

Network Working Group
Request for Comments: 1564
Category: Informational

P. Barker
University College London
R. Hedberg
Technical University Delft
January 1994

DSA Metrics
(OSI-DS 34 (v3))

Status of this Memo

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

Abstract

This document defines a set of criteria by which a DSA implementation may be judged. Particular issues covered include conformance to standards; performance; demonstrated interoperability. The intention is that the replies to the questions posed provide a fairly full description of a DSA. Some of the questions will yield answers which are purely descriptive; others, however, are intended to elicit answers which give some measure of the utility of the DSA. The marks awarded for a DSA in each particular area should give a good indication of the DSA's capabilities, and its suitability for particular uses.

Please send comments to the authors or to the discussion group
<osi-ds@CS.UCL.AC.UK>.

Table of Contents

1.	Overview	2
2.	General Information	3
3.	Conformance to OSI Standards	4
3.1	Directory protocols.....	4
3.2	Implementors' agreements and profiles	6
3.3	Protocol stacks.....	6
3.4	DIT structure	7
4.	Other protocols	7
5.	Extensions to the 1988 Standard	7
5.1	Schema	7
5.2	Support for replication.....	8
5.3	Support for access control	8
5.4	Miscellaneous	9
6.	Miscellaneous characteristics	10

7.	Management tools	11
7.1	Dynamic system management	11
7.2	Static system management	12
7.3	Data management.....	12
8.	Operational Use	12
9.	Interoperability	12
10.	Performance	13
10.1	Speed for various operations	14
10.1.1	Bind	14
10.1.2	List	15
10.1.3	Search	15
10.1.4	Read	16
10.1.5	Add entry.....	16
10.1.6	Modify entry	16
10.1.7	Modify RDN	16
10.1.8	Query rate	17
10.2	The results.....	17
10.3	Environment used for benchmarking	17
11.	Security Considerations	21
12.	Authors' Addresses	21

1. Overview

The purpose of this document is to define some metrics by which DSA products can be measured. Such metrics are valuable as whilst an X.500 DSA must conform to the specification in the standard - this is a sine qua non - protocol conformance is not in itself the hallmark of a usable implementation. A DSA must perform operations within a reasonable time; a DSA must offer good throughput of queries; a DSA must be able to handle a reasonable volume of data; if modification operations are provided, some sort of access control must be provided; a DSA and its data must be manageable.

In many respects, it is almost impossible to say that one DSA is better than other from looking at the responses to questions in this document. For some, the cost or level of support will be the key criterion. For another user, the flexibility of the schema management facilities, or the feasibility of running the DSA over an existing relational database, will be of prime importance. In many respects DSAs will just be different, rather than better or worse. However, all other things being equal, the look-up speed of a DSA is very obviously measurable, and there is a substantial number of questions on the speed of the various X.500 operations, and in particular on the look-up operations.

Throughout this document, some of the questions posed are annotated with a square-bracketed points score and an explanation as to how the points should be allocated. For example, a question might be

appended with "[2 if yes]", indicating score 2 points for an affirmative answer to that question. These points scores should be collated in Table 1 at the end of the document. The questions on DSA performance are judged to be important enough to have a separate table for those results: they appear in Table 2 (and optionally Table 3). Together, these tables constitute a measure of the DSA.

The metrics are on a section by section basis, which should help the reader who is seeking, for example, a DSA with fast look-up capabilities and extensive access control facilities, to focus on the critical aspects of a DSA for their particular requirement. No conclusions should be inferred from adding the scores together into one overall grand total and comparing such totals for different DSAs, as no attempt is made to assign weights to the different characteristics.

Whilst much of this document should usually be completed by the developers or suppliers of an implementation, the section on performance could be completed by anyone running the implementation. Indeed, it will be beneficial if several sets of performance figures can be gathered for each implementation, for a variety of hardware platforms.

2. General Information

This section contains general information about the implementation under discussion.

