

IP Address Assignment in Large Industrial Networks

Andy Swales, Network Vision, Inc. November 2003



Background

There has been much recent interest in automating the assignment of IP addresses when Ethernet is used as an industrial control network.

A number of large end-users and vendor user groups are attempting to determine the best solution.

This paper presents a software technique offering significant advantage over all alternative proposals.

Problem description

As use of Ethernet has increased in the industrial space, particularly for connection of simple sensors and I/O devices to PLC's and SCADA systems, customers are having to rethink device addressing techniques.

All devices on a TCP/IP communication network require assignment of an IP address before they can communicate with any other station.

These IP addresses must be assigned carefully, since a misconfigured IP address can cause significant disruption on a functioning network, particularly in cases where a new device is assigned an address already in use somewhere else on the network.

Industrial networks have the characteristic that they are expected to be continuously available, but individual devices have finite failure rates and the number of devices requiring replacement during the operating life of the network will be large.

The failure rate for an individual class of device may be sufficiently low that maintenance and restart procedures may be unfamiliar to the repair technician at time of replacement.

Any delay researching the correct replacement and restart procedures, or required coordination between maintenance technician and network services professional, will lead to increased 'Mean Time To Repair' (MTTR) and thus loss of production.

The replacement operations being considered must include not only the sensors and actuators themselves, but also major infrastructure components such as network switches and PLC's.

The device addressing strategy must be applicable to all common and likely styles of industrial network, including dedicated Layer 2, dedicated Layer 3, and mixed-use layer 2 or layer 3. (A dedicated network is one designed only to support the process control equipment, and not general purpose transient computers such as user laptops. A mixed-use network is one where the same wiring infrastructure is used to support both functions).

It must also be usable in environments where there is no direct communication from a central location to all devices on the industrial networks, such as where a vendor includes a private Ethernet network as part of a machine, only accessible through dedicated interface on the tool control computer or PLC.

Choices of mechanism

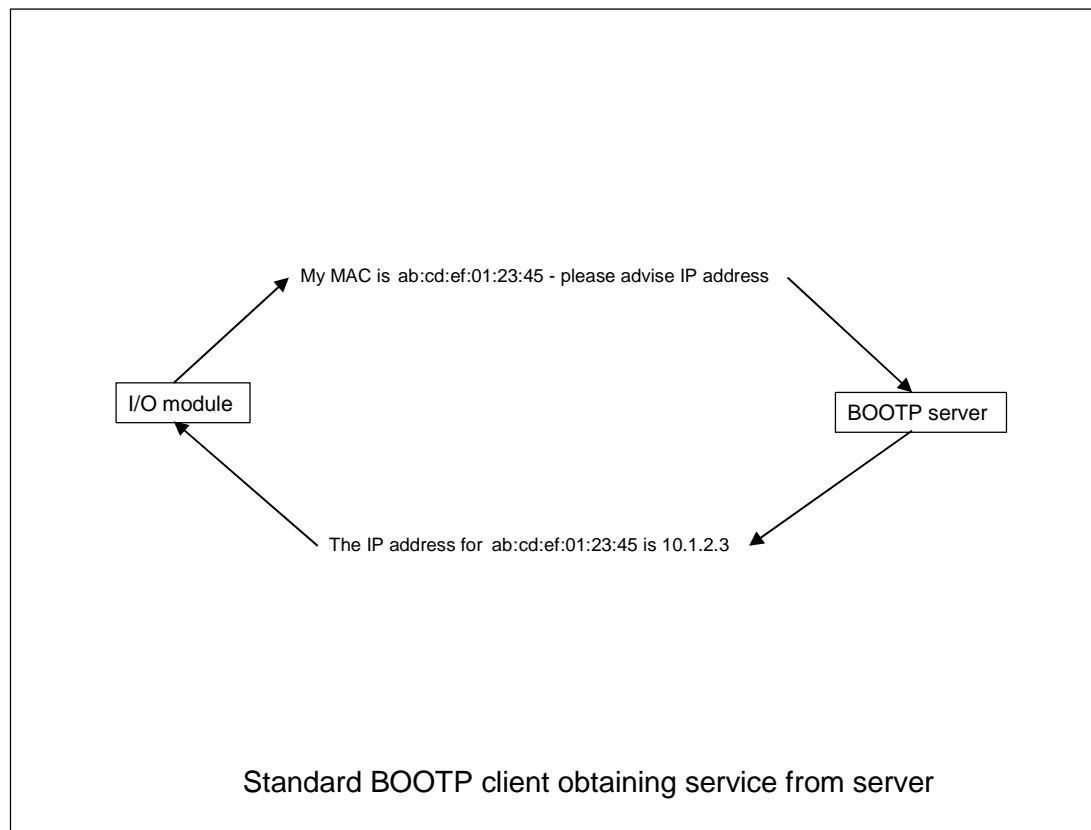
There are three standardized IP address assignment schemes in wide use already, and two more recent variations.

