Inter-X: Resilience of the Internet Interconnection Ecosystem

Summary Report – April 2011

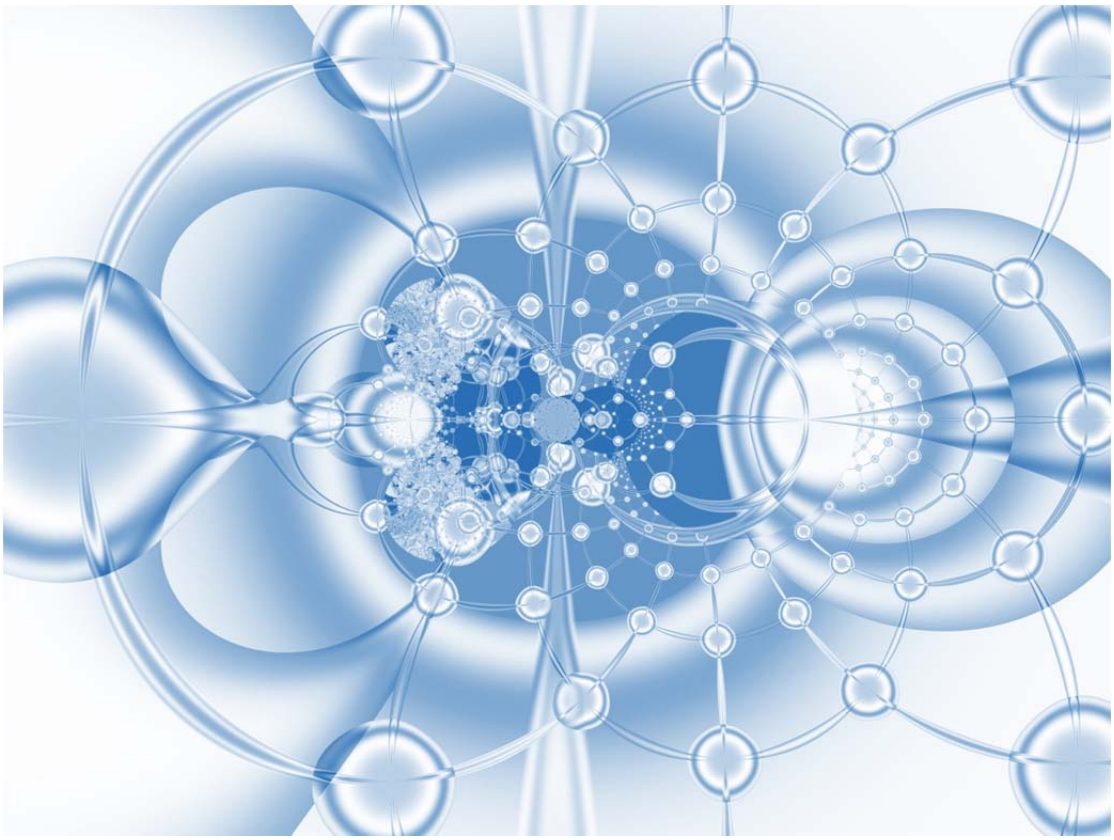

enisa
European Network and Information Security Agency

About ENISA

The European Network and Information Security Agency (ENISA) is an EU agency created to advance the functioning of the internal market. ENISA is a centre of expertise for the European Member States and European institutions in network and information security, giving advice and recommendations and acting as a switchboard for information on good practices. Moreover, the agency facilitates contacts between European institutions, the Member States, and private business and industry actors. Internet: *http://www.enisa.europa.eu/*

Acknowledgments:

While compiling this report, we talked extensively over a period of many months to a large number of technical and managerial staff at communications service providers, vendors, and service users. Many of our sources requested that we not acknowledge their contribution. Nonetheless we thank them all here. ENISA would like to express its gratitude to the stakeholders that provided input to the survey.

Editor: Panagiotis Trimintzios, ENISA

Authors:

Chris Hall, Highwayman Associates

Richard Clayton, Cambridge University

Ross Anderson, Cambridge University

Evangelos Ouzounis, ENISA

Contact

For more information about this study, please contact:

Dr. Panagiotis Trimintzios

panagiotis.trimintzios@enisa.europa.eu

Internet: http://www.enisa.europa.eu/act/res

18-Apr-2011 (b)

Table of Contents

Executive Summary

The Internet has so far been extremely resilient. Even major disasters, such as 9/11 and Hurricane Katrina, have had only a local impact. Technical failures have lasted only a few hours, and congestion has had a sustained effect only where the infrastructure is inadequate. The low cost and general reliability of communications over the Internet have led more and more systems to depend on it; we are now at the point where a systemic failure would not just disrupt email and the web, but cause significant problems for other utilities, transport, finance, healthcare and the economy generally. So the continued resilience of the Internet is critical to the functioning of modern societies, and hence it is right and proper to examine whether the mechanisms that have such an excellent track record in providing a resilient Internet are likely to continue to be as effective in the future.

