

STONESOFT

Whitepaper

Coping With ISP Failure

Table of Contents

Executive Overview	1
Few businesses have ISP failure contingency plans	1
ISP-level solutions	2
External load-balancing	2
A new alternative	2

Executive Overview

The general malaise and recession in the telecommunications sector seems to claim new victims every month. Recent revelations about WorldCom and Qwest show that even the largest suppliers are not immune. And as we all know when the mighty fall, they fall hard. Add to this the 'everyday' failures of connectivity including DNS attacks, viruses and other miscellaneous 'outages' that leave users without Internet connectivity for hours or even days and we see an increasingly uncertain environment.

In research carried out by Stonesoft 94% of respondents stated that 100% Internet availability was important.

44% of respondents had more than one ISP connection, only 32% used a connection from a different provider and 37% stated that when their ISP went down they simply waited for the ISP to fix the problem!

Amongst the debris of these collapses are the thousands of companies that relied on these providers to connect them to their suppliers, customers and partners via the Internet. Although many may have warning of the demise of their supplier and the time to arrange transfer to an alternative provider, many companies do go bust virtually overnight leaving their customers in the lurch.

Few businesses have ISP failure contingency plans

In this context many businesses should be examining their contingency plans to ensure that they can and will continue to operate in the event of a permanent or temporary halt to Internet service provision. Unfortunately, research shows that few are doing this and that many are not aware of the various solutions that can provide ISP fail-over that will allow business to continue as normal.

The small business or SME sector is perhaps the most exposed. Those companies without the resources to demand high level service agreements, or to employ dedicated in-house teams are at most risk from ISP failure, whatever its cause.

In a survey of our customer market* we found that 94 percent of respondents stated that 100% Coping With ISP Failure availability of Internet connectivity was important or very important – highlighting the impact of the 'New Economy' even following the crash of the dotcoms. Worryingly, very few of these companies had made contingency plans to cover the loss of service from their ISP either on a temporary or permanent basis. Only 44 percent of respondents had more than one connection and less than one sixth overall had a connection from an alternative supplier, leaving the vast majority vulnerable to ISP failure. Of these only a minority had automatic fail-over procedures established. There are a number of options that businesses can consider when their ISP goes down:

- Wait until the ISP fixes the fault themselves
- Deploy in-house resources to find and fix the fault
- Manually redirect to an alternative service
- Automatically re-direct

Amazingly our survey showed that 37 percent simply waited for the ISP to fix the fault itself. Meanwhile their businesses were losing money and customers. In cases of serious faults and permanent losses of service through bankruptcies this is clearly not an option.

Manual redirection of the IP address space requires time and is quite complex demanding a high-level of knowledge of the network architecture and the relationship with the ISP, locating and fixing faults also requires highly skilled (and therefore expensive) engineers. Only 16 percent of our survey used this method.

ISP-level solutions

When we look at automatic redirection the options initially look limited. The established solution for automatic redirection is Cisco's BGP4 routing server which, although effective in some situations, is a very expensive and complex solution. BGP or Border Gateway Protocol was designed as an ISP-level solution and as such is not designed to be implemented by end-users. It requires a great deal of technical expertise from specialist engineers throughout a multi-step process involving several activities that fall well beyond the normal bounds of software configuration.

For example, the implementation team must negotiate agreements between rival ISPs, acquire and configure sophisticated hardware and routing schemes, and must possess advanced BGP programming expertise. Moreover, businesses using this solution also need to own a Class B IP address – something that is not only beyond all but the largest companies, but which brings its own security and implementation risks.

BGP works by selecting packet routes from all available ISPs and thus sharing the load between them. Thus, if an ISP becomes unavailable for any reason packets are simply routed via one or more other established routes using other ISPs. As such it does provide high availability for outbound packets, but it cannot manage the re-routing of inbound packets, potentially a major issue for online businesses. Furthermore, BGP does not measure the performance of ISPs – it sees them as on or off. Thus it can slow-down performance and throughput in normal situations by sending packets via congested routes.

