

Avoiding the Frailties of the Internet with Intelligent Route Control

In any business climate, it is necessary to be prepared for contingencies that may affect business-critical operations. Business continuity is a philosophy that encourages a proactive approach to such business risks. In this paper, we discuss the philosophy of business continuity as it relates to the Internet operations of an enterprise, identify the weaknesses of the Internet, and describe the proactive approach of premise-based optimization solutions to address the risks to business continuity.

An Internap White Paper

©2009 Internap Network Services Corporation. All rights reserved.



Table of Contents

Executive Summary	3
Business Continuity Defined	4
Business Continuity and the Internet	4
Economic	5
Technology	6
The Border Gateway Protocol (BGP)	6
BGP Convergence	6
Invalid Updates	7
BGP Reachability	7
People	8
Environment	9
Business Continuity and Intelligent Route Control	10
Economic	11
Technology	11
People	12
Environment	12
Summary	13



Executive Summary

In any business climate, it is necessary to be prepared for contingencies that may affect business-critical operations. Business continuity is a philosophy that encourages a proactive approach to such business risks. In this white paper, we discuss the philosophy of business continuity as it relates to the Internet operations of an enterprise. We further identify the weaknesses of the Internet in four general categories (economic, technology, people and environment) and investigate the risks that are posed by these weaknesses. In conclusion, we describe the proactive approach of premise-based optimization solutions to address the risks to business continuity presented by the Internet in a manner consistent with the philosophy of business continuity.



Business Continuity Defined

Business continuity is a philosophy of proactive contingency planning and implementation, which anticipates conditions disruptive to the effective continuance of critical business activities. Disaster recovery, while related, is not precisely the same notion as business continuity. Disaster recovery consists of plans, procedures, facilities, etc., designed to affect recovery of normal business operation in the event of a disaster, while business continuity endeavors to prevent situations from having a negative impact on business. In other words, disaster recovery is best characterized as reactive, while business continuity prescribes a proactive approach to address inevitable failures of varying severity and that may even precipitate disaster.

A typical approach to developing business continuity strategies is to first assess the elements of the business that are vulnerable to failure and to further assess the likelihood and severity of potential failures. Once risks have been identified and categorized, contingency plans, procedures and automations are designed and put into practice. Since its plans are more likely to be applied regularly, business continuity is further set apart from disaster recovery with improved feedback capabilities of measurement and validation. Thus, feedback derived from experience provides the capability of improving responses to contingencies.

Business Continuity and the Internet

Computer networks and IT systems play an important role in almost all modern businesses. As such, they represent important strategic assets whose availability and efficacy are critical to business operations. Unfortunately, most businesses rely on network protocols and remedial redundancy in order to ensure availability of its network's critical elements.

The Internet does not lend itself through traditional means to proactive measures when it comes to vulnerabilities and potentially disruptive situations. Rather, the traditional protocols and operational practices used on the Internet are best characterized as reactive, requiring dramatic failure before being remedied.

To characterize the risks inherent to the Internet, we can create a chart with two axes: vertical, indicating severity, and horizontal, indicating likelihood. A necessary element in the risk assessment process is to quantify the dollar costs associated with the various contingencies that may be encountered. Therefore, increased severity corresponds to increased monetary cost. Risk assessment varies from business to business and also based on the degree to which a company is reliant on the network; however, a typical chart is provided in Figure 1 on the following page.



