Internet Analysis Report 2004
Protocols and Governance

FULL FINAL REPORT

Document Change Control

Version	Date	Authors	Summary of Changes
0.1	16/2/04	Ian Peter	
0.2	18/8/04	Ian Peter	Incorporates drafts
0.3	06/09/04	Ian Peter	Incorporating initial referee comments
0.4	11/09/04	Ian Peter	Final referee comments
1.0	13/09/04	Ian Peter	Copy for initial release
1.01	1/11/04	Ian Peter	Minor changes and clarifications

For more information contact:

Ian Peter

Ian Peter and Associates Pty Ltd

PO Box 10670 Adelaide St

Brisbane 4000 Australia

Tel 617 3870 1181 Mobile 614 1966 7772

Email ian.peter@ianpeter.com

Website www.ianpeter.com

Contents

1 PROBLEM DEFINITION 8
1.1 Introduction 8
1.2 Issues 9
1.3 Scope of study 10
1.3.1 Objectives 10

2 SUMMARY OF FINDINGS (EXECUTIVE SUMMARY) 11
2.1 History 11
2.2 User requirements and future needs. 13
2.3 Protocol issues 15
2.3.1 TCP/IP protocol 15
2.3.2 DNS and WHOIS 16
2.3.3 SMTP and Email protocols 17
2.4 Governance bodies 19
2.4.1 The International Engineering Task Force 19
2.4.2 Other standards bodies – ITU and W3C 21
2.4.3 ICANN related bodies 22
2.4.4 ICANN 23
2.4.5 Governance conclusions 24
2.5 Conclusions and Recommendations 24
2.5.1 Governance reform 25
2.5.2 Protocol reform 25
2.5.3 Change management issues with protocol reform 26
2.5.4 Technical Options 27
2.5.5 Risks and dependencies 28

3 STUDY BACKGROUND 29
3.1 Audience 30
3.2 References and Acknowledgements 30
3.3 Business Areas and Systems impacted 31
3.4 Related Projects 31

4 ANALYSIS 33

4.1 Examination of the history of the Internet as appropriate to considering its current purpose ... 34

4.1.1 Introduction ... 35

4.1.2 Early developments up to the 80s ... 35

4.1.3 "What God Hath Wrought" – the Internet infrastructure ... 37

4.1.4 Computers ... 38

4.1.5 Modems and LANs ... 39

4.1.6 Email beginnings ... 40

4.1.7 Early networking and the protocol wars ... 41

4.1.8 The World Wide Web ... 42

4.1.9 Conclusion ... 44

4.2 Examination of purposes and user requirements for today's Internet. ... 45

4.3 Examination of emerging trends and issues which will affect future usage and their impact ... 49

4.3.1 Size, scale, and speed ... 50

4.3.2 Global inclusiveness ... 50

4.3.3 ENUM and convergence with telephony ... 50

4.3.4 Wireless and mobility ... 51

4.3.5 Peer to Peer ... 51

4.3.6 Ongoing problems ... 52

4.3.7 Affordability ... 52

4.3.8 Conclusions on future needs ... 53

4.4 Examination of perceived problems with underlying technical standards in the light of current usage ... 54

4.4.1 Overview ... 55

4.4.2 TCP/IP ... 56

4.4.3 DNS and WHOIS ... 61

4.4.4 Mail systems and SMTP ... 68

4.5 Examination of governance structures as they have emerged (IETF, ICANN, RIRs, etc) 72

4.5.1 IETF ... 74

4.5.2 W3C and ITU ... 84

4.5.3 Root Servers and Regional Internet Registries ... 88

4.5.4 ICANN ... 92

5 RECOMMENDATIONS ... 97

5.1 Governance reform ... 98

5.2 Protocol reform ... 98

6 CHANGE MANAGEMENT ... 101

6.1 Factors in change ... 102

6.1.1 A compelling reason to change has to exist ... 102

6.1.2 Communication ... 102

6.1.3 Involvement of major players ... 102

6.1.4 Migration considerations ... 102

6.2 Stakeholder Analysis ... 103

6.2.1 Industry players ... 103

6.2.2 Current Internet bodies such as ICANN, IETF, etc ... 104

6.2.3 Government officials interested in Internet evolution ... 105

6.2.4 The academic and research sector ... 105

6.2.5 Corporate users of the Internet ... 105

6.2.6 Small business users on the internet ... 105

6.2.7 Non-commercial and individual users of the Internet ... 106

6.2.8 The 90% of the world's population who currently do not have Internet access. ... 106

6.3 Overall summation ... 106

7 TECHNICAL FEASIBILITY ... 107

7.1 Technical Assessment ... 108

7.2 Alternative Evaluations ... 108

8 RISK ASSESSMENT ... 110

8.1 Business Risks ... 111

9 NEXT STEPS ... 113

About the Author

Ian Peter became involved in the early beginnings of the Internet in Australia and Asia-Pacific from 1986. Currently a resident of Brisbane, Australia, he is involved with worldwide Internet organisations, provides strategic advice on Internet issues, and maintains a key interest in the early history of the Internet.

In 1993 he founded Ian Peter and Associates Pty Ltd, which works in the areas of strategy, policy, analysis and project management for a wide range of organisations. The company has a proud history of successfully managing change in large organisations, and of innovative approaches to business transformation across a wide range of areas including information technology management, human resource management, and communications management.

Past clients have included the United Nations Environment Program (UNEP), United Nations Conference on Environment and Development (UNCED), Internet Corporation for Assigned Names and Numbers (ICANN-DNSO), Asia Pacific Network Information Centre (APNIC), Telstra, Nortel, Ergon Energy, ABC-TV, Commonwealth of Australia and Queensland Government whole of government initiatives.

Contact details.

Ian Peter

Ian Peter and Associates Pty Ltd

PO Box 10670 Adelaide St

Brisbane 4000 Australia

Tel 617 3870 1181 Mobile 614 1966 7772

Email ian.peter@ianpeter.com

Website www.ianpeter.com

About the Internet Analysis Report – 2004

As someone who has been involved with the international growth of the Internet since the mid 1980's, and with Internet governance bodies in both professional and volunteer capacities, I have been concerned for some time that some of the organisations which are governing and managing the Internet have not evolved sufficiently to deal with the problems which have appeared as the Internet gets older and bigger.

I first heard the Internet protocols described as legacy systems in the 1990s, but, like most people, took no notice and got on with the business of growing it.

However, by the turn of the century cracks were beginning to appear. These cracks became larger and larger during 2003, to the point where it became clear to me and others, that something is wrong.

At the same time as some of us were becoming concerned with the technical issues, the World Summit on the Information Society was beginning to debate Internet governance – certainly not an unrelated topic. However, here another concern was evident to me: much of the governance debate was proceeding without full knowledge of what Internet governance actually involved.

My company, Ian Peter and Associates Pty Ltd, works professionally with large governmental and private sector organisations that have problems with organisational structures and old IT systems and uses established project and change management methodologies to analyse associated issues. It seemed to us to be appropriate to use these methodologies to look at today's Internet systems, governance, and user issues.

Searching the literature, we were surprised to find that traditional business analysis appears not to have been applied to Internet matters. As a result, we commissioned the Internet Analysis Report – 2004. We welcome your feedback.

Ian Peter

1 Problem Definition

1.1 Introduction

The Internet was developed in the 1970s and 1980s, initially as a means to connect mainframe computer systems for timesharing purposes. The system introduced for this fairly basic purpose has expanded to become a global multimedia information and communications system, connecting personal computers, phones, and hundreds of millions rather than the hundreds of devices originally foreseen.

Some of the significant developments not foreseen at the time of the original design include:

- The development and widespread use of networked personal computers
- The use of phones and portable devices on the Internet
- Broadband networks and processing power
- A network to be used for commercial purposes
- The World Wide Web.

Parts of the system are now over 20 years old, and the Internet is required to perform a number of important functions not included in the original design. New protocols have been developed, and various patches have been applied to base protocols, not always evenly. It seems appropriate to examine whether the current system, people, and processes are still appropriate.

Although it is clear that the system which has evolved is extraordinarily useful and needs to be continued, it is not clear whether the current Internet effectively meets its user's needs, on either a technical or a managerial level.

1.2 Issues

The Internet currently has some 600 millions users, or about 10% of the world's population. Although usage is still growing rapidly, both socio-economic and technical issues would appear to be slowing the overall growth rate. Some of the major user issues appear to be:

- The high incidence of viruses and worms arising from security weaknesses, giving rise to tens of billions of dollars of expenditure annually in an attempt to prevent damage
- A rapidly developing lack of trust in the system, caused by fraudulent use of addresses, non-existent companies, and the ease with which criminal activity can go unchecked
- A clogging of email systems with "spam", or junk email, most often emanating from fraudulent hosts and email addresses
- Issues as regards use that does not meet normal societal standards of behaviour
- Perceived "slowness" from a users point of view in accessing sites
- Limited availability in many parts of the world
- Affordability issues, particularly in developing countries
- Availability issues, particularly those pertaining to the ability to communicate via the Internet in native languages.

These and other significant issues are threatening the usefulness of the Internet. Preliminary analysis suggests that some significant improvements are both needed and possible, but not necessarily easily dealt within the current structure. Some of the reasons given include:

- Perceived inertia in the technical management and development by the Internet Engineering Task Force (IETF), giving rise to lack of faith among major players
- Perceived weaknesses in the management and governance structure
- Perceptions that the Internet has become a tool of US Government policy.

There is a lack of easy to read material analysing the current state of the Internet. There is also a lack of comprehensive user and business analysis. The Internet Analysis Report – 2004 addresses these needs.

1.3 Scope of study

1.3.1 Objectives

The key objectives of the Internet Analysis Report - 2004 are:

- To conduct a comprehensive business analysis of the Internet at the end of 2004, as an aid to determining future directions and strategies
- To produce a factual document giving the history and the current state of development of the Internet, to guide thinking about future directions
- To investigate the major user requirements for a 21st century Internet
- To initially analyse whether the current Internet is capable of meeting these objectives.

2 Summary of findings (Executive Summary)

2.1 History

Contrary to common opinion, neither the Pentagon nor 1969 hold up as the time and place the Internet was invented. A project which began in the Pentagon that year, called Arpanet, gave birth to the Internet protocols sometime later (during the 1970's), but 1969 was not the Internet's beginnings.

Larry Roberts, who was employed by Bob Taylor to build the Arpanet network, states that Arpanet was never intended to link people or be a communications and information facility. Arpanet was about time-sharing.

Time sharing tried to make it possible for research institutions to use the processing power of other institutions computers when they had large calculations to do that required more power, or when someone else's facility might do the job better. It never really worked as an idea – for a start, all the computers had different operating systems and versions and programs, and using someone else's machine was very difficult: but as well, by the time some of these problems were being overcome, mini-computers had appeared on the scene and the economics of time sharing had changed dramatically.

So it's reasonable to say that ARPANET failed in its purpose, but in the process made some significant discoveries which were to result in the creation of the first Internet. These included email developments, packet switching implementations, and development of the Transport Control Protocol – Internet Protocol (TCP/IP).

TCP/IP, the backbone standards which many people claim are the basis of determining what the Internet is, were developed in the 1970s in California by Vinton Cerf, Bob Kahn, Bob Braden, Jon Postel and other members of the Networking Group headed by Steve Crocker.

The sort of computers Arpanet was dealing with had very little power by today's standards. Only computer scientists used them. Computers with the power of modern day pocket calculators occupied whole floors of buildings. These monsters could only be afforded by large institutions.

It would take until the late 1970s for the personal computer to appear. Even so, for personal computers as well as mainframes, communication with other computers, and particularly other brands of computers, was an afterthought. It probably took the decade from 1983 to 1993 before anything like a sensible situation for computers to connect to the Internet emerged.

Ray Tomlinson is credited with inventing email in 1972, but in fact email is much older than that. Like many of the Internet inventors, Tomlinson worked for Bolt Beranek and Newman as an Arpanet contractor. He picked the @ symbol to denote sending messages from one computer to another. So then, for anyone using Internet standards, it was simply a matter of nominating <username>@<hostcomputername>. Email soon became the Internet's first "killer application".

In many ways, Internet adoption was about the path of least resistance. In the beginnings, governments wanted a completely different set of standards called OSI – but industry and governments could not agree on the details. There was a real mess out there, and no agreement on how to get out of it.

The dominant standards body that should have been interested in this problem was CCITT (Consultative Committee on International Telegraphy and Telephony) of the International Telecommunications Union (ITU), but they were essentially not interested in computers and software in the beginning, and when they did become interested, became committed to the ill-fated OSI track. So the Internet community had to devise its own way of dealings with standards.

This is probably where internet governance began to grow and formalise as a unique identity. A system called RFCs (Requests for Comment) was set up by Steve Crocker, and out of the network of engineers submitting and commenting on RFCs the Internet Engineering Task Force (IETF) evolved as a standards body.

Then the World Wide Web came along, and offered a much improved user interface and some substantial new applications. Every year from 1994 to 2000, the Internet saw massive growth. The Internet era had begun. The rest of the story is likely to be well known to most readers of this document.

