

# **Evolutionary Systems Design: Recognizing Changes in Security and Survivability Risks**

Howard Lipson

*September 2006*

**CERT**

Unlimited distribution subject to the copyright.

**Technical Note**  
CMU/SEI-2006-TN-027

The Software Engineering Institute is a federally funded research and development center sponsored by the U.S. Department of Defense.

Copyright 2006 by Carnegie Mellon University.

#### NO WARRANTY

THIS CARNEGIE MELLON UNIVERSITY AND SOFTWARE ENGINEERING INSTITUTE MATERIAL IS FURNISHED ON AN "AS-IS" BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.

Use of any trademarks in this report is not intended in any way to infringe on the rights of the trademark holder.

Internal use. Permission to reproduce this document and to prepare derivative works from this document for internal use is granted, provided the copyright and "No Warranty" statements are included with all reproductions and derivative works.

External use. Requests for permission to reproduce this document or prepare derivative works of this document for external and commercial use should be addressed to the SEI Licensing Agent.

This work was created in the performance of Federal Government Contract Number FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center. The Government of the United States has a royalty-free government-purpose license to use, duplicate, or disclose the work, in whole or in part and in any manner, and to have or permit others to do so, for government purposes pursuant to the copyright license under the clause at 252.227-7013.

For information about purchasing paper copies of SEI reports, please visit the publications portion of our Web site (<http://www.sei.cmu.edu/publications/pubweb.html>).

---

# Contents

<b>Abstract.....</b>	<b>v</b>
<b>1 Introduction .....</b>	<b>1</b>
<b>2 Assumptions Evolve, and So Must Software .....</b>	<b>2</b>
2.1 Assumption Mismatches and the Impact of Change .....	3
<b>3 Recognizing the Need for Evolutionary Design Activity .....</b>	<b>4</b>
3.1 Change Factors .....	4
<b>4 Evolutionary Design Activities .....</b>	<b>10</b>
<b>5 Research Needs in Evolutionary Design .....</b>	<b>12</b>
<b>References.....</b>	<b>14</b>



---

## List of Tables

Table 1:	Factors that Influence Evolutionary Design of Secure Systems.....	5
Table 2:	Possible Evolutionary Design Activities in Response to a Trigger Event .....	10



---

## Abstract

A fundamental truth of system design is that, in the absence of countermeasures, a system's security and survivability will degrade over time. Changes in the environment or usage of a system, or changes to the elements that compose the system, often introduce new or elevated threats that the system was not designed to handle and is ill-prepared to defend itself against. The first step in evolving to meet new threats to your system's security and survivability is to *recognize* the need to modify your system—that is, to recognize changes in security and survivability risks that trigger the need to enter the evolution phase of the system development life cycle.

It is essential that significant risk management resources be devoted to the ongoing evolution of any mission-critical system. The successful evolutionary design of a secure and survivable system is dependent on the continual monitoring of the system and its environment to detect changes that may affect the risk management assumptions on which the system's security and survivability are founded.





---

# 1 Introduction<sup>1</sup>

A fundamental truth of system design is that, in the absence of countermeasures, a system's security and survivability will degrade over time. This degradation occurs not because any bits "rust out" or because the system shows any other manifestations of physical aging. Rather, changes in the environment or usage of a system, or changes to the elements that compose the system, often introduce new or elevated threats that the system was not designed to handle and is ill-prepared to defend itself against. Since security and survivability are system-wide properties, successfully dealing with such changes often requires revisiting every phase of the system development life cycle (SDLC) to at least some degree and poses particularly critical challenges for the assembly and integration phase.

The system you have assembled and integrated from vendor and custom components must evolve in response to a myriad of environmental changes, but the first step in evolving to meet new threats to your system's security is to *recognize* the need to modify your system—that is, to recognize changes in security and survivability risks that trigger the need to enter the evolution phase of the SDLC. The example in the next section (although not about a security failure) illustrates the critical importance of recognizing the need for evolutionary design changes.

