access.

The effects of these parameters on the overall system will need to be taken into account.

7.7 Namespace.

Within X.500 the organisational namespace is very structured, typically in a country / organisation hierarchical schema, sometimes with locality levels as well. Whilst this provides for an easily managed system, it can make locating an organisation hard from a user perspective.

The 'power search' functionality of the NameFLOW-Paradise public access DUA has shown how users like the ability to search across a set of organisations. Extending this particular mechanism may not be cost effective, and solutions like index DSAs may provide a better end result.

7.8 Searchability.

If a DUI is monitoring the progress of user queries as suggested in some of the previous sections of this paper, it could potentially provide some measure of the success of typically search queries. I.e., measure the likelihood of a search operation succeeding.

7.9 1993 Protocols.

Using current DUA technology and an X.500(1993) model it is difficult to configure a system that allows a user to perform a one level search or list operation on the higher levels of the DIT, where the DIB is distributed across multiple DSAs. This is typically the case at the root level, or a country level. For example, a one level list or search operation below the c=GB node, for organisation name 'NEXOR' cannot easily be resolved by one single DSA.

You can only list or search a node if you hold the access control information relating to it. This access control information is stored in subordinate sub entries. A complete set of these sub entries will not typically be available in a DSA holding a country node, so the access rights cannot be determined locally.

The consequence of this is a one level search operation of a country node cannot normally be resolved by a single DSA. Projects like NameFLOW-Paradise have shown these high level search operations to be a key part of the overall service, allowing users to easily locate organisations within the DIT. Using the standard based mechanisms only, there are two methods by which this search can progress.

Either the DSA can perform multi-chaining. For each subordinate reference held, it would effectively chain a base object search operation to the appropriate DSA. Results will only be returned to the initiator when all DSAs have responded, or a time limit has been reached. Experience from the X.500 pilot projects shows it is rare for all DSAs to be available, so in practice, results will not be returned until the DUA requested time limit has been reached. With the user agent techniques currently used, relying upon these searches this is not a viable solution in a real-time environment.

Instead of multi-chaining, the DSA could return a set of continuation references to the initiator. However, few DUAs currently act upon continuation references. In time, this can be rectified, however it increases the DUA complexity.

It would be preferable for the DSA to be able to resolve the operations. NEXOR have proposed two minor modifications to the X.500(93) standard are made, which would allow a single DSA to resolve the queries.

For the list operation a DOP HOB agreement allows the entryACI to be passed from the subordinate DSA to the superior, suggesting this information can be used to resolve the list operation. However, it cannot as the prescriptive ACI information may also be required. The first defect report aims to resolve this, by passing the sub entry information in the HOB. This already happens in the HOB between superior and subordinate.

Secondly a HOB agreement between subordinate DSA and superior allows for attributes to be passed. Assuming the ACI is available using the mechanism in the previous paragraph, these attributes could be used for one level search operations when the 'copyShallDo' service control is set. However X.500(93) declares these attributes can only be used to resolve list operations, and not searches. The second defect report proposes this restriction is removed.

Monitoring the progress of these defects will be an important for the future of the NameFLOW-Paradise service.

8 Summary

In this paper, I have identified a set of clearly understood requirements for a service Directory system. I have reviewed how these requirements can be measured, to enable an assessment of whether the service provided meets the requirement. This way improvements and transitionary changes in the service can be tracked, allowing early detection of a lowering of the service level. This will lead to a more reliable Directory service, which in turn assists the Directory in becoming a key part of the infrastructure of the Information Super Hypeway. Only by providing a reliable, quality, Directory will the critical mass of data be achieved, at which point it will become a commercial necessity to have Directory access, firmly establishing the Directory as the backbone if the global information service.

Finally, the concepts presented in this paper reference X.500, but can easily be applied to other information services, such as the World Wide Web where unavailable and unmaintained servers are reducing the overall quality of the system thereby reducing the user perception of the service.

