

The Role of Standardization in Future Autonomic Communication Systems

John Strassner¹, Joel Fleck², David Lewis³, Manish Parashar⁴, Willie Donnelly⁵

¹Motorola Labs, john.strassner@motorola.com

²Hewlett-Packard, joel.fleck@hp.com

³Centre for Telecommunication Value Chain Research, Trinity College Dublin,
dave.lewis@cs.tcd.ie

⁴Rutgers University, parashar@caip.rutgers.edu

⁵Waterford Institute of Technology, wdonnelly@tssg.org

Abstract

This paper examines the potential problems of interoperability in the development and widespread future deployment of autonomic communication systems. It draws on experiences from previous standardization efforts in the communications domain and outlines the new challenges faced in building industrial strength networks that are self-configuring, self-healing, self-optimizing, self-protecting, and most importantly, self-governing. It then outlines the work-program being undertaken by the Autonomic Communication Forum and introduces the other papers in this session that present details of these challenges and directions being pursued.

1. Introduction

The communications industry, as it strives to reduce capital and operational costs, has enthusiastically embraced the concept of autonomic systems as articulated in autonomic computing [1][2][3]. This domain has applied a wide range of intelligent techniques to raising the level of automation in network operations, which provides a strong baseline for implementing the various features initially described in the vision of autonomic computing [1]. However, autonomic communication solutions will also have to work within the highly interconnected and complex environment of the communication industry. Not only will autonomic solutions need to interoperate with each other but also with legacy systems not designed or built using autonomic principles. They must be viable in an industry characterized by multiple equipment vendors, multiple communication technologies, multiple communication services and multiple service providers.

While significant progress has been made in specific autonomic solutions, there is a lack of coherence in this work. There is also a lack of rigor in the definition, scope and potential implementation of autonomic communications, computing, and

management. This has not been helped by the tendency to apply the word “autonomic” to every aspect of communications research. For example, the number of conferences that feature autonomies (and topics that are autonomically-inspired, such as self-* functionality) has grown dramatically year after year, in spite of the lack of any supporting standards activity. This is exacerbating the existing schism between industry and academia in the manageability area, since in order to provide self-* functionality, systems must be able to manage their health, configuration, and other details with minimal outside help [4]. Further, there are divergent views on the degree of autonomy and self-management which qualifies a communication system as autonomic. This is often influenced by the balance of self-organization and centralized control and management inherent in different communication domains (e.g., Internet and wireless ad hoc networks vs. connection oriented telecommunication networks).

As another simple example, consider the Simple Network Management Protocol (SNMP). It is well known that operators do not use SNMP to configure devices, and many don't use it for monitoring either. Yet, conferences worldwide discussing next generation network environments continue to talk about SNMP as a way to manage devices, because the mindset of industry is oriented towards building Management Information Bases (MIBs), and because neither industry nor academia has come up with a better solution. If this continues, then there is a danger of repeating the past problems that we have all experienced in network management by thinking in and building stovepipe technologies and, consequently, stovepipe standards [3] [5] [6].

2. Interoperability Issues for Autonomic Systems

In order to assess the interoperability requirements which arise when building autonomic communication systems, we must first identify the boundaries of such

systems. These are likely to form the most profitable areas for standardization work in the support of interoperable autonomic systems.

Essentially, an autonomic system manages resources in order to provide a set of services of value to some users while meeting the operational and business goals of those responsible for the resources and the provision of the services. The fundamental management element of an autonomic computing architecture is a control loop [1][2][3][4][6]. The elements of this control loop are responsible for monitoring the managed elements and other relevant data about the managed elements and the environment in which they are operating, analyze those data, and take action if the *state* of a managed entity and/or system is changed to an undesirable state. Note that undesirable means “non-optimal” as well as “failed”. For example, an action might not conform to expected guidelines as well as it could, or other services might suffer a small (but acceptable) degradation as a result of repairing another service. The FOCAL architecture has been designed to implement just such a control loop. A simplified version of the FOCAL control loop shown in *Figure 1*.




