Resilience of the Internet Interconnection Ecosystem

Chris Hall, Ross Anderson, Richard Clayton, Evangelos Ouzounis, and Panagiotis Trimintzios

Abstract In 2010 the European Network and Information Security Agency (ENISA) launched a study to investigate the resilience of the Internet's interconnection system and come up with policy recommendations. A large number of stakeholders were contacted, and their expertise has been reflected in the study. The formal outcome of the study was the publication by ENISA in early 2011 of a detailed technical report, '*Inter-X: Resilience of the Internet Interconnection Ecosystem'*. This paper presents a much abridged version of the ENISA report. In it, we present a summary of the problems that the Internet faces in keeping its interconnection system resilient, along with the recommendations proposed to policy makers.

Chris Hall
Highwayman Associates, Woodvill Cottage, Woodvill Road, Leatherhead, KT22 7BP, UK. e-mail: `chris.hall@highwayman.com`

Ross Anderson
Computer Laboratory, University of Cambridge, Cambridge, CB3 0FD, UK. e-mail: `ross.anderson@cl.cam.ac.uk`

Richard Clayton
Computer Laboratory, University of Cambridge, Cambridge, CB3 0FD, UK. e-mail: `richard.clayton@cl.cam.ac.uk`

Evangelos Ouzounis
European Network and Information Security Agency, PO Box 1308, Heraklion, 71001, Greece. e-mail: `evangelos.ouzounis@enisa.europe.eu`

Panagiotis Trimintzios
European Network and Information Security Agency, PO Box 1308, Heraklion, 71001, Greece. e-mail: `panagiotis.trimintzios@enisa.europa.eu`

1 Introduction

The Internet has been pretty reliable thus far, having recovered rapidly from most of the incidents that have occurred in its history. The effects of natural disasters such as Hurricane Katrina [21], terrorist attacks such as '9/11' [60], and assorted technical failures have all been limited in time and space [67, 68]. It is impossible to say whether we have just been lucky, and that sooner or later some event will catch us out. However it does appear likely that the Internet could suffer a systemic failure, leading perhaps to localised collapse and system-wide congestion, in certain circumstances:

- A regional failure of the physical infrastructure on which the Internet depends (such as the bulk power transmission system) or the human infrastructure needed to maintain it (for example, if pandemic flu causes millions of people to stay at home out of fear of infection).
- Cascading technical failures, of which some of the more likely near-term scenarios relate to the imminent changeover from IPv4 to IPv6. Common-mode failures involving updates to popular makes of router (or PC) may also fall under this heading.
- A coordinated attack in which a capable opponent disrupts the routing fabric, perhaps by broadcasting many thousands of bogus routes from a large ISP, or from a large number of compromised routers.

There is evidence that some implementations of the Border Gateway Protocol (BGP), the Internet's inter-domain routing protocol, are surprisingly fragile. There is evidence that some concentrations of infrastructure lack resilience and that significant disruption could be caused by localised failure. There is evidence that the health of the interconnection system as a whole is not a high priority for the networks that make up that system—by and large each network strives to provide a service which is reliable most of the time, at minimum achievable cost. The economics do not favour high dependability of the system as a whole as there is no incentive for anyone to provide the extra capacity that would be needed to deal with large-scale failures.

To date, we have been far from equilibrium: the rapid growth in Internet capacity has masked a multitude of sins and errors. However, as the Internet matures, as more and more of the world's optical fibre is lit, and as companies jostle for advantage, the dynamics may change.

There may well not be any immediate cause for concern about the resilience of the Internet interconnection ecosystem, but there is certainly cause for concern about the lack of good information about how it works and how well it might work if something went very badly wrong.

In 2010 the European Network and Information Security Agency (ENISA) launched a study to investigate the resilience of the Internet's interconnection system and come up with policy recommendations. A large number of stakeholders were contacted, and their expertise reflected in the study. The formal outcome was the publication by ENISA in early 2011 of a detailed

230 page technical report, '*Inter-X: Resilience of the Internet Interconnection Ecosystem'* [24] which goes into great detail on the security and economics of Internet routing, covering everything from the mechanics of route leaks to the financial conditions of the largest players (the Tier 1 transit providers).

In this paper we only present a brief canter through the landscape in Sections 2 to 14 before discussing the main recommendations in Section 15. The interested reader is referred to the full report for all the details.

2 Scale and Complexity

The Internet is very big and very complicated! The interconnection system we call the Internet comprises (in March 2011) some 37 000 'Autonomous Systems' or ASs (ISPs or similar entities) and 350 000 blocks of addresses (addressable groups of machines), spread around the world [37].

This enormous scale means that it is hard to conceive of an external event which would affect more than a relatively small fraction of the system—as far as the Internet is concerned, a large earthquake or major hurricane is, essentially, a little local difficulty. However, the failure of even a small fraction of the Internet may still have a significant impact on a great many people. When considering its resilience it is necessary to consider not only the global issues, but also a large number of separate, but interconnected, local issues.

The complexity of the system is not just a matter of its sheer scale and the number of interconnections between ASs, but is compounded by a number of factors:

- Modelling interconnections is hard because we only ever have a partial view. The Internet has a number of layers, each with its own properties and each interacting with the other layers. For example, the connections between ASs use many different physical networks, often provided by third parties, which are themselves large and complicated. Resilience depends on the diversity of interconnections, which in turn depends on physical diversity—which can be an illusion, and is often unknown.
  While it is possible to discover part of the 'AS-level topology' of the Internet (which ASs are interconnected) [59], it would be more valuable from a resilience perspective to know the 'router-level topology' (the number, location, capacity, traffic levels etc. of the actual connections between ASs). If we want to estimate how traffic might move around when connections fail, we also need to know about the 'routing layer' (what routes the routers have learned from each other) so we can estimate what routes would be lost when given connections failed, and what routes would be used instead. That also touches on 'routing policy' (the way each AS decides which routes it will prefer) and the 'traffic layer'—where end-user traffic is going to and from. This last is perhaps the most important layer, but very little is known about it on a global scale.

- The interconnection system depends on other complex and interdependent systems. The routers, the links between them, the sites they are housed in, and all the other infrastructure that the interconnection system depends on, themselves depend on other systems—notably the electricity supply—and those systems depend in their turn on the Internet.
- The interconnection ecosystem is self-organising and highly decentralised. Decisions to interconnect are made independently by the ASs, driven by their need to reach, and be reachable from, the entire Internet. The same holds at lower levels: the administrators of an AS configure their routers to implement their routing policy, then the routers select and use routes. But different routers in the same AS may select different routes for a given destination, so even the administrators may not know, a priori, what path traffic will take [27].
- The interconnection ecosystem is dynamic. Its shape changes all the time as new connections are made, or existing connections fail or are removed. At the corporate level, transit providers come and go, organisations merge, and so on. At the industry level, the recent rise of the content delivery networks (CDNs) has changed the pattern of interconnections [26, 30].
- The patterns of use are also constantly evolving. The rise of the CDNs also changed the distribution of traffic; and while peer-to-peer (P2P) traffic became a large proportion of total volumes in the early-to-mid 2000s, now video traffic of various kinds is coming to dominate in terms of both volume and growth [16, 44].
- The Internet is continuing to grow. In fact, just about everything about it continues to grow: the number of ASs, the number of routes, the number of interconnections, the volume of traffic, etc. [37, 38, 58, 64].

The scale and complexity of the system are hard to grasp. Resilience is itself a slippery concept, so the 'resilience of the interconnection system' is non-trivial to define—let alone measure!

3 The Nature of Resilience

There is a vast literature on reliability where engineers study the failure rates of components, the prevalence of bugs in software, and the effects of wear, maintenance etc.; the aim being to design machines or systems with a known rate of failure in predictable operating conditions [70]. Robustness relates to designing systems to withstand overloads, environmental stresses and other insults, for example by specifying equipment to be significantly stronger than is needed for normal operation. In traditional engineering, resilience was the ability of a material to absorb energy under stress and release it later. In modern systems thinking, it means the opposite of 'brittleness' and refers to the ability of a system or organisation to adapt and recover from a serious failure, or more generally to its ability to survive in the face of threats, including the prevention or mitigation of unsafe, hazardous or detrimental conditions that threaten its existence [34]. In the longer term, it can also mean evolvability: the ability of a system to adapt gradually as its environment changes—an idea borrowed from systems biology [41, 76].

For this study we follow Xie et al. [80] and define the resilience of a system as *the ability to provide and maintain an acceptable level of service in the face of various faults and challenges to normal operation*—i.e. the ability to adapt itself to recover from a serious failure and more generally to survive in the face of threats. A given event may have some impact on a system and hence some immediate impact on the service it offers. The system will then recover, service levels improve and at some time full service is restored.

Resilience therefore refers both to failure recovery at the micro level, as when the Internet recovers from the failure of a router so quickly that users perceive a connection failure of perhaps a few seconds (if they notice anything at all); through coping with a mid-size incident, as when ISPs provided extra routes in the hours immediately after the 9/11 terrorist attacks by running fibres across co-location centres; to disaster recovery at the strategic level, where we might plan for the next San Francisco earthquake or for a malware compromise of thousands of routers. In each case the desired outcome is that the system should continue to provide service in the event of some part of it failing, with service degrading gracefully if the failure is large.

