

Is Enterprise Big-Data Analytics for Real?

Stories of web-scale analytics on large (even petabytes) of data towards the cause of better search, social-networking and of course, targeted online advertising, are legend, and for the most part, quite real. However, when it comes to more traditional enterprises, the story is quite different. The term 'big-data analytics' is often overused, oversold, and interpreted at will to serve many purposes: "Big-data = Hadoop", "traditional BI is too expensive; so use the 'new' open source stack", "oh, we can actually do statistics in addition to slice-and-dice", "our business needs to be data-driven", "let's exploit social-media data", and the list can go on. In this talk I'll ask and attempt to answer some basic questions: (a) where data-management technology is headed and why; (b) what is 'big-data' beyond the by now standard volume, variety, velocity argument; i.e., how does 'big' data make any difference, if at all, to *how* one analyses (or 'mines') one's data; and finally (c) the difference between a technology/data-driven approach versus being 'outcome driven'.