These origins are important to our understanding because they help to explain how the Internet evolved. In particular, what we discover from a basic understanding of history is that the original protocols were introduced for a world which

- Had no personal computers
- Operated at very slow speeds
- Did not contemplate secure financial transactions
- Did not foresee non-English language users
- Was more concerned with computer to computer, rather than inter-personal communications
- Was willing to accept the trustworthiness of every participant
- Was for use by highly skilled and economically affluent people only.

It would be abnormal if protocols of this age and this difference of purpose were not, to all intents and purposes, legacy systems.

2.2 User requirements and future needs.

Our future Internet, rather than having 600 million users, may have close to 6 billion. So we may be only 10-20% of the way there, and there is a lot of growth to come. This will place new demands on the infrastructure and the way it is held together. Those experimenting with future networks with increased speed are already suggesting significant problems will exist coping with the increased size, scale and speed of the Internet.

The Internet Analysis Report - 2004 identifies both problems and future trends which need to be taken into consideration in looking at a future Internet.

Problems include:

- Identity fraud and other criminal activity
- Spam
- Viruses/worms
- Exposure of children to unacceptable material
- Hacking
- Speed
- Capacity to communicate in one's own language
- Affordability
- Accessibility in less economically developed areas and for socially disadvantaged groups

Some of the future issues and trends that need to be considered and which are outlined include:

- ENUM and convergence with telephony based systems
- the growth of wireless and mobility
- size, scale and speed issues
- the growth of peer to peer applications.

When these and other factors are considered, a statement begins to emerge about tomorrow's Internet.

The Internet is for everyone.

The Internet of the future must be

- *trustworthy*
- *reliable*
- *globally inclusive*
- *vendor neutral*
- *easy to use*
- *affordable*
- *able to change rapidly*
- *innovative and capable of significant expansion*
- *transparently and well managed*
- *involving industry, government and community stakeholders.*

This statement is to be built on to provide a more comprehensive overview of where the Internet must go.

2.3 Protocol issues

2.3.1 TCP/IP protocol

TCP (The Transport Control Protocol) in particular has come in for significant criticism, and a growing body of experts believe it will need to be replaced. Indeed, if it were easy to replace a fundamental Internet Protocol, this may have been done some time ago. It's the complexity of the change management problem that has delayed action rather than lack of recognised need for change.

Particular issues with TCP/IP include:

- Traffic prioritisation
- Unsuitability for financial transactions
- Security issues
- Performance issues with larger and faster networks.

The study concludes that TCP – if not TCP/IP - needs to be replaced, probably within a five to ten year time frame. The major issue to overcome is the migration issue which is discussed below.

2.3.2 DNS and WHOIS

Each host on the Internet can be located via an IP number. The Domain Name System (DNS) maps the numbers to names of hosts or websites (eg www.google.com, www.hotmail.com). Thus, when a user enters a name, the Internet knows which number to send the query to by looking up the DNS database.

It should be noted that the other widespread user of distributed network infrastructure, the telephone system, operates quite differently. It has no domain name equivalent with trade mark implications in normal uses – to contact a telephone address, you simply enter the number.

The DNS was introduced in 1984, several years before commercial traffic was able to be part of the Internet. At the same time, a public database called Whois was introduced, essentially to allow technical managers of hosts to contact their peers. This is the Internet equivalent of a telephone directory, but also serves a number of related purposes.

One issue with DNS is that it has not been possible to use native languages in email addresses, domain names, and the WHOIS database. This poses significant barriers to adoption for non-English speaking people.

The main problems here are that

- ASCII (the American Standard Code for Information Exchange) is used in the DNS. ASCII is incapable of supporting the complexities of foreign language
- No-one wants to substantially change the DNS**.**

Internationalised domain names (IDNs) have become a fundamental part of and an iconic symbol for the digital divide issue. ICANN has been criticised at its regular Public Forums for not giving the matter sufficient attention, failing to make significant progress, and being negative in its analysis of this issue. The Internet Analysis Report – 2004 examines this issue in detail.

Other issues with DNS include:

- Slow refresh rates, which pose particular difficulties with emerging standards such as ENUM and prohibit some interesting applications
- Issues with WHOIS and privacy
- Issues with security in DNS.

These are again problems that need to be addressed in a five year timeframe at the outside – some of them would be best handled more quickly if possible.

2.3.3 SMTP and Email protocols

To all intents and purposes, email is already broke, and must be fixed. The Internet's first and greatest killer application is now problematic.

In a survey examining email usage in 2003, the Pew Internet Project found that

- 25% of email users stated that the ever increasing volume of spam has reduced their overall use of email

- 70% of email users claimed spam had affected the quality of their on line experience
- 30% of users expressed fears that filtering approaches would cause loss of wanted mail
- 76% of users are bothered by offensive or obscene content in spam email
- 80% of users are bothered by deceptive or dishonest content in spam email.

Costs associated with spam have been estimated by various research firms at between $10 billion (European Union, 2004) and $87 billion (Nucleus Research, 2003) per annum. Spam volume is now estimated to exceed legitimate email volume; in May 2004, 76 percent of inbound e-mails scanned by email security provider MessageLabs Ltd were spam, up from 67 percent a month earlier.

ICANN claims spam issues as out of scope. *"... issues of concern to Internet users, such as the rules for financial transactions, Internet content control, unsolicited commercial email (spam), and data protection are outside the range of ICANN's mission of technical coordination"* (ICANN website). IETF has been very slow at doing anything in this field, preferring to leave investigation of the issues to a separate Internet Research Task Force (IRTF) group.

As a result, there is a general belief that nothing technical can be done to prevent spam. However, our analysis suggests that the existing protocols are significant contributors to the problem, and protocol reform could see spam volume drop by at least 80%.

SMTP, the basic email standard, is the online equivalent of borders without checkpoints and passports, or bank vaults without doors and locks. Some of the SMTP security weaknesses are:

- It allows anyone to connect with anyone without any system to say who they are
- It is simple to forge messages and pretend to be someone you are not with no checking whatsoever

- Not being one to one like telephone calls, it is easy to mass market to millions of email addresses at very low cost to the email sender.

These issues have been known for some time. Various attempts to provide improved protocols have been undertaken, but essentially have resulted in a mass of conflicting systems and standards. As a result, change is becoming more complex to initiate.

Email upgrades are complicated by

- Old systems which are never upgraded
- Incorrect applications of email systems
- The variety of applying protocols (eg http for webmail, smtp, nntp, pop etc)
- The ubiquitous nature of email
- IETF difficulties in handling big problems.

The Internet Analysis Report – 2004 analyses recent IETF work in this area and concludes that both governance issues and protocol reform need to be addressed to provide a more comprehensive solution.

2.4 Governance bodies

2.4.1 The International Engineering Task Force

Founded in 1986, the Internet Engineering Task Force describes itself as "a loosely self-organized group of people who contribute to the engineering and evolution of Internet technologies specifications". The IETF is unusual in that it is not a corporation and has no board of directors, no members, and no dues.

IETF's own internal analysis (RFC 3774 – IETF Problem Statement) has revealed significant problems, including:

- IETF is unsure about what it is trying to achieve

- Cannot determine what its scope should be
- Is unsure who its stakeholders are
- Cannot prioritise actions effectively
- Loses sight of overall architecture.

IETF governance contrasts substantially with the other two standards organisations involved with Internet standards. ITU has the strongest governance structure, being responsible eventually to member state representatives, and W3C standards work is determined and prioritized by a member organization.

So in this respect, IETF is peculiar. And this peculiarity brings with it certain problems because, in reality, few issues if any are purely technical and have no policy repercussions. This is shown out in case studies outlined in the Internet Analysis Report - 2004 where IETF of necessity has had to move outside its technical mandate but has not been effective in doing so.

Two of the case studies, covering DNSSEC and IPv6, also indicate extremely long time frames within IETF for protocol development and implementation. No-one can attribute these long time frames to technical complexity alone. Poor methodologies, volunteerism, under-resourcing, unprofessional behaviour and management issues have all contributed to delays, according to the IETF Problem Working Group.

IETF's decisions to address its problems in an open forum are to be applauded, as are its attempts to engage a wide global audience of engineers in its consensus based decision making structures. However, IETF is a classic technocracy. While it appears to be reasonably capable of managing the day to day concerns as regards maintenance of standards, it does not have the capacity to tackle major tasks or major change. To solve these problems, IETF would need to

- Have a clear relationship to a governing body with the competence and representative nature to deal with policy considerations which are outside the scope and expertise of IETF

- Have means of introducing non technical skills sets such as project and change management to its affairs
- Be resourced to provide dedicated rather than volunteer effort in working on major priorities
- Have a clear scope, mission, and relationship within an overall technical management structure for the Internet, and with other Internet bodies of a non-technical nature which fulfil complementary roles in determining policy and assisting adoption of standards
- Learn to communicate effectively with business, community and government stakeholders.

2.4.2 Other standards bodies – ITU and W3C

To an outsider, The International Telecommunications Union (ITU) appears to have all the efficiency and capacity to get things done that ICANN appears not to have. However, that perception may well be illusory; and anyone who has been involved in development of standards such as X.400 could point to equally problematic issues in ITU.

It should be remembered that, at the time IETF was established in 1986, telecommunications companies were not major players in the emerging Internet. They became more involved from 1990 on, as a commercial Internet got underway. It should also be remembered that in the 1970s the US telecommunications giant, AT&T, could not see a business case for involvement in the Internet. This perhaps is the most telling criticism of the staid and solid elder statesman that ITU is – it may find it difficult to be nimble or innovative in seeing future directions.

ITU, like IETF, is undergoing considerable internally driven reform to try to better cope with the demands of a rapidly changing communications technology landscape. There would appear to be room for some of the strengths of ITU to be better utilised alongside those of IETF and ICANN in the future, particularly as telephony and Internet based applications continue to converge.

W3C is the third “standards body”, and effectively addresses issues with the World Wide Web architecture. It separated from IETF in 1994 as it believed IETF to be incapable of dealing with its particular range of issues.

2.4.3 ICANN related bodies

It is important to realize that ICANN doesn't control everything in Internet technical co-ordination. An interesting history associated with the early growth of the Internet led to a number of quite independent structures being established. These include:

- Country domain administrators, who in many countries were early technical volunteers who have no formal relationship with national governments. In some developing countries, country administrators are located overseas and are not national citizens. Although a form of techno neo-colonialism remains in the administration of some country domains, and some hostility to co-operation with government authorities exists in others, most country domains are now forming appropriate locally defined relationships with their governments and their local constituencies.
- Regional Internet Registries (RIRs) such as APNIC (Asia Pacific) and RIPE (Europe), which were set up before the ICANN/US Government contract was in place and retain substantial independence while administering the IP numbering system.
- Root server operators, many of them volunteers. The central root server is administered by Verisign Ltd, and any changes need the approval of the US Dept of Commerce, not ICANN.

ICANN has a series of relationships with these separate bodies which it is attempting to formalise.

2.4.4 ICANN

The Internet Corporation for Assigned Names and Numbers (ICANN) exists in its current form largely because the US Government wanted it to be so.

Its structure is an evolving reactive mechanism. Anyone analysing its current structure without regard for the history of how it came to be would have to regard ICANN as

- eccentric in structure
- illogical in scope, and
- incomplete in terms of Internet governance.

The initial proposal for a body to administer the domain name system suggested establishment under Swiss law. However at the beginning of October 1998 the US Government's National Telecommunications and Information Administration (NTIA) responded to this proposal by announcing the Internet Corporation for Assigned Names & Numbers (ICANN). It would operate under an agreement with the NTIA with oversight by the US congress. The new body was asked to ensure competition in delivery of domain name services. Thus ICANN became a corporation under US law, with a contract to operate from the US government, despite concerns of many stakeholders.

ICANN claims its mission to be technical co-ordination. (ICANN website). However, because of the eccentricities and incomplete nature of Internet governance structures, ICANN has consistently worked in areas that cannot be regarded as technical co-ordination.

For instance, in 1999 it succeeded in establishing a Uniform Dispute Resolutions Policy (UDRP) for the top level domains; hardly a technical co-ordination task, but certainly a useful one for development of the new media.

Similarly eccentric is the role of ICANN in creating a competitive environment in DNS, part of its contract with US Department of Commerce. This would normally be seen as a regulatory body's responsibilities, not a technical co-ordination task.

Public policy matters where ICANN is active include intellectual property issues and security. Public policy matters where ICANN is not active include spam and consumer protection. Once again, the logic of involvement and non-involvement is not easy to follow.

Perhaps partially as a result of this mission confusion, ICANN does not handle public policy well or effectively. An example of this was its recent attempts to gain widespread public input in to the WHOIS database and privacy issues.

2.4.5 Governance conclusions

The problem with ICANN, and with IETF, is one of defining scope within a schema that effectively manages all needs of the 21st century Internet. No such schema exists, and that is why bodies such as ICANN and IETF are continually operating in areas outside of their level of competence in order to keep things afloat.

If there is a problem in Internet governance, it is the gaps between the competencies of existence governance bodies and the needs of Internet industry, governmental, and community users. As user needs in a broad sense do not come within the range of concern of any particular Internet governance body, it is inevitable that mistakes are being made and crucial issues are not being addressed.