Figure 1. Simplified FOCAL Autonomic Control Loop

In *Figure 1*, sensed data is first normalized. This is required because of the large variety of different types of management data that are required to provide an end-to-end service. The trend of converged networks is now exacerbated by the converging of different types of converged networks, as shown in *Figure 2*. This trend is required to provide seamless services [7] and pervasive presence to their end-users. Hence, different types of management data used for wired and wireless systems must be able to be harmonized to enable management systems to govern such systems.

After the sensed data is normalized, the current state of the Managed Element is compared to its desired state (which is pre-defined based on business goals [8][9]). If the states match, the process continues using the upper “maintenance” loop. If the states do not match, then the Managed Element needs to be reconfigured using the lower “adjustment” loop. The reconfiguration process uses dynamic code generation based on models and ontologies [2][3][4][6][10]. The

use of multiple control loops provides better and more flexible management.




Figure 2. Convergence of Converged Networks

Such an autonomic function typically requires a degree of machine reasoning and learning that adapts to changes in service usage, in the state of resources and in the operational goals as well as in the state of the operational environment – often referred to as context. This is shown in *figure 3*.

2.1 Points of Interoperability

The above discussion has identified several important points of interoperability: description and functionality of resources, services, context, goals, and the ability to dynamically adapt to changing business requirements, user needs, and environmental conditions.

Resources provide the underlying infrastructure and support for services consumed. Resource models must be explicit and extensible to support interoperability when more than one system (autonomic or not) is involved in managing a particular resource type. Examples are network capacity, computing services or access to data.

Services are the interfaces via which users of the autonomic systems derive value. In a composite autonomic system, these may represent end user services, or services provided to other parts of the operational support framework (which may or may not be autonomic systems themselves). Consistent with contemporary approaches to service oriented architecture, well defined service interfaces play an important role in building composite systems, where elements are provided from different software vendors and operated by different service providers.

Common Management Lingua Franca is a derived need of resources and services. One of the difficulties in current management systems is dealing with the complexity arising from technology-specific

and network-centric approaches that do not take business needs into account. For example, SNMP and CLI are currently unable to express business rules, policies and processes, which make it impossible to use these technologies to directly change the configuration of network elements in response to new or altered business requirements [3]. This disconnects the main stakeholders in the system (e.g., business analysts, who determine how the business is run, from network technicians, who implement network services on behalf of the business). Common management information, defined in a *standard representation*, is not available. There will most likely never be a single unified information model (just as there will never be one single programming language), but there *must be an extensible modeling basis for integrating these diverse data and knowledge*. This is the focus of the Modeling Expert Group of the ACF.

Context can be viewed in several different ways. Context can *restrict* functionality, as well as *mandate* a particular behavior. As shown in figure 3, the FOCALe autonomous architecture first determines context, so that the given context can *select* policies that are *applicable*. Policies then select *roles*, which indirectly determine the functionality of the system. When context *changes*, policies change, which change the roles and ultimately the functionality of the system. This presents a major challenge for interoperability, since the operational environment is likely to change in response to design authorities outside of the autonomous system engineering process. Therefore, a degree of run-time interoperability needs to be supported in order to cope effectively with context changes.




Figure 3. FOCALe Conceptual Autonomous Architecture

Goals represent business and end-user concerns and needs. In any autonomous system, the goals passed down for managing how resources are used to deliver value via services are subject to a set of specialist concerns from financial, marketing, security and data protection, legal/regulatory laws, and of course management decision makers. Thus, the expression of goals needs to be exchanged between these domains of expertise and the business systems they use. This interoperability requirement is exacerbated when

autonomic systems are participating in a value chain, which causes varying levels of decision-making authority to be distributed over multiple organizations.

Autonomic control is provided by a set of control loops, as described above. Currently, most autonomic systems research takes a homogenous approach to implementing this loop, and this is an area of intense innovation in the use of a variety of reasoning and learning techniques. A notable exception is FOCALe [3][5], which prescribes a set of *adaptive* control loops. In FOCALe, the adaptivity is derived from context awareness and guided by machine learning and reasoning. It is therefore likely that the interoperability within the control loop of an autonomic system will not present a stable problem domain for some time to come.

2.2 Interoperability Drivers

The key driver for interoperability is the need to address the total cost of ownership of communication systems. There are many ways in which autonomic systems can lower the total cost of ownership; the most common is the benefit derived from automating previous tasks that required human operators. In this scenario, savings arise both from the quicker execution of the task as well as from the savings of not using personnel. A derived benefit is that automation reduces or eliminates manual errors.