External load-balancing

An alternative option is to external load-balancers that test and monitor several connections on a continuous basis. As a true load-balancing solution these can fail-over 'dead' connections automatically. However, there are performance issues and this solution requires additional hardware beyond the firewall, such as a pair of firewalls and a pair of load balancers (for achieving HA on the load balancers), which adds to the cost of implementation.

External ISP load balancing equipment requires constant supervision, administration, and system updates, adding to maintenance costs. Administrators must also ensure the separate configurations of the firewall and the load balancing box are consistent, adding to the technical complexity of the management process. Lastly, external load balancers cannot effectively handle VPN traffic because load balancers encapsulate VPN traffic when transferring data, which affects data throughput, slowing performance.

A new alternative

However, a new alternative does now exist. Multi-link technology combines the capabilities of external loadbalancers with the firewall significantly easing implementation and maintenance.

Multi-Link determines the fastest ISP as part of the standard three-way TCP handshake without delaying network traffic. The outbound connection begins when the client sends the server a SYN packet that reaches the Multi-Link gateway, and is replicated to connected ISPs. Outgoing SYN packet source IP addresses are modified through NAT and identify the ISPs they are sent through. SYN packets reach the server, which responds (SYN-ACK) to the gateway from the same ISP. This process measures the RoundTrip Time (RTT). The ISP that delivers the first SYN-ACK response has the shortest measured RTT and is chosen for the connection.

Non-functioning ISPs (such as those disabled by a Denial of Service attack) are therefore automatically excluded and data packets are sent via the remaining routes – thus achieving automatic fail-over without complex re-engineering or negotiations with multiple ISPs. The flexibility of this solution even allows for connectivity to be maintained over a bank of ADSL connections thus providing a highly cost effective alternative to redundant leased lines or ISDN.

When an ISP fails Multi-Link sends a dynamic DNS update to the name server, which removes all IP addresses in the ranges of the failed ISP. This not only ensures that packets do not use this route, but also reduces the normal wait time by about twenty (20) seconds per potential connection. There are fewer timeouts because inbound packets reach functional addresses without waiting for non-functional addresses. Users should, therefore, experience no drop-off in response. If a previously failed ISP is reported functional, then the IP addresses for that ISP are reactivated on the name server.

Multi-Link also load-balances inbound connections to optimise high availability so customers can reach your site without delays or disconnects. Multi-Link accomplishes this by using special IP addresses for servers, which can be assigned multiple IP addresses from the IP address ranges of multiple ISPs. Return packets are then managed by way of policy routing and NAT.

When an external user requests the Web server's IP address, the DNS server returns all external IP addresses. The client software then chooses one address that dictates the ISP used for the connection. Typical clients can use the other IP addresses if the first one chosen is unreachable. Once the inbound connection reaches the firewall, the firewall translates (NATs) the external IP address to the internal Web server destination. Policy routing is used to route return packets through the same ISP that the initial SYN packet came through. Thus once again full redundancy is achieved protecting the business from the failure of an ISP.

So, to conclude, with so many businesses reliant on Internet connectivity for the survival of their businesses there is an urgent need for them to investigate ISP fail-over systems. Solutions now exist to match the requirements and budgets of virtually all customers. Even without the current difficulties in the telecommunications sector it would be wise to examine your organisation's contingency arrangements for ISP outages – it could be the difference between business survival and failure.

*Research carried out on behalf of Stonesoft in Q1 2002. 2000 security and network managers were polled by telephone.

STONESOFT

Stonesoft Corp.

Itälahdenkatu 22 A
00210 Helsinki
Finland
tel. +358 9 476 711
fax. +358 9 476 712 34

Stonesoft Inc.

1050 Crown Pointe Parkway
Suite 900
Atlanta, GA 30338, USA
tel. +1 770 6681 125
fax. +1 770 6681 131

Copyright 2006 Stonesoft Corp. All rights reserved. Registered or unregistered trademarks in this document are property of their respective owners. The products described in this document are protected by one or more of U.S. patents and European patents: U.S. Patent No. 6,650,621, European Patents No. 1065844, 1289202, and may be protected by other U.S. patents, foreign patents, or pending applications. Specifications subject to change without notice.