

# A Project for the Synthesis of Composite TCP/IP Networks During Emergencies

Selwyn Russell and Peter Croll  
School of Software Engineering and Data Communications,  
Faculty of IT, QUT, Brisbane, Australia  
and  
National ICT Australia Ltd. Queensland Laboratory, Brisbane \*

## Abstract

This paper describes a project being undertaken by National ICT Australia with a Safeguarding Australia theme.

History shows that the pattern of use of communications during a large scale emergency to be quite different from those of normal times. The emergency recovery team needs a communication network with capabilities and configurations which have a very different profile from everyday communications networks, to avoid network overload from enquiry bursts and to favour emergency management traffic. Ideally, this temporary network would involve a unification of all commercial networks which at other times are competitive. In transforming the network in a region, crucial decisions are which components of the network to modify and how. With the move towards TCP/IP based networks, the Internet offers a widely used and deployed platform, and the configuration of routing tables becomes the main focus in the quest for rapid modification of network properties to synthesize a composite network. Current router software is not easily modified simultaneously across many routers throughout a region in a way which suits emergency management.

This project seeks to provide a user-friendly transmission policy language for rapid definition of the required TCP/IP network behaviour, and the means to convert the policy to a list of target routers and their corresponding reconfigurations for automatic processing.

## 1 Introduction

This paper is concerned with TCP/IP communications networks as part of critical infrastructure protection and their use during times of large scale emergencies. Large scale emergencies can result from catastrophic natural events or from the actions or inactions of people. During these times, many people are badly affected and many others are attempting to rescue and assist. Good communications are essential to assist victims and related persons such as family, and enable emergency workers to go about their tasks in the most efficient and speedy co-ordinated manner.

During large scale emergencies, communications traffic between people takes on a very different profile from everyday communications, as discussed in section 2. There is a greater need for fault tolerance, for capacity to handle bursts of traffic, and prioritisation of communications within a class of users, viz the emergency personnel who themselves may be part of many co-operating groups.

There are many emergency services such as fire fighting services in cities and forest fire fighting units in rural areas in existence. Most of these have their own emergency communications, typically two way radio with communication between a base station and the mobile equipment, intended for operation over a relatively small area. The base station can communicate with other people and organisations via a commercial telephone utility, typically with a Quality of Service contract. The QoS might be meaningless during a large scale emergency because of damage to the utility's infrastructure. This type of emergency communications has worked well for the intended type of emergency, but may be ineffective

---

\*National ICT Australia is funded by the Australian Government's Department of Communications, Information Technology, and the Arts and the Australian Research Council through Backing Australia's Ability and the ICT Research Centre of Excellence program.

during a large scale event, as the recent New Orleans incident indicates. A base station is a single point of failure and a bottleneck for communications from one mobile unit to another. Using their two way radio networks, the team members in the field have to communicate with and through the base station and cannot communicate directly with one another. Furthermore, each of the different emergency groups in the same region, such as fire and police, will have its own dedicated base station and mobile units. Each group has its own allocated frequency band(s) and will have difficulty talking to anyone in another group using their two way radio. In such situations, another group is frequently added for coordination of the groups. This coordinated emergency services control team is critical for communications between the groups, and is another single point of failure and choke point for performance. To communicate without going through the central control, members of the rescue groups have to use the public carriers, which may have a poor quality of service at those times.

Experience at Kobe as outlined in section 2 indicated that the Internet is more resilient than other networks during such times [1]. The Internet has three important advantages. Firstly, the Internet was designed to be a connection between differing types of (digital data) networks [2], and today most other data networks have some way of interfacing to it, thus enabling at least some degree of interoperability and internetworking. Secondly, the Internet was designed for fault tolerance and fault recovery, to meet DARPA requirements [3]. Many Internet nodes today have provision for an alternate route to an important destination for use if a high priority link fails. With the growing use of emerging wireless technologies such as wireless broadband and free space optics, viable alternate links other than dial-up ISDN or satellite are practical. Another advantage for emergency communications is the network design has no central controller or base station and thus communications are more flexible, viz simultaneous one to one sessions, one to many, and broadcasts. Hence the interest in the Internet as a network useful during large scale emergencies.

The commercialisation of critical communications infrastructure previously owned by governments leads to a lower excess capacity, as shown in section 4. With the move towards TCP/IP based networks, discussed in 4.2, each carrier is approaching the long sought goal of a single integrated network. Even though these competing networks normally operate independently, for emergency response, our goal is to have them further integrate and give the appearance of a single network throughout the affected region. These competing networks normally have common gateway points where traffic from one is passed to the other. To function as a single composite network, the routing tables of at least the gateway routers need to be modified so the networks co-operate and share the loads, and give priority to traffic to/from emergency services devices. Current router software is based on the assumption of a stable network with only small and incremental changes. During an emergency, speed and accuracy are vital, but at those times there is more chance of human error. To simplify communications management throughout a region, it is desirable to be able to describe the ideal network behaviour at that time in a high level policy style specification, and to have the relevant network components automatically reconfigured in accordance with the regional specification.

To accomplish this goal, management must be able to simply specify a complete and unambiguous high level policy for the communications network, and accomplish this act in times of stress and time pressure with a minimum of delay. An RFC on policies for routing first appeared in 1989 [4] and languages for policy specification were used in RIPE-81 in 1993 [5]. Policies for routing are reviewed in section 6.1 but we conclude they are not the type of policies needed during an emergency. Prior work on policy languages is outlined in section 6.2, and we conclude the proposals are too complex for use by emergency workers during disaster management. We conclude that the outcomes of prior work on policies are not suitable for use in emergencies. This project is investigating means of rapidly specifying policy in ways meaningful to emergency management, as outlined in section 7.

