



Overview

Country or Region: United States

Industry: IT Services

Customer Profile

The Microsoft IT division supports the daily IT operations of Microsoft Corporation, which is headquartered in Redmond, WA .

Business Situation

Microsoft built a custom solution for identification and manual protection of sensitive data. This also required user interaction.

Solution

MSIT leveraged the File Classification Infrastructure (FCI) in Windows Server 2008 R2 with the AD RMS Bulk Protection Tool to enforce file retention, classify sensitive data and protect information at risk.

Benefits

- Reduced management costs
- Manage risk of file exposure
- Enabled simpler compliance
- Increased IT agility
- Persistent protection

Microsoft IT Uses the File Classification Infrastructure in Windows Server 2008 R2 to reduce costs and manage risks

“FCI improves security, efficiency and effectiveness, while reducing overall risk of inadvertent loss of highly sensitive business information such as personally identifiable information, health care information, financial information, legal information and intellectual property.”

Olav Opedal, Senior Program Manager for Information Security, Microsoft

To prevent the inadvertent disclosure of High Business Impact (HBI) information Microsoft IT leveraged the File Classification Infrastructure (FCI) in Windows Server 2008 R2 with the Active Directory Rights Management Services (AD RMS) Bulk Protection Tool to automatically enforce file retention limits, classify High Business Impact data and protect information at risk. This automation has reduced cost, mitigated risks, enabled compliance, and increased IT agility.



“FCI removes the broad approach of classifying a high level share which requires user input, to a fully automatic classification and protection solution for each individual file”

Olav Opedal, Senior Program Manager for Information Security Operations, Microsoft

Situation

The Microsoft IT division supports the daily IT operations of a large global corporation that has demands similar to those of many other organizations of the same size. These demands include the requirement to provide services for file storage and sharing for more than 130,000 users and 340,000 computers in hundreds of locations worldwide. At Microsoft, over 100 Terabytes of data is disbursed over 30,000 file shares across the company. This data includes files where High Business Impact (HBI) information could reside. Microsoft IT needed to create a solution that relied on technology to identify information that could be highly sensitive and then to help prevent the unauthorized disclosure—whether inadvertent or malicious—of this information.

Data leaks are common; since 2005 alone, 245 million records of U.S. residents have been exposed due to security breaches. “Loss of sensitive data was Microsoft’s number one operational risk, with two important audit issues reported to the board in 2008,” says Olav Opedal, Senior Program Manager for Microsoft IT Security.

Microsoft has long had content policies in place in accordance with a number of regulatory and corporate mandates. In 2006, Microsoft IT initiated a data loss prevention (DLP) project with the objective of addressing content security and compliance objectives at Microsoft regarding High Business Impact (HBI) information, while minimizing impact to business operations. More information on this existing solution can be found at http://www.microsoft.com/casestudies/Case_Study_Detail.aspx?casestudyid=4000003382.

To manage user identity and data-access rights, Microsoft IT Security also employs Active Directory® object user authorization. With the Active Directory directory service, the type of access granted to objects is determined by what user rights are assigned to the user and which permissions are attached to the objects. An object is a set of attributes that can include shared resources such as servers, shared volumes, and printers; network user and computer accounts; and domains, applications, and services.

The missing component was an automated system that pulled together the identification of HBI data and the protection process, including the Active Directory Rights Management Services (AD RMS) technology -- one that MSIT could rely on without having to maintain development staff for that purpose. Another major limitation of the existing solution was that shares were classified and protected as a whole. Any files that were identified as being located in a share with incorrect classification could not be protected on their own. Instead owners would be notified and the shares would be locked as a whole if action was not taken by them.

To improve the solution, Microsoft IT Security needed new technology.

Solution

With Windows Server 2008 R2, MSIT has the capabilities needed to address the

“The FCI Solution allows MSIT InfoSec Operations to maintain the day to day operations work by pulling reports as needed and allows the team be less dependent on operational infrastructure. FCI has less impact on customers as classification and security is automated on the backend for classification.”

Norresa Calos, Program Manager
for Information Security
Operations, Microsoft

issues described above as part of the Operating System. The File Classification Infrastructure (FCI) provides mechanisms to not only classify HBI files on a fileserver, but in conjunction with File Server Resource Manager (FSRM) can provide insight into the distribution of HBI data and automated enforcement of document retention policies. With the addition of the AD RMS Bulk Protection Tool, the identification, monitoring, and protection of HBI data on file servers is fully automated on a per file basis.