There are thus two edge cases of resilience:

1. The ability of the system to cope with small local events, such as machine failures, and reconfigure itself, over a time scale of seconds to minutes, in an essentially automated manner. This enables the Internet to cope with day-to-day events with little or no effect on service—it is reliable. This is what most network engineers think of as resilience.
2. The ability of the system to cope with and recover from a major event, such as a large natural disaster or a capable attack, on a time scale of hours to days or even longer. This type of resilience includes, first, the ability to continue to offer some service in the immediate aftermath, and second, the ability to repair and rebuild thereafter. The key words here are 'adapt' and 'recover'. This 'disaster recovery' is what civil authorities tend to think of when they use the term 'resilience'.

We are interested in the resilience of the ecosystem in the face of events which have medium to high impact and which have a correspondingly medium to low probability. We thus emphasise the second of these cases.

Robustness is an important aspect of resilience. A robust system will have the ability to resist assaults and insults, so that whatever some event is throwing at it, it will be unaffected, and a resilient response is not required. While resilience is to do with coping with the impact of events, robustness is to do with reducing this impact in the first place. The two overlap, and from the users' perspective the distinction may be a fine one; what the user wants is for the system to be predictably dependable.

Resilience is context-specific. Robustness can be sensibly defined only in respect of specified attacks or failures, and in the same way resilience also makes sense only in the context of recovery from specified events, or in the face of a set of possible challenges of known probability. We call bad events of known probability 'risk', but there is a separate problem of 'uncertainty' where we do not know enough about possible future bad events to assign them a probability at all. In the face of uncertainty, it is difficult to assess a combination of intermediate levels of service and recovery/restoration times, especially when what is acceptable may vary depending on the nature and scale of the event.

Moreover, no good metrics are available to actually assess the performance of the Internet or its interconnects. This makes it harder still to specify acceptable levels of service. For the Internet the problem is compounded by scale and complexity (see above) and by lack of information (see below), which makes it hard to construct a model which might be used to attach numbers to resilience. It is even hard to assess what impact a given single event might have—an earthquake in San Francisco of a given severity may have a predictable impact on the physical infrastructure, but that needs to be translated into its effect on each network, and hence the effect on the interconnection ecosystem.

Given these difficulties (and there are many more), service providers commonly fall back on measures which improve resilience in general terms, hoping that this will improve their response to future challenges. This qualitative approach runs into difficulty when the cost of an improvement must be justified on much more restricted criteria. For the Internet as a whole, the cost justification of investment in resilience is an even harder case to make.

4 The Lack of Information

Each of the ASs that make up the Internet has a Network Operation Centre (NOC), charged with monitoring the health of the AS's network and acting when problems occur. There is no NOC for the Internet.

In fact it is even worse than that. ASs understand their own networks but know little about anyone else's. At every level of the interconnection system, there is little global information available, and what that is available is incomplete and of unknown accuracy.

In particular:

- there is no map of physical connections—their location, capacity, etc.;
- there is no map of traffic and traffic volume;
- there is no map of the interconnections between ASs—what routes they offer each other.

The Internet interconnection system is, essentially, opaque. This opacity hampers the research and development communities in their attempts to understand how the Internet works, making it hard to develop and test improvements; it makes the study and modelling of complex emergent properties such as resilience even harder still.

The lack of information has a number of causes:

- **Complexity and scale**. To maintain accurate maps of the networks of fibre around the world might be a tractable problem. But many different logical connections run over these physical fibres, each of which will carry network traffic for numerous providers, which in turn support yet more providers' networks and circuits—rapidly multiplying up the combinations and permutations of overlapping use of the underlying fibre. Furthermore, much of this is dynamic—providers reroute existing networks and circuits as they extend or adapt their networks. Meticulous record keeping is required, but even within a single AS it is not always achieved. It would be even more complex to extend this mapping so as to measure the traffic volumes, given the sheer number of connections between networks.
- **The information hiding properties of the routing system** [6, 11, 15]. When mapping connections by probing the system from the outside, each probe will reveal something about the path between the two (inside and outside) points in the Internet at that specific time. But the probe reveals little about what paths may exist at other times, or what path might be taken if any part of the usual path fails, or what the performance of those other paths might be.
- **Security concerns**. Mapping the physical layer is thought to invite people with bad intentions to improve their target selection; so those maps that do exist are seldom shared [7].
- **The cost of storing and processing the data**. If there was complete information, there would be a very great deal of it, and more would be generated every minute. Storing it and processing it into a usable form would be a major engineering task.
- **Commercial sensitivity**. Information about whether, how and where networks connect to each other is deemed commercially sensitive by some. Information about traffic volumes is quite generally seen as commercially sensitive. Because of this, some advocate powerful incentives to disclose it, possibly in anonymised and aggregated form.
- **Critical information is not collected in the first place, or not kept up to date**. Information gathering and maintenance costs money, so there must be some real use for it before a network will bother to gather it or strive to keep it up to date. The Internet Routing Registries (IRRs) are potentially excellent resources, but are not necessarily up to date, complete or accurate, because the information seldom has operational significance (and may in any case be deemed commercially sensitive) [69].

- **Lack of good metrics**. While there are well-known metrics for the performance of connections between two points in a network, there are none for a network as a whole or, indeed, for a network of networks.

The poor state of information reflects not only the difficulty of finding or collecting data, but also the lack of good ways to process and use it even if one had it.

4.1 Incidents as a Source of Information

Small incidents occur every day, and larger ones every now and then. Given the lack of information about the interconnection system, the results of these natural experiments tell us much of what we presently know about its resilience. For example, we know the following:

- It is straightforward to divert traffic away from its proper destination by announcing invalid routes [8, 9, 18, 14, 62, 75], and in one well-known incident in February 2008 YouTube became inaccessible for a few hours [10, 53, 54, 66]. More publicity, and political concern, was raised by a 2010 incident during which China Telecom advertised a number of invalid routes, effectively hijacking 15% of Internet addresses for 18 minutes [4, 5, 19, 50, 52].
- Latent bugs in BGP implementations can disrupt the system. Most recently, in August 2010, an experiment which sent an unusual (but entirely legal) form of route announcement triggered a bug in some routers, causing their *neighbours* to terminate BGP sessions, and for many routes to be lost [67, 68]. The effects of this incident lasted less than two hours.
- In some parts of the world a small number of cable systems are critical. Undersea cables near Alexandria in Egypt were cut in December 2008 [29]. Interestingly, three cable systems were affected at the same time, and two of those systems had been affected similarly in January/February of that year [61, 72]. This seriously affected traffic for about two weeks [83, 84, 85].
- The Internet is critically dependent on electrical power and the August 2003 blackout in the Northeastern US and Canada had a widespread effect on connectivity, although the largest provider networks stayed up [20]. A large power outage in Brazil in November 2009 caused significant disruption, though it lasted only four and a half hours. Previous blackouts in Brazil had been attributed to 'hackers', suggesting that these incidents are examples of the risk of inter-dependent networks. However, this particular conspiracy theory has been refuted [71].
- The ecosystem can work well in a crisis. The analysis of the effect of the destruction at the World Trade Centre in New York on 11^{th} September 2001 shows that the system worked well at the time, and in the days thereafter, even though large cables under the buildings were cut and other facilities

were destroyed or damaged. Generally, Internet services performed better than the telephone system (fixed and mobile) [3].

These sorts of incident are well known. However, hard information about the exact causes and effects is hard to come by—much is anecdotal and incomplete, while some is speculative or simply apocryphal. Valuable information is being lost. The report *The Internet under Crisis Conditions: Learning from September 11* [60], is a model of clarity; but even there the authors warn:

> "...while the committee is confident in its assessment that the events of September 11 had little effect on the Internet as a whole ..., the precision with which analysts can measure the impact of such events is limited by a lack of relevant data."

5 Resilience and Efficiency

There are fundamental tensions between resilience and efficiency. Resilience requires spare capacity and duplication of resources, and systems which are loosely coupled (composed of largely independent sub-systems) are more resilient than tightly coupled systems whose components depend more on each other [51]. But improving efficiency generally means eliminating excess capacity and redundant resources.

A more diverse system is generally more resilient, but it will also be more expensive and complex. Diversity of connections is most efficiently achieved using infrastructure whose cost is shared by many operators, but collective-action problems can undermine the resilience gain. It is efficient to avoid duplication of effort in the development of software and equipment, and efficient to exploit economies of scale in its manufacture, but this reduces the diversity of the equipment that is used. It is efficient for the entire Internet to depend on one protocol for its routing, but this creates a single point of failure. Setting up and maintaining multiple, diverse, separate connections to other networks costs time and effort and creates extra complexity to be managed.