To reconfigure the network to meet the requirements of the policy, an initial analysis of the existing network structure and component capabilities would be used to generate a desired topology and the necessary transmission properties of switching nodes. A subsequent process would convert the synthesised requirements of individual nodes to a node-specific form which would be used to automatically reconfigure the device. These processes and algorithms are being investigated in this project.

There are many difficulties associated with achieving the overall goals and these are discussed in the paper.

## 1.1 Contributions and Structure of the Paper

In the next section we discuss some recent examples of communications during disasters and important lessons. The difficulties of rapid configuration of current IP networks, i.e. rapid simultaneous reconfiguration across a region of many routers, are considered with examples in section 3. The communications environment is continually changing and needs to be considered when planning networks. In section 4, we consider effects of the commercialisation of public wide area networks by governments over recent decades and the technological trends which are shaping communications. We outline the requirements for a rapid reconfiguration system with a simple policy based interface in section 5. Some previous research has been conducted in areas related to this project, and these efforts are outlined in section 6. Our project is outlined in section 7.

## 2 Recent Communications Experiences and Lessons

In this section we briefly examine some recent disasters with large scale damage to critical infrastructure. The first is the Kobe earthquake of January 1995, which has been used for preparation for future events by groups in Japan. The second is the World Trade Center in 2001. Hurricane Katrina and the Gulf Coast of the USA are also mentioned.

The Kobe earthquake of 17 January 1995 was an example of a natural disaster involving multiple administrations in a single country, a single city, with a large impacted area. Essentials such as gas, electricity and water were cut off for days. The communications network was overloaded, particularly by friends and relatives seeking to talk to people in Kobe or to learn of their fate. The mass media failed to adequately assist people locate the necessities of life: food, water, shelter. A number of very localised newssheets with relevant local information were set up and these were of great assistance.

The Kobe earthquake experience motivated preparation and planning for a more robust communications network [1] [6]. Because it was observed that Internet communications were more resilient during the disaster than others, Internet communications were chosen as the foundation, using satellite links between important nodes.

The Kobe network planning in 1995 had the implicit assumption that damage is caused by a natural disaster. When malicious saboteurs are also at large, the situation is more difficult because humans being inherently more unpredictable and innovative than nature.

Kobe struck without warning, so there was no possibility of heightened preparation in the days leading up to the event. Another unexpected event, but deliberately caused by people not nature, happened in New York, 11 September 2001. Kobe happened in a few minutes, whereas the New York event unfolded and worsened over a few hours, during which fast action was needed to prevent further loss of life. In both events, there were abnormally high demands on voice communications to and from those in the effected areas. In Kobe, there was widespread damage to the communications network, particularly wired ground links, but in New York the damage was confined to a comparatively small area.

In New York the routine operations of many major Fortune 500 corporations were adversely affected or threatened. Many had alternative computer installations set up after earlier attacks on the World Trade Center. In a disaster such as this, it is essential that commercially critical communications from outside areas to affected Internet sites be rapidly rerouted to the backup sites, i.e. neighbouring routers need to be reconfigured to send traffic to the correct remote fall back locations, rather than sending a “host unreachable” or “network unreachable” error.

Hurricane Katrina caused widespread damage in 2005 to many people and infrastructure in the Gulf Coast region. The category 5 winds caused intense damage to above ground structures, including aerial cabling and communications towers. The city of New Orleans gained the most attention in the media but it was not a critical node in the Internet and the Internet itself continued reasonably well [7]. The Internet infrastructure was typically based on below ground optical fibre and survived the hurricane quite well, however many ISPs suffered extended outages as a result of a loss of electricity.

### 2.1 Relevant Lessons

In the above disasters, and in disasters in general, wide damage to the communications infrastructure can be expected, thus making it difficult to build wide area emergency response networks in readiness.

In the Kobe rescue it was noticed that Internet facilities survived the earthquake better than other networks, and as a result planning began on an emergency network based on the Internet. The main purpose of this temporary network is to meet the demands of enquirers about the welfare of possible victims. There are two main components: a distributed data base with interfaces for the Internet, telephone, fax, and cellular phone; and a data network using satellite links between the ground nodes [1].

Satellite links have been recommended in Kobe to work around the likely damage to cables. Satellites would have been of little use at the World Trade Center because of the relatively small area and the transponder capacity needed for the thousands of people concentrated there. During Katrina, satellites would have been useless for high capacity links because of the water absorption of the signals and the likely destruction of outdoor dish antennae and towers from the intense wind gusts. Satellite ground stations near the sea were at risk from tidal surges and flooding. However, individual satellite telephones probably would have functioned indoors during the storm, depending on the sheltering enclosure, and would have been useful in the subsequent days. Thus it is difficult to select a single communications infrastructure as the solution for all types of catastrophes.

In 1995, communications in Japan were dominated by the Government carrier, but competitive carriers have been introduced. In the USA, as in most regions serviced by a modern communications network, there are probably multiple separate commercial networks in use, each effectively duplicating what its competitors do. Potentially, the duplication introduces some fault tolerance to the critical communications infrastructure, but during an emergency it would be much easier for rescuers if there were only one network to bring up, repair, and operate.

At Kobe, yearly drills using the Internet have been run around the anniversary of the earthquake. Reports on the first and the fifth have been published [1] [6]. When the first was conducted, many failures were experienced because of relatively minor issues, such as batteries in a flashlight not working because of age. The Japanese involved are adamant that a communications system that is not in constant use or regularly reviewed will fail during an emergency. The nature of the Twin Towers disaster makes communications drills for people at all high rise buildings throughout a nation impractical. However, drills by emergency services could be held regularly for simulated disasters.

