Support Group Application Note
Number: 283
Issue: 1.01
Author: DW

Acorn

TCP/IP Addressing, Subnetworking and Interoperability: an Overview

This document describes the different types of network address available in the IP world, and how they may be subdivided using IP netmasks. The concept of "special" IP addresses is also introduced, the "special" addresses which must not be assigned to machines are listed, and the procedures required to allow TCP/IP to interoperate with AUN are detailed.

Applicable Hardware :	All hardware attached to an IP network
Related Application Notes:	

Support Group
Acorn Computers Limited
Acorn House
Vision Park
Histon, Cambridge
CB4 4AE

Introduction

TCP/IP, the Transmission Control Protocol / Internet Protocol, is a world standard underlying protocol for communication between networks of computers; IP has been implemented on most operating systems (including RISC OS), and the Internet uses it exclusively.

As befits a protocol capable of supporting such a complex organism as the Internet, IP is itself not simple; a full internetworking "stack" comprises hardware control layers, IP, TCP, UDP and any combination of the wide range of communication protocols which themselves need the underlying IP system to operate. An example of a full internetworking protocol stack is pictured below; although we only intend to cover the IP layer and its interface to the hardware drivers and the internetwork in this Application Note, it is intended that some of the higher-level protocols, their configuration and use will be covered in a future document.

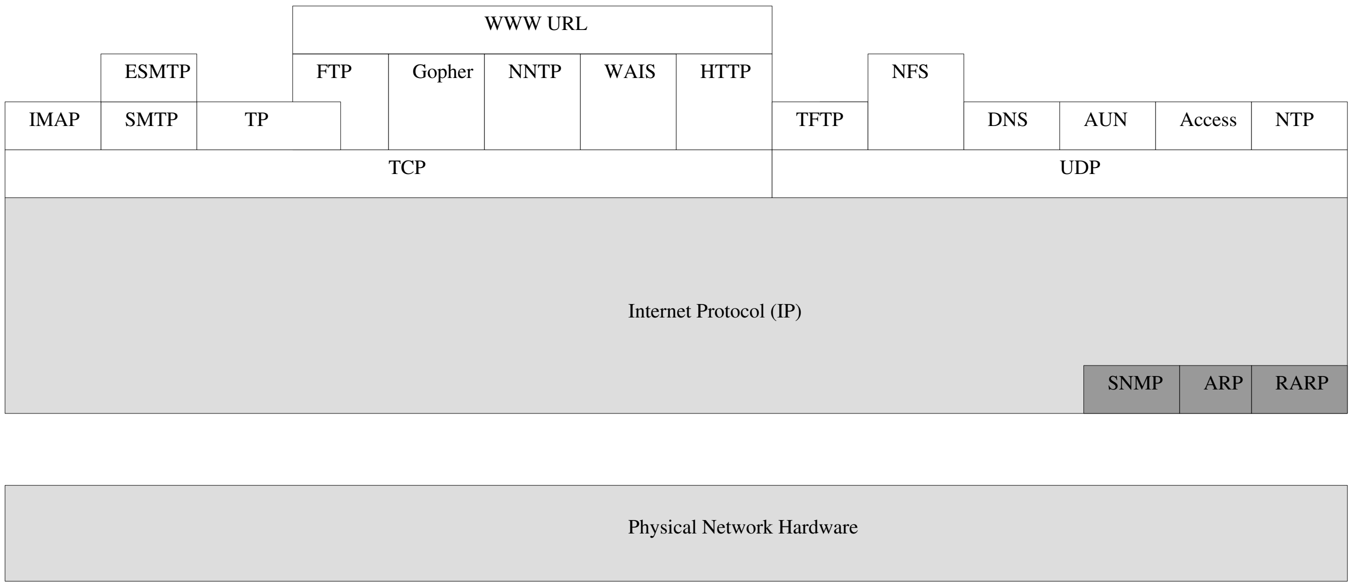

Figure 1: A Full Internetworking Protocol Stack

Internet Protocol

Background: Network Classes

IP addresses are divided into five address classes, categorised as A to E.

The first three classes, A, B and C, are available for normal allocation and for server-to-server communications. Class D is reserved for multicast systems (support within the DCI4 suite allows specifically addressed multicast packets to be passed, but has no mechanism provided to allow finer granularity than "all multicasts or none at all"), and Class E currently has experimental status only.

Class E is likely to become important, however, in that is is earmarked for future use with single-machine domains (ie home users, and educational establishments with just one machine connected to the Internet); the importance of Class E is likely to be emphasised significantly in the next revision of the Internet Standards documents.

