

Modern large-scale computing systems are typically a complex multi-vendor tangle of storage devices, databases, middleware, and other subsystems. If present trends continue, the difficulty of designing, configuring, and managing such systems will completely overwhelm software developers and system administrators within a few years, with the potential of bringing the information technology industry to a standstill. Foreseeing this looming crisis, IBM Research launched the autonomic computing initiative in 2001 to create computing systems that manage themselves in accordance with high-level guidance from humans. In the ensuing years, autonomic computing has become a key cross-disciplinary effort of over one hundred IBM researchers working in artificial intelligence, distributed and fault-tolerant computing, performance modeling and analysis, services computing, data management, storage systems, human-computer interfaces, and other areas.

**TECHNOLOGIES**

At the level of individual self-managing subsystems, improvements are being made to the capability of individual subsystems, like database and storage, to configure or optimize themselves, obviating the need for administrators to understand and manipulate hundreds of tuning parameters. More generic autonomic computing technologies — planning, optimization, modeling, feedback control, data mining, and machine learning, for example — that can be used within a broad array of autonomic subsystems are also being developed. IBM researchers are not only devising state-of-the-art algorithms, but are also overcoming major challenges inherent in situating these algorithms in real, complex computing environments. In addition, vendors are using IBM’s autonomic computing toolkit to construct their own autonomic subsystems.

Above the subsystem level, technologies are being developed that rely on interactions among many subsystems to achieve system-wide robustness and efficiency. In one such effort, several state-of-the-art technologies have been integrated, resulting in a system that intelligently inserts and removes probes to pinpoint failures based on automatically discovered system topology, and uses planning and knowledge-base technologies to work around or replace failed components.





Image at right was previously published in *IBM Systems Journal*, Volume 42, No. 1 (2003), cover.

**SYSTEMS**

To explore architecture and technology integration issues, prototypes of self-managing medium- to large-scale systems are being built. An autonomic data center prototype has already been constructed that assembles itself from subsystems that contain very little *a priori* knowledge, heals itself in the face of subsystem failures, and tunes its resource usage as workloads fluctuate. Finally, to assess system self-management capabilities, novel types of benchmarks are being developed that go beyond traditional performance measurements to quantitatively measure self-healing and self-configuration.

**THE HUMAN ELEMENT**

The human element is critically important in autonomic computing. Consequently, as part of IBM’s research program, ethnographic studies of system administrators’ behavior are being conducted, in an effort to learn how they work and what tools they most need. Furthermore, policy languages and tools are being developed that allow administrators to express their objectives naturally and precisely, and to understand the likely impact of objectives and actions upon system behavior.

**COLLABORATION WITH INDUSTRY AND ACADEMIA**

Recognizing that no organization can realize the ultimate vision of autonomic computing single-handedly, our researchers have helped IBM assume leadership positions in many pertinent efforts within the Organization for the Advancement of Structured Information Standards (OASIS), the World Wide Web Consortium (W3C), and other standards bodies. They have also fostered the growth of an active worldwide research community by collaborating with academic partners to establish numerous workshops and conferences, such as the International Conference on Autonomic Computing.