Such systems are made up of many components, in term of different technologies, network elements and network segments, often subject to different procurement cycles. Therefore, any autonomic solution must co-exist with other autonomic systems, possibly from other vendors, as well as be able to work with non-autonomic portions of the network. In practice, therefore, we must treat autonomic systems as *composite systems*.

A significant result of this is that a coherent set of operational goals can then be applied to the composite autonomic system with the aim of ensuring that different autonomic functions operate in harmony and that unwanted and unpredicted behavior does not occur that may require new skilled human intervention to correct. In other words, care must be taken to ensure that the additional system complexity introduced by autonomic solutions does not incur new human operator costs that may negatively impact the potential cost of ownership saving. Further, the additional engineering effort incurred by extra system complexity must also be controlled from the outset of any autonomic solution roll-out plan.

2.3 Existing standardization frameworks

The key challenge for developing interoperability standards for an autonomic system is that *all of the boundary elements*, namely resources, services, context and goals/policies, must be addressed within the same framework if composite autonomic systems are to be developed. It is this integration of interoperability points that is currently poorly addressed by current standards frameworks.

Further, autonomics is driven by the imperative to delegate decision-making from the more costly human sphere to the automated domain of autonomic systems. This will necessarily be a *progressive process* as human decision makers will need to have the confidence that the parameters of the decision making area is understood well enough to safely delegate it to automatons.

In addition, the ever-increasing pace of competition and technological innovation impose rapid changes on the decision-making domain with which managers are concerned. To successfully offer themselves as targets for the delegation of decision-making, autonomic systems must be able to respond rapidly and efficiently in dealing with the changing domain models for resources, services, context and goals. This implies that interoperability standards for autonomic systems must possess a level of usability and extensibility not required of previous generations of standards.

In considering the interoperability standards in terms of the boundary categories identified above, there would seem to be a large body of existing available agreement upon which the standardization of autonomic systems may draw.

Resource models for management purposes have been extensively standardized in the form of management information bases (MIBs) that conform to the manager-agent paradigm. This has been the accepted approach for network and systems management, with large bodies of MIBs developed for a large proportion of network and system equipment by bodies such as the IETF, the IUT-T and the Distributed Management Task Force (DMTF). Though these models are represented in different languages, the models themselves are mature. However, these models tend to focus on representing resources at the device level, since the primary aim is to support vendor independence in the development of manager-agent systems. Resource models for the type of higher-level management that autonomic systems aspire to have not been the subject of such comprehensive standardization, though some work has been conducted related to telecoms operations in the TeleManagement Forum.

Services are another area where considerable standardization effort has been undertaken. One of the earlier successful attempts was the CORBA set of standards of the OMG, but while they work well within enterprises, they were not well suited to deployment over the Internet in support of value chains. This role is now served by the adoption of web service technology using XML over HTTP. Through some variations exist, there seems to be some convergence onto the simple service format standardized by the WWW Consortium (i.e., the Web Services Description Language). Standardization support for composite web services and service discovery has also been addressed in various forms (e.g., the Universal Description, Discovery and Integration (UDDI) and Web Services Business Process Execution Language (BPEL) at OASIS: www.oasis-open.org).

The standardization of context is far less mature in comparison. In many application domains it is treated implicitly, but the need for explicit standards for context has been raised in recent years in the study of ubiquitous and pervasive computing. However, while the commercial market for pervasive computing remains immature and fragmented, there has been a strong drive from industry to develop a broad context interoperability standard. In the network and systems management domains, the models developed for monitoring, and the sophisticated event notification support these offer, may satisfy the needs of most practical context aware system in the near term.

The standardization of goal modeling again has not been subject to a broad industry move to standardization, most likely because the market in higher level management systems is more fragmented than the equipment market. There is, however, broad consensus that the use of policy rules is an appropriate way of capturing goals. Policy rules are already used extensively to provide improved flexibility in node management and access control, and a number of languages have been developed. Some have been standardized for integration with resources (e.g., in the IETF and DMTF) and services, e.g. the eXtensible Access Control Markup Language (XACML) OASIS. In particular, the translation of goals into changes to be made to the system has been a stumbling block for current management systems. The autonomic communications community is an important source for novel approaches based on the Policy Continuum [3][5] and model driven generation of code [2][3][5][9][11].

