Survivability — A New Technical and Business Perspective on Security

Howard F. Lipson & David A. Fisher

hfl@cert.org +1-412-268-7237
dfisher@cert.org +1-412-268-7703

CERT[®] Coordination Center
Software Engineering Institute
4500 Fifth Avenue
Pittsburgh, PA 15213

http://www.cert.org/research/

Abstract

In recent years, there have been dramatic changes in the character of security problems, in their technical and business contexts, and in the goals and purposes of their stakeholders. As a consequence, many of the assumptions underlying traditional security technologies are no longer valid. Failure to recognize the depth and breadth of these changes in combination prevents effective solutions to modern security problems. Survivability provides a new technical and business perspective on security, which is essential to our search for solutions. Moreover, our survivability approach expands the view of security from a narrow technical specialty, accessible only to security experts, towards a risk-management perspective that requires the participation of an organization as a whole (executive management, security experts, application domain experts, and other stakeholders) to protect mission-critical systems from cyber-attacks, failures, and accidents.

1. Survivability from a Technical Perspective

Survivability has been defined as the capability of a system to fulfill its mission, in a timely manner, in the presence of attacks, failures, or accidents [EFL99], where the term "system" is used in the broadest possible sense, and includes networks and large-scale "systems of systems." Although this definition captures our concept of survivability in a succinct form, it does not clearly expose the rationale and implications of this view.

While security traditionally has been focused on confidentiality of information, the problems of greatest concern today relate to the availability of information and continuity of services. Concern for continuity of critical services among infrastructure providers, their customers, and cognizant government agencies has led to the new field of infrastructure assurance. The commercial viability of companies depends on their ability to produce and deliver their products and services in a timely manner. These are mission goals that go beyond, and must therefore extend, the traditional scope of security.

Most security technology depends on certain underlying assumptions about the nature and structure of systems [Bla96]. Generally, these include assumptions that systems are closed, that they are under central (or unified) administrative control, and that administrators have the ability to observe any given activity within the system. These assumptions may have been appropriate when systems were isolated islands with highly controlled interfaces to the rest of the world. Today, however, systems are open, and no one person or organization has full administrative control. Observers, whether they are inside or outside the system, have only limited visibility into the structure, extent, or topology of the system. Lack of central administrative control and absence of global

[®] CERT and CERT Coordination Center are registered in the U.S. Patent and Trademark Office.

1999 New Security Paradigm Workshop 9/99 Ontario, Canada
© 2000 ACM 1-58113-149-6/00/0004...$5.00

visibility are properties of the Internet, of any distributed application residing on the Internet, and of any infrastructure in a deregulated industry. Now prevalent, these unbounded[1] systems with neither central control nor global visibility are incompatible with the assumptions underlying our extant security technologies.

Survivability often involves tradeoffs among several software quality attributes. This has led some people to incorrectly conclude that survivability is synonymous with dependability. Composite software quality attributes, such as dependability, involve tradeoffs among a fixed set of software quality attributes that are, in this case, determined by the definition of dependability. The tradeoffs involved in survivability are among the functional and nonfunctional requirements [Ebe97] determined by the mission. The software quality attributes critical to one mission or application may be irrelevant to another application. Also, a traditional methodology for achieving dependability (and security) is to ensure certain quality attributes in the components of a system, and then to rely upon a composition process that will preserve those qualities in the system as a whole. For survivability and safety, in general, no such composition process exists. In fact, a fundamental assumption of survivability is that no component is immune from compromise, accident, or failure. Instead, the functional and non-functional global properties of a survivable system must emerge by virtue of the composition process from unsurvivable components. We therefore recognize that survivability is an *emergent property*[2] that cannot be achieved at the level of atomic system components, because each component represents a single point of failure for its own survival. A well-known example of the concept of emergent properties is the creation of reliable systems from components that are less reliable than the composite system. This is somewhat analogous to the creation of survivable systems from unsurvivable components.

