Practical Autonomic Computing: Roadmap to Self Managing Technology

A White Paper Prepared for IBM
January 2006

ENTERPRISE MANAGEMENT ASSOCIATES®

Table of Contents

Practical Autonomic Computing: Roadmap to Self Managing Technology

Executive Summary

Today's Information Technology (IT) operations support model is rapidly becoming untenable. According to a recent IDC study commissioned by IBM Tivoli, approximately 70% of today's IT budget is labor. Because IT is increasingly fundamental to revenue in today's technology-driven businesses and will likely continue to grow in size and complexity, companies are facing a growth versus cost dilemma. IT complexity has to be simplified, or growth will begin to be hampered by the cost required to support it.

"Autonomic Computing" is IBM's term for an approach and blueprint including a set of products, tools and services that add self-managing capabilities to Information Technology systems. Its goal is to shift the burden of support tasks such as configuration, maintenance, and fault management from people to technology. Autonomic systems are designed to take over manually intensive IT operations tasks that IT professionals choose to delegate, many of which are routine and repetitive.

Autonomic capabilities are critical to businesses with large and complex IT environments, those using Web Services and/or Service Oriented Architecture (SOA) models, and those that leverage e-business or e-commerce. They are also key enablers for smaller businesses seeking to take advantage of current technologies, because they help mask complexity by simplifying infrastructure management.

Since first being proposed in 2001 by an IBM executive, the autonomic concept has been adopted by today's leading vendors and incorporated into their products. Aware that success is tied to interoperability, many are participating in the standards development necessary to provide the foundation for self-managing technological ecosystems, and are integrating standards into their technology.

Meanwhile, IBM is making a substantial investment in the autonomic concept and has released its first wave of standards-based components, tools and knowledge capital. This paper details IBM's approach and discusses ways in which it can be leveraged to control costs while delivering cost-effective service. The era of self-managing systems has arrived, and now is the time to begin to reap its benefits.

Introduction

Early computers, "electronic brains," relied on vacuum tubes rather than transistors to process information. Needless to say, maintaining those tubes and troubleshooting when one burned out meant that those early devices were out of service more than they were working. Today, when we think of some poor operations technician examining thousands of vacuum tubes to figure out which one to replace, the concept of supporting today's computing environments in the same way is unimaginable.

In a similar way, many of the tasks and processes associated with today's IT infrastructure support techniques are rapidly becoming outmoded. Gathering 15 or 20 people in a "war room" and poring over sniffer traces and log files to resolve a client-server issue is prohibitively expensive and too time consuming to address today's need for near perfect service uptime. Engineers of the future will look back on the effort we expend on simply keeping technology running and say, "When did they have time to get any real work done?"

IT Infrastructure is becoming exceedingly complex. Service Oriented Architecture (SOA) enables companies to execute software modules that lie outside their own infrastructure and control, for example at a business partner's site. Corporate websites require multiple, interconnected resources, including application, web and database servers as well as network equipment and software, to operate. New requirements for deployment and problem determination across geographically distributed technology introduce the need for functionality that is beyond the capability of yesterday's enterprise management tools.

Autonomic Computing is a term used to describe a broad range of standards, architecture, products, services and tools that enable technology systems to be self-managing. Just as today's computer systems contain transistors instead of vacuum tubes and have become extremely reliable, tomorrow's IT technology will largely manage itself without the need for humans to perform today's most time-consuming, resource-intensive and mundane management tasks. This isn't a "Star Wars" view of some other galaxy at another place and time, but rather an evolution that is already taking place. Autonomic technology and standards exist and continue to evolve,

and IT thought leaders are taking action on the need to integrate them into processes and products.

The promise of Autonomic Computing includes capabilities unknown in traditional products and toolsets. It includes the capacity not just to take automated action, but to do so based on an innate ability to sense and respond to change. Not just to execute rules but to continually normalize and optimize environments in real time. Not just to store and execute policies, but to incorporate self-learning and self-managing capabilities. It is a landscape that promises to ease the pain of taking IT into the future, by shifting mundane work to technology and freeing up humans for work that more directly impacts business value.

What is Autonomic Computing – Really?