1. Name of the information provider
2. Name of the implementation
3. Version number of the DSA described in this document
4. Are there plans to implement the additional features describe in the 1992/3 standard? [6 for full implementation, 4 if both access control and replication to be implemented, 2 for some 1992 features]
5. Name and address of supplier or person to contact

6. Describe the hardware and software platforms supported by the DSA
[up to 4 points may be awarded for this question]
 - (a) Hardware (If appropriate, can summarise as, for example
"generic UNIX platform")
 - (b) O/S (state version if critical)
 - i. UNIX) (be sure to indicate which flavour - e.g.,
SYSV [1], BSD [1], SUNOS, etc)
 - ii. VMS) [1]
 - iii. MS-DOS [1]
 - iv. Macintosh [1]
 - v. Other) [1]
7. Name any other software required to run the system which is not
supplied with the operating system or with the DSA software
itself. Examples might include a database package, or
communications software
8. Is this DSA an integrated part of a software package, and in such
case which ?
9. Is the software free? If the DSA needs other packages, are these
also freely available? [3 if completely free, 1 if requires
commercial software package]
10. Is commercial support available for this implementation? [3] ...
11. Is free, best effort support available from the developers? [2].
12. Is free support available via user groups or email lists? [2] ..
3. Conformance to OSI Standards
- 3.1 Directory protocols
13. Does the DSA implement DAP?
 - (a) Read ASE? [2]

- (b) Search ASE? [2]
- (c) Modify ASE? [2]
- 14. Does the DSA implement DSP?
 - (a) Chained read ASE? [2]
 - (b) Chained search ASE? [2]
 - (c) Chained modify ASE? [2]
- 15. Statement requirements according to section 9.2.1 in X.519.
 - (a) Supported application-contexts?
 - (b) Capable of acting as first-level DSA? [1]
 - (c) Chained mode supported? [1]
 - (d) Security-level(s) supported? [1 for strong + 1 for protected simple + 1 for simple authentication]
 - (e) All attribute types according to X.520? [1]
 - (f) All object classes according to X.521? [1]
- 16. Does the implementation meet the conformance clauses in section 9.2.2 and 9.2.3 of X.519?
Static requirements [2 if yes on all]
 - (a) Abstract syntaxes of application contexts
 - (b) Abstract syntaxes of information framework
 - (c) Minimal knowledge
 - (d) Support of root context
 - (e) Abstract syntax - attribute types
 - (f) Abstract syntax - object classesDynamic requirements [2 if yes on all]
 - (a) Mapping onto underlying services
 - (b) Distributed operations - referrals

- (c) DirectoryAccessAC - referrals
- (d) DirectorySystemAC - referrals
- (e) Chained mode

17. Please list all conformance testing work applied to the implementation (specify conformance test version number). [2 if any testing]

.....
.....
.....
.....

3.2 Implementors' agreements and profiles

Does the DSA conform to the following implementors' agreements? If so, state parts and version numbers.

18. EWOS? [1]
.....
.....

19. OIW? [1]
.....
.....

Does the DSA conform to the following profiles? If so, state which version numbers.

20. UK GOSIP? [1]

21. US GOSIP? [1]

State any other GOSIP profiles to which the DSA conforms

3.3 Protocol stacks

22. Which of the following transport and network layer protocols does the DSA support:

(a) TP.x over CONS (state transport class)? [2]

(b) TP.4 over CLNS? [2]

(c) TP.x over X.25(1980) (state transport class)? [2]

3.4 DIT structure

23. A suggested DIT structure, detailing an object class hierarchy, is presented in X.521. Does the DSA:

(a) Enforce this hierarchy?

(b) Allow the enforcement of this hierarchy?

24. Are structure rules optional or mandatory?

4. Other protocols

25. Not everybody uses OSI protocols at the network layer. Does the DSA support other "network" layer protocols?

(a) TP.0 over RFC1006 over TCP/IP [3]

(b) State any other options supported.

26. Does the DSA also run over any lightweight stack? If so, describe it with reference to the OSI seven layer model [1]

27. Can local DUAs access the DSA directly by some method of inter-process communications? [1]

5. Extensions to the 1988 Standard

5.1 Schema

28. Does the DSA fully support RFC1274, "The COSINE and Internet X.500 Schema"? [2]
If not, please supply a list of all those object classes, attribute types and attribute syntaxes in RFC1274 which are supported on a separate sheet. This might be summarised by saying, for example, "all those with standard attribute syntaxes", or "all except fooBar".

