



## DoubleVision and Internet Failover

### Introduction and Description:

The Net Integrator product line ships with a standard feature called DoubleVision. This product name is a bit misleading, since in reality it refers to two different configurations, one that provides high availability, and the other that adds bandwidth sharing.


The high-availability product (called **Internet Failover**) allows a Net Integrator to keep an Internet connection, even if one of its Internet links fails. This feature was originally added to provide to our customers a way for them to work around the chronically unstable DSL providers in the United States. In the original version, two or more PPPoE connections were supported, with the assumption that each were connected to different ISPs, which meant that a failure at one ISP would not lead to a complete loss of Internet connectivity for a company relying on a Net Integrator. Later, support for DHCP and static route failover was added, the latter being only available in our upcoming 3.71 Net Integrator OS (NIOS) release.

The bandwidth-sharing product (properly called **DoubleVision**) allows a company to use the bandwidth of Internet connections that are used for Internet Failover. It works by analyzing each outgoing connection for the route to the destination with the least latency (which is not necessarily the route with the highest bandwidth). It then uses that route for that particular connection to the destination. Since certain ISPs are usually better than others, this frequently leads to a majority of traffic always using a single connection, with the other being used sporadically or not at all. This is expected and normal behaviour.


### How to know if you are using Internet Failover:

You can set up multiple links in order of priority by setting some to *"Enable: As last resort"* instead of *"Enable: yes"* in the Dial-Up section of WebConfig. These links only get activated when the primary link(s) is/are marked "broken" by Net Intelligence.

You can see that you're using Internet Failover by checking the "<number>.<letter>." code next to your various Internet links on the status page of WebConfig. The different numbers imply the different backup priorities. For example,

Internet Connection:		"1.a. Indirect on eth1" "1.b. Indirect on eth2" "2.a. Modem"
----------------------	---	--

means that your primary links (1.x) are the first two and your secondary link (2.x) is the last (modem) link.




# NET INTEGRATION TECHNOLOGIES INC.


## Research and Development Centre

### How to know if you are using DoubleVision:

In the above example, if you have "1.a. Indirect on eth1", "1.b. Indirect on eth2", and "2.a. Modem", then if all links are non-broken, DoubleVision will split outgoing web traffic between the two "1.x." links. A couple of examples should make this clear:

Internet Connection:		"1.a. Indirect on eth1" "2.a. Indirect on eth2" "3.a. Modem"
----------------------	---	--

*Connections only using Internet Failover*

Internet Connection:		"1.a. Indirect on eth1" "1.b. Indirect on eth2" "2.a. Modem"
----------------------	---	--

*Connections using DoubleVision*

#### Note:


Some types of outgoing traffic either cannot be or should not be shared in this way, for example, FTP, ping, traceroute, and SMTP. Most of these protocols (FTP, ping, traceroute) use multiple TCP sessions for one "logical" session. SMTP is special because of spam-relay protection, which forces you to use a *different* outgoing SMTP server depending on which link is in use. To avoid these problems, we *only* use DoubleVision for outgoing web sessions; for other kinds of sessions, Net Intelligence chooses the best link as a "default" link and uses that for all outgoing non-HTTP traffic. In practice, this isn't much of a problem, since almost all high-bandwidth traffic comes from the web.

Incoming traffic is treated very differently from outgoing traffic: we accept connections on all non-broken links, but the DNS for your domain name is only registered to point at the "default" link chosen by Net Intelligence. This is because you can't actually tell client software to "use the best link" or "alternate between these two links" in a reasonable way, so we have to choose the best one and tell them to use that. Occasionally, the DNS-advertised "best" link starts to get too loaded down (probably because all the incoming traffic is using it!), so Net Intelligence decides to advertise the second-best link for a while instead. Of course, remote users may have a DNS cache of 5 minutes or more, so this change doesn't take effect immediately.

