Internet evolution

Dah Ming Chiu, Chinese University of Hong Kong

Abstract

Can Internet evolve, and where is it evolving towards? This article gives a plausible model for Internet evolution, and tries to explain its implications for different stakeholders.

To understand the Internet, it is not enough to only describe what it is, but also guess at how it is evolving. There seems to be less agreement about the latter than the former. This short article tries to provide one viewpoint for discussion.

How does and will Internet evolve?

Despite the tremendous success of the Internet, almost all agree that as new applications and services are developed, it needs to evolve to address new demands and requirements. One way Internet evolves is through planned transition. An example of planned transition is from IPv4 to IPv6. The original Internet protocol (IPv4) was designed to use 32 bit address, which proved to be inadequate given Internet's big success. After a non-trivial amount of effort, IPv6 was standardized, changing IP address from 32 bits to 128 bits, plus some other minor tweaks to the protocol. The ensuing transition from IPv4 to IPv6, however, took a long time (and is still on-going). The slow process of transition necessitated a whole suite of interim solutions to make sure IPv4 and IPv6 can co-exist and interoperate, which is more complicated than the new protocol itself! Some causes for the slow transition may be due to the introduction of a network address translation (NAT) mechanism that allowed non-public addresses to be used, relieving the most pressing issue with 32 bit address; while another reason may be the lack of incentive for existing network users to upgrade.

This led to discussions of Internet being ossified (or unable to evolve), and how to get out of that. A notable paper called for Internet diversification [1], proposing to let different architectures (resource management policies) to be operated and co-exist over the same physical substrate, allowing them to compete, resulting in Internet's evolution. We agree that when a system allows new mechanisms to emerge and compete for survival, it will likely evolve more effectively, compared to planned transition.

In [1], the proposed ingredient that allows Internet evolution is network virtualization. Via network virtualization, different overlays can be built, and co-exist under suitable policies. In this article, we describe an alternative de facto path for Internet evolution, called "Network as a Computer", and describe what it is and its implication.

Network as a computer

In recent years, more and more services and applications are implemented in the network (or the cloud). When we go online these days, it feels like we are connecting to a huge virtual computer. In other words, users view the network as a single entity providing services to their personal devices, much like a mainframe providing services to terminals. So the network has become the computer! Of course, we know that there is no single huge computer there – the virtual computer is made up from various networking and distributed computing technologies, such as Content Distribution Network (CDN), or Data Centers and Cloud Computing platforms. So "Network as a Computer" (NaaC) is just a model of (a part of) the current Internet, as shown in Figure 1. There is an important difference in the analogy to connecting to a mainframe computer, that is, the connection to the virtual computer via the Internet access networks, may support different bandwidths (or QoS) depending on how much we pay (we will return to this point later).

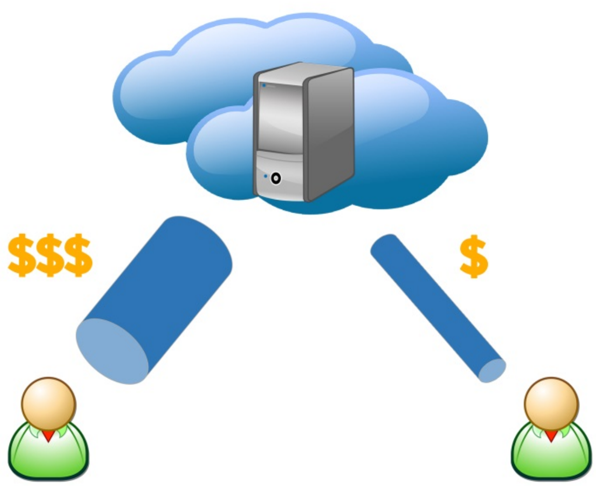

Figure 1 The "Network as a Computer" model of Internet

Besides the new services and applications (such as search and social networks), it can be argued that most traditional networking applications can be (re-)implemented in this computing model of the network, as long as we don't insist on having the user nodes keep states for the applications. Let us consider the simple example of a file transfer from user node A to user node B. In NaaC, node A first uploads the file into the computer, then node B gets notified and downloads the file from the computer. If one accepts this argument, then NaaC can be viewed as providing another form of network service, co-existing with the traditional Internet.