The Internet is a loosely coupled collection of independently managed networks. However, at its core there are a few very large networks, each of which strives to be as efficient as possible both internally and in its connections to other networks. So it is an open question whether the actual structure of the Internet is as resilient as its architecture would suggest. In the past it has been remarkably resilient, and it has continued to perform as it has evolved from a tiny network connecting a handful of research facilities into the global infrastructure that connects billions today. However, as elsewhere, past performance is no guarantee of future results.

6 Resilience and Equipment

A particular concern for the Internet interconnection system is the possibility of an internal technical problem having a systemic effect. The imminent changeover to IPv6 will provide a high-stress environment in which such a problem could be more likely to manifest itself, and the most likely proximate cause of such a problem is bugs in BGP implementations, which could be serious given the small number of router vendors. There have been a number of incidents in which large numbers of routers across the entire Internet have been affected by the same problem, when something unexpected triggers a bug in the software—and occasionally even in the specification of BGP [67].

No software is free from bugs, and universal dependence on BGP and universal connectedness makes bugs more serious. ISPs may test equipment before buying and deploying it, but those tests concentrate on issues directly affecting the ISP, such as the performance of the equipment and its ability to support the required services. Manufacturers test their equipment too. But both ISPs and manufacturers are mostly concerned that the equipment works well under normal circumstances. Individual ISPs cannot afford to do exhaustive testing of low-probability scenarios for the benefit of the Internet at large, while for their part the manufacturers balance the effort and time spent testing against their customers' demands for new and useful features, new and faster routers and less expensive software. Of lesser concern, unfortunately, is how secure routers and routing protocols might be against deliberate attempts to disrupt or suborn them.

A number of respondents to the consultation exercise which informed the writing of our report felt that money spent on more strategic testing of equipment and protocols would be money well spent.

7 Service Level Agreements and 'Best Efforts'

In any market in which the buyer has difficulty in establishing the relative value of different sellers' offerings, it is common for sellers to offer guarantees to support their claims to quality. Service Level Agreements (SLAs) perform that function in the interconnection ecosystem. From a resilience perspective, it would be nice to see ISPs offering SLAs that covered not just their own networks but the interconnection system too, and customers preferring to buy service with such SLAs. Unfortunately, SLAs for Internet access in general are hard, and for transit service are of doubtful value. In particular, the SLAs that are offered do not extend beyond the borders of that network, so whatever they do guarantee does not cover the interconnection system—the part between the borders of all networks.

Providers do not attempt to guarantee anything beyond their borders because they cannot. An end-to-end guarantee would require a back-to-back system of contracts between networks so that liability for a failure to perform would be borne by the failing network. That system of contracts does not exist, not least because the Internet is not designed to guarantee performance. It is fundamental to the current Internet architecture that packets are delivered on a 'best efforts' basis—the network will do its best but it does not guarantee anything. The Internet leaves the hard work of maintaining a connection to the end-points of the connection—the 'end-to-end' principle. The Transmission Control Protocol (TCP), which carries most Internet traffic (apart from delay-sensitive traffic such as voice), will reduce demand if it detects congestion—it is specifically designed to adapt to the available capacity, not to guarantee some level of performance.

The other difficulty with SLAs is what can (or should) be measured. For a single connection between A and B it is clear what can be measured, but it is not clear what level of performance could be guaranteed. But now consider a connection from A in one network to F in another network, which traverses four other networks (B, C, D, E) and the connections between them, as shown in Fig. 1.

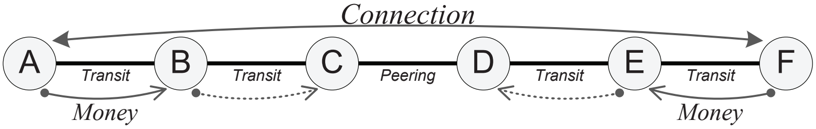

Fig. 1 How the flow of money fails to match the flow of traffic.

All these networks are independent, and have their own SLAs, each extending only as far as their borders. If we follow the money, A is paying directly for a connection to B and indirectly for packets sent from B to C. Similarly, F is paying directly for a connection to E and indirectly for the connection from E to D. C and D (who could well be global 'tier one' providers) have their own arrangements. But if E has low standards, or is having a bad day, to whom does A complain? B has a contract with A, and offers an SLA, but that does not extend beyond B. B also has a contract with C, with a different SLA, but even if B complained to C about its customer's problem we have come to the end of the money trail: C may not be able to hold D to account, as it peers with D, and it has no relationship with E.

Even if it were possible to establish an end-to-end SLA for this connection, and pin liability on the failing network, there are hundreds of thousands of paths between network A and the rest of the Internet. So whatever value SLAs may have, they do not offer a contractual framework through which customers can influence the resilience of the interconnection system, even if they wished to do this. Few customers understand the issue, or care to do anything about it. Generally the Internet is remarkably reliable, so a customer's principal interest in choosing a supplier is price—possibly moderated by the supplier's reputation.

8 Reachability, Traffic and Performance

While end-users care about traffic and performance, the basic mechanism of the interconnection system—the BGP protocol [65]—only understands reachability. Its function is to provide a way for every network to reach every other network, and for traffic to flow across the Internet from one network to another. All ASs (the ISPs and other networks that comprise the Internet) speak BGP to each other, and reachability information spreads across the 'BGP mesh' of connections between them. BGP is the heart of the interconnection system, so its many deficiencies are a problem.

The problems within the protocol itself include:

- There is no mechanism to verify that the routing information distributed by BGP is valid. In principle traffic to any destination can be diverted—so traffic can be disrupted, modified, examined, or all three. The security issues are covered further below.
- There is no mechanism in BGP to convey capacity information—so BGP cannot help reconfigure the interconnection system to avoid congestion [82]. When a route fails, BGP will find another route to achieve reachability, but that route may have insufficient capacity for the traffic it now receives.
- The mechanisms in BGP which may be used to direct traffic away from congestion in other networks—'inter-domain traffic engineering'—are strictly limited [13, 63, 81].
- When things change, BGP can be slow to settle down ('converge') to a new, stable, state [32, 33].
- The ability of BGP to cope, or cope well, under extreme conditions is not assured [78].

End-users expect to be able to reach every part of the Internet, so reachability is essential. But they also expect to be able to move data to and from whatever destination they choose, so they expect their connection with that destination to perform well. As BGP knows nothing about traffic, capacity or performance, network operators must use other means to meet end-users' expectations. When something in the Internet changes, BGP will change the routes used to ensure continuing reachability, but it is up to the network operators to tinker with the automated choices to ensure that performance is adequate, and take other steps if it is not.

Service quality in a 'best efforts' network is all to do with avoiding congestion, for which it is necessary to ensure that there is always sufficient capacity. The most effective way to do that is to maintain enough spare capacity to absorb the usual short-term variations in traffic and provide some safety margin. Additional spare capacity may be maintained to allow time (weeks or months, perhaps) for new capacity to be installed to cater for long-term growth of traffic. Maintaining spare capacity in this way is known as 'over-provisioning'; it is key to day-to-day service quality and to the resilience of the interconnection system.

Each operator constantly monitors its network for signs of congestion and will make adjustments to relieve any short term issues. In general the pattern of traffic in a network of any size is stable from day to day and month to month [1]. An operator will also monitor their network for long term trends in traffic. The management of capacity is generally done on the basis of history, experience and rules of thumb, supported by systems for gathering and processing the available data. The levels of spare capacity in any network will depend on many things, including how the operator chooses to balance the cost of spare capacity against the risk of congestion.

A key point here is that capacity is managed on the basis of actual traffic and the usual day-to-day events, with some margin for contingencies and growth. Capacity is not managed on the basis of what might happen if some unusual event causes a lot of traffic to shift from one network to another. If an event has a major impact, then the amount of spare capacity within and between networks will determine the likelihood of systemic congestion. So each individual network's degree of over-provisioning makes some contribution to the resilience of the whole—though it is hard to say to what extent.

If an event disables some part of the Internet, BGP will work to ensure that reachability is maintained, but the new paths may have less capacity than the usual ones, which may result in congestion. For many applications, notably web-browsing, the effect is to slow things down, but not stop them working. More difficulties arise with any sort of data that is affected by reduced throughput or increased delay, such as VoIP (voice over Internet Protocol) and streaming video. Congestion may stop these applications working satisfactorily, or at all [43, 77].

The important distinction between reachability and traffic is illustrated by considering what appears to be a simple metric for the state of the Internet: the percentage of known destinations which are reachable from most of the Internet, at any given moment. This metric may be used to gauge the impact of a BGP failure, or of the failure of some critical fibre, or any other widely felt event. But while the significance of, say, 10% of known destinations becoming unreachable is obviously extremely high for the 10% cut off, it may not be terribly significant for the rest of the Internet. We'd prefer to know the amount, and possibly the value, of traffic that is affected. If the 10% cut off is largely self-contained in traffic terms, or insignificant in economic terms, the failure has less impact. If the 10% cut off accounts for a large proportion of the remaining 90%'s traffic, the impact could be significant. So when talking about the resilience of the system, what is an 'acceptable level' of the 'best efforts' service? Are we aiming at having email work 95% of the time to 95% of destinations, or streaming video work 99.99% of the time to 99.99% of destinations? The answer will have an enormous effect on the spare capacity needed! Each extra order of magnitude improvement (say from 99.9% to 99.99%) could cost an order of magnitude more money; yet the benefits of service quality are unevenly distributed. For example, a pensioner who uses the Internet to chat to grandchildren once a week may be happy with 99%, while a company providing a cloud-based business service may need 99.99%.