#### Limitations:

With either DoubleVision or Internet Failover, unplugging any link should cause Net Intelligence to switch you over to a different, working one. If it doesn't, something is misconfigured or you've encountered one of the following limitations:

- Your DNS server always publishes the address of its "favourite" non-broken, high-



## **NET INTEGRATION TECHNOLOGIES INC.**

### **Research and Development Centre**

priority link. As a result, incoming traffic generally comes in on that address. Cached DNS users will continue to send to a “broken” link until their DNS cache times out. When testing, either set your DNS refresh constantly, or use the IP address.

- Incoming traffic is always *accepted* at the address of any non-broken link, even if DNS currently gives users no way of actually getting there.
- Outgoing web traffic always goes through all highest-priority DoubleVision links.
- Outgoing non-web traffic always goes through Net Intelligence's favourite highest-priority link.

#### **Code limitations that may be removed in some subsequent version:**

- There is no way to force an outgoing connection to use a particular link.
- You can't set the priority level of non-PPP links, i.e., cable modems are always at a fixed priority level, but you can change your PPPoE ADSL modem's priority if you want it to have a higher or lower priority than the cable. (To be fixed in NIOS 3.73.)
- Currently TunnelVision, like all non-web traffic, just uses Net Intelligence's "favourite" route. Eventually, it will support using all the highest-priority DoubleVision routes. This will likely be included in NIOS version 3.80.
- Net Intelligence's "favourite route" selection algorithm, while not bad, could be better. It should choose a different favourite route ASAP if the current one starts getting overloaded—right now it waits too long.
- Sometimes, transmitting a lot of data dramatically slows down receiving; this is counterintuitive, but it's a general problem for ADSL ("Asymmetric" DSL) and cable modem links, where uplink speed is much slower than downlink speed. This tends to confuse Net Intelligence's link selection, but it will be fixed nicely in 3.73 with the "active queue management" feature.

#### **How it all works:**

We detect broken links using an excellent (but not quite fool-proof) method called **Demi-Ping**. It detects most kinds of link failures to the Internet, although certain kinds of partial failures (eg. we can reach the Internet backbone but can't reach some countries) can't be detected. Net Intelligence should always notice if you unplug the physical connection to a link and automatically switch to your secondary link(s)—and this is the easiest way to check that it's working. Due to the fact that we are keying off packet timeouts, there will be a small lag in failing over to the other link. This is intentional and is meant to avoid us failing over unnecessarily when an upstream router loses a packet. We wait for a couple (the exact number varies between 2 and 5, depending on what Net Intelligence decides to do) of consecutive timeouts before we will fail over, and so the delay can be up to 45 seconds.



## **NET INTEGRATION TECHNOLOGIES INC.**

### ***Research and Development Centre***

The DNS server (including DynamicDNS) will publish one of the IP addresses for the most important non-broken link as the IP address for your domain. That is, if a #1 link is non-broken, then we'll publish its address; if all #1 links are broken, then we'll publish a #2 address.

Because incoming connections are usually addressed to your domain name, whichever IP your DNS is publishing is the one to receive most incoming traffic. However, if there's more than one non-broken link, any of those should be able to receive incoming traffic if the source knows the link's IP address; for example, you should be able to ping the link through its IP address.

### **Conclusion**

As indicated in the introduction, DoubleVision's load sharing works differently from typical load-balancing routers. It takes each individual session (such as a single web page) and assigns it to one Internet link or another. All packets for that session go through the same link. This is unlike the usual load balancing routers that split packets randomly across links—even packets belonging to the same session. This means two things:

- 1) You don't need both links to be through the same co-operating ISP that can decode a single session from multiple links.
- 2) If you only have one session at a time or your sessions are unluckily assigned to links, you get little-to-no performance improvement.

We believe that the first item in the list above greatly outweighs the limitation implied by the second item. DoubleVision brings highly available Internet access to a segment of the business community which cannot otherwise afford it. Since there are no OSPF or BGP routing tables to manage, or special agreements with an ISP to negotiate, and since it involves no extra cost or effort above those of a standard Net Integrator, we believe that, despite its limitations, it offers excellent value to anyone implementing a DoubleVision solution.

Authors: Avery Pennarun and Patrick Patterson

Copyright © 2003 Net Integration Technologies Inc.  
All rights reserved.