What is apparent from examining the range of standardization activities that have been undertaken that none integrate resources, services, context and goals/policies into a single framework suitable for developing composite autonomic systems. Though

there is some work on integrating pairs of these elements (e.g., services and resources in the Web Service Resource Framework at OASIS and the Open Grid Forum, policy and resources in the Common Information Model with the TMF, IETF and DMTF), none comes near to forming the full interoperability framework needed for autonomic systems.

Further, these existing standards do not support the form of rapid extensibility required for autonomic systems, nor do any have a detailed specification of behavior orchestration. The key to this may be in the management of explicit metadata, an area where rapid standardization is occurring in the Semantic Web domain. The ontological modeling that characterizes this area provides machine reasoning capabilities that are potentially important in realizing an autonomic control loop. More significant for interoperability, however, is the ability to annotate models with axiomatic metadata, such as constraints. This means that the assumption behind a particular standardized model can be embedded in a model in a way that can be automatically verified as attempts are made to extend the model. This provides the ability to more freely extend standardized models without unknowingly violating the specific conditions upon which the original interoperability agreement was formed [12].

Finally, any standardization effort cannot be driven forward solely in abstract models. Specific application areas must be pursued to motivate and verify the generated specifications. The convergence of communication and IT technologies offers many opportunities for new business models where lower capital costs will need to be matched by the lower operational costs that autonomic systems promise. They also provide opportunities to consider the interoperability between autonomic computing and autonomic communication solutions. Specifically, Next Generation Networks (NGN) based on a common IP infrastructure will become increasingly open, commonly shared and reliant on highly distributed components. Motivated by the potential to deliver valuable service mobility, personalization, transparency and immediacy, this creates attractive business opportunities to network operators and service providers. However, it also poses significant new challenges in many areas of communications and services management, especially when it comes to engineering solutions able to autonomously interoperate within and across distinct organizational domains.

3. Requirements in Organizing Standardization

We assert that by defining a unified set of Autonomic Standards, we can avoid both the divergence of different efforts within the autonomic community as well as help solve current and future management problems. By unified, we mean a set of standards that work together to apply autonomic principles and mechanisms to the entire life cycle of governing functionality provided by components, devices, software, and systems. However, standards are only beneficial if they are used and implemented. In order to achieve this, we need to energize and gain the support of academic and industry thought leaders. Indeed, the downfall of many standards is this lack of coordination and cooperation between academia and industry.

Therefore, any standardization activity in the area of autonomic communication must be led by industry and supported by academia, and should focus on mechanisms required to ensure interoperability (through a compliance certification suite). The coherence provided by standards will therefore be supported through research and development. Research, led by academia and supported by industry, provides coherence by aligning current and future research projects to the vision of interoperable autonomic systems as realized by such standards activities. Joint development between industry and academia will cement the cooperation between industry and academia, and can provide additional coherence through coordination of projects funded by organizations such as the EU and the NSF.

Working on standards for the sake of standards, while noble, is hard to justify for certain organizations. Furthermore, standards are meaningless unless they are tested, proven to work, and adopted. The obvious link is, of course, autonomic projects and programs under development in the EU (e.g., FP7) and the US (e.g., various NSF programs, such as FIND and GENI). All such programs require standards to ensure interoperability among the program participants. Developing a unified set of Autonomic Standards requires ensuring interoperability for programs that seek to use autonomics. A standards defining organization (SDO) is not a funding body, nor can it be solely created to pursue funding; however, an SDO "approved" stamp for proposals generated will differentiate such proposals from others that lack an internationally coordinated vision.

In order to support joint development, an SDO must have an open IPR policy. This policy will apply to all artifacts produced by the SDO, including standards,

software and documentation. Examples of such artifacts include information and data models, languages, protocols, and code. The JCP (Java Community Process) and/or the Eclipse community project provide sound, verified models for such IPR arrangements. Key elements of such IPR process will include specifications, reference implementations, and technology compatibility kits (i.e., a suite of tools, tests, and documentation used to test implementations for compliance).

