NORTEL NETWORKS™

Pre-Sales Engineering
Tips & Tricks

Volume 2, August 2003

Did You Know...

... there are specific equipping rules when provisioning unprotected DS1 services on the OPTera Metro 3500, using DS1 Mappers (not via a DSM)?

1. DS1 Mapper circuit packs should be inserted in the shelf from left to right starting in slot 4.

2. Protection Switch Controller (PSC) must occupy slot 2. Although the name is misleading the PSC is required because it is required for alarms, maintenance and provisioning.

Editor's Note:
This publication will evolve based on your content and information requirements, therefore please feel free to provide feedback to: pnf@nortelnetworks.com.

IN THIS ISSUE

IN THE SPOTLIGHT

i2002/i2004 DHCP w/ Auto VLAN Discovery

How exactly does DHCP with Auto VLAN discovery work with our IP phones? First of all, i2002/i2004 phones may be provisioned manually or may be configured to use DHCP. If provisioned manually, the user must provide the IP address, subnet mask and default route, as well as the IP address and UDP port of the IP-enabled PBX. With DHCP enabled, one has the choice of "partial" or "full" DHCP support. With "partial" DHCP support, the IP phone is dynamically assigned its own address information but is not assigned the address/port of the PBX. With "full" DHCP support, in addition to that which is provided with partial support, the address/port of a primary and, optionally, an alternate PBX, may also be dynamically provided. These VoIP specific parameters may be passed via DHCP private options 128, 144, 157 or 191.

So what about the VLAN? It only needs to be used if the phone and PC are to be connected to the same L2 switch port via a three port switch and the phone and PC need to be on separate subnets. First, the phone does not need to be in a different subnet than the PC connected to it, in which case, tagging is not enabled and no VLAN Id is configured on the phone. If the phone is connected to a L2 switch port which has VLAN tagging enabled, the VLAN Id may be manually provisioned at the phone or may be dynamically assigned by the DHCP server. It is important to remember BayStack release 2.5.x and future releases added the ability to support untagged and tagged traffic simultaneously on the same port (per VLAN tagging option). This allows traffic destined to the PC to be untagged while the traffic to and from the phone would be tagged. It is not mandatory to implement this feature but it allows complete VLAN traffic separation for end user and VoIP traffic once the VLAN ID membership is determined on the IP phones (Dynamically or manually).

So let's assume the phone and PC are connected to the same layer two switch port, the phone and PC need to be in different subnets, and "full" DHCP support and Auto VLAN discovery. Here's how it works: At initial boot, the phone broadcasts (layer 2 and layer 3 broadcast) a DHCP Discover message looking for a DHCP server. Since a VLAN Id has yet to be assigned, the initial Discover packet is sent untagged. Assuming the phone is connected to a layer 2 access switch, the switch will add the 802.1Q header, tagging the packet with the default VLAN Id, and forward the frame. The first layer 3-aware device must be configured for DHCP relay. (Alternatively a DHCP server could be installed on every subnet but this is not practical) Having received the DHCP Discover broadcast, it will forward the Discover packet to the DHCP server (address of the server must be manually provisioned on the L3 device), adding to the Discover packet the IP address of its ingress port corresponding to the default VLAN. It is this address which informs the DHCP server to which IP subnet to assign the phone. The server, having received the Discover packet, will respond with a DHCP OFFER message, assigning an IP address consistent with the default VLAN. In an Auto VLAN environment, this will be a temporary address. The server should also include in the OFFER message a list of the applicable VLAN Ids from which the IP phone can choose. The VLAN Ids are passed using one of the DHCP private options supported by the i2002/4 phones (options 128, 144, 157 or 191). Up to 10 VLANs may be passed to the client. Upon receiving the DHCP Offer with a list of suggested VLANs, the IP phone should accept the offered IP address (a temporary address) and proceed with the DHCP negotiations (i.e., a DHCP Request sent to the server for the assigned address and a DHCP ACK sent by the server to finalize the assignment). With no VLAN tagging, or with a manually provisioned VLAN Id, the DHCP process would be complete at this stage, but with Auto VLAN discovery the process continues. The IP phone will release the temporary IP address by sending a DHCP RELEASE to the server. Note that all packets sent by the IP phone to this point have been untagged. The phone will then choose the first VLAN Id included in the previous DHCP OFFER. Having released the address, the phone now restarts the discovery process, broadcasting another DHCP Discover message, but this time tagging all packets with the selected VLAN Id. The DHCP Relay agent, its ingress port presumably having an IP address on each of the offered VLANs, will update the DISCOVER packet with its IP address corresponding to the selected VLAN. Once again, it is the gateway routers IP address which informs the DHCP server to which IP subnet to assign the phone. The DHCP OFFER, DHCP Request and DHCP Ack continue as before, only now all packets are tagged with a specific VLAN Id.