Over the past 20 years, the maturing of Information Technology has been the primary driver of business change and growth. At the same time, IT management has become increasingly complex because the systems being managed have themselves become complex. Anticipating that support of technology will continue its upward trend of cost and complexity, some forward thinking IT architects have devoted considerable time and effort to devising solutions. It was out of this process that the idea of Autonomic Computing arose.

The self-managing concept comes as a result of a fundamental shift in device management that actually began back in the 1970's. In an effort to improve reliability, availability and serviceability, IBM applied technology to mainframe support by utilizing scripts to bypass failing components and to keep systems running while they were being replaced. In the 1980's, the notion evolved into centralized consoles that enabled management of multiple systems from one workstation and automatic resolution of common problems.

The term "Autonomic Computing" was actually coined in 2001 by Paul Horn, a senior vice president of research for IBM. He saw the trend towards smaller, cheaper and more powerful systems and correctly anticipated the result we are seeing today – that companies would begin to amass a proliferation of interconnected devices, then struggle to install, configure, manage and maintain them. He also predicted that, if new and radical solutions were not developed to address this problem, it would eventually limit the growth and evolution of technology-enabled business. Just as our modern communications system would not be possible if we relied on telephone operators to route every call and connection, the sheer complexity of the infrastructure management challenge would eventually grow so heavy that we would not be able to continue to progress technologically. The concept of Autonomic Computing was proposed to answer this need.

The name is derived from the body's autonomic nervous system. It controls activities, such as heartbeat, blood pressure and breathing that enable the body to self-regulate and adapt to changing conditions. The same concept is behind Autonomic Computing, although the two are different in terms of the way their activities are carried out. In humans, autonomic functions are largely involuntary, regulated by cells in the brain that process visceral information, correlate it, and issue specific regulatory instructions via the nervous system. In technology, autonomic activities process in real time the information arriving from the infrastructure, integrate that information, and then issue specific regulatory instructions guided by business oriented policies that direct specific IT actions.

In short, an autonomic system senses its operating environment, models its behavior in that environment, and takes action to change the environment or its own behavior. It has the following four properties:

- Self-Configuring
  - Characteristics that enable systems to adapt to changing conditions by changing their own configurations
  - Functionality that allows the addition and removal of components or resources without service disruption
- Self-Healing
  - Capacity to recognize and diagnose deviations from normal conditions and take action to normalize them
  - Capability to proactively circumvent issues that could cause service disruptions
- Self-Optimizing
  - Ability of the system to monitor its state and performance and proactively tune itself to respond to environmental stimuli

- Self-Protecting
  - Incorporation of intelligence to recognize and circumvent security threats
  - Facility of a system to protect itself from physical harm, such as excessive heat or motion

Since its inception, the idea of Autonomic Computing has evolved beyond the borders of IBM and is now a core concept in the products of most key technology vendors, although many have their own names for self-managing technology as it applies to their specific platform. Platform specific programs, while helpful, do not solve the larger problem of complexity that Autonomic Computing focuses on, that arises from heterogeneous, diverse IT infrastructures that include multiple platforms and vendors. Autonomic Computing is unique in its holistic focus – one that embraces the entire IT infrastructure with its diverse set of technologies. Autonomic Computing is unique in its holistic focus – one that embraces the entire IT infrastructure with its diverse set of technologies. Although the vision will continue to evolve over many years, we are seeing self-managing capabilities already in most product domains. Vendors are incorporating the concept of using technology to manage itself into their product designs, and, where these concepts support open standards, customers are starting to reap the rewards in terms of interoperability.

To facilitate this evolution, IBM has detailed an architectural blueprint for Autonomic Computing that describes the building blocks required for disparate products to work in concert. These building blocks are as follows:

- Autonomic manager: The autonomic manager manages a system resource by collecting information from a monitored resource, analyzing the information, using policies to plan responses, and executing the appropriate responses for a particular state. The functions of monitoring, analyzing, planning and execution are collectively called the MAPE loop.
- Touchpoint: A touchpoint, sometimes called a manageability endpoint, is a consistent, standard manageability interface for accessing and controlling a managed resource.
- Managed Resource: A managed resource can be any type of hardware or software resource, for example a server.
- Knowledge source: A knowledge source is a registry, dictionary or other repository that contains data relevant to the autonomic system. This knowledge source shares its information with the autonomic manager, giving it supporting data required for its activities.
- Manual manager: This is a user interface or console that enables a human to perform management activity. It enables the IT professional to perform tasks related to the Autonomic Computing infrastructure, such as policy definition and delegation of actions and tasks to autonomic managers.
- Enterprise service bus: The enterprise service bus coordinates and integrates the other building blocks by connecting them and directing their interactions.