Most security technologies derive from a fortress model in which there is a clear distinction between trusted insiders and other potential users and intruders. In the highly distributed applications and Internet-based systems of today there is little distinction between insiders and outsiders. Everyone who chooses to connect to the Internet is an insider, whether or not they are known to a particular subsystem. This characteristic is derived from the desire, and modern necessity, for connectivity. Companies cannot survive in highly competitive industries without easy and rapid access to their customers, suppliers, and partners. More and more, your partners on one project are competitors on the next, so that trust becomes an extremely complex concept. Trust relationships are continually changing, and in traditional terms may be highly ambiguous. Trust is especially difficult to establish in the presence of unknown users from unknown sources outside one's own administrative control.

A fortress model is only as strong as its weakest component. If a trusted insider abuses his or her authority, or an intruder finds an exploitable vulnerability in a security perimeter, the entire system can be compromised. In unbounded networks where everyone is an insider and often unknown, there are always numerous untrustworthy insiders. Furthermore, fortress models do not allow for graceful degradation, nor the fail-soft and fail-safe mechanisms demanded by availability and mission objectives.

So, the differentiating characteristics between survivability and traditional security are the purposes and goals, the technical context, the business context, the technical constraints, the underlying assumptions for applicable technology, and the potential effectiveness of individual technologies. We should not be too surprised that, in combination, the altered purpose, context, stakeholders, and assumptions render existing security technologies less than satisfactory.

Not surprisingly, advances in security tools, methods, and practices in recent years have been dominated by attempts to modify and adapt modern systems to conform to the assumptions of traditional security technology. Thus, despite the proliferation of open, unbounded systems where there is little trust, firewalls continue to be the primary mechanism for survivability. Despite the ample evidence that far too many vulnerabilities (both known and yet-to-be-discovered) inhabit proprietary commercial off-the-shelf (COTS) software, and the fact that these products are widely available for analysis by potential attackers searching for exploitable weaknesses, we increasingly embrace such software for solutions. We readily incorporate COTS software as components of larger systems, which then fall prey to attacks based on the COTS components' vulnerabilities. Despite diversity being the single most effective mechanism for security and survivability in networked systems today, we continue to

[1] An *unbounded system* is any system whose purpose or mission must be achieved in the absence of complete or precise information about some aspects of the system, in the absence of centralized administrative control, or in the presence of untrustworthy insiders. Examples would be the Internet, any system with distributed administrative control without central authority, any system with remote access, any system with unknown users, and any system containing commercial-off-the-shelf (COTS) software. We refine the definition of unbounded system both formally and informally in [FL99].

[2] *Emergent properties* are properties that arise or emerge from the combined actions and interactions of the various components of a system and often do not, or cannot, prevail within individual components of the system. For details on how emergent properties relate to survivability, see [FL99]. For examples illustrating the role of emergent behavior in the composition process, see [Hin97].

deploy identical single-vendor implementations of brittle standards. We purchase software based on features and initial purchase cost, rather than robustness and long-term or indirect cost.

Still worse, as long as we fail to explicitly recognize and embrace the fact that these combined changes of purpose, context, and constraints require us to lay a foundation of new assumptions upon which to build new classes of solutions, technology will remain in conflict with effective solution approaches. As a consequence, it will be impossible to exploit the properties of the new problem domain for radical and unanticipated solutions.

Although beyond the scope of this position paper, we have confidence and some limited evidence that effective solutions to survivability problems in unbounded networks can arise from revised assumptions coupled with advances in diversity, robustness, adaptability, and algorithmic solutions (which we call *emergent algorithms*[3]) that generate predictable nonfunctional global properties from simple local interactions [FL99]. One of the central points of this paper is that the focus of security has moved sufficiently in several dimensions to justify new foundational assumptions, but in the absence of these new assumptions the community has failed to exploit the unique characteristics of the revised problem space with compatible, rather than conflicting, solution paradigms. Survivability provides a new technical and business perspective on security that can guide us towards a better understanding of the nature and structure of modern, highly distributed systems and can lead us to solutions to security problems that today seem intractable.

