# A Comparative Performance Analysis of 7 Lightweight Directory Access Protocol Directories

E. J. Thornton
*IS Institute, University of Salford, The Crescent, Salford M5 4WT*
*e.j.thornton@pgt.salford.ac.uk*
*Tel No: 0161 2955664*
*Fax: 01484 532930*

D. Mundy
*IS Institute, University of Salford, The Crescent, Salford M5 4WT*
*d.mundy@salford.ac.uk*
*Tel No: 0161 2955664*
*Fax: 01484 532930*

D. W. Chadwick
*IS Institute, University of Salford, The Crescent, Salford M5 4WT*
*d.w.chadwick@salford.ac.uk*
*Tel No: 0161 2955351*
*Fax: 01484 532930*

## Abstract

This paper compares the performance of 7 different LDAP servers. The performance tests cover loading, searching (of both indexed and un-indexed attributes), adding, deleting and modifying the data. Tests were performed on directory trees containing up to 10 million entries. Overall, Critical Path had the best performance for up to 100K entries, and SunONE Directory Server had the best performance for over a million entries, although each directory, except one, showed it could achieve good performance in at least one test.

## Keywords

LDAP, performance, DirectoryMark

## 1. Introduction

We have designed and constructed a system for the Electronic Transfer of Prescriptions (ETP) [1] based around the storage of electronic medicinal prescriptions in a Lightweight Data Access Protocol (LDAP) directory [7]. The digitally signed prescriptions are created as X.509 attribute certificates [5], which makes LDAP directories ideal as the storage medium. Consequently we needed to find the top performing LDAP directory in order to carry out evaluative performance tests of our ETP system. The performance and scalability requirements of the ETP project are stringent; the system needs to hold millions of entries, support a large number of searches (well suited to LDAP directories) and a large volume of add/delete/modify operations (perhaps better suited to traditional RDBMSs)[3, 4].

This paper describes the results of the LDAP directory testing undertaken as part of initial tests for our ETP system. The first part of the paper provides an introduction to LDAP directories including an overview of the seven LDAP servers selected for our performance trials. Following this we describe the design of our performance tests and what we wished to gain from the results. Finally, the results

are displayed and critically analysed. We conclude with an evaluation of which directories are the best overall performers and which directory would be best suited for the prescription stores required within our ETP system.

## 2. Lightweight Data Access Protocol

### 2.1. The LDAP Directory

An LDAP (Lightweight Directory Access Protocol) directory is a specialised type of database typically read or searched more often than it is updated. The directory uses a hierarchical data model originally defined in the X.500 standard [6]. LDAP is the Internet standard protocol used for communicating with the X.500 directories. Originally used for relatively simple applications such as address book type lookups for telephone numbers, X.500/LDAP directories now provide far more complex functionality such as user authentication using Public Key Infrastructures [5].

### 2.2. Directories Selected

The following directories were available for evaluation:

| **Directory/Vendor** | **Operating System** | **Notes** |
|---|---|---|
| Critical Path InJoin Directory Server 4.0 | Windows 2000 Server | Loaned for evaluation from Critical Path @ http://www.cp.net |
| IBM SecureWay Directory 3.2.2 | Windows 2000 Server | Free full product download available at http://www-3.ibm.com/software/network/directory/ |
| iPlanet/SunONE Directory Server 5.1 (evaluation)* | Windows 2000 Server | Free trial download available at http://wwws.sun.com/software/products/directory_srvr/home_directory.html |
| Microsoft Active Directory | Windows 2000 Server | Integrated into Windows 2000 operating system. |
| Novell e-Directory 8.6 | Windows 2000 Server | Free full product download available at http://www.novell.com |
| OpenLDAP 2.0.23 | RedHat Linux 7.2 | Free to full product download and source code available at http://www.openldap.org/ |
| Syntegra Aphelion 2002 | Windows 2000 Server | Loaned for evaluation from Syntegra @ http://www.syntegra.com |

Table 1 – Directories Tested

It should be noted that Syntegra Aphelion 2002 was a brand new release and there is now a performance improvement patch available which was not available to us at the time of testing. IBM SecureWay and Novell e-Directory are free fully functional products, and OpenLDAP is also free as it is an open source solution. SunONE Directory Server was free to download as an evaluation version and Active Directory comes as part of the Windows 2000 operating system.

# 3. Testing Methodology

## 3.1. Measurement Tools

A generic testing client was needed to fairly evaluate each directory. MindCraft's DirectoryMark 1.2.1 [2] was used for both LDAP Data Interchange Format (LDIF) file generation and simulating all tests, as it is widely recognised as the industry performance benchmark for LDAP directories.

## 3.2. Directory Schema

The schema used for testing is depicted in Figure 2 below. The directory schema uses X.500 standard directory classes, implemented in all directories tested – country (gb), organization (Salford), organizationalUnit and organizationalPerson.