Microsoft IT Security team worked with stakeholders around the company to implement the new solution. Stakeholders include teams from File Share Operations, Rights Management Services and other Collaboration Services groups, various technical-support tiers, and Microsoft Legal and business-review groups. Stakeholder participation was important because applying any restrictions or removing any documents would affect production server service levels and other aspects of the IT infrastructure.

The team established three objectives for the initial project:

- Identify HBI content across the network
- Restrict access to HBI files wherever possible to those that may have a business need to access it
- Ensure that MSIT knows where HBI information is located
- Remove any documents that have passed standard retention time limits thereby reducing the volume of HBI content that could move across the network or be used on workstations.

With the policies above in mind, the following FCI classification properties were defined:

- Business Impact; with possible values of
 - HBI (High Business Impact)
 - MBI (Medium Business Impact)
 - LBI (Low Business Impact)
- Retention; with possible values of
 - SOX
 - Long
 - Short

The method to assign values to the properties for files fell into two categories:

- Location based
- Content based

The existing solution required that share owners determine a Business Impact classification for each share. This information was then centrally recorded and used to help identify files that were out of compliance with MSFT policies. The

“[FCI] allows us to satisfy compliance needs, accounting controls, etc. - document retention, expiration, legal requirements and ultimately lower costs enabled by these efficiencies.”

Dondi Vigesaa, Service Engineer
for the File Server Utility,
Microsoft

decision was made that this classification information for shares could be used further by providing default classification for any file that is not classified according to its content or not matching content-based policies searching for HBI files.

Solution Technical Approach

FCI was configured to assign a value to the Business Impact property based on a file's location. To do this three classification rules were setup:

- A rule using the Folder Classifier that assigned Business Impact to have the value HBI. Every share that was classified by the owner as HBI was added to the scope of this rule.
- A rule using the Folder Classifier that assigned Business Impact to have the value MBI. Every share that was classified by the owner as LBI was added to the scope of this rule.
- A rule using the Folder Classifier that assigned Business Impact to have the value LBI. Every share that was classified by the owner as LBI was added to the scope of this rule.

As part of the existing solution a series of patterns were defined that attempt to identify HBI files based on their content. The patterns were designed to search for information such as Personal Identifiable Information (PII) and Intellectual Property (IP). This logic was directly transferred to FCI classification rules using the Content Classifier mechanism.

All major content analysis logic was successfully transferred to classification rules using the Content Classifier provided by FCI. However, some of the more advanced content analysis logic enabled by the existing logic did not have an equivalent in the FCI Content Classifier. In the future, this will be addressed by classification extensibility modules that enhance FCI.

Once the files were classified, File Management Tasks were used to protect any identified HBI data. Specifically, two main protection actions were taken:

- Remove any rights from the individual files granted to Guest, Everyone, Domain Users, Authenticated Users on a daily basis using a custom action that calls ICACLS.EXE on each file
- Using the AD RMS Bulk Protection Tool called from a custom action, encrypt all HBI files to a standard template maintained by MSIT. All files are encrypted so that a central authority within MSIT has control over the encrypted file. This ensures that if the owner of the file is no longer with the company, the data is still accessible.

One of the primary requirements for MSIT, once files are classified and sensitive data is protected, is that proper accounting must be maintained. To address this requirement,

- A monthly report on “Files by Property” is scheduled with the File Server Resource Manager. This report includes aggregate data on classification which can be used for correct accounting.
- The File Management Tasks are configured to create a report on files that are encrypted.

Best Practices Learned from MSIT

Property selection

Most organizations immediately identify a few dozen properties when offered the chance to track meta-data. Attempting to come up with logic to assign values to all those properties is daunting. In general the best approach is to

- Identify the data management/protection policies to enable. For example
 - Ensure files containing personal information are encrypted
 - Restore business critical files first during disaster recovery
 - Restrict access to corporate IP to full time employees
- Let the policies dictate which properties are required. Considering the previous examples that could be
 - Personal Information
 - Business Critical
 - Corporate IP
- If there is no policy that requires a specific property, it should not be tracked (at least during initial deployment)

A deployment that starts with only four properties can easily add more properties or remove existing properties (from the schema and any files).