---

<sup>1</sup> This technical note is an updated version of an article [Lipson 06] originally published on the DHS Build Security In (BSI) Web site earlier this year, in the "Assembly, Integration, and Evolution" content area (<https://buildsecurityin.us-cert.gov/daisy/bsi/articles/best-practices/assembly.html>). These and any future updates will be reflected on the Web site.

---

## 2 Assumptions Evolve, and So Must Software

What has become a classic example of the catastrophic consequences that can result from a simple software error<sup>2</sup> was the explosion of the unmanned Ariane 5 rocket during the first minute of its maiden flight on the morning of June 4, 1996 [ESA-CNES 96]. The explosion destroyed the launch vehicle's payload (a set of scientific satellites) worth on the order of \$500 million.

The Ariane 5's flight control software reused design specifications and code from its highly successful predecessor, the Ariane 4 launch vehicle. In particular, one of the on-board modules, the Inertial Reference System, performed a data conversion of a 64-bit floating point value related to the horizontal velocity of the rocket and attempted to place the result into a 16-bit signed integer variable. This computation had never caused a problem with the Ariane 4, but the more aggressive flight path and much faster acceleration of the Ariane 5 produced a higher horizontal velocity and a corresponding data value that was too large for the 16-bit signed integer variable, causing an arithmetic overflow. A redundant backup process used the same software and failed in the same manner. The Inertial Reference System then generated some diagnostic output that was incorrectly interpreted as flight control data by other portions of the flight control system. Based on this faulty interpretation, the flight control system took actions that led to the self-destruction of the rocket.

Although arithmetic overflow is a very well-known and highly preventable error, the Ariane 4 design team did not add the exception-handling code necessary to check for arithmetic overflow and take appropriate remedial action. Based on the operating characteristics of the Ariane 4, the design team felt it was physically impossible to have a horizontal velocity large enough to cause an arithmetic overflow of a 16-bit signed integer variable. However, the reuse of this software in the Ariane 5 placed the code in a very different operating context in which the specific design assumption relating to horizontal velocity was no longer valid. Although the operating characteristics of the rocket had evolved, the underlying design assumptions based on those characteristics were not revisited by the design or testing teams, and so the software did not evolve to reflect its new operating environment. The assumptions on which the design was based no longer reflected reality.

In all there were 14 recommendations by the Flight 501 Failure Inquiry Board [ESA-CNES 96]. Two of these had especially strong implications for software evolution, software architecture, code reuse, and the design of commercial off-the-shelf (COTS) based systems: recommendation 5 (first two bullet items) and recommendation 12:

---

<sup>2</sup> A simple arithmetic overflow doomed the Ariane 5, but taken in its operational and software engineering context, the circumstances surrounding the error were complex.

**R5** Review all flight software (including embedded software), and in particular:

- Identify all implicit assumptions made by the code and its justification documents on the values of quantities provided by the equipment. Check these assumptions against the restrictions on use of the equipment.
- Verify the range of values taken by any internal or communication variables in the software.
- ...

**R12** Give the justification documents the same attention as code. Improve the technique for keeping code and its justifications consistent.

## 2.1 Assumption Mismatches and the Impact of Change

Mismatches in the basic assumptions (and in particular the risk-management assumptions) on which the design of a system is based have historically been a fundamental cause of countless security, safety, and survivability problems. Architectural mismatches among components are a nearly universal source of problems [Garlan 95] and are a direct result of mismatched assumptions. The security and survivability of COTS-based systems (and other forms of software reuse) suffer from mismatches between the assumptions made by the COTS software designers and the assumptions made by the system integrators [Lipson 02]. Invalid assumptions made by designers about the real-world operating environment are another cause of system failures.

