Automating NIST Cybersecurity Risk Assessment

In several locations the NIST Cybersecurity Framework version 1.0 designates risk assessment as a key part of the cybersecurity process. The Framework has five cybersecurity program functions. Risk Assessment is part of the Identify (ID) function, in the category ID.RA (risk assessment). Subcategory ID.RA-5 notes that "Threats, vulnerabilities, likelihoods, and impacts are used to determine risk". ID.RA-5 then calls out safeguard RA-3 from the current NIST Safeguards Inventory, NIST 800-53 rev 4. First published in 2005, NIST 800-53 is the standard Federal document for designating "Security and Privacy Controls for Federal Information Systems and Organizations".

The current version of security control RA-3 in NIST 800-53 rev 4 calls out protocols NIST 800-30 and 800-39. Security control RA-4, risk assessment updating, has been withdrawn and incorporated into RA-3, which now includes both quantitative risk assessment and periodic risk assessment updating. Good practice recommends risk assessment updating quarterly, although the safeguard states "at least annually".

Definition of Risk

Page 1 of NIST 800-30, first published in 2002, states that

> "Risk is the net negative impact of the exercise of a vulnerability (by a threat source), considering both the probability and the impact of occurrence."

Note that risk is much more complex than simple vulnerability. Vendors of automated vulnerability systems are prone to claiming that their programs produce a risk analysis. True cybersecurity risk includes vulnerability. It also must include threat source, impact and likelihood. Sadly, a high level of threat sources must be assumed in the current world.

An open window may be considered a **vulnerability** for uncontrolled access to a building. An open window in the server room on a back alley may be considered to create a **high risk** for theft of data by a malicious outsider. An open window in the server room on a tenth floor with no fire escape may be considered a **low risk** for malicious outsider data theft. The vulnerability and impact of data theft are the same in each case, but the risk is very different.

Note that low risk is not no risk. While a tenth floor window is not easily accessible, helicopters and window washers can provide access. There is no such thing as a zero risk, but there may be acceptable risks as defined by the NIST.

NIST 800-30 and NIST 800-39

Since the Clinger-Cohen Act of 1996, the National Institute of Standards and Technology has been required to set the standards for information security throughout the federal government. The publication of the NIST 800-30 risk assessment procedure in 2002 both eased and complicated the burden on organizations required to complete cybersecurity risk assessments. Although it established the procedures for assessing risk, the NIST 800-30 procedures are both voluminous and complex. Manually conducting a NIST compliant risk assessment remains problematic for many organizations. A partial copy of the applicable NIST risk assessment references is shown at right.

NIST information security risk management involves assessing risks, responding to risks by implementing safeguards and monitoring the results of the implementation. Since 1995, the best practices cycle time for periodic risk assessment has gone from every three years to every year to every quarter. With the March 3, 2011 final publication of NIST 800-39 risk assessment, response and monitoring have gone to continuous "near real-time" risk management. The process is shown in the diagram below, taken from page 32 of the NIST 800-39 protocol.