2.5 Conclusions and Recommendations

There are many problems facing the Internet at the current time. Those closest to each problem are making substantial efforts in isolation from other major issues to address their particular problems, often however seeking short term fixes rather than longer term solutions. A co-ordinated approach across all areas seems necessary for resolution.

IETF and ICANN have been unable to deal with the range of issues and concerns adequately. This is acknowledged as an issue by many people who are aware of the problems.

The question then becomes how to proceed to reform.

2.5.1 Governance reform

The WSIS governance debate will be ongoing until at least mid 2005, and at that point of time is only likely to consider what governance structures should be in the future. This study acknowledges that an appropriate forum exists to examine governance issues, and suggests that some of the information in the Internet Analysis Report - 2004 might assist in discussions of how to appropriately reform governance. It is not simply a matter of removing ICANN from the somewhat benign influence of the US Dept of Commerce; it is necessary to consider the range of issues involved in creating tomorrow’s Internet as an effective and equitable tool for human development. That requires some refocussing of discussion to address some wider issues and needs.

2.5.2 Protocol reform

What is not being discussed as comprehensively is how to reform protocols. The Internet Analysis Report – 2004 concludes that protocol reform is necessary.

Project management structures are both the normal and the most effective way to handle major change projects of this nature. A separate short term structure should be established for this purpose. It would consist of:

- A dedicated full time project team, containing the best engineering, project management, research, change management and communications expertise available to scope and deal with this problem. Other skills sets will also be necessary as the project progresses.
- A Project Steering Committee with major stakeholder involvement to review progress and determine direction.

The exact makeup of both the Project team and the Steering Committee needs further investigation with major stakeholders. A preliminary range of potential activities and involvement possibilities are outlined in the full study and at www.internetmark2.org.

The entire process is unlikely to be completed in less than three years. There are compelling reasons to suggest it must be completed within seven. At the end of a successful implementation, the project would hand over responsibility to an appropriate Internet governance structure which will most likely have emerged from WSIS and beyond.

It may be suggested that current Internet organisations should deal with this change. This study believes that would be ineffective, as:

- They have substantial operational issues to deal with which stretch their resources
- Their technical-only governance is inappropriate for the task
- Volunteer structures are inappropriate
- Their track record in handling major change in a timely fashion is not good
- They lack the range of disciplines and the skills base for handling change of this magnitude.

For these reasons the existing governance and standards bodies, although they should be encouraged to be active participants, should not be the only voices at the table for this future development initiative.

2.5.3 Change management issues with protocol reform

Knowledgeable insiders in stakeholder Internet industry companies and research institutions are aware to varying degrees of the need for improvement in the base protocols. However, executive decision makers at this stage are probably not aware of the need for reform. The Internet Analysis Report – 2004 is designed to assist in raising awareness, and contains a detailed analysis of the change management issues involved in protocol reform.

Some large industry players may initially perceive benefit in blocking or delaying change in order to protect dominant market positions, or to gain opportunities from proprietary approaches. There will doubtless be robust debates within some industry players about the market affect of change in this area.

However, market leaders will face crucial losses of markets if change does not occur within a reasonable time frame. Involvement in the new generation Internet will eventually be seen as a market plus; but the issues must be understood and clearly communicated for this to happen.

Some defensive actions can be expected in current Internet bodies, and some feeling of loss of control and that they should be in charge of change. The more positive involvement of individuals who initially react this way is more likely in a middle adoption phase, when the project clearly has impetus.

Government officials from a wide variety of backgrounds will have interest in new governance structures and in public policy issues which may emanate from projected changes. However they will be less interested in the protocol issues and content to leave that to others to develop.

Information is the key to success here. Unless the problems are recognised and acknowledged, and compelling reasons to effect change are understood, action to rectify the problems will not occur.

2.5.4 Technical Options

It appears necessary at this preliminary stage to look at some sort of gateway structure through which people pass to adopt “Internet Mark 2”. For a substantial period of time, Internet Mark 2 will need to co-exist and co-operate with the legacy Internet. However eventually people will pass through the gate to the other side.

And, although differing opinions exist, it also appears to be necessary to nominate a point in time at which the gates shut, and interoperability with the old system disappears.

Other options to be examined would include a layered super-structure on top of existing protocols, or an overall peer to peer architecture that effectively bypasses problematic protocols gradually.

2.5.5 Risks and dependencies

The risks associated with inaction on these issues substantially exceed those involved in moving forward. The risks associated with delay are also substantial.

Significant problems with base protocols are already evident. Given that the necessary changes cannot be implemented overnight, it is appropriate that a major effort to address these problems begin now.

Seed funding is necessary to move the project forward into planning phases Details will be developed with stakeholders. A preliminary range of funding, sponsorship and involvement possibilities are suggested at www.internetmark2.org.

The Internet Analysis Report – 2004 concludes that there are substantial inter-dependencies between protocol reform and governance. The dilemma facing the Internet community is that protocol reform has dependencies on governance reform, but cannot await completion of reform processes.

Protocol reform and governance reform must therefore be addressed in parallel, with some clear understandings among all parties concerned of the need to co-operate fully in ensuring that the next generation protocols and the next generation governance eventually come together in an appropriate manner.

3 Study background

3.1 Audience

This document is directed to all stakeholders in the future of the Internet. This includes:

- The Internet industry and related software development corporations
- Government officials interested in Internet evolution
- Current Internet bodies such as ICANN, RIRs, IETF, etc
- Other international organisations (including United Nations organisms) with an interest in Internet evolution
- Policy analysts and non government organisations.

3.2 References and Acknowledgements

The author makes no claim to uniqueness of any of these observations; rather they are widespread, and are synthesised here.

The involvement of referees and commentators has been very valuable in completing this study. It should be stressed that referees do not necessarily endorse or agree with all or any of the findings of the study. However, they have taken time to read and to comment, and their contribution is gratefully acknowledged.

Some referees have preferred to remain anonymous. Those who can be acknowledged publicly include:

- Dr Greg Adamson, Adamson Communication Consulting
- Avri Doria, Researcher, Lulea University of Technology
- Paul Vixie, Chairman, Internet Software Consortium
- Paul Wilson, Director-General, APNIC

The referees do not necessarily endorse all the findings of the study.

- Paavan Duggal, Advocate, Supreme Court of India assisted with legal advice.
- Scott Washington assisted with structural change issues.

Others provided substantial material represented in various sections of the report, particularly

- The Internet Society (www.isoc.org) and The Internet Tapes (www.theinternettapes.com) were substantial sources in the history section.
- The Work of Sally Floyd at ICIR (http://www.icir.org/floyd/) and Dr Greg Adamson (http://home.iprimus.com.au/greg_adamson/Pages/acc_research1.htm) were substantial sources in the analysis of TCP/IP.
- RFC 3774 – IETF Problem Statement (Elwyn B Davies) was a major source of analysis of IETF
- The work of Paul Vixie and Diane Davidowicz on DNS and DNSSEC is acknowledge
- The work of John Klensin and James Seng on International Domain Names
- The input of Phillip Hallam-Baker and Alan DeKok on SMTP and email issues.

3.3 Business Areas and Systems impacted

The issues covered in this report and impacted by its findings include

- Basic Internet protocols
- Internet governance organisations.

3.4 Related Projects

We are not aware of similar projects.

The Internet2 Project might come to mind. However Internet2 essentially does not cover governance issues, and essentially concentrates on performance and bandwidth issues. Useful though all this might be, it is entirely different in scope.

This study gives mention to various activities of IETF and ICANN in attempting to address some of the issues raised.

4 Analysis

4.1 Examination of the history of the Internet as appropriate to considering its current purpose

4.1.1 Introduction

The history of the Internet helps us to understand where we are in evolution and the nature of the issues we face.

4.1.2 Early developments up to the 80s

The idea of an Internet began in 1863 if not earlier, when author Jules Verne wrote of a future world where "photo-telegraphy allowed any writing, signature or illustration to be sent faraway – every house was wired".

The Internet couldn't have started without the big inventions of the 19^{th} century – electricity and the telephone. And, to a lesser degree, there was unlikely to be an Internet as we know it before there were the standard electronic broadcast media of radio and television. So the building blocks were the existing communications, broadcasting and energy technologies. Marconi, Alexander Graeme Bell, Tesla, and many other pioneers provided the inventions on which the Internet rode.

Many people have heard that the Internet began with some military computers in the Pentagon called Arpanet in 1969. The theory goes on to suggest that the network was designed to survive a nuclear attack. However, neither the Pentagon nor 1969 hold up as the time and place the Internet was invented. A project which began in the Pentagon that year, called Arpanet, gave birth to the Internet protocols sometime later (during the 1970's), but 1969 was not the Internet's beginnings. Surviving a nuclear attack was not Arpanet's motivation, nor was building a global communications network.

Bob Taylor, the Pentagon official who was in charge of the Pentagon's Advanced Research Projects Agency Network (or Arpanet) program, insists that the purpose was not military, but scientific. The nuclear attack theory was never part of the design. Nor was an Internet in the sense we know it part of the Pentagon's 1969 thinking. Larry Roberts, who was employed by Bob Taylor to build the Arpanet network, states that Arpanet was never intended to link people or be a communications and information facility.

Arpanet was about time-sharing. Time sharing tried to make it possible for research institutions to use the processing power of other institutions computers when they had large calculations to do that required more power, or when someone else's facility might do the job better.

What Arpanet did in 1969 that was important was to develop a variation of a technique called packet switching. In 1965, before Arpanet came into existence, an Englishman called Donald Davies had proposed a similar facility to Arpanet in the United Kingdom, the NPL Data Communications Network. It never got funded; but Donald Davies did develop the concept of packet switching, a means by which messages can travel from point to point across a network. Although others in the USA were working on packet switching techniques at the same time (notably Leonard Kleinrock and Paul Baran), it was the UK version that Arpanet first adopted.

However, although Arpanet developed packet switching, Larry Roberts makes it clear that sending messages between people was "not an important motivation for a network of scientific computers". Its purpose was to allow people in diverse locations to utilise time on other computers. It never really worked as an idea – for a start, all the computers had different operating systems and versions and programs, and using someone else's machine was very difficult: but as well, by the time some of these problems were being overcome, mini-computers had appeared on the scene and the economics of time sharing had changed dramatically.

So it's reasonable to say that ARPANET failed in its purpose, but in the process it made some significant discoveries that were to result in the creation of the first Internet. These included email developments, packet switching implementations, and development of the (Transport Control Protocol – Internet Protocol) or TCP/IP.

TCP/IP, (Transport Control Protocol – Internet Protocol), the backbone standard which many people claim is the basis of determining what the Internet is, was developed in the 1970s in California by Vinton Cerf, Bob Kahn, Bob Braden, Jon Postel and other members of the Networking Group headed by Steve Crocker. That Internet beginning was largely about large research institutions and universities. It would be some time before it became available to the rest of us. In fact, TCP/IP was not added to Arpanet officially until 1983.

Nobody knows who first used the word Internet – it just became a shortcut for "internetworking". The earliest written use of the word appears to be by Vint Cerf in 1974.

4.1.3 "What God Hath Wrought" – the Internet infrastructure

The Internet rides on the infrastructure of the global telephony system. This began with the telegraph system as it was called in the beginning. The first line was built in 1844 from Washington to Baltimore. By 1858 a transatlantic cable was in place, and by 1861 telegraph wires covered the USA.

From there on growth was profound. Within 100 years, the network had become the biggest single construction on the planet, wielding its way across continents and under oceans in a massive encircling web of cables, wires, satelliture, and wireless connections,

Marshall McLuhan suggested to us in his epic "Understanding Media" in the 1960s that "It is instructive to follow the embryonic stages of any new growth, for during this period of development it is much misunderstood, whether it be printing or the motor car or TV".

That's certainly true of the telephone system as much as it was for the Internet. In the early days, Alexander Graeme Bell thought the telephone would be good for broadcasting music. Then people thought it would be good for sending Morse Code messages. Sometime later people thought it would be good for people speaking with each other using "telephones". That idea stuck for quite some time as the dominant purpose of these networks, but by the 1990s we were beginning to see Internet networks emerge.

It's reasonable to assume that we are in very early stages of Internet development, and today's purposes for the Internet are preliminary rather than final. To the children of the digital age, the Internet is for information, communication, buying and selling, film, exchanging photos and music, sending messages, talking, all of the above and much more.

4.1.4 Computers

The sort of computers Arpanet was dealing with were monsters with very little power by today's standards. Only computer scientists used them. Computers with the power of modern day pocket calculators occupied whole floors of buildings. These monsters, or mainframes as they are called where they still exist these days, could only be afforded by large institutions. You have to wait another ten years until the late 1970s for the personal computer to appear.

The first personal computer, the Altair 8800, cost $US379 and was shipped by Micro Instrumentation Telemetry Services Company in January 1975. Over 1000 were sold. By 1977 The Radio Shack TRS 80, Apple 2, and Commodore PET were also on the market. IBM got the idea by about 1981 and released the first IBM PC.

The company that dominated the market in the early days, when they had 75% of the market, was Apple Computer. Apple was founded by Steve Jobs and Steve Wozniak just outside of San Francisco. The original Apple operating system was called CPM, and it was very like the competitor which was to overtake it and launch Bill Gates of Microsoft on the way to his fortune, MSDOS.