Figure 1 shows a Problem Determination sequence using the building blocks described above.

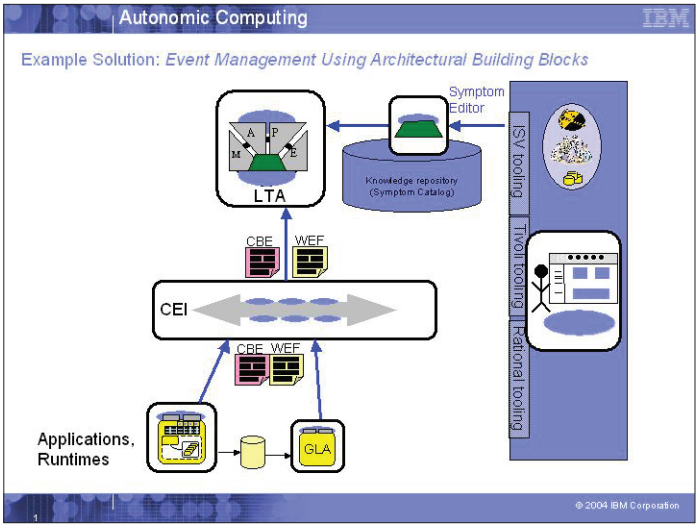

Figure 1: Problem Determination using Autonomic Computing Building Blocks

Products and services that incorporate Autonomic Computing capabilities do so within the framework of these building blocks, with different products and tools addressing different aspects of the architectural blueprint. We'll talk more about products and services in the following pages.

Getting Started with Autonomic Computing

Pain is often the driver for change. In the business world, pain arises from the need to save money, improve competitive positioning, enable growth or increase operational efficiency. All four are currently impacting the typical IT organization and autonomic concepts, products and tools can help to mitigate the negative impact of each.

The need to do more work for less money underlies all of these pain points, and organizations have reacted to this reality in a variety of ways ranging from layoffs to outsourcing. Obviously, investing in self-managing systems and building them into the organizational framework can result in considerable payoff in terms of maximizing utilization of personnel. Self-managing autonomic technologies free up staff to work on strategic, revenue impacting projects instead of tactical problem solving and fire fighting.

It's not hard to figure out an organization's pain points. The more difficult exercise is to figure out how to address them. How can an organization best prepare itself to leverage the payoffs of self-managing technology?

- **Evolution versus Revolution.**
  A key idea to keep in mind is that Autonomic Computing, like SOA, Web Services, Information Technology Infrastructure Library (ITIL) or any other emerging concept, is best approached as an evolution versus a revolution. The processes, standards and technology that comprise it are still early state and will continue to evolve over many years. For example, the Web Services Distributed Management (WSDM) standard version 1.0, with autonomic common base events as input, was approved by the Organization for the Advancement of Structured Information Standards (OASIS) in March, 2005. The first WSDM based tools for developers debuted on ibm.com/alphaworks in November, 2005. This standard describes the communication process that enables individual infrastructure components to interact with one another, allowing autonomic systems to execute functions such as managing and configuring. An OASIS technical committee was formed in May, 2005 to review another standards schema, the Solution Deployment Descriptor (SDD) specification, which specifies a software distribution standard that will enable autonomic configuration. Because these capabilities provide core functionality for Autonomic Computing, the conclusion we can draw is that enabling standards for the grand technology vision that is Autonomic Computing are now becoming available but are still evolving, and products will evolve with them.

  Does this mean that companies should wait 3 or 4 years until the market is more mature to begin integrating self-managing products into their environments? Absolutely not. Agile businesses, very likely competitors, are already factoring autonomic products and concepts into their IT architectures. The optimal way to move towards this new capability is to develop a plan that starts with small steps and proceeds incrementally, keeping pace with the marketplace.

