

# Streamlining Remediation Through Attack Simulation

By Gidi Cohen

## Overview

The complexity involved in remediation management is overwhelming for many organizations today. Balancing business and IT constraints with the need to protect themselves against an ever-increasing number of threats and attacks, enterprises need a new approach to remediation management that reduces both business risk and process cost. By automating the selective remediation approach, enterprises can achieve exponential payback in remediation efficiency and enterprise security. The use of attack simulation technology not only can determine critical exposures but also can assist in planning optimal remediation strategies and simulating planned changes before actual implementation. With such an automated process, enterprises can decrease overall business risk, reduce the costs involved in remediation, and protect the organization's assets without disruption to the business.

## Remediation Management Challenges

### Process Complexity

Remediation management is one of the most complex business processes in information security organizations. Its complexity stems

Approach:	<i>Reactive Remediation</i>	<i>Patch it all</i>	<i>Selective remediation</i>
<b>Key Concept</b>	Remediate systems only if actually attacked (reactive approach)	Patch every vulnerability automatically (using a patch deployment mechanism)	Manually find only the critical exposures and mitigate them in the most effective way (patching, changing configurations, etc.)
<b>Business Risk</b>	<b>Very High</b> due to frequent attacks on critical unprotected systems and regulatory noncompliance	<b>High</b> due to potential destabilization of critical business applications and processes and delayed solution (i.e., exposure to attacks) caused by patch availability and organizational constraints	<b>High</b> due to lengthy manual analysis process (i.e., exposure to attacks)
<b>Process Cost</b>	<b>Very High</b> due to disaster recovery costs after serious damages caused by an attack, e.g., a worm	<b>High</b> due to the number of systems touched during the process (i.e., handling costs)	<b>High</b> due to the cost of manual labor involved in analyzing and planning for the appropriate remediation

**Figure 1: A comparison of remediation management processes and their associated business risks and process costs.**

from the delicate balance required between the size of the IT environment, business considerations, and technology constraints on one hand and the need for an effective and quick remediation process on the other hand.

Today's typical IT environment consists of dozens of vulnerabilities per node. Considering that a large environment consists of tens of thousands—or even hundreds of thousands of nodes—the total number of vulnerabilities within the environment becomes exponentially large. While it is tempting to fix each one of these millions of vulnerabilities as soon as it appears, by using some automated solution for software distribution or patch deployment, this approach is impractical. And, in many cases, correct configuration of control mechanisms, such as firewalls and intrusion prevention systems, is more cost-effective than patching.

An additional layer of complexity is added to the remediation management process because the responsibilities are typically split between a few distinct organizations within an enterprise: The process is driven by the security organization, but changes are executed by other groups within IT, such as network or system management. This organizational structure not only creates a communication and control barrier, but also imposes business constraints on the remediation management process. For example, the timing of a required patch for an online trading system must be balanced with the uptime requirements of the system and the availability of IT resources to perform system changes.

The challenge is further complicated by technology constraints. Many vulnerabilities do not have an acceptable patch or remedy. For example, an operating system upgrade that remediates a specific vulnerability can cause another application, which is dependent on an older version of the operating system, to break. Such technical interdependencies between enterprise software and infrastructure components—including Web and application servers, databases, and operating systems—can seriously constrain enterprise remediation plans.

Finally, to ensure the protection of the enterprise network from fast-spreading worms, zero-day vulnerabilities, and application changes which may leave the network unknowingly exposed, the entire remediation process must take place within mere hours or, at the most, a few days. Today, the entire process can take weeks or even months, enabling attackers to cause irreparable harm before preventative action is taken.

## The Pain

The current complexity of remediation management process is accompanied by significant pain, both in terms of business risk and cost. Current remediation alternatives for organizations include: *reactive remediation*, *patch it all* (a.k.a., *over-patching*), or *selective remediation*. Figure 1 analyzes these three alternatives.

With the associated risks and costs outweighing the benefits of each of the existing approaches, none of them are compelling, long-term solutions. The remainder of this article will focus on a true solution to the remediation process: the automation of *selective remediation*. This solution automates the manual steps of the remediation management process, thereby decreasing business risk and process costs.

## The Solution—Automated Selective Remediation

### The streamlined remediation process

The streamlined remediation process is composed of five primary steps, illustrated in figure 2.