1. Bootstrap Protocol (BOOTP)

This leverages the fact that the physical Ethernet address (a.k.a. MAC address) of every device is unique. A computer somewhere on the network (the 'BOOTP server') is set up to listen for BOOTP request packets, which will be broadcast by each device when it starts up. The BOOTP request message contains information including the MAC address of the requestor. The BOOTP server contains a database matching each IP address (and other required supporting data) with the MAC address of the requestor. When the BOOTP request is received, and if there is a 'match' in the database, the server will respond with the correct IP address information.

BOOTP addressing has been defined since 1983, and is easy to implement on an I/O device. Almost all mainstream industrial control and network infrastructure vendors include BOOTP assignment as an alternative to their vendor-specific static addressing schemes.

Devices requiring BOOTP allocation can have their requests serviced by modern DHCP servers, such as those in Windows 2000 and Linux. Such a DHCP server will allow 'reservation' of IP addresses to a particular MAC address.



2. Reverse Address Resolution Protocol

An earlier definition than BOOTP, this method has almost totally been supplanted by BOOTP, because it cannot be used in situations where the server and target are in different IP subnets.

3. Dynamic Host Configuration Protocol (DHCP)

A later variation of BOOTP, designed to solve the problem of assigning temporary addresses to devices such as user laptop computers which frequently ‘appear’ and ‘disappear’ on a network.

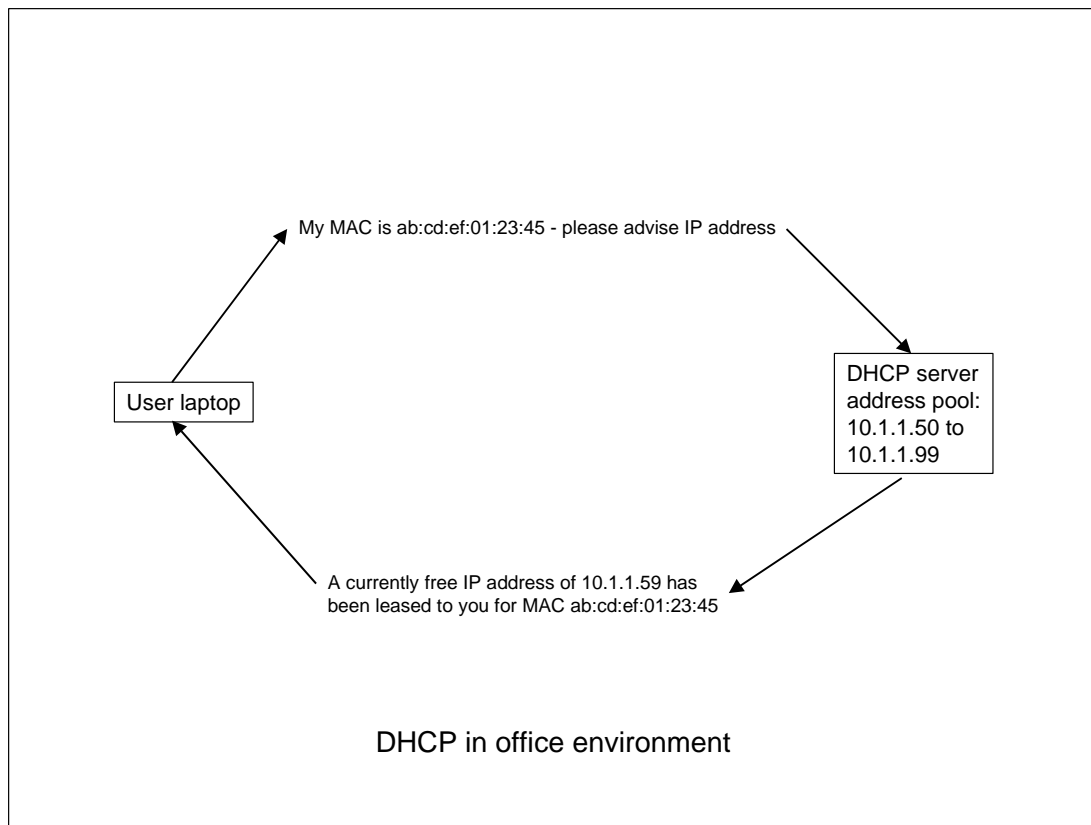
Instead of maintaining a rigid database where the MAC address must be pre-registered, as with BOOTP, a DHCP server instead keeps a ‘pool’ of addresses, each of which may be ‘leased’ or ‘free’ at any time. When a DHCP server receives a request for an address, it checks the table to see if an address is already allocated for that device (based on the MAC address and IP subnet), and returns it if so. If not, but there are free addresses in the ‘pool’, a ‘lease’ is created from one of the ‘free’ addresses, and this is allocated.

