NetworkCare
The knowledge behind the network

Lucent Technologies
Bell Labs Innovations



Minimizing E-Business Downtime: Considerations in ISP Multi-homing

By Alex S. Latchaw

Overview

With the demands of e-commerce, virtual private networks, and other Web services increasing at an exponential rate, the stability of Internet connections has become an area of concern for many organizations. Since downtime in these key applications is unacceptable, organizations are demanding strict service level agreements (SLAs) from their Internet service providers (ISPs). However, downtime due to weather, accidents, or equipment malfunctions outside the control of the provider can cause non-compliance with the SLA, placing the provider in a liable and potentially costly situation.

Many Tier 1 providers are presently offering 100% uptime SLAs with their service. These ensure that the service provider's clients are reimbursed for any loss of connectivity to the provider. Most SLAs do not cover an organization in the event of the provider's loss of access to other networks, or in the event of a loss of proper route aggregation outside the provider's network. While these SLAs remain unbreeched, the end client can still experience outages.

The only options an organization has are to passively rely on one provider and accept network outages as they occur, thus completely outsourcing its e-commerce environment, or to proactively reduce the likelihood of outages by multi-homing to several ISPs. Migration from a single to multi-homed environment presents several challenges and requires significant investment in equipment, leased services, and engineering man-hours.

Choosing ISPs

ISPs vary greatly. In choosing an appropriate provider, an organization must consider several variables, the importance of which vary from organization to organization. At the top of most lists are cost, SLAs, and the structure and design of the provider's network.

Cost

Costs accompanying Internet service can be separated into two categories: recurring and non-recurring. Recurring costs are the month-by-month or yearly contractual service charges. Non-recurring costs represent the initial capital required for equipment and setup.

Recurring costs are generally what a provider quotes when a particular inquiry is made. The monthly cost of a T1 connection for $1000, for example, can be a misleading figure. Generally, a service provider is separate from an organization's local exchange carrier and therefore does not include the leasing of the physical access loop circuit in its quoted cost. Depending on the geographic distance between an organization and the ISP's nearest point of presence, this cost can be considerable and must be negotiated

outside of the contract with the ISP. Other recurring costs will vary depending upon the services requested. For example, if an organization wishes to subscribe to Usenet feeds or wishes the provider to act as a primary or secondary DNS for its domain, subsequent monthly charges can be levied.

Non-recurring costs include the purchase of routers, CSU-DSUs, and installation charges for the physical circuit and the Internet service itself. The final design of an organization's access to the Internet may affect these initial startup costs and these changes must be weighed into the entire cost of the service. Many service providers will offer an inclusive package in which all of the equipment is leased and managed by the provider, and is simply allocated on the organization's premises. This option drastically reduces the non-recurring costs, but it introduces a higher monthly fee.

SLAs

SLAs are very important to most organizations when choosing a provider. Most providers have an open SLA for each of their customers, which is fully viewable on their Web site before an inquiry for service is made. These documents should be scrutinized carefully because many will offer high availability to the provider's network, not to the Internet as a whole. The more responsibility an organization assigns to its provider, the greater the need for an appropriate and individual SLA.

Network Structure

ISPs differ greatly in the geographic expanse of their networks and in their connectivity to other providers. There are many small ISPs that are able to provide their clients with high-speed service to their network at a low cost. Unfortunately, most of these ISPs typically over-subscribe their own circuits, which can create a bottleneck for organizations utilizing their services.

The "footprint" or physical range of a provider's network can affect many key applications. Unavailability in different regions can cause undue network latency when hops through separate networks result in delays in inter-branch VPN traffic and e-mail. E-commerce can be greatly affected by these multiple hops, obstructing a user's ability to access an organization's website. A limited geographic range may also necessitate more administrative overhead for monitoring multiple SLAs from multiple providers, each requiring a separate fee.

