

If we catch the flu will the Internet fall sick too?

Issue

1. In an increasingly risk-adverse society enhanced social distancing is likely to be a significant response to a *severe* flu pandemic as people attempt to insulate themselves from becoming infected by the virus. Coupled with an increasingly 'wired' society such a response is likely to lead to both an increased demand for services delivered over the Internet and changes to the pattern of access to services.
2. Concern has been expressed that under these conditions the Internet in the UK might not be able to support the demand created by changes to patterns in working, education, socialisation and shopping.

Observations

3. While the resilience of the technical structure of the Internet is not likely to present a problem, ineffective business continuity arrangements might result as a consequence of wide-scale use by unfamiliar users and ill-prepared organisations. [Annex A](#) provides a check-list for those preparing to use the Internet for business continuity arrangements during a flu pandemic. In particular it is important to arrange connections (such as broadband) to the Internet as these are unlikely to be able to be installed within the escalation timescales presented by a pandemic.

What is the Internet?

4. The Internet is essentially a network of computer networks that provides a means of gaining access to resources such as devices that host websites and email servers. Households usually gain access to the Internet over a 'broadband'¹ connection provided by an Internet Service Provider (ISP) whereas organisations with significant communications requirements may have a direct connection provided by their telecommunications service provider.

¹ Broadband is usually understood to mean an 'always on' connection, as distinct from a 'dial-up' modem connection, that offers a bandwidth at or above 512 kb/s (or kilo (1,000) bits per second).

The inherent resilience of the Internet

5. The Internet was conceived by the US DOD² to provide a means of communicating that was resilient in the event of a nuclear strike. Resilience is derived from two principle components: the inherent physical resilience that is a consequence of the highly interconnected network of resources³ and the protocols, or sets of rules, which are used to communicate information across the networks. Information that traverses the networks is divided up into ‘chunks’ or packets. Packets of information can take different routes across the networks depending on the availability of resources and the connections between them. This is in marked contrast to the way in which a telephone call is set-up over the PSTN⁴. Additional resilience is derived from the highly diverse nature of the Internet as key information is frequently stored at more than a single location.

But how resilient *is* the Internet?

6. Two reports^{5,6} have contributed very different pictures of the resilience of the Internet as a consequence of social response to a flu pandemic. The British Computer Society, responding to Cabinet Office, were relatively sanguine about the implications. However, when Booz Allen Hamilton brought together 30 CEOs and senior executives from leading corporations, private and public sector institutions, and governments at The World Economic Forum Annual Meeting in 2006 *“They concluded that the telecommunication infrastructure will be severely strained and likely overwhelmed early in the pandemic (some experts opined that the Internet would shut down within two to four days of the outbreak)”* [sic]. It seems worthwhile to reflect on these polarised responses in the light of recent evidence relating to the resilience of the Internet.

² The US Department of Defence commissioned Arpanet the forerunner of the Internet, in 1969. The network consisted of only 4 nodes. The Internet is now unimaginably complex.

³ Consisting of devices such as ‘servers’, where information is stored, and ‘routers’, networking devices that determine the next network point to which information should be forwarded toward its destination.

⁴ Public Switched telephone Network

⁵ Response to the Cabinet Office Business Advisory Group on Civil Protection on Internet Resilience and Pandemic Impact. British Computer Society, December 2006.

⁶ Influenza Pandemic Simulation, Implications for the Public and Private Sectors. Booz Allen Hamilton. January 2006.

7. The main threat to the resilience of the Internet is the loss of a part or parts of the network or their inaccessibility. This may come about as a consequence of the loss of a strategic resource (or 'node') or a connection to the node or as a consequence of congestion between nodes. Nodes may become unavailable as a consequence of malfunction or physical loss.

8. The most extensive test to-date of physical loss was the complete loss of New York's main connection to the Internet following the collapse of the World Trade Centre in 2001. However, even with the loss of a significant node local Internet connectivity was not lost entirely although services were extremely degraded. Viruses do not only infect corporate and home computers; the Slammer virus (January 2003) infected Internet routers and other computers running Microsoft SQL server⁷. Routers are an essential component of the Internet responsible for directing information from sender to recipient and their unavailability would result in degradation of service. Although Slammer was particularly virulent causing a number of headline corporate IT failures⁸ the effect on the Internet was not significant.

9. Resources might become inaccessible as a consequence of the loss of a physical connection between nodes. In December 2008 while sheltering from bad storms in the Mediterranean ships off the coast of Alexandria dragged their anchors severing three international cables. The cable breaks resulted in significant loss of connectivity between Europe and the Middle East, Pakistan and India. Such 'single points of failure' are rare. The UK has many diverse connections with the Internet by both cable and satellite.

