Jukebox: Logging

Table of contents

1. For The Impatient

Jukebox uses log4j.

2. Background

The logger component, created in early 1996, was utilizing double buffering. Double buffering was required in order to avoid slowing down the application (JDK 1.0.2 was **extremely** slow, and cost of I/O was tremendous).

The implementation was polished in 1998, when it has become a part of Apache JServ servlet engine - as you may guess, general purpose servlet engine is a very demanding application, and a lot of bugs were eliminated at that stage. However, the very foundation of the logger's double buffer, the queue, was implemented using `java.util.Vector`, and surprisingly, faster implementation wasn't required until much later (2002, to be exact) - when the same logging mechanism was used to support the high-load application at American Express. Only then, it was replaced by a specialized high-performance class.

(Historical note circa 2003, no longer correct) *Currently, the logger is stable (I'd guess, after 6 years in production), but honestly, I grew a little bit cold towards it - well, it was first (known to me, that is) logger implementation utilizing double buffering, but now everybody and their uncle is creating loggers, so this is not fun anymore. It's quite possible that the logger component will eventually be replaced with log4j - the project based on the same initial assumptions, but possibly of a better quality - just because the guys are dedicated to the logger, while for me the logger (at least today) is just a tool.*

3. Reason for the Switch

It didn't make sense to maintain a separate logger any longer. Drawbacks of incompatibility, uniquencess and necessity to evolve and maintain the code overweighted real or imaginable benefits of the architecture that existed at the time the decision to switch was made.