8.1 Traffic Prioritisation

In a crisis it is common to restrict access to some resources, to shed demand and free up capacity. For telephony a traditional approach is to give emergency services priority. But restricting phone service to 'obvious' emergency workers such as doctors is unsatisfactory. Modern medical practice depends on team working and can be crippled if nurses are cut off; and many patients who depend on home monitoring may have to be hospitalised if communications fail.

If capacity is lost in a disaster and parts of the system are congested, then all users of the congested parts may suffer a reduction in service, and some sorts of traffic (notably VoIP) may stop working entirely. If some types, sources or destinations of traffic are deemed to be important, and so should be given priority in a crisis, then serious thought needs to be given to how to identify priority traffic, how the prioritisation is to be implemented and how turning that prioritisation on and off fits into other disaster planning.

It is not entirely straightforward to identify different types of traffic. So an alternative approach may be to prioritise by source or destination. It may be tempting to consider services such as Facebook or YouTube as essentially trivial, and YouTube uses a lot of bandwidth. However, in a crisis keeping in contact using Facebook may be a priority for many; and shutting down YouTube—thus preventing the free reporting of events—would require solid justification. On the other hand, rate limiting ordinary users, irrespective of traffic type, may appear fair, but could affect essential VoIP use, and cutting off peer-to-peer traffic could be seen as censorship.

Hence, should a crisis occur, it would be inappropriate to expect ISPs to be the organisations to decide to discriminate between different sorts of traffic, or between customers of the same type (although premium customers at premium rates might expect to get better performance). It is not even clear that ISPs are, in general, capable of prioritising traffic on any given basis. So, if some traffic should be prioritised in a crisis, who will make the call, and will anyone be ready to act when they do?

8.2 Traffic Engineering

Traffic engineering is the jargon term for adjusting a network so that traffic flows are improved. In a crisis that would mean shifting traffic away from congested paths. This is less controversial than traffic prioritisation, but no less difficult.

When some event creates congestion in some part(s) of the interconnection system it would be convenient if networks could redirect some traffic away from the congested parts. When a network is damaged its operators will work to relieve congestion within their network by doing internal traffic engineering, adding temporary capacity, repairing things, and so on. One of the strengths of the Internet is that each operator will be working independently to recover their own network as quickly and efficiently as possible.

Where a network's users are affected by congestion in other networks, the simplest strategy is to wait until those networks recover. This may leave spare capacity in other networks unused, so is not the optimum strategy for the system as a whole. But, there are two problems with trying to coordinate action:

1. there is no way of telling where the spare capacity is, and
2. BGP only provides very limited means to influence traffic in other operators' networks.

In effect, if networks attempt to redirect traffic they are blundering around in the dark, attempting to make adjustments to a delicate instrument with a hammer. Their attempts to redirect traffic may create congestion elsewhere, which may cause more networks to try to move traffic around. It is possible to imagine a situation in which many networks are chasing each other, creating waves of congestion and routing changes as they do, like the waves of congestion that pass along roads which are near their carrying capacity.

With luck, if a network cannot handle the traffic it is sent and pushes it away to other networks, it will eventually be diverted towards routes with spare capacity, and after a while all the traffic will have been appropriately redistributed. It is impossible to say how much time would be required for this process; it would depend on the severity of the capacity loss, but it could be days or even weeks.

Strategic local action will not necessarily lead to a socially optimal equilibrium, though, as the incentives may be perverse. Since any SLA will stop at the edge of its network, a transit provider may wish to engineer traffic away from its network in order to meet its SLAs for traffic within its network. The result may still be congestion, somewhere else, but the SLA is still met.

8.3 Routing in a Crisis

Experience shows that in a crisis the interconnection system can quite quickly create new paths between networks to provide interim connections and extra capacity, and we've already discussed how the loss of facilities in New York City on 9/11 was quickly dealt with.

The interconnection ecosystem has often responded in this way, with many people improvising, and working with the people they know personally. This is related to traffic engineering, to the extent that it addresses the lack of capacity by adding extra connections to which traffic can be moved. The response of the system might be improved and speeded up if there were more preparation for this form, and perhaps other forms, of co-operation in a crisis.

In the end, if there is insufficient capacity in a crisis, then no amount of traffic engineering or manual reconfiguration will fit a quart of traffic into a pint of capacity. In extreme cases some form of prioritisation will be needed.

9 Is Transit a Viable Business?

The provision of transit—contracts to carry IP traffic to every possible destination[1]—is a key part of the interconnection system, but it may not be a sustainable business in the near future.

Nobody doubts that the cost of transit has fallen fast, or that it is a commodity business, except where there is little or no competition [56, 73]. In the USA, over the last ten to fifteen years, transit prices have fallen at rate of around 40% per annum—a rate which results in a 99% drop over a ten year period. In other parts of the world prices started higher, but as infrastructure has developed, and transit networks have extended to into new markets, those prices have fallen as well—for example, prices in London are now scarcely distinguishable from those in New York.

Where there is effective competition, the price of transit falls, and consumers benefit. In a competitive market, price tends towards the marginal cost of production. The total cost of production has also fallen sharply, as innovation reduces the cost of the underlying technologies and with increasing economies of scale. Yet every year industry insiders feel that surely nobody can make money at today's prices, and that there must soon be a levelling off. So far there has been no levelling off, though the rate at which prices fall may be diminishing.

The reason is simple: the marginal cost of production for transit service is generally zero. At any given moment there will be a number of transit providers with spare capacity: first, network capacity comes in lumps, so each time capacity is added the increment will generally exceed the immediate need; second, networks are generally over-provisioned, so there is always some

[1] Contrast transit with peering, where an ISP will arrange with another to carry traffic to and from each other's customers. Since peering with every network in the world is impractical, ISPs will purchase transit contracts to carry the traffic that cannot be handled by peering relationships. Note that transit traffic is invariably paid for, whereas peering (which is usually only with other ISPs of a similar size, or between 'content' and 'eyeball' ASs) is generally free.

spare capacity—though eating into that may increase the risk of congestion, perhaps reducing service quality at busy times or when things go wrong.

The logic of this market is that the price for transit will tend towards zero. So it is unclear how pure transit providers could recoup their capital investment. The logic of the market would appear to favour consolidation until the handful of firms left standing acquire 'market power'.

At a practical level, the provision of transit may be undertaken not to make profits but just to offset some of the cost of being an Internet network. For some networks the decision to offer transit at the market price may be increasingly a strategic rather than a commercial decision. Another significant factor is the recent and continuing increase in video traffic and the related rise in the amount of traffic delivered by the content delivery networks (see below). This means that the continued reduction in the unit price of transit is not being matched by an increase in traffic over transit links, so the transit providers' revenues are decreasing.

The acknowledged market leader, Level 3, lost $2.9 billion in 2005–2008, a further $0.6 billion in 2009, and another $0.6 billion in 2010 [47, 48, 49]. It is not possible to say what contribution their transit business made to this; industry insiders note that Level 3 did not go through bankruptcy as some other players did, and would make a small profit were it not for the cost of servicing its debt. However, the industry as a whole is losing large amounts of money.

10 The Rise of the Content Delivery Networks

Over the past four years or so, more and more traffic has been delivered by content delivery networks (CDNs). Their rise has been rapid and has changed the interconnection landscape, concentrating a large proportion of Internet traffic into a small number of networks [44]. This shift has been driven by both cost and quality considerations. With the growth of video content, of ever richer web-sites, and of cloud applications, it makes sense to place copies of popular data closer to the end users who fetch it. This has a number of benefits:

- local connections perform better than remote connections—giving a quicker response and faster transfers.
- costs are reduced because data-files are not being repeatedly transported over large distances—saving on transit costs. However, the key motivation for CDN customers is not to reduce the cost of delivery, but to ensure quality and consistency of delivery—which is particularly important for the delivery of video streams;
- the data-files are replicated, stored in and delivered from a number of locations—improving resilience.

This has moved traffic away from transit providers to peering connections between the CDNs and the end-user's ISP. In some cases content is distributed to servers within the ISP's own network, bypassing the interconnection system altogether.

One CDN claims to deliver some 20% of all Internet traffic [2]. Since the traffic being delivered is the sort which is expected to grow most quickly in the coming years, this implies that an increasing proportion of traffic is being delivered locally, and a reducing proportion of traffic is being carried (over long distances) by the transit providers.

Another effect of this is to add traffic at the Internet Exchange Points (IXPs)[2], which are the obvious way for the CDNs to connect to local ISPs. This adds value to the IXP—particularly welcome for the smaller IXPs, which have been threatened by the ever falling cost of transit (eating into the cost advantage of connecting to the IXP) and the falling cost of connecting to remote (larger) IXPs (where there is more opportunity to pick up peering traffic).