The focus of this report is the 'Internet interconnection ecosystem'. This holds together all the networks that make up the Internet. The ecosystem is complex and has many interdependent layers. This system of connections between networks occupies a space between and beyond those networks and its operation is governed by their collective self-interest – the Internet has no central Network Operation Centre, staffed with technicians who can leap into action when trouble occurs. The open and decentralised organisation that is the very essence of the ecosystem is essential to the success and resilience of the Internet. Yet there are a number of concerns.

First, the Internet is vulnerable to various kinds of common mode technical failures where systems are disrupted in many places simultaneously; service could be substantially disrupted by failures of other utilities, particularly the electricity supply; a flu pandemic could cause the people on whose work it depends to stay at home, just as demand for home working by others was peaking; and finally, because of its open nature, the Internet is at risk of intentionally disruptive attacks.

Second, there are concerns about sustainability of the current business models. Internet service is cheap, and becoming rapidly cheaper, because the costs of service provision are mostly fixed costs; the marginal costs are low, so competition forces prices ever downwards. Some of the largest operators – the 'Tier 1' transit providers – are losing substantial amounts of money, and it is not clear how future capital investment will be financed. There is a risk that consolidation might reduce the current twenty-odd providers to a handful, at which point they would start to acquire pricing power and the regulation of transit service provision might become necessary as in other concentrated industries.

Third, dependability and economics interact in potentially pernicious ways. Most of the things that service providers can do to make the Internet more resilient, from having excess capacity to route filtering, benefit other providers much more than the firm that pays for them, leading to a potential 'tragedy of the commons'. Similarly, security mechanisms that would help reduce the likelihood and the impact of malice, error and mischance are not implemented because no-one has found a way to roll them out that gives sufficiently incremental and sufficiently local benefit.

Fourth, there is remarkably little reliable information about the size and shape of the Internet infrastructure or its daily operation. This hinders any attempt to assess its resilience in general and the analysis of the true impact of incidents in particular. The opacity also hinders research and development of improved protocols, systems and practices by making it hard to know what the issues really are and harder yet to test proposed solutions.

So there may be significant troubles ahead which could present a real threat to economic and social welfare and lead to pressure for regulators to act. Yet despite the origin of the Internet in DARPA-funded research, the more recent history of government interaction with the Internet has been unhappy. Various governments have made ham-fisted attempts to impose censorship or surveillance, while others have defended local telecommunications monopolies or have propped up other industries that were disrupted by the Internet. As a result, Internet service providers, whose good will is essential for effective regulation, have little confidence in the likely effectiveness of state action, and many would expect it to make things worse.

Any policy should therefore proceed with caution. At this stage, there are four types of activity that can be useful at the European (and indeed the global) level.

The first is to understand failures better, so that all may learn the lessons. This means consistent, thorough, investigation of major outages and the publication of the findings. It also means understanding the nature of success better, by supporting long term measurement of network performance, and by sustaining research in network performance.

The second is to fund key research in topics such as inter-domain routing – with an emphasis not just on the design of security mechanisms, but also on traffic engineering, traffic redirection and prioritisation, especially during a crisis, and developing an understanding of how solutions are to be deployed in the real world.

The third is to promote good practices. Diverse service provision can be encouraged by explicit terms in public sector contracts, and by auditing practices that draw attention to reliance on systems that lack diversity. There is also a useful role in promoting the independent testing of equipment and protocols.

The fourth is public engagement. Greater transparency may help Internet users to be more discerning customers, creating incentives for improvement, and the public should be engaged in discussions on potentially controversial issues such as traffic prioritisation in an emergency. And finally, Private Public Partnerships (PPPs) of relevant stakeholders, operators, vendors, public actors etc is important for self-regulation. In this way even if regulation of the Internet interconnection system is ever needed after many years, policy makers will be able to make informed decisions leading to effective policies.

The objective of these activities should be to ensure that when global problems do arise, the decision and policy makers have a clear understanding of the problems and of the options for action.

There are local regulatory actions that Europe can encourage where needed. Poor telecommunications regulation can lead to the consolidation of local service provision so that cities have fewer independent infrastructures; and in countries that are recipients of EU aid, telecommunications monopolies often deepen the digital divide.

The aim of all these activities should be to ensure that the Internet is ubiquitous and resilient, with service provided by multiple independent competing firms who have the incentives to provide a prudent level of capacity not just for fair weather, but for when the storms arrive.