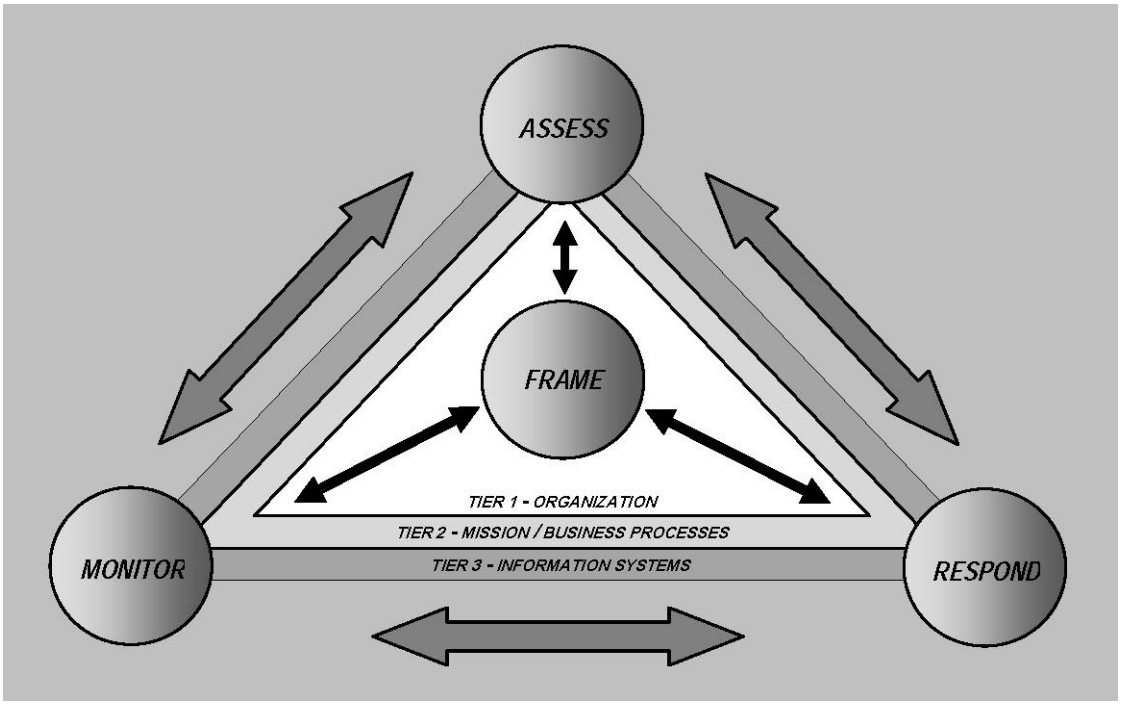

NIST 800-39 risk management process diagram

Automated Calculation of Risks

The NIST 800-30 documentation, while detailed, is well written. Page 8 states that "The risk assessment methodology encompasses nine primary steps…

Step 1 System Characterization (Section 3.1)
Step 2 Threat Identification (Section 3.2)
Step 3 Vulnerability Identification (Section 3.3)
Step 4 Control Analysis (Section 3.4)
Step 5 Likelihood Determination (Section 3.5)
Step 6 Impact Analysis (Section 3.6)
Step 7 Risk Determination (Section 3.7)
Step 8 Control Recommendations (Section 3.8)
Step 9 Results Documentation (Section 3.9)."

Manual risk assessment using the NIST protocols is a long process. A typical small organization could require up to 3 days of expert services from an experienced consultant who has mastered the NIST

protocols. However, many of these steps can be automated using standard expert system computer simulation methods well known to persons skilled in the field. The process for automating the risk assessment protocol is similar to that for automating tax return preparation and similar complex procedures.

1. **System Characterization (3.1)**

   Page 12 of 800-30 recommends questionnaires, document review, and automated scanning tools for system characterization. The Protected Information (PI) inventory taken before the risk assessment will form a large part of this task. In particular, ALL hardware and software in the organization needs to be inventoried. Only a complete inventory can ensure a systematic characterization of the system.

   Automated scanning tools are invaluable for system characterization, since they can provide detailed information at affordable cost. The National Vulnerability Database currently (May 2016) lists over 76,000 known vulnerabilities. Automated checklist scanners validated under the Security Content Automation Program (SCAP) can review hundreds of workstation configurations in minutes and report the information as pass/fail compliance with NIST recommended standards (Step 4 of the NIST 800-30 Risk Management process).

   Network scanning can be done in two general ways; widescale scanning of all workstations or individual scanning of representative machines. For preliminary assessments individual scans provide an inexpensive starting point. The use of network management tools to provide uniform configurations reduces the need for scanning of all workstations.

   As a rule of thumb, most initial workstation scans find very few passing configurations. As the network becomes hardened over time, it will be frequently be necessary to scan more workstations to confirm that all of the units are hardened to the same degree. A partially hardened network with one vulnerable point is a vulnerable network.