This method is very convenient for devices such as laptop computers which only run ‘client’ software such as web browsers, and do not care which IP address they actually are allocated so long as it works.

However this dynamic assignment method is not useful for most PLC and I/O devices, or network switches and web servers, which require to be allocated a well-defined IP address. Most PLC’s expect to transmit configuration changes and similar directly to an IP address, and must know in advance what IP addresses to use for each target.

Also it is very cumbersome to arrange for redundancy of a DHCP service so that addresses can be allocated by a standby server if the primary server is unavailable. So although the technique works reasonably well in an office environment, its value is limited in industrial control.

Most industrial control vendors have avoided use of DHCP protocol in favor of BOOTP, except for devices such as operator interfaces.



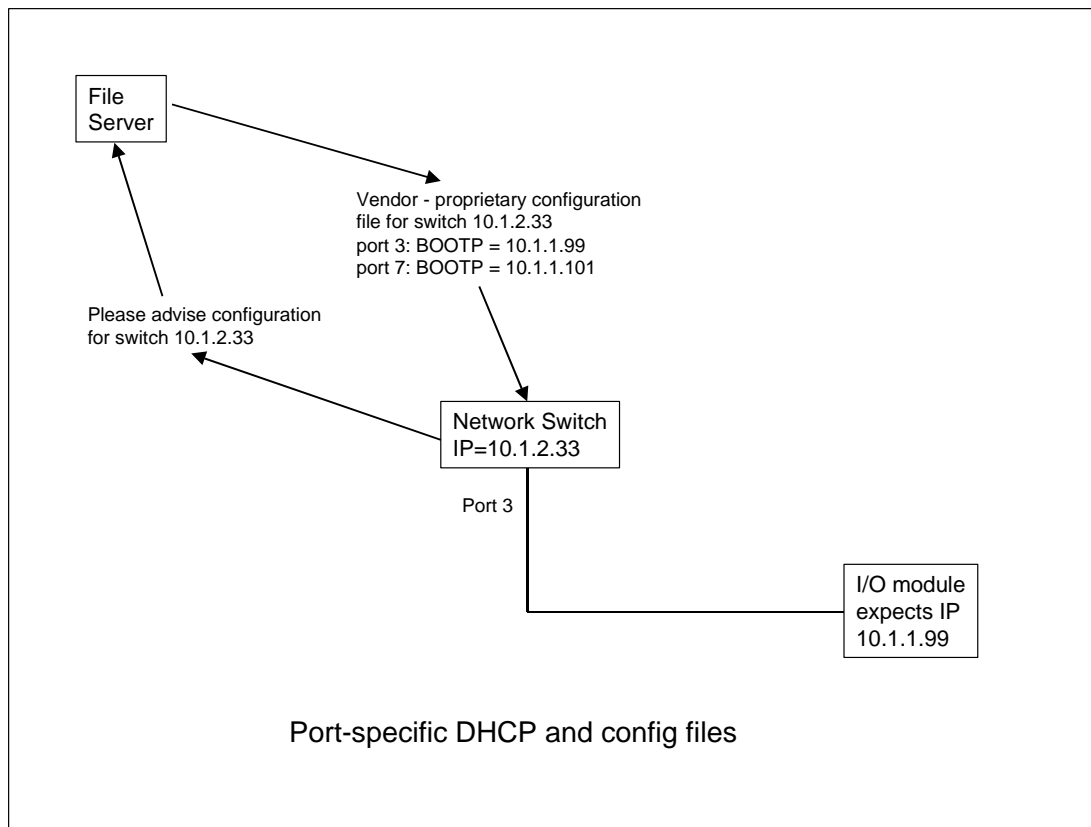
4. Port – specific DHCP

A number of vendors of Ethernet switch provide what is essentially distributed BOOTP/DHCP service. When a BOOTP or DHCP request is received on one of the ports of the switch, instead of the request being passed up the network to the ‘real’ BOOTP/DHCP server, it is answered locally.

In order to do this, the switch consults a local configuration table, associating port numbers with IP addresses.

The main problems with this approach are that the BOOTP database is held in the switch itself, and that there is no standardization between the databases of vendors. So although this mechanism makes it convenient to replace the end device, it is much more complicated if the switch itself fails and must be replaced.

Although it is common to require intricate configuration for routers (Layer 3 switches), the need for similar care in configuration and support for the layer 2 switches significantly increases their effective cost.



5. DHCP Option 82 Relay

Some Cable TV and DSL Internet Service Providers have made use of a DHCP option originally intended to allow end user devices to classify themselves to the DHCP server, and thus allow the DHCP server to allocate them different settings because of their particular needs.

The trick is performed by having the managed switch intercept an incoming BOOTP or DHCP request, just as in variant 4 above, and note which port the request came from. The DHCP request is then regenerated as a ‘DHCP forwarding request’ and at the same time an ‘Option 82 record’ is added which identifies which switch and which port received the query.

The DHCP request is forwarded to the ‘real’ DHCP server elsewhere on the network, where IP address allocations have been previously associated with the ‘Option 82’ data.

Some infrastructure vendors, in particular Cisco and Hirschmann, have proposed adoption of this technique for industrial IP address assignment.

