This special edition of the IEEE Data engineering Bulletin is devoted to Self-Managing Database Systems, an extremely important research topic for managing today's ever-growing datasets. Since the 1999 edition on Self-Tuning Databases[1], automatic information management has become increasingly critical not only for database systems, but also for storage systems as well as middleware and applications. In all these domains there is a common trend: hardware and software depreciates over time, while the human administration costs increase. It is typical for a company today to spend $20 on human administration and maintenance for every $1 spent on hardware and software. There is currently a vast research community working on the subject of automating data management tasks, varying from automatic configuration and organization to self-maintenance and healing.

The invited articles in this bulletin demonstrate the tremendous advancement in the field. The first three articles reflect the commercial state-of-the-art: an overview of the automated database management features in three leading commercial products (Microsoft's SQL Server, IBM's DB2, and Oracle). Research teams from these companies describe tools integrated in their products that automate physical database design as well as other daunting tuning tasks. The next article by Pautasso et.al. describes the design and development of an autonomic workflow engine, which can be used to compose large-scale system services. Moving on to system resource management, Narayanan et.al. present the design of a resource advisor which answers questions such as "how would the performance of my OLTP application change if doubled the main memory?". The last two articles reflect efforts on automating storage management: Qiao et.al. describe PULSTORE, an analytic framework to transparently alter storage configuration to satisfy a time-varying I/O workload while maintaining QoS guarantees. Finally, a large team from the Parallel Data Lab at Carnegie Mellon University report on the design and implementation of a self-* storage system, which can manage itself using performance predictions, thereby greatly simplifying tuning automation.

I would like to thank Sam Lightstone, Chair of the IEEE Data Engineering Workgroup on Self-Managing Database Systems, who kindly agreed to foreword this issue with a letter summarizing the current trends in the field and the workgroup's activities. I also cordially thank David Lomet and all the authors who graciously contributed their time and effort to make this special edition a reality. I hope that you will enjoy reading it.

[1] IEEE Data Engineering Bulletin Volume 22, Number 2, June 1999 Special Issue on Self-Tuning Databases and Application Tuning