Pulse Monitoring: Extending the Health-check for the Autonomic GRID

Sterritt, R. (2003). Pulse Monitoring: Extending the Health-check for the Autonomic GRID. In *Unknown Host Publication* (pp. 433-440). IEEE. <https://doi.org/10.1109/INDIN.2003.1300375>

[Link to publication record in Ulster University Research Portal](#)

Published in:
Unknown Host Publication

Publication Status:
Published (in print/issue): 01/08/2003

DOI:
[10.1109/INDIN.2003.1300375](https://doi.org/10.1109/INDIN.2003.1300375)

Document Version
Publisher's PDF, also known as Version of record

General rights
Copyright for the publications made accessible via Ulster University's Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy
The Research Portal is Ulster University's institutional repository that provides access to Ulster's research outputs. Every effort has been made to ensure that content in the Research Portal does not infringe any person's rights, or applicable UK laws. If you discover content in the Research Portal that you believe breaches copyright or violates any law, please contact pure-support@ulster.ac.uk.

Pulse Monitoring: Extending the Health-check for the Autonomic GRID

Roy Sterritt

*School of Computing and Mathematics,
Faculty of Engineering
University of Ulster
Northern Ireland
r.sterritt@ulster.ac.uk*

Abstract

This paper upon looking at the Autonomic Computing architecture and Grid Computing highlights the importance of health check mechanisms to achieve a reflex-healing duel strategy. This will provide new design options for the development of the Autonomic Grid. The resulting pulse monitor is based on extending the existing Grid heart-beat monitor with urgency or anxiety levels such as that used in the NASA beacon monitor. The paper concludes with a discussion that this health check mechanism may be utilized in the future to achieve the necessary sense of urgency within a system for affect and emotion intelligence.

influence on the development and progress of this new and exciting area. One particular area is that of dependable systems which already focus on the reliability, availability, safety, security, survivability and maintainability of systems [3]-[5].

This paper will first describe the general architecture of an autonomic system and its need for the health check mechanism as initially introduced in [23]. Previously published health check mechanisms are described with particular attention to the heart-beat and beacon monitoring. A new mechanism referred to as pulse monitoring is then introduced and considered for Grid Computing. The paper concludes with a discussion that this pulse mechanism may introduce the needed sense of urgency for the future development of affect and emotion intelligence within systems.

1. Introduction

Computing systems are becoming more and more complex. Real world system realities such as a high degree of networking, extreme number of nodes, boundary-less and concurrent evolution all result in the system becoming more ambiguous ensuring serious management challenges [24]. *Autonomic Computing*, launched by IBM in 2001 [1], is emerging as a valuable new approach to the design of effective computing systems.

The inspiration for Autonomic Computing has stemmed from the human body's autonomic nervous system. Computing systems that can adapt to changing environments and repair minor physical changes would be regarded as a major step forward in the way computer systems operate. This computing system would then behave very much like the autonomic behavior in the human body.

In order for Autonomic Computing to succeed it will need to draw on well-established resources within computing science. For instance, the collaboration of Software Engineering and Artificial Intelligence will be facilitated by Autonomic Computing [2]. Other major areas of computer science will also have a significant

2. Grid Computing

A grid infrastructure promises seamless access to computational and storage resources, and offers the possibility of cheap, ubiquitous distributed computing. Grid technology will have a fundamental impact on the economy by creating new areas, such as e-Government and e-Health, new business opportunities, such as computational and data storage services, and changing business models, such as greater organizational and service devolution [6][7]. The Grid is a very active area of research and development; with the number of academic grids jumping six fold in the last year [26].

The Grid historically arose out of a need to perform massive computation, the current direction demonstrates the potential to change the structure of electronic service provision and create a new grid service economy. The success of the grid will be founded on the development of new grid-enabled software systems and the evolution of legacy systems to grid-enabled systems. There are many middleware frameworks for distributed computing, many modeling techniques for software artifacts, and many development processes for

controlling the creation of new software systems and managing the evolution of existing software systems.