There are no practical distinctions in the way in which computer systems with class A, class B or class C addresses use these addresses. With some exceptions (covered in the "Special Addresses" section below), all addresses are equivalent for communication purposes.

The distinguishing feature of each class of address is the number of network numbers, and the number of host connections which each of those network numbers can support. These restrictions are detailed below:

Network Class	Max. number of network numbers	Max. number of hosts on each network number
A	126	16777214
B	16392	65534
C	2097150	254
D	Not applicable: Reserved for multicast systems	
E	Reserved for future use	

Table 1: Network Classes and Restrictions

Background: IP Addressing

An IP address takes the form of a 32-bit binary number, which must be unique throughout the whole of the connected network. To make the address more easily readable by humans, it is often represented in "dotted decimal" form as a.b.c.d, where each of a, b, c and d is a decimal number in the range 0-255.

The 32-bit IP address has two fundamental components; a network number and a host number. An optional subnetwork number may also be explicitly inserted in the address; this is dependent on the subnetwork mask, discussed later.

The network number appears at the left hand side (MSB) of the IP address, and contains some bits which can be used to deduce the class of the network associated with the number. In brief:

• If the most significant bit in the most significant byte is 0, the address is class A

• If the two most significant bits in the most significant byte are 10, the address is class B

• If the three most significant bits in the most significant byte are 110, the address is class C

• If the four most significant bits in the most significant byte are 1110, the address is Class D

• If the four most significant bits in the most significant byte are 1111, the address is Class E

It is worth clarifying what constitutes a "host" in the usual situation of a machine running one network stack and having one network card, a host can be thought of as a single machine. However, machines which have more than one physical network interface (eg systems which are configured as gateways) have as many host connections as they have network connections to gateway between, and a multiple-personality system running multiple IP stacks under different operating systems (eg a Risc PC using the DCI4 interface so that it can run !Internet from RISC OS and some other stack implemented in whatever OS is running in the 486 environment) has as many host IDs as it has IP stacks.

Unlike some other large-scale schemes which assign numbers to devices (a classic case being the public telephone network), IP addresses contain no information which gives any clues regarding the geographical location of the device which owns the IP address; IP addresses are not hierarchical. Given an IP address, the only thing you can deduce from it is a management auhority; that authority could manage one network which itself has global coverage and intermeshes with many other networks.

Choosing Network Numbers

The choice of network numbers is very important in relation to the way in which your network is routed, and there are a few simple rules which should be obeyed:

• Network connections with the same network number should communicate directly on the same physical LAN.

• Network connections with different network numbers do not communicate directly; they must use the services of a router. The router can either be directly connected to the LANs which supplied and which need to receive the packet, or it may be connected via a number of indirect steps (ie other routers) to a router which knows about the LAN which needs to receive the packet.

It is strongly recommended that, if you are planning to connect your own IP-running network to someone else's at any point, that you apply to have an address registered for your site.

"Special" Addresses

There are a number of IP addresses which are reserved for particular special functions. These addresses are listed below; such addresses must **NEVER** be assigned to a computer as the host IP address.

• 127.*n.n.n*, where *n.n.n* is any set of numbers. This is reserved for local software loopback. Note that IAB standards say that any address beginning with 127 must not be transmitted across the physical interface or appear on the cable. As not all software implementations check IP addresses before transmitting, some IP datagrams with source or destination address of 127.<*something*> may appear on the cable; this has been known to cause networks to fail.

• A network number of all 1s

• Host number 0 is reserved to refer to a particular network number. For example, the first class C address which can be allocated is 192.0.0.1

• The IP address 0.0.0.0 is reserved, and considered to be classless. It is used in two ways: as a secure address when the system does not know its genuine address (eg during the bootstrap phase of a UNIX workstation booting across a network, when the station does not know the true network address of that network) and by routers in a list of addresses to advertise the default route, this being the route to all networks which are not explicitly otherwise listed. Note that this "default route" is not a source or destination address, merely an entry in a table.

• The IP address 255.255.255.255 is reserved as a destination address to mean "broadcast this packet to all other systems on this network." 0.0.0.0 as a destination is an obsolete form of 255.255.255.255.

If one or more of these addresses are assigned to your networked machines, the errors which will occur are likely to be obscure, apparently intermittent and difficult to isolate; remember that not all network software checks to see that the IP address it is requested to use is a valid one.

You should also note that, once you start introducing subnetworking, other IP addresses become reserved, and assigning these addresses to machines can cause very similar problems; more details of this follow later.

