



SANS Institute

Information Security Reading Room

Case Study in Developing Fault Tolerant and Highly Available Systems with Secure Zones of Protection

Kevin Knox

Copyright SANS Institute 2021. Author Retains Full Rights.

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

GSEC v1.4b Option 2.

Case Study in Developing Fault Tolerant and Highly Available Systems with Secure Zones of Protection

I. Abstract

Process Control is the part of a company that controls the critical processes that company operations are dependent upon. It is part of the critical infrastructure of the company and the clients that it serves. Various parameters, status and measured values are constantly queried via Supervisory Control and Data Acquisition, (SCADA), to control the process.


Legacy mainframe systems that housed the Process Control System, (PCS), became too cumbersome and expensive to maintain prompting the move to a distributed 24X7 architecture. The distribution of process control, monitoring and alerting functions to various Unix and Windows servers via network connected devices forced us to realize that we were no longer isolated from the "world" and that securing our networks became objective number one.

This paper will discuss the processes and actions taken to provide 24X7 fault tolerant and highly available systems with physical as well as cyber security in the forefront.

II. Before

Prior to year 2000, Process Control performed process management with the aid of mainframe computers using Thicknet and a proprietary protocol called Network Operating System, (NOS). All interfaces to the mainframes for programmers and developers were hard wired RS-232 connections. Operators of the PCS were directly connected to the private network via Thicknet, 50 ohm coaxial cable, and all operator interfaces to the PCS occurred over these tapped connections. The private network design was of bus topology that made troubleshooting difficult. All remote connections from the remote terminal units housing our SCADA system were hard wired using directly connected phone lines and microwave connections over our private phone system. Process Control remained isolated from the rapidly developing corporate intranet/internet infrastructure and enjoyed physically isolated networks with no external interfaces. See Drawing #1.


Drawing #1



Access to information assets within the PCS was restricted physically and external threats were only of the form of natural disasters, operator error and equipment failures. The building structure housing Process Control consists of a “bunker” with physically protected redundant computer rooms, UPS, dual power feeds, environmental control and fire safety systems. Each isolated computer room is a mirror image of the other with dual mainframes, networks and device interfaces. The risk, initial vulnerability assessment and threat exposure were evaluated years previous to my involvement with Process Control. The common thread tying a PCS legacy system to a system utilizing distributed architecture is that the critical infrastructure is inviolable and that the same level of confidentiality, integrity and availability must exist in the old as well as the new.

The corporate demand for more and more data in the form of statuses, analog values, load forecasting, frequency control and much more came to the forefront. Our engineers, developers and programmers designed methods to present PCS data using the new technologies of the time, web based technologies employing HTML, Java, Perl and relational databases. This drove our move toward open systems and network development for

interfacing with our corporate architecture. It was at this time that we began to design our interfaces to the PCS using the “Star” physical network topology. This topology lent itself well to our desire to link our Open Data System, (ODS), with the legacy mainframe such that clients could gain access to PCS data through our ODS portal. See Drawing #2.




My predecessor and fellow sysadmin was instrumental in the initial design of the ODS. I became involved during this phase of the ODS design and became responsible for the HACMP design, failover capability, system administration, backup/recovery, Oracle DBA and network administration of the ODS system.


In Chapter 1 of the [HACMP for AIX Planning Guide](http://publib.boulder.ibm.com/pseries/hacmp/docs/431planning.pdf) located at: <http://publib.boulder.ibm.com/pseries/hacmp/docs/431planning.pdf>, the ultimate goal of HACMP is to eliminate single points of failure. The elimination of single points of failure in the PCS and ODS is necessary to provide the integrity and availability required of these environments. It

became clear that HACMP would make our servers, databases, applications and ethernets highly available. It also became clear that this high availability concept needed to be escalated to all aspects of our infrastructure. A single network deployed on a single switch defeated the concept of HA because a single point of failure was now introduced into the network.