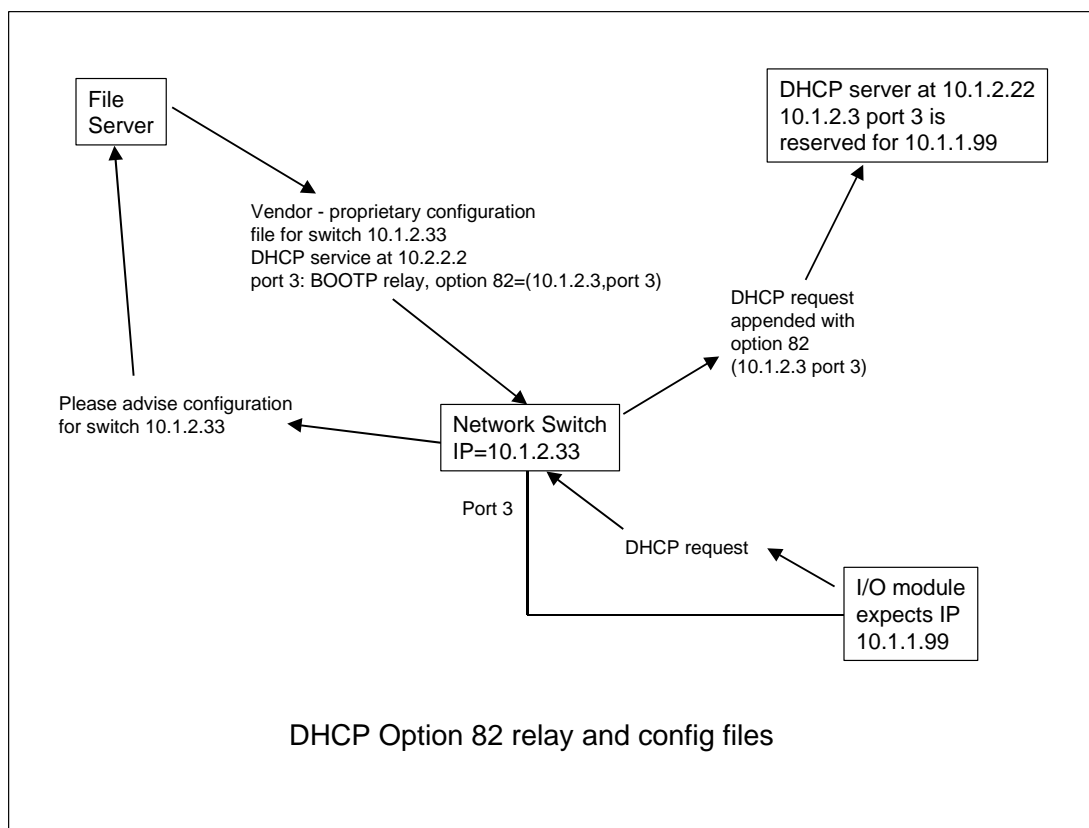
This method can be made to work, but requires special configuration at the switches, just as in variant 4. It is (at minimum) necessary for the switches to determine which of their ports support Option 82 forwarding, and which should leave the DHCP requests alone. Also the address of the DHCP server to which the requests must be

forwarded must be known. And the DHCP server must be a very modern one which understands Option 82 (ruling out the standard DHCP server on Windows 2000 for example)

In the event of replacement of a switch, all this information must be recreated. The methods by which switch vendors achieve this are usually to require maintenance at some central location of a file server containing images of the configuration files for each switch, and the switches are configured to query the server at restart to obtain their configuration.

To make provision for an additional I/O device on a network, it is now necessary to update two files – the central DHCP database, and the configuration file for the switch whose port will be attached

And of course, the end user must use switches and DHCP server software which support this use of Option 82 relay. Since the specification is very new, only the newest models of switch from each of the vendors is likely to support the facility, and that can imply costly infrastructure upgrades. And because of limited choice of vendor and model, the issues of vendor lock-in are significant.



6. Network Vision Auto-IP assignment

This technique combines the best features of variants 4 and 5, and significantly reduces the administrative burden while avoiding vendor lock-in by supporting almost all existing managed switches.

All managed switches support Simple Network Management Protocol (SNMP), and specifically the 'Bridge MIB' defined in RFC1493. They can therefore be asked by a central management station to report which port a particular MAC was found on.

The Auto-IP technique involves a modified BOOTP server which is capable of issuing SNMP queries to switches to identify the switch and port number. The server uses this information, along with knowledge of whether a station is currently up or down, to determine whether a newly-seen MAC represents a new device or a replacement for an existing one. The sequence is as follows:

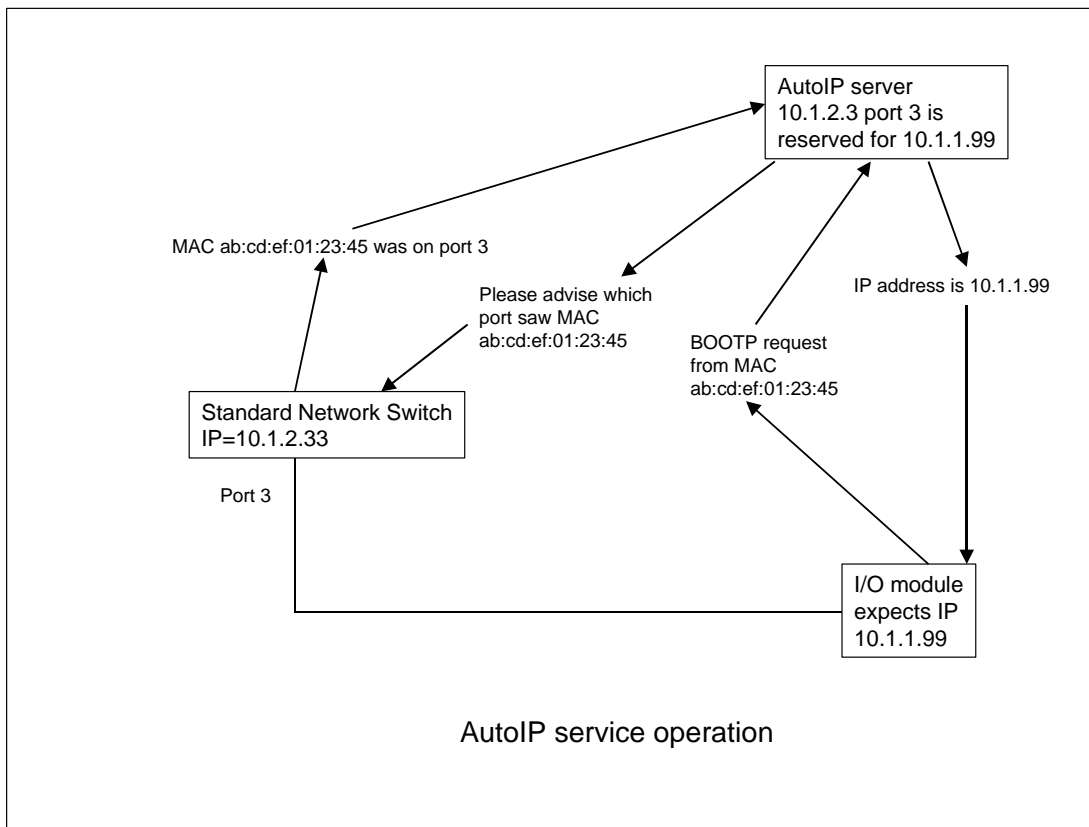
An I/O device issues a BOOTP request (which naturally includes its MAC address)

The Auto-IP server sees the request, and checks its database for a MAC match. If the MAC is already known, the IP address is returned just like in a conventional BOOTP service.

If the MAC address does not match, Auto-IP sends out SNMP queries to the switch(es) to determine which port of which switch sent the request.

If the switch address and port number matches one for which automatic assignment is configured, it is considered a 'potential replacement'.

If there is one and only one Auto-IP-configured device at that switch and port position which is currently 'down', the replacement is authorized. The IP address for the 'down' device is transferred to the new MAC and thus recorded in the database, and the correct BOOTP response is sent.