Note that if the DHCP server does not have an available address for the requested VLAN and does not respond to the DHCP Discover, the IP phone will resend the Discover broadcast, retrying up to 4 times and doubling its wait time between each successive attempt (approx. 4, 8, 15, 30 seconds). If the server still does not respond, the IP phone will give up on that VLAN and broadcast a new DHCP Discover using the next VLAN. As mentioned previously, up to 10 VLAN Ids can be included in the DHCP Offer.

Assuming that the DHCP server can service an address on one of the VLANs, it may also provide the VoIP specific parameters required for the phone to register to the Terminal Proxy Server (TPS) of the PBX. As with the VLAN Ids, the DHCP server will pass the VoIP-specific parameters in the DHCP Offer using any of the i2002/4 supported private DHCP options. These include the IP address and UDP port of the primary, and optionally an alternate, IP-enabled PBX. The phone finalizes the boot process by registering with the TPS.

Enter text string "VLAN-A=<voice VLAN ID 1>+<voice VLAN ID 2>;" into the "data subnet's" DHCP server as a Nortel vender specific option. (VLAN-A=<voice VLAN ID 1>;" for a single voice vlan) Enter the i200x full dhcp string into the "voice subnet's" DHCP server as a Nortel vendor specific option.

INTEROPERABILITY

Interworking Between Q-tagged and Untagged Frames

Network switches these days are designed with advanced and flexible hardware ASICs, which allow for a plug-and-play environment when it comes to Q-tagged verses untagged interfaces. For example, when connecting a BPS or one of the Stackable switches to a Passport 8600, one end of the link can be configured as a Q-tagged port and the other end of the link can be configured as an untagged port, and the link will work fine as long as the VLANs are properly defined.

With some of the older technology however, both ends of the link must configured the same (i.e. both ports must be configured either as Q-tagged, or both ports must be configured an untagged ports so that the switch can trigger on the appropriate Ethernet header for proper operation).

In other cases, such as with the Alteon products and other security products, it is important to understand that for security purposes, these switches require that both ends of the connection be configured the same (i.e. either both Q-tagged or both untagged, in order to work).

As a general design rule, the engineering recommendation is to have identical tagging configuration at both ends of the link whether it is required or not in order to maintain network consistency and simplicity for troubleshooting.

The Details on How SFFD Works

The Single Fiber Fault Detection (SFFD) feature on the Passport 8600 and the BPS/Stackable products provides an alternative to the Remote Fault Indication (RFI) feature.

SFFD is a proprietary implementation; it is available in release 3.5 on the Passport 8600, and in release BoSS 3.0 on the BPS and the Stackable switches. Since the intent of this feature is to provide an alternative to RFI, the feature is applicable only to GigE fiber ports.

At this time, the SFFD feature can only be enabled and disabled via CLI. SFFD is enabled/disabled on a port-by-port basis; by default it is disabled on all ports. The CLI command on the Passport 8600 is "config ethernet <slot#/port#> sffd [enable/disable].

Note 1: If you try to enable SFFD on a copper port on the Passport 8600, the command will be accepted and no error messages will be displayed; however the command will not do anything. On the BPS and the stackable switches, the CLI commands, in sequence, are: "enable", "config terminal", "interface fastethernet <switch#/port#>", "sffd [enable/disable]".

Note2: If you try to enable SFFD on a copper port on the BSP or on one of the stackable switches, the command will be rejected, and an error message will be displayed.