Deployment

Deployment of a classification solution automates some complex data management activities. Due to this, a planned rollout over a period of time is recommended. The period of time for the deployment of this solution in MSIT was from July 2009 through September 2009.

The initial step is to get the classification rules setup correctly. While modifying these rules and looking for the results, it is best to limit their scope to only a small set of representative data. The actual classification process is then much faster and it does not touch end-user data until the rules are correct. As classification rules are updated, it may become necessary to remove existing classification properties from the test data. A PowerShell script that calls the ClearProperty FCI API for the test files easily accomplishes this.

Once the classification rules have been identified, the policy mechanisms should be tested on the same classified data. This avoids and incorrect management of end-user files.

Since file management tasks and reports by Property force classification of files that are not yet classified, it is recommended that all existing files on the server are already classified before these policies are enabled. To do this, the classification rules are setup for the server and a classification schedule is setup to classify files during non-peak hours (10pm-6am in this case). Once classification of the existing data is complete (as indicated by classification completing without being canceled at 6am), new data will be classified on a nightly basis and should not cause policies to be slow.

Once the FCI configuration has been determined it needs to be pushed out to all servers. To do this there are two options:

- Create a script that sets up the desired configuration, or
- Use the Import/Export script provided on TechNet Script Center

Since the Import/Export script had not yet been finalized at the time of this deployment, MSIT created a PowerShell script that configured FCI. The script was then run on each individual server or cluster.

DFS-Replication

If a file server contains a DFS replica, only one side of the replication should be included in the FCI classification scope. The other replica should be excluded from FCI classification rules.

If the server has a read-only DFS replica, that scope should also be excluded from FCI classification rules.

Any policies that take advantage of classification properties can still be configured to run on these replicas, since the classification properties are replicated.

Regular Expression Complexity

Content based classification that uses regular expression patterns to identify the classification of a file, may slow down the classification process. Regular expressions can take longer periods of time to evaluate for a given file. A complex regular expression with a long file can easily drop classification speed from hundreds of files per second to one file every few seconds.

Promote a Culture of Content Protection and Awareness

Technology and policies alone will not protect an organization. The organization needs to continuously evangelize the importance of protecting sensitive content, and provide training on the do's and don'ts of sharing content. Establishing who within the organization has ownership over content is just the first step in promoting an attentive and vigilant culture of content security. Training and ongoing oversight are also key, and are just as important as the technical safeguards and solutions that organizations implement.

Benefits

"FCI in Windows Server 2008 R2 achieves all the major scenarios MSIT required for managing HBI data that previously required a customized solution" says Matthias Wollnik, Senior Program Manager in the FCI team. "The addition of the AD RMS Bulk Protection Tool takes the solution provided by Windows Server 2008 R2 beyond what could previously be achieved."

Reduce cost

In less than two months, Microsoft IT Security implemented an end-to-end information-security solution using standard components in Windows Server 2008 R2 and the AD RMS Bulk Protection Tool - and is scanning one-fifth of the company's file environment.

"FCI enables the Information Security team to move from in-house custom built solutions which enabled the classification of shares to ensure compliance with our

data handling standard, to a solution within the Server operating system that enables the classification of individual files” says Olav Opedal, a Senior Program Manager in the Information Security team at Microsoft.

Mitigate risk

Microsoft can help safeguard its most important information by applying controls based on data sensitivity, for targeted protection that meets industry best practices for regulatory compliance. Microsoft employees can stay compliant automatically with data handling standards that call for encryption of HBI documents, without the expense of extensive user training or restricting access to information too broadly.

“FCI enables Microsoft IT to classify and protect each individual file, instead of applying classification per share. This improves security, efficiency and effectiveness, while reducing overall risk of inadvertent loss of highly sensitive business information such as personally identifiable information, health care information, financial information, legal information and intellectual property. Currently 20% of Microsoft IT’s file server infrastructure is protected by FCI, slated to become 100% within the end of the fiscal year” says Opedal.

Enable compliance

Without end-user interaction, the solution restricts access, applies persistent safeguards, and provides proper accounting information according to the data sensitivity level for easier and less-costly compliance.