   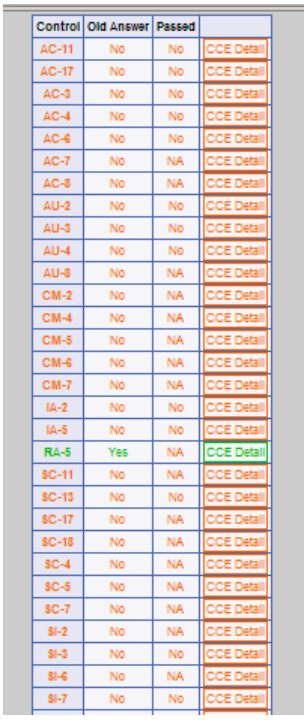

Control	Old Answer	Passed	
AC-11	No	No	CCE Detail
AC-17	No	No	CCE Detail
AC-3	No	No	CCE Detail
AC-4	No	No	CCE Detail
AC-6	No	No	CCE Detail
AC-7	No	NA	CCE Detail
AC-8	No	NA	CCE Detail
AU-2	No	No	CCE Detail
AU-3	No	No	CCE Detail
AU-4	No	No	CCE Detail
AU-8	No	NA	CCE Detail
CM-2	No	NA	CCE Detail
CM-4	No	NA	CCE Detail
CM-5	No	NA	CCE Detail
CM-6	No	NA	CCE Detail
CM-7	No	NA	CCE Detail
IA-2	No	No	CCE Detail
IA-5	No	No	CCE Detail
RA-5	Yes	NA	CCE Detail
SC-11	No	NA	CCE Detail
SC-13	No	No	CCE Detail
SC-17	No	NA	CCE Detail
SC-18	No	NA	CCE Detail
SC-4	No	NA	CCE Detail
SC-5	No	NA	CCE Detail
SC-7	No	NA	CCE Detail
SI-2	No	NA	CCE Detail
SI-3	No	No	CCE Detail
SI-6	No	NA	CCE Detail
SI-7	No	No	CCE Detail

   Uploaded scan results are translated into pass/fail results for specific NIST 800-53 recommended safeguards (see Step 4 of the 800-30 process.) Any network scanner that is SCAP validated can be used with this risk assessment program. There are more than two dozen validated scanners listed on the SCAP Validated Products page.

2. **Threat Identification (3.2)**

   Page 13 of NIST 800-30 lists natural threats, human threats, and environmental threats. An early 2000 information security paper by Jaisingh and Rees refers to the Microsoft classification of threats as being divided into natural disasters, Human Error, Malicious Insiders and Malicious Outsiders. Later papers, by Mintaka (2003) and Altoros (2007), include a more elaborate version of the Rees diagram.

   While there are other acceptable ways to identify threats, the dominance of Microsoft products in the marketplace indicates that the use of the Microsoft division of threats into Environmental, Human Error,

Malicious Insiders and Malicious Outsiders is both useful and widely acceptable. The Rees diagram is shown below.