Early personal computers featured a thing called a "command line". They didn't yet have a mouse, although joysticks for games machines were starting to appear.

The personal computer started the growth of the Internet as we know it. Without personal computers, the Internet would be very different indeed.

4.1.5 Modems and LANs

Modems enable the digital form of information that a computer uses to communicate using the analogue form of transmission of old style telephone systems. Modem is another term we are likely to forget soon in the digital age, but for many of us modems were where internetworking began, as we used modems to connect our computers to telephone lines.

There were apparently some early modems used by the US Air Force in the 1950's, but the first commercial ones were made a decade later. The earliest modems were 75 "baud" (or bits per second). That's about 1/750th of the speed of current modems, so they were pretty slow! Early commercial modems were 300 baud. Then came 1200 baud, and by 1989 2400 baud modems were popular.

By 1994, modems had got to 28.8 kilobits per second – which was just as well, because by then we were beginning to send more than text messages over the Internet and needed more bandwidth. That was thought to be an upper limit for phone line transmissions. But along came the 56k modem, and a new set of standards, so the speeds continue to push the envelope of the capacity of the telephone system.

So much so that many of have moved on, into wireless networks, and into "broadband" systems, which allow much faster speeds. But modems made the first critical link between computers and telephones and began the age of internetworking for most of us.

Another of the former Arpanet contractors, Robert Metcalfe, was responsible for the development of Ethernet, which drives most local area networks.

Ethernet essentially made a version of the packet switching and Internet protocols which were being developed for Arpanet available to cabled networks. After a stint at the innovative Xerox Palo Alto laboratories, Metcalfe founded a company called 3Com which released products for networking mainframes and mini computers in 1981, and personal computers in 1982.

With these developments in place, tools were readily available to connect both old and new style computers, via wireless, cable, and telephone networks. As the networks grew, other companies such as Novell and CISCO began to develop more complex networking hubs, bridges, routers and other equipment. By the mid 1980's, everything that was needed for an explosion of internetworking was in place.

4.1.6 Email beginnings

Before we started to want to network computers, there had to be a good reason – and the first one was email. Email is much older than Arpanet or the Internet. It was never invented, it evolved.

Before internetworking began, email could only be used to send messages to various users of the same computer. These computers might have up to one hundred users –often they used what were called "dumb terminals" to access the mainframe from their work desks. (Dumb terminals just connected to the mainframe – they had no storage or memory of their own, they did all their work on the remote computer).

Probably the first email system of this type was MAILBOX, used at Massachusetts Institute of Technology from 1965. This early email was just a miniscule advance on what we know these days as a file directory – it just put a message in another programmer's directory in a spot where they could see it when they logged in. Simple as that; just like putting a file in a folder. Another early program to send messages on the same computer was called SNDMSG.

Once computers began to talk to each other over networks, the problem became a little bit more like putting a letter in an envelope and addressing it. To do this, we needed a means to indicate to whom letters should go that the electronic posties all understood – just like the postal system, we needed a standard way to indicate an address.

This is why Ray Tomlinson is credited with inventing email in 1972. Like many of the Internet inventors, Tomlinson worked for Bolt Beranek and Newman as an Arpanet contractor. He picked the @ symbol to denote sending messages from one computer to another. So then, for anyone using Internet standards, it was simply a matter of nominating <username>@<hostcomputername>.

By 1974 there were about a hundred military users of email, because Arpanet eventually encouraged it. Email became the saviour of Arpanet, and caused a radical shift in Arpa's purpose. Things developed rapidly from there. Larry Roberts invented some email folders for his boss, a big advance. In 1975 John Vital developed some software to organize email. By 1976 email had really taken off, and commercial packages began to become available. Within a couple of years, 75% of all Arpanet traffic was email.

Email took us from Arpanet to the Internet. Here was something that ordinary people all over the world wanted to use.

It wasn't easy to use email in those days. Typing on line was awful – people had all sorts of problems with simple things like getting backspace keys to work. You had to set a "terminal type" for your computer to get the characters to appear properly – modems had to be programmed for the settings of the network – it really was pretty difficult to use those early networks; it says a lot for how much people wanted to network that they actually put up with those interfaces.

Later on email started to take on some pretty neat features. One of the first good commercial systems was Eudora, developed by Steve Dorner in 1988. Not long after Pegasus mail appeared

4.1.7 Early networking and the protocol wars

For personal computers as well as mainframes, communicating with other computers, and particularly other brands of computers, was an afterthought. It probably took a decade from 1983 to 1993 before anything like a sensible situation for computers to connect to the Internet emerged.

In many ways, Internet adoption was about the path of least resistance. In the beginnings, governments wanted a completely different set of standards called OSI – but industry and governments could not agree on the details. There was a real mess out there, and no agreement on how to get out of it.

The dominant standards body that should have been interested in this problem was CCITT (Consultative Committee on International Telegraphy and Telephony) of the International Telecommunications Union (ITU), but they were essentially not interested in computers and software in the beginning, and when they did become interested, became committed to the ill-fated OSI track. So the Internet community had to devise its own way of dealings with standards.

This is probably where internet governance began. A system called RFCs (Requests for Comment) was set up by Steve Crocker, another Internet pioneer, and out of the network of engineers submitting and commenting on RFCs evolved the Internet Engineering Task Force (IETF) as a standards body.

4.1.8 The World Wide Web

Before the World Wide Web the Internet really only provided screens full of text (and usually only in one font and font size). So although it was pretty good for exchanging information, and indeed for accessing information such as the Catalogue of the US Library of Congress, it was visually very boring.

In an attempt to make this more aesthetic, companies like Compuserve and AOL began developing what used to be called GUIs (or graphical user interfaces). GUIs added a bit of colour and a bit of layout, but were still pretty boring. Indeed personal computers were only beginning to adopt Windows interfaces – before that with MSDOS interfaces they were pretty primitive. So the Internet might have been useful, but it wasn't good looking.

Probably the World Wide Web saved the fledgling Internet. Not only did it change its appearance, it made it possible for pictures and sound to be displayed and exchanged.

The web had some important predecessors, perhaps the most significant of these being Ted Nelson's Xanadu project. But it was Tim Berners Lee who brought this all together and created the World Wide Web. The first trials were at the CERN laboratories (one of Europe's largest research laboratories) in Switzerland in December 1990. By 1991 browser and web server software was available, and by 1992 a few preliminary sites existed in places like University of Illinois, where Mark Andreesen became involved. By the end of 1992, there were about 26 sites.

The first browser which became popularly available to take advantage of this was Mosaic, in 1993. On April 30, 1993: CERN's directors made a statement that was a true milestone in Internet history. On this day, they declared that WWW technology would be freely usable by anyone, with no fees being payable to CERN. This decision – much in line with the decisions of the earlier Internet pioneers to make their products freely available – was a visionary and important one.

The browser really did begin to change everything. By the end of 1994 there were a million browser copies in use. In the same year Marc Andreesen founded Netscape Corporation.

1994 also saw the first Worldwide Web Consortium meeting at MIT in Cambridge, Massachusetts. This was a landmark also, because the W3C as they became known decided to split from the IETF standards body and create their own. Thus another standards body entered into the picture. In the beginning, freed from the IETF restraints, W3C was prolific in developing standards for the new medium.

Then the rapid growth began. Every year from 1994 to 2000, the Internet saw massive growth, the like of which had not been seen with any preceding technology. The Internet era had begun.

The rest of the story is likely to be well known to most readers of this document. With the World Wide Web came an explosion of global usage, which catapulted the Internet from relative obscurity to the mainstream of western communication media. Internet growth has been remarkable, and perhaps this rapid growth has been a factor in some of the problems which have emerged.

4.1.9 Conclusion

These origins are important to our understanding because they help to explain how the Internet evolved. In particular, what we discover from a basic understanding of history is that the original protocols were introduced for a world which

- Had no personal computers
- Operated at very slow speeds
- Did not contemplate secure financial transactions
- Did not foresee non-English language users
- Was more concerned with computer to computer, rather than interpersonal communications
- Was willing to accept the trustworthiness of every participant
- Was for use by highly skilled and economically affluent people only.

It would be abnormal if protocols of this age and this difference of purpose were not, to all intents and purposes, legacy systems.

4.2 Examination of purposes and user requirements for today's Internet.

In the ensuing years, technical developers have struggled in trying to scale a system built for an entirely different purpose and set of circumstances to meet the user requirements of today. Most core systems have been patched and repatched, optional addons have been myriad and unevenly applied, and the code and standards have become complex.

As a result, a number of user issues have emerged. These have been identified in a variety of surveys in many nations, and the most prominent are listed below.

Most of these issues are not being addressed adequately. These can largely be attributed to weaknesses in either technical standards or governance structures. The table below correlates user issues with areas where appropriate action might be taken.

Issue	**Technical Protocol issues**	**Governance/management**
Identity fraud	Ease with which identities can be changed or stolen is a basic technical issue which can be addressed	No clear responsibility for security in governance structures and no timely response
Spam	Technical solutions are available involving changes to SMTP and DNS which will eliminate over 90% of spam	No clear responsibility for spam in governance structures and no timely response
Viruses/worms	Technical solutions are available involving protocol changes which can eliminate a lot of virus/worm activity (particularly mass propagation via email)	No clear responsibility for security in governance structures and no timely response
Exposure of children to unacceptable material	Technical developments can assist in this area	No clear responsibility in governance structures
Hacking	Security can be tightened significantly and identity of	No clear responsibility in governance structures

	hackers can be more readily ascertained	
Speed	A variety of factors are involved, some of which are technical	Not essentially a governance issue
Capacity to communicate in one's own language	Can be made easier with changes to protocols	No clear responsibility in governance structures
Affordability and accessibility	Not essentially a technical issues although technical factors are involved	Can be addressed by digital divide initiatives

Later sections of this report contain a lot of additional detail to support these basic observations. On the surface at least, it appears that more attention could be paid to user issues. This report contains a lot of evidence that the user perspective is not a major consideration in Internet governance deliberations.

Leaving aside the issues that currently need to be addressed, it is possible to begin to define the characteristics which users require from tomorrow's Internet. The following statement provides a summary of what is required.

The Internet is for everyone.

The Internet of the future must be

- *trustworthy*
- *reliable*
- *globally inclusive*
- *vendor neutral*
- *easy to use*
- *affordable*
- *able to change rapidly*
- *innovative and capable of significant expansion*
- *transparently and well managed*
- *involving industry, government and community stakeholders*

This statement is to be built on to provide a more comprehensive overview of where the Internet must go.

4.3 Examination of emerging trends and issues which will affect future usage and their impact

The Internet is changing daily. Some of the more predictable and certain developments which will affect things in the next twenty years time frame are:

4.3.1 Size, scale, and speed

Our future Internet, rather than having 600 million users, may have close to 6 billion. So we may only be 10-20% of the way there, and there is a lot of growth to come.

It will also be potentially at least, much faster. As permanent connections replace dialup in many parts of the world, and as higher speed devices, processors and network backbones become available we can expect to see the same phenomenal growth which has characterised the last 20 years of the Internet. Already in Internet advanced countries, permanent connections from homes now exceed dialup connections.

This will place new demands on the infrastructure and the way it is held together. Those experimenting with future networks with increased speed are already suggesting significant problems will exist coping with the increased size, scale and speed of the Internet. A lot of the concern is with the TCP/IP protocol.

4.3.2 Global inclusiveness

But the growth will also bring other issues. One of these is multilingual domain names. As 80% of the people on this planet don't speak English as their first language, there is a natural desire to be able to use their own language on the Internet. Now that's difficult at present, because the core of the Internet finds all those difficult foreign characters hard to handle. But it's unlikely this issue will go away. It may involve some significant changes to the Internet.

4.3.3 ENUM and convergence with telephony

Digital convergence – the crossover and merging of separate communications and broadcasting media – has been on the agenda for some time. One interesting development in this area is called ENUM –

which is a new standard that allows every telephone number to become a world wide web address. So one day, in the not too distant future, we will have worked out easy ways to send instant messages from mobile phones to computers and back again. Messaging, and particularly instant messaging, something of a new genre for the Internet – are here to stay, and are only going to get better.

The other thing which will grow through ENUM and related developments is what they call voice over ip – or internet phones. Already we are seeing these being adopted both by large corporations in internal networks and by hobbyists in networks such as SKYPE – this isn't going to go away either, because at this point of time it may offer significant savings in running costs as compared with old fashioned telephony. It will take a while because telecommunications companies aren't exactly nimble, and the current DNS won't be able to handle this effectively, but one day the convergence of Internet and telephone futures will arrive.

4.3.4 Wireless and mobility

Wireless and mobility are again trends we can expect to see more of in tomorrow's Internet. We are already seeing the growth of wireless hotspots for mobile travellers in airports, hotels and other places, and we are seeing a growing range of mobile devices. We have talked about "anywhere, anytime" access for a long while – we can certainly expect to see that grow.

4.3.5 Peer to Peer

Another thing we can expect to see is a lot of new developments in what is called the 'peer to peer " space. If you know what Napster was, you can see what peer to peer is. Peer to peer is unlike a traditional network with a central computer through which all traffic passes – peer to peer allows almost direct communication on the network with any other computer, for tasks such as trading music and files. Napster spread like wildfire across the Internet, and since then we have seen many other similar developments.