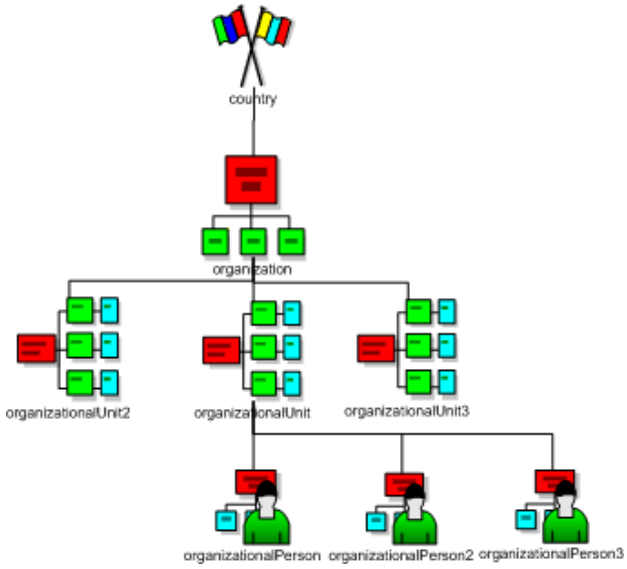


Figure 1 – Directory Schema for Testing

Ten organizationalUnits were added under the organization (Salford) entry in all tests. Four different numbers of organizationalPerson entries were tested:

- 10,000
- 100,000
- 1,000,000
- 10,000,000

## 3.3. The Hardware/Software Configuration

The server used for testing was a Dell PowerEdge 1400SC Server, with the following configuration:

- Intel Pentium 3 - 1GHz
- 512MB RAM
- 17GB SCSI Hard Drive (Fujitsu MAN3148MP)
- 36GB SCSI Hard Drive (Seagate ST336704LW)
- 120GB IDE Hard Drive (Maxtor 4G160J8)
- Microsoft Windows 2000 Server/Red Hat Linux 7.1 Dual Partitioned Operating System

The client used for testing had the following configuration:

- Intel Pentium 3 - 650MHz
- 128MB RAM
- 19GB IDE Hard Drive (Maxtor 32049H2)
- Microsoft Windows 2000 Operating System

## 3.4. The Requirements for ETP

The EPP system we have requires a specific set of LDAP directory operations.

**Requirement 1: Performance Of Full DN Lookup**

Prescriptions are stored in a LDAP directory under the tree OU=Prescriptions, O=EPP, C=GB and all are given their own Unique Identifier (UID) to form a Distinguished Name (DN) of UID=<unique number>, OU=Prescriptions, O=EPP, C=GB. Consequently all prescriptions can be retrieved via a read type operation i.e. a base entry search on the full DN.

**Requirement 2: Performance Of Prescription Addition**

Prescription entries need to be added to the LDAP directory.

**Requirement 3: Performance Of Prescription Deletion**

Prescription entries need to be deleted from the LDAP directory.

**Requirement 4: Performance Of Prescription Modification**

In particular circumstances a Pharmacy may only partially dispense a prescription and tell the patient to come back for the rest of their prescribed drugs at a later date. When this occurs the ETP system modifies the existing prescription entry to add a partially dispensed attribute.

**Requirement 5: Performance Of Directory With Large Datasets**

The ETP design encourages the NHS to provide regional repositories for the aggregation of prescriptions. Therefore for one of these directories to operate as a regional repository they must have consistently good performance at data entry levels from 500K to 20 million entries.

## 3.5. Test Design

A set of tests were defined based on the requirements of the ETP project – search, add, delete, modify and a set of mixed operations. In addition, the time taken to bulkload each directory was tested, using the proprietary LDIF loading tool included with each directory. It is important to note that LDIF bulkloading is not actually a requirement for the ETP system, and that LDIF is not necessarily the fastest bulkloading system for any given directory, as many suppliers provide their own purpose built proprietary mechanism. However, LDIF bulkloading is a standard mechanism supported by every system, therefore it can be used as a comparator. Further, bulkloading (LDIF or proprietary) will be required in any future ETP system in order to store information about prescribers, dispensers and patients. In Aphelion's case bulkloading is the only way of restoring a directory from its backup. However backup and restore operations were not timed in these sets of LDAP performance tests.

| Test No. | Description | Clients/Threads |
|---|---|---|
| S.1 | Simulated Read (i.e. base entry Search on distinguished name ) | 1/50 |
| S.2 | Full subtree Search with exact match on common name – an indexed attribute | 1/50 |
| S.3 | Full subtree substring Search on common name - an indexed attribute | 1/50 |
| S.4 | Full subtree substring Search on telephoneNumber - an un-indexed attribute | 1/50 |
| A.1a | Add organizationalPerson entry – indexed | 1/1 |
| A.1b | Add organizationalPerson entry - un-indexed | 1/1 |
| D.1a | Delete an entry – indexed | 1/1 |
| D.1b | Delete an entry - un-indexed | 1/1 |
| M.1 | Modify an indexed attribute (cn) | 1/10 |
| M.2 | Modify an un-indexed attribute (telephoneNumber) | 1/10 |
| X.1 | Mixed operations – Search, Add, Delete, Modify, ModifyRDN | 1/1 |

Table 2 – Description of Tests

Indexing was used only on the common name (cn) attribute (substring, equality and presence), and the un-indexed attribute used for testing was telephoneNumber. Each directory had different tuning options described fully in Appendix A.

# 4. Results

## 4.1. Bulkloading Results

The directories were bulk loaded from LDIF files, generated using DirectoryMark. The results from bulk loading can be seen in Tables 3 and 4. All directories' load times were tested under both indexed and un-indexed conditions, with the exception of Syntegra Aphelion, which defines its own indexes when none are defined by the administrator. Therefore Aphelion was only tested in the indexed state.

| | **10K** | **100K** | **1 million** | **10 million** |
|---|---|---|---|---|
| Critical Path IDS 4.0 | 00:01:32 | 00:22:31 | 11:00:34 | - |
| IBM SecureWay Directory 3.2.2 | 00:01:58 | 00:14:04 | 02:21:58 | - |
| iPlanet/SunONE Directory Server 5.1 (evaluation) | 00:00:10 | 00:01:22 | 00:12:13 | 29:54:13 |
| Microsoft Active Directory | 00:05:03 | 00:61:54 | 22:36:06 | - |
| Novell eDirectory 8.6 | 00:14:12 | - | - | - |
| OpenLDAP 2.0.23 | 00:00:37 | 00:08:36 | 13:12:36 | - |
| Syntegra Aphelion 2002 | 00:00:07 | 00:00:35 | 00:04:29 | 01:54:05 |

