Feature interview with Antonio Nucci: chief technology officer of narus, winner of the "CTO of the year"

Feature Interview with Antonio Nucci: Chief Technology Officer of Narus, winner of the "CTO of the Year"

Marco Mellia
Electronic and Telecommunication Department
Politecnico di Torino, Torino, Italy
mellia@tlc.polito.it

ABSTRACT

Dr. Antonio Nucci is the chief technology officer of Narus[1] and is responsible for setting the company's direction with respect to technology and innovation. He oversees the entire technology innovation lifecycle, including incubation, research, and prototyping. He also is responsible for ensuring a smooth transition to engineering for final commercialization. Antonio has published more than 100 technical papers and has been awarded 38 U.S. patents. He authored a book, "Design, Measurement and Management of Large-Scale IP Networks Bridging the Gap Between Theory and Practice", in 2009 on advanced network analytics. In 2007 he was recognized for his vision and contributions with the prestigious Infoworld CTO Top 25 Award. In 2013, Antonio was honored by InfoSecurity Products Guide's 2013 Global Excellence Awards as "CTO of the Year" [1] and Gold winner in the "People Shaping Info Security" category. He served as a technical lead member of the Enduring Security Framework (ESF) initiative sponsored by various U.S. agencies to produce a set of recommendations, policies, and technology pilots to better secure the Internet (Integrated Network Defense). He is also a technical advisor for several venture capital firms.

Antonio holds a Ph.D. in computer science, and master's and bachelor's degrees in electrical engineering from Politecnico di Torino, Italy.

1. INTERVIEW

1. *For the people that do not know you, could you introduce yourself, and describe what you have done in the past, and what you are doing now?*

   I graduated from Politecnico di Torino in Italy, where I got my Ph.D. degree in Electronics and Telecommunication Engineering in 2003. I worked in the Telecommunications Networks Group, a group that really helped me learn about communication technology, and to start building links outside of Italy. During my Ph.D. program, I had the opportunity to work with a lot of smart people in Ecole Ploytechnique Federale de Lausanne (EPFL) in Switzerland, in Université de Montréal in Canada, at Bell Labs in NY, and in Sprint Advanced Technology Labs (ATL), in California, where I spent one year solving very interesting networking problems. In 2004, I settled in the Silicon Valley, working at the Sprint Labs up to 2005, when Narus asked me to join them as the Chief Technology Officer (CTO). Since then I have been with Narus. The company was successful, launching several products in the area of traffic analysis and security. In 2010, we were acquired by The Boeing Company. Since then, I have served as the CTO of Narus, setting product strategies and vision of future Narus products. My responsibilities have expanded to cover innovation in the areas of cybersecurity and big data analytics.

2. *You left Italy more than 10 years ago. Do you think to come back and bring all your expertise to your country?*

   The reason why I decided to stay in California was for a mix of things. First, it was the opportunity to work with a lot of very smart people in Silicon Valley, and at universities such as Stanford, and various branches of the University of California, including Berkeley, Riverside, and San Diego. They are all very close, and host plenty of brilliant people to work with. I also enjoyed working with telecommunication network operators, applying everything that I have learned during my Ph.D. and to solve real problems. I kept on building on this, getting more and more responsibilities and dealing with bigger challenges. I do not know what the future is going to hold for me in Italy, and later in my career.

3. *Let us now focus more on some technical and research aspects that you think are interesting to focus on nowadays. According to your experience and current point of view: what are the hottest topics that you think will be very interesting now and possibly in the future?*

   I think that the concept of security is really hot. When I say "security", it is really a combination of several factors. It is not just about "training people" to take precautions online, or enabling people with a virtual desktop technology, so that they are given less chance to make a mistake. The concept of security is evolving. First of all, you can see many more events happening on the Internet. You browse the web, access social platforms, do banking, and buy online. More

[1] `http://www.narus.com`

and more, you do it from your mobile phone. Security has already become a sort of general, big problem.

4. *Big does not mean interesting from a research point of view.*

   Security is also getting more challenging than before: it is not just the awareness around the topic that is generating more interest from everybody and every company. A lot of the challenges come into play when you couple security (finding the needle in a haystack) with the concept of big data (with petabytes or exabytes to be analyzed). We are talking more about monitoring the entire network in real time, where you need to digest large volumes of data, several hundreds of petabits per second. The first big challenge is then a technological one.

   The second challenge is the sophistication of the security. Threats become much more sophisticated than ever, and that calls for more advanced security concepts to be applied. Which brings complexity. What intrigues me the most, is how do you take the innovation mostly developed in the academia into areas of distributed and parallel computing, and applying it to scale security algorithms.

   The last challenge comes from the variety of end-user devices. It is not just your PC that has to be protected, anymore. Now it is your PC, your tablet, your phone, your car, and your house. You buy a webcam, install it to watch you children play at home while you are at work. What if that webcam gets attacked? Consider an enterprise now. We are talking about millions of applications that enter the enterprise with people's phones. So, the problem is getting very, very complicated.