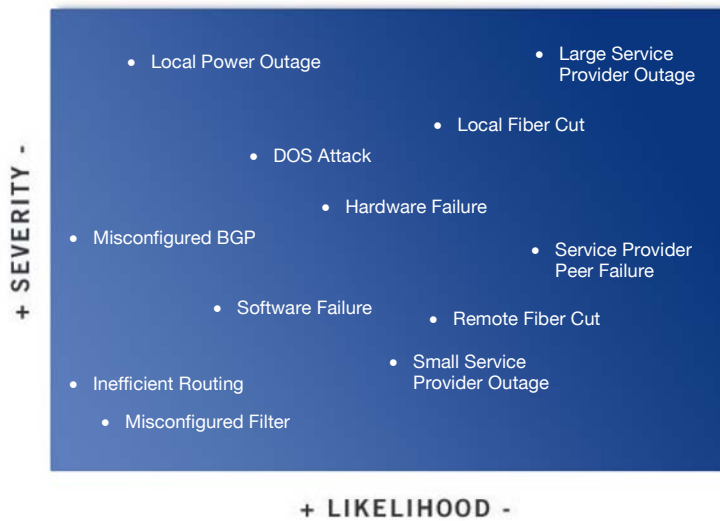


Figure 1

In developing a business continuity strategy, one must first address the risks that are most likely and will have the greatest impact on business. Some of the disruptive events that may affect a network (e.g. local loop failure, line card failure, etc.) may be addressed simply through redundancy and service level agreements with vendors and service providers. However, most of the disruptive cases that one expects to encounter on the Internet and that are most likely to occur are not simple to address. The causes of these difficult problems are many, but they lie primarily in the Internet's fundamental architecture and its openness.

The Internet is often described as a 'network of networks' interconnected through various means. The universality of interconnections that make up the Internet is its primary benefit, but it also represents the Internet's most striking weaknesses. These weaknesses can be generalized into four categories: economic, technical, people and environment. In the case of enterprises, the ability to address the risks presented by these weaknesses is limited to the company's own network and the circuits used to interconnect with its service providers. This is further complicated when enterprises communicate with remote users across the Internet as is the case with most VPNs and ASPs. In order to understand the Internet's weaknesses, and the risks they pose, it is useful to define the risks and consider the reasons that they exist, plus the impact they may have on an effective business continuity strategy.

Economic

Settlement-free peering and best-effort delivery are the cornerstones of the efficiency and relatively low cost of operating and connecting to the Internet.

The Internet is considered a tiered system, although the number of tiers and its proper placement within the hierarchy is debated. For the sake of argument, two general tiers are all that is necessary. The first tier is the core of the Internet, or the service providers. The second tier is comprised of those who connect to the Internet core via service providers. The economic distinction is that service providers generally do not pay for access to one another's networks. Those in the second (or access) tier pay only the providers they are physically connected to for



access to the entire Internet. In this scheme, an economic trade-off has been made, namely that of universal connectivity at the expense of direct economic relations with all who participate.

The efficiency and flexibility of the trade-off are clearly successful; however, these benefits do come at a price. Since there is no all-inclusive economic relationship among participants, there is no financial obligation to carry traffic between networks. The resultant philosophy of most service providers (and almost synonymous with IP) is described as 'best-effort' delivery of traffic. Best-effort delivery means that each service provider will make every effort possible to ensure traffic reaches its destination. However, since most destinations are likely on another network, and that network is not likely to have a financial obligation toward the sender, each service provider only takes responsibility for its own part of the path, with no guarantee for end-to-end transit.

Technology

Routing protocols (at the network layer) determine the paths along which Internet datagrams are forwarded. These routing protocols are divided into two categories: the Internal Gateway Protocols, such as RIP, OSPF, IS-IS, and the Exterior Gateway Protocol, BGP.

The two categories of routing protocols are used in different circumstances. Internal Gateway Protocols are used within an Autonomous System (AS). BGP, the only Exterior Gateway Protocol currently in use, communicates routing information between autonomous systems. Nonetheless, both categories of routing protocols provide IP reachability information and, to some degree, routing policy information.

The Border Gateway Protocol (BGP)

BGP is one of the most important protocols on the Internet since it is the only protocol capable of communicating routing and policy information across administrative boundaries. BGP works by establishing 'peer' relationships between routers. The peers exchange routing information in the form of route updates. Each update received from a peer is processed locally, and once a route has been selected as best, based on prefix length, policy and logical topology, the route is installed into the local routing table. Each router must process all BGP updates, and when new routes are installed or old routes are removed, the router generally propagates the newly selected routes to all of its peers except the peer that originally advertised the selected route.

BGP Convergence

Because BGP is an extremely flexible protocol and operates across administrative boundaries with varying routing policies, a large amount of time may lapse between the initial advertisement of a route and the route's installation into all of the BGP routers on the Internet. This process of reaching a final, stable state is called convergence. In some cases updates converge quickly, but on average Internet-wide convergence can take more than fifteen minutes. In pathological cases, BGP may take many hours to converge, and there are no guarantees that the protocol will fully converge every time for a given route.



The timers built into BGP and the route selection algorithms that it employs play the biggest role in the convergence of BGP. The Minimum Route Advertisement Interval (MRAI) is one of the most important; however, tuning this timer may benefit some, while negatively affecting others. In provider networks, where BGP has the largest impact on the Internet's stability, the emphasis is on being conservative with timers, advertisements, filters or anything that may cause instability.