Table 3 – Indexed Directory Load Times (HH:MM:SS)

| | **10K** | **100K** | **1 million** | **10 million** |
|---|---|---|---|---|
| Critical Path IDS 4.0 | 00:01:12 | 00:09:11 | 01:35:49 | - |
| IBM SecureWay Directory 3.2.2 | 00:01:49 | 00:12:57 | 02:08:12 | - |
| iPlanet/SunONE Directory Server 5.1 (evaluation) | 00:00:08 | 00:01:02 | 00:09:10 | 02:08:12 |
| Microsoft Active Directory | 00:04:55 | 00:54:44 | 21:01:33 | - |
| OpenLDAP 2.0.23 | 00:00:14 | 00:01:15 | 02:01:11 | - |

Table 4 – Un-Indexed Directory Load Times (HH:MM:SS)

IBM SecureWay Directory was a fast loader, however attempts to load 10 million entries were unsuccessful. IBM's manuals recommended splitting the LDIF file into smaller LDIF files of around 250,000 entries. Consequently the original LDIF was split into 40 small ones and each loaded from the command line using a batch script. Whilst the loads of 10K, 100k and 1 Million entries into IBM's directory had no complications, the load of 10 million repeatedly failed. Up to this point, IBM SecureWay had shown almost totally linear loading times in relation to each other (estimated loading time for 10 million LDIF file is approximately 22 hours).

SunONE Directory Server showed a linear increase in time between 10K, 100K and 1 million entries, however the load of 10 million entries took around 14 times longer than would have been expected from the previous loads. Microsoft's Active Directory showed the worst performance of all for 10K, 100K and 1 million entry loads. For this reason a load of 10 million entries was not even attempted as we estimated it could take the best part of a month (and certainly in excess of 20 days). The LDIF used to load Active Directory required modification, as some of the SAM account names were too long, and dots (.) and slashes (-) in the attributes also caused errors in the load. Although Critical Path loaded faster than Microsoft Active Directory, it was also felt that it would take too long to load 10 million entries (in excess of 14 days). Critical Path InJoin Directory Server does provide 2 other methods of bulkloading, however these are based on a proprietary format, not LDIF, so they were not examined.

A load of 10 million entries was attempted with OpenLDAP (estimated to be in excess of 37 days). Previous loads (10K, 100K and 1 million) had been done from a single LDIF file (with and without indexing). After several failed attempts at this for 10 million entries, the LDIF was split in the same way as for IBM SecureWay, and loads were attempted using a batch script. Indexing was turned off for the load.. Unfortunately and inexplicably the directory kept on failing at various points during the load, and even when the option to ignore errors was given the database was still corrupted and impossible to index.

Syntegra Aphelion was the fastest loader by far – 15 times faster than SunONE for 10 million entries, loading in just over 2 hours. However during the first attempts to load 10 million entries the process kept failing at just over 3,000,000 entries. This was because the Windows 2000 version of Aphelion had a 2GB limit on the size of the database that it supported, although the directory can support different databases to hold entries for different contexts. Consequently, 10 new databases were defined for each of the organzationalUnits, and the LDIF file was split accordingly for this into 11 files (1 for each organizationalUnit and another for the top level entries). Loading was then straightforward.

Due to the previously mentioned problems in loading four of the six directories with 10 million entries, only SunONE Directory Server and Syntegra Aphelion were tested with this number of entries. Another important point is that Syntegra Aphelion defines its own indexes when none are defined by the user – this therefore meant there was little point in timing the un-indexed loads, or tests A.1b and D.1b, which concern addition and deletion on entries from an un-indexed directory respectively.

Novell e-Directory achieved the worst performance for loading, almost 3 times slower than Active Directory, the next slowest directory at loading LDIF files. Based on this poor result, along with equally poor performance in the first test (S.1), Novell e-Directory was removed from further testing.

## 4.2. Search Results

**S.1 – Simulated Read (Base entry search on distinguished name (dn))**

Test S.1 examined the number of operations a second each directory could achieve by base entry searching on the exact distinguished names. Active Directory and Critical Path IDS showed consistent performance independent of the increasing data sizes. The results for the other directories, although showing small variance, are all arguably quite consistent within themselves.

Novell e-Directory performed the worst. It was so poor (over 4 times slower than the Critical Path IDS) that the directory was removed from further testing. Critical Path IDS performed the second worst in this test, whilst SunONE Directory Server performed the best, at over 600% better than Novell e-Directory and around 145% (800 operations a second) better than IDS. Interestingly, both Critical Path IDS and OpenLDAP showed marginally worse results at 100K entries than at 10K and 1 million entries, not only in this test, but also in the following three search tests. We cannot explain this.

| | **10K** | **100K** | **1 Million** | **10 Million** |
|---|---|---|---|---|
| Critical Path InJoin Directory Server 4.0 | 1562.5 | 1562.5 | 1562.5 | - |
| IBM SecureWay Directory 3.2.2 | 1666.7 | 1562.5 | 1666.7 | - |
| iPlanet/SunONE Directory Server 5.1 (evaluation) | 2173.9 | 2272.7 | 2381.0 | 2272.7 |
| Microsoft Active Directory | 2000.0 | 2000.0 | 2000.0 | - |
| Novell e-Directory | 342.5 | - | - | - |
| OpenLDAP 2.0.23 | 2272.7 | 1923.1 | 2173.9 | - |
| Syntegra Aphelion 2002 | 2173.9 | 2000.0 | 2083.3 | 2272.7 |