2. Survivability from a Business Perspective

Many businesses have contingency plans for dealing with business interruptions caused by natural disasters or accidents. Although the majority of cyber-attacks are relatively minor in nature, a cyber-attack on an organization's critical networked information systems has the potential to cause severe and prolonged business disruption, whether the business has been targeted specifically or is a random victim of a broad-based attack. If a cyber-attack disrupts critical business functions and interrupts the essential services that customers depend upon, then the survival of the business itself is at risk.[4]

Survivability is an emerging discipline [ISW97, ISW98] that blends computer security with business risk management for the purpose of protecting highly distributed information services and assets. A fundamental assumption is that no system is totally immune to attacks, accidents, or failures. Therefore, the focus of this new discipline is not only to thwart computer intruders, but also to ensure that mission-critical functions are sustained and a (situation-dependent) essential set of services is delivered, despite the presence of cyber-attacks. Improving survivability in the presence of cyber-attacks also improves the capacity to survive accidents and system failures that are not malicious in nature.

Traditional computer security is a highly specialized discipline that seeks to thwart intruders through technical means that are largely independent of the domain of the application or system being protected. Firewalls, cryptography, access control, authentication, and other mechanisms used in computer security are meant to protect an underlying application in much the same way regardless of the specific application being protected. In contrast, survivability has a very sharp mission focus, and is more akin to risk management than to the study of any technical aspect of software engineering, including individual or composite software quality attributes. Ultimately it is the mission that must survive, not any particular component of the system or even the system itself. The mission must go on even if an attack causes significant damage to or even destruction of the system that supports the mission. It is the shift toward risk management, an approach that is highly intertwined with the mission-specific features of the

[3] We define an *emergent algorithm* informally as an efficient distributed computation that generates and preserves those global system-wide properties that constitute the mission requirements for a system. These global properties include both functional and nonfunctional properties and are called "emergent properties" because they typically emerge from the combined actions and interactions of the various components of a system and because they often do not, or cannot, prevail within individual components of the system. Emergent algorithm is defined more formally in [FL99]. For survivability, we also consider only algorithms in which there are no single, nor constant number of, points-of-failure. Although necessary for survivability, this latter restriction precludes all algorithms that depend on centralized control, hierarchical decomposition, centralized data, or nodes with unique roles.

During the workshop discussion about our paper, Bob Blakley helped to clarify the meaning of "central control" in the context of our concept of emergent algorithms: "An emergent algorithm is survivable because you can pick a specified number of system components and destroy them, and the system will still fulfill its mission. This does not have central control as long as it doesn't matter which components you choose to destroy."

[4] One significant difference between disruptions caused by natural disasters and those caused by cyber-attacks (besides the notion of an intelligent adversary behind a cyber-attack) is that with a natural disaster there is a customer expectation of diminished service. A business disruption caused by a cyber-attack will likely be seen by a company's customers as a sign of incompetence. Unless the cyber-attack is widespread and well publicized, no customer sympathy will be forthcoming

application being protected, that is the most radical paradigm shift that is occurring as the new discipline of information survivability continues to emerge.

Survivability solutions are best understood as risk-management strategies that first depend on an intimate knowledge of the mission being protected. The mission focus expands survivability solutions beyond purely independent ("one size fits all") technical solutions, even if those technical solutions are broad-based and extend beyond traditional computer security to include fault tolerance, reliability, usability, and so forth. Risk-mitigation strategies first and foremost must be created in the context of a mission's requirements (prioritized sets of normal and stress requirements), and must be based on "what-if" analyses of survival scenarios. Only then can we look toward generic software engineering solutions based on computer security, other software quality attribute analyses, or other strictly technical approaches to support the risk-mitigation strategies.

