



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

Securing the Management Network Domain

GSEC Practical Assignment
Version 1.4b – Option 2

© SANS Institute 2003, Author retains full rights.

Prepared By:

Donald E. Bertier, Jr.

12 May 2003

Securing the Management Network Domain

Abstract

This paper will discuss considerations for implementation of a large-scale private network for managing service delivery in a secure manner. The service provider (NetProvider) currently provides private networking and Internet services to thousands of global locations via a private backbone and the deployment of virtual routing platforms for secure separation of customer networks.

Initially, network management was not cleanly separated from the corporate network and there were multiple points of exposure to and from insecure network access points. This was complicated by dividing a large component of the network into separate pieces as part of NetProvider spinning-off from its former parent company. A plan was formulated to identify the different elements of exposure and address both device and network based scenarios. The end result is a vastly improved implementation that fulfills the following goals:

- Secure customer locations from each other
- Secure customer locations from NetProvider
- Secure NetProvider from customers
- Secure NetProvider network infrastructure
- Secure management servers

In following sections, key components of the plan will be reviewed to provide insight as to how the management network has been implemented. Additionally, topics relating to reducing exposure associated with different user access methods will be briefly addressed. The author was the primary architect of the new NetProvider management domain and was responsible for the design and implementation described herein.

Evolution from Private Networking to Network Service Provider

NetProvider began as the networking arm of a content data provider with a privately managed global network. As the Internet boom continued late into the 1990's, the business decision was made to split off the network organization & assets into a separate company, and then outsource networking services from the new company. As part of this, an Internet Service Provider with similar network architecture was acquired and merged into the parent company for inclusion in the spin-off, which formed the foundation for NetProvider as a managed network services company.

Hence, NetProvider has networking roots from two different legacy elements – private network and Internet. In common were an integrated layer-2 backbone, core routing functions and management of connectivity, including the customer

local loop circuits. Both legacy architectures relied on management functionality that was integrated into the corporate network structure, including direct access from employee desktops. Customer premises equipment (CPE) was managed in-band as part of either the IP content delivery network or directly via Internet. Out of band dial into the CPE was also deployed in most locations. Management of Internet customers was somewhat secured via the corporate firewall gateway, whereas the private network connections offered minimal security by using the content servers as a buffer between the WAN and the customer LAN segments.

The business decision to evolve into a private networking services company drove some fundamental requirements for changing the network architecture. In order to support multiple private routed networks in a scaleable manner, the deployment of a virtual routing (VR) platform was instituted. This platform enabled multiple private IP networks to be supported on the same infrastructure with full IP routing isolation between different customers. For example, the same 10.10.10.1/24 can be used in more than one private network, without any routing conflicts in any of the networks.

The VR platform also provides IP services such as Network Address Translation, traffic shaping and stateful network-based firewalls. In 2001, Pamela Warren summed the overall advantages of this approach as follows:

With a network-based model to provide security with broadband access, service providers can achieve economies of scale while saving on the costs of hardware and software installation at each customer site by serving thousands of customers from a centralized location and a limited number of devices¹.

Moving the IP services into the VR platform also simplified the CPE configuration (in most deployments) to become a layer-2 bridging device. In this manner, customer LAN segments are bridged via the WAN link to the gateway VR device at NetProvider's infrastructure Hub location (a.k.a. Point of Presence or POP).

Since NetProvider historically managed devices in-band as part of the content delivery network, it was evident that the VR deployment would include a fundamental change in WAN management philosophy. This was also driven by the need to separate management of legacy network components from the former parent company network. A new "device management" VR private network was built and a new IP addressing plan created for global support of both infrastructure and customer premise networking equipment. Each CPE device would become managed via a fully separate layer-2 connection into the device management network -- isolated from the customer traffic. A migration process to move the legacy content delivery onto the VR platform was initiated. This process also included moving all existing network elements onto the device management VR network.

The initial management network deployment was erroneously considered secure based on the belief that route-based separation within the VR platform was sufficient. However, combined with the ongoing exposure to the legacy customer network and relatively ad-hoc management practices, there were several serious vulnerabilities to be addressed.

Exposure to legacy private network segments. The migration of the legacy content delivery network onto the VR platform was expected to take more than a year. Prior to this being complete, there remained exposure of Company networking devices/systems to the legacy network. Worms such as Nimda and Code Red previously impacted the content provider network and could inadvertently impact NetProvider or its ability to service other private network customers. Although the company spin-off was proceeding, there was to be a lengthy period of coexistence while the two entities (content & network) worked through the separation process.

Direct exposure to/from NetProvider resources. For the most part, routing was enabled between corporate desktops, new device management and the legacy private network. This posed a bi-directional risk that could result in service impact to customers originating from a user desktop or could alternatively expose NetProvider users to threats from the other insecure networks. It is becoming commonly understood that a majority of security incidents – both accidental and intentional – are initiated from inside the company boundaries. In a recent article, Chet Heath reviewed five widely-publicized security incidents from 2002 and commented that: *“in all these scenarios, the silent attack on critical data actually came from inside...”*² Current events continues to show the need to protect the company from the user.

Coexistence of infrastructure and device management. The device management network services both customer devices and infrastructure management. There were no protections initially deployed to isolate core routing/switching equipment from other points in the network (including the Corporate network). Since the core equipment aggregates substantial amount of customer traffic, these are critical devices to be protected. Although customer facing network devices were not configured for routing “through” the device into the device management network, most of the CPE remain deployed in unsecured locations and could be subject to physical compromise. With some effort, the WAN and IP connectivity could be hijacked, resulting in an unobstructed path into the NetProvider networks.

Other considerations: Beyond the exposure related to the initial deployment of the device management VR network, there also remained several other legacy issues that needed to be incorporated into an overall management strategy. User management, SNMP configurations and exposure to out-of-band dial connectivity also presented considerable challenges to be addressed.

Implementation of a Management Strategy

Development and implementation of a comprehensive management strategy was required to mitigate the many risks presented by the initial deployment of the new management domain. The overall concepts of security, simplicity, scalability and supportability were of paramount importance for consideration. The resulting management strategy plan was divided into several key areas that are addressed in following sections.

Network Classification and Separation. The ability to centrally monitor and manage the network architecture is critical. A key area to be resolved was that there was no clear separation from network segments used for internal/corporate use, network management systems or the segments used to manage the network devices. Several main classifications were internally defined with regard to the management domain:

- Corporate Network – designated for all user desktops and for servers that were not used in direct support of customer devices or service delivery management. The corporate network would not require direct access to network infrastructure or customer devices.
- Production Management LANs – designated as the network segments that would house all management servers and element management systems. All management of network devices would be focused on these segments.
- Device Management PN – consists of the private network deployed via the VR platform for the management of all customer premise equipment and remote network hub LAN segments.
- Legacy Networks – designated for the remainder of the content delivery private network that had not yet been migrated to the VR platform.
- Specialized Backbone Management – consists of networks dedicated for management of backbone routing and switching equipment that cannot reside directly on the VR platform.
- Internet – self-explanatory network interface required for management of legacy internet-based equipment not migrated to the VR platform.

From the network classifications, a centralized management model focused on the “Production Mgmt LAN” was designed.

Production Management LAN Gateway Design. As part of the separation of the Corporate Network from the Production Management LANs, the proposal to have a firewall as the gateway between the two areas was evaluated. Since the Corporate Network includes all of the user desktops, it is considered highly susceptible to malware. Hence, a cost effective Linux-based firewall was deployed in default deny fashion for improved security interfacing with the Corporate Network. Most of the enabled services were for using secure shell, MS Terminal Services Client and application-specific ports for network

management applications (e.g. – Xwindows for OpenView). Less secure services including TELNET, rsh and FTP were not authorized for use on internal servers.

A subsequent consideration for the Production Management LAN design was whether to have the segment utilize a single firewall gateway vs. multiple gateways. In order to separate the user-based traffic from the actual management traffic, a multi-gateway implementation was decided upon. This fully separates the traffic required to monitor and manage the network devices from the alternate gateway for user access from the Corporate Network.

Network based IDS using SNORT was also deployed for monitoring the Production Mgmt LAN gateways to the Corporate network. For high availability, firewall devices are deployed in a redundant fashion with dual diverse connections linking the Production Mgmt LAN into the device management VR network.

Deployment of Global Firewall Templates. A major advantage to deploying the device management network on the VR platform was that it enabled the creation of stateful firewall policies that could be distributed to each interface in the device management network. Virtual routing platforms support thousands of such firewall policies in a manageable fashion, as detailed by a CoSine Communications virtual services “application note” in 2002³. Nortel, Springtide, Corona, Celox and Unisphere are additional vendors who have engaged in the virtual routing competitive landscape in recent years.

Three separate “global templates” were defined for securing management of CPE connections, Hub LAN Segments and the Production Mgmt LAN gateways. Global templates allow for a single rule set deployment to multiple connections without rebuilding the policy for each instance. The rule sets define explicit services that are enabled between a given device management point (CPE or Hub LAN) and the Production Mgmt LANs. They prohibit any peer-to-peer connectivity or access to any other networks that might become accessible via the device management routing tables. The Hub LAN Segments policy is very similar to the CPE policy, with additional rules enabling some additional applications like syslog.

All device management traffic requires access through 2 firewall policy instantiations, one at the source and the other at the destination. For example, a user TELNET session to a network device requires access via the Production Mgmt LAN rule set and then again through the CPE connection rule set. However, a TELNET session is not allowed to originate FROM the CPE connection back to the Production Mgmt LAN in either of the two rule sets. The Production Mgmt LAN rule set has a higher degree of complexity, as it addresses ingress/egress traffic to the other Production Mgmt LANs in addition to the device management segments.

The firewall deployment for the CPE connection does not fully mitigate the issues with devices being deployed at insecure locations. However, the impact of a connection-hijacking scenario is dramatically reduced. As referenced by Shara Evans in April 2000, an encrypted implementation is the best defense against such concerns⁴. However, since most CPE network devices do not support direct encryption capability, the firewall approach is viable.

IP Planning and Route Considerations. An important component of the overall management strategy is a well-defined IP Address Plan. In addition to improved manageability and operational benefits, successful IP management simplifies firewall policies and route policies on server/network equipment. The IP plan requires public space for Production Mgmt LAN segments in order to guarantee no routing conflicts in the future. The remaining device management addressing allocates unregistered/private space that did not conflict with any space currently in use for Corporate or legacy networks.

With the control capabilities provided by modern firewalls, route management is frequently overlooked as another component to “defense in depth.” Basically, the NetProvider management strategy rule of thumb is that “default routes are evil!” By eliminating default routes (and all unnecessary routes) from the routing table, exposure to unwanted traffic (accidental or intentional) is substantially reduced. Network devices are configured only with static routes applicable to the Production Mgmt LANs. As necessary, route filters are applied to CPE to stifle NetProvider route propagation into routing tables of customers.

The route management concept is also very applicable to designated servers in the Production Mgmt LAN segments. For example:

- Data collection servers that perform SNMP gets to devices are only configured with the routes to those device management customer IP ranges.
- Element Management Server for provisioning backbone switches only requires a static route to the backbone switch management segment. This server does not have routes that enable communication to or from customer premise equipment. Should unauthorized access attempt to compromise this server from any other location, it will not have the ability to respond to a direct attack.
- SNMP trap receivers and event processing servers only receive events from the network. Hence, they are not supposed to be configured with any routes pointing into the device management network.

Route management is a good mechanism to help defend against a compromised network access point, in addition to the more prevalent misconfigurations or accidents that may occur in a large networking environment.

Isolation of Legacy Network. Addressing the legacy network in the management plan was fairly straightforward once the IP Address Plan was defined. A separate routed interface was added from the Production Mgmt LANs to the legacy content delivery network and a separate firewall applied which restricted the services allowed to/from the new company Production Mgmt LANs.

Segmentation of User Base within Corporate Network. In addition to separation of the Corporate Network from the management domain and legacy network, further classification of employees was performed to designate those who require operational access from those who do not. Most corporate office locations do not include personnel responsible for deployment/operational management of service delivery, and thus do not require access to the Production Mgmt LANs. By segmenting operational personnel on different LAN segments from the rest of the company, simple firewall policies manage the applicable access to the management domain.

User Authentication and Remote Access Considerations. User management improvements were also required as part of the overall management strategy. A centralized user database was deployed in support of NIS and RADIUS based authentication. Access into the Production Mgmt LAN requires separate logon from the Corporate Network with all remote shell/rlogin capabilities disabled on servers. For remote access (dial or VPN tunnel), authorized operational personnel utilize a separate RADIUS-based password for access to the Corporate Network than the one used for access into Production Mgmt LAN equipment. This policy reduces the risk of employees inadvertently exposing critical access passwords as part of remote access from an insecure location.

Wherever supported by network equipment, RADIUS has also been enabled to eliminate local password management issues on devices. The security team is responsible for management of non-RADIUS device accounts, as designated in corporate security policy.

Another remote access consideration is the out of band dial management connection to many of the network devices or terminal servers in Network Hub locations. Port-specific password protection is enabled where applicable. These passwords, along with the other network device passwords, are changed as required by corporate policy.


SNMP Cleanup and Standardization. SNMP is used substantially throughout the management networks for data collection, event notification, provisioning and troubleshooting activities. Historically, the SNMP management on many of the devices was insecurely set to the vendor default configuration. The management strategy plan incorporated mandates for community string management and locking down of authorized management nodes to those within Production Mgmt LAN segments. Although most of these devices continue to use SNMPv1, which has known security deficiencies, the management architecture deployment

isolates the SNMP exposure to the internal Production Mgmt LAN segments and should not present substantial risk. Nonetheless, patching is performed as necessary to address significant SNMP vulnerabilities such as widely publicized in February 2002.⁵

Since many of the company tools and software utilities incorporate the use of SNMP functionality, guidelines were established to eliminate the hard coding of community strings (and other passwords) in scripts and software source code.

Implementation Results

Execution of the management strategy over an 18-month period provided substantial improvement in the overall security posture of NetProvider. Some complexities arose during the implementation due to moving the Production Mgmt LAN into a new data center while also moving the corporate technology headquarters into a new location. Upon move completion, the resulting implementation is abstracted in the diagram below:



The “Device Management Private Network” cloud spans 50+ VR nodes globally. WAN connectivity includes 75+ Infrastructure Hub LAN segments and several thousand connections to CPE. Additionally, three Production Management LAN segments have been deployed globally in support of regional management and business continuity objectives. The global Corporate Network has been fully separated from both device management and the legacy network segment.

An area that did not receive much initial attention was the server configurations in the Production Mgmt LANs. An observation by James Teel states: *“today’s enterprise networks need security that extends from the server to all its end*

*points.*⁶ Focus has recently been given to servers in the global management architecture. Corporate security works closely with the sysadmin organization to verify patch currency and to perform regular nessus vulnerability scans. The nmap utility is also commonly used for scanning the different network segments. Only designated security systems in the Production Mgmt LANs are enabled through the firewalls for scanning purposes.

The dual gateway approach in the Production Mgmt LANs has been a source of mild contention with the sysadmin team. Although specific route profiles have been provided and kept current, there continue to be occasional inconveniences due to incorrect system configurations. Compared with previous experience that incorporated dual & triple-homed servers running combinations of RIP and RIP-II for managing route tables, the static route approach remains far superior and better from a security perspective.

Some additional considerations going forward will include:

- Incorporation of RSA securid into the user authentication picture (for both remote access and certain key network devices).
- Locking down of all Production Mgmt and infrastructure Hub LAN segment switch ports to authorized MAC addresses. Ports are generally to be locked down in the current policy, however the addition of MAC level security on a per-port basis will improve the profile for the LAN segments
- Deployment of SNORT on the device management gateway of the Production Mgmt LAN segment (in addition to the one deployed at the gateway to the Corporate Network).

These items will be integrated into the ongoing security management planning objectives.

Summary

Transforming a global private network into the foundation for a managed service provider takes considerable planning, patience and persistence. Working through the primary areas of exposure and adopting a divide-and-conquer strategy can successfully lead to a highly secure network management strategy.

Areas that are often overlooked in an growing company include the fundamental separation of critical network segments from the users of the network. IP planning, route management, user security, server management and strategic firewall deployment all make up important aspects of any management architecture plan. If such a plan is formulated thoughtfully and executed thoroughly, the core principles of confidentiality, integrity and availability can be achieved.

Bibliography

- ¹ Warren, Pamela, "Security – Nortel Takes it to the Edge", DSL Prime, March 20, 2001, URL: <http://www.dslprime.com/a/nortelsec.html>
- ² Heath, Chet, "Beware the Inside Job", Ziff Davis, URL: <http://security.ziffdavis.com/article2/0.3973.929232.00.asp>
- ³ Cosine Communications 2002, "Managed Network-Based Firewall Services", URL: http://www.cosinecom.com/downloads/virtualipservices/fw_an.pdf
- ⁴ Evans, Shara, "Defusing Frame Relay/ATM Scare Stories", Telsyte Publications, April 2000, URL: http://www.telsyte.com.au/standardswatch/fr_atm_security.htm
- ⁵ CERT[®] Advisory CA-2002-03 Multiple Vulnerabilities in Many Implementations of the Simple Network Management Protocol (SNMP), February 12, 2002, URL: <http://www.cert.org/advisories/CA-2002-03.html>
- ⁶ Teel, James, "Securing the Last Unprotected Area of the Network", SC Online Magazine, November 2002, URL: <http://www.scmagazine.com/scmagazine/sc-online/2002/article/52/article.html>