Figure 2: The five steps of a streamlined remediation process.

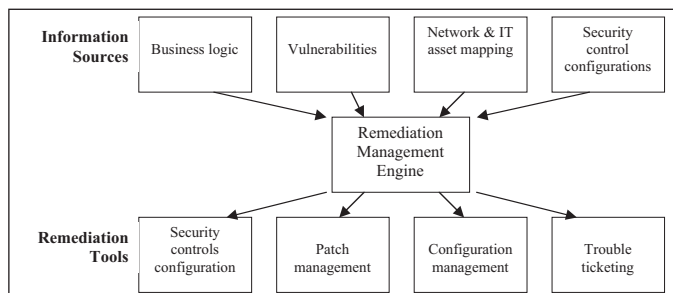


Figure 3: The three components of an automated selective remediation solution.

**Step 1: Finding critical exposures.** As stated earlier, there are many “remediation opportunities” due to the large number of vulnerabilities that exist. Luckily, only a very small fraction—typically 1% to 2%—of those vulnerabilities are actually critical to the business. Those critical vulnerabilities or infrastructure weaknesses are usually the only ones that require immediate attention and resources.

**Step 2: Planning optimal remediation.** For the 1% to 2% of all vulnerabilities that are critical, there are frequently several alternatives for remediation: patching, configuration changes, network alterations, and others. In addition, business and technical constraints, such as business risk, required availability and uptime, remediation tools, and software dependencies, must also be considered.

**Step 3: Simulating planned changes.** Before executing the remediation plan, which can be costly and error-prone, it is important to simulate planned changes to verify the viability of the plan.

**Step 4: Executing the plan.** Based on the approved plan, the appropriate changes are made to the network, system and application.

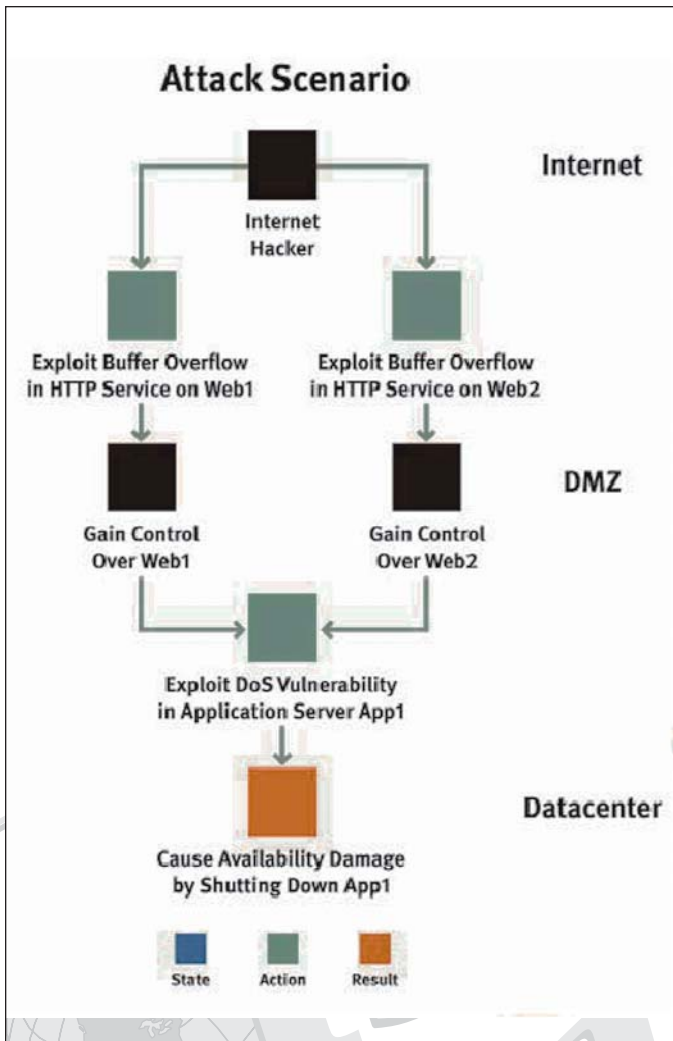
**Step 5: Verifying changes.** After executing the plan, it is mandatory to verify that the changes were successful and no critical component was unintentionally broken.