The primary components of the ODS consists of dual servers running Unix and HACMP with SSA using mirrored disks. See Drawing #2. We would “push” data from the legacy mainframe to the ODS through a couple of redundant 802.3 switches connecting the PCS Thicknet and the ODS HA ethernet interfaces. The corporate world would interface via a FDDI ring to the ODS and data was presented to our clients via a web server. Redundant firewalls were deployed in an attempt to protect the PCS from our external entities whom we viewed as hostile. The firewalls also gave us the security, access control, fault tolerance and availability that we desired. See Drawing #3. Notice the “backdoor” from the corporate through the ODS to the PCS. An intruder gaining access to the ODS would have access to the PCS if the intruder could figure out the username/password combinations that were typically “Smoking Joes”. Of course, the NOS and proprietary mainframe language were deterrents through obscurity, but the backdoor still existed. As we all know, the illusion of security is no security whatsoever. These FDDI links were quickly broken and new internal networks were deployed behind the firewall via redundant level 3 routers. See Drawing #3A.



As the replacement and maintenance costs of the mainframe skyrocketed, Process Control's design engineers selected a PCS vendor to replace the legacy system with a distributed system consisting of approximately 50 servers. Various PCS applications were deployed across these servers and networked using cascaded switches for speed, redundancy and availability. Each server is dual NIC'd with a primary and backup network. Highly available datalinks were re-established for porting data from the PCS to the ODS. See Drawing # 3A




As more and more corporate and external entities became dependent on our data, our networks began to grow. It was at this time that a network consultant was employed to aid in the design and development of our networks. I worked directly with the consultant to define our critical network segments and aid in developing the network architecture to be secure, fault tolerant and highly available. These design issues were presented to management for approval. Our general design philosophy

has been to provide redundant systems as long as the safety, business or economic needs could be proven. Management agreed.

III. During

Prior to the arrival of our network consultant, I aided in the development of a network drawing and identified our critical networks and systems. I identified the single points of failure and helped to develop a logical graphic depicting our zones of protection necessary for Defense in Depth and where our critical data resided. See Drawing #4.

Drawing #4



Drawing #4
6/23/03
K.D.Knox
GSEC Practical

Drawing #4 identifies our most valuable asset as being the Core PCS Data. It also depicts the relationship our data has with the surrounding elements, i.e.; the rest of the hostile world.

Upon the arrival of our network consultant, we further developed the logical as well as the physical components of our networks that needed to provide redundancy and fault tolerance. We identified the following critical components needing 24X7 operations and very high security:

- Process Control System.
- Open Data Systems.(Data Presentation & Client Access)
- Inter-Control-Center-Protocol.

In addition to the above, the following secondary systems were identified:

- Load Forecasting System.
- Backup and Recovery System.
- Vendor Remote Access System.
- Development and Application Programmers.

During this process it became evident that our clients had become highly dependent on the data that we provided. We quickly determined that any prolonged network outages would not be acceptable to our systems.