Severe overloading of the communications network, mainly by the public, implies that a separate network is desirable for the emergency organisations, or at least giving them priority for communications.

### **3 Difficulties with Current Reconfiguration Processes**

Although the composite network recommended in the previous section can in theory largely be achieved by reconfiguring the routers in the network, this task would be difficult and error prone during an emergency, and with current means it may be better to not even attempt it. In this section we show some of the difficulties involved in the rapid simultaneous reconfiguration of many routers. We begin with an illustrative scenario.

#### **3.1 Network Reconfiguration Scenario**

There has been an incident in the Central Business District resulting in the evacuation of a dozen high rise buildings. The telephone network is badly overloaded. The cellular wireless base stations cannot cope. The Internet in the CBD is becoming congested with email and attempts at VOIP. Emergency services communications which are based on the Internet are being impeded. A quick analysis of the situation indicates that if non-essential Internet traffic to and from the area were removed, the emergency personnel would cope well.

The recommendation is that only traffic directed to emergency services devices be allowed into the area, traffic from emergency services devices have priority within the area and for transmission from the area.

Implementation of this recommendation would involve knowledge of which networks are in the area, which routers are involved, and what to have in each of the routing tables.

## 3.2 Reconfiguring Links

Communications networks are designed to be fully connected, i.e. appear to a subscriber as having a connection from that subscriber to every other subscriber. Direct inter-connections are generally impractical and a network is built as having many switching centres or nodes, some of which are joined by direct links. Subscribers connect directly to edge nodes, and most nodes relay communications from one node to another, so that a message from one subscriber to another will pass through a sequence of nodes. Where a node should relay an incoming message is controlled by internal tables in the node. A node consists of multiple devices which are involved in the transmission process.

For efficiency, most links between nodes are fixed and not easily altered. Even microwave and free space optics links require visits by technicians, directional realignment (which might not be possible depending on the physical location, structures and mounting), equalisation and testing. Hence reconfiguration of the network between two subscribers is more easily accomplished by altering the routing tables at the nodes. We will now investigate whether this can happen with current practices.

## 3.3 Reconfiguring Routing Tables

In this section we point out the potential difficulties with rapid dynamic reconfiguration of a TCP/IP device by a network administrator, using commands for the Cisco router operating system IOS as illustrations. Some popular versions of open source firewall software have been designed to appear the same as the Cisco IOS commands, so the examples in this section apply to much of the existing worldwide IP network.

### 3.3.1 Access Control Lists

The IOS allows access control with a number of types of access control lists (ACLs) which can be assigned to an “interface” of a router (a physical connection through which bits enter and / or leave the router). An ACL is created as a group of single line statements. An ACL is processed from the first statement onwards until a match occurs, at which time action is taken for the packet in question, and no further statements of the ACL are read. If all ACL statements have been processed and no match has been found anywhere, the packet is dropped.

A particular ACL can be applied to one or more interfaces, and will be specified as applying to traffic coming into the router through that interface or departing from it.

The first example uses the earliest form of ACL, the “standard” ACL [8].

```
Router27(config)#access-list 2 deny 173.16.2.7
Router27(config)#access-list 2 permit 173.16.2.0 0.0.0.255
Router27(config)#access-list 2 deny 173.16.0.0 0.0.255.255
Router27(config)#access-list 2 permit 173.0.0.0 0.255.255.255
Router27(config)#access-list 3 permit 177.200.0.0 0.0.255.255
Router27(config)#interface a0
Router27(config-ip)# ip access-group 2 in
Router27(config-ip)# ip access-group 3 out
```

The prompt `Router27(config)#` indicates that the administrator is working with router named “Router27” and is in (privileged) configuration mode. The `access-list` command appends a statement to an existing ACL or creates a new ACL if not existing. The first statement denies passage to packets sourced from host 173.16.2.7 and the second statement allows packets from any host on the same network. Traffic from that host will match in the first statement and will not get to the second, whereas traffic on that network from other hosts will not match on the first and will be permitted by the second. Similarly, the third statement will block traffic from a particular subnet and the fourth will allow all traffic from all subnets of 173.0.0.0.

The `access-list` command syntax is

```
access-list <access list group number> <permit|deny> <source IP address> <optional IP mask>
```

Point to notice here are:

- Only the source address is considered. For many situations this form of access list is adequate and permits simple implementation.
- Access lists are numbered. All lines with the same number are considered to be part of the same ACL. New statements are appended to the ACL identified by the number.
- The ACL numbers are drawn from reserved groups which imply the associated protocol. Numbers 1 to 99 imply a standard ACL for TCP; numbers 600 to 699 imply the AppleTalk protocol. Thus a router can have only 99 ACLs defined in it.
- An ACL need not be in use, so ACLs could be entered and available for emergency situations.
- For correct operation, more specific matches must precede more general specifications.
- It is permissible only to append statements to an ACL. Statements cannot be re-ordered, altered or individually deleted. However an entire ACL can be deleted and recreated. This feature simplifies the router operating system and is tolerable because ACLs normally are stable for extended periods of time. However, the most likely change is to a more specific statement, e.g. deny traffic from a particular host or network which has been identified a source of spam, and these need to be at the start of an ACL. The current practice is to delete the ACL, go to `config` mode and rebuild the ACL. The burden is relieved by maintaining the ACL statements in a separate text file and pasting the updated file from an editor to the IOS command line.
- If the administrator makes an error in entering the statements, the statement is still accepted and later used, but may give strange network behaviour which is difficult to diagnose and repair.