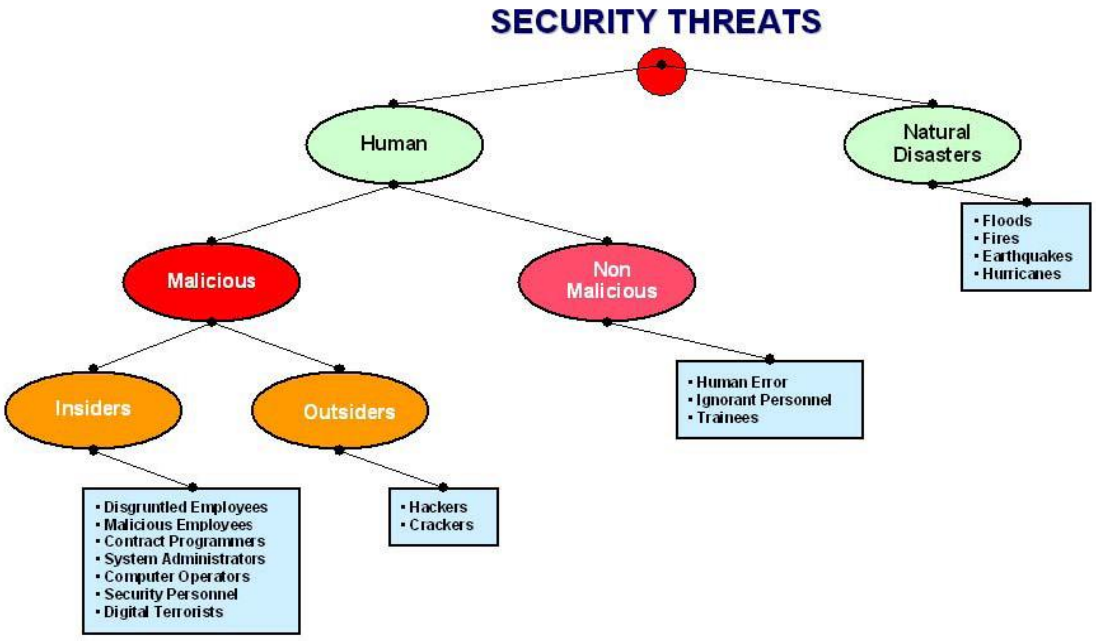

3. Vulnerability Sources (3.3)

In 2005, the NIST created the National Vulnerability Database (NVD), which superseded the I-CAT database referred to on page 16 of 800-30. The NVD is incorporated into the SCAP validated scanners that are an integrated part of the Automated Risk Management program from ACR. As of May, 2016 the NVD contains over 76,000 known vulnerabilities.

Page 18 of 800-30 notes that vulnerabilities in management, operational, and technical areas all need to be considered.

The Automated Risk Management program from ACR system further divides vulnerable areas into management (Procedure implementation and Internal controls), operational (Data acquisition, Data storage, Data retrieval, Data modification, Data transmission) and technical (System design). In addition, the environmental vulnerabilities of Wind (roof damage), Fire (and smoke) damage, Flood, Power loss (loss of operations), Power loss (Damage to building), and Vehicle collision (including car bombs) are included. It is believed that this division was taken from an early 800-30 risk assessment draft, but the original source has been lost.

Other division of areas of vulnerability could be made, but these are reasonably comprehensive and are easily assigned to particular 800-53 safeguards.

4. Control analysis (3.4)

The utility of the 800-30 risk assessment process was greatly enhanced by the 2005 publication of 800-53, “Recommended Security Controls for Federal Systems.” For the first time, a listing of

adequate safeguards to achieve an acceptable level of risk was made explicit by an authoritative source.

This frequently updated list is the basis for much of the Automated NIST Risk Management program from ACR. In addition to the any safeguards determined by optional scanning, additional safeguards are the subject of an online questionnaire, as shown below.

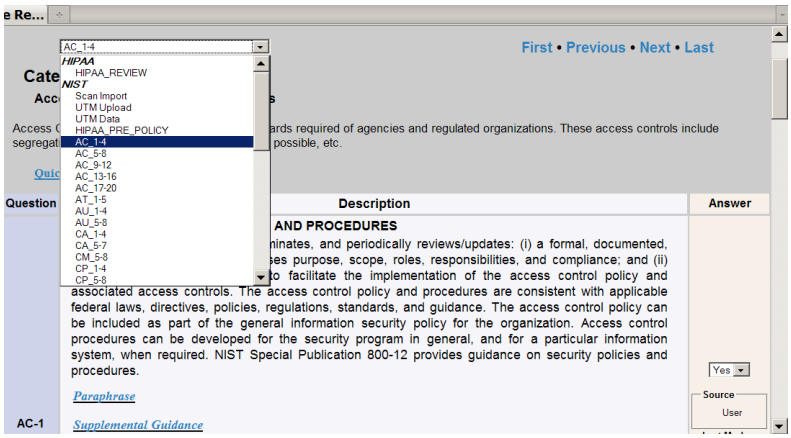

Two key elements in control analysis are anti-virus protection and intrusion protection. Both are highly important precautions, and the volume of virus and intrusion traffic is closely associated with the current security level of a network. A badly infected network will be both compromised and slow, as more and more network resources are misapplied by unauthorized uses.

Intrusion detection and anti-virus attack results are either uploaded or manually entered into the risk assessment program, depending on the specifics of the network system.