The benefit of viewing NaaC as a form of networking service is due to its modularity. Modularity always allows flexibility for evolution. In general, the modularity created by protocol layering allows new functions to be introduced in a lower layer without affecting higher layers and applications, leading to successful evolution. For example, Internet access protocols evolved from primarily fixed wired connections to mobile connections over the years.

NaaC provides a different form of (horizontal) modularization – it encapsulates a part of the Internet service to be provided by a virtual computer. Such modularization allows the technology implementing the virtual computer in NaaC to be replaced by better solutions when they emerge, not necessarily be based on exactly the current networking technology.

Innovation under NaaC

Though it is possible to fit earlier network applications and services in this NaaC model, the first large-scale technology that was rolled out under NaaC was Content Distribution Network (CDN). By replicating content at servers at the edge of the network, users, no matter where they are, receive content as if it comes from an omnipresent computer. The process of replicating the content is encapsulated away from the users; it can be carried out over the public Internet, or via privately-provisioned networks connecting the edge servers.

In recent years, much network research has shifted to data center and cloud networking problems. Besides various changes to existing network protocols to better support traffic inside data centers and clouds, this has also led to Software Defined Networking (SDN), which significantly changed how networking functions are provided. By removing the routing function from routers, routers become switches only responsible for forwarding, and routing policy can be configured centrally by software, providing considerable flexibility for service providers. NaaC service providers, in particular Google, quickly embraced SDN and implemented SDN in its data centers and privately deployed wide area networking connecting its data centers [2]. Although the separation of routing and forwarding is not conceptually new, as that was the case in good old telephone networks, it does open many possibilities for innovation in the context of data networking. Traditional telecom operators are also using this approach to re-invent its services, via the "Central-Office Re-architected as a Data-center" (CORD) initiative [3].

In the late 1990s, John Gage of Sun Microsystems created the term "the Network is the Computer" [4], which became the slogan of Sun in trying to redefine what computing is. Although NaaC seems to be a very similar phrase, it is actually a different idea (perhaps subtly so), because it is trying to describe what a network is. In this sense, NaaC is equivalent to saying "the Computer is the Network". NaaC is also closely related to Cloud Computing, but again, Cloud Computing is a general concept about computing, whereas NaaC is targeted at modeling a modular approach to evolve Internet, likely relying on various Cloud Computing technologies.

A NaaC-based solution for network QoS

How to provide quality of service (QoS) has always been a contentious issue in the design of Internet. The mechanism of packet switching allows limited network resources to be shared flexibly between applications and users, which in itself does not dictate whether QoS is provided or not. Fair congestion control, if implemented religiously, does imply a *best effort*

Internet service, hence no guarantee of QoS for any flow. There were many efforts over the years to add signaling in network protocols to allow different user/application flows to be able to receive different end-to-end QoS (e.g. Integrated Service and Differentiated Service), but these efforts were largely unsuccessful. The de facto solution for dealing with Internet QoS is via *over-provisioning* - network service providers try to provision more bandwidth to satisfy all traffic needs. When network technology can keep up with demand for network bandwidth at reasonable cost, the over-provisioning solution works quite well; but it is not a satisfactory solution, at least from an economic point of view, because it means high bandwidth cost and low value traffic will cause increased provisioning and make all users pay for it.

There are many reasons for the failure of previous proposals to solve network QoS, but one important reason may have been that the distributed nature of Internet control requires many stakeholders to buy in, which is very difficult to achieve. An evolutionary approach, if available, may lead to more effective solutions.

The NaaC model provides precisely this evolutionary mechanism. In NaaC networks, QoS is determined by both the (application and content) service provider and the user together. The user needs to pay for a network connection plan that can meet his/her most demanding QoS needs (this is possible in many cases, but not always). For example, if the user wishes to view streaming video at a certain bitrate, then his data plan needs to support that adequately. Based on NaaC, the content provider can also do its part in delivering the needed QoS, since its implementation of the virtual computer no longer needs to be based on best-effort service. So if the customer's QoS demand is justified business-wise, it can be delivered. Note, we are only suggesting that NaaC provides a potentially more economically viable alternative to over-provisioning; not that QoS is automatically available under NaaC, as QoS still requires both user and service provider to make their efforts respectively.