There is a positive effect on resilience, and a negative one. The positive side is that systems serving users in one region are independent of those in other regions, so a lot of traffic becomes less dependent on long distance transit services. On the negative side, CDNs are now carrying so much traffic that if a large one were to fail, transit providers could not meet the added demand, and some services would be degraded. CDNs also concentrate ever more infrastructure in places where there is already a lot of it. If parts of some local infrastructure fail for any reason, will there be sufficient other capacity to fall back on?

Finally, it is possible to count a couple of dozen CDNs quite quickly, but it appears that perhaps two or three are dominant. Some of the large transit providers have entered the business, either with their own infrastructure or in partnership with an existing CDN. There are obvious economies of scale in the CDN business, and there is now a significant investment barrier to entry. The state of this market in a few years' time is impossible to predict, but network effects tend to favour a few, very large, players. These players are very likely to end up handling over half of the Internet's traffic by volume.

11 The 'Insecurity' of the BGP Protocol

A fundamental problem with BGP is that there is no mechanism to verify that the routing information it distributes is valid [12, 39]. The effect of this is felt on a regular basis when a network manages to announce large numbers of routes for addresses that belong to others; this can divert traffic into what

[2] An IXP is a location to which many ISPs (and CDNs) connect. For the cost of a single link, they can exchange traffic (usually peering traffic, but sometimes transit) with all the other ISPs and CDNs who are IXP members.

is effectively a black hole. Such incidents are quite quickly dealt with by network operators, and disruption can be limited to a few hours, at most.

The great fear is that this insecurity might be exploited as a means to deliberately disrupt the Internet, or parts of it. There is also a frequently expressed concern that route hijacking might be used to listen in on traffic, though this can be hard to do in practice.

Configuring BGP routers to filter out invalid routes, or only accept valid ones, is encouraged as best practice [42]. However, where it is practical (at the edges of the Internet) it does not make much difference until most networks do it. Where it would make most immediate difference (in the larger transit providers) it is not really practical because the information on which to base route filters is incomplete and the tools available to manage and implement filters at that scale are inadequate.

More secure forms of BGP, in which routing information can be cryptographically verified, depend on there being a mechanism to verify the 'ownership' of blocks of IP addresses, or to verify that the AS which claims to be the origin of a block of IP addresses is entitled to make that claim. The notion of title to blocks of IP addresses turns out not to be as straightforward as might be expected. However, some progress is now being made, and the Resource Public Key Infrastructure (RPKI) initiative [45] should allow ASs to ignore announcements where the origin is invalid—that is, where some AS is attempting to use IP addresses it is not entitled to use [55]. This is an important step forward, and might tackle over 90% of 'fat finger' problems.

But the cost of RPKI is significant. Every AS must take steps to document their title to their IP addresses, and that title must be registered and attested to by the Internet Registries. Then, every AS must extend their infrastructure to check the route announcements they receive against the register. What is more, the problem that RPKI tackles is, so far, largely a nuisance not a disaster. When some network manages to announce some routes it should not, this is noticed and fixed quite quickly, if it matters. Sometimes a network announces IP addresses nobody else is using—generally they are up to no good, but this does not actually disrupt the interconnection system. So the incentive to do something about the problem is currently weak, although the number of incidents is expected to rise when IPv4 addresses are finally exhausted in late 2011.

Further, a route may pass the checks supported by RPKI, and still be invalid. A network can announce routes for a block of IP addresses, complete with a valid origin, but do so only to disrupt or interfere with the traffic (apparently) on its way to its destination. The S-BGP extensions to BGP (first published in 1997 [40]) try to solve the problem more completely, and there have been revised proposals since [12]; however they make technical assumptions about routing (traffic greed and valley-free customer preferences) that don't hold in today's Internet. Details of a new initiative, BGPSEC, were announced in March 2011. The aim is that this should lead to IETF standards by 2013 and deployed code in routers thereafter [46].

During the standardisation process in 2011–2013 a key issue will be security economics. ASs see the cost of BGP security as high, and the benefit essentially zero until it is very widely deployed.

Ideally, implementation and deployment strategies will give local, incremental benefit, coupled with incentives for early adopters. One possible mechanism is for governments to use their purchasing power to bootstrap adoption; another is for routers to prefer signed routes [31].

Technical issues that must the studied during the standardisation phase include whether more secure BGP might, in fact, be bad for resilience. Adding cryptography to a system can make it brittle. The reason is that when recovering from an event, new and possibly temporary routes may be distributed in order to replace lost routes, and if the unusual routes are rejected, because they do not have the necessary credentials, then recovery will be harder.

Finally, BGPSEC will not be a silver bullet, there are many threats [39], but it should tackle about half of the things that can go wrong after RPKI has dealt with origin validation.

To sum up, most of the time BGP works wonderfully well, but there is plenty of scope to make it more secure and more robust. However, individual networks will get little direct benefit in the beginning from an improved BGP, despite the significant cost. We will probably need some new incentive to persuade networks to invest in more secure BGP, or a proposal for securing BGP that gives local benefits from incremental deployment.

12 Exercises (‘War Games’)

The practical approach to assessing the resilience of the interconnection system is to run large scale exercises in which plausible scenarios are tested. Such exercises have a number of advantages and benefits:

- They start with real world issues. These exercises are not cheap, so there is an incentive to be realistic: planners consider what really are the sorts of event that the system is expected to face.
- They can identify some dependencies on physical infrastructure. By requiring the participants to consider the effects of some infrastructure failure, an exercise may reveal previously unknown dependencies.
- They can identify cross-system dependencies. For example, how well network operations centres can communicate if the phone network fails, or how well field repairs proceed if the mobile phone network is unavailable?
- They exercise disaster recovery systems and procedures. This is generally a good learning experience for everybody involved, particularly as otherwise crisis management is generally ad hoc.

Such scenario testing has been done at a national level [23] and found to be valuable, and recently an exercise at a multi-national scale has also been proved to be valuable.

On 4th November 2010 the European Member States organised the first pan-European cyber exercise, called CYBER EUROPE 2010, which was facilitated by ENISA. The final ENISA report [25] explains the importance of such exercises and calls for future activities based on the lessons learned.

13 The 'Tragedy of the Commons'

The resilience of the Internet interconnection system benefits everyone, but an individual network will not in general gain a net benefit if it increases its costs in order to contribute to the resilience of the whole.

This manifests itself in a number of ways:

- In Section 11 above, we discussed the various proposals for more secure forms of BGP from S-BGP in 1997 to BGPSEC in 2011, none of which has so far been deployed. There is little demand for something which is going to cost money to implement and whose direct benefit is limited.
- There exists best practice for filtering BGP route announcements [42], which, if universally applied, would reduce instances of invalid routes being propagated by BGP and disrupting the system. But these recommendations are difficult to implement and mostly benefit other networks, so are not often implemented.
- There is an IETF BCP[3] for filtering packets, to reduce 'address spoofing', which would mitigate denial of service attacks [28]. These recommendations also mostly benefit others, so are not often implemented.
- A smaller global routing table would reduce the load on all BGP routers in the Internet, and leave more capacity to deal with unusual events. Nevertheless, the routing table is about 75% bigger than it needs to be [37], because some networks announce extra routes to reduce their own costs. Other networks could resist this by ignoring the extra routes, but that would cost time and effort to configure their routers, and would most likely be seen by their customers as a service failure (not as a noble act of public service).
- The system is still ill-prepared for IPv6 [35], despite the now imminent (*c* Q3 2011 [36]) exhaustion of IPv4 address space.

It is in the clear interest of each network to ensure that in normal circumstances 'best efforts' means a high level of service, by adjusting interconnections and routing policy—each network has customers to serve and a

[3] An Internet Engineering Task Force (IETF) Best Common Practice (BCP) is as official as it gets on the Internet.

reputation to maintain. Normal circumstances include the usual day-to-day failures and small incidents.

The central issue is that the security and resilience of the interconnection system is an externality as far as the networks that comprise it are concerned. It is not clear is that there is any incentive for network operators to put significant effort into considering the resilience of the interconnection system under extraordinary circumstances.