The next example uses the more useful “extended” ACLs which allow inclusion of destination address, destination port, and transport type (`tcp` or `udp`). Their numbers must be in the ranges 100 – 199, 2000 – 2699.

```
Router27(config)#access-list 102 deny tcp host 173.16.2.7 any eq telnet
Router27(config)#access-list 102 permit tcp 173.16.2.0 0.0.0.255 any
Router27(config)#access-list 102 deny udp 173.16.0.0 0.0.255.255 any
Router27(config)#access-list 102 permit tcp 173.0.0.0 0.255.255.255
Router27(config)#access-list 143 permit tcp any 177.200.0.0 0.0.255.255 lt 1024
Router27(config)#interface e0
Router27(config-ip)# ip access-group 102 in
Router27(config-ip)# ip access-group 143 out
```

Note the following points

- The same limitations on editing apply to an extended ACL as apply to a standard ACL.
- An extended statement is needed if the destination address must appear.
- A specific protocol or port can be specified. If several are involved, then several statements would be used.
- The extended statement is more flexible but is more complex, with more opportunities for the administrator to make an error when pressed for time.

To replace an ACL, the current ACL is totally deleted, e.g. with a `no ip access-group 102` command, and then the new version is installed, as with `ip access-group` in the above example. During the time when the ACL is nonexistent, there are two possible modes of operation of the router: it blocks all traffic on the interface where the ACL had been installed ( a “fail safe” precaution), or it moves to the default mode of permitting all traffic. Probably neither of these is satisfactory during an emergency. In either case there will be at least a temporary instability in the routing behaviour until neighbouring routers in the region adjust.

### 3.3.2 Reconfiguring a Region

The tasks described above are for the case of a single router. During an emergency involving a region of many routers, it is important that they be reconfigured within a very small time window to avoid network instabilities which would otherwise result from “host unreachable” or “network unreachable” messages and from any conflicting automatic routing updates between routers. When there are many routers to be reconfigured, the task becomes correspondingly riskier.

Obviously the minimal impact will happen if the reconfigurations can be carried out by synchronised computer processes.

### 3.3.3 Reconfiguring Multi-Networks

Converting competing networks into a composite harmonious single network is another step up in complexity and difficulty. The exterior gateway routers interconnecting networks need to be reconfigured simultaneously along with the interior routers.

## 3.4 Summary

In this section we have demonstrated that reconfiguring routers to synthesise a composite critical communications infrastructure is itself difficult, assuming the new configurations are known, and too risky to be handled by a person under stress during a time of emergency. If it is to be done, it should be carried out by an automated system.

In section 5 we consider some of the requirements of such an automated system, but first we reflect on trends in public communications utilities which might influence our efforts during a disaster.

## 4 Communications Environment and Trends

In planning for critical communications infrastructure, the current environment needs to be taken into account, along with important trends which will affect the infrastructure. Because of the huge cost of extensive communications infrastructure, it is infeasible for governments to install an extensive dedicated network which is used only during emergencies. Rather they have chosen for commercial interests to provide most of the infrastructure and have it in daily use. In this section we consider the effects of this strategy on communications infrastructure for use during emergencies.

### 4.1 Effects of Commercialising Wide Area Networks

Wide area networks have historically been provided by government utilities, as in Europe, or by a government supported monopoly e.g. AT&T in the USA. These utilities provided good quality performance for the services they provided, but were frequently criticised as having high costs and charges higher than necessary.

In 1984, the USA government revoked the AT&T monopoly to introduce competition into the industry. Over the ensuing decades, other governments around the world have also sought to reduce the costs of business inputs in their countries by enabling competitive supply of telecommunications services. In most cases, the former monopoly has been corporatised and sold off to private shareholders, and has to fight for market share against competitors now licensed by its former owner.

Further competition in the supply of data services has appeared from some competitors which have their own large networks, such as electricity distributors and railway companies.

Thus the situation is that communications services have changed from a performance centred non-competitive single supplier slowly changing environment to intense competition where the CEO's future depends on current profits and share price.

#### 4.1.1 Effects on Spare Capacity

Heavy pressure on executives is to maximize Return On Investment to support the share price, to keep costs down and minimize investments. Whereas in the past, excess capacity was often provided because the government budget process required three to five years lead time for capital expenditure, excess

capacity is now kept small to minimize outlays for equipment which is suffering shorter operational lifetimes due to continual technological advancements. In today's environment, a three year budget lead time would allow a nimble competitor to snatch the market with a new product line.

There is pressure to minimize maintenance and staff to whatever is the lowest in the industry, the "world best practice" metric.

For disaster management, the impact of current commercial competition is to provide very little burst traffic capability, which is diametrically opposite to the experience at Kobe described in section 2.

#### **4.1.2 Lower Excess Capacity in Rural Areas**

Because of the higher density of potential customers, new entrants in communications infrastructure prefer to start in large cities and spread in time to smaller cities. The other negative factor against rural areas is the higher cost of deploying infrastructure due to higher transport costs, longer link lengths, less hospitable climate, etc. Services in rural areas have fewer suppliers, fewer alternatives, and have higher costs.

However, natural disasters such as floods and fires are more prevalent in rural areas than in cities. Disaster workers in these areas are dependent on communications for co-ordination and safety, and usually would have their facilities of their own, which connect into a commercial service.

#### **4.1.3 Lack of Single Managed Network**

Unlike the pre-privatisation network which had one owner and one manager, the current communications network consists of portions deployed by independent competing entities which control their own network infrastructure. Each new communications competitor has at least some network infrastructure which is for itself alone, ranging from extensive complete fibre and cable installations and cellular base stations to only wireless broadband access points connected by leased links.

The lack of a single manager for the entire network hinders the unified use of the resource at times of emergency, when national interest becomes dominant over private goals.