- **Standardize on standards-based products**
  Plan infrastructure build-outs and purchases with a view towards the future and look for vendors who are not only incorporating Autonomic Computing functionality into their products, but are also utilizing, or planning to utilize, open standards specifications such as WSDM and SDD. In this arena, products relying on proprietary or non-standard protocols and software are not the best choice. To maximize interoperability as your organization moves towards more complex technologies, make sure that the products you select are standards based.

- **Optimize and document IT support processes**
  Making best use of the technology is impossible without building a foundation to bring people and processes along with it. Although industry best practices such as ITIL have now entered the mainstream, with recent EMA research finding more than 47% of large companies adopting at least some of the ITIL functions, well designed and documented processes have become even more important with the advent of Autonomic Computing. In fact, processes can be considered a fundamental enabler of this technology. After all, IT professionals are responsible for determining which tasks can best benefit from autonomic self-managing technologies and developing the policies

that direct the autonomic function, and a policy often starts as a technical implementation of a defined manual process. From this perspective, one of the best preparations an organization can make is to invest in ITIL. Tivoli ITSM, based on ITIL, provides a pragmatic approach to automating IT processes.

The quickest way to do this is by partnering with experts in the area of IT Service Management (ITSM). ITSM is considered to be the "heart" of ITIL, and its processes address both day-to-day operation and support of IT as well as long term planning and improvement. IBM's approach to IT Service Management is considered to be one of the best in the business. It consists of IT Process Management products, IT Operational Management products, and the IT Service Management (ITSM) platform. This approach, which combines robust, workflow-based knowledge collateral and substantial solutions with a high degree of services expertise, can jumpstart the ITIL implementation process, building a solid foundation for implementation and support of today's complex IT architectures.

- **Bring IT staff at all levels into the plan**
  The people component in the equation cannot be overstated. Experience has shown IT professionals to be notoriously cautious of systems that automatically execute management activities, and rightly so. Their role is to keep systems running. Their caution is based on the fear that, with machines making decisions at a blinding pace, they have that much more potential to really make a mess of things. Once that happens, IT engineers are the ones who field complaints from customers and fix the problem. These IT professionals need to build up trust in the automated activities of a self-managing system if they are to willingly delegate management activities.

  This is another reason to adopt a steady, measured adoption plan for autonomic systems. IT engineers, not just architects and designers, should be instrumental in choosing technology investments and providing input to architecture decisions. In addition, they should be key contributors during the test and deployment phases. This not only enables them to begin to trust the decisions made by the automated management system, and hence

	Level 1 Basic	Level 2 Managed	Level 3 Predictive	Level 4 Adaptive	Level 5 Autonomic
Process	Informal, reactive, manual	Documented, improved over time, use of industry best practices, manual process to review IT performance	Proactive, shorter approval cycles	Automation of resource-management and transaction-management best practices, driven by service level agreements	Automation of IT Service Management and IT Resource Management best practices
Tools	Local, platform and product specific	Consolidated consoles, problem management system, automated software install, intrusion detection, load balancing	Role based consoles, product configuration advisors, real time view of current and future IT performance, automation of some repetitive tasks, common knowledge base of inventory and dependency management	Policy management tools drive dynamic change based on resource-specific policies	Costing/financial analysis tools, business and IT modeling tools, trade-off analysis, automation of some e-business management roles
Skills	Platform specific, geographically dispersed with technology	Multiple platform and multiple management tool skills	Cross-platform system knowledge, IT workload management skills, some business-process knowledge	Service objectives and delivery per resource, analysis of impact on business objectives	e-business cost-and-benefit analysis, performance modeling, advanced use of financial tools for IT context
Benchmarks	Time required to fix problems and complete tasks	System availability, time to close trouble tickets and work requests	Business system availability, service level agreement attainment, customer satisfaction	Business system response time, IT contribution to business success	Business success, competitiveness of service level agreement metrics, business responsiveness

Figure 2: Autonomic Computing Maturity Model

view autonomic technology as an ally rather than a threat, but also uses their front line perspective and experience to optimize the implementation process.

- **Develop a roadmap**
  Take some time to evaluate where your organization stands in terms of autonomic capabilities, and where it needs to be. Then come up with a step-by-step plan to get there. To facilitate this process, IBM has developed an Autonomic Computing maturity model that is used by its consultants to help consulting customers evaluate current readiness and capability. The model, shown in Figure 2, incorporates the aspects of Process, Tools, Skills and Benchmarks, describing what each looks like at the Basic, Managed, Predictive, Adaptive and Autonomic maturity levels. From this model, an organization can create a roadmap to autonomic capability, composed of activities and characteristics that describe each step.