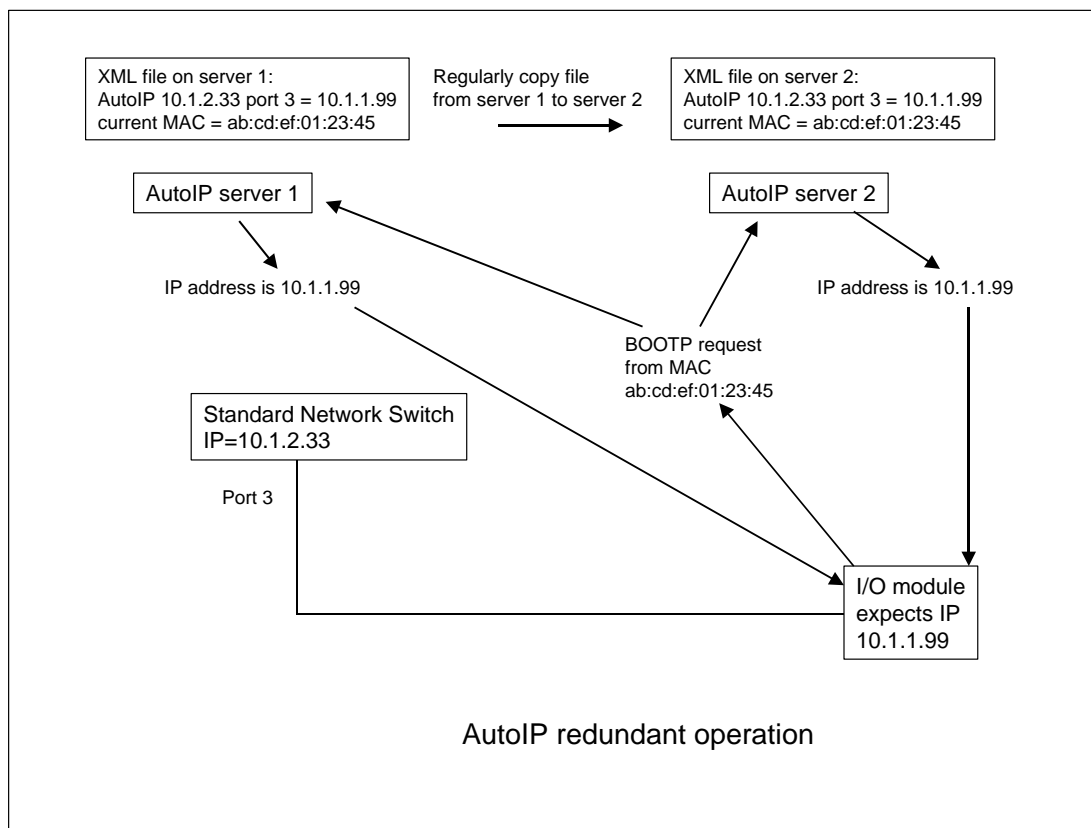


The 'one and only one' test is very important, because what it allows is for multiple Auto-IP – configured devices to be assigned to the same port of a managed switch. So long as all of the devices are normally operable (as would be the case in an industrial I/O situation), a single device which has failed can be replaced automatically. (The ASI device level bus has a similar facility for automatic determination of addresses, also requiring that all devices other than the one being replaced be operable.)

So it is no longer necessary to insist that all devices requiring IP assignment be directly connected to a managed switch. This can be a significant help in environments where an unmanaged switch or hub is more appropriate for embedding in equipment.

Determining whether a device is up or down is performed by a continuous 'ping scan' which attempts to communicate with each device under Auto-IP control every 30 seconds or so. Failure to respond to more than one such ping in sequence is considered failure. So devices can be replaced and complete the Auto-IP assignment process within 60 seconds, which is faster than most technicians can replace a device!

Since Auto-IP requires only that the replaced devices support BOOTP, and that managed switches support RFC1493, it can be easily retrofitted to almost all existing Ethernet – based control networks.



Auto-IP implementation

The Auto-IP product takes the form of a software package running on a Windows 2000 or XP computer, and which uses a web browser as its user interface for configuration and monitoring.

The Auto-IP service generates a very low 'load' in terms of memory and CPU usage, and almost always can coexist with a SCADA system or similar on the same computer.

The user interface allows an administrator to designate an Auto-IP device by assigning the following data

- IP address
- Netmask and Gateway
- Switch address
- Port number of switch
- SNMP 'community' (password) of switch

The data is entered and displayed in a 'spreadsheet – like' format which takes advantage of the likely similarity of settings for all devices on the same network.