Providers are conservative because they tend to have a large number of peers, and its networks tend to be large and have complex topologies. The number of peers is directly related to the volume of route updates, with the ISP usually receiving the same route update from a large number of peers. Since the same update is often learned from many peers, each update is typically evaluated before the ISP propagates the update within its own network and to its peers. In pathological cases, referred to as route oscillation, routes may be evaluated and propagated only to be quickly withdrawn as more of the same route updates are evaluated, thus delaying convergence. Situations such as this can create serious instability, which can cause Internet traffic to be delayed or even fail to reach its destination.

Invalid Updates

Other pathological cases can be caused by BGP's inherent sense of trust amongst peering routers. This implied trust means that all route updates are considered valid and are treated as such. Unfortunately, due to convergence delay, misconfiguration, external protocol interaction or a host of other reasons, not all updates are valid. Invalid updates in the worst cases can lead to 'routing loops' or 'black holes'. A 'routing loop' causes traffic to bounce between a pair of routers until its IP TTL expires, causing the traffic to be discarded. A 'black hole' is just what its name implies: a place where packets go never to return. Black holes are usually caused by a network advertising a prefix that it should not be advertising.

BGP Reachability

It is critical to understand that BGP provides only reachability information and makes route selections based on a composite of logical topology information (AS path attribute, prefix length, etc.) and policy (Local Preference, Multi-Exit Discriminator, AS path padding, etc.). BGP does not make decisions based on performance metrics, reliability or monetary cost.

BGP does not consider these factors because it was originally designed to provide loop-free paths between well-connected networks and to provide flexible policy mechanisms useful for steering traffic one way or another for administrative reasons. In other words, BGP's emphasis is on service provider interconnections and protocol stability. The other more subjective, but business-critical criterion, was intentionally traded off in favor of protocol stability.

Since BGP is focused on reachability and its own stability, traffic may only be rerouted due to hard failures or administrative policy changes. Hard failures are total losses of reachability as opposed to degradation. This means that even though service may be so degraded that it is unusable for an end-user, BGP will continue to assert that a degraded route is valid until and unless the route is invalidated by a total lack of reachability. BGP, as a dynamic routing protocol, is essentially reactive, and then only in cases of total failure.



People

The people who design, build and operate the Internet are the oft-overlooked glue that holds the system together. The day-to-day operation of the Internet relies heavily on people. These people are skilled at discovering and resolving problems, planning and implementing growth and technology roll-outs, managing and distributing traffic, coordinating peering, assigning and allocating IP addresses, and on and on. Unfortunately, these people aren't always able to prevent most problems from happening, and in some cases, they can be the cause of problems on the Internet.

There are a host of reasons that a problem may not be discovered or dealt with quickly, but the most common is that of the operator's vantage. For instance, given the physical and economic structure of the Internet, it is unlikely that an operator or engineer at one ISP would be aware of (or able to fix) problems on another provider's backbone unless someone (usually a customer with a greater interest in the end-to-end behavior or his own traffic) complains of trouble. The reason is not that operators do not care about their customers' traffic; quite the contrary, most operators just do not have the same view of the network and of the traffic that the enterprise has. Figure 2 is a representation of the differences between the vantage point of an enterprise and a service provider. Note the complexity of Internet connections and high level of traffic in the provider network vs. the relative simplicity and lower bandwidth levels of the enterprise.