SFFD sends several types of control packets as part of the link failure detection mechanism; these messages include heart-beat packets, link-presence

request packets, and link-presence response packets. The SFFD control packets use a destination MAC address of 0x010081 and a PDU-type of 0x7001, 0x7002, and 0x7003. It's important to ensure that these packets are not dropped by filters which are applied to ports with SFFD enabled.

The timer values and the counter values for the SFFD control packets are not user-configurable; the feature is simply either turned on or off by the user. The behavior and the time it takes for a link to be declared out of service and to recover after a fiber link failure/recovery are shown in the example below using a BayStack 470 and a Passport 8600.

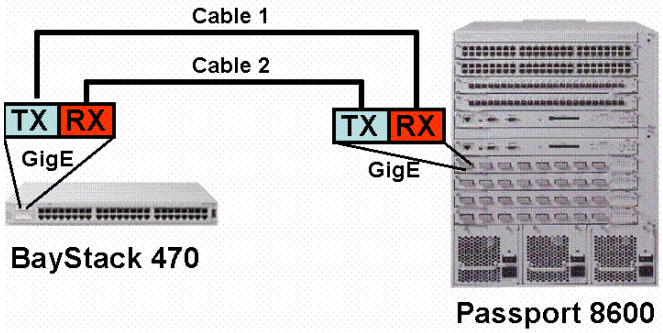

In the above diagram, when Cable 1 is physically disconnected to simulate a single-fiber fault, the following observations are made:

1. On the Passport 8600, the port goes down immediately because the laser signal is lost.
2. On the BS470, the operational link status will toggle up and down in 6 to 7 second intervals for about 50 seconds, then will go down permanently until the link is restored.
3. Note that although the link is operationally down after 50 seconds, the link light on the BS470 will still flash on and off in 5 to 6 second intervals as the switch tries to send SFFD control packets to the Passport 8600.

When Cable 1 is reconnected, the following observations are made:

1. The link goes into operational state about 18 seconds after the cable is restored.
2. Note that the link operational status comes up at the same time (after about 18 seconds) on both the BS470 and the Passport 8600.

When Cable 2 is physically disconnected to simulate a single-fiber fault on the other side of the link, the following observations are made:

1. On the BS470, the port goes down immediately because the laser signal is lost.
2. On the Passport 8600, the link will go out of service in about 10 seconds.
3. Note that although the Passport 8600 is out of service, the link light on the card will flash green every 5 or 6 seconds as the port tries to send SFFD control packets to the BS470.

When Cable 2 is reconnected to simulate a restore of the link, the following observations are made:

1. The link on the BPS recovers immediately.
2. The link on the Passport 8600 takes about 10 seconds to recover.

IP-PBXs / PBXs / KEY SYSTEMS

Meridian 1 Programming Tip – Removing Programmed Hosts

Ever programmed a new host in the Meridian 1/Succssion 1000 PBX switch using LD 117 and then later on wish to remove the host from the host list? There is a simple set of two steps to remove the host from the host list.

For example, one has programmed a host for a Succession Media Gateway 1 (SMG1) with the following information,

- LD 117
- New host CSEN_SMG1 10.22.2.71 1 (create a new host in SMG1 cabinet 1)
- Chg ELNK active CSEN_SMG1 (activate the ELNK and put CSEN_SMG1 into active mode)
- Chg MASK 255.255.255.0 1 (ensure that the mask is proper and using 24 bits)
- New route 0.0.0.0 10.22.2.1 1 (ensure a route is created in cabinet 1),
- Update DBS

The above steps will create a new host and the put it in active mode.

- LD 117
- Out route 1 1 (take out the default route 1 in cabinet 1 so that there is no link between the host and the route),
- Chg ELNK active CSEN_SMG1 1 (issue this command to "turn off" the ELNK linked to the host in cabinet 1),
- Update DBS

The switch will then display a prompt telling you that there is no host configured for the SMG 1. This will have removed the host from the host list.