29. Does the DSA implement the schema management defined in the 1992 standard? [2]

30. If not, is the schema stored in the Directory? In a distributed manner[2] or centralised[1] ?

31. Can a DSA manager extend the schema and add new

- (a) Attribute types with existing syntaxes? With compilation [1], or without compilation [2]
 - (b) Attribute syntaxes? With compilation [1], or without compilation [2]
 - (c) Attribute sets? With compilation [1], or without compilation [2]
 - (d) Object classes? With compilation [1], or without compilation [2]
32. Is it possible to add in or modify DIT structure rules, with compilation [1], without compilation [2]

5.2 Support for replication

33. Does the DSA support the replication mechanisms as described in the 1992 standard [2]?
34. Does the DSA support any other replication mechanisms?
- (a) Replication part of RFC1276 [2]
 - (b) Other (please give a reference to any description of the mechanisms, and indicate whether these mechanisms are used by any other implementations) [1 for any mechanism]
35. If the DSA supports replication, does it support:
- (a) Replication of a single entry? [2]
 - (b) Replication of a set of sibling entries? [2]
 - (c) Replication of a subtree? [2]

5.3 Support for access control

36. Does the DSA support access control as described in the 1992 standard [3]?
37. If not, does the DSA have any access control mechanisms at all?

[2]

38. If yes, does the access control scheme support the following:

- (a) Allow a user to maintain their own entry? [1]
- (b) Allow a user to maintain some attributes in their own entry, but not all attributes? [1]
- (c) Give management rights to a DSA manager in a fashion analogous to the privileges given to a UNIX super-user? [1]
- (d) Give management rights to a data manager on a per subtree basis? [1]
- (e) Give management rights (to an entry, group of entries, subtree, etc) to a group of users? [1]
- (f) Give access rights to users on the basis of the leading portion of their Distinguished Name? [1]
- (g) Is it possible to define a protection mechanism for each individual attribute type in one entry? [1]
- (h) Maximum number of Distinguished Names that can be defined for one access right to one attribute in one entry? If unlimited, state the constraints. [1 if more than 6 DNs are feasible] :
.....
- (i) Does the DSA support the extended access control techniques described in "An Access Control approach for Searching and Listing" by Hardcastle-Kille and Howes, in the Internet Draft, OSI-DS 21? [2]
.....
- (j) If there are features of the access control mechanisms which are not brought out by the above questions, please describe these additional features [up to 2 for wonderful additional features!]
.....
.....

5.4 Miscellaneous

39. Does the DSA fully support RFC1276, "Replication and Distributed Operations extensions to provide an Internet Directory using X.500"? [2] If not, please give a list of features that are supported.

.....

40. If the DSA uses RFC1006 and/or X.25(1980) at the network layer, does the DSA conform to RFC1277, "Encoding Network Addresses to support operation over non-OSI lower layers" [3]

6. Miscellaneous characteristics

41. Does the DSA use its own database, or can it be used in conjunction with a general-purpose database package such as Oracle? [1 for own, 1 for ability to map onto general purpose databases, 1 if any such mappings have been made]
42. If the DSA runs as a static server, state the start-up time for a DSA with a database of 20000 entries. If this varies widely according to configuration options, give figures for the various options.
43. What is the maximum number of simultaneous associations that the DSA may have open? [1 if more than 15 associations]
44. Maximum database size, in entries, megabytes, or as appropriate. If none, state what the constraints are. [1 if a database of more than 100,000 entries is feasible]
45. What is the run-time size of an entry as specified in section 10 (on performance)? This should be the marginal size of an entry and thus should include the overhead of default indexes, etc. ..
46. What is the on-disk database size of an entry as specified in section 10 on performance?
47. Does the DSA make of indexing? [2 if yes]

If so:

- (a) Can the database be fully inverted? [1]
- If not, state for which:
- i. attributes indexes are automatically built
-
-
- ii. attributes/attribute syntaxes indexes may be built
-

-
- (b) Does the index improve performance on:
 - i. Exact match [1]
 - ii. Leading substring match [1]
 - iii. Approximate match [1]
 - iv. Any substring match [1]
 - v. Trailing substring match [1]
 - (c) What is the increase in run-time size of an entry when adding an index?
.....
 - (d) What is the increase in on-disk database size of adding another index?
.....
48. What sort of approximate match algorithm does the DSA use?
Describe it briefly
.....
.....
49. Does the DSA attempt to use relay DSAs (which have access to more than one network) in order to achieve connectivity with DSAs which are not on the same network? [2]

7. Management tools

7.1 Dynamic system management

50. Are there tools for monitoring DSA activity, using:
- (a) DAP? [1]
 - (b) CMIP? [1]
 - (c) SNMP? [1]
51. Are there tools for controlling a run-time DSA? [2]

7.2 Static system management

52. If knowledge information is stored within the DIT, are there tools for knowledge management? [2]
53. Are there tools for checking that attributes with Distinguished Name syntax contain values of entries in the DIT (i.e., they do not contain "dangling pointers")? [1]

7.3 Data management

54. If the DSA doesn't use a general-purpose database package, what data management tools are available?
.....
55. Are there any tools for arboriculture - the moving, copying or deleting of DIT subtrees? [2]

8. Operational Use

The DSA may have lots of wonderful features -- on paper! But has the DSA been shown to work in practice? The following measures are intended to give some measure of confidence that the DSA's viability has been demonstrated.