A fundamental challenge is creating correct, robust, flexible and cost-effective grid-enabled software [8]. The Grid aims to be self-configuring, self-tuning and self-healing, similar to the goals of Autonomic Computing [25]. Its aim to fulfill the vision of Corbato's Multics [27]—like a utility company, a massive resource to which a customer gives his or her computational or storage needs [25]. As such, it is expected that Autonomic Computing will be required to provide some of the answers to achieve this vision.

3. Autonomic Computing Environment

The general properties of autonomic computing systems are summarized in Figure 1 [5]. The main goal of an Autonomic System is to be self-managing. This consists of four main objectives—self-configuring, self-healing, self-optimizing and self-protecting.

In achieving these self-managing objectives a system must have the following attributes: self-awareness, environment awareness, self-monitoring and self-adjusting.

Further information on autonomic computing properties can be found in IBM's autonomic 'manifesto' [1] and subsequent 'blueprint' [9]. Also in terms of dependability see [5] and for the initial discussion of the autonomic computing environment see [23].

An autonomic computing system is made up of a connected set of *autonomic elements*. Each element must include *sensors* and *effectors* [10]. The sensors are used to monitor the behavior of the system while the effectors are used to facilitate any actions that are required. The process begins with an assessment of the system using the sensors which compare the observed situation with that which is expected. The sensors then make a decision on whether an action is required. If an action is necessary, it is then executed by utilizing the effectors, thus creating a control loop [11].

This control loops is built into each autonomic element. Figure 2 shows a possible system architecture to support this model, as first discussed in [23]. Each autonomic element consists of a *managed component* and a corresponding *autonomic manager*. Hence the autonomic manager implements the required self-monitoring and self-adjusting, i.e. the control loop.

The state of the sensors (to the managed component) are constantly being observed by the *internal monitor* and evaluated and assessed for possible action by the *self monitor*. A system knowledge base exists for reference containing the expected states of the system which may then be used as a means of comparison with

the actual observed behavior of the system. Deviations are reported to the *self adjuster* for action, which may result in changes to the *managed component* through the effectors. Similarly, an *external monitor* observes the state of the environment via an *autonomic signal channel* and this also may trigger internal changes. The signal channel provides linkage to other autonomic managers. These may be virtual (in the same physical system), peer-to-peer, client-server [12] or Grid [28].

The *heartbeat* function through registration with external elements provides a security mechanism—'I am alive' monitoring.

The *pulse monitor* provides an indicator of the health of the system—analogy with measuring the actual pulse of a biological system instead of just checking the presence or absence of a pulse.

In Figure 2 communication with the external environment via the autonomic signal channel and heartbeat/pulse signals have been logically separated to facilitate the discussion in this paper. The physical reality is more than likely that they will share the same communications infrastructure, as presented in [23], unless large scale redundancy is built into the system.

Figure 3 suggests how autonomic elements are connected—again the logical distinction for discussion is made between the two types of communications.

Communication occurs asynchronously between the autonomic elements within a system (Figure 3) and the artifacts within an autonomic element (Figure 2) [9].

The design of the monitoring protocol must ensure that the monitoring activity and traffic are maintained at acceptable levels.

Figure 1 Autonomic Computing Tree

Figure 2 Potential Architecture of an Autonomic Element

Figure 3 Autonomic Computing Environment

There are a number of ways in which the 'health' of autonomic elements might be monitored. One is to have dedicated elements for that purpose. Another is to distribute monitoring responsibility among the autonomic elements so that they monitor each other using some form of gossip protocol. This should increase robustness but will make the system more difficult to manage.

4. Health Check Mechanisms

This section looks at the general concept of health check mechanisms, the duel strategy they offer and two specific manifestations of the mechanism.