### **4.2 Technological Trends**

Very early telephone networks were designed specifically for the type of traffic supplied by the customers at the time, viz. analogue voice. Over the decades, the analogue voice networks have been adapted to meet further development, such as analogue video for commercial television, but some more specific separate internal networks have been deployed in recent times, mainly to meet digital data demands from subscribers. The more established utilities are facing the management problems of maintaining multiple networks, including technically obsolescent, but functionally satisfactory, networks still within their designed economic lives. Competitive cost pressures are persuading large and established utilities to reduce their plethora of systems, and the concept of a single network providing integrated services, while not new, is being seriously considered by telcos. The popular means of implementing an integrated services network is TCP/IP, which of course is already an important part of many telco networks. With the growing popularity of "Voice over IP" (VOIP) Internet telephone products, the demand from subscribers for TCP/IP based services has no end in sight. (Although early observers expected non-telco VOIP products to detrimentally affect telcos, telcos have responded by offering their own versions and taking advantage of priorities in TCP/IP to negatively impact the performance of VOIP supplied by competitors [9]). Hence we can expect utilities' core networks to have an increasing major TCP/IP component.

The other trend is to wireless broadband and the services which can be provided thereon, thus enabling new industry entrants to service many customers without suffering the enormous costs of cabling to subscribers' premises. These wireless broadband networks consist of nodes and links similar to mobile telephone networks, with connections to the established communications networks. For simplicity of interfacing between the networks, a TCP/IP digital core makes sense for the wireless broadband entrants.

### **4.3 Summary**

Critical communications infrastructure for wide area networks is mostly based on infrastructure designed for daily commercial use by the public and operating in a highly competitive environment which encour-

ages cost cutting and discourages the installation of the excess capacity needed in times of emergency.

In well populated areas, the multiplicity of independently operated networks frustrates the central control of communications in the region during an emergency but provides the makings of potentially fault tolerant and adaptive network.

The move towards the use of TCP/IP which is itself potentially fault tolerant facilitates interoperability and resilience. The Internet offers an attractive basis for a critical communications infrastructure. In the present form there are difficulties in converting a collection of separate competing networks to appear as a single composite network, as outlined in section 3.

We discuss requirements of the reconfiguration process of routers in commercial TCP/IP networks in the next section, which must be brief because of space restrictions.

## 5 Brief Requirements of the Reconfiguration Process

We use the term “reconfiguration process” as a broad phrase to encompass all of the smaller processes which are involved in transforming the (TCP/IP) networks in the affected region into a composite network of co-operating units.

When there is a physical emergency, the authorities in the form of police or military or emergency services evacuate the healthy people from the area and then conduct a search and rescue operation in the area. The evacuation and take over of the area is partly so that the emergency workers can conduct their activities more effectively with fewer distractions. Similarly, during an emergency, Internet based communications could be altered so that important communications such as emergency services gain priority, and most other communications be diverted or set at a low priority.

The composite network is the outcome of the overall reconfiguration process, so we will briefly look at some requirements of the composite network. For use in times of emergency, robustness and availability are essential. At the minimum, the network needs to be able to convey messages which are for decision making and for control. A lower priority is acceptable for less urgent information such as enquiries about the welfare of friends. Hence, some form of prioritisation of traffic is needed. To reduce internal traffic in the region, redirection of certain traffic is required such as diversion to alternate sites of requests to business servers.

The specification of policies should be in a simple and clear format, and appear as plain English wording. Simple keywords are probably needed for fast and accurate processing, but they should be limited in number and easy to learn for those making the policies and easy to understand for those reading them.

During an emergency, damage may be continuing while rescues are proceeding, so the system should be adaptive.

Even though the intention is to have automatic processing, managers should be able to over-ride generated configurations.

Because lengthy configuration files for real routers are not easily understood by non-technical persons, an option is to create a generic language for specification of router configuration and have the initial reconfiguration for each router created using the generic language. The generic reconfiguration file for each router could be reviewed if necessary by a person who was familiar with the generic language but not necessarily expert on the particular router.

To control access and to allocate priorities, the final network needs to authenticate the source devices, and often the destination devices. It must be easy for the emergency handlers to specify devices and their privileges or capabilities.

It is probable that some of the DNS entries will also need to be altered. Policies should enable this to happen.

## 6 Previous Work

There has been quite of deal of effort expended in fields relating to these requirements. In this section we will consider work on policies and see that it is only obliquely related to our requirements. Policy specification languages have been researched, for a different purpose, and we will see they are too complicated for our requirements. In section 6.4 we will briefly look at outcomes of work on disaster communications.

## 6.1 Policies for Routing

An early work, “Policy Routing in Internet Protocols” [4], was published in 1989, when the Exterior Gateway Protocol was in use. This RFC was the first to address a problem which emerged as the Internet grew larger. The Internet connects autonomous networks or domains, now commonly referred to as “autonomous systems” (AS). Most of these are end domains, or “stub” domains, which generate their own traffic to and from other domains, but because of geographic dispersion, some domains act as relays or transit domains, accepting traffic from some domain and passing it through to another without using it inside the domain. As the Internet grew, with much of the traffic moving through transit domains, some of those who operated the transit domains developed concerns about what was moving through their domains, and users of the transit domains became concerned about who was handling their traffic. Specifications of constraints became known as “policies”. RFC 1102 proposed an approach whereby domain managers specified a policy restricting the type of traffic acceptable in that domain. The three metrics mentioned were class of customer, quality of service a domain could provide, and the cost recovery. To go from one place to another, a series of domain with acceptable policies would be determined, called a “Policy Route” in RFC 1102.