With innovative technology that increases the efficiency and automates steps one through three above, enterprises can shrink the time to remediation while simultaneously reducing the resources required to complete the process.

### Key Components of an Automated Selective Remediation Solution

An automated selective remediation solution is comprised of three primary components: information sources, a remediation management engine, and remediation tools.

- ▲ **Information Sources**—Any automated solution must be fed with various information sources, including business logic that describes the enterprise’s information assets and their associated criticality and business constraints; vulnerabilities that reside both in the critical information assets themselves, as well as in the adjacent and

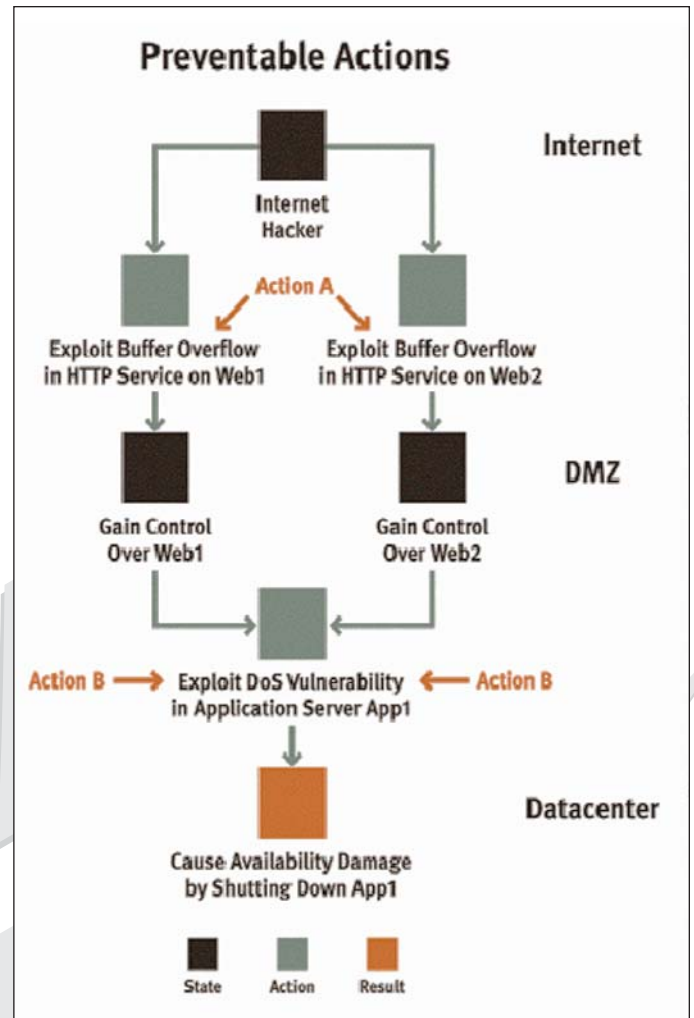


**Figure 4: A typical Internet hacker attack scenario.**

supporting infrastructure; network topology and asset mapping, including version number and patch level; and the configurations of security controls, such as firewall rules, intrusion prevention system configurations, and access control lists.

- ▲ **Remediation Management Engine**—The brains of the remediation management process reside in the remediation management engine. The engine automates all current manual steps to find critical exposures, plan for optimal remediation, and simulate the planned changes.
- ▲ **Remediation Tools**—In order to complete the remediation process, various remediation tools may be utilized in order to perform the desired changes and verify their success. The larger the toolbox, the greater optimization of the remediation process. A typical process may consist of a combination of remediation techniques, including patching, security control reconfiguration, system reconfiguration, and network re-architecture.

Such an automated, integrated process is a compelling idea, but is it really possible? The short answer to this question is YES. The long answer is that by using automated modeling and attack simulation techniques, the entire remediation process can be automated, bringing efficiency, cost and security benefits to organizations that implement these techniques.



**Figure 5: Two sets of preventable actions and their associated business risks.**

The next section will describe attack simulation and how it can be used to achieve a fully streamlined remediation process, saving up to 95% of the cost of other, less efficient approaches.

## Enabling Technology—Attack Simulation

### Attack Simulation Defined

An attack scenario is a step-by-step recipe of the sequence of actions that an attacker could take in exploiting vulnerabilities and infrastructure weaknesses to breach the Confidentiality, Integrity, or Availability (CIA) of critical information assets. Attack simulation is the technology that can automatically find all attack scenarios that can be initiated by all threats—internal or external—targeting any information asset.