4.1 Reflexes and Healing

Reflexes and healing is a duel strategy approach concept inspired by biological systems [13]. Animals have a reflex system, where the nerve pathways enable rapid response to pain. Reflexes cause a rapid, involuntary motion, such as when a sharp object is touched. The effect is that the system reconfigures itself, moving away from the danger to keep the component functioning.

The body will heal itself on a much longer timescale. Resources from one part of the system are redirected to rebuild the injured body part, including repair of the reflex response network. While this cannot help in the real-time response, directly after an event, it can prepare the system for the next event. In addition, it can readjust the system for operation with a reduced set of resources [13].

An example of the duel approach is being developed for high energy physics experiments with use of massive facilities to delve into the basic composition of matter [13]. In this case the data is so extensive that it is practically impossible to collect all data – decisions must be made in real-time as to whether or not an interesting event has occurred. During its several year lifespan only a small number of novel events are expected. Downtime for the computing system is not an option since this may be when a novel event occurs! Due to the expensive of the overall experiment duel or triple mode redundancy is precluded.

Essentially the design allows for non-critical applications to be overwritten upon fault conditions where the reflex reaction will cause a re-configuration to ensure the matter experiments are still being adequately (less than optimal) monitored. The healing approach then attempts to re-optimize the system with the remaining resources.

The logical difference between the pulse signal and general event messages has been highlighted in Figure 2 and Figure 3 since essentially the pulse provides the mechanism for a reflex reaction whereas the general event messages under fault conditions form part of the slower healing process—root cause analysis from the event stream.

4.2 Heartbeat Monitoring

The Globus OGSA (open grid services architecture) has a health-check facility referred to as the *Globus Heartbeat Monitor*. It is designed to detect and report processes that fail to provide a 'heartbeat' [14].

Figure 4 Heartbeat Monitoring

Within the Globus OGSA specification the HBM consists of

- HBM Client Library,
- HBM Local Monitor, and
- HBM Data Collector.

Essentially a Local Monitor runs on each host, checking and reporting on the status of the monitored processes and the actual system generating "I-am-alive" messages (heartbeats). The Data Collector receives heartbeat messages and identifies failed components from missing heartbeats. There may be one or more data collector per application [15].

The heartbeat monitor function ensures an element is operating (Figure 4). More information is required, however, to determine how well an element is performing, if it is necessary to improve its operation, consistent with the needs of autonomic computing.

4.3 Beacon Monitoring

New paradigms in spacecraft design are leading to radical changes in the way NASA designs the craft's operations [16]. Increased constraints on resource usage and greater focus on operations costs has led NASA to utilize adaptive operations and onboard autonomy [17]. NASA missions, particularly those to deep space, are

considering autonomic decision making to avoid the unacceptable lag time between a craft encountering new situations and the round-trip delay in obtaining guidance from mission control. Two of the first notable missions to use autonomy are DS1 (Deep Space 1) and the Mars Pathfinder [18].

The *Beacon Monitor* concept, first used in the DS1 mission work [19] automates the routine task of health monitoring and migrates the process from ground to the spacecraft [16]. With beacon monitoring, the spacecraft sends a signal to the ground that indicates how urgent it is to track the spacecraft for telemetry. This concept involved a paradigm shift for NASA from routine telemetry downlink and ground analysis to onboard health determination and autonomous data summarization [19].

In high-level concept terms, the beacon monitor is similar to the heartbeat monitor in grid computing, with the addition of a tone to indicate the degree of urgency involved. The following table summarizes the tone definitions [20]:

Nominal	All functions as expected no need to downlink.
Interesting	Interesting – non-urgent event. Establish comms when convenient.
Important	Comms need to take place within timeframe or else state could deteriorate.
Urgent	Emergency. A critical component has failed. Cannot recover autonomously and intervention is necessary immediately.
No Tone	Beacon mode is not operating.

