

## *Public Review: Advanced tools for operators at Amazon.com*

**Peter Bodík, Armando Fox, Michael I. Jordan, David Patterson, Ajit Banerjee, Ramesh Jagannathan, Tina Su, Shivaraj Tenginakai, Ben Turner, and Jon Ingalls**

*Public Reviewer: Thu D. Nguyen*  
*Rutgers University*

This paper addresses an important and timely topic: building infrastructure and tools to help human operators manage complex systems such as Internet services. This is a critical area that has typically been ignored (except by a very few research groups) in the recent rush toward autonomic computing. This is not surprising given that understanding human-system interactions is time-consuming and unfamiliar to systems researchers. This difficulty was amply reinforced in our recent effort to collect and analyze data on operator mistakes using human-factors experiments with volunteer operators. However, this is too important of an issue for system designers to ignore.

Autonomic computing, even if eventually successful, is probably quite a ways off in the future. Human operators need help now to deal with systems of unprecedented scale, complexity, and volatility. Further, regardless of the success of autonomic computing, humans will always be part of the installation and management of complex computer systems at some level. For example, humans will likely always be responsible for determining a system's overall policies and, given our demonstrated inability to produce bug-free code, for addressing unexpected behaviors or failures. Thus, it is critical that we build systems that are understandable and manageable by humans. Nothing is more aggravating than an automatic system that is either unpredictable or does the wrong thing in a manner that is opaque to the user or operator. What better way to ensure that our systems are understandable to humans by first building infrastructure and tools to help operators understand and better manage current complex systems.

Eventually, when we successfully build autonomic systems, the automatic management process may, algorithmically, be quite different from the human management process. However, to me, it seems critical to first understand how human operators currently manage systems. As system designers, rather than system managers, we probably don't understand the management process very well. Again, what better way to improve our understanding of the management process than to build infrastructure and tools to help op-

erators improve their efficiency. In exploring what works and what does not work to increase human management efficiency, we will undoubtedly learn a great deal about the management process.

Aside from the general direction, perhaps the paper's most important contribution is a demonstration of the necessary process to work in this area as well as some of the practical and research challenges. Specifically, to understand problems faced by human operators, systems researchers will have to spend time with operators. Also, a critical challenge in building tools is that they have to be comprehensive. Given the complex interdependency of today's systems, this is quite a challenge. We will have to address this challenge, however, because operators will likely not use a tool that only deals with a limited part of the system, making it quite difficult to evaluate a limited prototype. Finally, the paper shows two ways to help human operators, increasing understanding of the system and guidance based on past observation. The two tools address two important areas and serve to push the audience to think about possible improvements as well as other tools. The second tool, the recommendation system, also gives an interesting glimpse of how we might leverage machine learning along with current systems building practices to help human operators.