14 Regulation

Regulation is viewed with apprehension by the Internet community. Studies such as the one created by ENISA [24] are seen as stalking horses for regulatory interference, which is generally thought likely to be harmful. Despite having its origins in a project funded by DARPA, a US government agency, the Internet has developed since then in an environment that is largely free from regulation. There have been many local attempts at regulatory intervention, most of which are seen as harmful:

- The governments of many less developed countries attempt to censor the Internet, with varying degrees of success. The 'Great Firewall of China' is much discussed, but many other states practice online censorship to a greater or lesser extent. It is not just that censorship itself is contrary to the mores of the Internet community—whose culture is greatly influenced by California, the home of many developers, vendors and service companies. Attempts at censorship can cause collateral damage, as when Pakistan advertised routes for YouTube in an attempt to censor it within their borders, and instead made it unavailable on much of the Internet for several hours.
- Where poor regulation leads to a lack of competition, access to the Internet is limited and relatively expensive. In many less developed countries, a local telecomms monopoly restricts wireline broadband access to urban elites, forcing the majority to rely on mobile access. However the problem is more subtle than 'regulation bad, no regulation good'. In a number of US cities, the diversity of broadband access is falling; cities that used to have three independent infrastructures (say from a phone company, a cable company and an electricity company) may find themselves over time with two, or even just one. In better-regulated developed countries (such as much of Europe) local loop unbundling yields price competition at least, thus mitigating access costs, even if physical diversity is harder. Finally, there are few countries which impose a universal service provision on service providers; its lack can lead to a 'digital divide' between populated areas with broadband provision, and rural areas without.
- There has been continued controversy over surveillance for law-enforcement and intelligence purposes. In the 'Crypto Wars' of the 1990s, the Clinton

administration tried to control cryptography, which the industry saw as threatening not just privacy but the growth of e-commerce and other online services. The Communications Assistance for Law Enforcement Act (CALEA) was passed in 1994 to mandate the co-operation of telecommunications carriers in wiretapping phone calls. The EU has a controversial Data Retention Directive that is up for revision in 2011 and there is interest in both the UK and USA in how wiretapping should be updated for an age not only of VoIP but diverse messaging platforms. This creates conflicts of interest with customers, raises issues of human rights, and leads to arguments about payment and subsidy.
- Governments which worry about Critical National Infrastructure may treat Internet regulation as a matter of National Security, introducing degrees of secrecy and shadowy organisations, which does nothing to dispel concerns about motivation—not helped by a tendency to talk about the problem in apocalyptic terms [74].[4]

Whatever the motivation, government policies are often formulated with insufficient scientific and technical input. They often manage to appear clueless, and in some cases will make things worse. The present study is an attempt to help alleviate this problem.

We have identified a number of areas where the market does not appear to provide incentives to maintain the resilience of the interconnection system at a socially optimal level. However, any attempt to tackle any of the issues by regulation is hampered by a number of factors:

- The lack of good information about the state and behaviour of the system. It is hard to determine how material a given issue may be. It is hard to determine what effect a given initiative is likely to have—good or bad.
- The scale and complexity of the system. Scale may make local initiatives ineffective, while complexity means that it is hard to predict how the system will respond or adapt to a given initiative.
- The dynamic nature of the system. Content delivery networks have been around for many years, but their emergence as a major component of the Internet is relatively recent; this is a testament to the system's ability to adapt quickly (in this case, to the popularity of streamed video).

Up until now, the lack of incentives to provide resilience (and in particular to provide excess capacity) has been relatively unimportant: the Internet has been growing so rapidly that it has been very far from equilibrium, with a huge endowment of surplus capacity during the dotcom boom and significant capacity enhancements due to optical communications technology improvements since then. This cannot go on forever.

One caveat: we must point out that the privatisation, liberalisation and restructuring of utilities worldwide has led to institutional fragmentation in

[4] For a popular perception of the problems that government is grappling with see *Fight Cyber War Before Planes Fall Out of Sky* [79].

a number of critical infrastructure industries which could in theory suffer degradation of reliability and resilience for the same general microeconomic reasons we discuss in the context of the Internet. Yet studies of the electricity, water and telecomms industries in a number of countries have failed to find a deficit thus far [22]. In practice, utilities have managed to cope by a combination of anticipatory risk management and public-private partnerships. However it is sometimes necessary for government to act as a 'lender of last resort'. If a router fails, we can fall back on another router, but if a market fails—as with the California electricity market—there is no fallback other than the state.

In conclusion, it may be some time before regulatory action is called for to protect the resilience of the Internet, but it may well be time to start thinking about what might be involved. Regulating a new technology is hard; an initiative designed to improve today's system may be irrelevant to tomorrow's, or, worse, stifle competition and innovation. For example, the railways steadily improved their efficiency from their inception in the 1840s until regulation started in the late nineteenth century, after which their efficiency declined steadily until competition from road freight arrived in the 1940s [57].

The prudent course of action for regulators today is to start working to understand the Internet interconnection ecosystem [17].

The most important package of work is to increase transparency, by supporting consistent, thorough, investigation of major outages and the publication of the findings, and by supporting long term measurement of network performance. The second package is to fund key research in topics such as distributed intrusion detection and the design of security mechanisms with practical paths to deployment, and the third is to promote best practice, to encourage diverse service provision and to promote the testing of equipment. The fourth package includes the preparation and relationship-building through a series of Private Public Partnerships (PPPs) for resilience. Modest and constructive engagement of this kind will enable regulators to build relationships with industry stakeholders and leave everyone in a much better position to avoid, or delay, difficult and uninformed regulation.

Regulatory intervention must after all be evidence-based; and while there is evidence of a number of issues, the workings of this huge, complex and dynamic system are so poorly understood that there is not enough evidence yet on which to base major regulatory intervention with sufficient confidence.

15 Recommendations

The recommendations come in four groups. The first group is aimed at understanding failures better, so that all may learn the lessons.

Recommendation 1: Incident Investigation

An independent body should thoroughly investigate all major incidents and report publicly on the causes, effects and lessons to be learned. The appropriate framework should be the result of a consultation with the industry and the appropriate regulatory authorities coming into a constructive dialogue. Incident investigation might be undertaken by an industry association, by a national regulator or by a body at the European level, such as ENISA. The last option would require funding to support the work, and, perhaps, powers to obtain information from operators—under suitable safeguards to protect commercially sensitive information.

Recommendation 2: Network Performance Measurement

Europe should promote and support consistent, long-term and comprehensive network performance measurement. At present some realtime monitoring is done by companies such as Arbor Networks and Renesys, and some more is done by academic projects—which tend to languish once their funding runs out. This patchwork is insufficient. There should be sustainable funding to support the long-term collection, processing, storage and publication of performance data. This also has a network management / law enforcement angle in that real-time monitoring of the system could help detect unusual route announcements and other undesirable activity.

The second group of recommendations aims at securing funding for research in topics related to resilience—with an emphasis not just on the design of security mechanisms, but on developing an understanding of how solutions can be deployed in the real world.

Recommendation 3: Research Network Performance and Resilience

Europe should sponsor research into better ways to measure and understand the performance and resilience of huge, multi-layered networks. This is the research aspect of the second recommendation; once that provides access to good data, the data should help clever people to come up with better metrics.

Recommendation 4: Develop and Deploy Secure Interdomain Routing

Europe should support the development of effective, practical mechanisms which have enough incentives for deployment. This may mean mechanisms that give local benefit to the firms that deploy them, even where deployment is incremental; it may require technical mechanisms to be supplemented by policy tools such as the use of public-sector purchasing power, subsidies, liability shifts, or other kinds of regulation.

Recommendation 5: Research into AS Incentives

Europe should support research into economic and legal mechanisms to increase the resilience of the Internet. Perhaps a system of contracts can be constructed to secure the interconnection system, starting with the connections between the major transit providers and spreading from the core to the edges. Alternatively, researchers might consider whether liability rules might have a similar effect. If the failure of a specific type of router caused loss of Internet service leading to damage and loss of life, the Product Liability Directive 85/374/EC would already let victims sue the vendor; but there is no such provision relating to the failure of a service from a transit provider.

The third group of recommendations aims to promote best practice.

Recommendation 6: Sponsor Best Practice

Europe should sponsor and promote best practice in network management. Where best practice exists its adoption may be hampered by practical and economic issues. The public sector may be able to help, but it is not enough to declare for motherhood and apple pie! It can contribute various incentives, such as through its considerable purchasing power. For that to be effective, purchasers need a way to tell good service. The first three of our recommendations can help, but there are some direct measures of quality too.

Recommendation 7: Independently Test Equipment and Protocols

Europe should sponsor the independent testing of routing equipment and protocols. The risk of systemic failure would be reduced by independent testing of equipment and protocols, looking particularly for how well these perform in unusual circumstances, and whether they can be disrupted, suborned, overloaded or corrupted.

Recommendation 8: Regular Disaster Recovery Exercises

The consultation noted that these are effective in improving resilience at local and national levels. European Member States should consider involvement in regular Europe-wide exercises, which ENISA has already expressed a willingness to facilitate. These might provide an umbrella for a number of useful activities, such as investigating what extra preparation might be required to provide more routes in a crisis.

The final group of recommendations aims at engaging policymakers, customers and the public.

Recommendation 9: Contingency Plans for Transit Market Failure

It is possible that the current twenty-odd largest transit providers might consolidate down to a handful[5], in which case they might start to exercise market power and need to be regulated like any other concentrated industry. If this were to happen just as the industry uses up the last of its endowment of dark fibre from the dotcom boom, then prices might rise sharply. European policymakers should start the conversation about what to do then. Action might involve not just a number of European actors but also national regulators from other parts of the world (such as the US Federal Communications Commission). Recommendations 1, 2, 3, and 5 will prepare the ground technically so that regulators will not be working entirely in the dark, but we also need political preparation.