Table 5 – Simulated Read (Base entry search on distinguished name) (operations/second)

**S.2 – Full Subtree Search using exact match on common name (cn)**

Test S.2 examined the number of operations a second each directory could achieve by performing a full subtree search using an exact match on common name, an indexed attribute. All directories showed reasonably consistent performance, with Critical Path IDS showing absolutely consistent

performance independent of the number of entries. Critical Path IDS again performed at the lowest rate (50% worse than the best performer, Aphelion). IBM SecureWay performs consistently around 1/3 operations fewer a second than the top rating directories in this test. Active Directory experiences quite a large drop in performance between 100K and 1 million entries, becoming as slow as IDS. Syntegra Aphelion is marginally the best, although OpenLDAP 2.0.23 and SunONE Directory Server all perform well. It is interesting to note that 3 products performed marginally better at searching for an entry by its common name than being given the entire distinguished name of the entry (test S1).

| | **10K** | **100K** | **1 Million** | **10Million** |
|---|---|---|---|---|
| Critical Path InJoin Directory Server 4.0 | 1515.2 | 1515.2 | 1515.2 | - |
| IBM SecureWay Directory 3.2.2 | 1724.1 | 1612.9 | 1724.1 | - |
| iPlanet/SunONE Directory Server 5.1 (evaluation) | 2272.7 | 2173.9 | 2272.7 | 2272.7 |
| Microsoft Active Directory | 2272.7 | 2272.7 | 1562.5 | - |
| OpenLDAP 2.0.23 | 2381.0 | 1923.1 | 2381.0 | - |
| Syntegra Aphelion 2002 | 2381.0 | 2173.9 | 2272.7 | 2381.0 |

Table 6 – Full subtree exact match search on common name (operations/second)

**Test S.3 – Full subtree substring search on common name (cn) - an indexed attribute**

Test S.3 examined the number of operations a second each directory could achieve by performing leading substring searches on the common name, an indexed attribute. Compared with the searches in tests S.1, S.2 and S.4, the results of this test varied more greatly. Whilst most directories performed at approximately the same level, the performance of IBM SecureWay Directory dropped by over 1000 operations a second for all directory sizes, to fewer than 600 operations a second. In comparison to the other directories this is very poor, being less than 25% the performance of the best ones. Active Directory showed a similar decrease in performance between 100K and 1 million entries as in test S.2. Critical Path IDS again maintained consistent performance regardless of directory size, though still at 2/3 the level of the top performing directories.

OpenLDAP achieved the best performance in this test at 10K and 1 million entry levels with 2500 operations a second but experienced a blip at the 100K entry level. SunONE Directory Server also achieved 2500 operations a second at the 10 million entry level. Overall we can say SunONE Directory Server, OpenLDAP and Syntegra Aphelion produced consistently good results for this test.

| | 10K | 100K | 1 Million | 10Million |
|---|---|---|---|---|
| Critical Path InJoin Directory Server 4.0 | 1470.6 | 1470.6 | 1470.6 | - |
| IBM SecureWay Directory 3.2.2 | 595.2 | 581.4 | 588.2 | - |
| iPlanet/SunONE Directory Server 5.1 (evaluation) | 2381.0 | 2272.7 | 2381.0 | 2500.0 |
| Microsoft Active Directory | 2272.7 | 2272.7 | 1666.7 | - |
| OpenLDAP 2.0.23 | 2500.0 | 1923.1 | 2500.0 | - |
| Syntegra Aphelion 2002 | 2381.0 | 2173.9 | 2272.7 | 2381.0 |

Table 7 – Full subtree substring search on common name (operations/second)

**Test S.4 – Full subtree substring search on telephoneNumber - an un-indexed attribute**

Test S.4 examined the number of operations a second each directory could achieve by performing full subtree leading substring searches on telephoneNumber, an un-indexed attribute. Once again SunONE Directory Server, OpenLDAP and Syntegra Aphelion provide the best overall results. Active Directory also achieved high results in this test. Interestingly in most of the directory results for this test there is worse performance at 100K entries than at 10K and 1 million entries. Some directories also performed quicker than the equivalent test on the indexed common name attribute, whilst others performed worse. We are unable to explain these apparent discrepancies.

Critical Path IDS showed a minor drop in performance with 1 million entries, although only around 150 operations per second. IBM SecureWay Directory's performance had returned to similar levels as tests S.1 and S.2 after the large drop in performance in test S.3.

| | 10K | 100K | 1 Million | 10 Million |
|---|---|---|---|---|
| Critical Path InJoin Directory Server 4.0 | 1515.2 | 1515.2 | 1351.4 | - |
| IBM SecureWay Directory 3.2.2 | 1666.7 | 1562.5 | 1612.9 | - |
| iPlanet/SunONE Directory Server 5.1 (evaluation) | 2381.0 | 2173.9 | 2381.0 | 2272.7 |
| Microsoft Active Directory | 2173.9 | 2173.9 | 2000.0 | - |
| OpenLDAP 2.0.23 | 2381.0 | 1923.1 | 2381.0 | - |
| Syntegra Aphelion 2002 | 2173.9 | 1785.7 | 2173.9 | 2173.9 |

Table 8 – Full subtree substring search on telephoneNumber (operations/second)

## 4.3. Addition Results

**Test A.1a - Add organizationalPerson entry – indexed**

Test A1a examined the number of organizationalPerson entries a second that could be added to an indexed directory. All directories showed a decline in performance as the size of the database increased, however the degradation in performance varied greatly. Whilst being the best performer for a 10K entry directory, Critical Path IDS's performance rapidly tailed off. For 100k entry directory it was a factor of 12 worse, placing it next to bottom. . Active Directory had the next best performance, which was maintained between 10K and 100K, before dropping to a third at 1 million entries.

