Information Technology Sector

Risk Management Strategy – Internet Routing, Access and Connection Services

July 2011

Contents

Figures

Executive Summary

This report identifies and evaluates the responses that the IT community can take to counter threats and risks to the IT sector's Internet Routing critical function. In the 2009 IT Sector Baseline Risk Assessment (ITSRA 1.0), subject matter experts (SMEs) from government and private industry identified several risks and threats to the operation of six critical functions of the IT Sector. The six critical IT sector functions are listed in the box to the right.

Critical IT Sector Functions

- *Provide IT products and services*
- *Provide incident management capabilities*
- *Provide domain name resolution services*
- *Provide identity management and associated trust support services*
- *Provide Internet-based content, information, and communications services*
- ***Provide Internet routing, access, and connection services***

For the *Provide Internet routing, access, and connection services* critical function, ITSRA 1.0 identified three risks that could severely impair the operation of the critical Internet Routing function. Those risks are:

- A partial or complete loss of routing capabilities, either locally, regionally, or across large parts of the world, caused by deliberate or unintentional actions
- Natural disasters or manmade incidents that could impair the operation of concentrated routing facilities
- Ineffective or impaired responses to restoring routing operations after an outage or an incident

ITSRA 1.0 identified the most significant risks to the Internet Routing critical function, as well as to the other IT sector functions. This report identifies responses and recommended actions that could be taken by the IT community to respond to and manage the risks to the Internet Routing critical function. The risks and recommended risk management responses are listed in Table 1, and described in detail in this report.

Table 1. Risk and Mitigation Overview

Risk	ITSRA Likelihood and Consequence Ratings	Risk Mitigation Activities	Resulting Likelihood and Consequence Ratings[1]
Partial or complete loss of routing capabilities (Manmade Deliberate)	Low likelihood; High consequence	• Address verification • Secure routing protocols • Infrastructure diversity • Route flap dampening • Route consolidation • Real-time route leak detection • IPv6 transition testing • Multi-person change commit • Signed route announcements	Low likelihood; Medium consequence
Partial or complete loss of routing capabilities (Manmade Unintentional)	Medium likelihood; High consequence	• Signed route announcements • Route flap dampening • Real-time route leak detection	Low likelihood; Medium consequence
Concentration of facilities: physical loss (Natural or Manmade Deliberate)	Low likelihood; Medium consequence	• Secure, hardened, redundant facilities • Lower-profile operations • Insider threat mitigation	Low likelihood; Low consequence
Impair operations support and incident response (Manmade Deliberate)	Low likelihood; Medium consequence	• Incident management and incident recovery plans • Alternatives to Internet connectivity • Evaluate incident recovery limitations • Coordinated management of routing incidents • Federal government support for use of Internet security technologies	Low likelihood; Low consequence

1 Assumes complete implementation of the items noted in the Risk Mitigation Activities column

1 Internet Routing Risk Management Strategy

This section describes the risk management strategies that the IT Sector SMEs proposed for three of the Internet Routing function risks. Those risks, as identified in the ITSRA 1.0 report[2], are:

- A partial or complete loss of routing capabilities, either locally, regionally, or across large parts of the world, caused by deliberate or unintentional actions
- Natural disasters or manmade incidents that could impair the operation of concentrated routing facilities
- Ineffective or impaired responses to restoring routing operations after an outage or an incident

For each risk, the SMEs had the option of taking one of four approaches to risk mitigation:

- Avoid the risk
- Accept the risk and its potential consequences
- Transfer the risk to another sector or entity
- Mitigate the risk by preventative or proscriptive action.

The IT Sector SMEs used the *Mitigate the risk* as the selected response for all three risks. The following sections list and analyze activities that can reduce risks to the three specific risks identified in ITSRA 1.0.

1.1 Risk of Concern – Partial or complete loss of routing capabilities (Manmade deliberate or manmade unintentional)

1.1.1 Risk Overview

The operation of the Internet is based on the concept of hop-by-hop packet-switching. Packets of data, called Internet Protocol (IP) datagrams, which contain application-level data, such as a request from a user's browser to connect to an Internet web site, are switched across the Internet's communications links and networks by routers. IP routers are special-purpose devices that examine the destination IP address of each packet, and then map the IP destination addresses to the best next router to use to deliver traffic to the packet's ultimate destination. Each router forwards the IP datagrams on what it considers the best "next hop", until they reach the router or network to which the destination system is connected. The routers use routing protocols and routing algorithms to select the best available route to a destination, then build and maintain routing tables to select the best way to send traffic to a destination network. The routing tables are used to derive forwarding tables in the routers, and that information is used to switch packets through the network, on a hop-by-hop basis, towards the destination.

Router functions can be broken down into different "planes", which are the control, forwarding, and management planes. The router's control plane uses the information it receives through routing protocol exchanges and configuration commands to build the routing tables and the forwarding information base table (FIB), and to map routes to the router's communications interfaces. The forwarding plane is the router's packet-handling mechanism. It reads IP header destination addresses, looks up the next hop in the forwarding table, and forwards IP datagrams to their next hop. The management plane controls the router's operation. A router administrator uses the management plane to perform management and

[2] The ITSRA can be accessed via the following link:
http://www.dhs.gov/xlibrary/assets/nipp_it_baseline_risk_assessment.pdf

administrative tasks, such as entering commands, configuring routing policies, and initiating and examining logging. The router administrator accesses the router through a PC or terminal from the local network or the Internet. Figure 1 illustrates the logical components of a router, and the general relationship of the planes to each other.