Connectivity to other providers is crucial in determining the overall stability and reliability of a provider's network. Providers connect to each other in a hierarchical manner, referred to as tiers, to most effectively utilize the available bandwidth. Tier 1 providers generally have national or worldwide proprietary networks that connect multiple Tier 2 providers, which in turn sell this connectivity to end clients. Many providers offer network diagrams for further scrutiny upon request.

The availability of multiple points of presence (PoPs) is another factor to consider carefully. Most Tier 1 ISPs have access points into their network in many large cities. To create an additional level of redundancy, an organization can provide circuits to connect to separate access points within a provider's network.

Differing access points are not beneficial if local access loop paths do not differ as well. To create a fully redundant architecture, an organization must ensure that all paths between endpoints are unique and that no single point of failure is evident.

Internet Addressing Issues

Most organizations with a single Internet provider design utilize a block of addressing from the provider. This addressing is usually used as translational addressing at a firewall to allow access to internal resources. The service provider ensures route aggregation to all other networks. This hierarchical structure works well for single-homed sites, but once the decision to multi-home to providers is made, the Internet addressing scheme must be completely redesigned.

The Internet community is using a relatively new technology called Classless Inter-Domain Routing (CIDR) [i]. With CIDR, a particular block of addresses is grouped together and advertised as a single route. This drastically reduces the full Internet routing table to a manageable size. These blocks of addresses are labeled with a forward slash (/) and a number corresponding to the bits in the network portion of the subnet mask. Figure 1 shows examples of how CIDR is used.

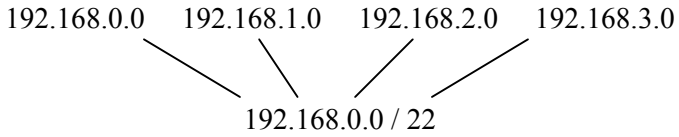

Figure 1

Due to the hierarchical nature of CIDR, one provider cannot advertise routes to the Internet for a portion of another provider's network. This would cause a discontinuation within a provider's advertised route that is unacceptable to most service providers. Therefore, once an organization decides to move to a multi-provider platform, it must relinquish its present addressing to be absorbed back into the original service provider it was allotted from. To take this step, an organization must acquire its own unique CIDR to advertise to the Internet.

IANA Registration

The Internet Assigned Numbers Authority (IANA)[ii] has full authority over all Internet Protocol number spaces utilized on the Internet. It allocates address space to Regional Internet Registries (RIRs) on an as-needed basis. The IANA has created a registry to cover each geographical area: APNIC for Asian-Pacific Region; RIPE NCC for Europe, the Middle East, and parts of Africa; and ARIN serving North and South America, the sub-Saharan Africa, and the Caribbean. In order to acquire address space, an organization must formally request an allotment or assignment of numbers from its Internet Registry.

The commonly accepted minimal CIDR block of address space needed to ensure route aggregation across the Internet is a /20 network, or 16 Contiguous Class Cs. Most organizations either do not own this much address space, or are utilizing some part of this block elsewhere. In this scenario, new address space must be acquired and used as the organization's Internet-advertised CIDR address block.

Addressing is allotted to organizations that plan to provide Internet connectivity for multiple other companies. An address range can be sub-netted as an organization sees fit and then allotted to a downstream entity to provide Internet access. Obviously the average allotment is much larger than a /20 network, since much more addressing is needed in order to ensure that all the organization's clients have adequate room for growth. Allotments are generally given to large service providers that have outgrown the addressing they were previously assigned.

Assigned addressing is primarily given to organizations that wish to home to several service providers, but do not plan to sell Internet connectivity to third parties. These assignments are restricted to a /20 or above, since any smaller CIDR block should be obtained from an upstream provider.

Due to the diminishing availability of address space, the IANA carefully investigates an organization's true requirements before granting an address assignment. The organization must provide significant justification based on its future need and past utilization of Internet space.

IANA is also the controlling entity for autonomous system (AS) numbers on the Internet. In order to run a dynamic routing protocol to a service provider, one must obtain a uniquely assigned AS number. The process for requesting these numbers is much the same as for requesting Internet addressing.