Glossary

DIT	Directory Information Base - the data used to build the DIT.
DIT	Directory Information Tree.
DMD	Directory Management Domain.
DUA	Directory User Agent - A process that interacts with a DSA using the Directory access protocol.
DUI	Directory User Interface - The tool a human user interacts with to access the Directory
DSA	Directory System Agent - A process that holds part of the distributed Directory, and services DUA requests.
ID	Internet Draft RFC.
MIB	Management Information Base used by SNMP.
NADF	North American Directory Forum
NSAP	Network Service Access Point - e.g, X.25, CLNS of TCP / IP address.
QOS	Quality of Service.
RFC	Request For Comments - Standards documents used within the Internet community.
SLA	Service Level Agreement
SNMP	Simple Network Management Protocol X.500 The OSI Directory Standard

References

[1] "Piloting A Researcher Directory Service in Europe", D. Goodman, UCL. PARADISE International Reports.

[2] "The ISODE User's Manual: Volume 5 - QUIPU", Colin Robbins, Steve Kille. July 1991.

[3] "The Directory - Overview of Concepts, Models and Service", CCITT Rec. X.500 1988 Edition | ISO / IEC 9594:1990 Edition.

[4] "The Directory - Overview of Concepts, Models and Service", ITU-T Rec. X.500 | ISO / IEC 9594-1:1993 Edition.

[5] "The World-Wide Web", Berners-Lee et al. Communications of the ACM, August 1994.

[6] "Strangers in Paradise", Bruno Koechlin, Katy Treca, Paul-Andre Pays, INRIA WWW Server, 1994.

[7] "Providing the X.500 Directory User with QOS Information", Paul Barker, University College London. Computer Communication Review, late summer, 1994.

[8] "X.500 Directory Monitoring MIB". G.Mansfiled, S.Kille. RFC 1567.

[9] "Handling QOS in the Directory". Steve Kille, 1991. Expired INTERNET-DRAFT.

[10] "DUA Metrics", P.Barker. RFC 1431.

[11] "DSA Metrics", P.Barker, R.Hedberg. RFC 1567.

[12] "Recommendations for an X.500 Production Directory Service", Russ Wright et al. INTERNET-DRAFT. November 1994.

[13] "Phasing Out The Root DSA", David Chadwick, Email message to OSI-DS, October 1994.

[14] "Principles of Software Engineering Management", Tom Glib, Addison-Wesley, 1988, ISBN 0-201-19246-2.

[15] "ISO transport services on top of the TCP: Version 3", M.T.Rose, D.E.Cass. RFC 1006.

[16] "Encoding Network Addresses to support operation over non-OSI lower layers", S.E. Hardcastle-Kille. RFC 1277.

[17] "DSA Benchmarks", NEXOR WWW server, URL=http: / / web.nexor.co.uk / users / cjr / dsamet / sun.html

[18] "Software Engineering", Martin L Shooman, McGraw-Hill. ISBN 0-07-057021-3.

[19] "Managing the International X.500 Directory Pilot", Colin Robbins, Proceedings of EurOpen Conference 1991.

[20] "Directory Navigation in the Quipu Directory System", Colin Robbins, Paul Barker. UNIX & Connectivity, November 1989.

[21] "Publishing Information on the Internet with Anonymous FTP", P. Deutsch, A. Emtage, M. Koster, M.Stumpf, INTERNET-DRAFT, September 1994.

[22] "The COSINE and Internet X.500 Schema.", P. Barker, S. Kille RFC 1274.

[23] "Counting the Directory Information Tree", S.E. Hardcastle-Kille, INTERNET-DRAFT, April 1992.

[24] "NADF Standing Documents: A Brief Overview", The North American Directory Forum, RFC 1417.

[25] "Service Description", SD-3, NADF.
[26] "Monitoring the Quality of Service of the NameFLOW-Paradise Directory.", Colin Robbins, NEXOR 1995.

Footnotes

(1) Delivering Advanced Network Technology to Europe Ltd.
(2) The tool a human user interacts with to obtain access to the Directory, as opposed to a DUA which is an agent the DUI and other processes use to access the Directory.
(3) Access is the generic sense, of first searching for it, then recovering the attributes
(4) DE is the specific DUI used to implement the QOS database.