Introduction to the Summary Report

This study looks at the resilience of the Internet interconnection ecosystem. The Internet is a network of networks, and the interconnection ecosystem is the collection of layered systems that holds it together. The interconnection ecosystem is the core of the Internet, providing the basic function of reaching anywhere from everywhere.

The Executive Summary above provides an abstract of the report's subject and broad recommendations.

The Full Report has four parts:

Part I Summary and Recommendations

This contains a more extended examination of the subject and a discussion of our recommendations in detail, followed by the recommendations themselves.

This part of the report is based on the parts which follow.

Part II State of the Art Review

This includes a detailed description of the Internet's routing mechanisms and analysis of their robustness at the technical, economic and policy levels.

The material in this part supports the analysis presented in Part I, and sets out to explain how and why the issues and challenges the report identifies come about.

Part III Report on the Consultation

As part of the study a broad range of stakeholders were consulted. This part reports on the consultation and summarises the results.

Part IV Bibliography and Appendices

There is an extensive bibliography and summaries of the financial statements of some of the major transit providers.

This Summary Report is Part I of the Full Report.

Two sections follow:

- Section 1 is a summary of the issues and challenges. It is intended to be read as an introduction to the recommendations, giving the background and the rationale for them. It serves also as an introduction to the rest of the Full Report.
- Section 2 contains our recommendations.

In the following, section number references to Sections 3 onwards refer to Part II the Full Report. References of the form **[C:xx]** refer to general points made in the consultation, while those of the form **[Q:xx]** refer to quotations from the consultation which made a particular, or a particularly apposite, point – those references point to Part III of the Full Report. References of the form [**1**] refer to the Bibliography, which is in Part IV of the Full Report.

This revised version of the report replaces the version published in December 2010.

1 Summary

The Internet has been pretty reliable so far, having recovered rapidly from most known incidents. The effects of natural disasters such as Hurricane Katrina, terrorist attacks such as 9/11 and assorted technical failures have all been limited in time and space. However it does appear likely that the Internet could suffer systemic failure, leading perhaps to local failures and system-wide congestion, in some circumstances including:

- A regional failure of the physical infrastructure on which it 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 BGP fabric by broadcasting thousands of bogus routes, either via a large AS or from a large number of compromised routers.

There is evidence that implementations of the Border Gateway Protocol (BGP) are surprisingly fragile. There is evidence that some concentrations of infrastructure are vulnerable and significant disruption can be caused by localised failure. There is evidence that the health of the interconnection system as a whole is not high among the concerns of 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 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 an equilibrium: the rapid growth in 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 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.

This section proceeds as follows:

- in Section 1.1 the challenges posed by the sheer scale and complexity of the Internet interconnection system are discussed.
- the nature of resilience and the difficulty of assessing it are discussed in Section 1.2.
- Section 1.3 discusses the information that we do not have, and how that limits our ability to address the issue of resilience, among other things.
- resilience and efficiency are antipathetic, which raises the challenges given in Section 1.4.
- the problems posed by the reliability of equipment, and the possibility for systemic failure are covered in Section 1.5.
- Section 1.6 examines the value of Service Level Agreements in the context of the interconnection system.

- all parts of the Internet must be able to reach all other parts, so 'reachability' is a key objective. However, being able to reach a destination does not guarantee that traffic will flow to and from there effectively and that expected levels of performance will be met. Section 1.7 discusses the challenges, with particular reference to the behaviour of the system if some event has disabled parts of it.
- every year the price of transit goes down, and every year people feel it must level off. The reason to believe that the price will tend to zero, and the challenges that poses are discussed in Section 1.8.
- the rise of the Content Delivery Networks (CDNs) and the effect on the interconnection system is discussed in Section 1.9.
- Section 1.10 tackles the insecurity of BGP.
- in Section 1.11 the value of disaster recovery exercises ("war games") is examined.
- a number of issues are related; tackling them would benefit everybody, but addressing them also costs each network more than they gain individually. This is discussed in Section 1.12.
- the contentious subject of regulation is raised in Section 1.13.

1.1 Scale and Complexity

The Internet is very big and very complicated **[C:1]**.

The interconnection system we call the Internet comprises some 37,000 'Autonomous Systems' or ASes (ISPs or similar entities) and 355,000 blocks of addresses (addressable groups of machines), spread around the world – as of March 2011 (see Section 3 of the Full Report).

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 a small fraction of the system may still have a significant impact on a great many people. When considering the resilience of this system it is necessary to consider not only the global issues, but a large number of separate, but interconnected, local issues.

The complexity of the system is partly related to its sheer scale, and the number of interconnections between ASes. This is compounded by a number of factors.