Border Gateway Protocol and the Autonomous System Concept

Once an organization has successfully acquired a unique CIDR from the IANA, it must ensure correct advertisement of this network to the Internet. The Border Gateway Protocol (BGP)[iii] is widely utilized for this purpose.

BGP utilizes an autonomous system (AS) to differentiate between separate organizations within a network. Each AS is seen by other ASs as the owner of the routes it aggregates. This type of peer-to-peer philosophy creates an open environment in which organizations can provide whatever services they wish to their clients.

Depending upon what traffic types are permitted, ASs can be categorized as follows:

- Stub AS: A stub AS retains a single connection to one other AS. This AS would only advertise its own CIDR to its peer, and would therefore carry only its own traffic.
- Transit AS: Transit ASs are designed to carry traffic between external peer ASs. These networks are commonly thought of as upstream service providers.
- Multi-homed AS: This AS would peer with multiple ASs but would only advertise its own CIDR. These networks could refuse to carry any external traffic.

BGP is a feature-rich protocol allowing for much customization. Load balancing can be configured between individual circuits or service providers on a best-route basis. BGP can be configured internally to provide redundancy in case of outages. In short, BGP's many features allow an organization to maintain control of network traffic to, from, and between connecting networks.

Next-Generation Internet Protocol

The next-generation Internet Protocol (IPv6)[iv] is a new standard in progress for the addressing of the future. Due to the reduction and eventual extinction of IPv4 address space, this new addressing scheme is being designed with a much greater initial address space. The current IPv6 design includes a rigid hierarchical structure that is not present in IPv4 addressing.

The hierarchy of IPv6 begins with the top-level service lroviders who are assigned top-level addresses for the sole purpose of propagating routes for their downstream customers. Each of these downstream entities is in turn assigned site-level addresses which are appended to each top-level address in order to differentiate between sites connected to high-level ISPs. Behind each site address are 24 bits of available host space for an organization to utilize as it sees fit.

The Internet benefits from this rigid hierarchy because it greatly reduces route tables; however, it may present some significant issues to overcome for multi-homed sites. Many working groups are in the process of resolving these issues.

Summary

More than 128 million English-speaking people and 88 million people who speak other languages are currently connected to the Internet[v].. These 216 million plus users expect organizations' websites to be available 24 hours a day, an unprecedented demand for a mass medium.

While inter-organizational communication is increasingly flowing through the Internet, outside individuals are also using e-mail more frequently as their primary means of contacting organizations. In addition virtual private networks over the Internet have drastically changed how employees and business partners communicate with each other. All of this communication depends upon the engineering of an organization's Internet connectivity.

The reality is that ISP networks go down, whether due to schedule outages or by mishap. Redundancy is an essential tool that allows organizations to reduce or eliminate potential failures in Internet communications. Without some type of multi-homing, an organization is completely at the mercy of its upstream provider, an unacceptable proposition for more and more organizations in today's competitive e-business market.

For further information, see NetworkCare's website at http://www.lucent-networkcare.com or call 1-888-767-2988 in the U.S. or 1-727-217-2303 outside the U.S.

References

[i] RFC 1518 An Architecture for IP Address Allocation with CIDR
RFC 1519 Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy
[ii] RFC 2050 Internet Registry IP Allocation Guidelines
RFC 1930 Guidelines for creation, selection, and registration of an Autonomous System (AS).
[iii] RFC 1771 A Border Gateway Protocol 4
RFC 1772 Application of the Border Gateway Protocol in the Internet
[iv] RFC 2373 IPv6 Addressing Architecture
RFC 2374 IPv6 Aggregatable Global Unicast Address Format
[v] Statistics provided by Global Reach - http://www.euromktg.com/globstats/

Copyright © 2000, Lucent Technologies NetworkCare

This is an unpublished work protected under the copyright laws. All rights reserved.
WP.EN.EBUS.0100