These concepts were further developed such as in RFC 1126 [10] “Goals and Functional Requirements for Inter-Autonomous System Routing”, in 1989. Around this time, RIPE [11] was formed in Europe. Being composed of numerous countries, one of the concerns of Europe was routing through different ASes. A RIPE online database was developed to hold the policy information supplied by each AS, and a description language [12] was developed.

In the USA, the Inter-Domain Policy Routing working group (IDPR) of the IETF completed its work in 1993 with a series of RFCs “IPDR as a Proposed Standard” [13], “An Architecture for Inter-Domain Policy Routing” [14], and “Inter-Domain Policy Routing Protocol Specification: Version 1” [15]. In this architecture, each AS specifies 1) a transit policy, indicating access restrictions, some measure of quality of service, and charging, and 2) a source policy for traffic from that AS, indicating preferred or unacceptable transit domains, desired quality of service, and charging methods. There are special entities added to the Internet to manage IDPR type functions [13]: policy gateways, path agents, route servers, mapping servers, and configuration servers.

More work followed, such as “Routing Policy System Security” [16] in 1999, “The COPS (Common Open Policy Service) Protocol” [17] in 2000, “Routing Policy System Replication” [18] in 2000, “Structure of Policy Provisioning Information (SPPI)” [19] in 2001, “Session Authorization Policy Element” [20] 2003, and an update for IP v6 in 2005 “Routing Policy Specification Language next generation (RPSLNg)” [21].

In spite of the work on policies for transmitting traffic, there appear to be few systems using them. The topic is not even mentioned in the Cisco CCNA course.

Note that these policies are intended to be used only for transmitting traffic and so are proposed for use only at the gateways of ASes. There are no policy based controls proposed for use within an AS.

Without going into further details, the reader can see that 1) this is not a simple system and requires a reasonably complex infrastructure with noticeably higher overheads, and 2) this is not the type of policy we proposed for use in an emergency in 3.1 and 5.

It seems fair to conclude that in spite of interest in using policies for some management of transit traffic in the Internet, it has reached a stage where further development is in the “too hard” basket. We are intending to use a different type of policy specification in a different way, as outlined in 7, but we realize this too will prove quite difficult.

## 6.2 Policy Specification Language

The IDPR team did not specify a way of writing machine readable policies. An option was the language of RIPE-81 developed in 1993 and updated by RIPE-181 in 1995 and also published as RFC 1786 [5].

Later, an IETF working group developed the Routing Policy Specification Language (RPSL) [22] in 1998, similar to RIPE-181. This RFC was replaced in the following year by RFC 2622 [23], a 68 page document with details of “classes”, data structures containing policy and administrative information. There are three classes relating to contact information: *mntner* class, *person* class, and *role* class. Other important classes are the *route* class, the *aut – num* class, the *dictionary* class, the *inet – rtr* class, and a number of *set* classes.

### 6.3 Policy Routing on a Single Router

In this subsection we will outline the Cisco router feature of “Policy Routing”, which is the only mention of policies for routing in the CCNP course. Hence Cisco Network Professionals are not trained in the policy routing for ASes discussed in 6.1.

Cisco describe their version of policy routing as “often a complicated task”. Three steps are needed to implement a policy routing instruction. An access statement such as `access-list` is needed to link an action such as permit to an IP address. The Cisco IOS `route-map` command [24] begins a multi-line entry to link an address to a physical interface on the router.

A `route-map` command has the form

```
route-map <name> <permit|deny> <entry number>
```

The `name` is the name shared by all entries for the single route map, similar to the access list group number used to distinguish ACLs in section 3.3.1. Unlike the ACLs discussed earlier where a new entry was always appended to an existing list, the location in a list of a new `route-map` entry is determined by the optional `entry number`. The default values are 10, 20, 30, ... A later entry giving an entry number of 23 would be inserted between the existing 20 and 30. The `route-map` command then moves from `config` mode to `config-route-map` mode. The number at the end of the `match` command specifies the access list to be linked to the route map. When the router is processing the list identified by the second parameter of the `route-map` command, if a match is successful, the router will act on the command in the next line. The `match / set` behave like “if match ... then ...”.

In the following example adapted from the Cisco training documents [24], router named `Router27` is connected to network 173.16.2.0 through interface `e1` and traffic from there is to be sent to `ISPa` through the router’s serial interface `0/1`, while traffic from network 173.16.3.0 enters through interface `e2` and is to depart to `ISPb` through interface `0/0`.

```
Router27(config)#access-list 2 permit 173.16.2.0 0.0.0.255
Router27(config)#access-list 4 permit 173.16.3.0 0.0.0.255
Router27(config)#route-map ISPa permit 10
Router27(config-route-map)#match ip address 2
Router27(config-route-map)#set interface serial 0/1
Router27(config-route-map)#exit
Router27(config)#route-map ISPb permit 20
Router27(config-route-map)#match ip address 4
Router27(config-route-map)#set interface serial 0/0
Router27(config-route-map)#exit
Router27(config)#interface e1
Router27(config-if)#ip policy route-map ISPa
Router27(config-if)#exit
Router27(config)#interface e2
Router27(config-if)#ip policy route-map ISPb
Router27(config-if)#exit
Router27(config)#
```

This is a simple example, but will not do the type of controls we seek.

### 6.4 Disaster Communications

As discussed in 2, the Kobe earthquake in 1995 led to the development of the I Am Alive system and the WISH satellite based network. These innovations were not taken up elsewhere, but some moves have been taken in the USA to put in place suitable networks for emergencies., such as the US Government Emergency Telecommunications Service (GETS). (These apparently failed to meet their objectives during the 2005 hurricane season in the USA.)