10. The ability of the Internet to carry traffic is being continually tested through both legitimate and malicious use. Public interest in the Clinton – Lewinsky affair (1998) and its disclosure on an Internet blog resulted in the highest levels of traffic that the fledgling Internet had seen. During the subsequent years of the dot-com bubble⁹ there was the largest investment in resources and connectivity in the Internet. Currently, it is unclear¹⁰ as to exactly how much installed capacity is actually

⁷ Structured Query Language, part of database software

⁸ Including the Seattle 911 (equivalent to the UK 999/112) service; Continental Airlines flights out of Houston and Newark were grounded as the airline was unable to reconcile passengers on their reservation and flight check-in systems and Bank of America and Royal Bank of Canada were unable to dispense cash from 13,000 ATMs

⁹ which 'burst' in the spring of 2000

¹⁰ See for example, <http://kn.theiet.org/magazine/issues/0906/into-the-light-0906.cfm>

in use. If market prices reflect the situation, connectivity between principle population centres are still very much commoditised. The most recent test at a country level was as a consequence of the cyber attacks on Estonia (May 2007). Although one of Europe's smallest countries Estonia has the highest broadband connectivity in Europe but relatively few connections with networks outside the country. This is the principle reason why the attacks were initially so successful.

But what about resilience as a consequence of a flu pandemic?

11. When the Internet is considered as an entity the available evidence indicates that is highly tolerant to faults. However, all communications networks invariably have 'pinch point' where services become throttled which can result in local congestion. A pinch point is very different to a 'single point of failure'. When communications systems are designed every effort is taken to remove single points of failure, namely a single location that has a significant effect on the entire system. However, communications systems invariably have capacity constraints either resulting from insufficient resources - as experienced in the aftermath of the loss of the World Trade Centre - or insufficient connectivity - as evidenced in Estonia's vulnerability to cyber attacks. When considering resilience as a consequence of society's response to a flu pandemic it is important to consider the effect of local pinch points.

12. The US Department of Homeland Security undertook a study¹¹ into the effect that a pandemic might have on communications networks. Their findings on congestion points are summarised in Figure 1. The study concluded that the pinch points of greatest concern (shown by red splashes) are at access points to the Internet. The implications of these pinch points for business continuity are summarised here at [Annex A](#).

¹¹ Pandemic Influenza Impact on Communications Networks Study. Department of Homeland Security, USA. December 2007.