However, even if all such assumptions were correct and were perfectly matched during the initial design, implementation, and initial deployment, the facts or circumstances on which some of these assumptions are based will invariably change over time. Any security (or other) problem arising from these changes can be considered to be a case of evolution failure—the failure of a system’s designers to evolve the system in a manner that properly reflects the impact of changes (e.g., in technology, operating environment, and business mission) on its underlying assumptions.<sup>3</sup> Whether creating new systems and components or reusing existing ones, designing a system for evolution is a key aspect of building security in (or building quality in) from the outset.<sup>4</sup>

---

<sup>3</sup> Hence, assumption mismatches occur not only across architectures, components, and systems, they also occur over time.

<sup>4</sup> Although this example is specifically about reuse, the lessons learned apply to the evolution of an aging component or system as well.

---

## 3 Recognizing the Need for Evolutionary Design

### Activity<sup>5</sup>

The strongest implication of the concept of evolutionary design is that the sustainment of any mission-critical system requires *perpetual design*. That is, at least to some extent, all SDLC activities must be perpetual if the quality attributes of a system are to be sustained over time. In addition to evolving a system and its components, it is also crucial for assurance cases (i.e., assurance arguments composed of artifacts and other evidence of assurance of desired system properties) to evolve as well. System characteristics that hinder or promote evolution are discussed in *Topics in Interoperability: System-of-Systems Evolution* [Carney 05]. Some fundamental “laws” of software evolution are described in “Rules and Tools for Software Evolution Planning and Management” [Lehman 01].

It is essential that significant risk management resources be devoted to the ongoing evolution of any mission-critical system. The successful evolutionary design of a secure and survivable system<sup>6</sup> is dependent on the continual monitoring of the system and its environment to detect changes that may affect the risk management assumptions on which the security and survivability of the system are founded.

Any significant change in system requirements can certainly affect the underlying risk management assumptions, but the effects of other changes might not be as obvious. Therefore, one of the most essential uses for risk management resources would be to support security and survivability monitoring to provide early warnings of emerging threats and increased risks to the system. The amount of resources to be devoted to this activity, and to those that conduct it, will depend on executive management’s risk tolerance and their perception of the cost/benefit ratio for this effort.

### 3.1 Change Factors

We use the term *risk assessment triggers* to refer to the elements of a system or its environment that should be monitored, looking for changes that can affect the risk management assumptions that underlie a system’s security and survivability properties.

---

<sup>5</sup> This is a revised and updated version of one of the sections contributed by Howard Lipson for “Managing Software Development for Survivable Systems” [Mead 01].

<sup>6</sup> A survivable system is one that continues to fulfill its mission despite an attack, accident, or subsystem failure. Survivability blends security and business risk management [Lipson 99] and is based on ensuring that the quality attributes that are critical to the success of an organizational mission, such as availability and reliability, are sustained. The overall level of service may gracefully degrade under stress, but a survivable system continues to provide the essential services that support the organizational mission.

Ideally, a best practice for system design would require that all such assumptions be explicitly specified in a design rationale document or in other system artifacts, but typically many such assumptions are merely implicit. Nevertheless, if during the lifetime of a system any of the assumptions on which its design was based no longer hold, the mission-critical properties of the system (in particular its security properties) must be reevaluated. It is therefore critical for management and the system design team to be made aware of any event or change that appears (or has the potential) to undermine one or more of those risk management assumptions. However, it is up to management and the system design team to determine whether a particular change or set of changes should trigger an evolutionary design activity and to decide on the extent of that activity.

Table 1 contains a representative set of risk assessment change factors (trigger elements) that might be tracked by an organization. Trigger events include changes in attack techniques, mission, management, staff, customers, and in the technological and legal environments. They also include changes to the elements that compose your system and changes to other systems with which your system is involved in a system-of-systems relationship.