While the work of the SDO should be open (i.e. with minimal IPR constraints, both for specifications as well as for code). The motivating goal of the IPR policy is therefore to strongly encourage the submission of enhancements and changes back to the SDO to foster and promote a common strong foundation and framework for all autonomic management. The open nature of such an IPR policy, coupled with developing software under an appropriate open source license, enables academia to freely participate and shape the future of Autonomics; at the same time, it ensures that industry participates in a vendor-neutral way. It also encourages individuals and organizations that are not officially part of the SDO to look at and use the work produced by the SDO, and hopefully join the organization.

4. Autonomic Communication Forum

The Autonomic Communication Forum (ACF) was established at the end of 2004 following an initiative by the EU-funded Autonomic Communication Accompanying Action project. Global interest quickly grew from both industry and academia through the Workshop on Autonomic Communication (now the Autonomic Networking conference) and further project funding from the EU.

At the Autonomic Networking conference in Paris in 2006, John Strassner of Motorola Labs, US, observed the lack of coherence in the wide range of research and development activities being conducted internationally, and proposed that the ACF should undertake the following:

1. Unify current thinking in autonomics by creating a new set of ACF sanctioned Autonomic Standards, focusing firstly on the management of systems, and secondly on computing and communications using autonomic mechanisms.
2. Building on the above, define an autonomic reference framework as well as a set of baseline compliance statements to guarantee interoperability.

3. Create an organizational structure that will empower academia and industry to work together in developing and maintaining the above goal.

The community responded enthusiastically to this call for action and has been building a new organizational structure for the ACF. This is driven by a Board of Directors (BoD), chaired by John Strassner, which consists of an Architecture Committee (co chairs Joel Fleck of Hewlett-Packard Office of Strategy and Technology and David Lewis of Trinity College Dublin), an EU-US Liaison Committee (co chairs Manish Parashar, Rutgers and Mikhail Smirnov, FOKUS) and an SDO Liaison Committee (Monique Callisti, Whitestein and Roy Sterritt, University of Ulster). The BoD has been guiding the formation of a number of chartered technical groupings. A key feature of all technical groups is that they possess both an industrial and academic co-chair, in order to maximize the flow of scientific results into usable autonomic standards.

The chartering of groups is overseen by the Architecture Committee and to date, the following technical groups have been formed:

- Architecture Expert Group
- Policy Expert Group
- Modeling Expert Group
- Semantics Working Group
- Service Composibility Management Working Group
- Assessment Working Group

The following are the groups currently active in the ACF:

Architecture Expert Group

The long-term goal of the Architecture Expert Group (AEG) is to develop a framework of methodologies, models, documents, implementations, and testbeds that collectively enable interested parties to develop solutions for Autonomic Communications. Its co-chairs are:

- Industrial Co-Chair – Dave Raymer, Motorola Labs
- Academic Co-Chair – Sven van der Meer, Waterford Institute of Technology

Policy Expert Group

The Policy Expert Group will enable autonomic communications by facilitating policy-based governance of communications systems. This will be accomplished primarily by borrowing, enhancing and extending existing information models (e.g., DEN-ng)

to represent more advanced policy features and capabilities. Its co-chairs are:

- Industrial Co-Chair – Greg Cox, Motorola Labs
- Academic Co-Chair – Joan Serrat, Universitat Politècnica de Catalunya

Modeling Expert Group

The goal of the Modeling Expert Group is to develop a networking “lingua franca” that enables vendor-specific management data and languages to be transformed into a common form. This will be accomplished primarily through enhancing the DEN-ng information model and developing semantic mappings from it. Its co-chairs are:

- Industrial Co-Chair – John Strassner, Motorola Labs
- Academic Co-Chair – Betty Cheng, Michigan State University

Semantics Working Group

The goal of the Semantics Working Group is to provide a locus for collecting and disseminating mathematical and other formal tools and techniques influencing the development of autonomic communications. Its co-chairs are:

- Industrial Co-Chair – in process
- Academic Co-Chair – Simon Dobson, University College of Dublin