There has been found to be some long-term drawbacks to this approach. Since one of the primary goals of beacon monitoring was to reduce the data sent to the ground (achieved by eliminating the download of telemetry data) operators lost the ability to gain intuition about the performance and characteristics of the craft and its components and the ability to run the data through simulations [16]. As such to fully benefit from beacon monitoring, the *fast loop* of real-time health assessment is supplemented by a *slow loop* to study the long term behavior of the spacecraft. This *engineering data summarization* is where the spacecraft creates a second set of abstractions about the sensor telemetry which is sent back to ground to provide the missing context for operators.

This duel approach has conceptually much in common with the reflex and healing approach.

5. Extending the Health-check Mechanism

5.1 Pulse Monitoring

A hybrid approach for the autonomic environment [21] is to use the urgency concept of the beacon monitor to turn the heartbeat monitor into a *pulse monitor*—so instead of just checking the presence of a beat, the rate is also measured (Figure 5).

The concept of the pulse monitor is based on extending the HBM construct. The HBM itself provides a means to ensure a vital process may be safeguarded. The lack of a heartbeat will alert the designated remote HBM that the process has died (or indeed the communications themselves have failed). This relative instant alert to the fact a process is no longer functioning enables immediate actions such as restarting the process and as such minimizing disruption.

Figure 5 Pulse Monitoring

Essentially the HBM provides a vital construct, without it the system is relying on another process noticing that the process has died with no guarantee on how much time will have lapsed before this occurs, if at all.

Yet, vital as it is, essentially the HBM only informs if a process is alive or dead (assuming comms are working) – not the processes actual health or state of being. Taking the biological analogy the rate of the heartbeat indicates the current conditions within which the biological 'system' is operating and determines strategies for 'components' within the system.

An important point to note from the HBM, and also from the Beacon Monitor, is the minimization of data sent – essentially only a 'signal' is transmitted. Any move towards sending more information must not compromise this reflex reaction. As such the tone or the beat must contain within it the urgency level.

This effectively provides a reflex reaction within the Autonomic Grid environment and adds a duel approach, sharing responsibility for environment monitoring and indicating increasing urgency levels.

Figure 6 Pulse Monitor's Logical Functionality

5.2 Pulse Monitor Functionality

The required functionality of the pulse monitor (Figure 6) will include the ability to register with another monitor. The Globus OGSA approach assumes a client-server relationship. In autonomic computing this relationship may be peer-to-peer as such a form of gossip protocol may be used and each local monitor must have the ability to register and monitor its neighbors.

The pulse monitor's functionality still includes the standard HBM function – the present or absence of a beat indicates the process is still alive. Yet the pulse monitor must also contain the pulse levels and rules concerning the changing between levels, the ability to resolve conflicts in perceived levels to enable the extended health check functionality.

The local internal environment must be alerted upon receiving a confirmed change in levels from the external environment while also being able to recommend a change due to changing circumstances it has detected

through its own self-monitoring that will have an effect on the external environment.

5.3 Towards Affect and Emotion

Looking further into the future, the Grid architecture would be more autonomous if only it had 'feelings' or emotions, a trait strongly suggested by some psychologists and AI researchers that is essential for intelligent behavior. This desirable feature is described in biology in organisms as the system of affect and emotion which evaluates the environment with respect to how it would affect the organism/system under consideration. It usually results in some sort of overall *feeling* of whether the situation is positive/negative, safe/dangerous etc.

In establishing the Autonomic Grid, it is important the system of affect and emotion be taken into consideration and how such a feature may be modeled and incorporated. Some preliminary work has been carried out by Norman et al. [22] who introduce a system of affect as running in parallel with a system of

cognition. Cognition is defined by [22] as the system which *makes sense of the world* by understanding and reflecting on things in the world. Affect, on the other hand, is defined to be more focused on *evaluation* of the events in the world and their value with respect to the organism/system. In other words the affect system gives a warning about possible danger and may lead to the body becoming more alert and ready for danger while potentially changing the cognitive strategy.