  For example, from the Process perspective, if an organization has only informal IT support processes and approaches each new incident in an ad-hoc, reactive way, they are at Level 1, or the Basic level. To get to Level 2, they must document support processes and start to adopt industry best practices, such as ITIL. This same approach can be taken to the Tools, Skills and Benchmarks perspectives, eventually charting a course towards becoming more mature as an organization and therefore better prepared to incorporate self-managing capabilities.

- **Start with a single pain point**
  Regardless of maturity level, the primary focus should be to address pain points and look for quick wins. For example, one of the biggest "black holes" in terms of staff time tends to be problem determination and troubleshooting. IBM has developed an autonomic tool called the IBM Autonomic Log and Trace Analyzer, which the Global Services organization has used to cut problem determination and troubleshooting time by 80% at some locations. This tool is available to vendors and is also available for free download. It is used to correlate, isolate and address problems that arise within diverse and heterogeneous IT environments. It leverages the end-to-end capability made possible by the newly ratified WSDM standard. This is one way to start with a small step that achieves tangible payoffs and immediate results.

- **Incorporate self-managing technology/ capabilities into your IT infrastructure**
  In addition to leading the industry towards standards development, IBM has developed an impressive array of partnerships, products and services that are described below. IBM products not only integrate autonomic technology, they support the steps leading up to its adoption. For example, Tivoli is building powerful capabilities into its CMDB (Configuration Management Data Base) to provide the modeling capabilities required for complex architectures. In addition, IBM's IT Service Management approach incorporates Tivoli tools that directly support ITIL processes. IBM has also published sample process flows as part of its ITSM strategy that can be used by any organization seeking to adopt ITIL-based best practice processes.

Products and Services and Where They Fit

IBM Products

Over the past four years, IBM's Autonomic Computing group has achieved the following landmarks:

- Architectural blueprint (published at http://www.ibm.com/autonomic/pdfs/AC Blueprint White Paper V7.pdf)
- Autonomic Computing Manageable Resource Interface Specification and Touchpoint Implementation Guide (published at http://dl.alphaworks.ibm.com/technologies/aide/ManageabilityEndpointSpecification.0.51a_draft.pdf)
- 550 patent applications
- 20+ specifications to standards bodies
- 450+ self-managing autonomic features in 70+ IBM products
- Services to accelerate customer implementation of Autonomic Computing self-managing technologies such as problem determination.

- 60+ business partners developing products according to Autonomic Computing standards and technologies
- More than 100 customer references

IBM technology and toolkits are available on developer portals. To help incorporate these key concepts into technology, IBM has developed a standards-based toolset used to build its own products and to provide to other vendors seeking to add autonomic capabilities. Currently, this toolset includes:

- Common Base Events: Standard XML structures for communicating events in a common way. Common Base Events is IBM's first implementation of the WSDM standard described earlier.
- Generic Log Adapter: Adapt log files from legacy applications to the Common Base Event format
- IBM Autonomic Log and Trace Analyzer: Simple implementation of a partial Autonomic Manager that assists administrators by displaying and correlating events (in the Common Base Event format) that are gathered from log files.
- Resource Model Builder: Generates data models for monitoring resources and taking action based on the monitored data.
- Automated Management Engine (AME): Using resource models generated with the Resource Model Building, AME collects information from a monitored resource, analyzes the information, uses policies to determine responses, and executes the appropriate response for a particular state. AME is an implementation of an autonomic manager, and hence is one instance of the MAPE Loop.
- Integrated Solutions Console (ISC): Web based console for consolidated management of Autonomic Computing systems. ISC operates in a portal-based environment.

IBM Tivoli products that incorporate MAPE Loop concepts into their technology:

- IBM Tivoli Monitoring: monitoring
- IBM Tivoli Intelligent Orchestrator: analysis
- IBM Tivoli Provisioning Manager: planning and execution
- Products specific to each resource perform the execution activities for each manageability endpoint

An extensive library of business process flows, called Service Accelerators, are available without charge for those companies seeking a guide to process improvement. This library represents extensive expertise developed in conjunction with IBM Global Services that, for most consulting organizations, represents proprietary knowledge capital that is unavailable to non-customers.