5. Likelihood determination (3.5)

For an 800-30 risk assessment, likelihood has a specific legal meaning, as follows;

> *High* - The threat-source is highly motivated and sufficiently capable, and controls to prevent the vulnerability from being exercised are ineffective.
>
> *Medium* - The threat-source is motivated and capable, but controls are in place that may impede successful exercise of the vulnerability.
>
> *Low* - The threat-source lacks motivation or capability, or controls are in place to prevent, or at least significantly impede, the vulnerability from being exercised.

Since the publication of 800-30 in 2002, cybercrime has exceeded illegal drugs as the leading criminal activity worldwide. Threat source motivation and capability can reasonably be assumed.

The Automated Risk Management program from ACR features safeguards from NIST 800-53. Mapping of these safeguards to the four threat sources (Environmental, Human error, Malicious insider and Malicious outsider) is done by inspection. For each threat source, the vulnerable areas of management (Procedure implementation and Internal controls), operations (Data acquisition,

Data storage, Data retrieval, Data modification, Data transmission), and technology (System Design) are also fairly obvious. With over 7,000 entries, this mapping is complex and time consuming, but fairly rigorous.

The validation of the safeguards map into an expert system computer program was done by observing experienced risk assessment consultants and tweaking the risk calculation engine to produce the same results using either a human expert or the expert system computer program. The development team had access to federally audited risk assessments, which greatly facilitated the validation process.

Information security risk assessments produced with this automated system have been audited by dozens of OCC, FDIC and HHS experts. No failed audits have been experienced.

6. **Impact analysis (3.6)**

Impact levels under 800-30 have very specific definitions.

> *High* - Exercise of the vulnerability (1) may result in the highly costly loss of major tangible assets or resources; (2) may significantly violate, harm, or impede an organization's mission, reputation, or interest; or (3) may result in human death or serious injury.
>
> *Medium* - Exercise of the vulnerability (1) may result in the costly loss of tangible assets or resources; (2) may violate, harm, or impede an organization's mission, reputation, or interest; or (3) may result in human injury.
>
> *Low* - Exercise of the vulnerability (1) may result in the loss of some tangible assets or resources or (2) may noticeably affect an organization's mission, reputation, or interest.

The calculation of impact levels is also mapped to 800-53 safeguards in a fairly obvious fashion. For example, a system that does not meet the requirements of safeguard CP-9, Information System Backup, will be much more impacted by Fire than a system which is compliant with CP-9 and has a well written contingency plan (CP-2) that includes training (CP-3) and testing (CP-4).

7. **Risk determination (3.7)**

The calculation algorithm for the risk assessment is given on page 25 of 800-30. Low, Medium, and High likelihoods of adverse events are scored at 0.1, 0.5 or 1.0, respectively. In the same manner, Low, Medium, and High impacts are scored at 10, 50 and 100 respectively. By multiplying the likelihood score and the impact score, a risk score from 1 (low) to 100 (high) is calculated.

8. **Control recommendation (3.8)**

The Gap Report gives a mapping of the recommended safeguards which are missing or underperforming, against the identified risks in order of impact. This report should be used to determine which safeguards need to be changed or updated.

The advantage of the Gap Report is that it can be used to prioritize the order of activities in the most cost effective manner possible. A frequent issue with security programs is the problem of treating all necessary changes as being equally important. The Gap Report prioritization gives a management tool for distinguishing between the important and the urgent changes to the system.

9. **Results documentation** Upon completion of the Automated Risk Management program from ACR risk assessment, the initial set of data will produce two reports, a "Baseline Report" showing the risk scores ordered by threat source and a "Risk Assessment Chart. "with the same risk scores shown in graphical form. Samples are shown below.