- Modelling the interconnection system is hard because we only have partial views of it and because it has a number of layers, each with its own properties and interacting with other layers. For example, the connections between ASes 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 **[C:7]**.

  While it is possible to discover part of the 'AS-level topology' of the Internet (which ASes are interconnected), from a resilience perspective, it would be more valuable to know the 'router-level topology', (the number, location, capacity, traffic levels etc. of the actual connections between ASes) **[C:2]**. 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 **[C:3]**. 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 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 electricity supply – and those systems depend in turn on the Internet. **[C:8]**, **[Q:3]** and **[Q:17]**.
- The interconnection ecosystem is self-organising and highly decentralised. The decision whether to interconnect is made independently by the ASes, driven by their need to be able 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.
- The interconnection ecosystem is dynamic and constantly changing. 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) changed the pattern of interconnections.
- 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 traffic in the early-to-mid 2000s, now video traffic of various kinds is coming to dominate both in terms of volume and in terms of growth.
- The Internet is continuing to grow. In fact, just about everything about it continues to grow: the number of ASes, the number of routes, the number of interconnections, the volume of traffic, etc.

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

This study attempts to provide some insight by describing the workings of the system and what we know about its resilience.

1.2 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 [**1**]. 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 [**2**]. 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 [**3**] [**4**].

Resilience of a system is defined as *the ability to provide and maintain an acceptable level of service in the face of various faults and challenges to normal operation*[1]. That is the ability to adapt itself to recover from a serious failure, or more generally to its ability 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 will improve and at some time full service and the system will be 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 collocation 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 equipment failures – and reconfigure itself essentially automatically and over a time scale of seconds to minutes. 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 a 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 of the system 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 as resilience.

This study is 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. It is thus biased toward 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 no resilient response is required. While resilience is to do with coping with the impact of events, robustness is to do with reducing the impact in the first place. The two overlap, and from the users' perspective these are fine distinctions; what the user wants is for the system to be predictably dependable.

[1] *following: James P.G. Sterbenz, David Hutchison, Egemen K. Çetinkaya, Abdul Jabbar, Justin P. Rohrer, Marcus Schöller and Paul Smith: "Resilience and survivability in communication networks: Strategies, principles, and survey of disciplines", Computer Networks, Volume 54, Issue 8, 1 June 2010, Pages 1245-1265, Resilient and Survivable networks.*

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. **[C:5]**

Moreover, no good metrics are available to actually assess the performance of the Internet or its interconnection system. This makes it harder still to specify acceptable levels of service. For the Internet the problem is compounded by its scale and complexity (see above) and by lack of information (see below), which make 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 system.

Given these difficulties (and there are many more), service providers commonly fall back on measures that improve resilience in general terms, in the hope 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.

1.3 The Lack of Information

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

In fact it is worse than that. ASes 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 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 ASes – 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 the workings of the Internet, and to develop and test improvements; it makes the study and modelling of complex emergent properties such as resilience harder still. **[C:2]**, **[Q:1]** and **[Q:2]**.

The lack of information has a number of causes:

- **Complexity and scale.** To map the networks of fibre around the world might be a tractable problem. Over those physical fibres run many different logical connections, 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, not all those things are fixed – providers

reroute existing networks and circuits as they extend or adapt their networks. To keep track, meticulous record keeping is required, but even within a single AS it is not always achieved. At a global level, measuring traffic volumes would be an immense undertaking, given the sheer number of connections between networks.

- **The information hiding properties of the routing system.** When trying to map connections by probing the system from the outside, each probe will reveal something about the path between two points in the Internet at the time of the probe. But the probe reveals little about what other paths may exist at other times, or what path might be taken if any part of the usual path is not working, or what the performance of those other paths might be.
- **Security concerns.** Mapping the physical layer is thought to be an invitation to people with bad intentions to improve their target selection so those maps that do exist are seldom shared.
- **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 information, and possibly in anonymised and aggregated form. **[C:23]**
- **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).
- **Lack of good metrics.** While there are some 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, a network of networks. ENISA has already started working in this direction, looking at resilience metrics from a holistic point of view[2].

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.

1.3.1 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 know about its resilience. **[C:4]**. For example, we know the following.

- It is straightforward to divert traffic away from its proper destination by announcing invalid routes. The well-known incident in February 2008 in which YouTube stopped working for a few hours is one example; see Section 5.6.4. More publicity, and political concern, was raised

[2] http://www.enisa.europa.eu/act/res/other-areas/metrics

by a 2010 incident in which China Telecom advertised a number of invalid routes, effectively hijacking 15% of Internet addresses for 18 minutes; see Section 5.6.9.