For additional information about IBM Autonomic Computing tools and products, go to:

http://www.ibm.com/autonomic/

http://developer.intel.com/technology/itj/2004/volume08issue04/art03_autonomic/p07_ibm_toolkit.htm

http://www.ibm.com/autonomic/wsdm/

http://www.ibm.com/developerworks/autonomic/library/ac-sivalue/

Self-Managing Autonomic Technology Mark

The IBM Self-Managing Autonomic Technology mark program is an advanced-level differentiator offered to IBM Business Partners who are members of the IBM Autonomic Business Partner Initiative. It is open to software vendors who have been certified by IBM as incorporating at least two of the following three IBM Self-Managing Autonomic Technology solutions into their solution sets:

- Problem determination/common logging technology
- Autonomic Management Engine technology
- Solution Installation technology

Vendors in this category include Corente, Macrovision, Network Physics, nLayers and Singlestep Technologies.

For additional information, go to:

http://www.ibm.com/autonomic/mark.shtml

Business Partners

Vendors from across the IT industry are working together to develop, adopt and incorporate autonomic concepts. For a current listing, go to:

http://www.ibm.com/autonomic/partners/solution_list.shtml

Services

IBM offers a wide array of service offerings, backed by methodology and tools, which enable and support the adoption of Autonomic Computing. One such offering is the Readiness Assessment. This is a Consulting Services offering tailored to each individual organization, during which an expert skilled in Autonomic Computing concepts helps a company crystallize its own business goals and objectives into a detailed roadmap. It is an excellent way to bring technology staff up to speed on the innovative processes and technology at the core of Autonomic Computing, and also helps to develop a vision for incorporating processes, products and services based on current technology, infrastructure, best practice, and business requirements. For details, go to:

http://www.ibm.com/autonomic/pdfs/wp-igs-autonomic.pdf

EMA's Perspective

Autonomic Computing is a concept whose time has come, and one that will gain industry focus over the next few years as business revenue continues to become increasingly intertwined with complex technology. Because support costs and infrastructure complexity are rapidly escalating, businesses should strongly consider alternatives to current support paradigms. Self-managing systems are the leading contender in this arena. Especially when combined with well-planned adoption of industry best practices, investment in tools incorporating self-managing characteristics should be weighed as a viable alternative to do this in a cost-effective way.

IBM's achievements in the area of Autonomic Computing are an instance, rare in the technology world, where actual product functionality is well ahead of marketing hype. The early products, such as the IBM Autonomic Log and Trace Analyzer, have shown their value in countless consulting engagements, and the emphasis on ITIL alignment and a robust CMDB can be considered to be major product differentiators. Where many vendors are simply repackaging older solution sets and revamping marketing messages to support current trends, IBM has heavily invested money and staff to deliver actual solutions well ahead of most of the industry.

IBM is still building out its Configuration Management Data Base solution to support autonomic capabilities, with general availability of a new, federated version targeted for June, 2006. The autonomic message still needs honing, as the architectural components and relationships, product-to-functionality mapping, and gaps in the product roadmap are still difficult for an outsider to trace. This would be a great help, both for potential customers and vendor partners. But overall, IBM should be proud of what it has accomplished and prospective customers should be aware that IBM's commitment to Autonomic Computing is already pervasive in today's solution offerings, and gaining traction for future products.

About Enterprise Management Associates, Inc.

Enterprise Management Associates, Inc. is the fastest-growing analyst firm focused on the management software and services market. EMA brings strategic insights to both vendors and IT professionals seeking to leverage areas of growth across e-business, network, systems, and application management. Enterprise Management Associates' vision and insights draw from its ongoing research and the perspectives of an experienced team with diverse, real-world backgrounds in the IT, service provider, ISV, and publishing communities, and is frequently requested to share their observations at management forums worldwide.

Corporate Headquarters:
Enterprise Management Associates
2585 Central Avenue, Suite 100
Boulder, CO 80301, U.S.A.

ENTERPRISE MANAGEMENT
ASSOCIATES®

Phone: 303.543.9500
Fax: 303.543.7687
info@enterprisemanagement.com
www.enterprisemanagement.com

1045.022706