Figure 7 Affect & Emotion determining the Cognition Strategy influenced by the Pulse Monitor

The three levels of Norman et al's model are named as reaction (lowest level), routine and reflection (adapted in Figure 7) [22]:

1. Reaction – lowest level where no learning occurs but immediate response to state information coming from sensory systems.
2. Routine – middle level where largely routinized evaluation and planning behaviors take place. It receives input from sensors as well as from the reaction level and reflection level. This levels assessment results in three dimension affect and emotion values: positive affect, negative affect and (energetic) arousal.
3. Reflection – top level receives no sensory input or has no motor output, it receives input from below. Reflection, a meta-process, where the mind deliberates about itself. Essentially operations at this level look at the systems representations of its experiences, its current behavior, its current environment etc.

The adoption of a pulse monitor by extending the Globus Heartbeat monitor offers the potential to begin

to build affect and emotion into the autonomic grid architecture.

6. Conclusion

Autonomic computing is gaining ground as a viable holistic approach to computer system development that aims to bring a new level of automation and dependability to systems through self-healing, self-optimizing, self-configuring and self-protection functions.

Grid Computing promises to change the structure of electronic service provision and create a new grid service economy. For this vision to be realized will require an effective Autonomic GRID.

This paper has discussed the general notion of a pulse monitor to provide a reflex means of observing the 'health' of each autonomic grid element. This concept extends the general heart beat monitoring construct that only indicates if a process is alive or not instead of its general health.

This concept could be further used to introduce affect and emotion into the grid architecture. A demonstration system to illustrate and further refine this mechanism is currently under development.

Acknowledgements

This work was undertaken through the Centre for Software Process Technologies, which is supported by the EU Programme for Peace and Reconciliation in Northern Ireland and the Border Region of Ireland (PEACE II).

References

- [1] P. Horn, "Autonomic computing: IBM perspective on the state of information technology", IBM T.J. Watson Labs, NY, 15th October 2001. Presented at AGENDA 2001, Scottsdale, AR. (available <http://www.research.ibm.com/autonomic/>), 2001
- [2] IJCAI Workshop, "AI and Autonomic Computing: Developing a Research Agenda for Self-Managing Computer Systems, Acapulco, Mexico, August 10, 2003, <http://www.research.ibm.com/ACworkshop>
- [3] A. Avizienis, J.-C. Laprie, B. Randell, "Fundamental Concepts of Dependability", UCLA CSD Report #010028, 2000.
- [4] B. Randell, "Turing Memorial Lecture – Facing Up to Faults", Comp. J. 43(2), pp 95-106, 2000.

