Pipeline
Technology for Service Providers.
www.pipelinepub.com Volume 13, Issue 8

Organic Networks: self-managing and self-growing networks

By: Scott St. John

Communications networks never stop growing. Since their inception, our global networks have always been a patchwork of old and current technologies. In an ongoing process, network practitioners stretch their minds to invent fresh ways of doing the same thing — essentially carrying bits to more places — faster, more reliably, more securely and at lower cost. But larger size, the enrichment of services, and escalating numbers of endpoints bring increasing complexity. Network capacity is expected to go on growing non-linearly and so as the number of possible interactions between elements increases non-linearly, complexity will increase at a multiple of those exponents.

Global telecommunications networks will eventually become too big and complex for humans to manage. We must automate more network functionality, and plan for how designers, operators, and intelligent machines will work together. In recent years, a lot of deep thought has been taking place about the future of networks: autonomic networks, intelligent networks, software-defined networks, network function virtualization, and organic networks. Clearly we are looking hard for a solution.

We see increasing speculation and serious research in the practical application of artificial intelligence, machine intelligence, and machine learning across the board: transport, health care, war, manufacturing and network security. Algorithms that allow machines to learn are being coupled to deep data stores fed by sensors that provide situational awareness. These learning programs create a model that can control tools that allow machines to direct actions. When placed in a feedback loop, this is the basic biological model: senses, memory, thought, and action. These machines will use their learning as responsibly as we teach them and as realistically as the data captures reality.

Sooner or later, probably sooner, our networks will become too big and too complex for individual humans to design, build, and manage. Daily existence in this complex age makes us more aware of the limitations of human beings. Daniel Khaneman and Amos Tversky's groundbreaking work on the limits of human expertise concludes that the "intuition" of some experts is illusory [Kahneman, D. (2011) *Thinking, Fast and Slow*.] Human cognition is limited in its ability to understand complexity, so humans have workarounds that allow us to cope and make decisions in a complex world, and those workarounds work, up to a point. Over much of our careers, these workarounds took the form of creating and using well-defined processes. But these very processes always include great friction generated by the humans in the process flows.

Machines will have cognitive limits, but they are limits that are different in degree and type from human limits. Handing things over to machines starts to look like a good idea. But humans have to design and build the machines in the first place, then live with them. If frail yet smart humans can create machines that are optimized for building, managing and optimizing networks for us humans to use, operators and customers will probably find themselves happy to work with those machines. But overcoming that frailty, that tendency to introduce a systematic error, is a great stumbling block in this highway to autonomic networks.

Humans providing tools and services for humans to use is nothing new. At one point in the history of telecom, it became necessary to replace human switchboard operators with electro-mechanical switches. If we hadn't done that, there would not have been enough eligible humans on the planet to switch the calls. At one time, phone billing systems consisted of lots of file cards, ledgers and people. We replaced human beings with a much smaller number of humans and massive billing systems with access to databases. Today we have computers in every aspect of the business, and fewer people per unit of operation.

We've talked a lot about autonomic networks, but vendors have not built anything beyond real time demonstrations, and the industry has not actually bought these products. However, as an industry, we have inserted a lot of simple proto-autonomic behavior into elements in our networks. Routers handshake, nodes report errors and anomalous behavior, elements identify themselves to network management tools, and those tools diagnose problems, all without much human intervention. The use of pattern recognition analytics from huge data sets is improving all the time, making it possible to predict network behavior and identify anomalies in ways that humans just can't do. Data scientists use this data to find useful correlations that drive greater reliability, increase efficiency, and create more customer satisfaction. So the industry is definitely making progress towards building networks that have many of the attributes of "intelligence". But the path to fully autonomous, self-managing networks is not yet mapped in any detail.