IBM SecureWay Directory and SunONE Directory Server both achieved good results throughout, with only a small degradation in performance with increasing size.. Syntegra Aphelion performed comparatively poorly, but its performance remained almost constant until 10 million entries. OpenLDAP performed very poorly throughout, and was the worst performer.. Overall, SunONE was the best performer, being capable of adding more entries a second to a 10 million directory than any other could to a 1 million directory.

| | **10K** | **100K** | **1 Million** | **10 Million** |
|---|---|---|---|---|
| Critical Path InJoin Directory Server 4.0 | 83.3 | 6.8 | 3.8 | - |
| IBM SecureWay Directory 3.2.2 | 20.0 | 16.7 | 11.5 | - |
| iPlanet/SunONE Directory Server 5.1 (evaluation) | 28.6 | 16.1 | 15.9 | 11.6 |
| Microsoft Active Directory | 31.3 | 32.3 | 10.4 | - |
| OpenLDAP 2.0.23 | 6.7 | 5.3 | 2.1 | - |
| Syntegra Aphelion 2002 | 8.4 | 8.5 | 7.0 | 2.8 |

Table 9 – Add organizationalPerson Entry to Indexed Directory (operations/second)

**Test A.1b - Add organizationalPerson entry - un-indexed**

Test A.1b examined the number of organizationalPerson entries that could be added to an un-indexed directory a second. It was not possible to test Syntegra Aphelion in this test, as Aphelion creates a set of default indexes when no user indexes are defined so its performance would be identical to its indexed performance. All systems performed better than when they were indexed, with the improvements ranging from 5% for IBM SecureWay to 240% for Critical Path IDS.

Critical Path IDS excelled, performing 5 times faster than any other directory at 10K and 100K. Even though its performance degraded by a factor of over 6 at 1 million entries, it was still the best (although it would be interesting to see how far IDS's performance degrades at the 10 million entry level). Whilst most performances degraded as the size of the directory increased, OpenLDAP stayed approximately the same at each entry level. Note however that OpenLDAP's performance was generally the worst.

IBM SecureWay showed only a gradual drop in performance, whilst Active Directory's was more dramatic. SunONE's performance actually increased from 10K entries to 100K, but then dropped quite rapidly.

| | **10K** | **100K** | **1 Million** | **10 Million** |
|---|---|---|---|---|
| Critical Path InJoin Directory Server 4.0 | 200.0 | 200.0 | 31.3 | - |
| IBM SecureWay Directory 3.2.2 | 21.3 | 19.6 | 15.6 | - |
| iPlanet/SunONE Directory Server 5.1 (evaluation) | 40.0 | 43.5 | 30.3 | 18.5 |
| Microsoft Active Directory | 34.5 | 17.5 | 10.9 | - |
| OpenLDAP 2.0.23 | 12.2 | 13.7 | 13.7 | - |

Table 10 – Add organizationalPerson Entry to Un-Indexed Directory (operations/second)

## 4.4. Deletion Results

**Test D.1a - Delete an entry – indexed**

Test D.1a examined the number of organizationalPerson entries that could be deleted from an indexed directory a second. OpenLDAP showed the poorest performance of all, with almost identical results to those in test A.1a for adding entries. Critical Path IDSachieved by far the best results at 10K and 100K, but because of its rapid drop in performance with directory size, by 1 Million it was the second to worst performer.

Syntegra Aphelion performed the best in this test overall, showing little performance variation with size. SunONE Directory, IBM SecureWay and Active Directory each showed approximately the same degradation of performance with size, dropping 50% between 10K and 1 million entries. Overall, SunONE was twice as fast as IBM SecureWay which was approximately 30% faster than Active Directory.

| | 10K | 100K | 1 Million | 10 Million |
|---|---|---|---|---|
| Critical Path InJoin Directory Server 4.0 | 100.0 | 71.4 | 3.1 | - |
| IBM SecureWay Directory 3.2.2 | 23.8 | 14.5 | 10.5 | - |
| iPlanet/SunONE Directory Server 5.1 (evaluation) | 41.7 | 21.7 | 20.8 | 10.1 |
| Microsoft Active Directory | 15.9 | 11.8 | 8.1 | - |
| OpenLDAP 2.0.23 | 6.0 | 5.1 | 2.2 | - |
| Syntegra Aphelion 2002 | 21.7 | 22.7 | 20.4 | 18.5 |

Table 11 – Delete organizationalPerson Entry from Indexed Directory (operations/second)

**Test D.1b - Delete an entry - un-indexed**

Test D1.b examined the number of organizationalPerson entries that could be deleted from an un-indexed directory a second. The results in this test show similarities with the results of test A.1b, (adding entries to an un-indexed directory). Critical Path IDS was the best, performing over 7 times better than any other directory for 10K and 100K, but again suffered another large performance drop at 1 million entries, by a factor of 10. Active Directory was the worst performer overall, and its performance also degraded quickly with size. IBM SecureWay again attained quite low results but only had gradual performance degradation, whilst OpenLDAP started with the worst performance, but this did not degrade at all (in fact performance increased slightly between 10K and 100K entries). SunONE performed second best overall, but from 1 Million entries upwards it performance degraded quickly. No results are shown for Syntegra Aphelion because it always contains indexed entries.