MULTIMEDIA COMMUNICATION SERVER (MCS)

MCS 5100– Server Patches

When upgrading a MCS 5100 system it is very important to read the release notes for any patches necessary. The most current release notes (FP1 1.1.12 build 470) make a reference to three patches for the MCS 5100 on the Sun VT100 platform:

- Solaris Kernel patch
- SNMP patch
- MED patch

The kernel patch will update the kernel on the server. The SNMP patch fixes a problem with the snmpd process. The MED patch will allow the simultaneous booting of the servers.

To verify if you need the kernel patch, log on to a server as root user. Run the command ***showrev***

At the end of the response there will be a line that reads:

Kernel version: SunOS 5.8 Generic 108528-18 November 2002. Look at the number after the dash if it does not read the number eighteen you will need to apply the patch. Make sure you download the correct file specified by the release notes.

These patches must be applied to each MCS 5100 server. The kernel patch takes about two hours to complete. The other two patches together much less time, in less than half an hour. Instructions on installing the patches are in the release notes.

The location to find these patches is (FTP server):

47.104.23.93

user = 1.1access

pwd = 1.1access

The kernel patch is in a folder called *solarispatches*. The MED patch is in one called *ntme*. The SNMP patch is in one called *snmpd*

Checking for required patches at each release of software will ensure the MCS 5100 is at a baseline configuration and will address known code issues.

NETWORK MANAGEMENT

Deleting Optivity NMS Databases

In a lab and demo environment where network configurations are constantly changing, it is useful to be able to clean out the Optivity NMS database so that you can do a 'clean' discovery when required. In Optivity NMS 10.1 there is a quick method of accomplishing this. Open up a console session on your NMS workstation and proceed as follows. The password required by the rm_appdb command is 'nms'.

- C:**rm_appdb smop**
- Password: ***nms***
- Are you sure you want to delete all elements in smop? **y**
- Removed 213 element(s).
- Removed 14 segment(s).
- C:\>**optivity_apps stop**
  (messages will appear for each stopped process)
- C:\>**optivity_apps start**
  (messages will appear for each started process)

ROUTERS / SWITCHES

Alteon Application Switches- Cookie Not Re-written

The Alteon Application and Web switches support the following modes for cookie operation.

1) Passive mode
2) Re-write mode
3) Insert mode.

In the re-write mode a cookie with a place holder value is set on the real server and the switch is configured to same value. When this cookie is seen by the switch it is re-written with a hash value of RIP/VIP and sent to the client. By default the switch is programmed to look for the place holder cookie, from the server, only at the first server response. If, for some reason, this cookie gets returned later in the session the switch will not re-write it. This can happen when HTTP/1.1 is used which is capable of returning multiple responses within the same session.

To work around this issue one has to ensure that

1) If using HTTP/1.1 this cookie is the returned in the first response.
2) Set the value of rport (/c/slb/virt x/service 80/ recount x) to a value when the cookie is returned by the real server.
3) Use HTTP/1.0

SECURITY / VPN

Alteon SSL Accelerator String-Based URL Load Balancing

To configure String-Based URL load balancing and to redirect all traffic for which there is an exact match to one set of backend servers and all remaining traffic to set of default server, we need to use a wildcard character "*". Right now, under SSL there is no "best match" concept and "*" will match on any string, including strings that are explicitly defined. To exclude explicitly defined strings from being matched by the "*", we need to use negative numbers when listing strings. In addition, load balancing option for default server, "lbop" needs to be set to "all" to match all specified strings.

Sting-Based URL LB example:

```
/cfg/ssl/server 1/adv/string 1
        match *.html
        location url
/cfg/ssl/server 1/adv/string 2
        match *.gif
        location url
/cfg/ssl/server 1/adv/string 3
        match *
        location url

/cfg/ssl/server 1/adv/loadbalancing
        type string
/cfg/ssl/server 1/adv/loadbalancing/backend 1
        lbstrings 1,2
        lbop any
/cfg/ssl/server 1/adv/loadbalancing/backend 2
        lbstrings 1,2
        lbop any
/cfg/ssl/server 1/adv/loadbalancing/backend 3
        lbstrings 3,-1,-2
        lbop all
```