The Reason for Netmasks: Subdivision of Networks

Subnetworking increases the network manager's control over the network address space, and provides a mechanism for using routers when only one or a small number of full network numbers is available. The mechanism for creating subnetworks is the subnetwork mask (or Netmask); this is another 32-bit number which is allocated at the same time as the IP address.

Subnetworking divides the normal range of host IDs (16 million for class A, 65534 for class B and 254 for class C) into a number of subnetworks, with a reduced number of systems on each subnetwork. The product of the number of subnetworks and the (number of stations + 1) on each subnetwork cannot exceed the number of stations which could live on the original, undivided network.

When planning a large new network installation (where "large" implies more than 200 stations on the network) which is intended to be connected to the Internet, it pays dividends to plan for subnetworking even if the facility is not used initially; in addition to the labour-intensive nature of changing the large number of addresses which will require changing if subnetworking is introduced to a network which has not been prepared for it, it is highly unlikely that your organisation would be allocated more than one network number.

Of course, there are limitations imposed by the basic structure of Class A, B and C network addresses; there cannot be a greater number of connections than the basic address type will permit, and in practice you will always have slightly fewer available addresses owing to certain addresses becoming reserved.

What you lose in raw numbers of connections, however, you gain in flexibility and control.

Small networks which do not use routers need not use subnetworking, if such networks are unlikely to grow beyond a few physical LANs and up to 200 stations. If you do not need to use a registered address and therefore do not mind how many IP network numbers you use, then there is less incentive to use subnetwork addressing.

Interaction of the IP Address and the Netmask

Like the IP address itself, the subnetwork mask is a 32 bit number. The interaction between the two numbers works according to the following rule:

• If there is a 1 in a bit position in the subnetwork mask, that bit forms part of the network number in the network address space

• If there is a 0 in a bit position in the subnetwork mask, that bit is part of the host IP address.

An example of the use of a subnetwork mask is given below:

	Binary	Dotted Decimal
IP Address	10000001.10000010.01001111.01010101	129.130.79.85
Subnet Mask	11111111.11111111.11111000.00000000	255.255.248.0
Network ID	10000001.10000010.01001000.00000000	129.130.72.0
Host ID	00000000.00000000.00000111.01010101	0.0.7.85

Table 2: Interaction of IP Address, Network ID and Netmask

The IP address AND NOT (the subnet mask) is the host ID, and the IP address AND the subnet mask is the network ID. In this case, the site has an assigned network ID of 129.130.0.0 (class B) and is using subnetworking to generate the network 129.130.72.0.

If you choose not to use subnetworking, you will still need to insert the default subnetwork masks which select the normal network numbers. The default masks are:

Address Class	Default mask (dotted decimal)	Default mask (hex)
Class A	255.0.0.0	FF000000
Class B	255.255.0.0	FFFF0000
Class C	255.255.255.0	FFFFFF00

Table 3: Default Subnet Masks by Network Class

Reserved Subnet Numbers

Health Warning: This section is somewhat mind-bending, and is only necessary for a complete understanding of the limitations in allocating subnet numbers; it may be glossed over on a first reading of this document.

Subnetting is further complicated by the fact that an address of all 1s is used for broadcasts to all subnets within a network. A host ID of "all 1s" is still the broadcast to all systems, either on a specific subnetwork or all subnets with this network. This reduces the number of available addresses to assign to hosts.

A subnet mask of 255.255.128.0 it is not useful, as 128.1.0.*n* refers to "system *n* on this subnetwork" (see Table 4) and 128.1.255.255, which would be a broadcast to all hosts on 128.1.128.0, also means "broadcast to all hosts on the network 128.1.0.0".

Additionally, 128.1.255.255 could be used to mean "broadcast to all hosts on the subnet 128.1.128.0"; this condition is illustrated in Table 5. As the two conditions are indistinguishable, this subnet cannot be used.

	Binary	Dotted Decimal
IP Address	10000000.00000001.00000000.00000011	128.1.0.3
Subnet Mask	11111111.11111111.10000000.00000000	255.255.128.0
Network ID	10000000.00000001.00000000.00000000	128.1.0.0
Host ID	00000000.00000000.00000000.00000011	0.0.0.3

Table 4: The Significance of 128.1.0.n

	Binary	Dotted Decimal
IP Address	10000000.00000001.11111111.11111111	128.1.255.255
Subnet Mask	11111111.11111111.10000000.00000000	255.255.128.0
Network ID	10000000.00000001.10000000.00000000	128.1.128.0
Host ID	00000000.00000000.01111111.11111111	0.0.127.255

Table 5: Anatomy of a Broadcast Address Clash

The first mask which *can* be used is the next mask of 255.255.192.0 giving four subnet values, but *only two usable* subnetwork numbers of 128.1.64.0 and 128.1.128.0. While 128.1.192.0 at first seems usable, the address 128.1.255.255 again means "broadcast to all hosts on network number 128.1.0.0". There is no separate way of defining a broadcast to all systems on the single subnet 128.1.192.0. Any address 128.1.0.*n* refers to system *n* on "this subnet," instead of on the specified subnet 128.1.0.0, so that subnet is also unusable.

