Global Information Assurance Certification Paper

GIAC Security Essentials Certification (GSEC)
Practical Assignment
Version 1.4 (amended April 8, 2002)

Process Re-engineering
A case study in improving confidential report distribution

By
Michael Toscano

Process Re-engineering
A case study in improving confidential report distribution

ABSTRACT

Every medium to large business employs Finance or Accounting departments. These departments are responsible for all of the company's finances and revenue. They track payments and collections as well as forecast the fiscal outlook for the company. Furthermore one of their main duties is to tabulate various pieces of the financial data and distribute it to the management of the company so that they may use that information to choose the direction of the company.

I work at a law firm that handles primarily business law. We have a moderate sized finance and accounting departments that handle the monetary transactions of the firm. They also distribute monthly, quarterly and annual reports to the firm's members and directors. These reports are distributed in a paper format and utilize enormous staff resources to generate. As an IT professional I am always looking at how I can apply my talents and skills to best help the organization. In this case I focused on how I can re-engineer the process for distributing these reports to make it easier for both the report recipient as well as the staff of finance and accounting departments. However no matter how much better I can re-engineer the process, I cannot sacrifice anything on security. In fact I want to increase the existing security of the distribution process. This paper is a case study that reviews the design process used to re-engineer the report distribution process while focusing on enhancing the overall security of the new process.

Process Re-engineering
A case study in improving confidential report distribution

In order to be successful, companies must devise a strategy to best meet their clients' needs while keeping the overhead or operating costs, as small as possible. They do this by properly balance the operating costs with their revenue. Finance and Accounting departments with in businesses not only handle the monetary transactions of the organization, but they also make sure that they distribute the balance of cost vs. revenue information to the company. This information is used to tabulate data that makes up snapshots of the financial outlook and success of the company.

At the law firm where I work, the finance and accounting departments deliver monthly reports by hand to the various practice group directors and members of the firm. Annually these reports used approximately 390,000 sheets of paper. For an employee to create and prepare their assigned portion of the reports for distribution, 240 hours annually where used on just the monthly reports. Once created and prepped, these reports where sent to a copy center to be duplicated and then sorted by recipient and once the all the reports where ready, they were then distributed by hand to every recipient. The whole process normally took between 3-4 days from the closing date to the report delivery date. During this time, the personnel resources of both the finance and accounting departments are completely utilized for the task of report creation. Conflicts would occur when someone would request specialized information for a client or other high priority requests were made to the staff. Because everyone is working to get the reports created and ready for copying, the finance and accounting departments are essentially shut down. Often a request would have to wait till the next week before it could even be evaluated, and many times by then it was too late, and the information was no longer applicable. The situation is much worse during the quarterly reporting times, as not only is the staff burdened with monthly reporting duties but the quarterly reports as well. You can imagine the nightmare that makes up the annual reporting done in December.

The finance and accounting personnel cannot be blamed for the bottleneck that is created at the start of every month as they prepare the reports that enable the various people at the firm to view their current fiscal snapshot. These are reports that have to be delivered each month into the hands of the recipients. Some of the reports contain sensitive information that should not be viewed except by the person the report is for. Next to delivering the reports, the security of the information is a vital piece to the process. By hand delivering these reports, they are ensured that the information is delivered to the right person on time. The reports are also organized and sorted for the recipient so that right away they know what they are looking at and how the information is applicable to them. Rather then having to search their entire stack of reports, each report is clearly marked and the information is easily obtainable. While this

process has flaws when thinking of the utilization of resources, it does provide unmatched security. It is also the way things have been done at this law firm since the beginning.

It is a standard practice for an IT professional to analyze and apply their knowledge to improve current business processes. Analyzing the report distribution process at my law firm, I can see many ways that the process can be improve. First the current process place a large demand on person-hours to create and prepare the reports, and also the material costs for the paper, printers and toner costs are key items that can be improved on. With respect to the current month end reporting method we ask, "What other way is there to securely distribute reports to multiple people each month other then printing and hand delivering it?"

