

Apache POI - Project History

by Andrew C. Oliver

1. Apache POI - Brief Project History

The POI project was dreamed up back around April 2001, when Andrew Oliver landed a short term contract to do Java-based reporting to Excel. He'd done this project a few times before and knew right where to look for the tools he needed. Ironically, the API he used to use had skyrocketed from around \$300 (\$US) to around \$10K (\$US). He figured it would take two people around six months to write an Excel port so he recommended the client fork out the \$10K.

Around June 2001, Andrew started thinking how great it would be to have an open source Java tool to do this and, while he had some spare time, he started on the project and learned about OLE 2 Compound Document Format. After hitting some real stumbers he realized he'd need help. He posted a message to his local Java User's Group (JUG) and asked if anyone else would be interested. He lucked out and the most talented Java programmer he'd ever met, Marc Johnson, joined the project. He ran rings around Andrew at porting OLE 2 CDF and rewrote his skeletal code into a more sophisticated library. It took Marc a few iterations to get something they were happy with.

While Marc worked on that, Andrew ported XLS to Java, based on Marc's library. Several users wrote in asking to read XLS (not just write as had originally been planned) and one user had special requests for a different use for POIFS. Before long, the project scope had tripled. POI 1.0 was released a month later than planned, but with far more features. Marc quickly wrote the serializer framework and HSSF Serializer in record time and Andrew banged out more documentation and worked on making people aware of the project

Shortly before the release, POI was fortunate to come into contact with Nicola -Ken-Barrozzi who gave them samples for the HSSF Serializer and help uncover its unfortunate bugs (which were promptly fixed). More recently, Ken ported most of the POI project documentation to XML from Andrew's crappy HTML docs he wrote with Star Office.

Around the same time as the release, Glen Stampoulziz joined the project. Glen was ticked off at Andrew's flippant attitude towards adding graphing to HSSF. Glen got so ticked off he decided to grab a hammer and do it himself. Glen has already become an integral part of the

POI development community; his contributions to HSSF have already started making waves.

Somewhere in there we decided to finally submit the project to [The Apache Cocoon Project](#), only to discover the project had outgrown fitting nicely into just Cocoon long ago. Furthermore, Andrew started eyeing other projects he'd like to see POI functionality added to. So it was decided to donate the Serializers and Generators to Cocoon, other POI integration components to other projects, and the POI APIs would become part of Jakarta. It was a bumpy road but it looks like everything turned out since you're reading this!

In Early 2007, we graduated from [Jakarta](#), and became our own Top Level Project (TLP) within Apache.

Copyright (c) @year@ The Apache Software Foundation All rights reserved.