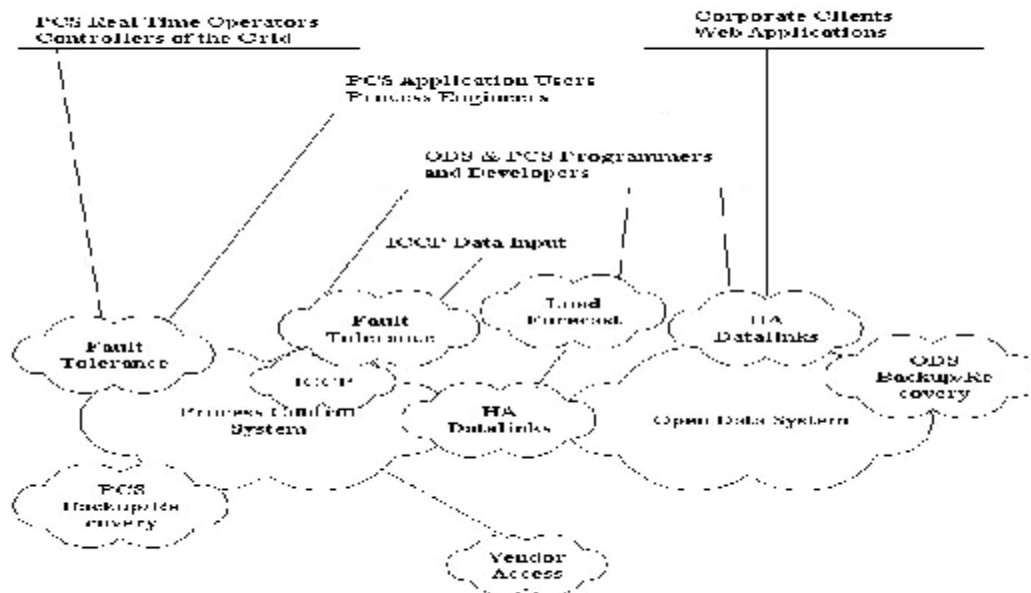
With the aid of our consultant, I helped design the network infrastructure which would provide 24X7 fault tolerant and highly available systems. The concept of Defense in Depth using network IDS, sniffers and the firewall rules base also came to the forefront in our design.

Our network consultant and I designed a basic network model that took into account the criticality of the component and application relative to maintaining 24X7 operational integrity. We decided that segmenting the networks by function and/or application provided the security and access control that we desired. The key elements incorporated into our design model to provide security, availability, extensibility and robustness were answering several key questions:

- How critical is the application?
- How critical is the data?
- Who needs access to what server?

See Drawing #4A for the logical separation of function and access control basic to our model.

Drawing #4A



Drawing #4A
6/23/03
K.D.Krus
6451001: Protected

One of the initial security steps taken was to sever the backdoor connection through the FDDI ring, to the ODS and onto the PCS. I severed the FDDI connections and established network segments via our redundant firewalls, routers and switches. These network segments were developed by application and/or function and consisted of the following:

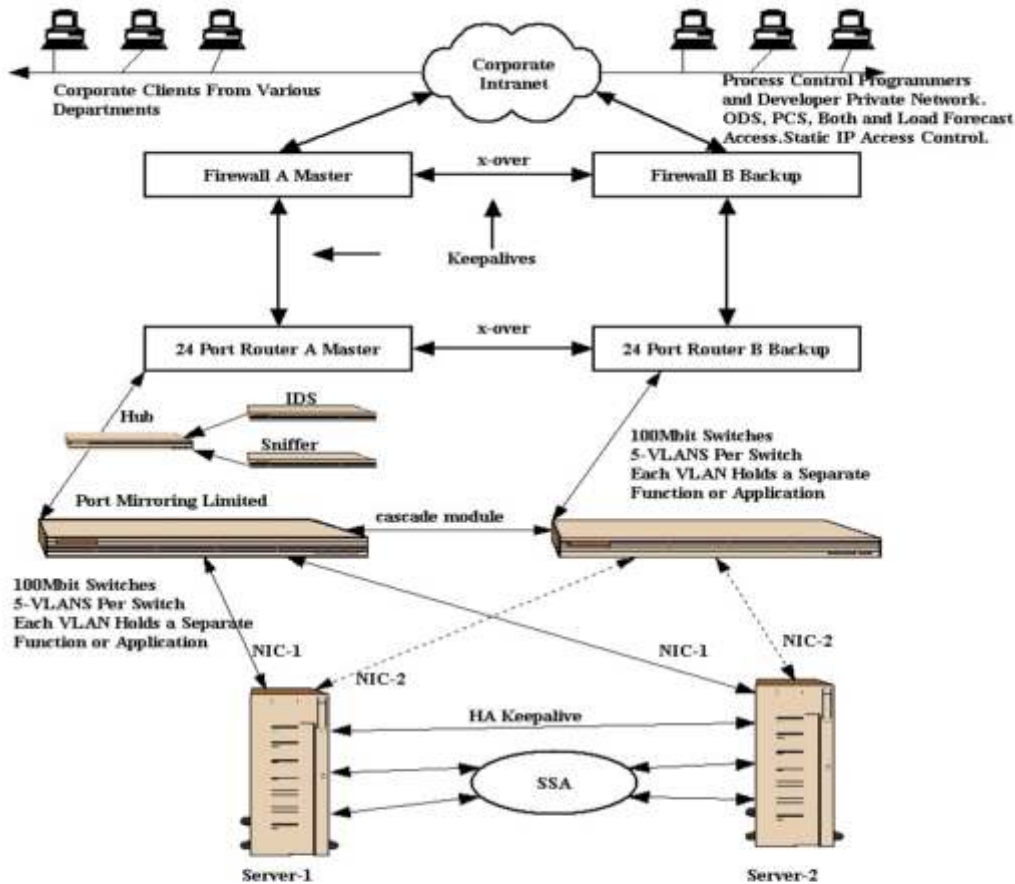
- Web Apps(Clients & Data Access)
- ICCP
- Backup & Recovery
- PCS
- Development
- Load Forecast
- Vendor Access