Recommendation 10: Traffic Prioritisation

If, in a crisis, some Internet traffic is to be given priority, and other Internet traffic is to suffer discrimination, then the basis for this choice requires pub-

[5] Just as the final version of the ENISA report was submitted, Level 3 (the transit provider with the largest market share) announced that it was acquiring Global Crossing (the second largest market share) for $3 billion ($1.1 of which is debt assumption), giving the merged company just over half of the market [86].

lic debate—and mechanisms to achieve it need to be developed. Given the number of interests seeking to censor the Internet for various reasons, any decisions on prioritisation will have to be taken openly and transparently, or public confidence will be lost.

Recommendation 11: Greater Transparency

Finally, transparency is not just about openness in taking decisions on regulation or on emergency procedures. It would greatly help resilience if end-users and corporate customers could be educated to understand the issues and send the right market signals. Further investigation is needed on mechanisms that can be developed to give the means to make an informed choice. This might involve combining the outputs from recommendations 2, 3, 5, 6 and 7 into a 'quality certification mark' scheme.

16 Conclusions

The Internet has been remarkably resilient up till now, and has shrugged off quite major incidents such as Hurricane Katrina and 9/11. However our civilisation has come to depend on the Internet, and it is now just as much a critical utility as water or electricity.

We have studied what sort of failures or attacks might cause significant service outages, and we have concluded that while the Internet is in pretty good shape, and does not need regulation in the way that electricity markets are regulated, there are still some things that the policy makers and industry might usefully do to improve its robustness and resilience. These include understanding failures better, funding research in resilience, promoting good practices and engaging with both public and private sectors in a constructive dialogue.

These activities will enable the decision makers in Europe and beyond to understand the Internet better, so that whether regulation is avoided, or is required in five or ten years' time, they will be in a position to propose informed and effective policies.

Acknowledgements We would particularly like to thank all the industry, academic and public sector stakeholders that took part in the survey we conducted as part of the study.

References

1. Agarwal, S., Chuah, C., Bhattacharyya S., Diot C.: The Impact of BGP Dynamics on Intra-Domain Traffic. SIGMETRICS/Performance'04, ACM, pp. 319–330 (2004)
2. Akamai Inc.: Facts & Figures. (2010) http://www.akamai.com/html/about/facts_figures.html#6
3. Alleman, J., Liebenau, J.: Network Resilience and its Regulatory Inhibitors. E. Bohlin et al. (Eds): Global Economy and Digital Society, Elsevier Science Ltd, pp. 379–394 (2004)
4. Alperovitch, D.: U.S.-Based Internet Traffic Redirected to China. McAfee Blog 17 Nov (2010) http://blogs.mcafee.com/mcafee-labs/u-s-based-internet-traffic-redirected-to-china
5. Alperovitch, D.: April Route Hijack: Sifting through the confusion. McAfee Blog 19 Nov (2010). http://blogs.mcafee.com/mcafee-labs/april-route-hijack-sifting-through-the-confusion-2
6. Augustin, B., Krishnamurthy, B., Willinger, W.: IXPs: Mapped? IMC'09, pp. 336–349 (2009)
7. Blumenfeld, L.: Dissertation Could Be Security Threat. Washington Post (2003) http://www.washingtonpost.com/ac2/wp-dyn/A23689-2003Jul7
8. Blunk, L.J.: New BGP analysis tools and a look at the AS9121 Incident. (2005) http://iepg.org/march2005/bgptools+as9121.pdf
9. Bono, V.J.: 7007 Explanation and Apology. NANOG Mailing List (1997). http://www.merit.edu/mail.archives/nanog/1997-04/msg00444.html
10. Brown, M.A.: Pakistan hijacks YouTube. Renesys Blog 24 Feb (2008) http://www.renesys.com/blog/2008/02/pakistan_hijacks_youtube_1.shtml
11. Bush, R., Maennel, O., Roughan, M., Uhlig, S.: Internet Optometry: Assessing the Broken Glasses in Internet Reachability. IMC'09, pp. 242–253 (2009)
12. Butler, K., Farley, T.R., McDaniel, P., Rexford, J.: A Survey of BGP Security Issues and Solutions. Proc IEEE, 98(1), 100–122 (2010)
13. Caesar, M., Rexford, J.: BGP Routing Policies in ISP Networks. IEEE Network, 19(6), 5–11 (2005).
14. Chadd, A.: Murphy's Law Strikes Again: AS7007. Murphy's Law Mailing List, 16 Aug (2006) http://lists.ucc.gu.uwa.edu.au/pipermail/lore/2006-August/000040.html
15. Chang, H., Willinger, W.: Difficulties Measuring the Internet's AS-Level Ecosystem. Information Sciences and Systems, pp. 1479–1483 (2006)
16. Cisco Systems: Cisco Visual Networking Index: Forecast and Methodology, 2009–2014. White paper (2010) http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-481360.pdf
17. Claffy, k.c.: Ten Things Lawyers Should Know About the Internet. (2008) http://www.caida.org/publications/papers/2008/lawyers_top_ten/
18. CNet News Staff: Router glitch cuts Net access. CNet News, 25 Apr (1997) http://news.cnet.com/2100-1033-279235.html
19. Cowie, J.: How To Build A Cybernuke. Renesys Blog, 22 Apr (2010) http://www.renesys.com/blog/2010/04/how-to-build-a-cybernuke.shtml
20. Cowie, J.H., Ogielski, A.T., Premore, B.J., Smith, E.A., Underwood, T.: Impact of the 2003 Blackouts on Internet Communications. Renesys (2003) http://www.renesys.com/tech/reports/Renesys_BlackoutReport.pdf
21. Cowie, J., Popescu, A., Underwood, T.: Impact of Hurricane Katrina on Internet Infrastructure. Renesys (2005). http://www.renesys.com/tech/presentations/pdf/Renesys-Katrina-Report-9sep2005.pdf
22. De Bruijne, M.: Systems that Should Have Failed: Critical Infrastructure Protection in an Institutionally Fragmented Environment. J Contingencies Crisis Management, 15(1), 1–29 (2007)

23. ENISA: Good Practice Guide on National Exercises. ENISA Technical Report (2009) http://www.enisa.europa.eu/act/res/policies/good-practices-1/exercises/national-exercise-good-practice-guide
24. ENISA: Inter-X: Resilience of the Internet Interconnection Ecosystem. ENISA Technical Report (2011) Available at: http://www.enisa.europa.eu/act/res/other-areas/inter-x.
25. ENISA: Cyber Europe 2010 — Evaluation Report. ENISA Technical Report (2011) Available at: http://www.enisa.europa.eu/act/res/cyber-europe-2010/cyber-europe-2010-report
26. Faratin, P., Clark, D.D., Bauer, S., Lehr, W., Gilmore, P.W., Berger, A.: The Growing Complexity of Internet Interconnection. Communications Strategies, 72, pp. 51–71 (2008)
27. Feamster, N., Winick, J., Rexford, J.: A Model of BGP Routing for Network Engineering. SIGMETRICS/Performance'04, ACM, pp. 331–342 (2004)
28. Ferguson, P., Senie, D.: Network Ingress Filtering: Defeating Denial of Service Attacks which Employ IP Source Address Spoofing. RFC 2827, RFC Editor (2000)
29. Fried, M., Klemming, L.: Severed Cables in Mediterranean Disrupt Communication. Bloomberg, 19 Dec (2008) http://www.bloomberg.com/apps/news?pid=newsarchive&sid=aBa0lTN.dcoQ
30. Gill, P., Arlitt, M., Li, Z., Mahanti, A.: The Flattening Internet Topology: Natural Evolution, Unsightly Barnacles or Contrived Collapse? Ninth Passive and Active Measurement Conference (PAM) (2008)
31. Gill, P., Schapira, M., Goldberg, S.: Let the Market Drive Deployment: A Strategy for Transitioning to BGP Security. Boston University Computer Science Technical Report BUCS-TR-2011-003 (2011) http://www.cs.bu.edu/~goldbe/papers/sbgpTrans.html
32. Griffin, T.G., Premore, B.J.: An Experimental Analysis of BGP Convergence Time. ICNP'01, pp. 53–61 (2001)
33. Griffin, T.G., Shepherd, F.B., Wilfong, G.: The Stable Paths Problem and Interdomain Routing. IEEE ACM T Network, 10(2), 232–243 (2002)
34. Hollnagel, E., Woods, D.D., Leveson, N.: Resilience Engineering: Concepts and Precepts. Ashgate Publishing (2006)
35. Huston, G.: The ISP Column: Is the Transition to IPv6 a "Market Failure"? (2009) http://www.potaroo.net/ispcol/2009-09/v6trans.pdf
36. Huston, G.: IPv4 Address Report. http://www.potaroo.net/tools/ipv4/
37. Huston, G., et al.: CIDR Report. http://www.cidr-report.org
38. ISC: The ISC Domain Survey. http://www.isc.org/solutions/survey
39. Kent, S.: Threat Model for BGP Path Security. Working Draft (2011) http://tools.ietf.org/html/draft-kent-bgpsec-threats-01
40. Kent, S., Lynn, C., Seo, K.: Secure Border Gateway Protocol (S-BGP). IEEE J Sel Area Comm, 18(4), 582–592 (2000).
41. Kitano, H.: Systems Biology: A Brief Overview. Science, 295, 1662–1664 (2002)
42. Kuhn, R., Sriram, K., Montgomery, D.: Border Gateway Protocol Security — Recommendations of the National Institute of Standards and Technology. NIST (2007) http://csrc.nist.gov/publications/nistpubs/800-54/SP800-54.pdf
43. Kushman, N., Kandula, S., Katabi, D.: Can You Hear Me Now?! It Must Be BGP. SIGCOMM Comput Commun Rev, 37, 75–84 (2007)
44. Labovitz, C., Iekel-Johnson, S., McPherson, D., Oberheide, J., Jahanian, F.: Internet Inter-Domain Traffic. SIGCOMM'10, ACM, pp. 75–86 (2010)
45. Lepinski, M., Kent, S.: An Infrastructure to Support Secure Internet Routing. Working Draft (2011) http://tools.ietf.org/html/draft-ietf-sidr-arch-12
46. Lepinski, M., Turner, S.: An Overview of BGPSEC. Working Draft (2011) http://tools.ietf.org/html/draft-lepinski-bgpsec-overview-00