Why Napster was able to grow so quickly, when changes to protocols such as the attempted introduction of IPv6 has been so slow, is instructive, and at least two reasons emerge.

The first is technical. Because Peer to Peer is an end application that involves no change to the central protocols, there were no significant issues involved in putting it onto the network (and, in essence, no need to consult with any technical body to gain approval to do so). IPv6, on the other hand, affects a core protocol and therefore raises a number of adoption and interoperability problems.

But there is another significant reason. Napster adoption was driven by free music (and later video) downloads. Legal or illegal, this was a significant incentive to do something. For most users, however, IPv6 offers no such incentive.

4.3.6 Ongoing problems

Not that Napster was popular with the authorities! There are many things people would like to control on the Internet, including illegal software, music piracy, and pornography. Although there are some technical advances which can assist, it would be naïve to expect these to completely disappear or be completely under control in the next twenty years (pornography, for instance, has thrived for centuries in various forms despite attempts to control it) .

It is also naïve to expect that use of the Internet for illegal purposes will disappear (just as it is unlikely that all criminal activity using telephones can be controlled). Much can and should be done in these areas, and technology and better protocols can help, but they are not going to create a perfect world when the world contains imperfect people.

4.3.7 Affordability

Another thing that will become apparent as this all happens is the necessity for access for all people in all countries, at affordable rates. Once a medium goes so far with penetration of usage, it starts to become an economic necessity – and at the same time, a human right to have access.

The human rights argument for a right to communicate has been in UN circles since the 1980s, and in more recent years has begun to include access to the Internet. At both global and local levels we can expect the arguments for affordable and inclusive access to continue to be advanced.

4.3.8 Conclusions on future needs

What does that say for the protocols and for governance? Well, the only thing that is for sure is that things will continue to change. We can expect to see a larger role for the United Nations, and we can expect some major protocol changes.

4.4 Examination of perceived problems with underlying technical standards in the light of current usage

To examine the issues which prompted this study further it helps to understand the nature of some of the core systems, what they do, and their origin.

4.4.1 Overview

This is of necessity a cursory examination. It is not meant to be a detailed technical study for engineers. The overall situation is far more complex and the numbers of standards involved are in their hundreds, if not thousands. Nevertheless, for the purposes of this analysis, it seems important to understand that the core systems, rather than the periphery, are contributing substantially to today's Internet problems. This is important to understand because there is a great deal of understandable reluctance at this point of time to deal with the core protocols.

4.4.2 TCP/IP

4.4.2.1 Introduction

TCP/IP (or TCP and IP) have been described as the pair of protocols that make the Internet what it is. To purists, use of TCP/IP is the fundamental distinction that describes the Internet.

Essentially, TCP/IP describes system-neutral protocols for transportation of data across the internet between different systems.

Invented in the 1970's, largely adopted in the late 1980s, TCP/IP hit its first big problem in the early 1990s when it became apparent that the numbering system (see DNS) was going to run out of numbers in the foreseeable future.

Therefore in 1995, after several years of work, TCP/IP V6 was released to solve this problem. Adoption has been very slow – Indeed, it was only in the latter half of 2004 that TCP/IP V6 was completely added to the Internet root servers.

It appears that TCP/IP's main advantage is its capacity to scale backwards to existing old systems. Apart from that it appears to be in need of fairly significant modification for scaling to the future, where voice traffic, Internet television and other factors may demand a more sympathetic base protocol.

TCP (The Transport Control Protocol) in particular has come in for significant criticism, and a growing body of experts believe it will need to be replaced. Indeed, if it were easy to change a fundamental Internet Protocol this may have been done some time ago. It's the complexity of the change management problem that has delayed action rather than lack of recognised need for change.

4.4.2.2 Issue – traffic prioritisation

Although TCP/IP has proven to be remarkably robust, it may not scale to the future. In particular, TCP/IP does not know how to differentiate between traffic priorities (eg visiting a website requires a fairly immediate response as soon as we click on it, email delivery can wait a few seconds). This lack of prioritisation is one of the major causes of the "slowness" of the Internet as perceived by users (real speed is something quite different and has a lot of other factors).

4.4.2.3 Issue – unsuitability for financial transactions

As pointed out by Dr Greg Adamson,

"A financial transactions architecture must be deterministic: the result of the transaction in the overwhelming majority of cases has to be what was meant, and when it is not there should be evidence of what went wrong. The design of the Internet protocol suite TCP/IP is non-deterministic. It aims to achieve overall reliability in a network, not necessarily individual reliability for each segment of that network. This concept of 'best effort' is core to the Internet's design and to an understanding of the Internet's flexibility. While the telephone network will reject an attempt to connect if the destination is unavailable (a busy signal), the Internet will send information out in the hope of success, by design (best effort). There are many methods for overcoming the limitations that this creates, including within the TCP protocol itself, but the design choice of 'connect if a full service available' or 'make every effort to get any part of the message through' remains. The original design specification states:

> *[The Internet Protocol has] no mechanisms to augment end-to-end data reliability, flow control, sequencing, or other services commonly found in host-to-host protocols. The internet protocol can capitalize on the services of its supporting networks to provide various types and qualities of service (Postel 1981)."*

4.4.2.4 Security issues

There are also security issues with TCP/IP, with researchers warning of vulnerabilities that need to be addressed. In April 2004, a major alert was issued to deal with a fundamental vulnerability.

4.4.2.5 Performance issues

Users of large scale sites are already experiencing problems with the protocol, which tends to suggest that ordinary users will become affected in the near future, as bandwidth and processing availability continues to grow.

4.4.2.6 Conclusion

TCP – if not TCP/IP - needs to be replaced, probably within a five to ten year time frame. The major issue to overcome is the migration issues (see below)

4.4.2.7 Migration Considerations

The problem of a new TCP is as complex (if not more so) that the TCPIPv4/v6 changeover which the Internet community has found very hard to deal with.

However, the factors in slow IPv6 deployment all revolve around the fact that there is no communicated compelling reason to change. Given that a point of time will arise when changes to TCP are necessary for basic performance, it can be expected that, if a migration is conducted with appropriate change management planning, the adoption will be far quicker and far smoother than the IPv6 changeover. However, some basic factors need to be taken into account.

4.4.2.8 A compelling reason to change has to exist

No major change to core Internet protocols is going to be undertaken unless a compelling reason exists.

4.4.2.9 A compelling reason to change has to be communicated

Similarly, a major communication project is needed for a change of this magnitude. Trying to do it without formal change management planning would not succeed.

4.4.2.10 Involvement of major players will help

In particular, Microsoft and Cisco should be engaged, and planning for total adoption should take into account their adoption schedules. Once again compelling reasons to change have to be communicated: the cost to the market leaders to make the change across all product ranges is substantial, and the economic reasons to stay with current configurations may be compelling.

We may be talking about changes to every connected device on the planet here. This could exceed 1 billion devices.

4.4.2.11 Co-existence with old systems has to be allowed but only for a limited period of time

There cannot be a one day changeover. This is addressed elsewhere in the Internet Mark2 analysis, but for a quite substantial period of time the new system must co-exist with the old.

However this cannot be forever. There must be a date where the new Internet is able to move forward unencumbered by legacy systems. That is essential to forward development.

4.4.2.12 Conclusion:

The efforts to change TCP/IP must continue, and in the medium term appear to need to take on additional momentum.

Change should be considered in association with other core protocol issues identified in other drafts.

4.4.3 DNS and WHOIS

4.4.3.1 Introduction

Each host on the Internet has a range of IP numbers. The Domain Name System (DNS) maps the numbers to names of hosts or websites (eg www.google.com, www.hotmail.com). Thus, when a user enters a name, the Internet knows which number to send the query to by looking up the DNS database.

It should be noted that the other widespread user of distributed network infrastructure, the telephone system, operates quite differently in several respects. It has no domain name equivalent with trade mark implications in normal uses – to contact a telephone address, you simply enter the number.

The DNS was introduced in 1984, several years before commercial traffic was able to be part of the Internet.

4.4.3.2 WHOIS

At the same time, a public database called Whois was introduced, essentially to allow technical managers of hosts to contact their peers. This is the Internet equivalent of a telephone directory, but also serves a number of related purposes. WHOIS stores details of the names and addresses of domain owners and technical contacts. It was named after a UNIX operating system command (whois) which gave basic details about system users.

4.4.3.3 Issue - Multilingual domain names

4.4.3.3.1Introduction

More than 80% of people on this planet do not use English as their primary language, and even at this early stage of adoption over 60% of Internet users are not primarily English-speaking. That percentage will grow, but already the demand to be able to communicate on the Internet in ones own language has been expressed very clearly.

It has been possible for some time to include most languages in the content of web pages or email messages, using tagging techniques which indicate what language set is being used.

However it has not been possible to use native languages in email addresses, domain names, DNS listings, and the WHOIS database. This poses significant barriers to adoption for non-English speaking people.

The main problems here are that

- ASCII (the American Standard Code for Information Exchange) is used in the DNS. ASCII is incapable of supporting the complexities of foreign language.
- No-one wants to make major changes to the DNS.

To these must be added the fundamental issue in the evolution of Internet use which has seen functions originally designed for system unique identification purposes become part of corporate branding and presence on the Internet. This has led to significant additional complexity – a brand may require components which are far more complex than simple unique identification. This adds substantial complexity to the demands for multilingual domain names.

A significant amount of work to rectify this problem has been ongoing since at least 1998. In 2000 the IETF set up an IDN Task Force to work on appropriate standards. In 2001 ICANN joined in with a committee addressing policy issues in this area.

The most promising technical direction would appear to be adoption of UNICODE as part of DNS – however, IETF, being reluctant to change the DNS, is pursuing a range of other solutions. Most of these resolve around the idea of PUNYCODE, which essentially creates a gateway for understanding UNICODE in ASCII. However, this partial solution leaves open the possibility that 2 different domain names in different languages could return results with the same unique identifier in ASCII.

Thus an associated issue arises. The problem is much simpler if purely local solutions are adopted within each country using character sets suitable for the local situation that don't have to apply elsewhere. However, if these are to be identifiable and unique as part of a global database, another set of problems emerge.

And this of course is complicated by the complexities of written language itself. Some languages read left to right, others right to left; most languages have regional variations (even English, with different spellings in European and USA use); some character sets give rise to vastly different groups of languages. Thus, for a system which has become part of corporate branding in a multitude of languages with regional variations, the question of uniqueness on a global scale combined with local interpretation in the exact form that populations are used to seeing the brand, becomes quite complex.

Associated issues surround the use of WHOIS data. English users are used to being able to find out something about a domain name in a language they understand. To extend this facility to every user's language of choice would immediately remove the possibility of either machines or users of other languages being able to read the information.

In this context, it is necessary to address the issue that the World Summit on the Information Society (WSIS) in its Declaration of Principles (2003) stated

"The creation, dissemination and preservation of content in diverse languages and formats must be accorded high priority in building an inclusive Information Society, paying particular attention to the diversity of supply of creative work and due recognition of the rights of authors and artists. It is essential to promote the production of and accessibility to all content—educational, scientific, cultural or recreational—in diverse languages and formats. The development of local content suited to domestic or regional needs will encourage social and economic development and will stimulate participation of all stakeholders, including people living in rural, remote and marginal areas."

Internationalised domain names (IDNs) have become a fundamental part of and an iconic symbol for the digital divide issue. ICANN has been criticised at its regular Public Forums for not giving the matter sufficient attention, failing to make significant progress, and being negative in its analysis of the issue.

The lack of progress gives rise to some significant questions, eg

- How much different would progress be if a team of experts were employed full time to work on the problem, rather than relying on volunteer effort?
- All attempts to solve the problem are predicated on the IETF desire not to change the dns infrastructure – what would happen if that was not a consideration or we had sufficient other reasons to change DNS as well?
- Is there a way quite separate to our current understanding of DNS and use of domain names to create an architecture which meets the branding and unique memorable identifier needs of Internet corporations and users while retaining the routing needs of a global Internet?
- Will it be better in solving this problem to look at a body related to but separate from IETF and ICANN to bring in funds and devote dedicated expert resources to solving this problem?
- To what extent is the revised Verisign test bed actually solving the basic problems anyway and what dangers (if any) does the revised test bed present?

4.4.3.4 Issue – refresh and size

DNS was never supposed to scale to its current size of over 600 million entries. One of the issues with DNS which will get worse as the Internet grows is “refresh” rates:

The DNS operates via the co-ordination of servers all over the world which update each other and provide the latest details on where a domain is located. However, because of the size and distributed nature of the DNS, a new or changed site address may take some days to propagate across the Internet to the various DNS servers that direct traffic.

Although this significant delay in DNS refresh times is a feature of Internet management, it is hardly a desirable one or one which would be deliberately architected. One of the issues arising from this arises in the implementation of ENUM – a standard which links the telephone system to the Internet system, thus potentially allowing SMS messages (a telephone system application) and instant messaging (an Internet application) to be able to interchange across devices. That is not useful in a situation were it may take days for the DNS to realise that a user has gone to the shop and would like to receive instant messages on their mobile phone for the next half hour or so. Increasing mobility in Internet use will demand resolution of the refresh issue.