I began to design our VLANs to accommodate enough ports for network IDS, sniffer and additional server connectivity. I built multiple VLANs on redundant switches to reduce cost and separate networks from each other providing a measure of network security by isolating traffic based on application, data and access.

Drawing #5 is the basic structure I helped design for our networks, which provides the necessary levels of fault tolerance and high availability required by our department. Each component of the system can be

modified according to its criticality. In other words, based on my network template, some systems can have redundant networks, such as the ODS and PCS, while other systems do not need that capability such as the Backup and Recovery system or the Load Forecast system. Also, certain applications can be made highly available while other applications need fault tolerance or simply manual failover.

Drawing #5



Drawing #5
Basic Network Segment Model
for HA and Fault Tolerance

Drawing by K.D.Knox 06/03/03

I was able to identify the who, what and why security questions concerning access. I worked with our Information Technology group and the Network Design Manual --- IP 101: All About IP Addresses to segment our external network, i.e.; the department network outside the process control firewall. A single static segment and 3 dynamically assigned segments were defined. This design manual is located at the following URL:

<http://www.networkcomputing.com/netdesign/ip101b.html>

The static and dynamic segments were defined as follows:

- 10.50.50.0 netmask 255.255.255.192.

The first subnet begins at 10.50.50.0 with the potential to have 64 host addresses, 0 through 63. With the first and last address reserved, this allowed for defining 62 different hosts. This segment was removed from the DHCP server by I.T. and designated as statically assigned. The remaining subsegments of our Class C network began at 10.50.50.64, 10.50.50.128 and 10.50.50.192 with each having the ability to have 62 dynamically assigned hosts. This provided the margin for growth necessary.

Specific users such as developers were given static IP addresses and user accounts on the ODS and PCS such that their access was restricted via access control lists on the servers, firewall rules and username/password combinations. These static IP addresses became entries in our firewall rule base allowing them the secure shell service only. This limited their ability to run any other service and eliminated the "clear text" username/password problems associated with telnet, ftp and xterm. . My policy was to deny all and only allow specific IP addresses and specific port access. Other users such as our Load Forecast engineers were given access to the Load Forecast application. The application was segregated from other systems and networks and again, firewall rules were implemented to allow the users access based on my access model. These users were isolated onto a different network segment than our core trusted users. This allowed me to build the firewall rule base such that 5 groups of users were defined; PCS access, ODS access, corporate access, Load Forecast access and both PCS and ODS. Those clients not needing access, but resident on the same network segment as the core access users were limited by having dynamic IP addresses assigned. These IP addresses were not allowed through the firewall. For example, the 10.50.50.0/24 network was defined with a netmask of 255.255.255.192. This limited the access to only the first 62 IP addresses, 10.50.50.1 through 10.50.50.62. These static IP addresses were given to our most secure and trusted users. The remaining portions of the host addresses were designated as DHCP and those addresses were denied access by our firewall rule base. Our corporate intranet clients were only given access in the firewall rule base to port 80 of our web server and other socket pairs necessary for our Java Applets.