| | 10K | 100K | 1 Million | 10 Million |
|---|---|---|---|---|
| Critical Path InJoin Directory Server 4.0 | 333.3 | 333.3 | 30.3 | - |
| IBM SecureWay Directory 3.2.2 | 22.7 | 16.9 | 13.0 | - |
| iPlanet/SunONE Directory Server 5.1 (evaluation) | 43.5 | 45.5 | 27.8 | 16.9 |
| Microsoft Active Directory | 15.6 | 12.7 | 6.0 | - |
| OpenLDAP 2.0.23 | 11.8 | 13.5 | 13.5 | - |

Table 12 – Delete organizationalPerson Entry from Un-Indexed Directory (operations/second)

## 4.5. Modification Results

**Test M.1 - Modify an indexed attribute (cn)**

Test M.1 examined the number of operations a second each directory could achieve by modifying common names (cn), an indexed attribute. OpenLDAP attained the lowest results of all in this test, with just 1.3 operations a second on 1 million entries. Critical Path IDS was by far the best performer, being over 100 times faster than OpenLDAP at one point. Interestingly IDS increased in performance by a factor of 2 between 10K and 100K entries, before rapidly decreasing in performance by a factor of over 5 between 100K and 1 million entries. Syntegra Aphelion was the second worst performer in this test, though it was still 5 times faster than OpenLDAP. The sharp drop in performance for Aphelion between 1 Million and 10 Million entries may be due to the modification in the setup of the directory at the 10 million data entry level, highlighted previously in the bulkloading section of this paper.

The second best performer was, Microsoft's Active Directory, but it suffered from a 2/3rds performance drop at 1 Million entries. Although Active Directory performed well, the test performed had an important difference to the others; in the original tests there was a 50/50 split in the operations between adding an additional value to the CN attribute and replacing the value. Active Directory would not allow multiple values for the CN attribute, so all scripts were changed to simply replace the value.

IBM SecureWay and SunONE were approximately 3 times faster than Aphelion, with SunONE being slightly the better of the two as its performance did not tail off as rapidly. By 10 Million entries, SunONE was nearly 5 times faster than Aphelion.

| | **10K** | **100K** | **1 Million** | **10 Million** |
|---|---|---|---|---|
| Critical Path InJoin Directory Server 4.0 | 188.7 | 333.3 | 59.9 | - |
| IBM SecureWay Directory 3.2.2 | 40.3 | 34.0 | 23.0 | - |
| iPlanet/SunONE Directory Server 5.1 (evaluation) | 40.0 | 37.3 | 30.2 | 13.0 |
| Microsoft Active Directory | 96.2 | 98 | 32.8 | - |
| OpenLDAP 2.0.23 | 3.3 | 2.4 | 1.3 | - |
| Syntegra Aphelion 2002 | 12.5 | 12.2 | 9.4 | 2.7 |

Table 13 – Modify indexed attribute cn (operations/second)

**Test M.2 - Modify an un-indexed attribute (telephoneNumber)**

Test M.2 examined the number of operations a second each directory could achieve by modifying the telephoneNumber attribute, an un-indexed attribute. OpenLDAP showed the worst performance

throughout, with results between 5 and 14 times lower than the next worst, Syntegra's Aphelion. Interestingly, Aphelion's results were almost static up to 10 million entries, when they increased. This contrasts greatly with the test M.1 results that showed a drop at 10 million entries. The increase in performance is almost certainly due to the altered database configuration that allows Aphelion to update 10 databases separately without needing to make any modifications to the indexing.

Performance of Critical Path IDS was very good, being 3 times faster than its nearest rival, Active Directory. However another sharp drop in IDS performance occurred between 100K and 1 million entries (a factor of 6), making them approximately equal at 1 Million entries. Active Directory required modification to the scripts (as in the previous test) as it would not allow attributes to contain multiple values. IBM SecureWay also achieved good results, although it was not quite as fast as Active Directory at 10K entries, its performance decreased less rapidly, making them equal at 1 Million entries. SunONE Directory Server's performance was mediocre, though its performance decreased the least (excluding Aphelion) so we suspect it would have been second highest by 10 million entries.

| | **10K** | **100K** | **1 Million** | **10 Million** |
|---|---|---|---|---|
| Critical Path InJoin Directory Server 4.0 | 312.5 | 277.8 | 45.2 | - |
| IBM SecureWay Directory 3.2.2 | 78.1 | 70.4 | 48.1 | - |
| iPlanet/SunONE Directory Server 5.1 (evaluation) | 50.8 | 51.8 | 36.5 | 21.0 |
| Microsoft Active Directory | 99 | 95.2 | 48.1 | - |
| OpenLDAP 2.0.23 | 5.8 | 4.2 | 1.6 | - |
| Syntegra Aphelion 2002 | 26.5 | 27.5 | 23.1 | 37.5 |

Table 14 – Modify un-indexed attribute telephoneNumber (operations/second)