Paradoxically, given that the Internet is considered a decentralised system and the telephone system is perceived as centralised, the telephone system has a separate and quite decentralised architecture, with nothing equivalent to the DNS root infrastructure. Each telephone operator maintains their own authoritative routing info, or subcontracts that function to somebody.

DNS is a first generation system that could be improved in several respects to meet new needs which are emerging.

4.4.3.4.1Issue – WHOIS and Privacy

Whois has in the past been a completely open system, but again has not scaled to current uses and needs. Many of the current days Whois issues are to do with privacy; once again, ENUM comes to the fore as a system which brings new needs to DNS, particularly the need to have the "silent number" concept of telephony carried over into the world of the Internet. Other problems associated with the current system include the ease with which contact details are used for spam mailing lists, and the nuisance domain renewal business which exploits the openness of the database. These factors in particular are leading to calls for changes to WHOIS so that certain personal details are kept private.

4.4.3.4.2Issue (non-technical) – Use of domain names

DNS also gives rise to a raft of complicated non-technical issues. Because names are used instead of numbers to direct users to sites, use of particular names and combinations of letters give rise to legal disputes about usage e.g who is entitled to use Woolworths.com. A surprising industry has sprung up around domain names and their usage, with accompanying international law and governance needs. It should be realized that, for a functioning Internet, none of this is actually necessary, but is a result of design decisions made in the 1980s. Dotcom or dotanything is not a necessary or functional part of a useful Internet, but a design decision with substantial ramifications from which an industry and regulation and governance structure has emerged.

4.4.3.4.3Issue – security and DNSSEC

As Diane Davidowicz and others have pointed out, the distributed management of the DNS, and support for the redundancy of DNS zones across multiple servers, give it a robustness which is highly desirable. However, the original DNS protocol specifications did not include security. Without security, the DNS is vulnerable to attacks of various types, and the accuracy of the information contained within the DNS is vital to many aspects of Internet based communications.

The threats that surround the DNS are due in part to the lack of authenticity and integrity checking of the information held within the DNS, and in part to problems with other protocols that use DNS host name information.

The various problems which arise carry colourful names such as rogue servers, DNS spoofing, cache poisoning, and client flooding.

4.4.4 Mail systems and SMTP

To all internets and purposes, email is already broke, and must be fixed. The Internet's first and greatest killer application is now problematic.

In a survey examining email usage in 2003, the Pew Internet Project found that

- 25% of email users stated that the ever increasing volume of spam has reduced their overall use of email
- 70% of email users claimed spam had affected the quality of their on line experience
- 30% of users expressed fears that filtering approaches would cause loss of wanted mail
- 76% f users are bothered by offensive or obscene content in spam email
- 80% of users are bothered by deceptive or dishonest content in spam email.

The problem is far more serious in personal email accounts than work email accounts,

Costs associated with spam have been estimated by various research firms at between $10 billion (European Union, 2004) and $87 billion (Nucleus Research, 2003) per annum. Spam volume is now estimated to exceed legitimate email volume; in May 2004, 76 percent of inbound mails scanned by email security provider MessageLabs Ltd were spam, up from 67 percent a month earlier.

Email users world wide are subjected to a wide range of preventable fraud scams, viruses, spam, etc which is rendering the medium useless.

Meanwhile, ICANN claims spam issues as out of scope. *"... issues of concern to Internet users, such as the rules for financial transactions, Internet content control, unsolicited commercial email (spam), and data protection are outside the range of ICANN's mission of technical* *coordination"* (ICANN website). IETF has been very slow at doing anything in this field, preferring to leave investigation of the issues to a separate Internet Research Task Force (IRTF) group.

As a result there is a general belief that nothing technical can be done to prevent spam. However, our analysis suggests that the existing protocols are significant contributors to the problem, and protocol reform could see spam volume drop by at least 80%.

4.4.4.1 SMTP issues

SMTP, or the Simple Message Transfer Protocol, is the basic standard for email, and exists since the 1980s when the Internet was small and honest.

Perhaps more than any other system on the Internet, email has seen a number of improvements and different protocols, each of which have been adopted by only part of the Internet email community. This capacity not to adopt standards is a feature of the Internet, making dealing with change more difficult than it otherwise might be.

SMTP comes from an innocent age, when no-one thought it would be necessary to prove that the person sending a message was who they said they were. The basic flaws in smtp authentication are now causing significant problems, particularly the ease with which email sender details can be forged. This helps the transfer of some viruses and a lot of the worst spam, and makes Internet fraud a lot easier than it might otherwise be. It must be said that not all viruses and spam can be attributed to problems with protocols, but better protocols would make the situation substantially better.

The basic reason spam and viruses exist most often in email is SMTP. SMTP is the online equivalent of borders without checkpoints and passports, or bank vaults without doors and locks. Some of the SMTP security weaknesses are:

- It allows anyone to connect with anyone without any system to say who they are

- It is simple to forge messages and pretend to be someone you are not with no checking whatsoever
- Not being one to one like telephone calls, it is easy to mass market to millions of email addresses at very low cost.

These issues have been known for some time. Various attempts to provide improved protocols have been undertaken, but essentially have resulted in a mass of conflicting systems and standards. As a result, change is becoming more complex to initiate.

Email upgrades are complicated by

- Old systems which are never upgraded
- Incorrect applications of email systems
- The variety of applying protocols (eg http for webmail, smtp, nntp, pop etc)
- The ubiquitous nature of email
- IETF inertia in handling big problems.

4.5 Examination of governance structures as they have emerged (IETF, ICANN, RIRs, etc)

The emergence of Internet governance and management as been evolutionary rather than a response to any examination of structural needs in a typical business sense. The primary building block had been volunteerism.

No one body runs the Internet. In fact, it is not clear whether the Internet has a governance structure at all at this stage; rather, it has a series of interrelated bodies co-ordinating parts of the system.

However, the following bodies play major roles.

4.5.1 IETF

4.5.1.1 History

The Internet Engineering Task Force (IETF) is the standards body for the base Internet protocols. It is one of three relevant standard bodies in Internet governance: The first is the ITU, which originally rejected a role in computer standards development, and the other is the Worldwide Web Consortium, which separated from IETF in 1994 because it considered IETF as incapable of the standards development needed for the World Wide Web.

Founded in 1986, the Internet Engineering Task Force describes itself as "a loosely self-organized group of people who contribute to the engineering and evolution of Internet technologies specifications". The IETF is unusual in that it is not a corporation and has no board of directors, no members, and no dues.

The first IETF meeting was held in 1986 with 21 attendees. The 7th IETF in 1987, was the first meeting with over 100 attendees. The form of management has remained basically unchanged since 1992.

IETF's strength is its capacity to draw together in a standards-making situation stakeholders from all the major industry players. Its weaknesses however are substantial.

4.5.1.2 Issues

IETF is to be praised for discussing its problems openly, and seeking to address them through consulting its participants. It's own Problems Working Group came up with the following analysis in 2003.

"The IETF started off as a small, focused organization with a clearly defined mission and participants who had been working in this area for a significant period of time. Over the period 1989-1999 the IETF grew by a factor of ten or more in terms of number of participants, and volume of work in progress. The effects of this growth have been compounded by the extension of the scope of the IETF which makes the work much more varied. Also during this period, the Internet has become more complex and the requirements placed on it by a far larger user community have

changed as the network has come to have a pivotal role in many areas of life.

Many of the problems and symptoms appear to be fundamentally caused by the organization failing to fully adapt its management structure and processes to its new larger size and the increased complexity of the work. The IETF has also failed to put in place a clear definition for its future mission now that the initial mission has been completed or outgrown."

The analysis goes on to conclude:

- IETF is unsure about what it is trying to achieve
- Cannot determine what its scope should be
- Is unsure who its stakeholders are
- Cannot prioritise actions effectively
- Loses sight of overall architecture.

4.5.1.3 Practices

IETF has also been accused of not being aware of standard project management practice, including:

- Stakeholder interests are rarely if ever analysed
- Change management plans are rarely prepared
- Standard risk analysis and task management etc seem to be out of scope
- Dependency identification, process tracking and other standard project management approaches are not rigorously used.

This lack of project discipline and knowledge is a major operational constraint which, when combined with the volunteerism inherent in IETF approaches, leads to long time frames for adoption of standards, poor takeup of standards when they are adopted, and a generally scattered approach. Again quoting from the IETF Problems Group,

”The acceptable deadlines for finishing a piece of work, and the criteria used to determine them, are rarely, if ever, documented. Also the estimates of the time required to complete the work often differ widely from the time actually taken. The combination of these factors makes determining the feasibility of delivering within the required timeframe and then adjusting the scope of the work to fit the timeframe requirements extremely difficult.

Another concern is the difficulty IETF has in dealing with large or complex problems. Interdependencies seem to be difficult to manage. Although the establishment of the Internet Research Task Force (IRTF) has assisted in complex analysis, complex problems remain unaddressed.

4.5.1.4 Governance

As mentioned above, IETF governs itself , notwithstanding the existence of internal committees known as the Internet Engineering Steering Group (IESG) and the Internet Architecture Board (IAB). It makes its own decisions without any particular mechanism for seeking input from non-technical users. Again quoting IETF,

“IETF meetings are of little interest to sales and marketing folks, but of high interest to engineers and developers.”

Most participants at meetings are from North America.

Difficulties in getting results are compounded by the fact that IETF is essentially responsible to no-one, and is a classic technocracy in an era when such structures have been largely abandoned (the Klinger-Cohen Act in USA, for example, mandated establishment of structures within the USA Government for management of technical facilities to include the involvement of business and user interests).

IETF governance contrasts substantially with the other two standards organisations involved with Internet standards. ITU has the strongest governance structure, being responsible eventually to member state representatives, and W3C standards work is determined and prioritized by a member organization. So in this respect, IETF is peculiar.

And this peculiarity brings with it certain problems because, in reality, few issues if any are purely technical and have no policy repercussions.

4.5.1.5 Case study – Email problems, MARID and intellectual property rights

IETF had one of its regular meetings in Minneanapolis in late 2003, when the world was reeling under the effect of spam and viruses, IETF didn't discuss spam as part of its meeting agenda. The reasons appear to be

- IETF is a voluntary association of engineers with no formal(or even informal) avenues for user input in determining priorities

- The nature of IETF working groups and their attempts at "rough consensus" make it easy to block any forward progress (in the case of spam, there is a multi million dollar spam industry and a multi million dollar spam filtering industry both financially affected if progress is made).

However by early 2004, when it was obvious that major players such as Microsoft, Yahoo, Sendmail and AOL were going to implement new standards with or without IETF involvement, IETF convened its MARID Working Group to examine the problem. As a result of participation from two of the major proponents, Microsoft with its Caller-ID proposal and Meng Wong's SPF proposal, a merged proposal emerged called Sender-ID.

However, Microsoft wished to retain certain rights to development work under licence, while making it available to others. Thus, in addition to deciding on technical issues, the work group was faced with some complex licensing situations.

As this article was being written the complexities of this development were being played out within the technical working group. But whatever the outcome is, the issue is that IETF as a volunteer technical standards organization, has neither the expertise nor in reality the scope to arbitrate on licensing and intellectual property matters (or to make policy decisions in isolation on what sort of licences should or should not be granted). This lack of expertise and lack of structure to deal with issues such as this is a major structural flaw.

Sender ID will help the situation if it survives the last throes of IETF approval, but the IETF process has resulted in various compromises to achieve consensus. Participants, knowing that something is better than nothing in this area, have co-operated and compromised to get a result.

However the result is far from optimal. Indeed it is unlikely that spam and virus issues will be addressed in any significant way until both SMTP and DNS are changed.

4.5.1.6 Management of standards adoption

Similar flaws exist in other areas. IETF's lack of change management expertise, for instance, has derailed and delayed needed standards implementations. IETF members are not to be blamed for this – a group of very talented people are giving of their time to do what they know best – but the overall structure is clearly problematic.

There is an associated issue as regards adoption of standards. Standards adoption is not compulsory, and standards can tend to overlap each other. Therefore IETF outputs (some thousands of standards) have mostly been adopted by small parts of the Internet community only, making any new changes increasingly difficult to deal with. Complexity is rising each year and making further change more difficult at a time when some basic changes seem necessary.

4.5.1.7 Case studies – Ipv6

The record on large projects can be particularly traced in the stories of IPv6 and DNSSec. Both have taken many years to develop and have experiences significant difficulties as regards propagation.

IPv6 had its formal beginnings in 1991.Although other issues were involved, a major catalyst for activity was the realization that the Internet would run out of IP addresses (numbers which are allocated to specific Internet sites) in the foreseeable future.

Work then began on the next generation Internet protocol, and many issues were put on the table for discussion but eventually dismissed as the protocol took on a manageable form. Traffic prioritization was one thing that went by the board.

The first recommendations for the new protocol emerged in 1994. After some discussion they were finalized in 1995. It took till April 1996 to move much further forward, and to define the transition phases as to how the network would adopt the new protocols.