Benefit of understanding Internet evolution

What is described so far is but a conceptualization of some recent Internet developments in terms of Internet evolution. At this point, some may ask: what is the value of such a discussion of otherwise common knowledge? I think it has implications for (a) how we teach and do research on computer networking; (b) how we engineer and develop new functions for the Internet; and (c) how it may affect different stakeholders of the Internet business, as elaborated below.

Implication for academia

The inspiration for writing this article came from the experience of teaching a computer networking course for MSc students. Most MSc students may already have exposure to how the traditional protocols of the Internet work. They need to be also exposed to the latest developments of Internet using a single logical framework; they also need to understand some of the research ideas for future Internet with a critical mind. The discussion of Internet evolution seems to fit this bill. In this context, we can also discuss different application and user requirements, and in which direction Internet will more likely evolve. Perhaps now is a good time to write a new textbook on computer networking based on the theme of Internet evolution.

Hopefully the discussion of Internet evolution can also contribute to networking research. In the early days of Internet development, there used to be active (sometimes heated) discussion about network architecture issues. Perhaps due to maturation and "ossification" of the Internet, and the future Internet research communities being more fragmented; such discussions seem to have quieted down. Will the topic of Internet evolution ignite more architectural level discussions?

Implication for standardization

Internet Engineering Task Force (IETF) is often credited as the effective standardization process that led to Internet's success. IETF is open and is able to involve different stakeholders (academia, service providers, equipment vendors, users as well as governments), so it serves to bring in different ideas, have the ideas debated, till broad level buy-in is reached. Having Internet protocols specified and standardized enables network protocols to be implemented and run on different hardware platforms, which allows market competition and technology evolution, as new hardware platforms replace old.

In the development of NaaC technology in recent years, Open Source software played important roles, for example in the development of SDN technology. The use of software, in particular open source (so that developers can build on each other's efforts), to roll out new technology can significantly cut down the lead time needed to follow the standardization-manufacturing-deployment cycle. This is possible again because of the evolutionary NaaC model, and often the virtual computer is under a single administrative control. This does not mean protocol specification and standardization is not needed, but it can happen in parallel with innovation and development of the technology.

Implication for networking businesses

The NaaC model implies that Internet service is now provided by two parallel networks, one is the traditional Internet as we know, and the other is the NaaC network, the inside of the latter may not be running all traditional Internet protocols. What fraction of the whole Internet service is running as the first network versus the second network is driven by forces externally, hence determined by the process of evolution. If significant traffic and services move from the first network into the second network over time, the businesses running the first network may need to transform themselves to also provide the service of the second network, if they do not want to become marginalized.

In the traditional Internet business, the equipment makers may have more control of their business if they establish leadership positions in IETF. In the world of the NaaC Internet, the rules of the game have changed. You cannot focus only on IETF, and rely on that to manage your product mix. Popular networking technology may be developed from the Open Source community first.

Societal and user considerations

What will this evolution process lead us to, and will it be for better or for worse? We don't have a magic glass. But some ramifications are apparent:

a) Innovation is more likely to occur from the NaaC side of the Internet, as that side is driven by more new business needs, and is likely to have more funding;
b) The issue of net neutrality does not go away; in fact it may become more complicated, especially if the virtual computer continues to provide more and more traditional networking services;
c) We assumed much innovation in NaaC is based on open source that lead to new standardized open protocols and systems. If much of the developments were closed, it could slow down the evolution of the network as a whole;
d) Network surveillance and censorship becomes easier.

Despite of potentially more governance issues, we are confident that we are on a path to more innovation leading to a better Internet for us all. My personal guess is that the Internet will not become dominated by the NaaC service model, but reach a suitable equilibrium.

Concluding remarks

We posit that Internet is evolving, explainable by a "Network as a Computer" model. Compared with the ossified period, we are seeing innovative solutions to previous challenges, such network QoS. It is important for students, engineers and various stakeholders to understand the Internet in light of its evolution, in order to address its new challenges, whether technical or in relation to its governance. It is full of new opportunities.

References

[1] JS Turner and DE Taylor, "Diversifying the Internet", IEEE Globecom 2005.

[2] Jim Duffy, "Google opens up on its SDN", Network World, 2015

[3] CORD, http://opencord.org/

[4] https://en.wikipedia.org/wiki/John_Gage