The subnet 0 should not be used, since 0 refers to the network as a whole, and is inherently associated with the broadcast address, even though this usage is now deprecated.

When represented in dotted decimal, it can be very difficult to recognise the broadcast address for a single subnet; it is often better to do most of the working out of subnet allocations in binary or hex, and convert back to decimal if required as the last stage.

Choosing a Netmask

Address management is much simpler if the same subnet mask is configured on all systems which share a particular network number. When you choose your netmask, you must bear in mind that it defines the point in the final IP address where host numbers are replaced by network numbers; to avoid complications later on, it is strongly recommended that you choose your netmask to be a contiguous block of 1s followed by a contiguous block of 0s, as in the example in Table 2. Hosts on different subnetworks require an intervening router, so changing the subnet mask may equally be thought of as exchanging repeaters for routers. Before choosing the subnet boundary, you need to know:

• how big the largest subnetwork is which you can manage using repeaters alone

• what the cost, performance and management overheads are in using routers in select places on your network in place of repeaters

The importance of these issues is that, for a network with a single subnet mask, all subnetworks will have the same maximum size. For a class B network, for example, most managements would be likely to choose from:

• 254 subnets of 254 hosts
• 126 subnets of 510 hosts
• 62 subnets of 1022 hosts
• 30 subnets of 2046 hosts

These figures produce an acceptable balance between repeaters and routers; the same subnetting principles apply to the smaller scale of a class C address, although there will of course be a smaller number of hosts addressable.

TCP/IP and AUN

AUN (without the full TCP/IP suite in coexistence) at DCI2 level operates using the InternetA module; essentially, an AUN packet propagating across Ethernet comprises an Econet packet wrapped up in an IP header. AUN takes the approach of using class A IP addresses with a class B subnet mask, with the result that the IP address of an AUN station looks like:

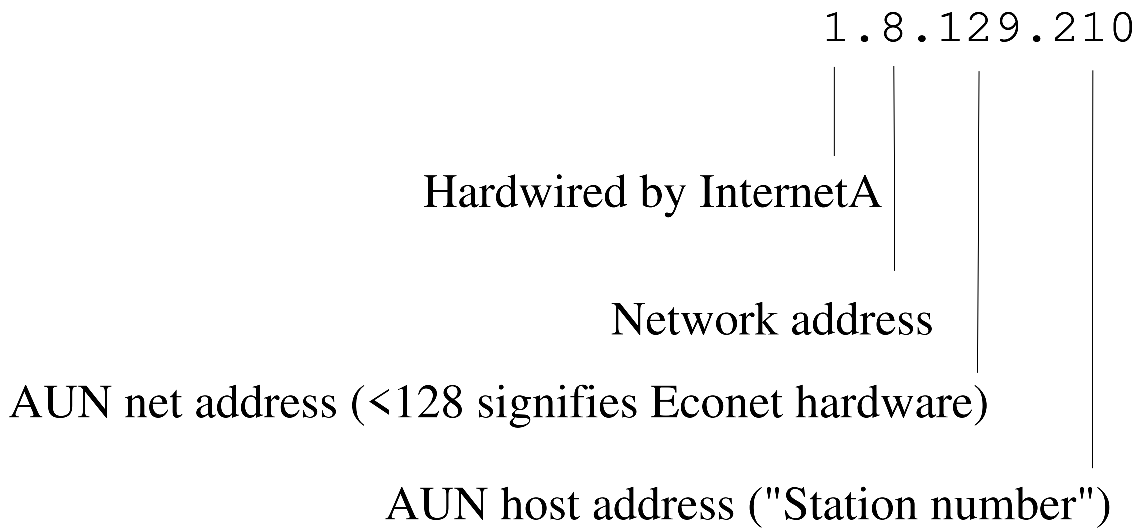

Figure 2: Composition of an AUN address as seen by TCP/IP

The "network address" is related to the number of the individual LAN containing the required station, and which is bridged into the single visible network.

