

## ***Public Review: Walking toward moving goalposts: agile management for evolving systems***

**Richard A. Golding and Theodore M. Wong**

*Public Reviewer: Alistair Veitch*  
*Hewlett Packard Laboratories*

This paper is important to read, for several reasons. Lets first examine the basic premise of the paper, which is that we will only ever get to true self-managing systems if the management is built into the entities comprising the system, rather than "bolted on" through the use of some external application(s). This in itself is a fine discussion topic, although I sincerely hope it's not true, since there are, and will continue to be, huge numbers of "dumb" devices that will not then be able to be part of a self-managed system unless they are proxied somehow. In a world where the cost of management is increasing (because it's dominated by human costs), and the hardware is getting cheaper, not having such self-managing systems is increasingly painful, and the need for research on them is increasingly important - and is the first major reason you should read this paper, since it describes such research.

The main body of the paper describes K2, a new distributed storage system being built for the express purpose of experimenting with autonomic management techniques. Central to the system is the idea of building pools of allocatable resources from the collection of underlying devices, and matching those to storage pools (which are the application-visible unit of resource allocation), and using the devices themselves to manage each pool. Coming to an understanding of how such systems might be used to solve general management problems is the second major reason to read this paper.

While the simple model outlined above would appear to have the potential for working well, the devil is in the details. In particular, I'm sure there are several gotchas waiting for the authors as they continue to refine their system. For instance, performance in K2 is measured only in IOPS, with little or no mention given to other important metrics such as latency or bandwidth, nor thought on how the multiple, and very non-linear relationships between them, will be used or incorporated into the hill-climbing optimization-engine described in the paper. Similarly, reliability is given as simply MTDL, without mention of more sophisticated, and (more importantly) widely used metrics,

such as recovery-time objectives and recovery-point objectives, nor the many mechanisms by which reliability is achieved other than encoding schemes. These aren't reasons to not read the paper however, but more warnings to the unwary (not the authors, they know this already), that storage systems are incredibly complicated beasts, and don't give in readily to simple management solutions.