Service Composibility Management Working Group

The goal of the Services Composibility Management Working Group is to define a Services Composition Management Framework to allow the dynamic composition of services in an NGN environment. Its co-chairs are:

- Industrial Co-Chair – Jose A. Lozano, Telefonica I+D
- Academic Co-Chair – Thomas Magedanz, Fraunhofer FOKUS

Assessment Working Group

The goal of the Assessment WG is to provide a comprehensive catalog of approaches, supporting the production as well as the operational lifecycle in all phases, including design paradigms, engineering methodologies, verification, validation, and testing methods, and performance and quality indicators. Its co-chairs are:

- Industrial Co-Chair – Antonio Manzalini, Telecom Italia
- Academic Co-Chair – Peter H. Deussen, Fraunhofer FOKUS

In addition, there are a number of expert and working groups in formation. Three of these that are close to being approved by the Board are the Security Expert Group, the Autonomic Communications Working Group, and the Peer-to-Peer Working Group.

5. Related Activities

There are several activities in the US and in other countries whose mission overall and/or complement the objectives of the ACF. These include large research and infrastructure efforts such as GENI, Cyberinfrastructure, TeraGrid in the US.

For example, GENI (www.geni.net) is a industry/university “clean slate” effort aimed at re-conceptualizing of the Internet architecture. Specifically, GENI will be an experimental platform that will be designed to support both research and deployment, effectively filling the gap between small-scale experiments in the lab, and mature technology that is ready for commercial deployment. At the physical level, GENI will consist of a collection of physical networking components, including links, forwarders, storage, processor clusters, and wireless subnets. These resources are collectively called the GENI substrate. On top of this substrate, a software management framework will overlay network experiments on the substrate, where each experiment is said to run in a slice of the substrate. Clearly, autonomics is extremely relevant to all aspects of this effort and ACF as already started interacting with the GENI community to establish synergies.

The NSF FIND (Future Internet Network Design, see <http://find.isi.edu>) is a new long-term initiative of the NSF NETS (Networking Technology and Systems, see http://www.nsf.gov/funding/pgm_summ.jsp?pims_id=12765&org=CNS) research program. FIND is another type of clean slate technology that is more futuristic (e.g., 15 years out) than GENI. However, FIND addresses broader issues, such as how to elicit trust from society in managing network services, and how to support ubiquitous and pervasive computing, regardless of whether the device connecting to the Internet is a simple mobile or a powerful supercomputer. In particular, FIND proposes a holistic use of multiple different approaches and technologies to ensue that the Internet grows commensurate with business needs from various constituencies.

In the EU, a cluster of projects has been funded around the topic of situated and autonomic computing. More broadly, the Open Grid Forum (OGF, formerly GGF – see www.ogf.org) is focused on accelerating grid adoption to enable business value and scientific

discovery by providing an open forum for grid innovation and developing open standards for grid software interoperability. It is community of users, developers, and vendors leading the global standardization effort for grid computing. Once again, autonomies is of great relevance to all aspects of Grid computing and a synergistic relationship between ACF and OGF can be of significant mutual benefit. It is the objective of the ACF to explore and foster collaboration with these and other such activities.

Another interesting EU program is the Situated and Autonomic Communications (referred to as SAC – see <http://cordis.europa.eu/ist/fet/comms.htm>) initiative. Its objective is to promote research in the area of new paradigms for communications and networking systems that react locally to context changes). SAC envisions such programs as autonomously controlled, self-organizing, massively distributed, and technology independent. This would enable communications to become task- and knowledge-driven and fully scalable.

The main objectives are to (1) define a self-organizing communication network concept and technology that can be situated (i.e., react locally) in multiple and dynamic contexts, and (2) to study how various social, business and other needs impact future communication paradigms. This is similar to the other activities, in that SAC aims to support the evolving needs of society and economy through enabling a service-oriented communication network. SAC uses the EC Integrated Project and Networks of Excellence (NoEs) to realize SAC goals. IPs are expected to address both objectives in an integrated way, while NoEs are expected to promote and federate European basic research in networking.