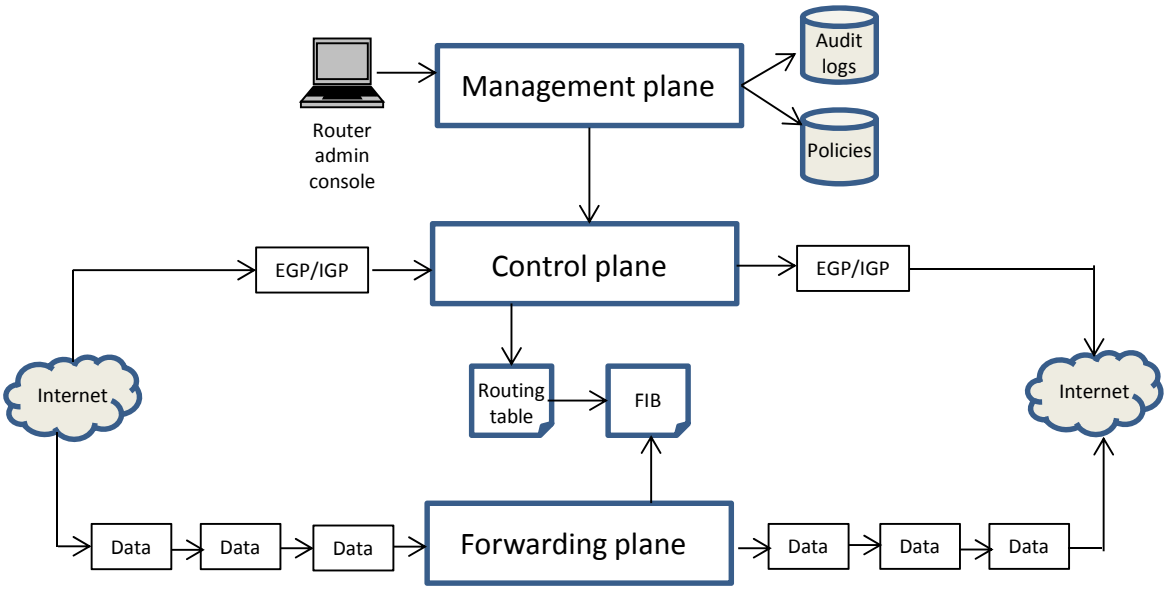

Figure 1: Router Management, Control, and Data Forwarding Planes

Each router in the path from source to destination makes its own independent forwarding decisions, based on its current view of the Internet in its routing table, about the best interface and data path to use to forward each packet. Assuming a complete path exists from source to destination, and there are no firewalls or packet filters that would block it, the packet will arrive, most likely after being forwarded by one or more IP routers.

Other networks, such as the telephone network, use similar techniques to determine the route traffic traverses between two end-points. One of the differences between the phone network and the Internet is that in most cases, the phone network sets up a temporary, dedicated connection for the duration of the call. Internet routing, by contrast, is "connectionless", as each IP datagram is routed independently, but the datagrams that constitute a single message or transaction are re-assembled in the correct order by the destination system.

Internet routers operate independently, except for exchanging information with neighboring routers about which networks are reachable through them. The Internet operates without a central authority or controlling mechanism. The decentralized nature of the Internet makes Internet routing particularly resilient, highly robust, and resistant to complete failure. A complete failure of routing throughout the Internet is extremely unlikely, but a partial or localized outage or failure is possible. It could be caused by a concerted logical or physical attack on the routing or communications infrastructure, a natural disaster, actions by a nation-state to restrict Internet access, or other actions. Routing instabilities introduced by unstable, misconfigured, or misbehaving routers could also cause system-wide problems.

Control of Internet routing is distributed throughout the Internet, so it is unlikely that all routers would be affected by any single incident. However, localized or widespread Internet routing failures or outages could occur, either because of manmade deliberate or manmade unintentional actions. These can have significant impacts on Internet routing. Some of the possible causes of routing outages include:

Router hardware or software failures – Routers are specialized computer systems running routing software on specialized router hardware. Like all computer systems, routers may suffer hardware or software problems that degrade or impair their operation. Router software problems can be caused by flaws in the router code, or by viruses or malware accidentally or deliberately introduced into the software that affects the control plane or the processing resources of the router.

Improperly configured routers– Routers use routing protocols to communicate with other routers to advertise the networks that are attached to or are reachable through them. Router software uses these routing protocols, which influence other routers' announcements about destination reachability to build entries in their routing tables. The tables indicate where to direct traffic to reach other networks in the Internet. The router software automatically selects the best available route to a destination network, but the parameters that influence the selection are largely determined by how the operator configures the routing features and the routing policies within a routing domain. A "routing domain" is a network or group of networks that is administered under a common set of routing policies. A router administrator could configure the routing features incorrectly, causing mistakes in policy which result in overriding desired routes, or advertising reachability for routes for which they do not provide connectivity. Such errors could occur either inadvertently or maliciously.

Incorrect routing announcements – A router's routing policies could be misconfigured, so that the routing tables mask the existence of otherwise reachable networks. Misconfiguration could also introduce a route with a more desirable set of attributes, which may draw traffic to that router, even though there may be shorter or more direct routes available. A router could also announce to neighboring routers that it provides reachability to networks to which it does not connect. These problems may be caused by a routing administrator making an unintentional routing configuration error (causing "leaking routes"), overriding the normal operation of routing, or incorrectly configuring routing options so that routing policies are not executed properly. They could also be the result of a deliberate action to block or divert traffic ("route hijacking") destined for certain Internet sites or systems. Many deliberate route hijacking or inadvertent traffic blocking incidents that have been detected are resolved relatively quickly, but some intentional incidents can last several hours or days. Route hijacking that diverts rather than blocks traffic can be difficult to detect, and it is often dependent on the local network topology. It may be a legitimate way to re-route traffic around an Internet outage, or it may be an attempt by an Internet Service Provider (ISP) or a nation-state to monitor or examine traffic before it is delivered to its destination.