One option would be to distribute the reports via email. Each user would receive an email containing their reports in an attachment. The main benefit of an email distribution method is that email is used by everyone in the firm and would not require any new skill. Also the report author could include additional text in the email body to help identify their reports. Alas, the number and size of these reports eliminate email distribution immediately. Anyone who has used email will tell you that creating a bottleneck on the email server once a month is a bad thing. In today's business world email has become a vital method of communication, perhaps even more so then the telephone. If we were to distribute the reports via email, it would put an enormous strain on the email sever, which would severely impact the performance of the machine, perhaps even disabling it. Also email is not as secure of a method as hand delivering the report. Unless some form of encryption is used, there is no way to ensure that the person who views the report is the intended recipient. The encryption would help make sure that the document is readable only by the recipient. However if a person where determined to get at the information, there are decryption utilities out today that given enough time, can read the contents of the encrypted file. Imagine what would happen if one of the reports contained salary information and that got into the wrong hands?

Another option would be to store the reports on a data file server and have each recipient access their reports via the network. Security would be handled on the file system side, granting users access only to those reports they should see. Storing the vast amount of reports on a server would not incur a significant expense, as the costs of hard disks have decreased over the years. However while it is possible to store these reports in a secure environment, this option does complicate the way a person would locate and retrieve their desired report. What previously was available at hand is now off on some server located within the company LAN, and a level of separation from the person and the data is introduced. This separation is a breeding ground for confusion. It is very common for IT people to take for granted the mundane task of directory navigation, but in

the real world not everyone is as computer savvy as we would like them to be. Often simple file navigation can be overwhelming for the novice user. The confusion could be limited by using meaningful directory names and organizing the reports logically on the file server.

However, another drawback to this solution is the maintenance. There are over 100 individual reports that get distributed monthly at the firm, with many having distribution lists in double digits. The only way to manage the security of the files is by individual file and individual user. The maintenance would be drastically easier if a pattern existed that would enable group based security. Due to the large amount of reports and distribution lists a grouping pattern that would decrease the maintenance load does not exist.

The optimal solution would be something that is easy to use, easy to maintain, and secure. The ease of use and maintenance cannot of course compromise the security. The goal is to provide the best possible delivery mechanism without exposing any of the confidential information. The current system while labor intensive, is secure and reliable. Whatever system is developed must also be secure and reliable. The solution that we chose to implement takes the data file server solution and wraps it inside a method that is as accessible and available as email. What we did was create an intranet website that enables the distribution of the reports we store on a data drive on our network. The website would be the firm's utility for accessing their reports, while the report creators would use the network to place the files on the data drive. Utilizing Active Server Scripts, HTML, and JavaScript we created a visually appealing and friendly environment with which the users could access their reports. Adding a database to the scripting we are able to add additional security to the reports as well as create a manageable infrastructure for the site and reports. By adding layers of security as well as alternative security checks we further protect the data. In this case we have actually exceeded the previous level of security.

The first step in creating our reporting site was to prepare the web server machine. The firm is comprised of a windows based network, so naturally our OS of choice is Windows 2000 Server. Furthermore we chose to use Internet Information Server 5.0 as our web server backend and SQL Server 2000 for our backend database. Before we even began coding, we applied all of the relevant security patches to the operating system, web server and sqlserver. Also we a ran IIS Lockdown, a tool provided by Microsoft that closes the known security flaws and disabling seldom used features in IIS 5.0.

Distribution of these reports is for internal use so there is no need for external access to the site. This should reduce the need for constant security monitoring. However, we still treat the machine as if it is exposed to the real world. Every security alert and patch that is released is applied and the machine is routinely scanned for any vulnerabilities.

The first layer of security for our report website is implemented thru the IIS application. We disable Anonymous access and turn on the Windows Authentication security method. Because in our network, each user must log into the Windows domain to access the network resources, we should never have anonymous access to our system so by disabling it we would not impact the report distribution, but rather we make it that much more secure. "…Integrated Windows authentication is a secure authentication method that doesn't transmit usernames or passwords. Instead, it relies on a cryptographic exchange with the server." ("Windows Authentication Methods"). This means that even though there is an authentication check, no data is passed in the clear during the authentication session. One of the drawbacks to using this type of authentication is that it is not cross browser compatible, it only works on Internet Explorer. However, this is an intranet website and the firm standard browser is Internet Explorer, so it is acceptable to use this for our authentication.