- Latent bugs in BGP implementations can disrupt the system. Most recently, in August 2010, an experiment that 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. The effects of this incident lasted less than two hours; see Section 5.6.5.
- 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. 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. This seriously affected traffic for perhaps two weeks. See Section 5.6.6.
- The system is critically dependent on electrical power. A large power outage in Brazil in November 2009 caused significant disruption, though it lasted only four and a half hours; see Section 5.6.6. Interestingly, previous blackouts in Brazil had been attributed to 'hackers', suggesting that these incidents are examples of the risk of inter-dependent networks. This particular conspiracy theory has been refuted.
- 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 11th 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). See Section 5.6.10.

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*", [**5**] 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.*"

1.4 Resilience and Efficiency

There are fundamental tensions between resilience and efficiency. **[Q:5]** Resilience requires spare capacity and duplication of resources, and systems which are loosely coupled (made up of largely independent sub-systems) are more resilient than tightly coupled systems whose components depend more on each other. But improving the efficiency of a system generally means eliminating excess capacity and redundant resources.

A more diverse system is generally a more resilient one, but diversity adds to cost and complexity. 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 **[C:7] [Q:9]**. 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 equipment used **[C:9]**. 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 **[C:6]**.

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 in other areas, past performance is no guarantee of future results.

1.5 Resilience and Equipment

A particular concern for the interconnection system is the possibility of an internal technical problem that could have 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 equipment vendors for this kind of equipment. **[C:9]** There have been a number of incidents in which large numbers of routers across the entire Internet have been affected by the same problem, something unprecedented and unexpected which exposes a bug in the software, and occasionally in the specification of BGP.

No software is free from bugs, but the universal dependence on BGP makes bugs there 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 as part of their development process. But the interests of both ISPs and manufacturers are for the equipment to work well under normal circumstances. Individual ISPs cannot afford to do exhaustive testing of low-probability scenarios for the benefit of the Internet at large. The manufacturers for their part balance the effort and time spent testing against their customers' demands for new and useful features, new and faster routers and less expensive software. Also of concern is how secure routers and routing protocols are against deliberate attempts to disrupt or suborn them.

A number of respondents to the consultation felt that money spent on testing equipment and protocols would be money well spent. **[C:10]**

1.6 Service Level Agreements (SLAs) 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 **[C:20]**. In particular, where an operator offers an SLA, it does not extend beyond the borders of their network **[C:19]**; so whatever their guarantees are, they do not cover the interconnection system – the part between the borders of all networks.

The SLAs offered to end-customers by their ISPs reflect the SLAs that ISPs obtain from their transit providers and peers. The standard SLAs offered to end-customers may be published, but the SLAs offered between networks may be part of contracts that are kept confidential. Given how little such SLAs are generally thought to cover, it is an open question how much information is being hidden

here – but it is another aspect of the general lack of information about the ecosystem at all levels. (The consultation asked specifically about inter-provider agreements, see Section 9, Question 8.)

Providers do not attempt to guarantee anything beyond their borders because they cannot. Any such 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, that is, 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, will reduce demand if it detects congestion – it is designed to adapt to the available capacity, not to guarantee some level of performance.

The other difficulty with SLAs is what can and what 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, or by whom. Consider a connection from *a* in one network to *b* in another network, which traverses four other networks and the connections between them:

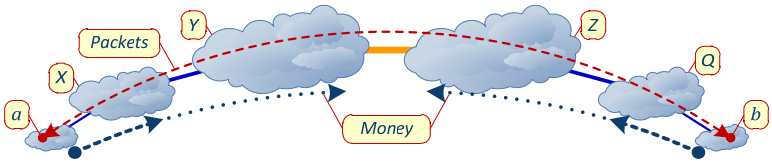

Figure 1: Connection between *a* and *b*

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 and indirectly for packets to and from the connection between networks *Y* and *Z*. Similarly, *b* is paying for packets to and from the mid-point on the other side. If network *Q* has low standards, or is having a bad day, to whom does *a* complain? Network *X* has a contract with *a*'s network, and offers an SLA, but that does not extend beyond *X*. Network *Y* has a contract with *X*, with a different SLA, but even if *X* complained to *Y* about its customer's problem we have come to the end of the money trail: *Y* cannot hold *Z* to account for the performance of *Q*. Suppose *a* were to demand a strong SLA from their provider: *X* certainly has no way of imposing some standard of service on *Q*, and simply cannot offer to make any guarantee.