Consider the analogy of a village farmer with the mission of supplying food to a village. The farmer may have a fence around the crops to keep out deer, rabbits, and other intruders (traditional security). The farmer may have an irrigation system to be used in the event of insufficient rainfall (redundancy). He or she may plant a variety of crops so that even if environmental conditions (e.g., pests) adversely affect one crop, others will survive (diversity). All of this is well and good. But even if all the crops fail and no food is grown, the mission can still succeed if the farmer has an alternate strategy based on the mission of providing food — *not* necessarily growing food using the local ecosystem. If the crops fail, the farmer may turn to hunting or fishing to provide the life-sustaining mission fulfillment that fellow villagers depend upon. Is hunting a security, reliability, or fault tolerance strategy? No — it is outside the system for growing food. This is a risk-management strategy that can only be formulated with an intimate understanding of the mission that must survive. Detailed technical expertise on fence-building, or even agriculture, is helpful but inadequate compared to strategies based on an intimate knowledge of the mission requirements.

Survivability depends not only upon the selective use of traditional computer-security solutions, but also upon the development of effective risk-mitigation strategies based on scenario-driven "what-if" analyses and contingency planning. "Survival scenarios" positing a wide range of cyber-attacks, accidents, and failures aid in the analyses and contingency planning. However, to reduce the combinatoric explosion of possibilities inherent in creating representative sets of survival scenarios, these scenarios focus on adverse effects rather than causes. Effects are also of more immediate situational importance than causes, because you will likely have to deal with (and survive!) an adverse effect long before a determination is made as to whether the cause was an attack, an accident, or a failure. Awaiting the outcome of a detailed post-mortem to determine the cause, before acting to mitigate the effect, is out of the question when dealing with the survival of most modern mission-critical applications.

Contingency (including disaster) planning requires risk-management decisions and economic tradeoffs that only executive management can make (preferably with guidance from technical experts in the application domain, in computer security, and in other software engineering or related disciplines). Survivability depends at least as much upon the risk-management skills[5] of an organization's management as it does upon the technical expertise of a cadre of computer security experts. This is certainly appropriate from an organizational perspective, because business risk management is a primary function of executive management, and not the role of computer-security experts or other technical gurus. Expertise in risk management and the organization's mission resides with that organization's management. The role of the experts in security, the application domain, and other technically relevant areas is to provide upper management with the information necessary to make informed risk-management decisions.[6] Thus, the preparatory steps necessary for survivability must be taken by an organization as a whole, rather than by security experts alone.

Let's consider the Galaxy-4 satellite that spun out of control on May 19, 1998, interrupting up to 90% of the pager service in the United States, along with some television network feeds, and some credit card verification services. Even though a cyber-attack was not to blame (though "an international hacker attack" was on an early list of speculative causes), the example is quite illustrative. In fact, the cause (or at least a partial cause — crystals forming on tin-plated relay contacts, and an unexplained failure of a backup system) was not determined until long after service was restored [Reu98].

Dealing with adverse events such as this one, without waiting for a definitive determination of the cause, is central to the survivability paradigm. Successful handling

[5] Here we are not referring to an abstract technical skill in the science of risk management, but rather to the ability to manage risk in the context of the specific business mission and goals.

[6] The primary role of the technical experts (in the security and application domains) is to ensure that the solutions that support alternative risk-management strategies are technically sound. For example, a ship's lifeboats provide a life-saving alternate means of buoyancy in the event that the ship sinks, but the lifeboats must be seaworthy and able to safely hold the expected number of passengers. Otherwise this survivability strategy is fatally flawed. Executive management's expertise in risk-management cannot replace technical expertise, but rather must build upon it.