IT agility

FCI allows MSIT to change the solution as their needs change without having to rebuild the whole solution. If the logic to identify Personal Information needs to change, the FCI classification rule can simply be updated with the new logic. All existing policies continue to function as before. If a policy needs to be changed (for example, documents retained long term need to now be retained for 10 years instead of 7), the policy can easily be updated without requiring a complete rescan of the files on all servers. Such flexibility simplifies the solution and ensures that costs are minimized.

“On our MSIT File Server Utility clustered systems, we’ve leveraged FSRM extensively to gain insights around usage with the basic reporting features and other mechanisms inbox” says Dondi Vigasaa Service Engineer in the File Server Utility. “With FCI now included in FSRM on R2, we have rich inbox reporting that is highly configurable and provides actionable intelligence in an automated fashion, which allows us to satisfy compliance needs, accounting controls, etc. - document retention, expiration, legal requirements and ultimately lower costs enabled by these efficiencies.”

Persistent Protection

The integration of Rights Management Services with FCI reduces cost and increases efficiency. Microsoft can centrally apply targeted and persistent rights, access policies, and safeguards to information based on its sensitivity level, wherever it resides—personal computers, servers, databases, applications, and

For More Information

For more information about Microsoft products and services, call the Microsoft Sales Information Center at (800) 426-9400. In Canada, call the Microsoft Canada Information Centre at (877) 568-2495. Customers who are deaf or hard-of-hearing can reach Microsoft text telephone (TTY/TDD) services at (800) 892-5234 in the United States or (905) 568-9641 in Canada. Outside the 50 United States and Canada, please contact your local Microsoft subsidiary. To access information using the World Wide Web, go to: www.microsoft.com

This document is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS OR IMPLIED, IN THIS SUMMARY.

© 2009 Microsoft Corporation. All rights reserved.

Microsoft, Active Directory, and Microsoft Windows Server are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. The names of actual companies and products mentioned herein may be the trademarks of their respective owners.

Document published September 2009

more.

Olav Opedal says, "We get automatic, persistent, and targeted protection of sensitive information as the solution scans for it. And if a hacker should get in, or we have an internal breach, our information is protected with Rights Management Services and/or encryption. Now, we can automatically detect sensitive information, apply safeguards, and the system notifies the owner that no further action is necessary. Data owners no longer have to classify their file shares or manually encrypt their HBI documents." Automation also reduces the risk of data owners not applying policies properly.

Future Plans

The Content Classification mechanism provided with FCI only allowed for some of the content analysis mechanisms provided by the existing solution which used RSA Data Loss Prevention. Recently, as a proof-of-concept, RSA has produced a plug-in for FCI to take advantage of the advanced content analysis mechanisms provided by RSA Data Loss Prevention. This will allow MSIT to use the full fidelity of RSA's content analysis mechanisms while leveraging the policies and integration mechanisms provided by FCI in Windows Server 2008 R2 without having to modify the FCI based solution.

One fundamental aspect of managing data is that the amount of data being managed should be limited to what is necessary. FCI with the File Server Resource Manager provides a mechanism to correctly identify files that can be archived and removed from active management. MSIT has already identified the properties need to enforce this mechanism. Test deployment of the policy is already active, but currently limited to reporting of the stale data. In the near future, MSIT will enable automatic file management tasks to expire data that is not HBI and older than 7 years.

As part of ongoing content analysis efforts, MSIT is working to define a set of classification rules to identify files

- That should be retained for the minimal amount of time
- That are required for SOX compliance

Once that is complete, the automatic expiration mechanisms will ensure that SOX required files are removed after the correct period of time and that files that should be retained for a short amount of time only will be expired after 1 year.

Windows Server 2008 R2

Windows Server 2008 R2 is the latest version of the Windows Server operating system from Microsoft. With Windows Server 2008 R2, you can create custom

Software and Services

- Microsoft Server Product Portfolio
 - Windows Server 2008 R2
- Technologies
 - File Classification Infrastructure
 - File Server Resource Manager
 - Active Directory

Third Party Products

RSA DLP Datacenter 7

solutions that are easier to plan, deploy, and manage, than with previous versions of Windows Server. Building on the features, security, reliability, and performance provided by Windows Server 2008, Windows Server 2008 R2 extends connectivity and control to local and remote resources. This means that your organization can benefit from reduced costs and increased efficiencies gained through enhanced management and control over resources across the enterprise. For more information, go to www.microsoft.com/WindowsServer2008R2/