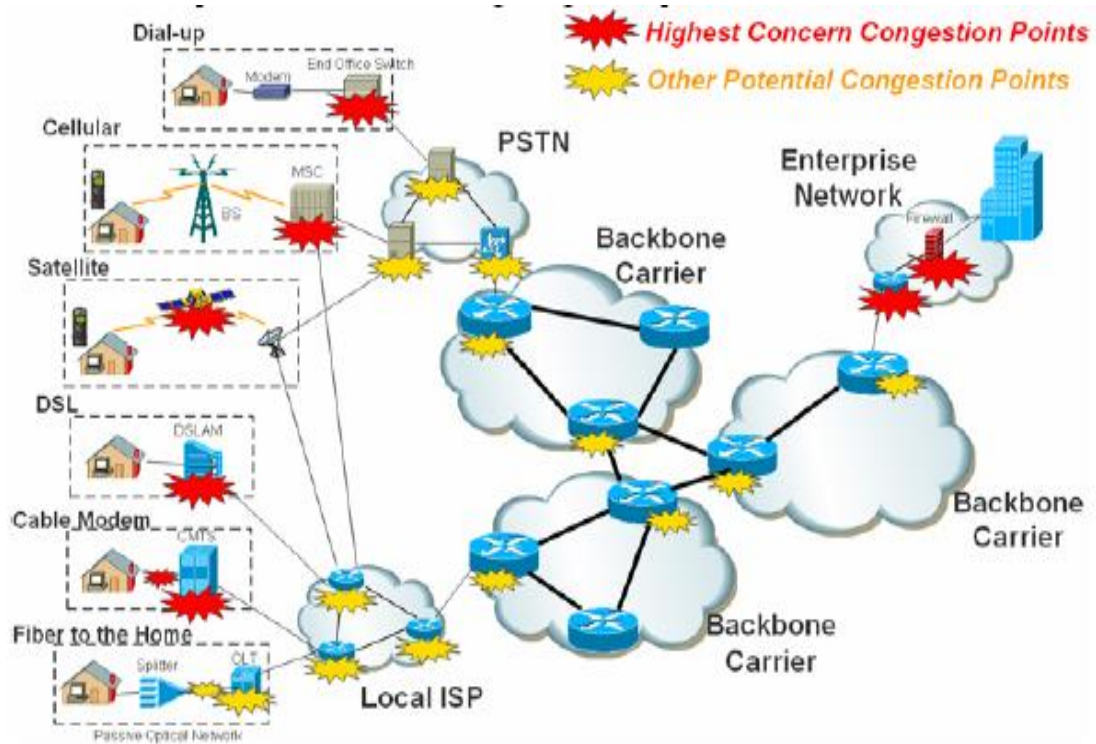


Figure 1. Potential pinch points in connecting across the Internet¹¹

DSL = Digital Subscriber Line roughly equivalent to a 'broadband' connection

ISP = Internet Service Provider

Societal response to a flu pandemic

13. It seems likely that societal response to a *severe* flu pandemic will be to take measures to decrease the likelihood of transmission of the virus through:

- increased social distancing and
- reduction in face-to-face communication.

This may be realised as:

- An increase in working away from the office environment, principally at home, possibly augmented with caring for infected family members and children not attending school;
- Children who would normally attend nursery or be educated out of the home being at home;
- Students in higher education potentially engaging to a greater extent in distance learning and

- Reluctance to leave the home to shop and to be entertained and socialise outside the immediate family.

This is likely to result in:

- increase in voice and data communications, entertainment and distance learning delivered over the Internet, shopping, browsing, blogging etc.

14. Anecdotal evidence¹² from the outbreak of winter vomiting virus¹³ in January 2008 which at its peak affected more than 100,000 people a week revealed that remote logins to corporate systems increased by 18% at the height of the outbreak. Additionally, there is likely to be an increase in requirements for information and news and the Internet is an ideal resource for gathering specific, focussed information to supplement broadcast information.

Effect on the Internet

15. The overall effect of society's response to a severe pandemic on the Internet is that traffic will increase and the period of highest traffic¹⁴ will broaden as users seek out times for a satisfactory experience to their requirements. The increase is likely to be led by bandwidth intensive services¹⁵ and intensive commercial applications. In comparison, services such as email, blogging and downloading web pages are less intensive but the overall volume is likely to increase. The overall effect of the increase in traffic will be to generally slow the response times to requests for information which will be reflected in slower delivery of emails and web pages and degradation in the quality of voice calls made over the Internet¹⁶. In addition to the connectivity pinch points, considered above, resources (such as web servers) hosting news and information sites perceived to be of high merit¹⁷ are likely to become congested resulting in slow delivery of pages or at worst inaccessibility.

¹² Reported by Signify a supplier of secure authentication solutions based in Cambridge, England.

¹³ Norovirus

¹⁴ Widely accepted to be between 20:00 and 22:00 during week days.

¹⁵ Such as those that stream video eg video on demand services eg BTVision and youtube; BBC iplayer (the highest bandwidth intensive site hosted in the UK); and Google earth.

¹⁶ Frequently referred to as VOIP (Voice Over Internet Protocol), Skype is a prominent VOIP service.

¹⁷ Such as principle broadcast and newspaper sites

Response by the telecommunications industry

16. In response to bandwidth intensive services there is evidence that Internet Service Providers (ISPs) are starting to throttle traffic¹⁸ *"in order to optimise the experience for all customers"*¹⁹. This is likely to be more prevalent the further an ISP is away from the top tier²⁰. As ISPs themselves increasingly offer bandwidth intensive services, such as high definition TV and video such policies are likely to become more entrenched. However, such policies are likely to have little effect on 'text-based' services¹⁹.

17. The industry response to a reduced labour force, as a consequence of a pandemic, would be to seek to contain the impact on services by prioritising fault repairs over routine maintenance and installation of new services²¹. New services provided during a pandemic are likely to be focused on urgent requests from Category 1 and Category 2 responders²².

Conclusions

18. The available evidence indicates that while the Internet is highly unlikely to fail, services may suffer degradation, though this is unlikely to be sufficiently severe to render them useless. These consequences stand a good chance of being mitigated by scheduling the time of access.

19. Largely as a consequence of the highly heterogeneous nature of the Internet any degradation of service is unlikely to be uniform.

20. Access to the Internet is likely to present the greatest opportunity for congestion. This can be mitigated by identify potential 'pinch points' in telecommunications arrangements and taking appropriate actions.

21. Plan ahead, it is unrealistic to expect that Internet connections can be installed on-demand during the escalation phase of a *severe* pandemic.