The Windows Authentication security transaction is the first interaction between the users and the website security. One of the additional benefits of using the Windows Authentication is that while the challenge/request is performed, it is done so without the user being involved. Because all of the machines are members of the same domain, the challenge / request is done behind the scenes. In the event that the challenge request fails the user is redirected to an invalid login page and the attempt is logged for review by the network team.

The next layer of security we add is to the individual sections of the website. To do this we need to create a strong yet flexible access control list that we will store inside our database. We will use ASP scripts to make calls to the security tables in our database to create the navigation and validate access to the various pages and reports that make up the website. To facilitate easy navigation and enable us to organize the different areas of reporting that we will have on our site we developed the infrastructure of the site into 4 tiers. Tier 1 is the highest level. This level contains a high level of distinction between the reports. Either they are Accounting reports or Finance Reports. Tier 2 is the subcategory level, which separates the reports by types, Monthly, Quarterly or Annually.

The Tier 3 is the most important tier. It represents the actual page or resource level. This level is where we assign our permissions too. If a user has permission to a Tier 3, they must be able to pass thru tiers 1 and 2. Really the

first two tiers are visual helpers for the users to navigate the site and to provide a way to organize the reports.

The fourth and final tier we built in is not really there for the report distribution on the site, but instead it is used for on demand reports. By meeting with our user community we were able to learn that there are other report types that are handled by the accounting and finance departments. These reports are specialized to show a particular bit of data or are collaborative reports that span multiple data sources and are used to show a very high level snap shot of the firm. They do not have a set schedule as to when they are distributed so they are created and sent out on demand. By adding these types of reports to our site, we are now centralizing the financial data reporting for the entire Firm.

Tier 4 works as a subset of Tier 3. If a user can get to Tier 3, they can get to all of the Tier 4's that are attached to that Tier 3. An example of this is your common HTML form submits to a results page. The form page handles the user inputs; the results page processes the inputs and then formats and displays it. The HTML form page is the Tier 3, the results page is the Tier 4. If a user can get to the HTML form, they should of course be able to submit the form.

It is important to go over the architecture of the site because it is what drives the security model. We want to keep the model as simple as possible so that it is easy to maintain but without losing any security. Because we knew we needed to make the interface to the site easy to use we built in the top two tiers. These tiers are used to organize the data on screen for the users. The middle or third tier is where we assign the access rights to. The bottom tier is just a way to link additional pages to the middle tier. The links are essential pieces to the security model we implement.

We chose to implement the page security into five types. The first type is PUBLIC. This type is intended for a page that is accessible by any person who has access to the site. This type is seldom used to grant permissions, but it does allow for an easy way to deny all users access to a section.

The second type is by Practice Group. In a law firm a practice group refers to a group of lawyers all working in the same legal area. But it also means departments, for example accounting and finance are both practice groups, even though they aren't comprised of lawyers. This type allows us to grant or deny access to large groups of people. The third type of permission is the Organization. An organization is a smaller group of people that exist under a practice group. For example the practice group of accounting has organizations

of billing and receiving. Typically a person in billing does not need to see the same information as a person in receiving, so this type of permission allows us that level of granularity.

The fourth type of permission is Custom groups. The second and third types of permissions are comprised of data that is already maintained as part of another application that we are just leveraging off of. However in some cases we need to assign multiple items the same permissions, and these permissions cross practice groups and departments. So we created a helper application that allows an administrator to create a custom group and add users to that group. That custom group can then be assigned access to a resource on the website.

The fifth and final type of permission is the Individual access. This type of permission is a direct assignment of permission for a resource to an individual user. This is useful for such reports that have only one recipient or if a useful way to group the distribution is not prevalent.