*Table 1: Factors that Influence Evolutionary Design of Secure Systems*

<b>Change Factors (Triggers)</b>	<b>Examples of Trigger Events</b>
<b>Business and Organizational</b>	
<i>Mission, essential services, essential quality attributes, key information resources and assets</i>	The organization's mission has changed or the system will be purchased and deployed by other organizations with different missions.
<i>Business strategies and tactics</i>	Changes in business strategies or tactics may require new types of data to be collected and processed, as well as increased connectivity among elements of the system, imposing new security requirements.
<i>Management</i>	New executive managers may differ in their tolerance for risk and their risk management strategies.
<i>Organizational staff</i>	Turnover may result in a lowering of staff expertise, which reduces the organization's ability to handle the human processes associated with security, such as properly configuring systems for optimal security. Moreover, in a rapidly growing organization, new staff may be less trustworthy than previous staff (e.g., there may be less time for background checks, or there may be more remotely stationed employees).
<i>Workflow and processes</i>	Changes in organizational processes to which the system contributes may affect the overall survivability of the mission. There may be new ways to attack the system or its human-machine interface.

<b>Change Factors (Triggers)</b>	<b>Examples of Trigger Events</b>
<i>Customers</i>	New customers may be less known (and hence less trustworthy), may require more extensive access to information resources and assets, or may require a higher quality of service (e.g., higher availability) than previous customers.
<i>Collaborators</i>	A new or existing collaborator may require a deeper level of integration with your business processes than your system currently supports. Or your partner on one project may become your competitor on another, requiring a more complex trust model.
<i>Competitors</i>	Business competitors may offer new services to your customers that your system currently cannot provide.
<i>Usage, functionality, access, or quality of service</i>	User requests for a new means of access to a system (e.g., wireless networking), new ways of using an existing system, the introduction of a service, or improvements in the quality of an existing service have security and survivability implications that need to be considered in any design activity undertaken in response to those requests. For example, a manufacturing plant that will now be handling a new and particularly volatile chemical ingredient requires an evolutionary redesign to improve the security and safety of the plant's control systems.
<b>Threat Environment<sup>7</sup></b>	
<i>Attack techniques</i>	A new attack technique or variation has been discovered for which the system cannot adapt automatically or through routine maintenance (e.g., simply by adding a new rule for resistance, recognition, or recovery).
<i>Malicious adversaries</i>	Awareness of industrial competitors engaging in espionage (or sabotage), growth in criminal activity, or increases in nation-state-sponsored cyber terrorism may require additional system resources to be devoted to security and survivability.

<sup>7</sup> This category is meant to represent the malicious threat environment only. Hence, ordinary business competitors who essentially engage in fair play are considered to be part of the Business and Organizational category. Nonetheless, they can certainly threaten the long-term survival of other businesses through aggressive marketplace competition that stays within the rules and the law.

<b>Change Factors (Triggers)</b>	<b>Examples of Trigger Events</b>
<b>Operating Environment</b>	
<i>Technology environment</i>	Changes in the technological environment in which the system operates (e.g., changes in the systems environment, the availability of new security tools and techniques, technological advances in the state of the practice for the application domain, and the increasingly widespread dissemination of detailed knowledge about the application domain and its supporting technologies) can trigger the need for evolutionary improvements in system security and survivability.
<i>Physical environment</i>	The migration from wired desktops situated in a physically secure environment to laptops and other wireless mobile devices that routinely traverse insecure and even hostile environments increases the possibility of both physical and cyber theft. This intensifies the need to strengthen the protection of enterprise-sensitive data resident on those devices, as well as access to enterprise databases and services through those devices.
<b>Economic Environment and the Acquisition Marketplace</b>	
<i>Cost, profit, and affordability</i>	Changing cost factors may threaten or improve a system's survivability because they change the cost/benefit ratio associated with various survivability solutions (e.g., risk mitigation strategies). Affordability is a primary factor that is traded off against security and survivability. For instance, new technology could provide a replacement for an existing component at a much lower price. Greatly reduced component cost could trigger an evolutionary redesign, using multiple instances of the new component (possibly from multiple vendors) to provide increased redundancy and diversity, thereby supporting greater survivability. As another example, increased stockholder demands for short-term profits may tilt the security and survivability requirements toward higher risk, which may be reflected in cutbacks in security administrators or vendor maintenance contracts.
<i>Vendors and contractors</i>	A new vendor for a system component may require remote maintenance and trusted access.
<i>COTS products</i>	You may have to replace a COTS component that is no longer supported with a new component whose contribution to system security and survivability needs to be evaluated.