¹⁸ see for example <http://www.itproportal.com/portal/news/article/2009/6/11/bt-wants-youtube-bbc-iplayer-pay-broadband-access/>

¹⁹ <http://news.bbc.co.uk/1/hi/technology/8077839.stm>

²⁰ Tier 1 is the name given to the top tier of ISPs that provide connectivity and resources that only peer with another tier 1; allegedly they do not charge one another but ensure an equitable share is maintained of the Internet. Their business model relies on charging lower tier ISPs for carriage. Indication that an ISP is charging a content provider does not bode well for that ISP's business model.

²¹ Response from the EC-RRG communicated by the Secretary 14 February 2008.

²² As defined in the Civil Contingencies Act 2004.

Further information contact Nigel P Brown (nigel.brown@cabinet-office.x.gsi.gov.uk).

Civil Contingencies Secretariat

September 2009

revised October 2011, December 2011

Annex A. Implications for business continuity arrangements

Access from remote locations

A1. **Does your local connection have the capacity required for the intended use?** Email and web browsing is not very demanding on bandwidth however running a transaction platform has very much more demanding data requirements. Engage with your ISP (Internet Service Provider) now to ensure that your connection meets your requirements.

A2. **Is contention on your local connection too high?** Internet connections to domestic properties are contended, or shared, before they reach your ISPs (Internet Service Providers) connection into the Internet. The bandwidth (frequently referred to as 'speed' [sic]) in a domestic ISP contract is the highest likely to be available. In practice the bandwidth will be significantly lower. On a normal day there may only be a few people in your street using the Internet connection, in a pandemic it may be everyone. Children could also be doing their academic work on line. The increased contention will have the effect of throttling local speed. Business connections are usually uncontended.

A3. **Does your ISP place a cap on the amount of information that can be downloaded?** Some cheap Internet access contracts frequently contain a cap on the amount of information that can be downloaded over a billing period. When the limit is reached the service is either severely throttled or in the extreme cut off. Increased use for recreational activities may cause the cap to be reached rendering the connection useless.

A4. **Does your ISP enforce a traffic management policy?** Some Internet access contracts contain a clause that limits bandwidth (usually during the evening). This can have a considerable effect on the user experience for streaming video.

A5. **Have alternative means of accessing the Internet been considered?** Access to the Internet can be achieved by a wide range of means both over fixed line (cable (eg Virgin Media), ADSL (Asynchronous Digital Subscriber Line) achieved over a PSTN (Public Switched Telephone Network), and 'dial-up' (over the PSTN)) and wireless (satellite (eg Inmarsat BGAN service), 3G and WiFi). The bandwidth that is realised is highly dependent on geographical location and local contention.

A6. If response time produces an unsatisfactory experience consider time-shifting access to the Internet. Those downloading large files frequently report that the best speeds are achieved in the early hours of the morning.

Access to corporate resources

A7. Do employees know where to seek advice on the organisation's pandemic response arrangements?

A8. Are corporate resources scaled for a large proportion of staff gaining remote access? The response to a flu pandemic may be the most severe test of remote working facilities an organisation has experienced. Corporate resources, such as remote access servers (RAS), may not have sufficient capacity for the number of concurrent connections²³ that are required for remote access. At one extreme this will prevent additional users from gaining access at the other extreme corporate systems may become overwhelmed and deny access to all users. Organisations planning to make wide use of home working are recommended to take action to increase the capacity of their RAS and if possible to carry out tests to understand how they respond to a high number of concurrent connections.

A9. Are corporate users trained in using remote access? If they are not they will probably need to resort to support services, if these are delivered through a call centre, response may be constrained as a consequence of operatives not being able to work as a consequence of the pandemic. Ensure that key users are familiar with remote access through regular use.

A10. Do remote users have everything that they need to work remotely from their normal place of work? Does the remote connection enable access to fall the necessary file storage locations (information could potentially be held on inaccessible servers), corporate contact directories, employees calendars etc.?

A11. Are policies in place to manage expectations that will arise from prolonged home working?

A12. Are corporate users familiar with remote collaborative working practices? If they are not remote working could become hugely ineffective.

²³ RAS are normally scaled for a percentage of the workforce, not all of it. This is typically 25% but can vary from as low as 10% to as high as 50%. CPNI (Centre for the Protection of National Infrastructure).

A13. **Are remote access policies and procedures in place?** For example there could be potential health and safety implications (inappropriate posture and lighting conditions) resulting from long-term remote access. Regular telephone conferences with team members might be appropriate to help reduce isolation by those that do not routinely work out of the supportive office environment.

A14. **Are information resources sized for changes to access patterns resulting from a pandemic?** If your enterprise is providing web-based information ensure that the equipment and access has sufficient capacity for the anticipated demand.

A15. **Are employees familiar with security arrangements for disposal of print outs?**