More general Emergency Telecommunications Networks (ETS) have been considered by ANSI and the ITU, but not closely associated with TCP/IP networks. In 2004 the IETF published several RFCs on ETS. RFC 3689, “General Requirements for Emergency Telecommunication Service (ETS)” [25], defined a set of system requirements so that five existing ETS-related standards, one from ANSI and four from

the ITU, could be supported on the Internet. Solutions for the support are not attempted. The other RFC, “IP Telephony Requirements for Emergency Telecommunication Service (ETS)” [26] extends the set of system requirement to support IP telephony as an end to end application across networks which may comprise the Internet and other networks. Again, solutions were left to others at a later time.

## 6.5 Summary

Prior work in the policy field has not been on the topic of alteration of critical communications infrastructure to assist disaster management.

# 7 Outline of the Research

At this stage, the program is waiting for suitable researchers to join, so this section is an outline of the anticipated direction of the research. Artificial intelligence techniques are likely to be involved.

## 7.1 Policy Specification

Policy specification is the starting point for a user of the final system. One line of research is the acquisition from potential users of their preferred ways of specifying policies for operation during an emergency. There are many possible types of emergencies in very different parts of the world, e.g. summer forest fires in outback Australia versus winter alpine avalanches in Switzerland, so a flexible technique with a low chance of error is essential. A menu system may be useful but may be too slow for some experts. The input system should be able to detect ambiguities from what people tell it and ask for further differentiation as required. Techniques from Artificial Intelligence might be useful.

The outcome of this stage might be the (relatively) simple policy statement, which would contain at least one statement to the effect that priority be given to traffic originating from or going to emergency services devices. The raw specification might be usable directly or might have to be translated to a more machine friendly form for later processing.

## 7.2 Network Discovery

The state of the potential composite network needs to be known before decisions can be made as to how it ought to be configured to best suit the current emergency. If there are different separate networks operated by competing companies or government agencies, information on operational details might be difficult to obtain. A company might be reluctant or even refuse to divulge details of its network for fear of leaking to a competitor.

Even if all parties agree to the release of information, there is the problem of how to obtain a comprehensive and consistent view of possibly heterogeneous systems and their topologies, addresses of interfaces of routers and links, routing configurations of all routers and gateways, and other details which might prevent a composite network from operating optimally.

Furthermore, during an emergency, some damage of network components can be expected. In TCP/IP networks, components use routing protocols to inform other components of changes they observe, but there can be serious delays in this process, possibly leading to instability. Gathered information might be inaccurate, leading to poor reconfiguration decisions.

In the case of disruption to important commercial services, as in the World Trade Center, the addresses of alternative sites need to be obtained so external traffic attempting to enter the region can be diverted pending restoration of normal services.

In the case of I Am Alive type facilities located outside the affected area, the address of the portal needs to be known, so enquiries can be diverted to minimise traffic in the disaster zone.

The identities and addresses of all participating emergency services devices need to be an input here to ensure they receive priority.

The output of this stage is a machine processable form for input to the next process, generating a preferred network configuration.

### 7.3 Network Configuration

Assuming an overall picture of the composite network can be built, the next step is to determine if and how part or all of it can be used to assist in communications for emergency services. Depending on the situation, a minimal alteration or disruption to current non-emergency services might be a requirement, consistent with ensuring emergency services function to their peak.

Particular capabilities or relative deficiencies of a device need to be taken into account, for example to ensure that an older node with low processing capabilities is not tasked with computationally intensive authentication functions as a gateway.

This process will be trivial for a small region with only a single router, but is likely to be very difficult for a large region with dozens of competing networks and hundreds of routers. Some algorithms might be available or developed to enable a satisfactory temporary solution.

A topology of network links, nodes and their generic configurations will be generated at this stage. The conversion of generic configurations to configurations for specific routers could be done at the point of generation of the desired topology, or it could be deferred to the deployment stage. If done at the same point, there would be a delay while the series of individual translations run. If the generic configurations are dispersed to the appropriate routers, the translations could run in parallel. The option for distributed translation requires that all sites involved in the translation have the correct software. Here we might heed the experience and recommendation of the IAA team in Japan, that unless such software is tested regularly, it will fail when needed in a genuine emergency. The lower risk strategy is for translations to be run by the control centre where the composite network is synthesized. In some situations, distributed translation may be chosen, but not in other situations. Sometimes a combination may be best, particularly if some of the routers are older models with lower switching capacity and lower computational resources.

### 7.4 Deployment

In the deployment stage, for a particular router, if translation is distributed, a process at the local router will take the relevant generic configuration created at the time of synthesis of the composite network and will produce a set of custom configuration commands, taking into account any special requirements of the router and its immediate environment. If translation is at the synthesis location, then the required configuration will be transmitted to the local router.

In either case, there is a need to reconfigure the router with the new configuration data. Obviously the reconfiguration needs to happen in as quickly as can be arranged, with minimal problems in the composite network. It can be expected that various unforeseen difficulties will arise in actual networks when this operation is attempted.

The outcome of this stage is the set of reconfigured routers and links forming the temporary composite network, along with any configuration information needed to restore the separate networks to their pre-emergency state when the emergency has passed.

### 7.5 Summary

The whole process of reconfiguration to form a composite network is technically complex, and there may be considerable commercial and legal obstacles to forming one. We have outlined the networking stages needed and which will be researched. Investigation of legal aspects is outside the scope of the present project.

## 8 Conclusions

The trend in wide area networks is to have a number of competing for-profit suppliers with separate networks servicing the same region. Individually, each may lack excess capacity to handle bursts of communications during an emergency, and have too many potential points of failure. If these could be combined into a single composite network for the duration of the emergency, and access to the area controlled by a policy set by the emergency services management, the burst capability would be improved and duplicate network entities would provide some capability for adaptive recovery from damage. This

paper has outlined a project to achieve that goal by converting separate independent TCP/IP networks in a region into a temporary single cohesive network in a time of emergency.

