

Towards Fool-Proof IP Network Configuration Assessment

Abstract

IP networks have come of age. They are increasingly replacing leased-line data infrastructure and traditional phone service, and are expected to offer Public Switched Telephone Network (PSTN)- quality service at a much lower cost. As a result, there is an urgent interest in assuring IP network security, reliability, and Quality of Service (QoS). In fact, regulators are now requiring compliance with IP-related mandates. This paper discusses the complex nature of IP networks, and how that complexity makes them particularly vulnerable to faults and intrusions. It describes regulatory efforts to mandate assessment, explains why many current approaches to IP assessment fall short, and describes the requirements for an effective solution to satisfy business, government, and regulatory requirements.

1. Introduction

IP networks throughout the public and private sectors are now mainstream. Every day, they are responsible for transporting real-time and critical voice, video, and data traffic. As a result, it is no longer acceptable for IP networks to deliver “best-effort” service. They are expected to perform at carrier-grade level. However, it is enormously challenging to deploy IP networks and assure consistent, high quality service delivery, given that they are such complex and dynamic environments.

IP networks are comprised of devices such as routers, switches, and firewalls that are interconnected by network links. These devices are not “plug-and-play.” Rather, they must be provided with specific instructions, also known as scripts or configurations, which indicate exactly how they are to interact with each other to provide the correct end-to-end IP network service. This is why we refer to IP device configurations as the DNA of the network — they literally control the network’s behavior.

Unfortunately, there is nothing simple or standard about these configurations. Each one must be manually programmed into the network devices, and every vendor uses a different configuration language for its devices. Furthermore, device configurations change virtually every day in response to new application deployments, organizational or policy changes, new device or technology deployments, device failures, or any number of other reasons. Device configurations have an average of 1000 lines of code per device. A Fortune 500 enterprise that relies on IP can easily have over 50 million lines of configuration code in its network. But numbers of devices and lines of code are only part of the problem. Configurations can contain parameters for about 20 different IP protocols and technologies that need to work together. Those protocols and technologies must satisfy various, constantly changing service requirements, some of which are inherently contradictory, such as security and connectivity with the Internet. So configuration errors can easily occur due to entry mistakes, feature interaction, poor process, or lack of a network-wide perspective.

The labor-intensive and constantly changing nature of IP network operations is analogous to software development. The key difference, as illustrated in the figure below, is that software development has matured to the point where errors are significantly reduced by having different people responsible for requirements, code writing, and testing. More importantly, testing in software development is a well-established process, while there is no similarly rigorous process in IP network deployment and operation. The impacts of configuration errors are well documented. BT/Gartner has estimated that 65% of cyber-attacks exploit systems with vulnerabilities introduced by configuration errors [1]. The Yankee Group

has noted that configuration errors cause 62% of network downtime [2]. A 2006 Computer Security Institute/FBI computer crime survey [3] conservatively estimates average annual losses from cyber-attacks at \$167,000 per organization.

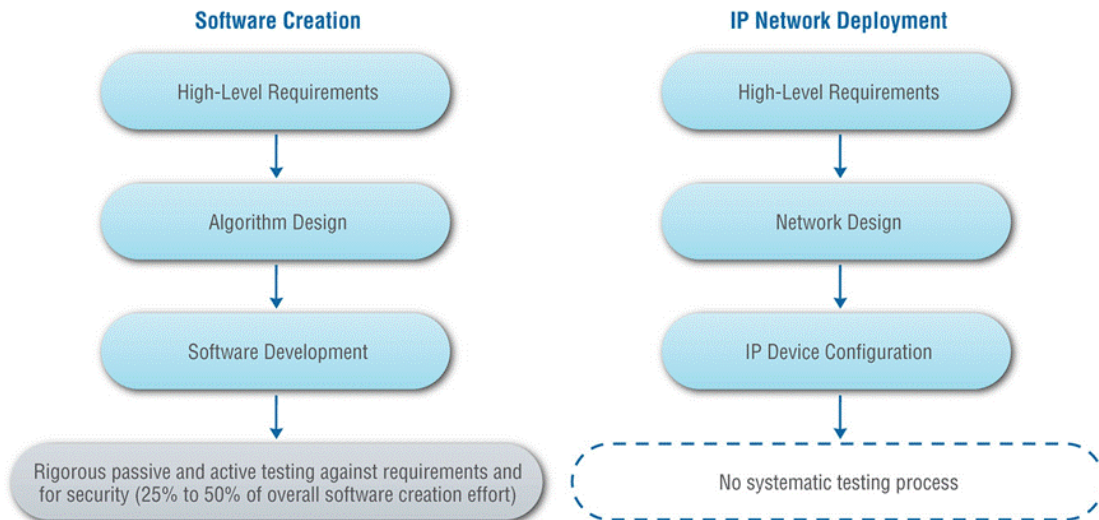


Figure 1: Inadequate Testing in IP Network Deployment Compared to Software Development

2. Configuration errors found in operational IP networks

IP network configuration errors are hard to detect since they can require validation of multiple protocols and device configurations simultaneously. These errors typically remain latent until they are exploited by cyber-attackers, discovered by auditors or result in network failures. Below are specific examples of configuration errors, and their potential impact on the organization. Many of these errors have been discovered in operational networks while performing configuration assessments.

2.1. Reliability

Organizations that depend on the IP network to provide a very reliable service have to ensure that there are no single points of failure in the network. It is not sufficient to just provide redundant network devices and links at the physical level. It is also critical to ensure that the configurations of the network devices make use of the available redundant physical resources, and that the redundancy is ensured across multiple layers. Examples of misconfigurations that result in single points of failure include

- Mismatched device interface parameters. This mismatch prevents devices from establishing logical connectivity even though physical connectivity exists.
- Hot Standby Routing Protocol (HSRP) inconsistently configured across two routers that are expected to mirror each other. The standby router will not take over when the main router fails.
- Access control lists (ACLs) or firewall rules stop specific application traffic on a path. So even if the path provides redundancy in general, the ACLs/rules still are cause for a single point-of-failure to exist for the specific application traffic.
- Use of a single Open Shortest Path First (OSPF) area border router (ABR). The OSPF areas that are connected by the ABR will become isolated if the ABR fails.

- Multiple VPN connections sharing a single physical link or device. Redundancy expected from the multiple VPN connections is not provided due to their dependence on a single physical resource.

In addition to errors that introduce single points of failures described above, other errors in configuration of IP routing protocols such as OSPF, Border Gateway Protocol (BGP), Multi-Protocol Label Switching (MPLS), and Intermediate system to intermediate system (IS-IS), can also impact network reliability. Examples of such errors include

- Inconsistent routing parameters such as OSPF Hello and Dead interval across multiple routers. OSPF will not function efficiently if such parameters are inconsistent, resulting in ephemeral traffic loops and poor network performance.
- Best-practices proposed by vendors and experts for routing protocols, such as use of a full-mesh to connect all internal BGP (iBGP) routers, and OSPF route summarizations include IP addresses of all interfaces except the loopback interface of a router, are not followed. Not adhering to best-practices generally results in an unstable network that will have intermittent connectivity issues that are difficult to debug.
- Use of inappropriate IP addresses, such as addresses assigned to other organizations, or private addresses in parts of the network directly exposed to the Internet. Such networks will start advertising routes for IP addresses they do not own, resulting in Internet routing issues.

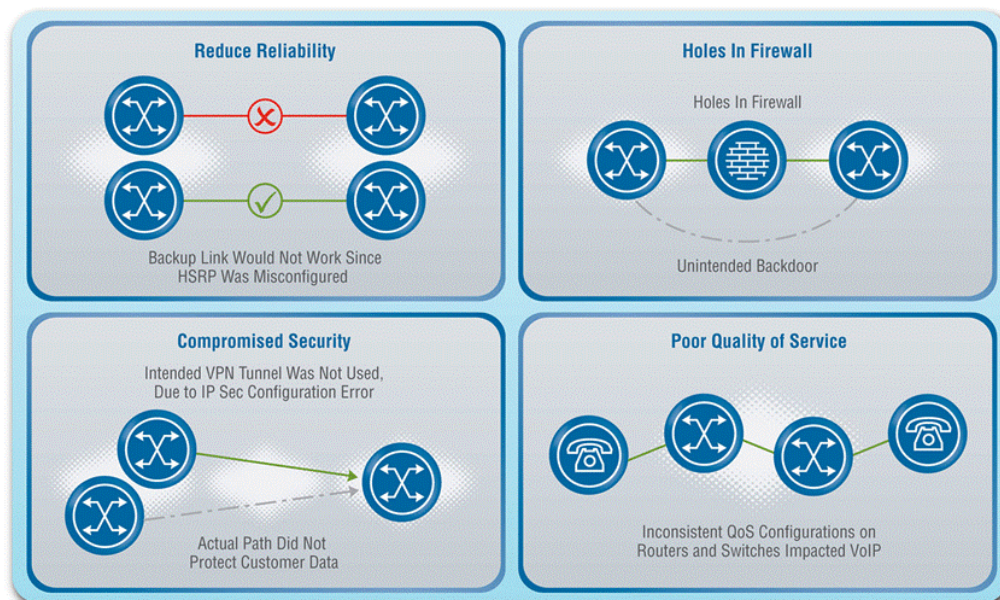


Figure 2: Configuration Errors in Operational IP Networks

2.2. Security

The most obvious configuration errors in this category can be found in firewalls, in the form of “holes” inadvertently left in firewall configurations. These “holes” are rules permitting specific application traffic to pass temporarily through the firewall, and then not removed after they were not needed. Cyber-attackers scanning enterprise networks discover these holes and craft their attacks on the

enterprise infrastructure through these holes. Apart from the obvious firewall holes, there are several other examples of errors that impact security, such as

- Static route on device does not direct application traffic into IPSec tunnel. This results in sensitive traffic remaining unprotected as it transits the network instead of flowing through the secure IPSec tunnel.
- Best-practices for Virtual LAN (VLAN) security, such as disabling dynamic-desire, and using root-guard and BPDU-guard on switch access ports, are not followed. Leaving the dynamic-desire VLAN feature enabled in a switch allows an attacker that connects to the switch to monitor all traffic passing through the switch.
- Link left active between devices. If the devices belong to network segments that are not meant to have a direct connection, then a backdoor has been introduced that can be exploited by attackers.
- Mismatched IPSec end-points. This results in sensitive traffic remaining unprotected as it transits the network instead of flowing through the secure IPSec tunnel.
- Adequate authentication is not used between devices for exchanging routing protocol information. An attacker can connect to a network device and extract or inject spurious routing information.

2.3. Quality of Service

IP traffic with demanding network latency and packet-loss rate requirements, such as Voice over IP (VoIP) and financial services applications, requires appropriate Differentiated Services and other Quality of Service (QoS) configurations in the network devices. In a large network, it is easy to make errors in the QoS configurations. Examples of such errors include

- Incorrect bandwidth or queue allocation on device interfaces for higher priority traffic. During high-load periods, higher priority traffic will not receive its due bandwidth or queue, resulting in higher latency or packet-loss.
- Inconsistent QoS policy definitions and usage across multiple devices. The same QoS policy may be implemented differently across multiple devices, resulting in application traffic receiving different treatment at the different devices, which can impact latency and packet-loss during periods of high-load.

3. Regulators expect compliance

The world's growing reliance on IP and the highly networked nature of government computing environments have also motivated a wave of regulations to improve security, reliability, and QoS.

In the United States, the Federal Information Security Management Act (FISMA) of 2002 [4] requires federal agencies to develop, document, and implement security programs. Office of Management and Budget (OMB) Circular A-130 (an implementation guideline for FISMA) [5] establishes, among other things, a minimum set of controls to be included in automated, inter-connected information resources. The National Institute of Standards and Technology (NIST) has promulgated security requirements [6] for protecting the confidentiality, integrity, and availability of federal information systems and the information handled and transmitted by those systems. NIST's "Guideline on Network Security Testing" [7] recommends that security testing be a routine part of system and network administration. It also directs organizations to verify that systems have been configured based on appropriate security mechanisms and policy. In addition, laws such as the Sarbanes-Oxley Act and Health Insurance Portability and Accountability Act, among others, are fueling the push for network protection.

Outside the U.S., organizations such as the British Standards Institute (BSI), International Organization for Standardization (ISO), and Information Technology Infrastructure Library (ITIL) recognize the complexity of IP networks and the importance of security. BSI, which published a book [8]

in 2006 entitled “Delivering and Managing Real World Network Security,” explains that networks must be protected against malicious and inadvertent attacks, and “meet the business requirements for confidentiality, integrity, availability, non-repudiation, accountability, authenticity and reliability of information and services.”

4. Many assessment approaches prove deficient

Many of the solutions for IP network configuration assessment have proven woefully inadequate. They fall generally into three categories: manual assessment, invasive systems, and non-invasive systems.

4.1. Manual assessments

IP device configurations in many organizations are large, complex software systems that depend on a hands-on, highly skilled administrator base for creation, updates, and troubleshooting. Given the size of many networks and the costs of labor, the manual approach has obvious limitations. One large U.S. federal agency, for example, has 10 five-person teams handling manual analysis of device configurations for its 120 locations.

4.2. Invasive systems

Invasive scanning solutions, such as ping, traceroute, and their commercial variants, send traffic to devices in the network and use the responses to assess compliance. Such approaches work for simple usage, such as demonstrating IP connectivity between network nodes and identifying the software version on the devices. However, when it comes to rigorous assessment, they have serious shortcomings, including:

- No root-cause analysis. They can detect problems, but offer little, if any, help in diagnosing the configuration errors that caused them.
- Non-scalable. They cannot deliver “all” or “none” results, which generally require a huge number of tests. For example, to confirm that “There is connectivity between all pairs of internal subnets,” connectivity tests are required for the number of subnets squared.
- No testing requirements on contingencies. Contingencies may be security breaches, component or link failures, changes in traffic conditions, or changes in requirements themselves. It is impractical to simulate contingencies on a network that supports real-time and critical services. For example, to detect the existence of a single point of failure, one would have to fail each device and check whether the end-to-end requirement still holds.
- Potential to disrupt network operations. Invasive scanning can introduce malware into a network, or inadvertently exploit a vulnerability that brings down a device.

4.3. Non-invasive systems

Non-invasive solutions include network simulation tools and Network Change and Configuration Management (NCCM) systems, which are analogous to software version control systems like Concurrent Versioning System (CVS) and Source Code Control System (SCCS). Such systems tend to treat configurations as “blobs,” and support IP device configuration backups, upgrades, controlled rollbacks, and maintainability of device configurations. While these capabilities are important, they are not sufficient for detecting issues that must be proactively resolved to ensure that the network continuously satisfies service requirements.

Non-invasive assessment is preferable to manual and invasive assessment because it does not impact ongoing network operations, but many of the existing systems have limitations, including:

- Individual-device assessments. Configuration management tools assess individual devices in isolation using a template-based approach, even though structural vulnerabilities are often created by interactions between protocols across multiple devices. Even non-security protocols can interact improperly to create structural vulnerabilities. For example, if redundant tunnels traverse the same physical router and that router fails, all tunnels fail.
- Non-scalable. Certain types of requirements, such as reachability, can be assessed by network simulation tools; however they can take hours to compute reachability for networks with more than 50 devices, because they simulate each and every transition of the state machine of each protocol, whether it is for routing, security, reliability, or performance.

5. Toward a solution that works

Building on our long history of involvement in assuring all types of communications networks, Telcordia has spent years researching the issues of IP network reliability and security. This research has yielded important insight into the features and functions that an effective IP network configuration assessment solution must have, and all are capabilities that are achievable today.

5.1. Desirable features of a solution

A scalable and effective solution for performing IP network assessments to detect configuration errors needs to possess the following features:

- Automatic and proactive network-wide, multi-device, and multivendor assessments against a comprehensive and updatable knowledge base that considers the network in its entirety and not just at a per-device level. The knowledge base should include rules for best current practices, regulatory compliance, and customer-specific requirements.
- Findings should visualize non-compliant rules and devices down to the “root” cause, eliminating speculation about cause.
- Non-intrusive, detailed, multilevel visualizations for physical connectivity, IP subnets, routing, VLAN, VPN, and MPLS. These visualizations, and the service reachability analysis mentioned below, can be computed using graph theory algorithms on data from the configurations.
- Service reachability analysis that visualizes path and single points-of-failures between network devices without generating traffic on the network.
- Network change impact analysis using the rules knowledge base, so new or changed configurations can be analyzed to detect errors before deployment to devices.
- Automated reconciliation of configuration and inventory information to identify and eliminate inconsistencies and errors.

5.2. Configuration extraction

Network and security administrators are generally reluctant to share IP network device configurations because they include sensitive information such as passwords and IP addresses. IP address anonymization and password obfuscation tools are of limited benefit since their usage tends to result in critical information being removed and lost from the configurations. The loss of this information makes the configuration assessments less effective. So for any configuration assessment solution to obtain complete configurations from administrators, it needs to provide assurances that their configurations will be adequately protected.

The most effective approach for acquiring configurations is for the configuration assessment solution to have direct read-only access to the IP network devices for extracting the configurations using device

vendor-supported technologies such as secure ftp or remote copy. This direct approach ensures that the most current configuration information is securely retrieved without modification by administrators, and any other device-specific data relevant for validation can also be retrieved. Another approach is to rely on backups of device configurations from a file-system. Most organizations maintain versions of their IP network device configurations on a file-system as backups, to be used to recover a device after its failure or for rolling-back configurations after an unsuccessful configuration change. The configuration assessment solution can acquire these backed-up configurations automatically, either periodically or every time new configurations appear in the backup file-system.

5.3. Configuration adaptors

Supporting the features listed in Section 5.1 requires detailed information from the device configurations. Since every IP network device vendor has their individual configuration language, software adaptors are needed that can extract the detailed information from vendor-specific format (e.g. Cisco IOS, Checkpoint, etc), and convert the information into a vendor-neutral representation. Based on our experience, the adaptors need to extract as many as 750 attributes from a single configuration to support the desired features, as compared to less than 100 that are extracted by NCCM systems.