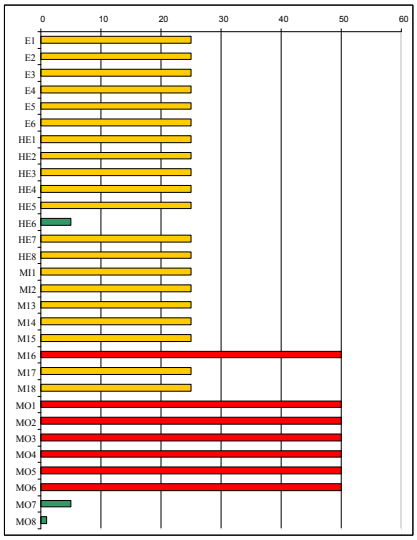

	Threat Source	Vulnerability	Likelihood	Impact	Baseline Score
E1	Wind	Roof damage	M	M	25
E2	Fire	Smoke damage	M	M	25
E3	Flood	Facility damage	M	M	25
E4	Power loss	Loss of operations	M	M	25
E5	Power loss	Damage to building	M	M	25
E6	Vehicle collision	Facility damage	M	M	25
HE1	Human error	Data acquisition	M	M	25
HE2	Human error	Data storage	M	M	25
HE3	Human error	Data retrieval	M	M	25
HE4	Human error	Data modification	M	M	25
HE5	Human error	Data transmission	M	L	25
HE6	Human error	System design	M	M	5
HE7	Human error	Procedure implementation	M	M	25
HE8	Human error	Internal controls	M	M	25
MI1	Malicious insider	Data acquisition	M	M	25
MI2	Malicious insider	Data storage	M	M	25
M13	Malicious insider	Data retrieval	M	M	25
M14	Malicious insider	Data modification	M	M	25
M15	Malicious insider	Data transmission	M	H	25
M16	Malicious insider	System design	M	M	50
M17	Malicious insider	Procedure implementation	M	M	25
M18	Malicious insider	Internal controls	M	H	25
MO1	Malicious outsider	Data acquisition	M	H	50
MO2	Malicious outsider	Data storage	M	H	50
MO3	Malicious outsider	Data retrieval	M	H	50
MO4	Malicious outsider	Data modification	M	H	50
MO5	Malicious outsider	Data transmission	M	H	50
MO6	Malicious outsider	System design	M	L	50
MO7	Malicious outsider	Procedure implementation	M	L	5
MO8	Malicious outsider	Internal controls	L	L	1

A periodic information security risk assessment is required under several sets of regulations, including HIPAA, FISMA, GLBA and the Cybersecurity Framework. However, manual preparation of a compliant assessment is likely to be outside the experience of most security officers. The burden these regulations place on organizations can be eased by the use of an Automated Risk Management program.

Using the Risk Assessment to Design a Risk Management Program

The key elements of the risk management process are the Gap report and risk assessment program.

The Gap Report is an online report that lists missing or underperforming safeguards in order of their impact on overall risk scores. A sample is shown on the page following.

The key to cost effective risk management is to review the top 15-25 missing safeguards, correct the quick and inexpensive items, input the new information and rescore the risk assessment. The risk assessment program is purchased on an annual license and can be updated as often as daily. By making the process iterative and starting at low expense, the organization can build management support by showing reasonable progress towards full HIPAA Security Rule compliance.

Gap Report

Assessment ID:	05-09-16-1462797253
Site URL:	http://www.multiple.riskassess.complianceobjects.com
Finalized date:	May 9, 2016 - 5:34 am
Report date:	May 17, 2016 - 7:22 am

1. CP-2 CONTINGENCY PLAN

Comments and Implementation Schedule

Safeguard not current and in place.

High Risks Affected:

E1 E2 E3

Medium Risks Affected:

E4 E5 E6 HE1 HE2 HE3 HE4 HE5 HE7 HE8 MI2 MI3 MI4 MI6 MI7 MI8 MO4 MO6

Low Risks Affected:

HE6 MI1 MI5 MO1 MO2 MO3 MO5 MO7 MO8

Official Wording:

The organization develops and implements a contingency plan for the information system addressing contingency roles, responsibilities, assigned individuals with contact information, and activities associated with restoring the system after a disruption or failure. Designated officials within the organization review and approve the contingency plan and distribute copies of the plan to key contingency personnel.

References

- NIST 800-12, 800-14, 800-34, 800-66

2. CP-3 CONTINGENCY TRAINING

Comments and Implementation Schedule

Safeguard not current and in place.

