Autonomy

At IETF 90 in Toronto, there was a well-attended Birds of a Feather meeting entitled "Use Cases for Autonomic Networking (UCAN)." The associated mailing list is called "Anima," standing for "Autonomic Networking Integrated Model and Approach." So what exactly is autonomic networking?

The dictionary defines "autonomic" recursively as "relating to, affecting, or controlled by the autonomic nervous system," which doesn't help much. The autonomic nervous system is an important aspect of an animal's body that takes care of vital functions such as breathing and swallowing without conscious control. This inspired the concept of autonomic computing introduced by IBM in 2001, with the intention of making computing systems self-managing to the extent possible. Autonomic networking, which has been an active area of research in recent years, refers to the application of such ideas to networks. One forum for these discussions has been the Network Management Research Group in the IRTF.

What exactly is autonomic networking (AN)? One way to look at it is "plug and play for the ISP" or "plug and play for the enterprise network." This is a step forward from the original concept of plug and play for home networks, which has long been recognised as a vital requirement (see, for example, the work of the IETF Homenet working group).

The goal of self-management includes self-configuration, self-optimization, self-healing and self-protection. AN puts operational intelligence into algorithms at the node level, to minimize dependency on human administrators and central management. Nodes participating in AN will discover information about the surrounding network and negotiate parameter settings with their neighbours and other nodes. Ideally, autonomic nodes use stable closed-loop control methods to achieve self-management, instead of using more traditional top-down network configuration and monitoring tools to set and verify their parameters. Nodes may also have learning and cognitive capability, i.e., the ability to self-adapt their decision-making process based on information and knowledge sensed from their environment. In the most sophisticated case, advanced data analytics may be part of the input to the autonomic mechanisms.

Is this more than science fiction? Certainly; many aspects of small networks have been self-configuring for years – otherwise there would be no unmanaged home or small office networks, for example. Numerous existing protocols have a flavour of autonomic properties (for example, the spanning tree algorithm needs no manual configuration in order to operate, and some routing protocols require very little configuration). Recently, prototypes and initial products of explicitly autonomic protocols have emerged from companies such as Huawei and Cisco. However, it is clearly necessary to have some basic standards in place if AN is to become relevant to large multi-vendor networks.

Why is this important in 2014? The main motivation is not new: large network operators, both ISPs and enterprises, are suffering more and more from the problems and difficulties caused by the central configuration of hundreds or thousands of network elements. Now, however, after some years of research and discussion, ideas about how to achieve autonomic networking are becoming concrete. Fortunately, it is now economic to provide enough computing power and memory in network elements to support AN. The time is therefore ripe for a standardisation effort.

What sort of network parameters might be set by autonomic methods? A number of use cases for large networks were proposed in the UCAN BOF at IETF 90: network address and prefix management, optimisation of mobile backhaul links, risk-aware routing, and detection of SLA violations. Certainly many other examples will emerge. Two very fundamental aspects of AN can also

be viewed as uses cases in themselves: securely bootstrapping new devices, and creating a secure autonomic control plane for use by specific AN applications.

How does management keep control? While it is obviously desirable to reduce the need for tedious human interventions, it is essential that network managers can still ensure that the network does what is needed and remains fully secure, even if many nodes are configuring and managing themselves. For this reason, the model for AN must include a mechanism for communicating the *intent* of human managers to all self-managing nodes, for matters such as resource control, service requirements, and security policy. At the same time, in real networks, AN mechanisms will need to co-exist with traditional top-down management and monitoring tools for many years, so it must be possible to introduce AN technology in small steps.

What next? At the time of this writing, a possible IETF working group called Anima (Autonomic Networking Integrated Model and Approach) is under discussion. A complete solution for autonomic networking would be a very ambitious goal. The scope of the proposed effort is much more modest: it is to define a minimum set of specific reusable infrastructure components to support autonomic interactions between devices, and to specify the application of these components to one or two elementary use cases of general value. The main goal is therefore to develop common infrastructure components for distributed functions. The infrastructure should be capable of providing the following services to those distributed functions:

- a common way to identify nodes
- a common security model
- a discovery mechanism
- a negotiation mechanism to enable closed-loop interaction
- a secure and logically separated communications channel
- a consistent autonomic management model

Some important topics are intentionally not included in these initial goals as they are considered separate matters that should be considered later:

- Mechanism for distributing policy intent to autonomic nodes
- Use of data analytics by autonomic nodes
- Other external information sources
- System-wide integration of autonomics

Further reading: draft-irtf-nmrg-autonomic-network-definitions and draft-irtf-nmrg-an-gap-analysis.
Mailing list: anima@ietf.org