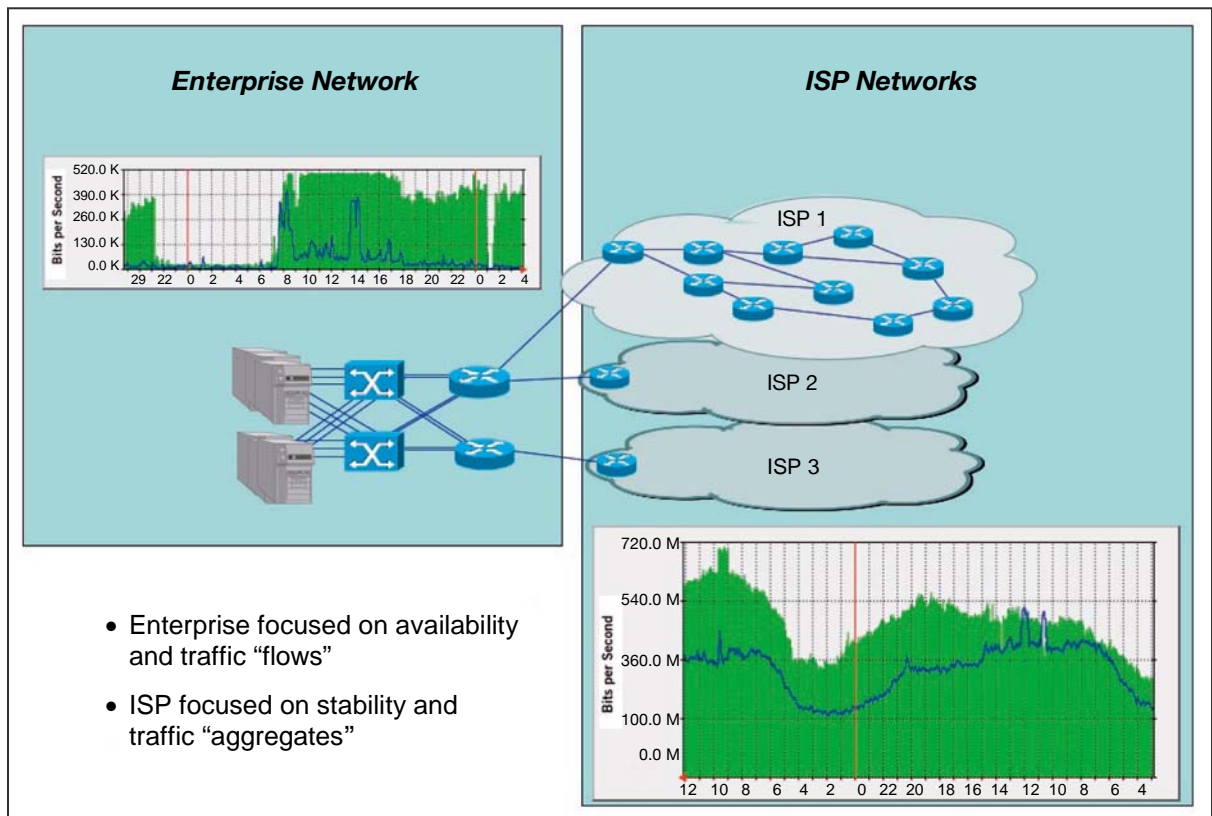


Figure 2



Operators and engineers generally view traffic as an abstract aggregate that is distributed among the various backbone links and peers that make up their network. ISPs generally view performance as a measure within its own network. Traffic and routing policy is distributed among peers and links with the aggregate, rather than each individual customer in mind. For example, congestion may be avoided at a peer in the West by moving some fragment of the aggregate to the Midwest or to the East. While beneficial to the aggregate, this could dramatically degrade performance for individual customers. This steering of traffic, in turn, could also cause congestion or inefficient routing on other networks that are beyond the operator's visible network horizon. But, because the vantage point of the operator is imperfect, there is little that can be done on their end, short of waiting for complaints to filter in.

In addition to steering traffic and managing customer complaints, ISP operators are responsible for the security, configuration and maintenance of the circuits and equipment that make up the ISP's network. None of the tasks mentioned is trivial; given that any security blunder, any configuration error or any piece of failed equipment within an ISP network could result in degradation or complete loss of service to a great number of Internet users, both customers and non-customers alike.

For example, it is common for ISPs to maintain route filter lists. These lists are used to ensure the accuracy of routing information carried within their network and between customers and peers. Even the slightest error in a single entry of these lists could cause an entire subnet to be withdrawn from the global routing table. While checks are usually in place to safeguard against such problems, and they are indeed rare, configuration errors do occur to varying degrees of disruption.

Also, consider that ISPs must perform routine maintenance such as daily filter list updates and weekly or monthly semi-major maintenance such as software upgrades. Even emergency maintenance activities due to software or hardware defects happen with striking regularity. These maintenance windows are where human error is most likely to affect the Internet. This is because during maintenance activities (no matter how well-planned) the complex interaction of people, machines, protocols and traffic makes it difficult to account for every scenario and to plan appropriately. Some of the most widely publicized and disastrous events in the history of the Internet have been the results of unexpected interactions during routine maintenance activities. In many cases, poor planning and a lack of safeguards have dramatically increased the impacts of such events.

