

Network Address Translation (NAT)

In January of 1993, Tony Eng and I published the first paper to propose Network Address Translation (NAT) in CCR based on work done the previous summer during Tony's internship with me at Bellcore. Early in 1994, according to Wikipedia, development was started on the PIX (Private Internet Exchange) firewall product by John Mayes, which included NAT as a key feature. In May of 1994 the first RFC on NAT was published (RFC 1631). PIX was a huge success, and was bought by Cisco in November of 1995.

I was asked by the Sigcomm Industrial Liaison Board to write the first of a series of articles under the theme "Examples of Research Affecting the Practice of Networking." The goal of the series is to help find ways for the research community to increase its industrial impact. I argued with the board that NAT isn't a very good example because I don't think the lessons learned apply very well to today's environment. Better to start with a contemporary example like Openflow.

Why isn't NAT a good example? It's because I as a researcher had nothing to do with the success of NAT, and the reasons for the success of NAT had nothing to do with the reason I developed NAT in the first place.

Let me explain. I conceived of NAT as a solution to the expected depletion of the IPv4 address space. Nobody bought a PIX, however, because they wanted to help slow down the global consumption of IP addresses. People bought PIX firewalls because they needed a way to connect their private networks to the public Internet. At the time, it was a relatively common practice for people to assign any random unregistered IP address to private networking equipment or hosts. Picking a random address was much easier than, for instance, going to a regional internet registry to obtain addresses, especially when there was no intent to connect the private network to the public Internet.

This obviously became a problem once someone using unregistered private addresses wished to connect to the

Internet. The PIX firewall saved people the trouble of trying to obtain an adequate block of IP addresses, and having to renumber networks in order to connect to the Internet if they had already been using unregistered addresses. Indeed, PIX was conceived by John Mayes as an IP variant of the telephone PBX (Private Branch Exchange), which allows private phone systems to operate with their own private numbers, often needing only a single public phone number.

The whole NAT episode, seen from my point of view at the time, boils down to this. I published NAT in CCR (thanks Dave Oran, who was editor at the time and allowed it). For me, that was the end of it. Some time later, I was contacted by Kjeld Egevang of Cray Communications, who wanted to write an RFC on NAT so that they could legitimize their implementation of NAT. So I helped a little with that and lent my name to the RFC. (In those days, all you needed to publish an RFC was the approval of one guy, Jon Postel!) Next thing I knew, NAT was everywhere.

Given that the problem that I was trying to solve, and the problem that PIX solved, are different, there is in fact no reason to think that John Mayes or Kjeld Egevang got the idea for NAT from the CCR paper.

So what would the lesson learned for researchers today be? This: Solve an interesting problem with no business model, publish a paper about it, and hope that somebody uses the same idea to solve some problem that does have a business model. Clearly not a very interesting lesson.

I agree with the motivation of this article series. It is very hard to have industrial impact in networking today. The last time I tried was five or six years ago when I made a serious effort to turn ViAggre (NSDI'09) into an RFC. Months of effort yielded nothing, and perhaps rightly so. I hope others, especially others with more positive recent results, will contribute to this series.

Paul Francis

Max Planck Institute for Software Systems