What we can do now

One realistic path is opening up. It will involve another systematic revolution in OSS. Just like with the introduction of NGOSS/Frameworx, a stepwise path exists to achieve this revolutionary change. Again, like NGOSS, it will come by applying a novel overall vision to the rudiments of new tools that are naturally bubbling up in our inventive industry. The principle new seeds are: SDN/NFV unbundling, openAPIs, big data, social network link models, cloud infrastructure, IoT, machine learning, accessible intelligent agents with available open APIs, a marketplace for innovation with apps, intelligent edge-based security appliances, and a vast ICT network that links endpoints to cloud computing enters. For the last decade, we called this evolving new vision **Autonomic Networks**.

Future networks will contain autonomic elements, along with a collection of SDN-like control nodes. Michael Beringer explains that autonomic networking and SDN are complementary concepts. Networks will certainly contain ever more virtualized elements and functions: SDN and NFV being early manifestations of this. The first vendors are already beginning to piece these parts together. MYCOM OSI is linking big data and analytics into a manager of manager model, where the heritage OSS systems both feed data and act as the hands of the MYCOM OSI automation. It also reaches out directly to NFV devices. Elsewhere, gen-E is using social network models for the assessment and presentation of big data analytics to organize consoles and reporting more in line with naturalistic, human reasoning. But these are still piecemeal, somewhat blind, steps.

Briefly, the forthcoming near-term future architecture will be NFV enabled network devices controlled by software-defined networks. Sensors in the network, environment, and customer devices will track network performance and customer behavior — concurrently feeding both nonSQL data stores and intelligent model-driven controllers. Machine learning systems will continuously tap the deep data stores and train the model used by these intelligent controllers. The intelligent controller will reach down to the NFV enabled devices and optimize the network in conjunction with provisioning resources, security, and interconnects in virtualized data centers — placing services into these provisioned resources based on real-time consumer demand.

At first there will be many different and specialized intelligent controllers. They will intercommunicate with openAPIs such as those in development by the TMForum. Each will be an expert in its own domain. In the beginning, these domains will be function based. But we should evolve out of our function-based architectures. Eventually the natural clustering of the data itself will segment the domains of specialization. For example, some might specialize in specific device families and others in consumer service clusters.

We will remain in this distributed controller intelligence world — but with ever greater degrees of integration and coordination. For example, security controllers will start by just focusing on identifying the patterns they associate with abnormal behavior and acting to damp that down. But, over time, these will interlink with network optimization controllers to insure that the actions taken for security are consistent with maintaining overall network stability and service continuity. As these security controllers look for any departure from normal network and service behavior, this discovered information can also feed controlled, adaptive provisioning of network and cloud infrastructure.

The intelligent controllers will draw from all the theoretical work and practical experiments in AI: machine intelligence and machine learning. It now appears these controllers will be specifically trained service identities of the major intelligent agents in development by IBM (Watson), Microsoft (Cognitive Serves), Apple (Siri), Facebook, Amazon (Alexa), and others. Both Microsoft and IBM are in a market race to capture new users and new skills. As such, they have opened up their cognitive intelligence APIs to use by almost all comers. Every successful OSS and BSS company should be putting in place R&D teams to utilize these APIs.

It is likely that the self-organizing, self-managing, self-protecting networks of the future will emerge gradually. ICT companies will add chunks of greater and greater intelligence, and one day Autonomic Networks will be normality and most of us will take it all for granted. But this is still a bit short of what we really mean by "organic". This needs to include self-reproducing, self-building networks.

Where we should go

There is not yet any common understanding of the true nature of the architecture of these networks of the future. We like to call them "organic networks" because they will almost certainly emerge with a range of characteristics that parallel biological systems. A very early notion of "organic networks" was set out by Andrew Lippman and Alex Pentland. That was some time ago, but their concepts of "viral networks" and "influence networks" still make some sense.

Organic networks should have, for all practical purposes, unbounded scalability. We have no idea how big this global network of networks will need to be to meet the demands of human beings. But we must insure that it can survive, however big it gets. As Lippman and Pentlan point out, "… lack of scalability [in today's networks] is perhaps the most stringent restriction on system design, regulatory practice and economic development."