56. How many entries in the largest DSA in use in operational use? :
57. What is the largest set of DSAs supporting an organisation?
58. What is the estimated number of organisations using this implementation for service use? [8 if more than 100 organisations, 5 if more than 50 organisations, 3 if more than 20 organisations, 2 if more than 5 organisations, 1 if more than 1 organisation]
59. Is this DSA used commercially with an installed base of more than 10 customers? [2]

9. Interoperability

The X.500 Directory is the OSI Directory. OSI stands for Open Systems Interconnection -- DSAs have to be able to inter-operate. They also have to be seen to interoperate.

60. Is this DSA in use in X.500 pilots?
- (a) Is this DSA in use anywhere in the COSINE/Internet Pilot? [3]

.....
(b) Is this DSA in use in any other major pilot? [2]

61. Name any other systems which you believe the system to
interoperate with. (It is not sufficient to say "any system
which supports the conformance clauses ...")

62. Please name all interoperability testing applied to the
implementation, specify test suite and what other implementation
that was used [1 per implementation, up to maximum of 5]

10. Performance

This section should give an outline to the expected performance of the DSA. A number of operations are timed in order to give a feel for the DSA's speed and throughput. Note that all operations should be resolvable within a single DSA. Chaining and referral are not assessed, although it should be possible to infer, albeit approximately, the speed of distributed operations.

- i. The tests should be made against an organisational database of 20000 entries. Some tests are against subsets of this data, and so the database should be set up according to the following instructions.

Create an organisational DSA with 20000 entries below the organisation node. Sub-divide this data into a number of organisational units, one of which should contain 1000 entries, another of which should contain 100 entries, and a third which should contain just 10 entries. The entries, which should differ, should be created with the following attributes:

- (a) Common Name
- (b) Surname
- (c) Telephone number
- (d) Postal Address (of 100 characters)

(e) Object class

- ii. In all the tests, two timings should be taken. In order to normalise the test results as much as possible, it is suggested that these tests be undertaken on an otherwise lightly loaded machine.
 - (a) A typical "cold start" reading should be given. In this case the system will not have the advantage of any benefits that derive from operating system paging, or caching.
 - (b) A best possible figure should be given, which indicates the upper limit of DSA performance.
- iii. The timings should relate to the default set-up, and should be entered in Table 2. If significant performance gains can be made by use of configuration options, such as building extra indexes to support searches, measures of the improved performance may also be given, and should be entered in Table 3. Attention should be also drawn to any optimisations, heuristic or otherwise, which are not evidenced in the following tests.
- iv. Please note that the tests should be made using a DUA and DSA with full 7-layer stacks, rather than some lightweight protocol.

10.1 Speed for various operations

The tests are described, one subsection per operation. The results should be entered in Table 2 (and Table 3 if a non-default set-up is also measured).

10.1.1 Bind

The time it takes for a DUA to bind to the Directory. This time should include all the initialisation time a DUA process needs before it can query the Directory: e.g., reading of tailor files, schema information, etc. Give the bind time for each of the following levels of authentication. State "n/a" if the implementation does not support a particular level of authentication.

63. Anonymous

64. Simple

65. Simple protected

66. Strong

10.1.2 List

Give the time for listing a set of organisational unit sibling entries.

67. 10 entries

68. 1000 entries

10.1.3 Search

In this section, two sets of search operations should be performed on the DSA.

i. A single level search of 100 entries within an organisational unit.

ii. An organisation subtree search, on the subtree of 20000 entries.

The following searches should be tried. Unless otherwise stated, the "XXX" or "YYY" part of the search filter should be chosen in such a way as to return a single result. Unless stated otherwise the results should return all attributes for the entry.