In order for attack simulation to find all possible attack scenarios on a continuous basis, it must work on a virtualized model of the IT environment. In this way, attack simulation is an exhaustive yet non-intrusive process.

Figure 4 illustrates a typical attack scenario. The attack, started by an Internet hacker, causes a critical availability breach to the target enterprise's datacenter infrastructure—in just two simple steps.

Each of the attack scenarios begins at one of the threat origins and includes a sequence of one or more attack steps that can be performed by an attacker in order to cause the business loss. An attack step can be

either an exploitation of vulnerability or a 'legitimate' usage of a service (e.g., telnet, and ftp).

Attack scenarios are extremely useful in several steps in the remediation process, including finding critical exposures, planning optimal remediation, and simulating planned changes.

## Finding Critical Exposures

Critical exposures are the ones that may lead to an attack by threatening a critical information asset. Therefore, the only way to selectively find critical exposures is by finding all possible attack scenarios that may lead to a serious business impact.

Business risk is a good way to measure serious business impact. Risk to an information asset depends on two factors: 1) the likelihood of a successful attack on the information asset and 2) the potential damage that could be caused by the security loss. In a mathematical format, this risk would be represented by the equation:

$$\text{Risk} = \text{Likelihood} * \text{Impact}$$

The computation of the attack scenario likelihood should take into account the likelihood that an attack will be initiated from the threat origin (as estimated based on threat intelligence data), the number of attack steps in the attack path, and finally, the success likelihood of each of the attack steps. The success likelihood of exploiting a vulnerability is determined by the degree of difficulty in exploiting the vulnerability, the skill of the attacker, as well as the popularity or availability of the vulnerability. A vulnerability which is known to be popular among hackers carries a higher probability of success.

The business impact is determined by analyzing the specific business process at risk and considering other impacts, such as regulatory incompliance. The remediation management engine typically receives this information as part of the business logic provided.

## Planning Optimal Remediation

The key to mitigating an attack scenario is determining the *minimum* set of attacker's actions which, if prevented, would mitigate the entire attack. Figure 5 illustrates two sets of preventable actions. The first action is an exploitation of the buffer overflow vulnerability in the HTTP service (Action A). The second action is the exploitation of the Denial of Service (DoS) vulnerability in the application server (Action B). Preventing either of these actions is enough to prevent the entire attack.

Of course, not all actions are created equal. Each action might have different prevention alternatives. Figure 6 illustrates that Action A can be prevented by blocking access from the Internet to the HTTP service (Remedy A1), or alternatively by patching the HTTP services in order to remove the buffer overflow vulnerability (Remedy A2). Action B can be eliminated by adding a firewall rule that prevents access from the DMZ to the management interface of the application server (Remedy B1). As an alternative, the application server could be patched in order to eliminate the DoS vulnerability (Remedy B2).

Deciding on the *right* remedy requires the consideration of many factors. For example, if the organizational policy requires access from the Internet to the HTTP services in the DMZ, the option to block such access doesn't make sense, i.e., Remedy A1 is irrelevant.

Another example is the prevention of Action B. Upgrading the application service (Remedy B2) may require shutting down the e-banking application for eight hours, preventing customers from accessing their