of such events is far more dependent on prudent risk management and contingency planning by executive management than on any specific technical approach by security experts or other gurus. For instance, a "perfect" technical solution (i.e., having a diversely redundant, immediately available backup satellite) would be economically infeasible. The practicality of many technical solutions can only be evaluated in the full business context. Executive management, through its contingency planning, would consider business solutions that might transcend purely technical solutions. One approach would be to have an agreement in place with another communications company to provide the needed capacity upon, say, six hours' notice (with the backup company dumping its own lower-priority customers) in exchange for an annual fee. An alternate approach (using lawyers rather than technologists) would be to have a disclaimer in the contract agreement with customers telling them that the customer must bear the risk of service outages. This would put the customers on notice that they need to prepare to provide their own redundancy, whereas in the previous approach the service provider took care of redundancy through an agreement with an alternate provider. (Because it raises customer awareness of some of the risks inherent in the delivery of service and possibly increases the perceived value of uninterrupted service, the "legal disclaimer" approach might even generate some customer interest in asking the original service provider to provide redundancy for an additional fee.) The "legal disclaimer" approach is not one that technical experts would likely come up with, but it is quite effective in assuring the survivability of the business mission and goals. As this example illustrates, the risk-management viewpoint supports an "economics of survivability" that allows businesses to successfully prepare for and overcome the adverse effects of cyber-attacks, accidents, and failures with approaches that can transcend those offered by technical experts alone.

Contrast this new perspective with current management practice with respect to security. Upper management's primary decision-making role, from a traditional security viewpoint, is predominantly to determine how much direct funding and other resources to grant to the organization's security experts for the rather loosely defined purpose of "beefing up security" to some vaguely articulated industry standard level of practice. In the minds of management, the perceived link between security funding and the business mission (and the business bottom line) is tenuous at best. "If I spend more money on computer security my risk of intrusion will likely go down. But will this reduce any significant risks to my business mission? What risks will be reduced, and by how much?" With no clear benefit visible to management, the resulting security funding is typically inadequate to meet even the limited technical goals of the security experts. For the most part, what is sorely missing is an in-depth analysis of threats to the organization's mission and a corresponding cost-benefit analysis for risk-mitigation strategies and contingency planning. The computer-security experts, isolated from management's intimate understanding of the business mission, are in no position to perform the necessary threat analyses, except from the narrow perspective of their technical specialties.

As an example, consider the new government programs that are meant to assure that our nation's critical infrastructures will continue to operate despite cyber-attacks, accidents, or failures. Government concern for critical infrastructure assurance [Pre97] is helping to fuel the current interest in survivability, but this interest is not being driven by the businesses (such as those in energy, transportation, banking, and telecommunications) that would benefit from such protection. The government is asking industry to participate in critical-infrastructure assurance programs, with the motivation that these programs are in the best interests of the nation, the industries, and their customers. But none of these communities are willing to pay for the increased costs. Real investment in critical-infrastructure protection will occur only when executives understand that these changes are essential to their competitiveness and profitability. Unfortunately many of the businesses involved see these programs as mandating technical solutions, which would be at odds with their customers' needs and their own profitability. Greater awareness is needed of the business risk-management aspects of survivability, so that the organizations that operate our nation's critical infrastructures would be motivated by self interest to assure their own survivability. Critical-infrastructure assurance can then be based on risk-management tradeoffs that depend on overall business missions and goals, not solely on technical fixes that are independent of those goals.

In summary, there has been a revolutionary technical shift in business applications from stand-alone, closed systems over which organizations exercised complete control, to highly distributed, open, COTS-based systems over which only very limited control and limited insight are possible. Not only are most Internet services outside of the control of the businesses that use them, but so are the functionality and software quality attributes of the COTS-based software used to build business applications. This technical shift has taken us so far that we can no longer solve security problems entirely in the technical domain.

From the traditional computer-security perspective, executive management has never been sufficiently engaged. The security experts simply present a bill or funding request to management for generic technical solutions, independent of threat analyses that are specific to the mission being protected.

Executive management must be concerned with threats to the business mission, and must be intimately involved in formulating mission-specific risk-mitigation strategies. Moreover, technical experts need to be aware of the business issues that lead to the technical issues they face. Only then can they contribute effectively to the risk-management approaches that are needed to assure survivability of highly distributed mission-critical applications, operating in unbounded domains, in the face of cyber-attacks, accidents, and system failures.