The organic infrastructure will have a high degree of openness to innovation and participation. To meet the scalability ambition, it is clear that the organic global network will, as today, result from a connected set of components provided by a multitude of contributors. Hence, initiatives like the TMForum's openAPI are mainstream enablers on the autonomic roadmap. However, one of the lessons from the Internet and the explosion of the personal computer market is that our traditional idea that 'innovation and growth supplied entirely from defined service providers' is outdated. Users will invent and implement new functionality from the network edge. It is suggested they should be able to schedule specific kinds of capacity to facilitate these edge services. There must be a mechanism for these innovations to flow throughout the network, gaining market base much like YouTube videos or Facebook "stories". Aps and app markets are a good start at this.

The organic network will have focused network functional intelligence. Boundless scalability probably requires some discipline about where the smartness will direct the network. The simpler the functionality in the network, the more likely it is that the network can successfully look after itself. If the simple rules result in survival-oriented behavior. So the intelligence in the network needs to be focused on building and running the network. The smartness that adds value to the communications — services and applications — will likely have to live at the network edge and clouds. Lippman and Pentlan: “Viral communications derives directly from the end-to-end principle on which the Internet is based — the intelligence is in the end nodes,….”

But here, we must augment Lippman and Pentlan. Future networks could be micro-optimized for the user and the service riding the network during the specific-use-transit-time needed by the customer’s use of the service. By communicating the service characteristics and customer wishes into the network optimizing equipment, an edge requestor can propagate wishes into micro decisions on network path characteristics. The network will control the path and aggregation based on the rich bounty of edge-communicated service goals. *Using our biological analogy, the edge will provide the nutrients (services and customers) that the organic network needs to thrive. The network will self-seek these resources and consume them, gaining a sense of vitality by doing so.*

The Internet has taught us another lesson: human beings are not all on the same team. The prevalence of malicious human behavior is now an accepted fact. For those of us who just want to use the network for business and pleasure, we need an organic network that recognizes and autonomously develops immunity to assault and disease. And for those that do not, automatic routing into playpen game environments.

Unbounded scalability does not imply that the network must maintain a static pre-determined architecture and technical foundation. We expect the opposite: organic networks need to be self-renewing and self-healing**.** In networks, as in biological systems, things go wrong. Components fail, accidents happen. Sometimes people manage to do damage, by mistake or deliberately. Extensive virtualization of all network functions will make it possible for the network to reprogram its own capability when needed for repair or renewal.

Lastly, the network must be able to grow. If it can recognize where its resources (services and consumers) are located, it should be able to grow the network to access those resources. At first, this seems impossible. Fiber and network electronics are quite complex devices and require specialized manufacture. But two future industries in their early childhood will allow networks to grow themselves. First is robotics in construction and transport. It will be possible to create machines which order infrastructure and cable from robotic factories, have it transported by autonomous devices, and lay that infrastructure in place by robots. Second is the development of 3D printers. As these gain multi-material functionality and the ability to do complex fabrication, networks can be grown into the places that need them; just as planes, houses and bridges can be printed today.

Autonomous evolution

Perhaps for the first time, our networks will be complex and pervasive enough to undergo some form of unplanned Darwinian-style evolution. This will either replace or be some as yet unknown offshoot of market competition. While that evolution will be unplanned, it should not be unconstrained. The fitness of the evolving network must be measured by its value to humans, not merely by the system's ability to survive and grow. Somehow the direction of evolution needs to result in the organic network being both unselfish and self-controlled. It should put human interests first, follow international law, and accommodate the inadequacies of human behavior. This differentiates organic network evolution from biological evolution, in which the only driver is fitness for survival and procreation. That’s a big responsibility for our ICT industry.

As part of the wider AI community, more than one metaphor will be used to describe the network’s self-organizing behavior, for example, human health, ant colonies, bee hives, human communities, plant/insect symbiosis and so on. Any biological system that generates emergent large-scale behaviors from relatively simple components may provide insights. However, simple slime molds build distributed connected structures that, when photographed, look just like city network transportation systems. Slime mold is an example of how nature can create complex and fluidly changing forms from simple components following simple rules. So we can build our organic network controller from the emergent intelligence of slime molds.

SlimeNet. You heard about it here first.