That might be regarded as the beginnings of an adoption phase – some five years after the need was agreed to. However, in 2004 – some 13 years after the problem was first worked on- implementation is at a low level. Not even the central root servers of the Internet had adopted the protocol by June 1994.

No detailed analysis as to why adoption has been so slow appears to have been undertaken. It is not clear whose responsibility such a study or adoption would be in any case – that would be beyond the scope of most standards bodies, but there is no other body to take up the slack. (and no body to review the performance of IETF on a regular basis).

This of course has a significant effect on Internet performance.

4.5.1.8 Case study - DNSSec

A similar story emerges for DNSSEC, the much needed security architecture for the Domain name System (DNS).

The Internet was designed without a security layer. By 1990 it was obvious that changes would need to be made.

However it was five years later, in 1995, that DNSSEC (as it was to become) first became a topic within IETF.

RFCs began to appear in 1997, and in 1999 RFC2535 was produced. Little happened in the way of deployment, and by 2001 it was acknowledged that there were significant deployment problems. So the standard was rewritten and expanded.

In 2003 substantial tests suggested that DNSSEC was ready for deployment. However another rewrite of the standards began, and current expectations are that the work may complete in 2005 and deployment may begin.

Thus, 15 years after a problem became obvious, and 10 years after work began, DNSSEC will not have commenced deployment on one of the most serious flaws in the Internet protocols. Given the long slow and ineffective deployment of IPv6, it's impossible to guess how long DNSSEC deployment might take.

No-one can attribute these long time frames to technical complexity alone. Poor methodologies, volunteerism, under-resourcing, unprofessional behaviour and management issues have all contributed to delays, given the importance of the tasks to the performance of the Internet. Retro-fitting security was always going to be complex, but the time being taken and the difficulties being experienced are only partially due to the technical issues themselves.

IETF participants have raised a number of concerns about the way that individuals can sabotage group efforts – indeed individuals well up in the hierarchy of IETF have been accused of this at times.

4.5.1.9 Acknowledgements

The positive work of IETF and its members is acknowledged. Without them the Internet we know today could not exist.

However these volunteers now labour under an ineffective structure that is not conducive to innovative and timely developments for the Internet.

4.5.1.10Conclusion

IETF's decisions to address its problems in an open forum are to be applauded, as are its attempts to engage a wide global audience of engineers in its consensus based decision making structures.

However, IETF is a classic technocracy. There needs to be a mechanism for determining priorities that include user perspectives, and for prioritizing workload to ensure that the most urgent issues are dealt with.

IETF appears to be reasonably capable of managing the day to day concerns as regards maintenance of standards, but does not have the capacity to tackle major tasks or major change. This is not uncommon: the vision of those charged with day today operations often cannot extend to future directions.

To solve these problems, IETF would need to

- Have a clear relationship to a governing body with the competence and representative nature to deal with policy considerations which are outside the scope and expertise of IETF
- Have means of introducing non technical skills sets such as project and change management to its affairs
- Be resourced to provide dedicated rather than volunteer effort in working on major priorities
- Have a clear scope, mission, and relationship within an overall technical management structure for the Internet, and with other Internet bodies of a non-technical nature which fulfil complementary roles in determining policy and assisting adoption of standards
- Learn to communicate effectively with business, community and government stakeholders.

As other parts of this study deal with issues regarding other Internet technical co-ordination bodies, other non-technical issues, and with protocol issues currently administered by IETF, it is inappropriate to consider this section in isolation in determining more appropriate structures.

4.5.1.11 Migration Considerations

Many within the IETF community will regard this report as both timely and important. However, there is something of the "old boys club" within IETF, with its own peculiar pecking order and structures and traditions. It would be unusual if this report was not criticised by some current IETF participants.

There will be resistance to change. However that does not negate the necessity for establishment of more appropriate Internet governance structures able to deal with the realities of a 21st century Internet.

4.5.2 W3C and ITU

4.5.2.1 Introduction

Some mention needs to be made of the other standards bodies involved in Internet Governance, ITU and World Wide Web Consortium. These organisations sit parallel to IETF (and IEEE) in the technical standards development area.

4.5.2.2 ITU

Unlike IETF, which is a loose body anyone can join, ITU's origins is an agreement between nations, and emerged from the first International Telegraph Convention in 1865. It became a UN specialist agency in 1947.

As technology developed, ITU's scope expanded progressively from its original wired networks scope to include wireless and satellite networking.

ITU may have had a major role in Internet in the early days. However, in the era of analogue telephone networks, and before computers had any effect on telecommunications networks at all, it appeared disinterested in what were effectively computer communication standards. ITU's disinterest in emerging technology effectively outside of its level of expertise was a major factor in the mergence of IETF as the main Internet standards body.

It should be remembered that, at the time IETF was established in 1986, telecommunications companies were not major players in the emerging Internet. They became more involved from 1990 on, as a commercial Internet got underway. It should also be remembered that in the 1970s the US telecommunications giant, AT&T, could not see a business case for involvement in the Internet. This perhaps is the most telling criticism of the staid and solid elder statesman than ITU is – it will find it difficult to be nimble or innovative in seeing future directions.

Nevertheless, unlike ICANN, ITU has taken a role in examining spam issues. It is also becoming increasingly active on a number of other issues in the lead up to WSIS governance decision making.

An area of interest where ITU and IETF have worked together is the emerging ENUM standard. ENUM bridges the world of telephony and the world of the Internet by making every phone number potentially an internet-addressable entity.

To an outsider, ITU appears to have all the efficiency and capacity to get things done that ICANN appears not to have. However, that perception may well be illusory; and anyone who has been involved in development of standards such as x.400 could point to equally problematic issues in ITU.

ITU, like IETF, is undergoing considerable internally driven reform to try to better cope with the demands of a rapidly changing communications technology landscape. There would appear to be room for some of the strengths of ITU to be better utilised alongside those of IETF and ICANN in the future, particularly as telephony and Internet based applications continue to converge.

4.5.2.3 W3C

The other Internet body involved in protocol development is the World Wide Web Consortium (or W3C) W3C, unlike IETF, is a membership organization that has more control over directing its activities than the more open IETF. It was formed by Web inventor Tim Berners Lee in 1994 to expedite web standards development and move it out of the more transport focused IETF.

W3C specifically addresses the architecture of the World Wide Web. Its area of coverage includes HTML specifications, metadata, etc. It worked closely with IETF in development of the http protocol, and it is at the level of http that the boundaries of concern of the two organisations intersect.

W3C does not technically produce standards as such, but co-operates with the International Standards Association (ISO) and other bodies to do so. It is able to produce results reasonably quickly because it is able to more effectively control activities, but it suffers from a lack of openness.

W3C's mission and goal support a vision of the Web which is open, trusted, interoperable, decentralized, universally accessible and vendor neutral.

4.5.3 Root Servers and Regional Internet Registries

4.5.3.1 Introduction

It is important to realize that ICANN doesn't control everything in Internet technical co-ordination. An interesting history associated with the early growth of the Internet led to a number of quite independent structures being established. ICANN has a series of relationships with these various separate bodies which it is attempting to formalise.

IETF we have looked at separately. But not all work was standards work. There were naming systems and numbering systems that needed administration, and a growing base of root servers (the central computers) which were important to the operations of the Internet.

4.5.3.2 ccTLDs – Country Domain Administrators

As the Internet spread beyond the USA, individuals in the technical community began to undertake management of domain numbers and the new top level domains (such as dot uk and dot au) within various countries on a voluntary basis. The only qualification was technical knowledge. These early country administrators did a great job, but as the Internet grew systems were needed to hand over the management to more representative organisations of various types.

Until 1998, any national sovereignty issues involved in administration of a domain name were virtually unacknowledged. However in the process of establishing ICANN, the role of governments was partially acknowledged for the first time.

Nevertheless, many developing countries still have their domains administered offshore by non-nationals; a sort of techno neo-colonialism.

Changeover was not always able to be conducted smoothly, and in many countries the old volunteer system remains. In cases such as Australia, change to a representative body had to be accomplished after legal proceedings.

Indeed, it has been suggested that ICANN effectively holds countries to ransom over use of their ccTLD domain name, insisting on recognition of ICANN as the ultimate authority on domain name issues. Redelegation has required recognition of the ICANN role, but the clauses are unsatisfactory to many countries.

However, most country domains are now forming appropriate locally defined relationships with their governments and their local constituencies to address these issues.

4.5.3.3 RIRs – the Regional Internet Registries

One of the most time consuming functions was the numbering system. Each Internet site has a unique number or set of numbers (known as an IP address) to identify it to other computers on the network. Once the number of Internet hosts became too many to count on one hand, some sort of system of registering host computers had to be devised. In 1988, as the task grew bigger, the Internet Assigned Numbers Authority (or IANA**)** took over as the core body for the central tasks. Under contract from the US government, it oversaw the allocation of IP addresses to Internet Service Providers. , and a number of other tasks associated with the technical management of the Internet.

The US government maintained control of this system through a contract with IANA during the 1980s. However, in Europe an organisation called RIPE NCC was set up independently in 1992, as the first Regional Internet Registry. Unlike its US counterpart, it was not government run or government funded, but was an association of major stakeholders in the region. In 1993 APNIC was established to provide a similar function in the Asia-Pacific region. The US based Regional Internet Registry, ARIN, emerged as a separate identity a few years later, in 1997.

The RIRs are non profit membership organisations, with their membership coming mainly from ISPs and telecommunications companies in their respective areas.

4.5.3.4 The remaining IANA functions

IANA continued various other co-ordinating tasks until the year 2000, when the remaining functions and what remained of IANA became part of the Internet Corporation for Assigned Names and Numbers (or ICANN), who we will talk about shortly.

However, as we have just learnt, some of the original IANA functions, most importantly those of the regional internet registries outside of the USA, had already been internationalised and set up separate structures of their own.

4.5.3.5 Root server management

Another separate function was management of the root servers.

In order to accommodate a coming commercial Internet, the US government entered into a contract with Network Solutions in 1991 to manage the root server system.

Verisign Ltd still maintains this facility under contract to the US government in 2004; indeed, no change to the Internet root occurs without approval of the US Dept of Commerce.

However, 2004 also saw the growth of the root server system infrastructure so that, for the first time, more root servers existed outside of the USA than inside. Most of the root servers are maintained by volunteers.

In addition to the number registries, there are also name registries, who administer dot com, dot org, or country domains like .us, .uk and dot au. Around these central functions, usually carried out by a non profit body, a growing business has emerged of competing registrars who sell domain names on the Internet. There was a need for co-ordination, and particularly for a means to resolve disputes in allocation of names.

This is the role of ICANN (see next section)

4.5.4 ICANN

4.5.4.1 Introduction

The Internet Corporation for Assigned Names and Numbers (ICANN) exists in its current form largely because the US Government wanted it to be so.

It's structure is an evolving reactive mechanism. Anyone analysing its current structure without regard for the history of how it came to be would have to regard ICANN as

- eccentric in structure
- illogical in scope, and
- incomplete in terms of internet governance.

It may be the best we have got, but it's a long way from the best we could have.

Once the web started to grow and commercial forces came on board in a large way, it was obvious that concerns about name disputes, use of names, etc would have to be controlled in some way.

In 1996 a committee was established to address these issues, including representatives of the Internet Society (ISOC), Internet Architecture Board (IAB), Internet Assigned Numbers Authority (IANA), International Telecommunications Union (ITU), World Intellectual Property Organization (WIPO) and International Trademark Association (INTA). The group initially looked at issues with top level domains (.com, net, org), plus country top level domains (ccTLDs, eg.us, uk, jp, au etc). Right from the beginning there was controversy about the relationship with the US government – not desired by some proponents but stoutly defended by others.

The initial proposal saw establishment under Swiss law. However at the beginning of October 1998 the US Government's National Telecommunications and Information Administration (NTIA) responded to this proposal by announcing the Internet Corporation for Assigned Names & Numbers (ICANN). It would operate under an agreement with the NTIA with oversight by the US congress. The new body was asked to ensure competition in delivery of domain name services.

Thus ICANN became a corporation under US law, with a contract to operate from the US government, despite concerns of many stakeholders. Those concerns remain at the beginning of 2004 and are central to the WSIS concerns about Internet governance.

ICANN Mission Statement

According to ICANN,

"The Internet Corporation for Assigned Names and Numbers (ICANN) is the private-sector body responsible for coordinating the global Internet's systems of unique identifiers.

The mission of ICANN is to coordinate the stable operation of the Internet's unique identifier systems. In particular, ICANN:

1. Coordinates the allocation and assignment of three sets of unique identifiers for the Internet:

- *Domain names (forming a system referred to as "DNS");*
- *Internet protocol (IP) addresses and autonomous system (AS) numbers; and*
- *Protocol port and parameter numbers.*

2. Coordinates the operation and evolution of the DNS's root name server system."

All of that appears to be in line with ICANN's often stated mission of being a technical co-ordination body. ICANN states specifically that its scope is limited to technical matters:

"Other issues of concern to Internet users, such as the rules for financial transactions, Internet content control, spam, and data protection are outside the range of ICANN's mission of technical coordination".

However, because of the eccentricities and incomplete nature of Internet governance structures, ICANN has consistently worked in areas that cannot be regarded as technical co-ordination.