3. Conclusion

In recent years there have been dramatic changes in the nature and structure of information systems. Traditional security solutions are not sufficient to deal with the modern security problems associated with highly distributed mission-critical systems, which have neither central administrative control nor global visibility. New foundational assumptions and new solution paradigms are necessary to protect such systems from cyber-attack.

Survivability is an emerging discipline that blends computer security with business risk management for the purpose of protecting highly distributed information services and assets. A fundamental assumption is that no system is totally immune to attacks, accidents, or failures. Therefore, the focus of this new discipline is not only to thwart computer intruders, but also to ensure that mission-critical functions are sustained and essential services are delivered, despite the presence of cyber-attacks, failures, and accidents.

Survivability solutions are best understood as risk-management strategies that first depend on an intimate knowledge of the mission being protected. The mission focus expands survivability solutions beyond purely generic ("one size fits all") technical solutions, even if those technical solutions are broad-based. Risk-mitigation strategies first and foremost must be created in the context of a mission's requirements (prioritized sets of normal and stress requirements), and must be based on "what-if" analyses of survival scenarios.

Survivability depends at least as much upon management's real understanding of their objectives as it does upon the technical expertise of their security experts. This is appropriate from an organizational perspective, because business risk management is a primary function of executive management, and not the role of computer-security experts or other technical staff. This in no way implies that survivability problems are not amenable to technical solutions, only that solutions require a partnership that can integrate technical and business considerations. It also means that survivability requires technical solutions that consider the true context in which a system must operate, and that any application-independent solutions will be inadequate.

Acknowledgments

We are grateful to the chairs, organizers, program committee, and participants of NSPW'99 for creating a workshop atmosphere in which new ideas could flourish.

References

[Bla96] Bob Blakley, "The Emperor's Old Armor," *Proceedings of the 1996 New Security Paradigms Workshop*, Lake Arrowhead, California, September 17-20, 1996, Association for Computing Machinery, 1997.

[Ebe97] C. Ebert, "Dealing with Nonfunctional Requirements in Large Software Systems," *Annals of Software Engineering*, Volume 3, September 1997, p. 367-395.

[EFL99] R.J. Ellison, D.A. Fisher, R.C. Linger, H.F. Lipson, T.A. Longstaff, N.R. Mead, "Survivable Systems: An Emerging Discipline," *Proceedings of the 11th Canadian Information Technology Security Symposium (CITSS)*, Ottawa, Ontario Canada, May 10-14, 1999, Communications Security Establishment, 1999.

[FL99] D. A. Fisher and H.F. Lipson, "Emergent Algorithms — A New Method for Enhancing Survivability in Unbounded Systems," *Proceedings of the 32nd Annual Hawaii International Conference on System Sciences*, Maui, Hawaii, January 5-8, 1999 (HICSS-32), IEEE Computer Society, 1999. Available at: http://www.cert.org/research/

[Hin97] Heather M. Hinton, "Under-Specification, Composition and Emergent Properties," *Proceedings of the 1997 New Security Paradigms Workshop*, Langdale, Cumbria UK, September 23-26, 1997, Association for Computing Machinery, 1998.

[ISW97] *Proceedings of the 1997 Information Survivability Workshop*, San Diego, California, February 12-13, 1997, Software Engineering Institute and IEEE Computer Society, 1997. Available at: http://www.cert.org/research/

[ISW98] *Proceedings of the 1998 Information Survivability Workshop*, Orlando, Florida, October 28-30, 1998, Software Engineering Institute and IEEE Computer Society, 1998. Available at: http://www.cert.org/research/

[Pre97] Presidential Commission on Critical Infrastructure Protection (PCCIP), *Critical Foundations — Protecting America's Infrastructures*, The Report of the Presidential Commission on Critical Infrastructure Protection, October 1997.

[Reu98] Reuters, "Pager Glitch Cause Lost in Space," Wired News, Wired Digital Inc., August 11, 1998. Available at: http://www.wired.com/news/news/technology/story/14355.html