The fault tolerant component of the system resides in the hardware. Refer to Drawing #5. If a switch, router or firewall fails then connectivity is almost instantly re-established between the client and server, such that the failed switch, router or firewall is seamless to the client or server.

The highly available component of the design resides in the SSA and HACMP design on the servers. In the event of a server failure, the corresponding cluster component is in “mutual takeover” configuration where the failover mechanism requires the adjacent server to start the corresponding application services and resources of the companion server. This usually will take 5 minutes to mount the shared disks, start the databases, reconfigure the IP addresses and start the associated processes. The majority of the failover is handled by HACMP but my predecessor and I developed Perl and Korn shell scripts to start and stop the various applications and databases such as Oracle. These were deployed on each server within the cluster.

In a like manner, the failure of an individual server’s service NIC will be recognized by the configuration and HACMP to be a “network down” event and force an IP address takeover of the standby NIC, resulting in restored functionality of the requisite services. In a like manner, the failure of a switch will be sensed by HA and cause an IP address takeover of the service NIC by the standby NIC connected to the corresponding redundant switch. This action maintains the integrity and availability of our systems.

As the network rollout wound down I began to focus my attention on other security initiatives. Using the ISS Scanner tool, I scanned our ODS network for TCP and UDP vulnerabilities and discovered much vulnerability in our ODS cluster. The tool discovered snmp, smtp, sendmail, r-utilities, Smoking Joes and other vulnerabilities. Each of these vulnerabilities were addressed and removed in subsequent weeks. Naturally, the elimination of these vulnerabilities progressed as the testing and implementations were allowed in a Real Time and Near Real Time environment.

In the O’Reilly book by Simson Garfinkel and Gene Spafford, Practical Unix & Internet Security, steps were taken to harden our systems based on the following guidelines:

- Eliminate Unnecessary Services.
- Implement ssh to eliminate “Trust” relationships and “clear text” login information.
- Place Restrictions on User Passwords
- Remove Setuid and Setgid directories and files.
- Eliminate Root Login Remotely
- Eliminate Generic Accounts
- Setup the ftpusers file.
- Eliminate booting from cdrom without a password
- Remove the Telnetd, snmp, sendmail and the r-scripts.

Part of the environment a PCS has always maintained is for it's systems to be available, of high integrity and that the data is assured to be accurate and not questionable. By designing our networks with redundant and highly available servers, switches, NICs, disks, networks, routers, access controls and firewalls, we were able to maintain the fundamental components of secure systems, confidentiality, integrity and availability, necessary for the 24X7 operation required of a PCS.

The next step in achieving our vision of Defense in Depth was to implement IDS. Working with our 3rd party vendor we selected a network IDS solution.

I placed network sensors at the junction of each network segment and it's connection to the router. See Drawing #5. The default Attack policy was deployed and each network sensor was tuned to control "false positives". I deployed the Attack Policy on the network sensors based on the Stephen Northcutt book, Network Intrusion Detection, An Analysts Guide and the recommendation of our network consultant. An intruder will typically probe a site over time before an attack actually occurs. The Attack Policy specifically searches for different types of portscans that are precursors to an actual attack. I setup alert paging through the Management Console with the aid of our network consultant using the shell scripting language readily available in Unix. Paging was activated for all alerts captured by the IDS.

As our security initiative progressed, further hardening of our physical facilities to provide Defense in Depth was instituted by my managing engineer. Access control to the computer rooms was established using Smart cards. Video monitoring of the computer rooms was also established with central logging and video retention. Logging mechanisms were put in place to document and control all individuals entering and leaving. 24x7 processes were developed to control physical access to our computer rooms.