## 4.6. Mixed Operations Results

**Test X.1 - Mixed operations – Search, Add, Delete, Modify, ModifyRDN**

Test X.1 examined the performance of each directory whilst performing a mix of the operations used during the previous tests – search (50%), add (12%), delete (10%), modify (16%) and modifyRDN (12%). Based on the results of the previous tests, the results in Table 15 below show many of the same trends. OpenLDAP again performed the worst, due to it's poor performance on add, delete and modify operations. Critical Path IDS showed the best performance with 10K and 100K

entries, over 3 times faster that of its nearest rival, but again it had a staggering drop in performance at 1 million entries, by a factor of 18. IBM SecureWay and SunONE came joint second at 10K entries, but SecureWay's performance declined more rapidly than that of SunONE. For directories above 1 Million entries, SunONE Directory Server was the best performer in the mixed operations testing. Syntegra Aphelion's performance was poor for 10K entries (under half that of SunONE), but it managed to maintain the same performance up to 1 Million entries, by which time it had become respectable. However, it dropped substantially for 10 million entries. The suspected reason for this loss is the need to update 10 databases whilst maintaining a single index. Microsoft Active Directory performance was mediocre throughout.

| | **10K** | **100K** | **1 Million** | **10 Million** |
|---|---|---|---|---|
| Critical Path InJoin Directory Server 4.0 | 151.0 | 151.0 | 8.4 | - |
| IBM SecureWay Directory 3.2.2 | 50.3 | 27.5 | 21.6 | - |
| iPlanet/SunONE Directory Server 5.1 (evaluation) | 50.3 | 37.8 | 30.2 | 21.6 |
| Microsoft Active Directory | 33.6 | 27.5 | 18.9 | - |
| OpenLDAP 2.0.23 | 9.4 | 7.4 | 2.7 | - |
| Syntegra Aphelion 2002 | 23.2 | 21.6 | 20.1 | 8.6 |

Table 15 – Mixed operations (operations/second)

## 5. Results Summaries

Critical Path InJoin Directory Server 4.0 – Overall performed extremely well up to the 100K entry level. Indeed it can be argued that overall in our test results Critical Path IDS 4.0 was the top performing directory up to that level. After 100K entries its performance drops off rapidly for all modification type operations, so that by 1million entries it is only average.

IBM SecureWay Directory 3.2.2 – Consistently average performance over all of the tests means that one must have other reasons than performance for selecting this product.

iPlanet/SunONE Directory Server 5.1 (evaluation) – Over 1 Million entries this was the top performing directory in our set of tests. Below the 100K level its performance is amongst the top performing directories for Search type operations, and so dependant on application requirements users may choose to deploy iPlanet/SunONE Directory Server 5.1 over Critical Path InJoin Directory Server 4.0.

Microsoft Active Directory – Consistently average performance over all of the tests, similar in many respects to IBM's SecureWay. There were also some noticeable drops in performance between 100K and 1 Million entry databases.

OpenLDAP 2.0.23 – Performance was extremely strong in the search tests we carried out but in all the modification type tests it was by far the worst performing directory.

Syntegra Aphelion 2002 –For the search results it is amongst the top performers but is well below average for the modification type tests that we carried out. However, two of its redeeming features are its extremely fast bulk loading times, and its relatively low performance reduction as directory size increases. This makes it a directory worthy of consideration for very large data sets of 10 Million entries or more.

In conclusion, for use in ETP, the Critical Path IDS appears to be the best choice for prescription repositories up to 100K entries, whereas SunONE appears to be the best choice for prescription repositories of over a 1 million entries.

## 6. Directory Analysis from an Administrative Perspective

All the LDAP directories tested were easy to install, with the exception of OpenLDAP. We wished to configure Berkeley DB as the backend to OpenLDAP, but the documentation for this was quite thin, and sometimes out of date. There is no administrative interface for OpenLDAP, and all tasks must be undertaken by either editing text files manually or issuing commands in the shell. It must be said however that once the initial barriers are overcome, OpenLDAP was relatively easy to administer, and users of Linux and other similar operating systems will be used to (and often prefer) this more controlled approach to administration.

SunONE/iPlanet Directory Server used a Java based interface, which was highly intuitive, well laid out and easy to administer. The same is also true of Syntegra Aphelion, IBM SecureWay and Critical Path IDS, however these directories all used a HTML/CGI style interface accessed via a web browser. It was found that the web based interfaces sometimes produced errors. In all 3 of the previously mentioned directories, a function was provided to start and stop the directory service, and when there was an error (for example a failed load of the directory because of a corrupted the data store), they would sometimes state that the directory service was started, when it was impossible to make a connection or changes to it. Critical Path also used another application for administration in addition to

the web browser, which meant it was necessary to use 2 interfaces, and also a text file required editing. However compared with Active Directory this was still simple. Active Directory was the hardest to administer, which was surprising, as Microsoft normally provide reasonably good user interfaces for their software. In order to administer Active Directory, such as making changes to the schema, it was necessary to install separate 'snap-ins', giving the impression Microsoft do not trust most of their directory administrators to make simple changes.

OpenLDAP did not provide a backup and restore facility, although the directory could be backed up and restored by either exporting the directory to a LDIF file, or simply stopping the service and copying (or replacing) the files. Active Directory was easy to backup using ntbackup, however to restore the directory, it was necessary to shutdown and reboot the server into 'Directory Services Restore Mode'. Once this was completed it was necessary to reboot again, making restoration arduous. SunONE/iPlanet Directory Server, Critical Path and IBM SecureWay all provided easy backup and restore, in proprietary format (also true of Active Directory). Syntegra Aphelion's backup and restore facility worked using LDIF files, which meant the backup could be easily edited, and the facility was extremely fast.

In summary, all interfaces were good, with the exception of Active Directory, which provides a poor interface, and a highly awkward backup and restore.

## Acknowledgements

This research was funded by the UK EPSRC under grant number GR/M83483. The authors would also like to thank the AccordIT for support with the SunONE LDAP directory, Syntegra for the free evaluation software and support and finally CriticalPath for their free evaluation software and support.