<b>Change Factors (Triggers)</b>	<b>Examples of Trigger Events</b>
<b>Political, Social, Legal, and Regulatory Environment</b>	
<i>Legal environment</i>	New laws, increased enforcement of existing laws, and lawsuits can change the risk equation and threaten the mission. For instance, use of the system in a new and stricter jurisdiction may increase the risk of liability, which might be mitigated by strengthening certain aspects of the system's security.
<i>Government regulation</i>	Changes in government regulations that mandate increased privacy, security, safety, competition, or quality of service may trigger the need to modify a system's design to ensure that these new requirements are specified and satisfied.
<i>Certification requirements or standards</i>	Customers, regulators, and insurers may expect a system to be modified to the extent necessary to comply with new (or changed) standards or certification requirements, so as to reduce the actual or perceived level of risk associated with operating a system in a particular domain or environment. For example, business interruption insurance rates that include cyber attack may depend on a certification of the security and survivability of a system (i.e., the presentation of sufficient evidence to demonstrate that the system meets a given standard).
<i>Political and social environment</i>	Changes in privacy concerns, trust relationships, or the risk tolerance of a society will affect the security and survivability requirements that systems are expected to satisfy.
<b>Relationships to Other Systems and Infrastructures</b>	
<i>Dependencies and interdependencies</i>	New interconnections among the systems within an enterprise can create single points of failure, such as multiple systems relying on a single service. Increased dependency on a system may also be brought about by the elimination of manual processes, staff positions, or legacy systems, which means there is no longer an alternative if the system fails. Increasing the interdependencies within an enterprise may mean that a failure is more likely to have pervasive effects (e.g., cascading failures). Moreover, a mismatch in security models among the interconnected systems can readily cause a violation of security requirements.
<i>Usage relationships</i>	Changes to systems that depend on your system (and of course changes to any system that your system depends on) may require evolutionary changes to your system to sustain the security and survivability of the overall system of systems.



Change Factors (Triggers)	Examples of Trigger Events
<b>System Feedback (Lessons Learned)</b>	
<i>System instrumentation and audits</i>	System logs allow the operations team to monitor and improve the security and survivability of the system while it is in use (e.g., through configuration changes). Further analyses of this data may be used to identify survivability issues and to improve security and other system quality attributes in future releases of the system. This analysis may also identify gaps in system instrumentation and the need for improvements in the quality or coverage of system logs and in the frequency or quality of audits.
<i>Operational experience (attacks, accidents, and failures)</i>	Feedback from the field may lead to the discovery of new threats to a system's security and survivability or may reveal existing deficiencies.
<i>Results of periodic security and survivability evaluations</i>	Troublesome results from regularly scheduled penetration testing or other security and survivability evaluations can trigger awareness of the need for evolutionary improvements.
<i>Technical society meetings, security courses, seminars, journals, news reports</i>	Awareness of lessons learned by others' system failures and compromises can trigger improvements in your own system.

---

## 4 Evolutionary Design Activities

A change in one or more of the trigger elements can initiate any of a broad range of evolutionary design activities described in Table 2, from no action at all, to performing one or more system development life cycle activities, to abandonment of the system. The organizational unit responsible for monitoring for changes in risk management assumptions would initiate the consideration of an evolutionary design activity, but management and the system design team would be responsible for evaluating the impact of any trigger event that was identified and for determining the scope of any subsequent design activity in response to that event.