Environment

The environmental challenges posed by the Internet extend beyond the physical environment in which it exists. The Internet environment is also comprised of policies, technologies, traffic characteristics and the presence of malicious elements such as 'hackers' or self-replicating 'worms.'

The physical environment does present challenges to the Internet in that the representative networks themselves contain physical elements that can be disrupted. The most vulnerable aspects of the Internet are the cables and fiber in the ground and stretched overhead between poles. These vital conduits are vulnerable to weather, natural disaster, civil unrest and human error. In fact, some of the most common environmental disruptions that occur on the Internet



are severed cables or fiber, usually the result of construction accidents. These incidents can cause vital backbone capacity to suddenly disappear. Since many ISP backbone circuits actually follow the same conduit, the worst cases can cause congestion, delays and complete loss of connectivity across the entire Internet.

As noted in the previous section on technology, the Internet is vulnerable to conflicting policies that may cause routing inefficiencies and these same conflicting policies can even cause technology incompatibilities that are very difficult to troubleshoot. Incompatible policies can cause outages and traffic delays that are very difficult to contend with. For example, one ISP may filter at different prefix lengths than another, or the BGP implementations themselves may be incompatible due to the use of proprietary attributes, early implementations of experimental code, etc. The only solution to these problems is often for people from different organizations to agree on a common remedy. In many cases, it can take many hours or even days to identify the people who are capable of affecting a reasonable solution, during which Internet traffic and end-users suffer.

Perhaps one of the most prevalent and alarming aspects of the Internet environment is the increasing occurrence of malicious activities. Individuals and organizations effectively protect themselves from malicious activities with firewalls, intrusion detection systems and other security products. However, there is a growing threat of malicious activity directed not only at individual organizations, but at the Internet infrastructure itself. Denial of Service (DOS) attacks are the most common malicious activity directed at the Internet infrastructure. The goal of these attacks is to make the Internet unusable for as many end users as possible, and because of this widespread impact, the most enticing target is the Internet infrastructure itself. These attacks generally cause large amounts of invalid traffic to be directed at elements (often routers) of the Internet infrastructure. Since most of the targets are not optimized to deal with traffic directed toward them, but rather through them, and since the traffic generally appears valid, traffic delays, routing instabilities and congestion often result. Since these attacks may not be targeted at individuals or a particular company, traditional security approaches do nothing to protect end-users from the effects of infrastructure attacks. Due to the openness of the Internet, the appearance of legitimacy of the traffic produced, and the improved ability of malicious elements to disguise the origin of attacks, these problems are only likely to increase in frequency and severity.

Business Continuity and Intelligent Route Control

The Internet, through its ubiquity and openness, clearly benefits efficient communication and commerce. On the other hand and for the reasons mentioned in the sections above, the risks of relying on the Internet are great, and therefore must be addressed. In some cases, the reactive philosophy of disaster recovery addresses these risks via infrastructure redundancy. However, redundancy alone cannot affect a business continuity strategy. Business continuity demands a proactive approach dealing with risks that do not always manifest themselves as failures, but as service degradations that negatively affect business operations and that can even precipitate business-critical disasters.

Intelligent route control solutions are effective in dealing with the economic, technology, people and environmental risks inherent to the Internet. The solutions address the risks in a proactive



manner with flexible policy controls and reporting mechanisms necessary to maintain a fluid business continuity strategy.

Economic

Intelligent route control solutions are designed to allow the operator flexible performance policies while optimizing cost. This is made possible by measuring performance to all relevant destinations across each transit while also taking into consideration the cost of each transit.

Cost and performance are both more relevant to an enterprise than the logical topology or the policies carried in BGP. While there is little correlation between the topology presented by BGP and performance, additional measurements can augment BGP routing decisions in order to determine the best path to use.

There does tend to be a relation between topology and cost. Since BGP prefers routes that are logically close (via AS Path attribute), it is common for the largest and most extensively peered networks to be 'close' to most destinations on the Internet. It is also common for the largest networks to be the most expensive in terms of bandwidth. As a result, BGP tends to direct traffic to networks that are often the most expensive.

