

Interview with OpenLDAP's Howard Chu

Who is Howard Chu?

Howard Chu is the Chief Architect of the OpenLDAP project and its main corporate sponsor Symas Corporation. OpenLDAP is a free, open source implementation of the Lightweight Directory Access Protocol (LDAP) which provides an enterprise with shared address books, single sign-on functionality, automount of home directories and file sharing for Linux, Unix, Mac and Windows clients.

Q: Can you tell us a little bit about what you do, the OpenLDAP project, its relationship with Symas?

Well, as Chief Architect for the OpenLDAP Project I occasionally make decisions about what technical features should or should not be integrated into the code. For the most part though, developers in the OpenLDAP community simply work on whatever they choose, whatever scratches their itch. I wrote my first contribution in 1998 and was invited to join the core team shortly after that. Under Kurt Zeilenga's leadership, most of the early development in OpenLDAP was focused on cleaning up portability issues and implementing LDAPv3. The more radical evolution of the code since its UMich origins has been at my instigation and most of that is my code. I've been working full time on the project since 1999 as a Founder of Symas which has chosen to invest in this technology through funding my participation.

Q: How do you compare OpenLDAP with proprietary directory services technologies like Active Directory or SunOne?

Active Directory is fundamentally flawed in so many areas it barely deserves mention. It is grossly non-compliant with the LDAP specifications, breaking interoperability. And its database design is so broken it can barely get out of its own way. Our recent assessment of AD and Active Directory Application Mode (ADAM) as LDAP servers and the benchmarks that show it to be 3 to 5 times slower convinced us that enterprise strategies based on that as a production enterprise directory are headed for trouble. See Symas' enterprise assessment whitepaper of Microsoft Active Directory's Application Mode versus OpenLDAP and the report's update.

SunOne was, for some years, probably the leading directory technology in the industry. However, the original development teams walked away from the code base years ago and it's showing its age, with numerous well documented stability and maintainability issues. Today OpenLDAP has a significant lead in performance, scalability, and reliability. Unfortunately we can't publish benchmark results against SunOne due to a restriction in their end user license. It's worth nothing, however, that SunOne is being replaced by OpenDS, an open source directory project written in Java. The reign of Sun's proprietary directory service is over; SunOne has reached the end of the line.

OpenLDAP is unmatched by any other directory service, proprietary or open source. Of all the others available, the proprietary ones are just hiding their dirty laundry and all of them are just a waste of time and money.

Q: We are seeing the emergence of an Open Source software stack upon which it is perfectly possible to run an enterprise. Where do you see OpenLDAP's position in this 'Open Source enterprise stack'?

I think the best answer to that is to point to HP's Open Source Investment Portfolio and Open Source Middleware Stack. The selected OpenLDAP as the directory technology (and Symas as the support partner).

Several smaller ISVs have also adopted OpenLDAP as their directory technology of choice (Ventyx and Zimbra the most notable) and we expect more announcements of that type.

Q: It seems every OSS project has its own different LDAP schema, for example Samba's schema is very different from those used by GOSa or Kolab. What's your solution to the problem of schema proliferation and associated problems of incompatibility and complexity?

Schema proliferation in LDAP directories is really quite manageable. The main point is to get these various teams to publish formal specifications of their schema for public review, to aid in their adoption. As an example, we're working with the Samba team and IETF Kerberos working group to develop a standard LDAP schema for Kerberos KDC information. We're creating a rational superset of the schemas currently used by Heimdal and MIT, which can be consistently implemented by both and then relied upon by Samba and other applications that need to work closely with Kerberos and LDAP. It's simple really: the people and teams have to interoperate, in order to ensure that the software will interoperate.

It's a bit surprising that this is even considered a problem in the LDAP space, because it's generally so easy to address. You very rarely run into truly incompatible schema definitions. Usually you just find that the published standard schema are incomplete or inadequate for a specific application you had in mind. That's to be expected, since most of the published schema are only intended as starting points, and they're meant to be extended and mixed and matched with other schema. In contrast, schema management in relational databases is a truly intractable problem. There are no shared definitions in the SQL world like there are in X.500/LDAP. In fact there isn't even a single SQL in the first place, there are a variety of subtly incompatible dialects without any authoritative reference. Even such fundamental concepts as elementary data types (integer, Int64, etc.) lack a standard definition across various implementations.

Of course we do run into situations where in depth education is needed. We do a lot of formal and informal consulting for enterprises moving to OpenLDAP. Some compatibility concerns occasionally pop up but they're quickly addressed as technical staff gets up to speed with early "LDAP University" classes that Symas teaches.

Q: How can the OSS community work better towards encouraging the use of OpenLDAP by enterprises?

It's first about selecting LDAP as the technology for directory data. We see enterprises and OSS projects implementing directory data stores with other technologies and they rarely scale, perform, or administer adequately for enterprise deployments. LDAP offers a superior and readily available database for directory data. Second, take the time to qualify your LDAP use against OpenLDAP. Having invested in LDAP capable code, you should test it against the most standards-compliant LDAP technology and offer your users the chance to easily deploy on OpenLDAP. Third, they should benchmark the OSS directory technologies using the proposed schema, representative data samples and workloads, and at numbers of entries similar to what enterprises might need. These benchmarks are simple to do and Symas can help a project get started with the OSS benchmarking technology we use constantly. Those three steps will quickly convince OSS developers to endorse OpenLDAP as their recommended OSS directory technology.