*Table 2: Possible Evolutionary Design Activities in Response to a Trigger Event*

Evolutionary Design Activity	Example
No action needed or taken	Conclude that greatly increased hiring activity does not pose a new threat to the system's mission because all new hires are subject to thorough background checks.
No action taken, but increase monitoring of this trigger (or set of triggers)	Increase resources devoted to monitoring feedback from the field in response to evidence from operations indicating a performance slowdown resulting from a rare combination of customer actions.
Further analysis needed to determine next activity, if any	Generate scenarios that reflect the discovery of a new type of cyber attack. Use these scenarios as input for a security analysis, the results of which may drive additional evolutionary design activities.
Perform a portion (delta) of one or some of the system development life cycle activities	A small change to the system architecture increases resistance to a new attack scenario.
Perform a portion (delta) of each of the full set of life cycle activities	A modification to the mission touches all life cycle activities to one extent or another.
Do a full redesign	A major change in the technology of the application domain, coupled with sweeping improvements in defensive technology, cannot be incorporated by evolutionary design activities alone.
Abandon the system	A drastic change in the mission makes the system obsolete or unnecessary.

For example, a computer security expert involved in risk assumption monitoring learns of a new attack technique that might threaten the security and survivability of the existing system. Let's assume that this new attack technique cannot be countered by straightforward maintenance activities such as applying a security patch to a system component or adding a new rule to a firewall. Based on the new attack technique, the security expert generates a set of attack scenarios to be used as input for a security and survivability analysis of the existing system. If deficiencies in the system's resistance to this new attack (or in the system's ability to recognize or recover from the attack) are discovered, then one or more life cycle activities, such as a modification of the system architecture or a change in security and survivability requirements, will be necessary.

The completion of one life cycle activity may trigger the need for another. Adjustments in the design tradeoffs with other system quality attributes may also be called for. For example, a specific architectural change meant to improve security may have unanticipated adverse effects on some other system quality attributes. These implicit tradeoffs can be systematically evaluated and explicitly adjusted using the results of an architecture tradeoff analysis [Kazman 98]. The point at which the evolutionary design process stops is dependent on the risk tolerance of the organization, and the perceived cost/benefit ratio, with respect to the particular set of trigger events. If evolution is not feasible, the organization may tolerate the risk or seek other alternatives that transcend the system.

It is essential that the evolutionary design activities take place in the context of full access to a comprehensive set of artifacts of the design process (such as descriptions of the rationale for tradeoffs made during the last design cycle). Continuity of at least the core members of the design team is particularly crucial for the evolutionary design of survivable systems so that the mission-specific design expertise can be sustained throughout the life of the system. Otherwise, the evolutionary design process will likely degenerate into patching, which can never support the long-term security of systems. Just as security must be designed into a system from the beginning and not tacked on later as an afterthought, long-term security cannot be sustained through patching or routine maintenance but only through the continual incorporation of new security and survivability solutions through a principled evolutionary design process. The development and promulgation of a suite of best practices to support this process would be a fundamental contribution to the software engineering profession.

---

## 5 Research Needs in Evolutionary Design

While evolutionary design is a critical aspect of building security in, there are few best practices for evolution that are supported by ample evidence and general consensus in the software engineering community. Those practices that do exist are typically classified under software maintenance.

Moreover, many aspects of evolutionary design are not yet well understood by the software engineering community. For example, it is not practical (i.e., not economically feasible) for a system to evolve along all of the possible dimensions outlined in Table 1. How to decide during the initial design of a system which dimensions of change are most likely and how to make them amenable to low-cost redesign or automated upgrades is an important topic for further research and investigation. Some limited success in this area has been achieved through automatic upgrades of firewall rules and databases of attack signatures for detecting and eliminating viruses and other malware—essentially rapidly evolving a system in response to evolving security threats.

One of the most problematic aspects of the evolution of a secure and survivable system is how to recognize changes in security and survivability risk that arise as a result of your system's dependence on components, services, and systems (in a system-of-systems environment) whose design, implementation, and operation are not in your direct control. For example, changes in a vendor's development processes, underlying technology, personnel, ownership, or suppliers may increase security risks to your own system. Risks introduced or increased by changes in the elements that your system is dependent upon can be addressed in a variety of ways—often at the architectural level—but, of course, only if you are aware that external changes may have affected your earlier risk management assumptions.

