The Unreliable Internet

Geoff Huston
April 2001

Any history of the Internet, and there are quite a few of them these days, will cite that one of the major motivations of the original DARPA packet switched network project was a focus on research into highly resilient packet switching architectures. The Internet started out as a project with the clear objective of defining a network architecture that could withstand various forms of failure. The loss of a circuit, the loss of a router, the temporary overloading of a network component, or any other similar event should not bring down the entire network. As a greater challenge, in the face of such failure, the network should be able to heal itself, and use alternate paths that detour around the damage site quite automatically. To up the stakes even further, the objective was to make this healing function so transparent to the network's applications that any application that had traffic flowing across the damage site should not see any disruption to the transport service, and the application should not be aware of the network outage and subsequent automatic repair. As far as the Internet is concerned, the detection of failure and the subsequent healing around the failure is seamless and invisible to the applications that use the network. As is often said, the Internet sees damage and routes around it.

If the Internet can do all this then why talk about the unreliable Internet? Surely this is one of the more resilient communications systems that could ever be engineered, given that the network is able to silently and efficiently adapt to changes in its internal connectivity. While the theory is still correct, the practice is somewhat different. Resiliency can still be engineered in Internet networks, but, increasingly, this is a function performed by systems at the outer edge of the network rather than relying on resiliency within the core of the network. Like so many other role responsibilities in the Internet, resiliency is becoming the customer's responsibility.

The provision of Internet service is an excellent example of an open market. There are competitive enterprises offering Internet access services in almost every part of the Internet. The result is that for any potential ISP customer there is a choice of ISPs who can offer Internet access. The ISPs in such a market compete on price and quality. Within the dial-up access market, quality, to a certain level, is immediately apparent to any consumer. Quality is expressed as the ability to complete a call, rather than encounter constant busy tones from the modems. Quality may be the ability to roam through multiple locations, or the availability of web servers or other value added services used to complement the basic dial access. But for the fixed access provider quality is a somewhat different metric. Here the customer of the service is an enterprise network, and the metrics of quality are somewhat differently expressed. Here quality, or a service level specification, often refers to the level of congestion within the ISP network, and the extent of connectivity of the network, and the latency of the paths within the network. But for the end customer these network-specific metrics are not of particular relevance. The end customer may be accessing sites reachable across many network transit hops, and the perceived quality of the end-to-end transactions is not only dependant on the quality of the local ISP, but also dependant on the quality of a large collection of remote ISPs. These are all too often factors well beyond the immediate control of the local access ISP. In such a market, where the metrics of quality are so uncertain to customers, the only differentiator that ISPs can use is price. Competitive pressure in price-based markets leads to a commodity market, where the most efficient provider is able to set the market prices for all other competitive providers.

Such a market forces ISPs to become highly efficient in the carriage of IP traffic. There is not a generous margin for over-engineering of the provider's network, as the drive for ever higher efficiencies within the network leads to networks which offer adequate service levels, but have little margin to withstand exceptional events such a failure of various network switching or transmission components. The Internet routing protocols will certainly detect and heal such network failures, but in so doing will divert traffic to alternate paths. If such paths were already carrying an optimal traffic load, this additional diverted load will cause various forms of network congestion. While the network itself has managed to heal itself, the quality of the delivered service has nevertheless suffered under such conditions. The conventional way to address such service issues is to increase the overall capacity within the network, so that diversion of traffic onto alternate paths will not cause undue disruption to existing traffic flows. At the extreme end of such engineering is the fully redundant network, where every active traffic-carrying circuit is backed up by a standby idle circuit. SUch a network offers excellent levels of resiliency and constant quality levels under various common failure conditions, but at any point more than half of the total network capacity is idle. This is not a highly efficient network. In the drive to achieve ever greater efficiency of operation, one of the first things to look at is this margin of over-provisioned capacity, and the inevitable outcome is that this margin is gradually eroded in order to drive down network costs and increase operating efficiency.

If this were the only factor regarding network resiliency then it may well be a short-lived concern. The energetic deployment of fibre optic cables has created a market for capacity where supply is outpacing demand, and the unit cost of network transmission capacity is inevitably coming down. As these cost decline it becomes feasible to increase the margins of over-provisioning within the network and address this issue service quality during periods of component failure.

There is one remaining essential vulnerability of the entire system. The critical point of failure is the network's routing protocol. In the face of component failure it is the routing protocol that undertakes the role of healing the network. If the routing protocol fails then there is nothing to back it up, and parts of the network will be isolated. Service will be disrupted. As a network scales in size and complexity, the routing protocol carries more information across a wider network domain. There are various limits inherent in the memory capacity of individual routers, processing capacity and even limits within the protocol itself. Also routing traffic is carried in IP packets along with normal data. If a segment of the network becomes chronically congested the routing protocol itself cannot communicate through congestion point. Routing failure is a rare occurrence in most networks, but it can and does occur from time to time in almost every network. The network is not perfectly resilient.

However, this is a network-centric view of resilience. To look at the same problem from the perspective of the fixed line customer there is yet another point of failure where there is no automatic backup and recovery. This is the access line itself and the network components which terminate this line in the provider's network. If these particular components fail, the rest of the network may recover quite perfectly, but for the affected customers, the service has been effectively suspended.

For some fixed line customers this model of service may represent an acceptable compromise between price and service resilience. For other customers this risk of service failure is simply unacceptable. As the Internet becomes an integral part of more and more enterprises, the number of customers for whom resiliency is absolutely necessary is increasing. This applies not only to end customers, but even more so to ISPs themselves. An ISP who purchases upstream transit services from a single transit ISP is also dependant on the resilience of their chosen transit service provider. Failure of the transit service becomes failure of their service.

One potential response is to purchase two distinct access services from a single provider, preferably using two or more physically separate access lines and connecting to distinct network access points. While this addresses the risk of failure of the access system, there is still the residual risk of failure of the provider's routing system, or failures of upstream provider networks.

Given that there is a competitive market in the supply of access services and a competitive market in the supply of transit services, a common approach to this risk is to purchase two or more services from distinct providers. The intent of this arrangement is that in the event of failure of any form in one provider's network, the end customer will still receive service from the other provider. This configuration is termed "multi-homing", and refers to a customer's network being attached, or "homed" in two or more upstream provider networks.

Multi-homing is very widespread within the ISP sector, and it is now a common situation that where an ISP purchases transit services from one transit provider, it will invariably undertake a similar purchase from a second transit provider. In that way the ISP is attempting to safeguard its service against failure in any single transit service.

The more recent trend is that of multi-homing of fixed line end customers, or corporate or enterprise networks. The same drivers that have prompted ISPs to multi-home are now extending to the end customer, and the same solutions are just as effective for end customers as they are for ISPs. But this has a curious twist. As more customers turn to multi-homing as a means to provide resiliency of their Internet service, there are fewer customers who are critically dependant on the resiliency of any single service provider. In essence, the customer base is placing less value in the resiliency of any single provider and directly investing in resiliency through the purchase of multiple access services. The message this passes to service providers is becoming increasingly clear. Customers are placing less trust, and less value on the resilience of any single provider's network. This in turn reduces the motivation of service providers to extensively over-engineer their networks. The investment in such efforts to engineer significant over-capacity into the network is not one which will allow a service provider to operate at an increased unit price. Through their actions, customers are passing the message to the provider market that service resilience is fast becoming a customer responsibility, rather than a required and highly valued network attribute.

For the Internet, a network originally designed with resilience as a major objective, this is indeed a most curious development.