## References

[1] Mundy, D.P. and Chadwick, D.W. "A system for secure electronic prescription handling", Second International Conference On The Management Of Healthcare And Medical Technology On: The Hospital of the Future Bringing Together Technology, Health Care and Management. Abstract and Main Paper on accompanying CD-Rom, Stuart Graduate School of Business, Center for the

Management of Medical Technology, Illinois Institute of Technology, Chicago, Illinois, USA, July 28-30, 2002 @ http://sec.isi.salford.ac.uk

[2] DirectoryMark @ http://www.mindcraft.com

[3] Lewis, Jamie. In Evaluating Directories, Be Aware They're Not Right For All Apps. Internet Week, 10th January 2000 @ http://www.internetweek.com/columns00/lewis011000.htm

[4] Oracle. Directories and Relational Databases Compared @ http://download-west.oracle.com/otndoc/oracle9i/901_doc/network.901/a90153/dir_conc.htm

[5] ISO/ITU-T Rec. X.509(2000) The Directory: Authentication Framework

[6] ISO/ITU-T Rec. X.501(1993) The Directory: Models

[7] Wahl, M., Howes, T., Kille, S. "Lightweight Directory Access Protocol (v3)", RFC 2251, December 1997 @ http://www.ietf.org

## Appendix A – Software Configuration

Note: For all directories, the data store was created on a SCSI drive separate from the SCSI drive hosting the OS and the software. All indexing was also removed prior to testing except for tests which required indexing on cn.

**Critical Path InJoin Directory Server 4.0**

Loading was done using ldapmodify as follows:

Ldapmodify -D cn=manager -w (password) -f data10k.ldif -a

The hexadecimal default value (200000) in ds.D.00 was changed for different database sizes i.e. 7A1200 for 2GB, F42400 for 4GB etc. It was necessary to modify some of the DirectoryMark scripts for Critical Path because of the schema. For tests 5a and 5b (which used the same scripts), the format of the seeAlso attribute was changed from seeAlso: Babs Jenson to seeAlso: cn=Babs Jenson. For test 8, the modification script was changed on the telephone number attribute, as 'this has been ldap modified' was unsupported by Critical Path's telephone attribute rules (i.e. +1 234 567-8901).

**iPlanet/SunONE Directory Server 5.1 evaluation**

Loading was done using ldif2db:

ldif2db –n userRoot –i f:\backup\data10k.ldif

Transaction logging and durable transactions (stored on disk) were turned off in dse.ldif. Access and error logs were used, but set to renew every day or when over 100 or 50 MBs respectively. The nsslapd-allidsthreshold attribute in dse.ldif was set to 10,000 for 10k, 100,000 for 100k, 1,000,000 for 1 million and 10,000,000 for 10 million to help improve search performance as recommended in the manuals. Adjusting this value did not appear to provide any gain or loss in performance.

**IBM SecureWay Directory 3.2.2**

Loading was done using bulkload from the DB2 command (DB2CMD) command prompt:

Bulkload –i data10k.ldif

The DB2 optimise function was used after each bulkload, as recommended in the manual, located in the web administration interface.

**OpenLDAP 20.2.3**

ldapadd was used to add the initial 2 entries to the directory. For all other loading, the bulk loader slapadd was used:

slapadd –l /data10k.ldif –f /usr/local/etc/openldap/slapd.conf -d 0

The cache size was adjusted (in slapd.conf) for each different data point, i.e. 10,000 for 10K, 100,000 for 100K and 1,000,000 for 1 million. In order to run tests 7, 8 and 9 it was necessary to modify the main schema of the directory (in core.schema). The rules added for telephoneNumber were:

EQUALITY – telephoneNumberMatch
SUBSRTING MATCH - telephoneNumberSubstringsMatch

Also, the equality rules for postalAddress were modified from caseIgnoreListMatch to caseIgnoreMatch:

EQUALITY - caseIgnoreMatch

Sleepycat's Berkeley DB 4 was chosen as the backend, as this is the recommended backend[1]. A symbolic link was created from /usr/include/db.h to db4/db.h (the Berkeley database location).

[1] Since the tests on OpenLDAP were run, OpenLDAP has a new backend hierarchical database which is reputed to be much faster for modification operations than the Berkeley database that we used in these tests.

**Microsoft Active Directory**

The directory was loaded using LDIFDE from the DOS command prompt:

Bulkload –I -f data10k.ldif

Index adjustments were made using the ADSIedit snap-in.

**Syntegra Aphelion 2002**

The directory was loaded using import from the DOS command prompt:

import –f \usr\var\lde\test\test_lde.conf –c c=gb –l \usr\var\lde\test\data10k.ldif

The log level was set to 'debug packet handling' (2), with schema check off and access check off. Substring indexing was defined on 'word'. It was not possible to test Aphelion totally un-indexed, as when no indexes are user-defined, Aphelion uses a set of defaults.

**Novell e-Directory 8.6**

The directory was loading using the ICE facility from the command prompt:

Ice –S LDIF –f data10k.ldif –a –D LDAP –d CN=admin,O=salford, c=gb –w <password>

Novell SLP Directory Agent, Console One and Novell Directory Services were installed (ZenWorks and Workstation manager were not). Entry Cache and Block Cache were both set to use 50% each of the available memory. Dynamic Caching was also used, and all cache refresh intervals were set to 15 seconds. It was also necessary to extend the search timeout from 60 seconds to 120 seconds, as around 50% of operations in Test S.1 timed out after 60 seconds.