For instance, in 1999 it succeeded in establishing a Uniform Dispute Resolutions Policy (UDRP) for the top level domains; hardly a technical co-ordination task, but certainly a useful one for development of the new media.

Similarly eccentric is the role of ICANN in creating a competitive environment in DNS, part of its contract with US Department of Commerce. This would normally be seen as a regulatory body's responsibilities, not a technical co-ordination task.

ICANN has experimented with various structures to accommodate its various governmental, industry based and community stakeholders. It attempts to provide forums for industry, registrars, registries, business non commercial users, standards bodies and governments to have input. That's a difficult task, which most people would agree that ICANN has done well, but not perfectly. People refer to the ICANN circus, with its 3-4 meetings per annum in exotic locations.

In analysing ICANN, three components of Internet governance have been suggested:

- Technical engineering functions that allow different components of the Internet to interact
- Technical co-ordination of the key protocols and addresses and names that underpin the technical functioning of the Internet
- Public policy matters

Public policy matters where ICANN is active include intellectual property issues and security. Public policy matters where ICANN is not active include spam and consumer protection. Once again, the logic of involvement and non-involvement is not easy to follow.

Perhaps partially as a result of this mission confusion, ICANN does not handle public policy well or effectively. An example of this was its attempts to gain public input in to the WHOIS database and privacy issues.

ICANN set up three task forces to examine the subject, and at its Rome 2004 meetings devoted a day to obtaining input. Due to poor response, deadlines were extended for input. The methodology attracted input from ICANN insiders (registrars, registries, etc), but very little from Internet users. This methodology to come up with the obvious was convoluted and time consuming.

4.5.4.2 Conclusion

The problem with ICANN, and with IETF, is one of defining scope within a schema that effectively manages all needs of the 21st century Internet. No such schema exists, of course, and that is why bodies such as ICANN and IETF are continually operating in areas outside of their level of competence in order to keep things afloat.

If there is a problem in Internet governance, it is the gaps between the competencies of existence governance bodies and the needs of Internet industry governmental and community users. As user needs do not come within the range of concern of any particular Internet governance body, it is inevitable that mistakes are being made and crucial issues are not being addressed.

5 Recommendations

There are many problems facing the Internet at the current time. Those closest to each problem are making substantial efforts in isolation from other major issues to address their particular problems, often however seeking short term fixes rather than longer term solutions. A more co-ordinated approach across all areas seems necessary for resolution.

IETF and ICANN have been unable to deal with the range of issues and concerns adequately. This is acknowledged as an issue by many people who are aware of the problems.

The question then becomes how to proceed to reform.

5.1 Governance reform

It is clear that the WSIS governance debate will be ongoing until at least mid 2005, and at that point of time is only likely to consider what governance structures should be. This study acknowledges that an appropriate forum exists to examine governance issues, and suggests that some of the information in this study might assist in discussions of how to appropriately reform governance. It is not simply a matter of removing ICANN from the somewhat benign influence of the US Dept of Commerce; it is necessary to consider the range of issues involved in creating tomorrow's Internet as an effective and equitable tool for human development. That requires some refocussing of discussion to address some wider issues and needs.

5.2 Protocol reform

What is not being discussed as comprehensively is how to reform protocols. This study concludes that protocol reform is necessary.

Project management structures are both the normal and the most effective way to handle major change projects of this nature. A separate short term structure should be established for this purpose. It would consist of:

- A dedicated full time project team, containing the best engineering, project management, research, change management and communications expertise available to scope and deal with this problem. Other skills sets will also be necessary as the project progresses.

- A Project Steering Committee with major stakeholder involvement to review progress and determine direction.

The exact makeup of both the Project team and the Steering Committee needs further investigation. A further study is suggested with major stakeholders to determine some effective ways to pursue this.

It will be the job of the project team to

- examine options for protocol reform
- examine options for technical development
- suggest timetable and budget for a series of interdependent projects
- determine resources needed
- determine migration and change management strategies
- work with the Steering Committee to obtain and manage funds
- on approval of plans, commence implementation and testing phases.

At the end of a successful implementation, the project would be handed over responsibility to an appropriate Internet governance structure which will most likely have emerged from WSIS and beyond.

The entire process is unlikely to be completed in less than three years. There are compelling reasons to suggest it must be completed within seven.

It also appears necessary at this preliminary stage to look at some sort of gateway structure through which people pass to adopt Internet Mark2. For a substantial period of time, Internet Mark2 will need to co-exist and co-operate with the legacy Internet. However eventually people will pass through the gate to the other side. It may also be necessary to nominate a point in time at which the gates shut, and interoperability with the old system disappears.

This is of course a substantial undertaking. It demands the best minds of our time, both in governance and in implementation.

6 Change management

6.1 Factors in change

The following are essential to realise in addressing this problem.

6.1.1 A compelling reason to change has to exist

No major change to core Internet protocols is going to be undertaken unless a compelling reason exists. It is not well understood yet, but the case is mounting.

6.1.2 Communication

The reason for change doesn't only have to exist, it has to be known. A major communication project is needed for a change of this magnitude. Trying to do it without formal change management planning would not succeed.

6.1.3 Involvement of major players

In particular, large players such as Microsoft and Cisco will need to be approached sooner rather than later, and planning for total adoption will have to take into account their adoption schedules. Once again compelling reasons to change have to be communicated: the cost to the market leaders to make the change across all product ranges is substantial, and the economic reasons to stay with current configurations may be compelling.

We may be talking about changes to every connected device on the planet here. This could exceed 1 billion devices.

6.1.4 Migration considerations

Co-existence with old systems has to be allowed, at least for a limited period of time.

There cannot be a one day changeover. For a quite substantial period of time the new system must co-exist with the old.

However this cannot be forever. There must be a date where the new Internet is able to move forward unencumbered by legacy systems. That is essential to forward development.

6.2 Stakeholder Analysis

The key stakeholders have been identified as:

6.2.1 Industry players

This group includes:

- large Monopoly players such as Microsoft and Cisco
- large ISP companies such as AOL, Earthlink
- the Open Source movement
- large online business such as Google, Yahoo, Amazon and Ebay,
- hardware providers and chip companies
- Registry/Registrar interests, eg Verisign, GoDaddy
- Telcos and Cablecos
- On line small businesses.

So we see a wide and varied industry affected by the level of changes that may be necessary.

Knowledgeable insiders in many of these companies and research institutions are aware to varying degrees of the need for improvement in the base protocols. However, executive decision makers at this stage are probably not aware of the need for reform. This publication may assist in raising awareness.

Some large industry players may initially perceive benefit in blocking or delaying change in order to protect dominant market positions, or to gain opportunities from proprietary approaches. There will doubtless be robust debates within some industry players about the market affect of change in this area, which could be accompanied by inaction in assigning staff to work on the issues.

However, market leaders will face crucial losses of markets if change does not occur within a reasonable time frame. Involvement in the new generation Internet will eventually be seen as a market plus, but the issues must be understood and clearly communicated for this to happen.

6.2.2 Current Internet bodies such as ICANN, IETF, etc

Some feeling of loss of control and that they should be in charge of change should be expected from these bodies. However, this study believes that would be inappropriate because:

- They have substantial operational issues to deal with which stretch their resources
- Their technical-only governance is inappropriate for the task
- Volunteer structures are inappropriate
- Their track record in handling major change in a timely fashion is not good
- They lack the range of disciplines and the skills base for handling change of this magnitude.

The incumbents need to be engaged as stakeholders, even significant stakeholders in terms of direction. However they should not be operationally in charge of this future development initiative, nor should they be the only voice at the table.

It is to be expected that some people in these areas – including prominent figures - will see the project as impossible, and raise a raft of reasons not to proceed. The debate will have to proceed however on the merits of the proposal, the consequences of doing nothing, and the likelihood of current actions and approaches being able to solve the problems in a timely and appropriate manner.

The more positive involvement of individuals who are initially unsupportive is more likely in a middle adoption phase, when the project clearly has impetus.

6.2.3 Government officials interested in Internet evolution

Government officials from a wide variety of backgrounds will have interest in new governance structures and in public policy issues which may emanate from projected changes. However they will be less interested in the protocol issues and content to leave that to others to develop.

6.2.4 The academic and research sector

This section of the population has had an historical role in Internet developments, and can tend to be at the leading edge in next generation usage. Individuals within this area will make major contributions during early stages.

6.2.5 Corporate users of the Internet

Corporate users need to be informed of changes and what they mean from the outset. Corporations may, for instance, need to schedule any changes to TCP/IP or DNS over a time frame of up to 5 years. The only similar event in corporate computing history would by the Y2k preparations.

The motivation for corporate users will be the consequences of doing nothing as the world changes, and the benefits of a more secure environment, less expenditure on firewalls and spam prevention etc. This group will mainly be involved in middle adoption phases.

6.2.6 Small business users on the internet

Small business users will be middle and late phase adopters.

6.2.7 Non-commercial and individual users of the Internet

These users with the exception of computer hobbyists will have their needs taken care of by those who provide them with Internet services, and will be middle and late phase adopters.

6.2.8 The 90% of the world's population who currently do not have Internet access.

One of the least affected groups, and on the whole late phase adopters (and ultimate beneficiaries). One of the advantages for this group is to come on board after change has been affected rather than have to go through it.

6.3 Overall summation

The analysis suggests that a transformative style is needed. It will need to be both charismatic and directive, depending on stakeholder groups involved.

In this style of change, major change must be affected within a limited period of time with multiple stakeholders. It follows that it will be essential to take on an early adopter group and move ahead, while being careful to leave room for others to come on board.

Early adopters will come from all stakeholder groups. However the most motivated early movers are likely to be industry players and research organisations.

Information is key to success here. Unless the problem is recognised and acknowledged, action to rectify it will not occur.

7 Technical Feasibility

7.1 Technical Assessment

Reform of the base protocols has been avoided in Internet technical circles because of the substantial change management ramifications. However, not to reform the base protocols has now become equally problematic, as the current protocols will not scale to the future or perform adequately. It's only a matter of time before the Internet collapses under its problems unless this is done.

Indeed any new protocols, however well introduced, will suffer the same problem within probably a lesser time frame than the 20 years during which the first Internet has served us well. Therefore it seems important to envisage structures whereby releases can progressively be introduced. That is more complex than it might seem, and will involve some tradeoffs. The alternative however is an Internet which becomes useless.

A project structure and dedicated staff is necessary. Major technical stakeholders may need to fund the developments.

The path to change however can be easier than imagined. A few potential directions will ease this problem.

7.2 Alternative Evaluations

The methodology suggested at this stage would see protocol reform, with a gateway structure through which all Internet users would eventually pass to the new and improved system.

Two alternatives may be available. One which has been suggested would essentially involve an overlay super-structure to create a next generation Internet. This would basically "pave over" the existing Internet.

The other alternative suggests a peer-to-peer style structure whereby problem protocols are eventually bypassed altogether.

Technical analysis will consider these directions as well.

8 Risk Assessment

The project contains risks that will require monitoring to ensure that they are minimised and do not impact the success of the project.

8.1 Business Risks

The Business Risks involved in inaction are identified below:

Risk	Likelihood	Severity	Notes on mitigation
Issue not understood	**Medium**	**High**	**Education of stakeholders must be a key component of project so that necessary actions and consequences of inaction are understood**
Issue not communicated	**High**	**High**	**Communication is necessary at all phases and particularly during implementation and migration**
Reaction too slow	**Medium-High**	**High**	**Project management techniques and dedicated staff are necessary to manage resource input and ensure timely response**
No co-operation to address problem	**Medium**	**High**	**Unlikely at a technical level. Best mitigated against by thorough communication and education campaigns.**
Resistance from incumbent governance bodies	**Medium**	**High**	**Existing bodies and many individuals within them are to be honoured and appreciated for what they have done in addressing complex issues and for their detailed knowledge as an input to future developments.**
Large players see opportunity for proprietary monopoly if collapse of current	**Low**	**Medium**	**Those few bodies with sufficient market strength to create proprietary monopolies**

system is allowed			**should be engaged and worked with. The issue should be understood from the outset and mitigated against by positive developments.**

9 Next steps

The risks associated with inaction on these issues substantially exceed those involved in moving forward. The risks associated with delay are also substantial.

Significant problems with base protocols are already evident. Given that the necessary changes cannot be implemented overnight, it is appropriate that a major effort to address these problems begin now.

The Internet Analysis Report – 2004 concludes that there are substantial inter-dependencies between protocol reform and governance. The dilemma facing the Internet community is that protocol reform has dependencies on governance reform, but cannot await completion of reform processes.

Protocol reform and governance reform must therefore be addressed in parallel, with some clear understandings among all parties concerned of the need to co-operate fully in ensuring that the next generation protocols and the next generation governance eventually come together in an appropriate manner.

Seed funding is now necessary to begin to move the project forward into planning phases. These phases will need to produce, at a minimum.

- Project Plan
- Business Case
- Implementation Strategy
- Change Management and Communication Plan

All readers of this report are invited to submit comments and suggestions by email to feedback@internetmark2.org.

Details of future phases and activities will be developed with stakeholders. A range of funding, sponsorship and involvement possibilities are outlined at www.internetmark2.org. Feedback and comments can also be made at this site.