5. *Clean slate approach or incremental approach, for security?*

   We really tried so many times to patch the broken toy, right? This is not just applied to security; this is applied to the entire Internet. For how many years did we to try to patch BGP? And we have not adopted it. So, it is a general problem. The human tendency is to go and try to fix the problem. But I think that we really need to look at everything that we do, by taking security among one of the first steps in the chain. Is patching the system cheaper at the end?

6. *Thus you think that security is a requirement that must be considered from scratch and from the initial design.*

   Yes. Today, it is all flipped. First you push an application on the market. Then, you think about security. But we shall flip and say: I need to build this application according to these security standards, and maybe I am going to see fewer side effects about when the application is actually used. I'm thinking of a holistic vision, where all precautions are taken. Then, you need to keep your system updated, and protected from new and evolving threats.

7. *Let us suppose that there is a young student looking for a good topic for his Ph.D. You definitively suggest security as the good topic. What about the scenario that he/she has to consider?*

   One of the scenarios I suggest is mobile computing. According to Gartner[2] and other analysts, we are officially entering the era of mobile computing. We have mobile O.S., like Android(s) and iOS, plus social platforms. In 2015, the number of mobile device accessing the Internet will be double the number of fixed line devices accessing the Internet. Of course, when we talk about security, big data, automation, machine-to-machine, we need to understand this scenario. And here, "mobile" is the key topic that I recommend a young student to choose as the topic of the Ph.D. study.

8. *Let us go back to something more personal. Let us go back to when you were studying and reading papers: What are the papers that you remember as particularly inspiring, and that you would suggest to other students to read?*

   When you start, you look for the experts in the area that you are focusing on. So, I started with the deterministic and stochastic optimization, artificial intelligence, and operational research. I grew up with the papers from luminaries at EPFL like Jean-Yves Le Boudec [2] and Patrick Thiran [3]. I then had the chance to work with them, while I was here in California. When I was picking and choosing my Ph.D. topic, I read papers on Network Calculus by Rene Cruz [6, 7]. They gave me a strong methodological background, and that is really important even when you work on more practical problems. Knowing the theory is fundamental in any field. Technology evolves, but the concepts are typically very similar.

9. *What about your papers? Which is the one that you like most?*

   I think that the paper that was more interesting, because of the complexity of the problem and because of the interest around this specific problem, was "Googling the Internet: profiling Internet endpoints via the World Wide Web", which appeared at SIGCOMM'08 [4]. Similarly, I like the paper "Mosaic: quantifying privacy leakage in mobile networks", which appeared also at SIGCOMM this year [5]. Both papers address very practical topics, and tackle real problems. They are not perfect from a technical and scientific point of view, but from a topic and market perspective, those papers are very interesting to me. Merging real problems with smart solutions is very important for a company research group.

10. *Was it there a follow-up from those works?*

    We are a company, and we invest in research to design better products. The ideas behind "Googling the Internet" paper have been refined, and they made it into a product that aims at eliminating all those tedious and manual operations that have to be performed when dealing with traffic analysis. New applications and protocols are born every day. You need first to spot them, and afterwards derive a signature so that you can recognize them. Analysts have traditionally performed this second part by manually looking

[2] `http://www.gartner.com/newsroom/id/2408515`

into the data and searching for patterns,. We aim at making this as much as possible automatic. That's the idea behind [4] after all. The research idea has been fleshed out into Narus' product line. To me, this is the best result for an Engineer.

2. REFERENCES

[1] InfoSecurity Products Guide's 2013 Global Excellence Awards as "CTO of the Year" http://www.infosecurityproductsguide.com/world/index.html.

[2] Jean-Yves Leboudec, personal homepage, http://people.epfl.ch/jean-yves.leboudec.

[3] Patrick Thiran, personal homepage, http://icapeople.epfl.ch/thiran/.

[4] I. Trestian, S. Ranjan, A. Kuzmanovic, A. Nucci, "Googling the Internet: Profiling Internet Endpoints via the World Wide Web," Networking, IEEE/ACM Transactions on , vol.18, no.2, pp.666,679, April 2010.

[5] N. Xia, H. Song, Y. Liao, M. Iliofotou, A. Kuzmanovic, Z. Zhang, A. Nucci, "MOSAIC: Quantifying Privacy Leakage in Mobile Networks", ACM Sigcomm'13, Hong Kong, HK, August 2013.

[6] R. L. Cruz, "A Calculus for Network Delay. Part I: Network Elements in Isolation," Information Theory, IEEE Transactions on, vol.37, no.1, pp.114,131, Jan 1991.

[7] R. L. Cruz, "A Calculus for Network Delay. Part II: Network Analysis," Information Theory, IEEE Transactions on, vol.37, no.1, pp.132,141, Jan 1991.