47. Level 3 Communications, Inc.: FORM 10-K For the fiscal year ended December 31, 2006 (amended). http://lvlt.client.shareholder.com/secfiling.cfm?filingID=1104659-07-88481
48. Level 3 Communications, Inc.: FORM 10-K For the fiscal year ended December 31, 2008. http://lvlt.client.shareholder.com/secfiling.cfm?filingID=1047469-09-2002
49. Level 3 Communications, Inc.: FORM 10-K For the fiscal year ended December 31, 2010. http://lvlt.client.shareholder.com/secfiling.cfm?filingID=1047469-11-1410
50. Leyden, J.: China routing snafu briefly mangles interweb — Cockup, not conspiracy. The Register, 9 Apr (2010) http://www.theregister.co.uk/2010/04/09/china_bgp_interweb_snafu/
51. Longstaff, P.: Security, Resilience, and Communication in Unpredictable Environments Such as Terrorism, Natural Disasters and Complex Technology. (2005) http://pirp.harvard.edu/pubs_pdf/longsta/longsta-p05-3.pdf
52. Magnuson, S.: Cyber Experts Have Proof That China Has Hijacked U.S.-Based Internet Traffic: UPDATED. National Defense Blog, 12 Nov (2010) http://www.nationaldefensemagazine.org/blog/Lists/Posts/Post.aspx?ID=249
53. McCullagh, D.: How Pakistan knocked YouTube offline (and how to make sure it never happens again). CNet, 25 Feb (2008) http://news.cnet.com/8301-10784_3-9878655-7.html
54. McPherson, D.: Internet Routing Insecurity :: Pakistan Nukes YouTube? Arbor Networks Blog, 25 Feb (2008) http://asert.arbornetworks.com/2008/02/internet-routing-insecuritypakistan-nukes-youtube/
55. Mohapatra, P., Scudder, J., Bush, R., Austein, R. (Eds.): BGP Prefix Origin Validation. Working Draft (2011) http://tools.ietf.org/html/draft-ietf-sidr-pfx-validate-01
56. Norton, B.: Internet Transit Prices — Historical and Projected. (2010) http://drpeering.net/white-papers/Internet-Transit-Pricing-Historical-And-Projected.php
57. Odlyzko, A.: Collective hallucinations and inefficient markets: the British railway mania of the 1840s. (2010) http://www.dtc.umn.edu/~odlyzko/doc/hallucinations.pdf
58. Odlyzko, A.: Minnesota Internet Traffic Studies (MINTS). http://www.dtc.umn.edu/mints/
59. Oliveira, R., Pei, D., Willinger, W., Zhang, B., Zhang, L.: The (In)Completeness of the Observed Internet AS-level Structure. IEEE ACM T Network, 18, 109–112 (2010)
60. Partridge, C., et al.: The Internet under Crisis Conditions: Learning from September 11. The National Academies Press, Washington (2002)
61. Popescu, A.: Deja Vu All Over Again: Cables Cut in the Mediterranean. Renesys Blog, 19 Dec (2008) http://www.renesys.com/blog/2008/12/deja-vu-all-over-again-cables.shtml
62. Popescu, A.C., Premore, B.J., Underwood, T.: The Anatomy of a Leak: AS9121. NANOG 34 (2005) http://www.renesys.com/tech/presentations/pdf/renesys-nanog34.pdf
63. Quoitin, B., Pelsser, C., Swinnen, L., Bonaventure, O., Uhlig, S.: Interdomain Traffic Engineering with BGP. IEEE Comm Mag, 41(5), 122–128 (2003)
64. Radovcic, S.: European Internet Exchange Association 2010 Report on European IXPs. (2010) https://www.euro-ix.net/resources/reports/euro-ix_report_2010.pdf
65. Rekhter, Y., Li, T., Hares, S.: A Border Gateway Protocol 4 (BGP-4). RFC4271, RFC Editor (2006)
66. RIPE NCC: YouTube Hijacking: A RIPE NCC RIS case study. (2008) http://www.ripe.net/news/study-youtube-hijacking.html

67. Romijn, E.: RIPE NCC and Duke University BGP Experiment. (2010) http://labs.ripe.net/Members/erik/ripe-ncc-and-duke-university-bgp-experiment
68. Romijn, E.: Re: Did your BGP crash today? NANOG Mailing List, 27 Aug (2010) http://www.merit.edu/mail.archives/nanog/msg11505.html
69. Siganos, G., Faloutsos, M.: Analysing BGP Policies: Methodology and Tool. INFOCOM'04, IEEE, pp. 1640–1651 (2004)
70. Smith, D.J.: Reliability, Maintainability and Risk. 7th edn, Elsevier (2005)
71. Soares, M.: Brazilian Blackout Traced to Sooty Insulators, Not Hackers. Wired (2009) http://www.wired.com/threatlevel/2009/11/brazil_blackout/
72. Telegeography: Four international cable breaks in a week. 4 Feb (2008) http://www.telegeography.com/cu/article.php?article_id=21567
73. Telegeography: IP transit prices continue their downward trend. 16 Nov (2010) http://www.telegeography.com/cu/article.php?article_id=35206
74. UK Cabinet Office: Cyber Security. Fact Sheet 18 (2010) http://download.cabinetoffice.gov.uk/sdsr/factsheet18-cyber-security.pdf
75. Underwood, T.: Internet-Wide Catastrophe — Last Year. Renesys Blog, 24 Dec (2005) http://www.renesys.com/blog/2005/12/internetwide_nearcatastrophela.shtml
76. Wagner, A.: Robustness and Evolvability: A Paradox Resolved. ProcBiolSci, 275, 91–100 (2008)
77. Wang, F., Mao, Z.M., Wang, J., Gao, L., Bush, R.: A Measurement Study on the Impact of Routing Events on End-to-End Internet Path Performance. SIGCOMM'06, ACM, pp. 375–386 (2006)
78. Wang, L., Zhao, X., Pei, D., Bush, R., Massey, D., Mankin, A., Wu, S.F., Zhang, L.: Observation and Analysis of BGP Behavior under Stress. IMW'02, ACM, pp. 183–195 (2002)
79. Wilson, G.: Fight Cyber War Before Planes Fall Out of Sky. The Sun, 19 Oct (2010) http://www.thesun.co.uk/sol/homepage/news/3186185/Security-chiefs-warn-Britain-must-protect-itself-against-cyber-warfare-amid-government-cuts.html
80. Xie, L., Smith, P., Banfield, M., Leopold, H., Sterbenz, J., Hutchinson, D.: Towards Resilient Networks using Programmable Networking Technologies. IFIP IWAN 2005 (2005)
81. Yang, Y.R., Xie, H., Wang, H., Silberschatz, A., Kroshnamurthy, A., Liu, Y., Li, L.E.: On Route Selection for Interdomain Traffic Engineering. IEEE Network, 19(6), 20–27 (2005)
82. Yannuzzi, M., Masip-Bruin, X., Bonaventure, O.: Open Issues in Interdomain Routing: A Survey. IEEE Network, 19(6), 49–56 (2005)
83. Zmijewski, E.: Mediterranean Cable Break—Part II. Renesys Blog, 31 Jan (2008) http://www.renesys.com/blog/2008/01/mediterranean-cable-break-part-1.shtml
84. Zmijewski, E.: Mediterranean Cable Break—Part III. Renesys Blog, 2 Feb (2008) http://www.renesys.com/blog/2008/02/mediterranean-cable-break-part.shtml
85. Zmijewski, E.: Mediterranean Cable Break—Part IV. Renesys Blog, 7 Feb (2008) http://www.renesys.com/blog/2008/02/mediterranean-cable-break-part-3.shtml
86. Zmijewski, E.: A Closer Look at the 'Level 3 + Global Crossing' Union. CircleID, 14 Apr (2011) http://www.circleid.com/posts/20110414_a_closer_look_at_the_level_3_and_global_crossing_union/