An emerging area of research and engineering known as *assurance cases* has great potential for helping system designers, integrators, architects, developers, acquisition personnel, and other stakeholders deal with the changing risks associated with the evolution of components, services, and systems.<sup>8</sup> This would include a typical situation in which components are being integrated into a system by an organization that is not in control of the components' development or evolution. Assurance cases use *arguments* that refer to bodies of *evidence* to demonstrate that specific *claims* about a system (or component) are satisfied (e.g., with respect to a particular quality attribute). Assurance cases have been used successfully in the safety domain (where they are known as *safety cases*) and efforts are underway in the research community to extend the use of assurance cases to the security domain [Bloomfield 06a, Bloomfield 06b].

---

<sup>8</sup> A new content area on assurance cases is being developed for the BSI Web site (<https://buildsecurityin.us-cert.gov/>) and should appear by the end of this year.

If the development of assurance cases for security were integrated into the system development life cycle as a best practice, then assurance cases would evolve along with the component or system being developed, and stakeholders could be notified of any change in an assurance case. One can envision an environment in which assurance cases would be the key to tracking the evolution of a system, and the elements that the system is dependent upon, with respect to the ongoing satisfaction of security claims. Assurance cases could then be used as a means of recognizing those evolutionary changes that impact the system's security and survivability risks. Research is needed not only to create and refine methods and tools for developing assurance cases for security but also to ensure that assurance cases can be developed in a cost-effective manner in order to provide practical support for evolutionary design.

In summary, further research and reduction to practice is urgently needed on a broad range of evolutionary design considerations for building secure and survivable systems, including evolutionary design principles and candidate best practices solicited from the software engineering community. The goal is to begin to formulate a set of demonstrably useful, highly actionable best security practices to support evolutionary design. Ensuring that you can effectively evolve the security capabilities of a system in response to a changing risk environment is an essential part of building security in—that is, building in the capability to evolve and improve the security and survivability of a system over its full lifetime of use.

---

## References

*URLs are valid as of the publication date of this document.*

- [Bloomfield 06a]** Bloomfield, R. E.; Guerra, S.; Masera, M.; Miller, A.; Saydjari, O. S. *Assurance Cases for Security Workshop Report*, Version 01c. Workshop on Assurance Cases for Security, June 13-15, 2005, Arlington, VA.  
[http://www.csr.city.ac.uk/AssuranceCases/Assurance\\_Case\\_WG\\_Report\\_180106\\_v10.pdf](http://www.csr.city.ac.uk/AssuranceCases/Assurance_Case_WG_Report_180106_v10.pdf) (2006).
- [Bloomfield 06b]** Bloomfield, R. E.; Guerra, S.; Miller, A.; Masera, M.; & Weinstock, C. B. "International Working Group on Assurance Cases (for Security)." *IEEE Security & Privacy* 4, 3 (May-June 2006): 66-68.
- [Carney 05]** Carney, David; Fisher, David; & Place, Patrick. *Topics in Interoperability: System-of-Systems Evolution* (CMU/SEI-2005-TN-002). Pittsburgh, PA: Software Engineering Institute, Carnegie Mellon University, 2005.  
<http://www.sei.cmu.edu/publications/documents/05.reports/05tn002/05tn002.html>
- [ESA-CNES 96]** European Space Agency (ESA) and National Center for Space Study (CNES) Inquiry Board (Prof. J. L. Lions, Chairman). *ARIANE 5 – Flight 501 Failure – Report by the Inquiry Board*. Paris: ESA and CNES, July 19, 1996.  
[http://en.wikisource.org/wiki/Ariane\\_501\\_Inquiry\\_Board\\_report](http://en.wikisource.org/wiki/Ariane_501_Inquiry_Board_report)
- [Garlan 95]** Garlan, David; Allen, Robert; & Ockerbloom, John. "Architectural Mismatch: Why Re-use Is So Hard." *IEEE Software* 12, 6 (November 1995): 17-26.
- [Kazman 98]** Kazman, R.; Klein, M.; Barbacci, M.; Longstaff, T.; Lipson, H. F.; & Carriere, S. J. "The Architecture Tradeoff Analysis Method." *Proceedings of the Fourth IEEE International Conference on Engineering of Complex Computer Systems (ICECCS 1998)*. Monterey, CA, USA, August 10-14, 1998. Los Alamitos, CA: IEEE Computer Society Press, 1998.  
<http://www.sei.cmu.edu/ata/iceccs.pdf>