It is worth noting here that the InternetA module, which forms part of the DCI2 !Bootnet application (supplied as part of AUN Level 4) has the first field of its dotted decimal representation hardwired to 1. If you wish to have a link of any form between such a network and the Internet, you will need to gateway the system in such a manner that the Internet (or whatever other IP-supporting network you are connecting to) cannot see the full IP addresses of the machines on the network running InternetA. Alternatively, you can replace InternetA with the full Internet module (contained as part of the new DCI4 suite, which has been released by Acorn as freeware and is available by anonymous ftp from ftp.acorn.co.uk and mirror sites) and hence remove the hardwiring problem. This may be done in a boot sequence by simply disabling the AUN client software and running !Internet and !Bootnet in that order. However, with this software loaded, it is not possible to use the !Gateway application.

It is necessary to ensure that mapping between the two-byte *net.station* addresses used by AUN and the four-byte *site.network.net.station* IP addresses is set up correctly; this is achieved by explicitly mapping AUN nets to IP net addresses in the !Bootnet.Files.AddMaps file, detailed in the "AUN and TCP/IP" appendix of the AUN Manager's Guide. The format of this file consists of a series of lines having the syntax

`addmap` *byte1.byte2.byte3*`.0` *net*

where *byte1, byte2* and *byte3* are the first three bytes of an IP network address and *net* is the number of the AUN net which is to correspond to that network. Individual station numbers are preserved, hence the command

`addmap 1.1.129.0 130`

would cause the AUN address of 130.57 to be translated to the IP address of 1.1.129.57

Note that, as addmap maps one net number onto another net number, the final byte in the IP address referenced in an addmap command must be 0. Any other value will produce unexpected results.

Addmap is capable of supporting IP addresses beyond the InternetA network address restrictions, so all classes of IP network can have AUN addresses associated with them.

Example: Subnetting a Class A Network for IP and AUN

The example of a class A network is being used for two reasons, these being that many sites who decide that they will never need to connect their internal network to any external systems will gain more flexibility from using a class A addressing scheme, and because (as detailed above) AUN imposes a class A scheme by default.

An AUN network gives the sysadmin the base network address of 1.0.0.0 to work from; the class B subnetting AUN uses gives a default netmask of 255.255.0.0 to work with. If you look at this from a purely IP viewpoint and consider the default netmask, then the second-most significant byte in the network address appears as though it can never change; however, in AUN the software integral to !Gateway determines the value of this byte (and for Ethernet networks, the third byte too), independently of the rest of the system. If Econet is still in use and the station in question is connecting using Econet or another Econet-like protocol (indicated by an AUN net address <128), the second byte is determined by !Gateway and the third is determined by an Econet bridge.

If you plan to use !Gateway, you must ensure that the only IP packets which will be passing through the !Gateway machine have an address of 1.*n.n.n*; otherwise you will have to use a third-party gateway system which supports the range of IP addresses appropriate to your network.

Once the subnetting has been set up in the IP universe (the netmask is usually left at FFFF0000, and any necessary changes in the second byte are determined by gateways), AUN routing can be set up using **`addmap`** as above. IP routing, where applicable, needs to be done using the **`route`** command; this is documented in the TCP/IP User Guide and Application Note 284.

Integrating TCP/IP and Acorn Access

This requires the building of a two stage boot sequence which includes a file which is accessed before the desktop is active and a file which is called when the desktop becomes active.

If a !Boot file; of type Desktop, already exists then rename it as Desktop.

Using !Edit create an Obey file which contains the following lines (amending the application pathnames where necessary):

```
RMKill ShareFs
RMKill Freeway
RMKill AccMsgs
RMKill InternetA
Filer_Run ADFS::4.$.Network.!Internet
Filer_Run ADFS::4.$.Network.!BootNet

RMReinit AccMsgs
RMReinit Freeway
RMReinit ShareFs
Desktop -f ADFS::4.$.Desktop | Used if there is a Desktop file
|Desktop | Used if there isn't a Desktop file.
```

Note: The pathnames to the applications and files may vary according to the disc structure used. The order in which the modules are **`*RMReinit`**-ed is vital.

Save this file as `!Boot` and reset the machine.

Required Reading

Acorn Computers Ltd, "TCP/IP Protocol Suite Installation Guide"
The DCI4 equivalent of this document is supplied in plaintext form in the "Docs" supdirectory which comprises part of the DCI4 stack distribution archive, downloadable from ftp.acorn.co.uk

Further Reading

Washburn and Evans, "TCP/IP, Running a Successful Network" (Addison-Wesley) 1993
Hegering and Läpple, "Ethernet - Building a Communications Infrastructure" (Addison-Wesley)
O'Reilly, "TCP/IP Network Administration" (O'Reilly)