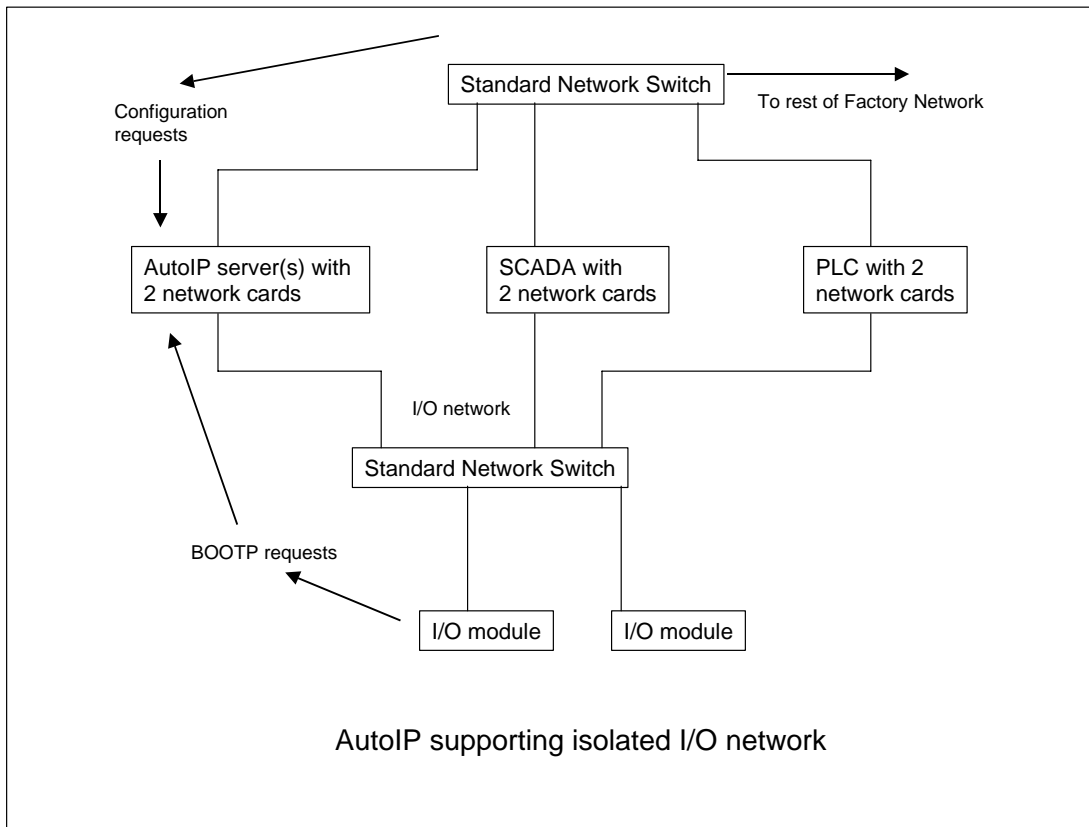
There is also a facility allowing automatic data entry using Network Vision's graphical 'IntraVUE' package, where the administrator can simply designate an existing station for Auto-IP assignment by clicking on the node and setting a checkbox in its properties. This again minimizes the possibility of misconfiguration.

There is absolutely no problem in arranging redundancy on the IP assignment mechanism. Because the software uses BOOTP and not DHCP, there is no concept of a 'pool' of addresses whose leases must be coordinated. So long

as multiple Auto-IP servers are operating from the same database, they will issue the same IP address data to any client requesting it. It is often wise to arrange at least two servers to cover the same plant area, one being close to the targets, and one being close to the 'center'. That way, if there is any sort of network disturbance during the time that BOOTP service is required (such as one of the switches going through a reset), there will be less delay at the target.

Auto-IP uses simple text files structured in XML format as its database, and normal file mirroring techniques can be used to ensure that all servers responsible for a given plant area remain synchronized.

Use on isolated I/O networks



In many cases, customers deliberately separate the Ethernet networks used for I/O from those used for general-purpose supervision. This may be for security reasons, or simply to allow an OEM to pre-configure and test a tool before delivery, and not have to alter any IP addresses or similar on installation.

These arrangements can be a challenge for server-based schemes such as Option 82 or distributed DHCP, where it is important to be able to access central servers containing configuration data for the switches or the targets.

In the case of Auto-IP, the simplest solution is to arrange for the Auto-IP workstation to have twin Ethernet interfaces, so that it can 'see' the I/O and enterprise network concurrently. Auto-IP will communicate correctly with the switches and targets on the I/O network, and still be able to service web requests and file mirroring on the main network. And users requiring redundancy can arrange two such computers if necessary.

For details, contact
info@intravue.net
or visit our website,
www.auto-ip.com

IntraVUE™ is a product of Network Vision, Inc.
Newburyport, MA, USA 1.978.499.7800

Copyright 2004, Network Vision, Inc.