- [Lehman 98]** Lehman, M. M. "Software's Future: Managing Evolution." *IEEE Software* 15, 1 (January/February 1998): 40-44.
- [Lehman 01]** Lehman, M. M. & Ramil, Juan F. "Rules and Tools for Software Evolution Planning and Management." *Annals of Software Engineering* 11, 1 (November 2001): 15-44.
- [Lipson 99]** Lipson, Howard & Fisher, David. "Survivability—A New Technical and Business Perspective on Security," 33–39. *Proceedings of the 1999 New Security Paradigms Workshop*. Caledon Hills, Ontario, Canada, Sept. 22–24, 1999. New York: Association for Computing Machinery, 2000. <http://www.cert.org/archive/pdf/busperspec.pdf>
- [Lipson 02]** Lipson, H. F.; Mead, N.; & Moore, A. P. "Can We Ever Build Survivable Systems from COTS Components?" *Proceedings of the 14th International Conference on Advanced Information Systems Engineering (CAiSE' 02)*. Toronto, Ontario, Canada, May 27-31, 2002. Heidelberg, Germany: Springer-Verlag (LNCS 2348), 2002.
- [Lipson 06]** Lipson, Howard. "Evolutionary Design of Secure Systems – The First Step Is Recognizing the Need for Change." Build Security In Web site, April 2006.  
<https://buildsecurityin.us-cert.gov/daisy/bsi/articles/best-practices/assembly/467.html>
- [Mead 01]** Mead, N. R.; Linger, R. C.; McHugh, J; & Lipson, H. F. "Managing Software Development for Survivable Systems." *Annals of Software Engineering* 11, 1 (November 2001): 45-78.





<b>REPORT DOCUMENTATION PAGE</b>			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave Blank)		2. REPORT DATE September 2006		3. REPORT TYPE AND DATES COVERED Final
4. TITLE AND SUBTITLE Evolutionary Systems Design: Recognizing Changes in Security and Survivability Risks			5. FUNDING NUMBERS FA8721-05-C-0003	
6. AUTHOR(S) Howard Lipson				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Software Engineering Institute Carnegie Mellon University Pittsburgh, PA 15213			8. PERFORMING ORGANIZATION REPORT NUMBER CMU/SEI-2006-TN-027	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) HQ ESC/XPK 5 Eglin Street Hanscom AFB, MA 01731-2116			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES				
12A DISTRIBUTION/AVAILABILITY STATEMENT Unclassified/Unlimited, DTIC, NTIS			12B DISTRIBUTION CODE	
13. ABSTRACT (MAXIMUM 200 WORDS)  A fundamental truth of system design is that, in the absence of countermeasures, a system's security and survivability will degrade over time. Changes in the environment or usage of a system, or changes to the elements that compose the system, often introduce new or elevated threats that the system was not designed to handle and is ill-prepared to defend itself against. The first step in evolving to meet new threats to your system's security and survivability is to <i>recognize</i> the need to modify your system—that is, to recognize changes in security and survivability risks that trigger the need to enter the evolution phase of the system development life cycle.  It is essential that significant risk management resources be devoted to the ongoing evolution of any mission-critical system. The successful evolutionary design of a secure and survivable system is dependent on the continual monitoring of the system and its environment to detect changes that may affect the risk management assumptions on which the system's security and survivability are founded.				
14. SUBJECT TERMS evolutionary systems design, computer security, system survivability, risk management, system development life cycle, SDLC			15. NUMBER OF PAGES 25	
16. PRICE CODE				
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL	