



Interested in learning more about security?

SANS Institute InfoSec Reading Room


This paper is from the SANS Institute Reading Room site. Reposting is not permitted without express written permission.

Computer Security Considerations in Disaster Recovery Planning

The successful implementation of a disaster recovery plan is contingent upon the effectiveness of its design. This paper focuses on specific computer security considerations to be included in disaster planning and recovery strategies. The scope of this discussion will include the following: Neglecting to implement the proper computer security considerations into disaster recovery planning can make an already critical situation spiral into one that suffers considerable setbacks. For example, following a significant serv...

Copyright SANS Institute
Author Retains Full Rights

AD



Computer Security Considerations in Disaster Recovery Planning

GIAC Security Essentials Certification
Practical assignment V1.4b
Option1

By
Doug Velliquette

November 15, 2004

© SANS Institute 2005, Author retains full rights.

Abstract

The successful implementation of a disaster recovery plan is contingent upon the effectiveness of its design. This paper focuses on specific computer security considerations to be included in disaster planning and recovery strategies. The scope of this discussion will include the following:

- Perimeter defenses
- IDS network and host based
- Virus protection
- Patches and host configurations
- Vulnerability surveillance

Recommendations are made to include generalized aspects of each area into a disaster recovery plan. With this in mind, careful consideration to the site's computer security policy will allow for greater customization and exact detail to the individual company's disaster recovery needs. Addressing computer security, within disaster recovery planning, is vital to insuring efficient and successful recovery of operations.

Discussion

Disaster recovery closely parallels computer security operations in several functional areas. Threat evaluation, risk assessment, mitigation, and service priorities are but only a few of the items that are on the event horizon. Traditional disaster recovery procedure looks at the varying aspects of planning and implementation from an administrative perspective, focusing primarily on physical infrastructure, backup and restoration procedure, staffing, logistical operations, and connectivity. Attention to computer security must be given at all levels of recovery to ensure the integrity of the system(s).

Neglecting to implement the proper computer security considerations into disaster recovery planning can make an already critical situation spiral into one that suffers considerable setbacks. For example, following a significant service interruption, network operations may have the proper work instruction to follow for the four-hour vendor replacement program, or for the restoration of file servers from the Exabite Octopus. However, after rebuilding the first several workstations, someone restores an old system file from an unauthorized backup and introduces a Sasser variant that may have been exhaustively dealt with once before. By neglecting to include the proper mitigating factors, the disaster recovery plan has already proven inadequate in regard to secure recovery.

Perimeter defenses

Starting with a strong perimeter defense and then working toward effective defense-in-depth, border protection should be evaluated for effective recovery operations. The external firewall, border router, and VPN services all have configuration and placement considerations as the type and amount of network traffic that is allowed to traverse your company's boundaries may vary.

The firewall is a key component in the network security infrastructure. The individual site's computer security policy should refer to, or include, a firewall policy that details hardware and software specification, configuration, redundancy, physical security, etc. However, the setup and configuration of firewall services will depend on the recovery strategy that the company implements.

With this in mind, plan an approach that meets the primary service needs of the company's operations during disaster recovery. Be sure that management has signed off on the list of prioritized services and the limited functionality based wholly on those services during recovery operations. As recovery approaches "normal operations" status, the perimeter defense configuration should be back to its pre-disaster state. Develop a firewall contingency plan based on your current policy with consideration given to the recovery strategy (e.g. service priorities). Of course, recommended best practice is advised in effectiveness testing as well. Depending on the recovery strategy being implemented, the type and frequency of network traffic may change considerably. A CERT Security Improvement Module offers a good approach to testing your firewall disaster policy.

While it is theoretically possible to exhaustively test a firewall policy by generating and monitoring network traffic, it is practically infeasible. Therefore, a traffic sampling technique must be used. Two possible approaches are to capture or replay existing traffic or generate simulated traffic [1] (CERT Coordination Center)

Therefore, implement a method to capture traffic from a pre-disaster state so that the post-disaster configuration can effectively be tested with data that is relevant to the company's operation.

For the non-application specific vulnerabilities, a packet filtering border router can provide some basic protection against TCP/IP vulnerabilities. By filtering lower OSI level packets, the border or boundary router can prevent a distributed denial of services (DDoS) attack caused by a directed flood of Internet Control Message Protocol (ICMP) traffic. This high-speed packet filter will mitigate some of the risks associated with external network traffic prior to reaching the main firewall.

If there are any off-site access points to the network, chances are a virtual private network gateway is in operation. VPN access can securely connect remote offices, workers on travel, customers, and telecommuters by using the Internet as a backbone. This being the case, in a high panic disaster situation, the old adage may hold true, “out of site out of mind”. Again, remote connectivity restoration may not be first on the list of service priorities outlined by the disaster recovery plan. However, the recovery strategy should include a mechanism for notifying those individuals utilizing remote access there by limiting the burden on the customer support service during this time of crisis.

If not already in use, a VPN gateway could prove to be an ideal solution to displaced workers during the recovery period, given proper planning and implementation. A quick, cost-effective implementation of a VPN during recovery would be the use of a product that supports secure socket layer or SSL VPN. Utilizing the common desktop browser, a remote user can gain access to web servers and web-based applications through a secure tunnel after authentication, allowing for limited utilization of such resources as email or document management services. For even greater utilization of VPN capabilities during the recovery period (and a bit more security), selecting a VPN solution that implements both SSL and L2TP/IPsec would prove to be much more effective. By providing downloadable access clients, there are some VPN solutions that can offer greater access to much needed resources simply by following a link and running an installation package. VPN management primarily varies within two basic configurations. An enterprise or “site-to-site” configuration, where remote networks are connected to one another, and the remote access configuration, where users can connect via an 888 number or local Internet service provider, which may require additional client software to be loaded.

An important consideration in the “remote access through VPN” strategy is the incompatibility between network address translation (NAT) and IPsec protocols. Many home networks today are set up with gateway devices incorporating NAT. This incompatibility should be addressed when utilizing remote access via telecommuting as part of a recovery solution. One possible solution to this is presented in Thomas W. Shinder’s work on ISA VPN / Firewall configuration on Windows 2003 server [2]. Dr. Shinder provides a detailed approach on encapsulating the encrypted data within UDP headers to allow the data to pass through the NAT. This paper also serves as a tutorial for setting up the necessary packet filters on the Windows 2003 ISA Firewall/VPN.

Today there exist many appliances that serve as both firewall and VPN routers, as in the case of the Windows 2003 ISA Server. However, if the company’s border protection is already established and you are selecting to implement a VPN device, then another consideration in dealing with the VPN is the placement. Again, depending on the site’s recovery strategy, the placement of the VPN gateway may pose issues concerning vulnerabilities if not properly placed in relation to the firewall. As listed in the article *The 8 Hurdles to VPN*

Deployment by Christopher M. King [3], the following are good general VPN gateway placement rules:

- Do not compromise the overall network security policy.
- The VPN gateway placement must not be a single point of failure.
- The VPN gateway must accept only encrypted traffic from the untrusted network.
- The VPN gateway must accept non-encrypted and encrypted traffic from the trusted network.
- The VPN gateway must defend itself from Internet threats.
- The overall architecture must filter traffic after VPN decryption.

General setup and configuration issues aside, implementing a VPN gateway into the company's recovery strategy can provide a fast safe cost effective way of reestablishing connectivity to customers and displaced workers.

Now that boundary protections have been given consideration and packets begin to proliferate, it becomes crucial to employ intrusion detection techniques.

Intrusion Detection

In a dated online article from 2002, Marc Ranum [4], security consultant, aptly predicts impending threats to new systems going online.

Another thing I think is going to be really critical in the next couple of years is the new availability of mass rooters. ... So, I think we're going to see an awful lot more indiscriminate hacking, and a lot of systems that are getting compromised within minutes or, in some cases, seconds of connecting to the Internet [4].

Marc's words ring true in recalling the onslaught Kelz, Sobig, MyDoom, and Sasser brought to us in recent years. With recovery in progress, and the perimeter defenses being fortified, everyone on the recovery team is as busy as they've ever been. From a security standpoint, this can be the most critical time of the recovery process. In a partially restored network, configurations are being restored and tested and displaced workers or clients are calling in for VPN or temporary dialup access. Any possible exploit on any number of vectors could potentially jeopardize the set time frame for recovery. The advantage to initially establish and maintain a network Intrusion Detection System (IDS) will prove positive beyond any conjecture.

The Network based IDS is crucial in dealing with the known and often unknown threat vectors. A vulnerability assessment specifically designed to reflect a "worst case" recovery scenario should also be implemented into the IDS design for disaster recovery. This would include a risk assessment based on the

potential of any exploit by any threat vector during the recovery period. As recovery approaches normal operations, and transitional vulnerabilities no longer exist, standard IDS procedures for the site security policy may be resumed. However, if the site's IDS change and test procedure has not been maintained, now may be a good time to review, before a disaster is declared.

A viable combination of configurable considerations for IDS solutions is Snort and a Cisco IDS appliance. Multiple layers of security and functional redundancy are always best practice. The Cisco IDS appliance provides for equipment interoperability and dedicated hardware on the network. It also offers control configuration and some powerful data handling tools. Depending on the necessary requirements, most Cisco IDS appliances support multiple subnets and VLANs [5]. Snort, on the other hand, has several useful features that make it the choice of many IDS stewards. The first is Snort's portability. Snort can be configured to run on any number of platforms including BSD, Linux, Solaris, SunOS, IRIX, HP-UX, and Win32, to name a few [6]. Having grown out of necessity in the open source community, Snort has a large user base. Updates to signature files are usually more readily obtainable than any commercially available IDS software on the market. Of course, all of the perks of an open source project are also available (tweaks, custom configurations, etc.) to Snort users.

As the recovery effort progresses beyond setting up the "heavy guns" (NIDS), host based intrusion detection systems (HIDS) should be strategically deployed. There is high probability that someone has exploited some vulnerability along the way. For example, press coverage spreads the word that your main office has lost its primary data center, reducing your network to a make shift mix of dialup access points. Not considering the ramifications, the public relations officer is publicizing the number to call for the displaced workers to request the remote connectivity accessibility. Gaining remote access through social engineering during all of the chaos is an easy target for the wily hacker. Even though the "bad guy" is in the house, no one may know. However, host-based IDS that has been set to learn the "recovery network" in consort with proactive network IDS and firewall solutions, will aid in maintaining the integrity of your enterprise even after recovery efforts have come to fruition. Despite all of the monitoring, detection, and hardening, during the recovery process, the fortifying of your IT security operations following a major disaster will probably be the most intensive task of all.

Virus Protection

The latest onslaught of worms and trojans traversing the Internet via multifaceted vectors has been enough to make anyone's hair thin. Consideration of the appropriate virus protection scheme for the site during recovery should include both network and host-based approaches.

Being that the majority of viruses appear through email systems (on exchange servers), a product such as Scanmail by Trend Micro is an invaluable resource. By providing real-time detection and removal of viruses at the server level, all incoming mail traffic can be content filtered and assured safe for delivery to intended recipients [7]. With automatic virus definition file updates and configurations that allow the filtering of email attachments, integration of this type of product will provide a high level of protection to your network's integrity during this crucial time of recovery.

Many host-based virus protection programs have proven track records of protecting individual systems, providing that proper configuration is established to keep virus definition files current. Therefore, during host restoration, consideration should be given to properly configuring the virus protection program to receive current virus definition updates. Settings should be configured such that the virus definition file be automatically updated with a frequency that is consistent with the vendor's release schedule. With some vendors, it may be possible to centrally manage host-based virus detection. By "pushing out" virus definition files and initiating individual host scans from a centralized virus management server, the organization may save precious time in setting initial configurations. This proactive method of centrally managing host based virus protection can also provide valuable alerts and usage statistics, which can be used to improve the fortification process as you return to normal operations.

Patches and host configurations

Patch management has always been somewhat of a quandary in the network arena. In consideration of service packs, security patches, bug fixes and individual application updates, it's a good thing that a fair bit of the patching process has been web-enabled. In a perfect world, there would be one update to the host per month that would cover all of the bases. However, there exists no such animal. Patch management is as varied a process as there are Microsoft patches. There has been some help from the commercial industry in offering applications that help to track and administer patches across the enterprise. Nonetheless, it is still quite a job.

Carnegie Mellon's Computer Emergency Response Team (CERT) reported 3,784 computer vulnerabilities in 2003 [8]. That's an average of about 15 every working day which is a lot to read about, much less deal with.

It is possible to simplify patch management by implementing such a tool as Shavlik's HFNetChk. This software utilizes patch scanning to evaluate systems on the network for most recent patch installations [9]. If systems are found to be un-patched, then the patch server pushes out the most recent patch(s) to the delinquent host. Such a solution can provide a cost effective approach to the daunting task of patch management while eliminating reliance on the individual system administrators to stay current. The implementation of a patch test

procedure for an enterprise-wide patch management strategy is strongly encouraged. Patch implementation at the operating system level can result in the introduction of new vulnerabilities. Duplicate platforms can be maintained in a secure stand-alone environment while new patches and security configurations are tested for the potential presence of newly introduced vulnerabilities.

Host images should be maintained with the most current configurations for individual service priorities and stored with off-site backups. In the event that many services (departments) are lost during disaster, a highly customized “base” image will allow for an expedited restoration of systems. This is an effective way to ensure proper security policy is set on new hosts coming online during recovery. Using the analysis of consequence of loss based on the levels of confidentiality, integrity, and availability for the data each system processes, specific security requirements can be addressed in the system’s local security policy. For instance, in a Windows 2000 system, the local security policy can be customized for account access, passwords, auditing, and user rights. The company’s computer security program plan should already have these defined based on the appropriate protection level and operating platforms in use. The alternative to this is the grouping of systems requiring similar protection levels into separate virtual local area networks (VLANS). In an Active Directory environment individual Organizational Units (OU) can be administered to separately. The benefit of this type of segregation is the management of security policies by Group Policy Object (GPO) or Desktop Authority [10]. Domain administration of security policy ensures that the appropriate settings are made to the local system upon login.

Keep in mind that during the recovery process, the time that it takes to patch the varied hosts may save the company a lot of incident recovery work down the road. Also, having the foresight into pre-planning for host security policy and configuration will ensure successful mitigating efforts in protecting vital information. Another volley in support of effective initial patch and configuration management is the consideration of recovery priorities. The business rules that are written into the site’s disaster recovery plan may give credence to the recovery of a “higher priority” function (e.g. the finance department). Therefore, given the low visibility of computer security related tasks, the computer security department may not have the vulnerability-scanning engine in place for some time. Nor will they have the time it takes to scan, review, notify, and patch whatever holes are found. It would, therefore, be in the best interest to fortify the highest priority services through patching and configuration controls, and then work toward establishing the vulnerability surveillance system.

Vulnerability surveillance

Although many of these security functions can be conducted in parallel, assessing vulnerabilities after recovery has reached near completion will make

best use of valuable time. Within the scope of vulnerability scanning, the majority of the time will be devoted to the repetitive process of configuring, scanning, and analysis, followed by the repair of security holes. Now that the defense perimeter is up, intrusion detection is in place, precautions are being taken against past and future viruses and worms, and patch and system configuration management is adequately addressed, it should be stressed that this stage of your defense-in-depth will be the most time consuming.

While there are plenty of COTS (commercial off-the-shelf) packages that are available, one industry proven security scanning application that is highly configurable and free is Nessus. Nessus provides the security professional with a complete set of scanning options that utilizes one of the most comprehensive vulnerability databases available. Some key features to Nessus are scalability, SSL and PKI encryption support, smart service recognition, and open source community development and support [11].

Vulnerability and threat analysis could not be mentioned without consideration being given to Internet Security Systems (ISS). ISS is the current commercial leader in vulnerability analysis, providing detailed products and services designed to provide a full range of security solutions [12]. While most of the features found in one scanning application are provided for in others, running ISS in conjunction with Nessus will ensure greater coverage in assessing the vulnerability of the entire enterprise. While there exists excessive overlap in the analysis capability between the two, there are those few exploits that one may identify over the other. This results in a highly comprehensive and redundant approach, allowing for multiple scanning engines to mitigate the risks of a compilation of vulnerabilities, while enforcing your site's security policies.

Through vulnerability surveillance, the many elements involved in an effective computer security program are well substantiated. By considering all past vulnerabilities as well as the daily onslaught of new exploits, vulnerability scanning and analysis activities during disaster recovery enable quality control and continuous improvement of the computer security program.

Conclusion

Careful consideration of disaster recovery planning in the areas of perimeter defense, IDS, virus protection, patches, host configurations, vulnerability scanning and remediation will enable security personnel to effectively prepare for recovery operations while keeping the foremost threat mitigating practices in mind. This effective defense-in-depth approach to the disaster recovery planning process will ensure that best practices for computer security are not overshadowed or even worse, ignored during recovery operations.

References

- [1] "Acquire Firewall Hardware and Software." CERT Coordination center, Carnegie Mellon Software Engineering Institute. July 1, 1999. <http://www.cert.org/security-improvement/practices/p054.html> (July 10, 2004)
- [2] Shinder, Thomas W. "Configuring Windows Server 2003-based ISA Server Firewall/VPN Server to Accept inbound NAT-T L2TP/IPSec Calls." ISA Server. Org. July 22, 2004. <http://www.isaserver.org/tutorials/natt2003.html> (July 25, 2004).
- [3] King, Christopher M. "The 8 Hurdles to VPN Deployment." Information Security. 1999. <http://infosecuritymag.techtarget.com/articles/1999/vpn.shtml> (July 26, 2004).
- [4] Ranum, Marcus J. "IDS at the crossroads." Information Security. June 2002. <http://www.infosecuritymag.com/2002/jun/cover.shtml> (July 26, 2004).
- [5] Cisco intrusion detection. 1992-2004. Cisco Systems, Inc. <http://www.cisco.com/warp/public/cc/pd/sqsw/sqidsz/index.shtml> (August 10, 2004).
- [6] Caswell, Brian. Roesch, Marty. "About Snort." Snort.org. August 15, 2004 <http://www.snort.org/about.html> (August 15, 2004)
- [7] ScanMail for Microsoft Exchange. 1989-2004. Trend Micro, Inc. <http://www.trendmicro.com/en/products/email/smex/evaluate/overview.htm> (September 6, 2004)
- [8] "CERT/CC Statistics 1988-2004." CERT Coordination center, Carnegie Mellon Software Engineering Institute. October 19, 2004. http://www.cert.org/stats/cert_stats.html (October 19,2004)
- [9] Shavlik HFNetChk Pro. 1997-2004. Shavlik Technologies, LLC. http://www.shavlik.com/hfn_windows.aspx (October 19, 2004)
- [10] Cavalancia, Nick. "Desktop Authority and AD Group Policy Objects." Script Logic. August 2004. <http://www.scriptlogic.com/whitepapers/davsad.pdf> (October 25, 2004)
- [11] Deraison, Renaud. "Features." Nessus. 2000-2004. <http://www.nessus.org/features.html> (October 27, 2004)

[12] System Scanner. 1994-2004. Internet Security Systems, Inc.
http://www.iss.net/products_services/enterprise_protection/vulnerability_assessment/scanner_system.php (October 27, 2004)

© SANS Institute 2005, Author retains full rights.



Upcoming SANS Training

[Click Here for a full list of all Upcoming SANS Events by Location](#)

Hong Kong Advanced Forensics Seminar	Hong Kong, Hong Kong	Nov 09, 2009 - Nov 14, 2009	Live Event
SANS Sydney 2009	Sydney, Australia	Nov 09, 2009 - Nov 14, 2009	Live Event
SANS Vancouver 2009	Vancouver,	Nov 14, 2009 - Nov 19, 2009	Live Event
SecurityByte 2009	New Delhi, India	Nov 17, 2009 - Nov 20, 2009	Live Event
SANS Geneva CISSP at HEG 2009 Autumn	Geneva, Switzerland	Nov 23, 2009 - Nov 28, 2009	Live Event
SANS London 2009	London, United Kingdom	Nov 28, 2009 - Dec 06, 2009	Live Event
SANS WhatWorks in Incident Detection Summit 2009	Washington, DC	Dec 09, 2009 - Dec 10, 2009	Live Event
SANS CDI East 2009	Washington, DC	Dec 11, 2009 - Dec 18, 2009	Live Event
SANS WhatWorks in Data Leakage Prevention and Encryption Summit 2010	New Orleans, LA	Jan 07, 2010 - Jan 12, 2010	Live Event
SANS Security East 2010	New Orleans, LA	Jan 10, 2010 - Jan 18, 2010	Live Event
SANS AppSec 2010 and WhatWorks in AppSec Summit	San Francisco, CA	Jan 29, 2010 - Feb 05, 2010	Live Event
SANS San Francisco 2009	OnlineCA	Nov 09, 2009 - Nov 14, 2009	Live Event
SANS OnDemand	Books & MP3s Only	Anytime	Self Paced