Routing protocol complexity – Complexity in routing protocols and their application can lead to errors in comprehension and mistaken configuration. Internet routers typically use several protocols to communicate with each other to exchange routing information. The most widely used router protocol is the Border Gateway Protocol (BGP), which is widely used for inter-domain routing on the Internet. Other routing protocols, such as the Routing Information Protocol (RIP), Intermediate System to Intermediate System (IS-IS), and Open Shortest Path First (OSPF) are also used. Generally speaking, RIP, IS-IS, and OSPF are used for routing within smaller ISP and corporate networks, and are referred to as interior gateway protocols (IGPs). Internal BGP (iBGP) is used for exchange external routes internally within an Autonomous System (AS). External BGP (eBGP) is used between ASes, and it is the standard, de facto exterior gateway protocol (EGP) for advertising network reachability information between ASes, larger networks, and among carriers on the Internet backbone. BGP is a standard protocol, but its implementation and configuration syntax vary somewhat between different router manufacturers and open source versions. BGP is designed to support extensive routing policies. It allows flexibility and choice in network interconnectivity, which is usually based on the business relationships between network operators. As networks using BGP become more complex, configuring BGP properly can demand sophisticated network engineering experience. Configuring and tuning BGP can be particularly

problematic in networks that have multiple connections to different carriers and ISPs, and that need to manage outgoing traffic carefully to direct traffic across the most cost-efficient, preferred, or optimal links. BGP and other routing protocols use the same communications path to pass data packets and control plane information (i.e., routing information), so they may be vulnerable to attacks that manipulate the data plane to attack the control plane or exhaust data plane resources.[3]

Insufficient routing protocol security – The original design of the Internet was based on principles of mutual trust. The early members of the Internet community were well-known to each other, and the operation of the network was optimized for connectivity, performance, and function. At that time, hacking was a good thing, because the hackers were known, trusted members of the community, and their "hacks" were intended to improve the Internet and create new capabilities. Security was not a consideration in the design of the protocols of the Internet infrastructure, including the routing protocols and DNS. Consequently, the routing protocols lack many of the controls that might provide transport connection and integrity protections. Consequently, this leaves them vulnerable to attacks that would affect routing announcements and other router-to-router communications. As new security standards are introduced to routing protocols, there is a chance that unforeseen problems within the new feature sets (in design, implementation, or configuration) could cause intermittent routing interruptions. They could also introduce new dependencies or attack surface elements that do not exist in today's systems.

Communications link outages – Routers depend on communications circuits, such as leased lines, local area networks, and microwave, fiber optic, and satellite links, to connect to neighboring routers and networks. Most of these links are operated by communications common carriers, which are usually not the same organizations that run the routers. These links may fail or go out of service periodically because of technical problems, cable cuts, or other issues beyond the control of router operators. An outage in a communications carrier's core network may not affect Internet routing very much, as the carrier would most likely have other routes available to which the traffic could be switched. An outage in the "last mile" of a circuit serving a router facility may cause a connectivity problem, because there may be no alternate path to the Internet until the circuit is fixed. Additionally, the ongoing convergence of legacy data communications infrastructure (e.g., PSTN, Frame Relay, ATM, and X.25) to IP and MPLS-based network via technologies such as pseudowires may introduce new dependencies that aren't easily identified, and that introduce new shared risks.

Outsourced router components – The router industry was developed largely by U.S. companies, which still dominate the market. U.S. router vendors design most of their systems, software, and chips, but much of the chip fabrication and product assembly is now done less expensively overseas. In recent years, companies in Asia have moved into designing and building routers, transmission system elements, and telephone switching equipment, and they have been gaining market share. Concerns have been raised about the potential risk that overseas manufacturers could embed malicious code or logic bombs in router hardware and software, telephone switching gear, and other infrastructure equipment. If this were done, the malicious code could permit adversaries to eavesdrop on or interrupt communications, as well as disrupt SCADA and industrial and process control systems. Furthermore, supply chain management and integrity issues may not be able to identify counterfeit or compromised devices or open source software in the network that purport to have been developed by legitimate vendors,

IPv6 Adoption - Recently, the Internet Corporation for Assigned Names and Numbers (ICANN) issued the last set of Internet Protocol Version 4 (IPv4) addresses to the Regional Internet Registries (RIRs). Figure

[3] *Losing Control of the Internet, Using the Data Plane to Attack the Control Plane*, Hopper, N, Kim, Y, et al; http://www.cs.brown.edu/people/jes/papers/2011_NDSS_AttackingBGP_Schuchard.pdf

2 (below) depicts the Internet number resource hierarchy, by which IP addresses and ASNs are allocated. IP address blocks and ASNs are delegated by the Internet Assigned Numbers Authority (IANA) to RIRs, such as ARIN and APNIC, which may in turn allocate them to National Internet Registries (NIRs), which assign them to ISPs, which assign them to their customers. Once the RIRs have allocated the remaining IPv4 address space, only IPv6 address space can be issued. This means that very soon, ISPs and businesses will only be issued IPv6 addresses.

The currently running IPv4 address space will be used for many years to come, but networks and systems must be able to communicate using both addressing schemes. The adoption of IPv6 is occurring at a slow pace. Most updated router software has, in most cases, supported both IPv4 and IPv6 addressing for several years, and the plan for networks to add IPv6 addresses usually involves supporting both IPv4 and IPv6 addresses, as well as adding more hardware, memory, and computing power to accommodate IPv6. DNS services are involved in the adoption as well, so they must be configured with IPv6 AAAA resource records that point to the IPv6 addresses of hosts and servers, and must also be available to both IPv4 and IPv6 transit, in order to answer queries coming from either network layer protocol.

For the augmentation to work, DNS servers must respond with IPv6 addresses, routing announcements passed between routers must include IPv6 addresses, and Internet routers must see a path and have connectivity to destination IPv6 networks. IPv6 routing problems will not necessarily cause a routing outage, but they could cause networks or hosts using IPv6 addresses to be cut off from some hosts on the Internet. IPv6 systems and networks share the same vulnerabilities as IPv4 systems to hijacking, transport connection DoS, and other attacks. Once IPv4 address allocation ceases, unprepared networks may not be able to communicate to new systems or networks, causing a connectivity balkanization of the Internet.

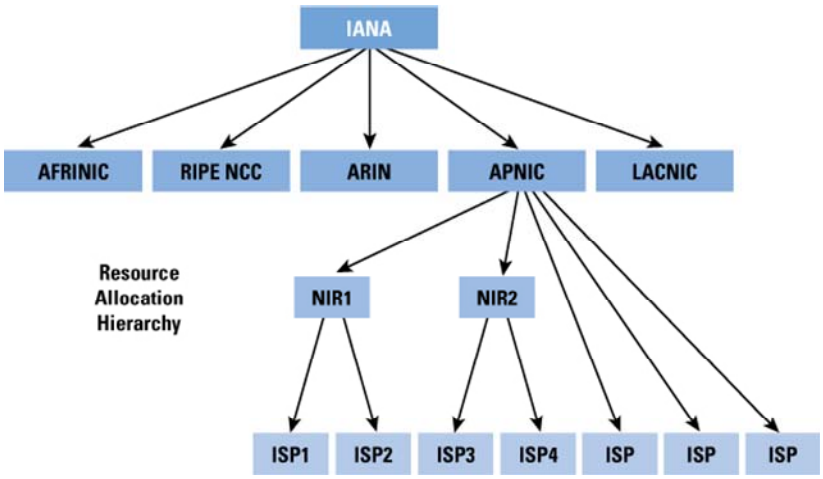

Figure 2: Internet Number Resource Allocation Hierarchy

32-bit Autonomous System Numbers (ASNs) - Not only are the RIRs running out of IPv4 addresses, they are also running out of 16-bit ASNs. The solution is for RIRs to assign new ASes longer, 32-bit ASNs. However some BGP routers and other Internet components may not yet understand 32-bit ASNs, just as some Internet components and applications can't handle IPv6 addresses. An AS is a group of IP

networks controlled by a network operator, such as an ISP, and an AS typically employs a common routing policy across the entire administrative domain. AS numbers serve other important routing functions in addition to identifying administrative domains. The list of ASNs a given route announcement has traversed is attached to a route announcement. This information helps routers detect routing loops, and apply routing policies.

Rapid Growth of Routing Tables – To determine where to send traffic destined for a host on a specific network, routers build routing tables in memory that map networks to router connections, and then use that information to generate a FIB. Routing tables have grown as the number of Internet networks has expanded. Routing tables that map all reachable networks in the Internet may have more than 360,000 network prefixes and more than 35,000 AS entries[4], and each prefix in the Internet may be reachable via multiple paths. Routing information may change when new networks and access circuits are added to the Internet, all of which must be conveyed via BGP update messages, and reflected in updated routing tables. Maintaining and stabilizing large routing tables can be an ongoing challenge for Internet backbone providers. The resources necessary to maintain that information, and the frequency of changes associated with it, continue to increase considerably. The primary risk is that routing tables could become so large that they exceed routers' memory or processing capabilities, or that the number of unique prefixes in the routing system will exceed the capacity of a FIB. That may cause some routers to drop parts of the table, leaving some networks unreachable through those routers. It may also trigger wider system instability because the protocol algorithms may take excessive time to converge when the routing tables are so large. It is also possible that attackers could introduce route flaps into Internet core routers, which could cause instability in the routing tables, reducing Internet routing capabilities until the tables stabilized again.

1.1.2 Risk Response

In the risk analysis, the sector partners determined that some of the steps that could be taken to mitigate this risk are:

Routing Policy - In the risk analysis, the sector partners noted that some of the steps that could be taken to improve the resilience of Internet routing are embedded in a fundamental capability of the BGP protocol.

- BGP allows router administrators to create local routing policy, which governs how the router handles traffic for or from the routes announced in BGP messages. Policy allows router administrators to ensure that routers handle traffic properly, to tune router performance, and to ensure that router operation complies with business and administrative requirements. The policy boundary is defined by the AS, which determines what is governed by the local policy, and what is governed by policies in other ASNs.

- Formulating and applying appropriate local routing policy is an important response to managing and controlling routing risks. Those risks include improperly coded software, malicious changes to policy attributes, and policy configuration errors. Another risk response is raising the level of assurance that routing announcements are correct. Resource Public Key Infrastructure (RPKI) address verification is a relatively new tool that can provide more assurance for address verification.

[4] NANOG51 http://www.nanog.org/meetings/nanog51/presentations/Tuesday/ahmed-churnevolution.pdf

- Most BGP software presumes a default local policy of accepting all route announcements from peers and announcing all routes learned or configured from inside the AS. This promiscuous default policy encourages a broad exchange of routing information, but it has several negative consequences. Software errors can cause excessive route churn, and configuration errors may promote questionable or unsafe behaviors throughout the Internet. Efforts to mitigate this risk have focused on standardizing ways to express local policy, and adding capabilities to allow other routing domains to track or identify policy configuration choices that can affect the performance of routers.

- The first efforts to express local policy used brute-force prefix filtering, which blocked unwanted prefixes. Other efforts to express policy used BGP attributes, which were coded "hints" to domains more than one hop away about the policy intentions of the authorized prefix holder. A programmatic language, the Routing Policy Specification Language (RPSL), provided a rudimentary means for documenting local policy and publishing it for peer and third party review.

- These steps helped with risk mitigation to the extent that other routing domains could check another domain's local policies. However, this left the problem of determining the true AS origin or authorized intermediate AS from the authorized prefix holder. The more comprehensive solution to this problem is RPKI, which is discussed below. The addition of address verification in RPKI simply provides the peering engineer with more tools to detect discrepancies between desired local policy and what is actually seen at the router. Much of the automation remains in the formulation of local policy, but risks remain due to the manual work of installing the policy on routers.

Address verification – The Internet operates on an implied level of trust, and one aspect of that trust is that networks actually have the right to use the addresses their routers announce. That is, those networks are authorized to assert reachability for the relevant prefixes from the specified ASNs. Internet routers announce routes to reachable networks, but address verification may add another layer of "defense in depth" to Internet routing. Router software problems, malware, Internet hacking, and other problems would remain, but this may help solve some of the problem.

- ARIN and other RIRs have recently implemented or are planning to implement RPKI, which supports the use of digitally signed Route Origin Authorizations (ROAs). An ROA associates a set of IP network addresses with an AS that serves those network addresses, and is authorized to assert reachability for those networks. It indicates an address holder's authorization for an AS to announce that network. ROAs attest to route ownership, tying a network prefix to an ISP or an organization that has been assigned that address by an RIR.

- RPKI adds an extra layer of assurance in Internet routing, but it has been difficult to implement. ROA repositories and the digital certificate authorities must be established by the RIRs or some other trusted party, and ROA validation only establishes a loose association between the network prefix and the address holder.

    - For it to work properly, RPKI checking must be done throughout the Internet, and there must be a single, coordinated ROA repository that everyone believes. It would be aligned with the current Internet number resource allocation hierarchy.

    - Ideally, it would be singly-rooted, as any other case would leave relying parties without adequate information to resolve collisions in the system.

- Achieving such a goal may prove to be difficult, but formal policy support and adoption of RPKI by federal government agencies would encourage its wider use.
- However, while RPKI seems to have momentum at the moment, it does introduce an entirely new party into the routing system and as a result, it expands the attack surface.
- Alternatives to the hierarchical RPKI architecture should be evaluated as well, such as using the DNS with number resource certificates that can be validated through DNSSEC, or non-hierarchical models which mimic and then secure the existing trust relationships in the routing system.

Secure routing protocols – The BGP routing protocol is vulnerable to route manipulation, route re-direction, and denial-of-service attacks[5]. These threats may be mitigated by authentication between neighboring routers, filtering routing prefixes, and limiting the size of Autonomous System paths to destination networks, as well as object-level integrity checks for network layer reachability information and associated attributes. Secure routing protocols, such a Secure BGP (SBGP), allow routers in different networks to determine that they are receiving route announcements from authorized routers, and not from unauthorized sources.

- The purpose of Secure BGP is to help detect intentional traffic diversion and blocking. Routers using the Secure BGP protocol digitally sign their route announcements, so that neighboring routers can authenticate them.
- However, for both carriers and content providers, Secure BGP remains a controversial subject.
    - Getting all BGP speakers to sign BGP announcements is a significant undertaking, and real-time verification of the signed announcements adds operational complexity.
    - Some ISPs and Internet Exchange Points (IXPs) contend that while Secure BGP would increase the level of assurance of route announcements, it would not necessarily mean that those route announcements would be correct. They advocate better route filtering and policy application as a more effective solution than Secure BGP to reduce or eliminate routing problems.
- Recent research[6] has indicated that the most effective defense against route hijacking is to use both Secure BGP and route filtering. Filtering routes from adjacent networks prevents attacks that originate in those networks.
- Attaching digital signatures to address announcements would help authenticate the originators of route announcements, which may make identifying the source of incorrect route announcements easier. ISPs and carriers should announce the networks they serve correctly, and not deliberately block or divert traffic by announcing addresses they do not serve.

[5] Protecting Border Gateway Protocol for the Enterprise, Cisco Corporation, http://www.cisco.com/web/about/security/intelligence/protecting_bgp.html

[6] Goldberg et al, NANOG 49: http://www.nanog.org/meetings/nanog49/presentations/Tuesday/HowSecure_NANOG_print.pdf

- In recent years, the governments of some countries have sought to block their citizens' access to web sites the governments have found objectionable, for political, moral, or other reasons. Authenticated route announcements may help identify situations in which ISPs or nation-states block access to other networks through spurious route announcements.

Infrastructure diversity – Diversifying the IP network and underlying infrastructure that supports routing, such as long-haul communications circuits, decreases the risk of an external problem affecting the global Internet routing system.

- ISP's or IXP's routers may be working properly, but if their links to other routers are down or inaccessible, they and their customers may be cut off from the rest of the Internet.
- Route and equipment redundancy and diversity are well-established principles of communications systems design, and most ISPs, IXPs, and organization routing centers have at least some measure of infrastructure diversity.
- ISPs and IXPs must stay apprised of changes that carriers make in the links that connect them to the Internet, to make sure that any changes that are made to their carriers' external circuits do not affect diverse routing, or expose them to unnecessary risks.

Route flap dampening – Maintaining the most current versions of router software provides routers with better "route flap" suppression (dampening) routines.

- Route flaps occur when the router software repeatedly updates a routing table entry from one "best route" to a network to another, and repeatedly advertises or withdraws reachability for a given prefix.
- Route flaps occur because of router configuration errors, incorrect routing announcements, or other causes, but they waste router processing resources, and can introduce instability in other routers. When networks receive updated routing announcements from other routers, the routers should employ route dampening techniques to minimize the number of changes that are made to the routing tables.
- Reducing the number of updates to routing tables improves router performance, enhances router stability, and reduces the frequency of routing update traffic between routers.

Route consolidation – AS operators can help control the size of Internet routing tables by consolidating the networks they announce into as large a prefix as possible, instead of announcing individual networks or small blocks of networks.

- Each individual network, such as a /24 (Class C) network that is announced by an AS eventually becomes an individual entry in the routing tables of the Internet's core routers.
- Announcing a block of 256 /24s as a single /16 condenses what might have been 255 separate entries into one entry in the full Internet routing tables, reducing the routers' processing loads and potential for route flaps.
- The operator of an AS that announces large prefixes can still conduct traffic engineering to manage traffic to and from an upstream provider.

- AS operators, particularly those with relatively few networks and slow growth patterns, should invest in and utilize industry best practices to help keep routing table growth in check.

Real-time route leak detection – Systems that detect unintentional or deliberate attempts to hijack routes to networks could be improved if they were able to detect "route leaks" soon after they occurred, rather than after Internet traffic has been blocked or intercepted.

- Experimental systems have been proposed[7] to detect large route leaks by identifying abnormal types of route announcements. A route hijacking prevention enhancement to the BGP protocol, "Pretty Good" BGP, quarantines new route announcements for 24 hours, as long as an alternate route to a newly announced prefix is available.

- Neither method is perfect. The former requires constant monitoring and analysis, while the latter has the disadvantage of temporarily embargoing most bona fide route changes.

- Government and industry may elect to support the development of better route leak detection systems, to protect the Internet from traffic blocking and hijacking.

IPv6 Adoption Testing – Router software has supported both IPv4 and IPv6 for a number of years, but that does not mean that transitioning networks, applications, and the Internet infrastructure to include IPv6 will be simple. The adoption is proceeding slowly but deliberately, and in many cases, invisibly to Internet users.

- Part of ISPs' and IXPs' plans for the adoption must be detecting, analyzing, and resolving IPv6 adoption issues.

- Recovery testing must be incorporated into organizations' planning and management of the transition from IPv4 to one that includes IPv6.

- A key element in the process is developing and testing recovery procedures, so that the mechanism can recover from routing problems in either the IPv4 or IPv6 part of the adoption process.

- There is no "flag day" for a complete cutover to IPv6, so coexistence of both IPv4 and IPv6 network layer protocols is expected to last a decade or more.

Multi-person change commit – To prevent both malicious and unintentional network configuration changes, companies that control major network installations should deploy procedures that require peer review of any modifications to the routing system, as well as requiring more than one person to commit any network changes, based on multi-party multi-factor authentication. Similar procedures are used in missile silo launch sequences, so that the actions of two or more people are required to commit changes. This reduces the ability of an inexperienced, rogue, or coerced employee to affect a major configuration change without authoritative approval.

[7] Zhang and Khare, NANOG 49; http://www.nanog.org/meetings/nanog49/presentations/Tuesday/LRL-NANOG49.pdf

1.2 Risk of Concern – Natural disasters or manmade incidents that could impair the operation of concentrated routing facilities (Natural or Manmade Deliberate)

1.2.1 Risk Overview

The second risk of concern identified in the baseline assessment for the Internet Routing function is that natural disasters or a manmade incident could destroy or disable a data center or telecommunications facility that houses a number of routers. A natural disaster may also damage or disable power and communications lines at or near the router facility, which could leave the routers intact, but leave them without a way to communicate with the Internet. The majority of Internet and other communications outages involve isolated equipment failures, communications circuit outages, or faulty routing announcements. In the event of a natural disaster or a terrorist attack, restoring Internet routing operations may be hampered by problems accessing the routing facility, physical damage to communications lines, or getting replacement hardware, software, or routing updates.

The routers that provide an organization with connectivity to the Internet may be housed in a data center or in a telecommunications facility. Locating critical communications equipment such as routers and web servers in a number of geographically dispersed locations, and establishing processes and procedures for other locations to act as fail-over and backup operations sites, are well-established principles in data communications and telecommunications operations. These practices are fundamental elements of the resilience of the IT Sector.

The major IXPs, which concentrate Internet routing and communications facilities for ISPs and carriers, can be particularly vulnerable to natural disasters. An IXP may have a number of local ISPs or Internet backbone carriers co-located to exchange Internet traffic. An IXP puts routers from a number of local, regional, and backbone carriers in the same building, where they can exchange traffic directly, instead of being in separate locations connected by high-speed communications lines.

Within the IXP, the routers can send traffic to other carriers or ISPs over high-speed local area network or direct links, avoiding the expense and potential delay of routing traffic through other parts of the Internet. The IXPs are frequently located in or near big cities, where they concentrate traffic originating from local ISPs across high-speed links to other parts of the Internet. IXPs in Europe, for example, may send traffic destined for web sites in North America directly across the Atlantic to another IXP on the East Coast, where it is re-directed to other IXPs or ISPs for delivery to its destinations.

IXPs are major hubs for routing Internet traffic, so they can be particularly vulnerable to natural disasters and electric power and communications outages. IXP operators take many of the same precautions as data center operators to reduce these vulnerabilities. They usually have emergency power supplies and communications redundancy and diversity, and take other measures to ensure continuous operations.

The likelihood of an incident that would destroy or impair the operation of a major routing center may be fairly low, but the concentration of routing facilities in relatively few major routing centers does increase their vulnerability. Some of the possible causes of a natural disaster or a manmade incident that would affect routing are:

Natural disasters – A natural disaster that affects a relatively large area, such as a flood, earthquake, or a hurricane, could destroy or impair the operation of a major routing center. There are a number of major routing centers on the West Coast of the U.S., which has a well-chronicled history of natural disasters,

such as earthquakes and volcanic eruptions. Many of the major routing centers on the East and Southeast parts of the U.S. are in locations relatively close to sea level, or that are in areas that are periodically hit by hurricanes, tropical storms, and blizzards. A natural disaster affecting undersea communications cables may also affect communications from a routing center.