Some of the major technical problems have been discussed. The review of related unclassified work indicates that this problem is quite difficult and is not actively being researched and published. Earlier work on policies for routers assumed the continuing independence and autonomy of all of the separate networks; our work seeks to merge the separate ASes into a single composite AS.

Requirements for a solution have been outlined, along with a research project involving the distributed reconfiguration of routers.

Our system is neutral to the type of links involved, so a network with all radio or satellite links, as proposed in Kobe, is within our scope.

Although initially oriented towards management of emergencies, the outcomes also apply to networks in normal situations. Setting up a large ad hoc network consisting of nodes and wireless links would be simplified with our system. Within a single Autonomous System, the whole network could be set up and maintained using high level policies, with potential improvements in performance, fewer network configuration errors, and a more enjoyable lower stress workplace for network managers.

## References

- [1] Yoichi Shinoda, Tomomitsu Baba, Nobuhiko Tada, Akira Kato, and Jun Murai, "Forethought and Hindsight: Experiences from the First Internet Disaster Drill", 1996, INET 1996. Accessed 24 November 2005.
- [2] Barry M. Leiner, Vinton G. Cerf, David D. Clark, Robert E. Kahn, Leonard Kleinrock, Daniel C. Lynch, Jon Postel, Larry G. Roberts, and Stephen Wolff, "A Brief History of the Internet", Tech. Rep., Internet Society, Oct. 2005.
- [3] Internet Society, "A Brief History of the Internet and Related Networks", Tech. Rep., Internet Society, Oct. 2005.
- [4] D. Clark, "Policy Routing in Internet Protocols", May 1989, RFC 1102.
- [5] T. Bates, E. Gerich, L. Joncheray, J-M. Jouanigot, D. Karrenberg, and M. Terpstra. J. Yu, "Representation of IP Routing Policies in a Routing Registry (ripe-81++)", Mar. 1995, RFC 1786 and RIPE-181.
- [6] Nobuhiko Tada, Yukimitsu Izawa, Masahiko Kimoto, Taro Maruyama, Hiroyuki Ohno, and Masaya Nakayama, "IAA System ("I Am Alive"): The Experiences of the Internet Disaster Drills", 2000, Proceedings of INET 2000. Accessed 24 November 2005.
- [7] James Cowie, Alin Popescu, and Todd Underwood, "Impact of Hurricane Katrina on Internet Infrastructure", 9 September 2005.
- [8] Cisco Networking Academy, "Cisco Certified Networking Associate, Module 2. Version 3.1", Mar. 2005.
- [9] Chris Griffith, "Telcos set to pull plug on cheap Net calls", p. 12, 5 January 2006.
- [10] M Little, "Goals and Functional Requirements for Inter-Autonomous System Routing", Oct. 1989, RFC 1126.
- [11] RIPE (Reseaux IP Europeens), "RIPE Terms of Reference", 29 November 1989, RIPE 1.
- [12] Tony Bates, Jean-Michel Jouanigot, Daniel Karrenberg, Peter Lothberg, and Marten Terpstra, "Representation of IP Routing Policies in the RIPE Database", Feb. 1993, RIPE-81.
- [13] M. Steenstrup, "IDPR as a Proposed Standard", July 1993, RFC 1477.
- [14] M. Steenstrup, "An Architecture for Inter-Domain Policy Routing", June 1993, RFC 1478.

- [15] M. Steenstrup, "Inter-Domain Policy Routing Protocol Specification: Version 1", July 1993, RFC 1478.
- [16] C. Villamizar Avici, C. Alaettinoglu, D. Meyer, and S. Murphy, "Routing Policy System Security", Dec. 1999, RFC 2725.
- [17] J. Boyle, R. Cohen, S. Herzog, R. Rajan, and A. Sastry, "The COPS (Common Open Policy Service) Protocol", Jan. 2000, RFC 2748.
- [18] C. Villamizar, C. Alaettinoglu, R. Govindan, and D. Meyer, "Routing Policy System Replication", Feb. 2000, RFC 2769.
- [19] K. McCloghrie, M. Fine, J. Seligson, K. Chan, S. Hahn R., R. Sahita, A. Smith, and F. Reichmeyer, "Structure of Policy Provisioning Information (SPPI)", Aug. 2001, RFC 3159.
- [20] L-N. Hamer, B. Gage, B. Kosinski, and H. Shieh, "Session Authorization Policy Element", Apr. 2003, RFC 3520.
- [21] L. Blunk, J. Damas, F. Parent, and A. Robachevsky, "Routing Policy Specification Language next generation (RPSLNg)", Mar. 2005, RFC 4012; Updates: 2725, 2622.
- [22] C. Alaettinoglu, T. Bates, E. Gerich, D. Karrenberg, D. Meyer, M. Terpstra, and C. Villamizar, "Routing Policy Specification Language (RPSL)", Jan. 1998, RFC 2280.
- [23] C. Alaettinoglu, C. Villamizar, E. Gerich, D. Kessens, D. Meyer, T. Bates, D. Karrenberg, and M. Terpstra, "Routing Policy Specification Language (RPSL)", June 1999, RFC 2622; Obsoletes: 2280.
- [24] Cisco Networking Academy, "Cisco Certified Networking Professional, Module 1. Version 3.1", Mar. 2005.
- [25] K. Carlberg and R. Atkinson, "General Requirements for Emergency Telecommunication Service (ETS)", Feb. 2004, RFC 3689.
- [26] K. Carlberg and R. Atkinson, "IP Telephony Requirements for Emergency Telecommunication Service (ETS)", Feb. 2004, RFC 3690.