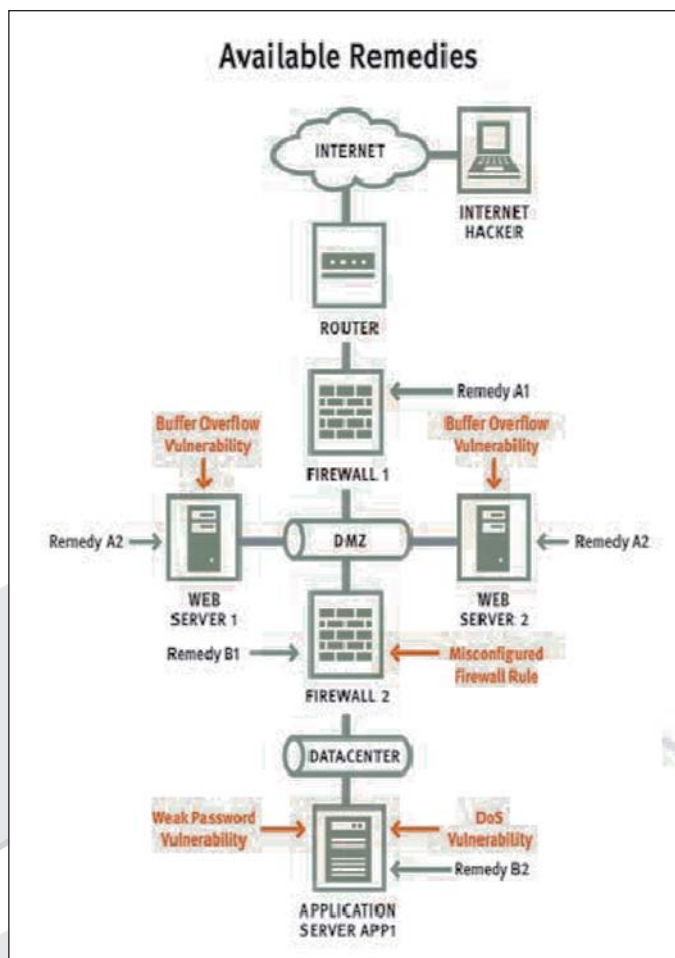


Figure 6: Available remedies for the two sets of actions.

bank accounts. The organization may prefer to add a firewall rule (Remedy B1) instead.

Before choosing the optimal set of remedies, the term "optimal" must be defined. Optimal is relative to the goal function, which can vary between organizations and even within the same organization for different events. Common goal functions include:

- ▲ Shortest time to remediation (temporary or permanent)
- ▲ Minimum process costs
- ▲ Maximum risk reduction
- ▲ Minimum downtime and safest implementation
- ▲ Smallest dependency on other departments

By quantifying parameters, such as time and cost for each remediation action, remediation alternatives can be provided with total cost and total time scores. Therefore, finding the optimal plan given the selected goal function is now possible.

More advanced remediation strategies can be employed as well:

**Defense in depth**—Selection of a few complementary sets of remedies in order to form several lines of defense.

**Defense in steps**—Selection of several sets of remedies according to different goal functions. For example, if a quick workaround is required to mitigate a fast-spreading worm, then one set of actions can be selected for the shortest time to remediation (e.g., by changing one or two firewall rules), and another set of actions can be selected for longer-term maxi-

mum risk reduction, e.g., patching all relevant servers, which takes much longer to complete.


## Simulating Planned Changes

Changes to networks—and the servers and applications within them—can be catastrophic to business operations if not done with caution. A remediation plan may look very compelling from a cost or time perspective, but if not planned appropriately, it can result in damage incurred rather than in time or money saved.

After selecting a remediation course, the planned changes must be validated in terms of their actual risk reduction—and to ensure that the changes do not break any business or technical constraints. By applying the planned changes to a futuristic, virtualized model of the IT environment, the effects of planned actions can be visualized, in a ‘what-if’ manner. First, the enterprise can use the remediation management engine to simulate the remedial actions—such as firewall configuration changes or software upgrades due to patching—and create a futuristic model of the IT environment. Then, based on this new model, the enterprise can utilize the attack simulation technology to determine whether risk has actually been reduced to an acceptable level and whether the various components within the network are still functioning appropriately.

Predicting whether an application will be “broken” can be achieved by comparing the model to a given set of business rules, such as “e-banking must be available for use from the Internet” and “Service pack 3 is not allowed in the DMZ.”

### Summary

Automated selective remediation is clearly the necessary approach for streamlining today’s complex and often overwhelming vulnerability remediation process. Using attack simulation technology to not only find critical exposures but also plan optimal remediation and simulate planned changes, enterprises can ensure that the right steps are taken to remediate threats and vulnerabilities—quickly and effectively—while balancing the enterprise’s business and IT constraints. Implementing this automated process for remediation is critical in decreasing business risk and reducing both the hard and soft costs of the remediation process. 

---

*Gidi Cohen, CEO and co-founder of Skybox Security, is a recognized expert on the topic of security, risk management and analytic technologies. He is an experienced speaker on the topic of security and analytic technologies, most recently presenting at InfoSecurity 2004, TiEcon 2004 Security Spotlight and New York Information Security and Technology Exchange 2004.*