6. Summary

The need to control and reduce operational costs in the face of increasing systems connectedness and complexity provides a strong motivation for the development of autonomic systems. The focus on convergence of traditional network and IT services means that interoperability must encompass both autonomic computing and autonomic communications solutions. However, any positive impact on the total cost of ownership from introducing autonomic systems will be quickly eroded if the implementations resulting from autonomic architectures do not provide agile solutions that adapt readily to new stimuli, if the resulting implementations are not interoperable, and if the resulting implementations preclude the extended use of embedded legacy systems. Further, since management systems are currently undergoing significant re-engineering, extensible approaches for

ensuring interoperability must be established quickly if the development of autonomic systems is to avoid the high integration costs that “silo management” implementations have for years imposed on the telecommunications industry.

The Autonomic Communication Forum has therefore undertaken a program of work that brings together leading R&D experts from both academia and industry to start proposing open solutions to autonomic system interoperability. This initiative will ensure the problem of interoperability is addressed early enough for open solutions to become embedded into the future engineering culture of autonomic systems.

The ACF is in the process of formalizing its membership, IPR and operational procedures and welcomes all parties interested in influencing the direction of the management industry by becoming active contributors.

This session presents a detailed snapshot of some of the work being carried out in the forum currently. In [15] the crucial work of the ACF Architecture WG toward a common technology-neutral system architecture for autonomic systems is introduced. In [13] the related policy-based management model being examined by the Policy EG for such an architecture is described. In [14], the problem being addressed by the Service Composability Management WG as it examines how autonomic systems can be composed in NGN deployments.

Acknowledgements

This work is partially sponsored by Science Foundation Ireland under grant number 03/CE3/I405.

7. References

- [1] IBM, “An Architectural Blueprint for Autonomic Computing”, v7, June 2005
- [2] S. Dobson et al., “A Survey of Autonomic Communications”, ACM Transactions on Autonomic Adaptive Systems, Vol. 1, No. 2, Dec. 2006, pp. 223-259.
- [3] J. Strassner, “Autonomic Networking – Theory and Practice” [tutorial], Proc. 2008 IEEE/IFIP Network Operations and Management Symp. (NOMS 2008), IEEE, April 2008.
- [4] B. Jennings, S. van der Meer, S. Balasubramaniam, D. Botvich, M. Ó Foghlú, W. Donnelly, and J. Strassner: “Towards Autonomic Management of Communications Networks.” IEEE Communications Magazine, Vol. 45, Issue 10, October, 2007, pp 112-121.
- [5] J. Strassner, “Policy Based Network Management”, Morgan Kaufman Publishers, Sept. 2003, ISBN 1-55860-859-1

- [6] J. Strassner, N. Agoulmine, E. Lehtihet, "FOCALE - A Novel Autonomic Networking Architecture". The Latin American Autonomic Computing Symposium (LAACS 2006) Campo Grande, MS, Brazil, July 18-19, 2006.
- [7] <http://www.motorola.com/content.jsp?globalObjectId=6611-9309>
- [8] J. Strassner, "A New Paradigm for Network Management – Business Driven Device Management". In: SSGRRs 2002 conference (summer session)
- [9] J. Strassner, D. Raymer, E. Lehtihet, S. van der Meer, "End-to-end Model-Driven Policy Based Network Management", In: Policy 2006 Conference
- [10] J. Strassner and D. Raymer, "Implementing Next Generation Services Using Policy-Based Management and Autonomic Computing Principles", NOMS 2006
- [11] <http://www.omg.org/mda>
- [12] Lewis, D., O'Sullivan, D., Feeney, K., Power, R., Ontology-based Engineering for Self-Managing Communications, 1st IEEE International Workshop on Modeling Autonomic Communications Environments, Dublin, Ireland, 25-26 Oct 2006
- [13] Cox, G., Serrat, J., Strassner, J., de Souza, J.N., Raymer, D., Samudrala, S., Davy, S., Barrett, K., An Enhanced Policy Model to Enable Autonomic Communications, EASe 2008, Belfast, Northern Ireland
- [14] Magedanz, T., Lozano, J.A., Schreiner, F., Gouveia, F., González, J.M., Towards Autonomic Communication Mechanisms for Service Composability Management
- [15] Raymer, D., van der Meer, S., From Autonomic Computing to Autonomic Networking: an Architectural Perspective, EASe 2008, Belfast, Northern Ireland