69. Exact match for a surname:

surname=XXX

70. Leading substring match for a common name:

commonName=XXX*

71. Any substring match for a common name:

commonName=*XXX*

72. Trailing substring match for a common name:

commonName=*XXX

73. Approximate match for a common name:

commonName"=XXX

74. More complex filter, searching by object class and two other attribute types:

objectClass=person AND
(commonName=XXX* OR telephoneNumber=*YYY)

75. Search returning all entries (i.e., 100 entries in the single level search, and all 20000 entries in the subtree search:

objectClass=*

In this case, no attribute values should be returned in the result set.

10.1.4 Read

76. A single read operation, returning all attributes.

10.1.5 Add entry

77. Add an entry beneath an entry which has:

- (a) 0 children
- (b) 10 children
- (c) 1000 children

10.1.6 Modify entry

Modify an attribute value, other than an RDN value, for an entry which has

- 1. 10 siblings
- 2. 1000 siblings

78. Modify an entry

- (a) Add description attribute
- (b) Remove description attribute

10.1.7 Modify RDN

Modify an RDN value for an entry with the following number of siblings.

79. Modify RDN

- (a) 10 siblings
- (b) 1000 siblings

10.1.8 Query rate

As the time taken for a single read will usually be negligible, the following list and set of reads should give a clearer indication of the query rate.

- 80. A list to return 100 entries for persons, and then a read of each entry returning all attribute values.

10.2 The results

The results of the tests just described should be entered in Table 2 (and optionally Table 3), at the end of the document.

10.3 Environment used for benchmarking

Date of test.....
Name of tester
The results will be directly correlated to the test set-up used, and in particular, the hardware. Please answer the following questions to describe the test environment:

- (a) Processor (make and model)
- (b) Processor speed (MIPS)
- (c) Primary memory available
- (d) If disk-based DSA, disk I/O interface and disk speed
- (e) O/S version
- (f) Network type and bandwidth (e.g., 10 Mbit Ethernet)
- (g) Protocols in transport layer and below (e.g., TP 0, RFC1006, TCP/IP)
- (h) How/where timings obtained?
 - o C procedural interface
 - o DUA shell (e.g., Quipu's DISH)

Section			Points	
No.	Description		Maximum	Scored
2	General Information		20	
3	Conformance to OSI		35	
4	Other protocols		5	
5	Extensions	Schema	16	
	to the	Replication	10	
	1988	Access Control	15	
	standard	Miscellaneous	5	
6	Miscellaneous characteristics		15	
7	Management tools		10	
8	Operational use		10	
9	Interoperability		10	
10	Performance		see table 2	

Table 1: DSA Metrics

Operation	Cold DSA		Optimum Performance	
Bind				
--Anonymous	
--Simple	
--Simple Prot	
--Strong	
List				
-- 10 entries	
-- 1000 entries	
Search	single level	subtree	single level	subtree
-- exact
-- leading sub
-- any sub
-- trailing sub
-- approx
-- complex
-- return all
Read	
Add				
0 siblings	
10 siblings	
1000 siblings	
Modify				
10 siblings	
1000 siblings	
Modify RDN				
10 siblings	
1000 siblings	
Query rate	

Table 2: Speed of operations - default set-up

Operation	Cold DSA		Optimum Performance	
Bind				
--Anonymous	
--Simple	
--Simple Prot	
--Strong	
List				
-- 10 entries	
-- 1000 entries	
Search	single level	subtree	single level	subtree
-- exact
-- leading sub
-- any sub
-- trailing sub
-- approx
-- complex
-- return all
Read	
Add				
0 siblings	
10 siblings	
1000 siblings	
Modify				
10 siblings	
1000 siblings	
Modify RDN				
10 siblings	
1000 siblings	
Query rate	

Table 3: Speed of operations - non-default set-up

Security Considerations

Security issues are not discussed in this memo.

Authors' Addresses

Paul Barker
Department of Computer Science
University College London
Gower Street
London
WC1E 6BT
United Kingdom

Phone: +44 71 380 7366
Fax: +44 71 387 1397
EMail: P.Barker@cs.ucl.ac.uk

Roland Hedberg
Rekencentrum
Delft Technical University
Michiel de Ruyterweg 10-12
Postbus 354, 2600 AJ Delft
The Netherlands

Phone: +31 15 785210
EMail: Roland.Hedberg@rc.tudelft.nl

OR

Roland Hedberg
Umdac
University of Umea
s-901 87 Umea
Sweden

Phone: +46 90 165204
EMail: Roland.Hedberg@umdac.umu.se