- [5] R Sterritt, DW Bustard, "Autonomic Computing-a Means of Achieving Dependability?", Proceedings of IEEE International Conference on the Engineering of Computer Based Systems (ECBS'03), Huntsville, Alabama, USA, April 7-11 2003, pp 247-251.
- [6] D. De Roure, N. Jennings, N. Shadbolt, "A Future e-Science Infrastructure" aka Research Agenda for the Semantic Grid, EPSRC/DTI Core e-Science Programme, December 2001.
- [7] I. Foster, C. Kesselman, J.M. Nick, S. Tuecke, "The Physiology of the Grid – An Open Grid Service Architecture for Distributed Systems Integration", <http://www.globus.org/research/papers/ogsa.pdf> June 2002.
- [8] M. Atkinson, J. Crowcroft, C. Goble, J. Gurd, T. Rodden, N. Shadbolt, M. Sloman, I. Sommerville, T. Storey, "Computer Challenges to emerge from e-Science" aka UK e-Science Research Agenda, National e-Science Centre : http://umbriel.dcs.gla.ac.uk/nesc/general/news/uk_e-Science_agenda.html
- [9] IBM, "An architectural blueprint for autonomic computing", April 2003.
- [10] A.G. Ganek and T. A. Corbi, The dawning of the autonomic computing era, IBM Systems Journal, Vol 42, No 1, 2003, pp. 5-18
- [11] Autonomic Computing Concepts, IBM White Paper, 2001.
- [12] D.F. Bantz, C. Bisdikian, D. Challener, J.P. Karidis, S. Mastrianni, A. Mohindra, D.G. Shea, M. Vanover, Autonomic personal computing, IBM Systems Journal, Vol 42, No 1, 2003, pp. 165-176
- [13] T. Bapty, S. Neema, S. Nordstorm, S. Shetty, D. Vashishtha, J. Overdorf, P. Sheldon, "Modeling and Generation Tools for Large-Scale, Real-Time Embedded Systems", 10th IEEE International Conference and Workshop on the Engineering of Computer-Based Systems, Huntsville, Alabama, USA, 7-10th April 2003, pp11-16.
- [14] The Globus Heartbeat Monitor Specification v1.0, http://www-fp.globus.org/hbm/heartbeat_spec.html
- [15] P.F. Stelling, I. Foster, C. Kesselman, "Fault Detection Services for Wide Area Distributed Computations", Retreat 1998.
- [16] M.A. Swartwout, "Engineering Data Summaries for Space Missions" SSDL, 1998.
- [17] J. Wyatt, R. Sherwood, M. Sue, J. Szijarto, "Flight Validation of On-Demand Operations: The Deep Space One Beacon Monitor Operations Experiment", 5th International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS '99), ESTEC, Noordwijk, The Netherlands, 1-3 June 1999
- [18] N. Muscettola, P. P. Nayak, B. Pell, and B. Williams, "Remote Agent: To Boldly Go Where No AI System Has Gone Before".. Artificial Intelligence 103(1-2):5-48, 1998.
- [19] J. Wyatt, H. Hotz, R. Sherwood, J. Szijarto, M. Sue, "Beacon Monitor Operations on the Deep Space One Mission", 5th Int. Sym. AI, Robotics and Automation in Space, Tokyo, Japan, 1998
- [20] R. Sherwood, J. Wyatt, H. Hotz, A. Schlusmeyer, M. Sue, "Lessons Learned During Implementation and Early Operations of the DS1 Beacon Monitor Experiment," Third International Symposium on Reducing the Cost of Ground Systems and Spacecraft Operations, Tainan, Taiwan, 1999.
- [21] R. Sterritt, "Towards Autonomic Computing: Effective Event Management", Proceedings of the 27th Annual IEEE/NASA Software Engineering Workshop, Greenbelt, MD, Dec. 2002, pp 40-47.
- [22] D. A. Norman, A. Ortony, and D. M. Russell, Affect and machine design: Lessons for the development of autonomous machines, IBM Systems Journal, Vol 42, No 1, 2003, pp. 38-44.
- [23] R. Sterritt, D.W. Bustard, "Towards an Autonomic Computing Environment", 1st International Workshop on Autonomic Computing Systems at 14th International Conference on Database and Expert Systems Applications (DEXA'2003). Prague, Czech Republic Sept. 1-5, 2003.
- [24] Y. Tohma, "Fault Tolerance in Autonomic Computing Environment", Proceedings of the 2002 Pacific Rim International Symposium on Dependable Computing (PRDC'02), IEEE Computer Society, 2002.
- [25] J. Ledlie, J. Shneidman, M. Seltzer, J. Huth, "Scooped, Again" in Proceedings of IPTPS 2003 conference. Berkeley, CA. February 2003.
- [26] O. Malik, "Ian Foster = Grid Computing", Grid Today, Oct. 2002.
- [27] F.J. Corbato, V.A. Vyssotsky, "Introduction and Overview of Multics System", Proceedings of AFIPS FJCC, 1965.
- [28] G. Deen, T. Lehman, J. Kaufman, "The Almaden OptimalGrid Project", Proceedings of the Autonomic Computing Workshop, 5th Int. Workshop on Active Middleware Services (AMS 2003), Seattle, WA, pp 14-21, June 2003.