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 *a*'s network and the rest of the Internet. The problem is intractable. So whatever value SLAs have, they do not offer a contractual framework through which customers can influence the resilience of the interconnection system, even if they wanted to. In addition, few customers understand the issue, or care to do anything about it. Generally the Internet is remarkably reliable, so customers' principal interest in choosing a supplier is price – possibly moderated by the suppliers' reputation. **[C:18]**

1.7 Reachability, Traffic and Performance

While end-users care about traffic and performance, the basic mechanism of the interconnection system – BGP – only understands reachability **[Q:11]**. 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 ASes (the ISPs and other networks that make up 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. **[Q:16]**

The problems with 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. These security issues are discussed separately in Section 1.10.
- there is no mechanism in BGP to convey capacity information – so BGP cannot help reconfigure the interconnection system to avoid congestion. **[Q:12]** When a route fails, BGP will find another route to maintain reachability, but that route may not have sufficient 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.
- when things change BGP can be slow to settle down ('converge') to a new, stable state. **[C:12]**
- the ability of BGP to cope or cope well under extreme conditions is not assured.

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 ensure that the result will perform adequately, and take other steps if it does 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. 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 on the interconnection system, 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 and streaming video. Congestion may stop these applications working satisfactorily, or at all.

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 that 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 would prefer to know the amount, and possibly the value, of traffic that is affected. 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% to 99.9%) 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% or even 90%, while a company providing a cloud-based business service may need 99.99% or more.

1.7.1 Traffic Prioritisation

In a crisis it is common for access to some resources to be restricted, 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 will suffer a reduction in service, and some types of traffic (notably VoIP) may stop working effectively. 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. **[Q:19]**

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. Moreover, shutting down YouTube in a crisis – thereby 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.

So it is inappropriate for ISPs 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 in a crisis). **[Q:21]** It is not even clear that ISPs are, in general, capable of

prioritising some 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?

It is clear that this challenge entails both technical and policy aspects. The former are related mainly to the mechanisms that should exist in network equipment to support traffic prioritisation. The latter refer mainly to the policies that specify what traffic should be given priority. It is very important to tackle both aspects of the problem.

1.7.2 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 its 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. However, there are two problems with trying to coordinate action:

1. there is no way of telling where the spare capacity in the system is;
2. BGP 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 be diverted towards spare capacity elsewhere. Given enough time the system would adapt to a new distribution of capacity, and a new distribution of traffic. It is impossible to say how much time would be required; 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, but the SLA is still met.

1.7.3 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 – for example, in the aftermath of the '9/11' attack, as discussed above.

The interconnection ecosystem has often responded in this way with many people improvising, and working with the people they know personally. **[C:13]** This is related to traffic engineering, to the extent that it addresses the problem 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 cooperation in a crisis. **[C:14]**

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 would be needed.

1.8 Is Transit a Viable Business?

The provision of transit – the service of carrying traffic to every possible destination – 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. In the US, over the last ten to fifteen years transit prices have fallen at rate of around 40% per annum – 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 – 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 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 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 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 (CDNs, see below). This means that the continued

reduction in the unit price for transit is not being matched by an increase in transit traffic, so transit providers' revenues are decreasing.

The acknowledged market leader, Level 3, lost $2.9 billion in 2005-2008 and a further $0.6 billion in 2009, and another $0.6 billion in 2010. 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 many others did, and would make a small profit if it were not for the cost of servicing its debt. However, the industry as a whole is losing large amounts of money (we summarise some of the major providers' financial statements in Appendix II).

1.9 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. 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 ones – giving quicker response and faster transfers.
- costs are reduced because the data is not being repeatedly transported over large distances – saving on transit costs. However, the key motivation for the customers of CDNs 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 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. 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), 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 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 serving users 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 the Internet's traffic by volume.

1.10 The "Insecurity" of BGP

A fundamental problem with BGP is that there is no mechanism to verify that the routing information it distributes is valid. In principle traffic to any destination can be diverted – so traffic can be disrupted, modified, examined or all three. **[C:11]** The effect of this is felt on a regular basis when some network manages to announce large numbers of routes for addresses that belong to other networks; this can divert traffic into what 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. It is worth remembering that the operational layer is part of the ecosystem, and not all problems require technical solutions.

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. However, as discussed in Section 3.1.11, 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 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. **[Q:13]**

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, under the name RPKI (Resource Public Key Infrastructure). The RPKI initiative should allow ASes 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. This is an important step forward, and might tackle over 90% of 'fat finger' problems (outages caused by mistakes rather than deliberate attempts to disrupt). **[Q:14]**

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 weak, although the number of such incidents is expected to rise when IPv4 addresses are 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) address the issue more completely, and there have been other proposals since; 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.

