

Issues in Networked Data Management in the Next Millennium

Michael L. Brodie

GTE Laboratories

brodie@gte.com

Surajit Chaudhuri

Microsoft Research

surajitc@microsoft.com

1. Premise

The next generation of computing will involve vast numbers of devices and humans interacting to achieve organizational, enterprise, and any number of other personal objectives. The explosive growth of the web and the emergence of a new generation of personal gizmos are two extreme examples: the web is a universal publishing platform and the PalmPilot™ is a hand-held database of contact and scheduling information. In the emerging networked world, data will reside not on a few data servers but on millions of servers and devices distributed worldwide in connected and disconnected modes. Conventional database concepts, tools, and techniques apply in the abstract but the networked world will present several discontinuities. Some fundamental database assumptions will no longer apply. What are the data management requirements in the future networked world? What are the current and future requirements and challenges for networked data management? In this industrial session, three industry leaders with major commitments to networked data management will present their views and will respond to questions.

2. A Sampling of Issues

We motivate the discussion among the panelists by raising a few key specific technical issues that are sure to play an

important role in the networked data management architectures.

A key requirement in the networked data management is to recognize the inherent *dynamic* nature of the network. Such dynamism requires mechanisms that *make universal plug and play* possible. How would such a plug and play architecture be supported? Several issues that need to be addressed in this context are role of *meta-data*, and the inherently difficult aspects of *administration* and models of *data consistency*. The networked data management infrastructure would require support for *novel programming and execution environment*. For example, the gizmos should be able to *auto-install applications* from the network with ease. The support for multiple models for data consistency will need to be complemented by support for *degrees of completeness* in the querying subsystem. Such architectural decisions will impact the nature of *caching strategy* in the network. What are the new sets of tools that we will need in the context of a networked data management? Will there be collateral changes? In particular, can current operating systems adequately support future networked requirements? Finally, we need to consider the effect of XML and other industry initiatives on networked data management.

Panelists, Christopher J. Pound from BT, Rafiul Ahad from Oracle and James Hamilton from Microsoft will express the challenges that they see and help us find answers to these complex technical issues. Dr. Pound expresses the view of a major Telcom operator, a consumer of database technology and a major owner of the networked data management problem. Ahad and Hamilton express the views of major database technology providers with different approaches to the challenges ahead.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the VLDB copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Very Large Data Base Endowment. To copy otherwise, or to republish, requires a fee and/or special permission from the Endowment

Proceedings of the 25th VLDB Conference, Edinburgh, Scotland, 1999.