For each of these five types of permissions, only two access permissions exist, Grant and Deny. If the grant permission is assigned then the user may access the permission. If the deny permission is assigned then a user is not able to access the resource and the resource is never displayed to them. However because our permission types do overlap with each other it was necessary to add a rule to the security model that would prevent a collision of overlapping permissions. If we grant individual access to John Doe but deny access to the practice group John Doe is in, how should our security model enforce this? The answer to this is easier then it looks, we simply have to ask ourselves which is more harmful to our system, someone not getting a report they should have access to, or someone having access to a report that they should not. Many of the reports in our system contain sensitive information that should only be viewed by the intended recipients. With this in mind we built the rule into our security model that makes the deny permission at any level override any of the grants. So to deny at the practice group level and to grant at the individual level would be interpreted by the system as deny. This allows us to have overlapping security but assures us that information access is granted as long as there is no overlapping denial.

We have built a simple application security model consisting of grant and deny access with 4 tiers that enable us to manage and group the permissions easily and efficiently. The next stage of security we built in is designed to obscure the report file information from the user so that in the event that they were ever able to access the actual data file server, they would not be able to locate the files or even be able to make a reasonable guess as to a files contents. The people in charge of the reports have a network drive access to the data file server so that they may store the created reports in an area that the web site can access them. We secured the network share so that only those people who should be able to access the report share can, but we also know that by creating

an additional access method to the reports, we need to take steps to secure it as well. What we have done is obscure the actual report file's path and name. The report recipient sees a text link description of the report thru the web site. When they click it the report is then opened up inside the browser as a file called "reportfile" plus the extension. So a word document would be "reportfile.doc". On the file server the report file might be called "c12345a.doc", a name that would only make sense to the author of the report. If someone was able to access the network folder they would have no idea what the file contained or what information was inside this. This is a bit of security by obscurity with a little bit of flair. We just want to make it as difficult as possible for someone to gain useful information if they were ever to penetrate this backdoor.

The primary goal of any security system is to make the effort needed to bypass the security worth the value of the information they would retrieve. So far the first step in our overall security plan is the Windows Authentication. That was enabled within the IIS server. The 4 tier application security requires us to build both the access control lists as well as the validation code. We store our access control lists with in our SQL Server database and use ASP sub routines with SQL queries to process the security request. Then we obscure the file names and secure the network file share used by the report writers.

At this point we have done a good job of securing the site from the casual user. However we would be in error if we assumed that all of the access to this system would be exactly as intended and that there would never be an attempt to circumvent the security as it is. We want to ensure that the site is as secure as we can make it and in order to do this we had to spend extensive time into trying to circumvent the security and the re-engineering it to close any holes we found.

The navigation for the site is driven off of the ACL database and only resources that a user has access too are displayed. Each resource has an identity key. A resource identity key is passed from the navigation to the ASP scripts and is used by the security code to check the permissions for that resource. These ID's are unique to each page and report (resource), they are the link from the resource to the permission. The security validation code first checks to see if the resource has public permission with deny or grant permission assigned, then practice group, then organization, and finally individual. If any deny is found for that resource then the navigation for that resource is not displayed. This is the first way our custom security model prevents unauthorized access.

We then tried to circumvent the security by attempting to spoof the resource identity in an attempt to gain access to a report that we don't have the grant permission on and was not part of our navigation. We pass the system a

resource identity key of something we have access to, but with the page name of something we don't. Basically we are tricking the security check to run against something we have access too to get to a page we don't have access on. This test resulted in us getting past the security system. To prevent this from happening again, we added another layer of validation to the security code. Before it checks for the grant / deny status for a page, the resource identity key is compared to the name of the resource we are checking and displaying. Below is the pseudo code used to validate a page:

```
sub ValidatePage()

MyArray = split(Request.ServerVariables("SCRIPT_NAME"),"/",-1,1)

Set theKey = Getidentitykey()

Sql Query for identity key
 set Results = Conn.Execute(SQL)

if Results.EOF then
    Invalid Key/ Page Combo
end if

'We are at the top level
if Results("level") = "1" then
     sql query for the first child page
     set Results = Conn.Execute(SQL)
end if

'Do the page names match
if lcase(MyArray(ubound(MyArray))) <> lcase(Results("path")) then
    sql = Query for the rest of the children
    set childResults = Conn.Execute(SQL)
'set the flag
foundpage = false
do until ChildResults.EOF
   if lcase(MyArray(ubound(MyArray))) = lcase(childResults("_path")) then
       'Found a match
        foundpage = true

    end if
   ChildResults.MoveNext
loop
if foundpage = false then
   Invalid Key Page Combo
end if
```

```
end if
end sub
```

If they key/page combo do not match then we immediately deny access and log the attempt. Of course not every instance of this action is malicious in nature, and in most cases it could be just a typo, but it is still important to log these attempts and review these logs periodically.

All of these layers of security we created are applicable both to the reporting side and the on-demand query side of our website. There is one more layer of the security that applies to just the reporting side. We talked earlier about how we created a network file share for the report writers to be able to store the reports where the website could get them. We then secured the shares so that only the report writers could get to the shares. What we did not mention at the time was that the permissions on the report file are also used as part of our security system. As previously discussed the site will only display the navigation for items that there is a grant privilege without any overlapping deny. This applies to both reports and on-demand pages. With reports, there is one more check that is done. If the user's windows account does not have read permission on the report then the report is not shown. This is done to provide an extra layer of security as well as eliminate confusion. Only reports that exist are displayed because the ASP script uses the file system object to read the properties of the report to verify its existence. These properties cannot be read if the report does not exist, or if the user's windows account does not have read access. Because the security is controlled by the website and the application security model, the windows file security is the last line of defense to prevent unauthorized access.

So much of the focus of the design of this website was on the security of the reports that the actual distribution piece almost took a back seat in the design process. What is special to this system is that all of the security checks and validations are done without any user interaction that from their perspective it is often overlooked. However because the security is not intrusive does not imply that it is weak. In fact because each layer overlaps each other there is no single point of failure. If someone is able to by pass one layer of security does not mean that they will be able to by pass them all.

The purpose of this site was to improve upon the current report distribution process without a loss in security. The report website that was designed both meets and exceeds that goal. It has been in production for over a month and has received rave reviews from every level of the firm. The largest impact of this site is how much it saves the firm in both people and material resources. Annually these reports used approximately 390,000 sheets of paper. For an employee to create and prepare their assigned portion of the reports for distribution, 240 hours annually where used on just the monthly reports. Once created and prepped, these reports where sent to a copy center to be duplicated and then

sorted by recipient and once the all the reports where ready, they were then distributed by hand to every recipient. The whole process normally took between 3-4 days from the closing date to the report delivery date. Now the entire reporting process takes 1 day. There is no paper used to distribute the reports, however if someone wished to print their report they could do so from their office.

The staff of finance and accounting are not spending exorbitant amount of time generating this reports but rather can now spend time on more important analysis or urgent requests. All of these improvements are second to the major improvement in security. Before the reports were hand delivered to the recipient and at that point the responsibility for security was shifted to the recipient. With the website, the report is always stored on the site and the security is always in place, so the only way to incur a shift in the responsibility of the security would be to print it out. . It is not a realistic goal to make a system completely secure, it is just not possible. Instead we want to ensure that the effort needed to gain unauthorized access outweighs the value of the information that they are trying to access. The knowledge I received from the SANS GIAC course aided in the overall design of this system because during each step I was able to look at the overall picture of the site and think of all the different ways to gain unauthorized access. It also helped me not become overwhelmed with aspects of security. By simply weighing the cost of securing the site vs. the impact of a security breach I was able to apply the right balance of security to the site.

References

"IIS Lockdown Tool", URL: http://www.microsoft.com/technet/treeview/default.asp?url=/technet/security/tools/locktool.asp (1-31-04).

Princeton University, Department of Computer Science, "Web Spoofing", URL: http://www.cs.princeton.edu/sip/WebSpoofing/ (2-10-04).

Reselman, Bob Active Server Pages 3.0 by Example. Indianapolis: QUE Corporation, 2000.

Scambray, Joel et. Al. Hacking Exposed Second Edition. Berkely" McGraw-Hill, 2001

"Windows Authentication Methods", URL: http://msdn.microsoft.com/library/default.asp?url=/library/en-us/dnw2kmag00/html/iis5auth.asp (2-1-04)