High Risks Affected:

E1 E2 E3

Medium Risks Affected:

E4 E5 E6 HE1 HE2 HE3 HE4 HE5 HE7 HE8 MI2 MI3 MI4 MI6 MI7 MI8 MO4 MO6

Low Risks Affected:

HE6 MI1 MI5 MO1 MO2 MO3 MO5 MO7 MO8

Official Wording:

The organization trains personnel in their contingency roles and responsibilities with respect to the information system and provides refresher training [Assignment: organization-defined frequency, at least annually].

References

- NIST 800-34, 800-50, 800-84

Most of the safeguards recommended by NIST are policy and procedure related. Implementation costs are low. Once the inexpensive issues are dealt with, the effectiveness of more technical changes can be calculated and documented using the risk assessment program. Periodic changes to the NIST recommended safeguards are automatically updated in the risk assessment program.

Improvements over Time

The risk assessment status of a small medical organization (7 people) over time is shown below. The initial assessment showed a large number of unacceptable risks, shown in red. Over a six month period the system was brought into full NIST information security compliance using a series of incremental improvements.

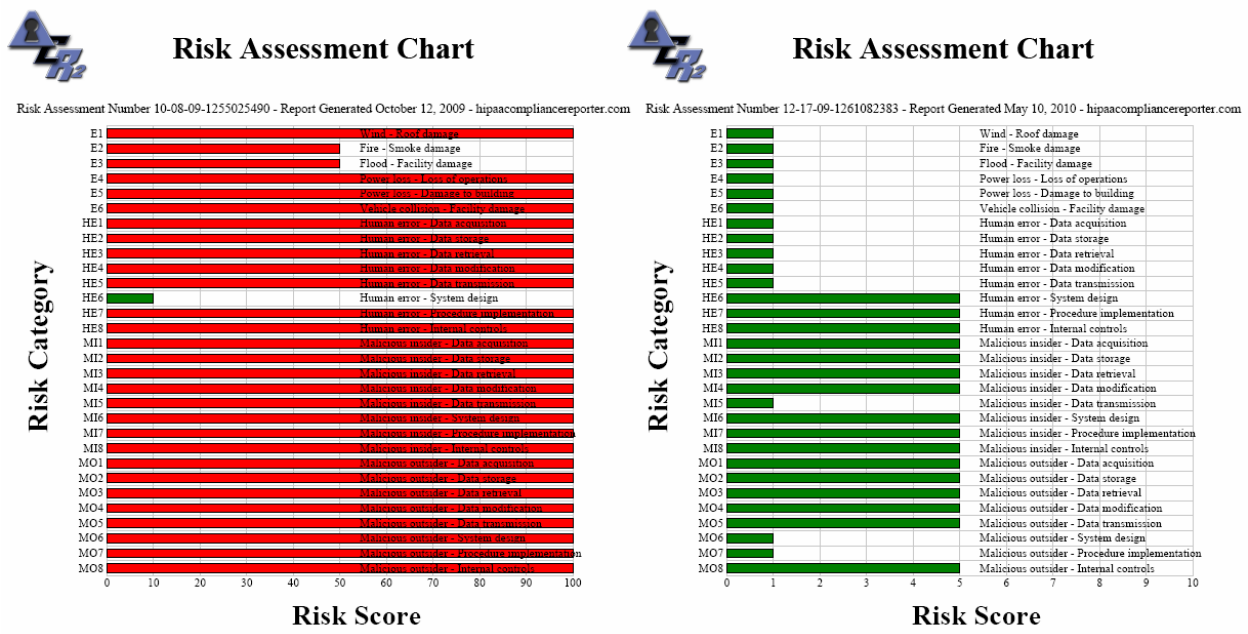

Ongoing Maintenance

Once a system has been brought into cybersecurity compliance, maintenance is an ongoing task. Every change to the system needs to be recorded into the system inventory. Every change to NIST recommendations needs to be responded to and the system upgraded as necessary. However, the advantage of the risk assessment software is that the difference between important changes that need immediate response and changes that can be phased in over time is easily documented through changes in risk scores and in the Gap Report.