Physical attack on a routing center – Many ISP and IXP routing centers are located in buildings in large cities or in major suburbs of big cities. Many are in office buildings, and while they are usually unmarked, their locations can usually be found on the Internet. One web site about IXPs features a series of aerial photographs of the office building that houses one of the major East Coast IXPs. Natural disasters do not select their targets, but terrorists could use information gleaned from the Internet to target a routing hub.

Insider threat or social engineering attack on a routing center – The threat of an insider attack is always present in any organization, no matter how unlikely it may seem. Routing center insiders who have access to sensitive systems, facilities, and information are highly desirable targets for attackers, who may launch social engineering and phishing attacks against them.

Logical attack on a routing center - Routing centers, as well as other parts of the Internet infrastructure, such as DNS services, are vulnerable to logical attacks, such as Denial of Service (DoS) attacks. An attacker can flood a router with a high number of service requests, which may overwhelm the router's processing capacity, or prevent legitimate service request from reaching the router. DoS attacks are more commonly directed at specific networks or carriers, and not necessarily specific routing facilities.

1.2.2 Risk Response

The sector partners determined that some of the steps that could be taken to mitigate this are:

Secure, hardened, and redundant facilities – Routing center operators take extensive steps to secure their facilities from physical attacks and natural disasters.

- Many of the same precautions that have been used to protect major data centers have been applied to routing centers. These measures include guarded, physically isolated buildings, controlled-entry access points, and close control over access to different parts of the facility.
- Other measures include locating the facilities in locations that are less susceptible to natural disasters, establishing backup or mirrored operations in geographically separate sites, and hardening utilities and communications facilities to withstand physical attacks or natural disasters.
- The sites usually have backup generators that can provide emergency power for several days, as well as provisions for operations personnel to run the facility during a power outage or a natural disaster.

Lower-profile operations – Operators of routing centers should try to operate with as low a profile as possible, so that the location of their operations is not so visible, and so that they are less vulnerable to terrorist attacks.

- The effects of a routing center's routing operations should be apparent to Internet users, but the physical location of the routing center should not.
- Specially-trained security personnel, physical barriers, and incident management procedures can prevent terrorist or other physical threats to the facilities.

- It is not uncommon for corporate and government data centers, which usually occupy much larger facilities than routing centers, to be in unmarked buildings or located within the buildings in a protected office campus.

- The locations of some of the major IXPs in the United States are freely available on the Internet. The Internet makes it particularly difficult to maintain a high level of anonymity of data and routing centers, but wherever possible, ISPs and IXPs should attempt to operate their routing and network control facilities at as low a profile as possible.

- Maintaining a low profile for ISP and Inter-Exchange Carrier (IXC) facilities would not necessarily affect their operators' abilities to run higher-profile business, marketing, and promotional activities to support the routing operations.

Insider threat mitigation – Routing center operators can reduce the threat of insider attacks by running comprehensive background checks on employees and contractors, monitoring and auditing system activity logs, and requiring two or more people to approve and implement software modifications or upgrades.

1.3 Risk of Concern – Ineffective or impaired responses to restoring routing operations after an outage or an incident (Manmade Deliberate)

1.3.1 Risk Overview

When the elements of the computer communications network that became the Internet were developed in the 1970's, communications security was not as great a concern as it is today. Today, security is an important Internet concern, so security has been retrofitted, sometimes with less than perfect results, into Internet operations. The developers of the Internet also recognized that it would be impractical to establish a central control over the system. It was purposely designed to be decentralized, because that meant it could be flexible, and expand and contract as needed. The Internet's founders recognized that in order for it to be most useful, the Internet should be open to new members. One of their few requirements was that all of the participants in the network abide by the same rules, and use common communications protocols.

A key element of the early Internet was that its participants trusted one another to abide by the rules, and to "do the right thing", That meant helping to improve the system by solving problems of mutual concern, and devising new protocols and software that expanded the network's capabilities. The element of trust also meant that other participants were also acting for the good of the network.

Today, many Internet operations are still based on trust. For example, Internet routing is based on the assumption that routes announced by neighboring routers are correct. However, the actions of computer hackers, identity thieves, nation-states, cyber terrorists, and others have distorted the trust of the Internet. Today the simple trust model has been changed to "trust, but verify". The increased scale and scope of the Internet has presented greater challenges to developing and implementing new and innovative mechanisms for establishing trust relationships in Internet transactions

Verification has added a layer of security that is essential for establishing trust relationships across the Internet infrastructure. Public Key Infrastructure (PKI) digital signatures are now widely used to verify users' identities. The root of the DNS system has been signed, as have many of the Top Level Domains

(TLDs), such as .com, .net, .org, and .edu. Signed route announcements have been proposed to add a greater level of assurance to routing.

These steps increase the security of the Internet, but they also make it less resilient to address or recover from an incident or an attack. The security infrastructure that has been added may make the Internet more "brittle", making incident recovery more difficult and more complicated, and potentially lead to unintentional operator errors. For example, if PKI assurance were to be widely implemented in Internet operations, recovering from an incident or an outage may mean that IXPs must download or check certificate revocation lists (CRL) to verify that digital certificates are still valid. These steps would be in addition to other steps to bring the physical plant, power, and other utilities back on-line. ISPs, IXP operators, and other communications service providers have well-established procedures for incident response, crisis management, and service restoration, but adding new measures to establish verification and authentication may affect how quickly or how easily routing services may be restored.