With all the activity involved in rolling out networks in Real Time and Near Real Time systems and keeping connectivity established for our production systems, I began to tune our firewall rule base. As I became familiar with the rule base and added more and more rules tuning our access control, I discovered a couple interesting rules which were activated in our initial firewall deployment. These rules were necessary during the implementation to expedite connectivity without denying service and allow connection troubleshooting.

The following rules were uncovered:

- Net 172.0.0.0 Any Source Any Destination Any Service

- Net 10.0.0.0 Any Source Any Destination Any Service

These rules became to be known as the “Big Bad Rules” and were quickly deleted and tuned. These rules effectively left our networks wide open to multiple vulnerabilities and exploits from external entities. My concern was that a worm, virus or trojan infecting our external entities could potentially find it’s way through our firewall and into our networks.

One of the interesting and dangerous lessons learned when I detected these rules was that other subsequent rules were rendered ineffective by the Big Bad Rules. Case in point. I added the following PCS firewall rule to access our web server at 10.30.30.7:

- Source 10.50.40.5 Destination 10.30.30.7 Service http Action allow.

This rule was below the Big Bad Rule in the rule base and because of the hierarchy associated with the rule base, the rule had absolutely no effect. I came to find out that as I made the rule base more and more precise and granular and the Big Bad Rules were deleted, clients and users alike began losing portions of their connectivity. It became necessary to monitor the logs over time and activate only those services necessary to provide the appropriate port access for our specific clients.

The outcome of this exercise and further research exhibited that firewalls and IDS are key components to Defense in Depth, but that extra effort needed to be put into the security of our various applications. The below listed link at the e-Eye Security web site further explores this need.

<http://www.eeye.com/html/Research/Papers/DS20010322.html>

Our group is presently evaluating an application firewall and scanner for our PCS and ODS applications. It is important to note that our engineers recognized the pattern of hackers targeting specific applications and that action needed to be taken. Our understanding is that an evolution is occurring in security products and that more needs to be undertaken to guarantee the availability and integrity of our internally developed applications.

Let me address the issue of fault tolerance surrounding the redundant firewall and routers. I aided in this design with our network consultant and our firewall administrator. In Drawing #5, you see the Master configuration in that Firewall A and Router A are the Master controllers. Firewall B and Router B are the hot-standby or Backup devices.

Keepalive connections have been established between Firewall A and B and Firewall A and Router A.

The failure of Firewall A will activate Firewall B as Master because the firewalls have lost their keepalive connection. This is seamless to our networks. The appropriate personnel are notified via paging when this occurs. The recovery consists of restoring Firewall A and manually bringing Firewall A back to Master.

Now a failure of Router A will escalate to a Firewall A failure via the keepalives between Firewall A and Router A. This brings Firewall B to Master and Router B to Master via Virtual Router Redundancy Protocol, vrrp, between the two routers. A static route is employed on Router A and B to Firewall A which also uses vrrp. If the interface on Firewall A is 10.50.50.1 and Firewall B takes over that address and Router B becomes Master, then Router B will pass traffic to Firewall B because of the static route. The recovery is to restore Router A which is designed to always be Master by it's priority level being maximum in relationship to Router B having a lesser priority level. When Router A re-establishes its Master status it becomes necessary to manually bring Firewall A to Master. A problem has reared its ugly head in this situation. In the event of a simple power on reset or glitch where Router A resets and comes back up Master after re-boot, Firewall A has failed to B, Router B returns to Backup when Router A becomes Master and manual intervention becomes necessary to bring Firewall A backup to Master. I am presently working with the manufacturer to address this issue.

This design keeps our connectivity fault tolerant and available for our 24x7 operations. It provides the integrity and availability necessary for our operations. The segregation of the network segments maintains the confidentiality, security and access control to protect highly critical segments from the less critical segments. This segmentation also provides an added component to our Defense in Depth strategy.