5.4. Telcordia IP Assure

Telcordia IP Assure meets the needs for improving the security, reliability and regulatory compliance by extracting up to 750 parameter values from IP device configurations and inventory and providing the above necessary capabilities. It is scalable software that can be used to analyze and improve IP networks of all sizes, ranging from a few tens to many thousands of network devices such as routers, switches and firewalls. IP Assure provides value to both the network and the security groups in an organization. For organizations that purchase IP Assure, we can help to deploy it, train users, and customize the large, built-in rules knowledge-base to include your proprietary requirements. Once installed, users can program it to upload network data automatically, or retrieve data at any time, on demand, from the system’s web-based GUI. Enterprises that choose not to license the system can still access its capabilities as part of an as-needed consulting service. Our experts will bring the system on-site to assess your network, issue a comprehensive report, and discuss the findings. Finally, a mid-range hosted service is available for customers that want to use the software themselves but want to avoid the cost of system deployment at their location.

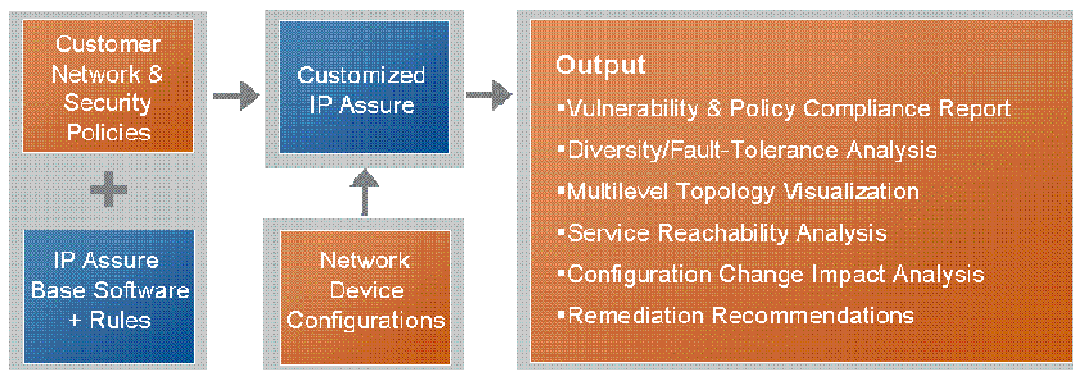


Figure 3: Telcordia IP Assure Information Flow

6. Summary

IP networks are no longer optional throughout the business and government sectors. This fact, along with the emergence of international regulations on security, reliability, and QoS, means that IP network assessment is also no longer an option. Many existing solutions on the market, including troubleshooting by skilled administrators, traffic-based vulnerability and penetration testing, NCCM software, and network simulation tools, do not (and cannot) fulfill the world's increasingly rigorous objectives. However, the technology exists today for a non-intrusive and comprehensive IP network assessment solution. Such a solution can provide auditable validation of regulations, eliminate IP network downtime caused by configuration errors, and stop the cyber-attacks that exploit those errors. Telcordia IP Assure is such a solution that is available today, further details can be obtained by visiting the product web-page at <http://www.telcordia.com/products/ip-assure>.

7. References

- [1] British Telecom/Gartner study, "Security and business continuity solutions from BT," http://www.btnet.cz/business/global/en/products/docs/28154_219475secur_bro_single.pdf.
- [2] Zeus Kerravala, "The Road to a Five-Nines Network," Enterprise Computing and Networking, Yankee (February 2004).
- [3] L. Gordon, et al, CSI/FBI Computer Crime and Security Survey, 2006.
- [4] Federal Information Security Management Act (FISMA) of 2002, <http://csrc.nist.gov/policies/FISMA-final.pdf>.
- 4) "Security of Federal Automated Information Resources," OMB Circular A-130, Appendix III http://www.whitehouse.gov/omb/circulars/a130/a130appendix_iii.html.
- 5) "Minimum Security Requirements for Federal Information and Information Systems," FIPS-200 published by NIST, <http://csrc.nist.gov/publications/fips/fips200/FIPS-200-final-march.pdf>.
- 6) "Guideline on Network Security Testing," SP800-42, published by NIST <http://csrc.nist.gov/publications/nistpubs/800-42/NIST-SP800-42.pdf>.
- 7) "Delivering and Managing Real World Network Security" published by British Standards Institute <http://www.bsi-global.com/en/Standards-and-Publications/Industry-Sectors/ICT/ICT-standards/BIP-0068/>.