Near Real-Time Risk Assessment

The issuance of NIST 800-39, the "flagship document" of the NIST IT security publications marks a significant change in information security risk management philosophy. It explicitly replaces the 2002 NIST 800-30 periodic risk assessment with "near real-time management of risk" and puts a high emphasis on widespread responsibility for information security across the organization.

In 2009 ACR 2 Solutions demonstrated an NIST 800-39 near real-time updated risk assessment program at the computer laboratory at Clarkson University. The automated continuous risk management system demonstrated at Clarkson University dealt only with the Tier 3 Information System level as shown in the NIST 800-39 process diagram. The components of the demonstration system are shown on the following page.

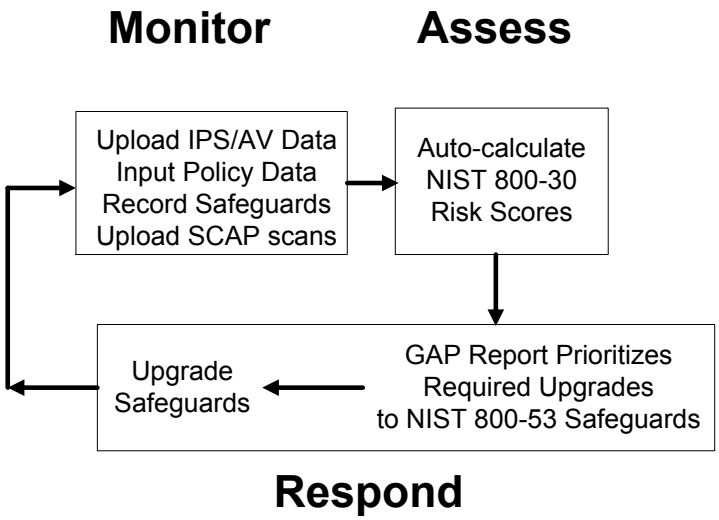

The "feedback loop for continuous improvement in the risk-related activities of organizations" now recommended by NIST 800-39 is essentially identical to the process demonstrated by ACR 2 at Clarkson University. During the demonstration period, the **Monitor** step included IPS and AV data, policy data, changes in published NIST 800-53 safeguards and SCAP vulnerability scans of all of the test PCs in the Clarkson computer laboratory. The Intrusion Prevention System (IPS) and Anti-Virus (AV) information were provided by a Fortinet UTM appliance. SCAP vulnerability scanning was provided by a Threat Guard SCAP validated scanner. The ACR 2 Solutions Risk Calculation Engine took the monitoring data and translated it into risk scores **(Assess)** using the NIST 800-30 protocol. The Engine then issued a list of proposed upgrades to the current safeguards. Upgrading the safeguards (**Respond**) provided information to the **Monitor** step and the system cycle repeated.

The practical frequency of updating of IPS and AV data depends on the level of activity in the system. At the fairly active Clarkson University site it was determined experimentally that daily reporting of IPS and AV data provided a long enough sampling period that random spikes were eliminated. A large site with much higher bandwidth might be able to use a shorter cycle time. A very small site might only be able to cycle the risk assessment on a weekly or even monthly basis without reporting huge fictional swings in apparent risks caused by a single attempted intrusion.

The key calculation to determine what the practical interval can be for "near real-time monitoring" is the background level of attempted intrusions and virus attacks. Ideally, a monitoring interval should be long enough that a real increase in attacks will be reported but a statistically insignificant increase will be ignored. During the Clarkson demonstration typical 24 hour event totals averaged 2.5+/- 0.3. Hourly monitoring intervals would have given a range of zero to two, making statistical analysis problematic. A real attack effort – apparently from a Chinese IP address - increased 24 hour event totals from around 2 to over 100. The automated program had an automatic notification to system administrator.

A useful rule of thumb is that monitoring period to period variation should always be less than 50%, and ideally less than 5%. For large organizations with high bandwidth this could yield near real-time monitoring intervals of days, or even hours. For small organizations with limited traffic, minimum monitoring intervals could be weeks or even months. This value must be determined experimentally, using the data from a Unified Threat Management (UTM) device or dedicated Intrusion Protection System (IPS).