Some of the risks that may cause an ineffective or impaired response to restoring routing after an outage are:

Brittleness of the Internet – In the interests of increasing the security and reliability of Internet communications, some groups have advocated adding more information assurance capabilities to Internet routing. For example, DNS can be made more secure by using DNS Security Extensions (DNSSEC). RPK) can provide a degree of assurance that networks advertised by an AS actually belong to it. Secure BGP can be used by Internet routers to verify the authenticity of routes that other routers advertise. The objective is to prevent fraudulent misdirection or blocking of Internet traffic by adding some method of verifying that addresses, DNS replies, or routing information originated from an authoritative source.

As worthwhile as these efforts to improve Internet reliability may be, some ISPs and IXPs are wary of them. They contend that they may add overhead and delays to Internet routing and DNS services, without adding significant benefits. Still, the implementation of some of these security measures is proceeding, even if their adoption has been slow. Some of the RIRs, which allocate IP address space to ISPs and carriers, have established RPKI services. The American Registry for Internet Numbers (ARIN), which allocates IP addresses and AS numbers for North America, is establishing an RPKI service for the address space it has allocated to its customers. The ARIN RPKI will allow ISPs to provide assurance that the ARIN address space they announce has been assigned to them, or to their customers. The objective is to reduce the threat of address hijacking and traffic diversion. Signed DNS roots for some domains exist, and ISPs and DNS registries are testing the use of signed DNS replies. Secure BGP, which passes digitally signed route updates between BGP speakers, is also being implemented, if only slowly, because of concerns about how digitally signed route update messages may affect routing table update performance.

Adding these additional capabilities may have the unintended consequence of making the Internet more "brittle", in that they may make it more difficult to recover from an Internet outage. Many Internet information assurance measures assume the existence of reliable Internet communications to work. If part of routers' Internet outage recovery procedure is to get signed BGP route announcements from neighboring routers, but the Internet outage blocks the routers from retrieving a CRL to check the validity of the digital certificates, the routes may be rejected.

Understanding service dependencies – One of the implications of the potential 'brittleness" of the Internet is that ISPs, carriers, and organizations that run their own Internet access networks must understand the dependencies in their Internet operations. In order to recover from an outage, an ISP or a carrier may

have to reload DNS zone updates or broadcast and receive routing information from other routers. Recovery may be complicated by requirements to download or access CRLs to verify the validity of PKI certificates. Network support personnel who expect to use the telephone to troubleshoot routing problems may find that long-distance calls that are normally routed over VoIP networks don't work at all or are difficult to complete. Internet phone services such as Skype and VoIP SIP phones may not work either. ISPs and network operators must understand the ways in which their operations depend on the availability of Internet communications, particularly in situations in which parts of the Internet may not be available.

Incomplete recovery planning and testing – ISPs and carriers can recover from Internet outages more quickly and with fewer problems if they have comprehensive recovery plans in place. Most reputable organizations have recovery plans, but the risk is whether those plans incorporate recovery processes for recent changes that have been made in the network or its services. Another risk to recovery planning is the extent to which recovery plans have been tested. Testing uncovers deficiencies in plans, processes, and procedures, but running full-scale tests of those plans takes time, money, and resources that some organizations may not be willing or able to commit.

1.3.2 Risk Response

The ITSRA determined that some of the steps that the Internet Routing sector partners may take to mitigate the risk of impaired operations support and incident response are:

Comprehensive incident management and incident recovery plans – There must be a comprehensive incident management plan in place that can be followed to respond to incidents, even if unforeseen events or circumstances block or delay the response to an incident.

- A comprehensive plan describes the recommended recovery steps for incidents that can be foreseen, such as a fire, power failure, or a physical attack on the operations center. The plan may describe only general incident handling procedures for incidents that are unlikely to occur or which cannot be foreseen, such as an airplane crashing into the building.

- Incidents that may have catastrophic effects, such as a fire or a nuclear attack, should incorporate procedures for a backup or a mirror site to take over operations.

- Whenever contingency and recovery plans are updated, they should be tested to make sure they will work. The tests should exercise the network equipment and services that will be recovered or brought back into service, and they should be conducted by the personnel who will conduct and manage the recovery.

- Recovery testing and contingency planning should also include backup personnel, in case a primary router administrator, system administrator, or network engineer is not available during a real outage.

Determine alternatives to Internet connectivity – Router operators must determine which of their incident recovery procedures depend on Internet connectivity, and develop processes and procedures to recover from an outage without those dependencies. This problem speaks to the "Internet brittleness" issue, in that procedures for recovering from an Internet outage should not assume that the Internet will be available to get files necessary to recover from the outage.

Evaluate incident recovery limitations – Router operators may have to accept certain limitations to their responses to outages and incidents, particularly if the incident restricts or blocks physical access to the routing facility. A fire or a natural disaster may mean that router operators or emergency personnel cannot reach the routing facility. In that case, the router operator may have to accept a delay in executing local recovery actions in the incident response plan, or provide for an automatic re-start capability, all of which are factors in packet-switched communications emergency management.

Coordinated management of routing incidents – Internet governance agencies, ISPs, and IXPs should work together more closely to coordinate their Internet outage incident response plans. Internet routers will direct traffic around a network affected by a router outage, but ISPs and IXPs should work together more closely to develop and coordinate incident response plans.

Federal government support for use of Internet security technologies – Internet security technologies, such as RPKI, DNSSEC, and Secure BGP are being adopted slowly by ISPs and IXPs.

- The federal government depends extensively on the continued operation of the Internet, yet it does not have a formal policy of supporting or encouraging the use of these Internet security technologies.
- The .gov DNS TLD has been secured using DNSSEC, but only a third of the sub-domains under .gov have been signed.
- Creating policies will not necessarily change network practices, but they will signal the intention of the federal government to adopt security technologies that benefit itself and the rest of the Internet community.