IV. After

The security of our systems became paramount after the events of September, 2001. Management "buy in" to our security initiatives was instrumental in acquiring the necessary resources to harden our systems and follow the Defense in Depth strategy. The following list details some of our accomplishments, but we realize that the process is circular and ongoing:

- We have recognized many of our vulnerabilities and management is fully behind our endeavors to protect our confidentiality, integrity and availability.
- Fault tolerant and highly available systems have been developed.
- Password policy has been established.

- Servers have been hardened.
- Firewall Rule Base has been hardened and consolidated.
- Access policy has been established.
- Defense in Depth has been established and will continue to be addressed.
- Process control has been implemented to address IDS, firewall and server alerts.
- 24x7 on-call specialist respond when alerted and notify the response team if necessary.
- Continuous and automated monitoring of critical networks are in place.

The security initiative taken by Process Control after September 2001 was and is multi-pronged. The physical, cyber and intellectual components of security came forward and have been addressed. The culture and awareness of security within our group has been implemented. Programmers, developers, sysadmins, clients and operators alike have been indoctrinated to the needs for security.

In a small shop like Process Control, it becomes necessary to handle multiple tasks and employ help when necessary. The team effort necessary to move from legacy systems to distributed architecture was phenomenal. By employing our network consultant and working side by side, I was able to acquire additional knowledge concerning networks and network devices. I was introduced to host and network based IDS along with the firewall rule base. Before fall 2001, IDS and the firewall rules base were something someone else did. Since then, I have become directly tied to those functions.

Defense in Depth does not simply mean that one throw money at a problem. It also implies that personnel within a department or organization understand the threats, risks and vulnerabilities associated with their functions or processes. Staying up to date with the latest developments in technology becomes critical.

Highly available and fault tolerant systems add significantly to the complexity of networks, servers and applications. These systems are set up with specific methods and simple changes can destroy the integrity of the systems. It pays to keep a trained and knowledgeable eye on the system as a whole, maintain "Change Management Policy" and test, test and test some more.

V. Impact:

This paper addressed the issue of fault tolerance and high availability with secure zones of protection. I touched upon the multi-faceted security

initiatives undertaken by Process Control. I explained my actions and activities in providing Defense in Depth. The bottom line is that all the aforementioned security initiatives aided in providing highly available and fault tolerant systems. By hardening our clusters, we removed certain threats and vulnerabilities that had the potential to make our systems unavailable. By implementing IDS with 24x7 alerting and response mechanisms we increased the availability of our systems. By fine tuning our firewall rules and segmenting our networks by application and function we reduced the potential that an outage of a single system would affect the whole. The physical measures in place and further enhanced, served to provide assurances that rogue or disgruntled employees or contractors would be identified and documented for prosecution if that became necessary. These physical deterrents aid in maintaining the system integrity and availability.

I work in a process control environment. Security is a process to be controlled. The Process Control group will control the security of our installation. Security is a closed loop process where threat, vulnerability and risk assessments will be continually made and action taken.

VI. References:

1. IBM. HACMP for AIX Planning Guide. 2nd Edition Version 4.3.1. SC23-4277-01. Jul. 1999. URL:
<http://publib.boulder.ibm.com/pseries/hacmp/docs/431planning.pdf>
2. Lewis, Chris. "*Network Design Manual --- IP 101: All About IP Addresses*". URL:
<http://www.networkcomputing.com/netdesign/ip101b.html>
3. Garfinkel, Simson, Spafford, Gene. Practical Unix & Internet Security 2nd Edition. Cambridge: O'Reilly & Associates, Incorporated, Apr. 1996.
4. Northcutt, Stephen. Network Intrusion Detection, An Analysts Handbook. Indianapolis: New Riders Publishing, Jun. 99.
5. eEye Digital Security Web Site Permech, Ryan. "The Use of Application Specific Security Measures in a Modern Computing Environment" URL:
<http://www.eeye.com/html/Research/Papers/DS20010322.html>