Intelligent route control solutions allow the user to specify their own policy in terms of performance and cost. Therefore, it is possible for an operator to specify a policy in which the low-cost network is used for all traffic unless there is a performance problem, at which time a more expensive candidate may be used. These decisions are made based on three criteria: user-defined policy, aggregate bandwidth utilization for each transit provider, and empirical performance metrics.

Intelligent route control solutions have the ability to overcome best-effort delivery shortfalls and high-network costs by allowing users to specify relevant cost and performance policies rather than relying on the arbitrary metrics offered by BGP.

Technology

The technology risks covered earlier in this paper are addressed by intelligent route control solutions in much the same way that economic concerns are addressed. This is because intelligent route control solutions do not depend on BGP to make route selections.

Instead of relying on BGP, intelligent route control solutions take empirical measurements that are used to determine current, real-time performance characteristics of the Internet. The real-time observation of traffic and path characteristics allows intelligent route control solutions to identify problems and shift traffic quickly without the need to wait for slow BGP convergence. In the event of a disaster where reachability is lost, the route control solution is capable of making a routing change nearly twice as fast as BGP converges end-to-end. In pathological cases, such as a wide scale ISP outage, the factor is much higher. In fact, BGP will not ever converge to a new path when the current path is merely degraded, because BGP has no insight into path performance.

Alternate path analysis also allows an intelligent route-control solution to gauge path stability and can avoid sending traffic down paths that are problematic due to route oscillations, black



holes or routing loops. BGP is unable to identify such pathologies since there is an implied trust between routers, and most advertisements are assumed valid with very few exceptions.

Since intelligent route control solutions constantly measure path characteristics, pathologies are quickly identified, and the offending paths are eliminated from possible selection.

People

The risks posed by the people who operate the Internet are caused by a focus on the Internet's core, an imperfect vantage point and a tendency to react slowly to problems that affect small amounts of traffic.

Situated at the edge of an enterprise network, intelligent route control devices maintain the vantage point that is most relevant to the enterprise and focuses only the enterprise's traffic, as opposed to the vast aggregate of traffic considered downstream by their service providers.

Furthermore, the maintenance and traffic engineering activities of downstream network operators does not go unnoticed by the route control device's measurement facilities. With a keen eye on the network's performance, the intelligent route control solution is able to address problems quickly and reports network trouble to the operator, so that it can be shared with upstream providers in order to resolve persistent or recurring troubles.

The people risk to the Internet is largely unpredictable and even more difficult to quantify; however constant measurement, reporting and leveraging of all available service providers can eliminate that risk.

Environment

The environmental risks involved in relying on the Internet are many and varied; however they do all have one thing in common: they are all significant threats to business continuity. Most issues related to the Internet environment such as fiber cuts, DOS attacks, etc., happen to occur outside of a company's network and may not even be directly related to the company itself. As such, it is very difficult to identify and address such issues from a company's vantage point.

Since such issues generally appear as routing instabilities or dramatic increases in latency or packet loss, intelligent route control solutions can easily identify and remedy the problems and the affected prefixes.

ISP redundancy and intelligent route control solutions may be the only viable business-continuity solution in the face of ever increasing DOS attacks on the Internet's infrastructure. Since such attacks are not directed at an individual company, local firewalls and related security technologies have no affect. ISPs themselves may be able to contend with such attacks; however, their responses tend to be slow (depending on the size and impact of an attack) and at times disruptive, causing even legitimate traffic to suffer.

Intelligent route control offers a solution by steering traffic away from the affected infrastructure. This is not a foolproof approach. However, given sufficient diversity (two or more service providers) there is usually a path around even the most massive attacks.



Summary

While the Internet offers a tremendous opportunity for enterprises, the inherent architecture of the infrastructure and the environment in which it operates contains a great deal of potential risk. Route control systems mitigate that risk by adding intelligence to multi-homed routing decision. Intelligent route control devices take into account business objectives and service risks by maintaining routing policies and service demands as defined by an enterprise. The systems leverage redundant infrastructure, proactively measure service levels and alternate paths, and assert routing entries consistent with the proactive principles of business continuity.

Importantly, in addition to proactively addressing the risks involved in relying on the Internet, route control systems provide feedback in the form of performance, distribution and economic reports so that an enterprise may validate policies, and further tune them. Route control systems are the business continuity solution for enterprises reliant on the Internet for business-critical functions. Anything else is an acceptance of risk and an invitation for disaster.