During the standardisation process in 2011-2013 a key issue will be security economics. ASes 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 early adoption; another is for routers to prefer signed routes. Technical issues that must be studied during the standardisation phase include whether more secure BGP might, in fact, be bad for resilience (as was pointed out in the consultation, **[Q:15]**). 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, but it should tackle about half 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 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. **[Q:20]**

1.11 Cyber Exercises on Interconnection Resilience

The practical approach to assessing the resilience of the interconnection system is to run large-scale exercises in which plausible scenarios are tested. **[C:16]** Exercises can test both operational and technical aspects as well as procedural, policy, structural and communication aspects.

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 can network operations centres communicate if the phone network fails, or how well can field repairs proceed if the mobile phone network is unavailable? **[Q:17]**

- 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. **[C:15]**

Such scenario testing has been done at a national level and found to be valuable[3]. Something at a larger 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 evaluation report published by ENISA[4] proves the importance of such exercises and calls for future actions based on the lessons learned.

1.12 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. **[C:21]**

This manifests itself in a number of ways.

- In Section 1.10 above, we discussed the various proposals for more secure forms of BGP, from S-BGP in 1997 to BGPSEC in 2011, none of which have so far been deployed (see Section 3.1.12). There is little demand for something which is going to be difficult to implement and whose direct benefit is limited.
- There exists best practice for filtering BGP route announcements, which, if universally applied, would reduce instances of invalid routes being propagated by BGP and disrupting the system (see Section 3.1.11). But these recommendations are difficult to implement and mostly benefit other networks, so are not often implemented.
- There is an IETF BCP[5] [**6**] for filtering packets, to reduce 'address spoofing', which would mitigate denial of service attacks (see Section 5.8.3). 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 as about 75% bigger than it needs to be, because some networks announce extra routes to reduce their own costs (see Section 3.1.9). 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, despite the now imminent (circa Q3 2011) exhaustion of IPv4 address space. **[Q:10]**

[3] *Good Practice Guide on National Cyber Exercises,* ENISA Technical Report, *2009. Available at: http://www.enisa.europa.eu/act/res/policies/good-practices-1/exercises*

[4] *CYBER EUROPE 2010-Evaluation Report,* ENISA Report 2011*. Available (after 15/04/2011) at: http://www.enisa.europa.eu/act/res/*

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

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 reputation to maintain **[C:17]**. Normal circumstances include the usual day-to-day failures and small incidents **[Q:7]**.

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. **[Q:18]**

1.13 Regulation

Regulation is viewed with apprehension by the Internet community. Studies such as this are seen as stalking horses for regulatory interference, which is generally thought likely to be harmful. **[C:22]** 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 telecommunications 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, few countries 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' on 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 Clinton administration passed the Communications Assistance for Law Enforcement Act (CALEA) in 1994 mandating the cooperation of telecommunications carriers in wiretapping phone calls. The EU has a Data Retention Directive that is up for revision in 2011 and there is interest both in the UK and the USA in how wiretapping should be updated for an age not only of VoIP but also of diverse messaging systems. 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[6].

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

This study has 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. CDNs have been around for many years, but their emergence as a major component of the Internet is relatively recent; it is 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 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 a number of critical infrastructure industries that 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 reliability deficit thus far [**7**]. In practice, utilities have managed to cope by a combination of anticipatory risk management and Public-Private Partnerships (PPPs). 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 fall-back 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

[6] *See [**236**] UK Government, Cabinet Office Factsheet 18: Cyber Security. And for the popular perception of what government thinks see [**237**] "Fight Cyber War Before Planes Fall Out of Sky".*

nineteenth century, after which their efficiency declined steadily until competition from road freight arrived in the 1940s [**8**].

The prudent course of action for policy makers today is to start working to understand the Internet interconnection ecosystem. 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 we recommend 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 good 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 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 yet enough evidence on which to base major regulatory intervention with sufficient confidence.

2 Recommendations

Our 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. Incident correlation and analysis may lead to assessment and forecast models. The appropriate framework should be the result of a consultation with the industry and the appropriate regulatory authorities. 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. The implementation of Article 13a of the recent EU Telecom Package[7] may provide a model for this.

Recommendation 2 Data Collection of Network Performance Measurements

Europe should promote and support consistent, long-term and comprehensive data collection of network performance measurements. At present some real-time monitoring is done by companies such as ArborNet 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 into Resilience Metrics and Measurement Frameworks

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.

[7] *Directive 2002/21/EC of the European Parliament and of the Council, of 7 March 2002, on a common regulatory framework for electronic communications networks and services (Framework Directive), as amended by Directive 2009/140/EC and Regulation 544/2009.*

Recommendation 4 Development and Deployment of Secure Inter-domain 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 that Improve Resilience

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 transit provider.

The third group of recommendations aims at promoting good practice.

