



The following paper was originally published in the
Proceedings of the Fifth Annual Tcl/Tk Workshop
Boston, Massachusetts, July 1997

xmb - dancing in the tar pits
Using Tcl/Tk to build a CASE environment

Henry R. Tumblin
CertCo L.L.C., New York, NY
Charles E. McElwain
Open Market, Inc., Cambridge MA

For more information about USENIX Association contact:

1. Phone: 510 528-8649
2. FAX: 510 548-5738
3. Email: office@usenix.org
4. WWW URL: <http://www.usenix.org>

xmb - dancing in the tar pits

Using Tcl/Tk to build a CASE environment

Henry R. Tumblin Charles E. McElwain

CertCo L.L.C. *Open Market, Inc.*

55 Broad Street *245 First Street*

New York, NY 10004 *Cambridge, MA 02142*

tumblin@certco.com mcelwain@openmarket.com

Managing software across multiple build platforms and environments can be thought of as dancing around a tar pit; if you slip and fall then you may not be able to get out for a long time! There's far too much knowledge about the local build environment required to enable software developers and release engineers to build, test, and release products. We intend to show that by using **Tcl/Tk** as a CASE integration engine, the amount of time and training required to bring a new developer, test or release engineer online is greatly reduced. Product build environments become standardized and the cost of adding new platforms is reduced.

The challenge we faced was one of release engineers having to build on multiple platforms, developers needing to build and test on a subset of these platforms, all working from a variety of individual environments and experience levels.

The problems that had to be solved were:

- distributing builds across the multiple platforms
- centralizing knowledge of build configurations in a tool (for repeatability)
- giving feedback on system resource usage and sharing
- providing better feedback and control (via a UI) and minimize training
- evolving a CASE tool towards greater maintainability, extensibility and reusability.
- adding multiple projects, and co-ordinating and analyzing inter-project dependencies

We looked at solving these problems through a three layer tool containing: a point-and-click GUI, a CASE kernel, and the remote tasks. The tool was based on a "build flow" which ran through:

Select Project -> Select Platform(s) -> Select Options -> Distribute -> Analyze

The problem was then how to turn this into a programmatic flow using **Tcl/Tk**.

We built a small kernel representing these tasks, whose behavior is controlled by and changeable via config files which are themselves **Tcl**. Scripts to execute on the remote hosts were generated, then launched on the remote host. From within these scripts, host-specific config files were sourced to enforce the build environment and policies, as well as supplemental information, including compilers, compiler flags, paths, etc.

These scripts captured statistical information on each distributed build in the form of a **Tcl**




Figure 2: Selecting different projects, **omx**, **Tcl**, and **Toolshed** and multiple tags, *main_line*, *OMX_2_0_a_RC2*, *OMX_1_0_2_RC1*, within one of them.

Figure 3: Selecting build options for the current project **omx** and tag *main_line*.

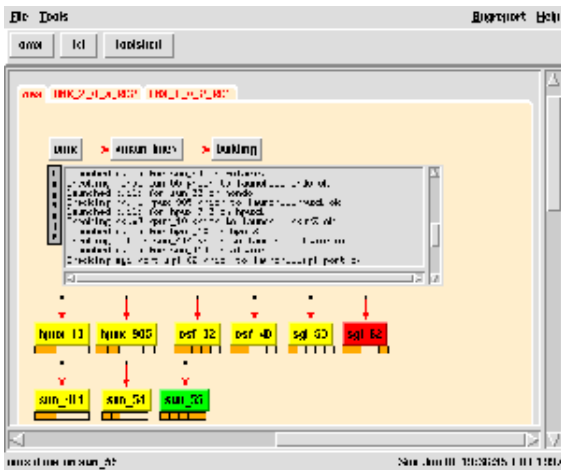
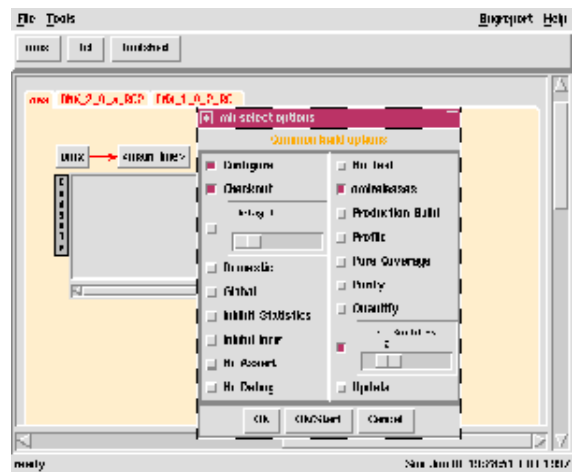


Figure 4: Watching a build in progress across multiple platforms. One has succeeded, *sun_55*, one has failed, *sgi_62*, and the others are still in progress.