

Appliantized Massively Parallel High Performance R

Justin Lindsey

CTO, Netezza
jlindsey@netezza.com

To overcome performance limitations in both time (elapsed time of a computation) and space (total amount of information to operate on) general purpose computing has left the confines of single core computing. This departure has created powerful and significant systems that are multi-core (multithreaded) and multi-processor (parallel or distributed computing). The challenge with these paradigms is they have added significant complexity to the task of controlling and instructing the machine. A path forward using a parallel R appliance will be presented with a detailed discussion of a particular implementation and a simpler model for controlling the machine. The presentation will highlight accomplishments of the present and suggest areas for future work toward a common parallel R standard.