Recommendation 6 Promotion and Sharing of Good Practice on Internet Interconnections

Europe should sponsor and promote good practice in network management. Where good 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. Such information sharing should include modest and constructive engagement of industry stakeholders with public sector in relationship-building strategic dialogue and decisions through a series of PPPs for resilience.

Recommendation 7 Independent Testing of Equipment and Protocols

Public bodies at national or European-level 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 Conduct Regular Cyber Exercises on the Interconnection Infrastructure

The consultation noted that these are effective in improving resilience at local and national levels. The efforts at this level should continue in all countries in Europe as ‘we are as weak as the weakest link’. ENISA will support the national efforts. In addition regular pan-European exercises should be organised by European Member States in order to test and improve European-wide contingency plans (measures, procedures and structures). These large scale exercises will 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; thus effectively becoming part of improving the pan European cyber preparedness and contingency plans.

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

Recommendation 9 Transit Market Failure

It is possible that the current twenty-odd largest transit providers might consolidate down to a handful, 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 agencies but also national regulatory authorities. Recommendations 1, 2, 3, and 5 will prepare the ground technically so that policy makers will not be working entirely in the dark, but we also need political preparation.

Recommendation 10 Traffic Prioritisation

If, in a crisis, some traffic is to be given priority, and other traffic is to suffer discrimination, then the basis for this choice requires public 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 – Towards a Resilience Certification Scheme

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. In the long term efforts, including ENISA's, should focus on what mechanisms can be developed to give them the means to make more informed choices. This might involve combining the outputs from recommendations 2, 3, 5, 6 and 7 into a 'quality certification mark' scheme. Such scheme may prove an important tool to drive the market incentives towards enhancing the resilience of the networks and more generally of the interconnection ecosystem.

Respondents to the Consultation

We thank all those who gave their time to respond to the consultation and help us with this study. Some chose to contribute anonymously, but we can thank by name the follwoing:

Olivier Bonaventure	Professor	UCLouvain, Belgium
Scott Bradner	University Technology Security Officer, Office of the CIO	Harvard University
Bob Briscoe	Chief Researcher	Networks Research Centre, BT Group plc
kc claffey	Principal Investigator	CAIDA
Andrew Cormack	Chief Regulatory Adviser	JANET(UK)
Jon Crowcroft	Marconi Professor of Communications Systems	Computer Lab, Cambridge University
John Curran	CEO	ARIN
Dai Davies	General Manager	Dante
Nicolas Desmons	Chargé de Mission	ARCEP, France
Amogh Dhamdhere	Post-Doctoral Researcher	CAIDA
Giuseppe Di Battista	Professor of Computer Science	Roma Tre University
Nico Fischbach	Director, Network Architecture	Colt
Mark Fitzpatrick	Engineer	Federal Office of Communications, OFCOM, Switzerland
David Hutchison	Professor of Computing	Lancaster University
Malcolm Hutty	Head of Public Affairs	LINX
Christian Jacquenet	Director of the Strategic Program Office	France Telecom Group
Balachander Krishnamurthy	Researcher	AT&T Labs Research
Craig Labovitz	Chief Scientist	Arbor Networks
Ulrich Latzenhofer		Rundfunk und Telekom Regulierungs, Austria
Simon Leinen	Network Engineer	SWITCH
Tony Leung	Global Internet and Network Convergence Manager	REACH
Kurt Erik Lindqvist	CEO	Netnod
Neil Long	Researcher and Founder	Team Cymru Research NFP
Patricia Longstaff	David Levidow Professor of Communication Law and Policy James Martin Senior Visiting Fellow, Oxford Martin School Visiting Scholar	Syracuse University Trinity College, Oxford
Paolo Lucente	Architect/Designer	KPN International
Bill Manning		USC/ISI

Maurizio Pizzonia	Assistant Professor, Computer Science	Roma Tre University
Andrew Powell	Manager of Advice Delivery to the Communications Sector	UK Centre for the Protection of National Infrastructure
Edwin Punt	Product Manager	KPN International
Bruno Quoitin	Assistant Professor	University of Mons
Anders Rafting	Expert Adviser	Swedish Post and Telecom Agency
Jennifer Rexford	Professor	Department of Computer Science, Princeton University
Stefan Stefansson	Network Security Specialist	Post and Telecom Administration in Iceland
David Sutton	Director	tacit.tel (Telecommunications and Critical Infrastructure Technologies) Limited.
Guy Tal	Director of Strategic Relations	Limelight Networks
Rob Thomas	CEO and Founder	Team Cymru Research NFP
Nigel Titley	Head of Peering and Transit Strategy	Easynet/Sky
Andreas Wildberger	Generalsekretär	Internet Service Providers Austria (ISPA)