



Fostering Business Resilience

An Interisle Consulting Group Whitepaper

For more information contact:

Christopher Owens, Principal
Interisle Consulting Group, LLC
39 S Russell Street
Boston, MA 02114
+1 617 413 3734
<http://www.interisle.net>
chris@interisle.net

Executive Summary

This paper explores the concept of *business resilience*: the ability of an enterprise to function effectively in the face of disruptive events. It focuses on the role of communications networks and distributed computing infrastructure, at the same recognizing that achieving business resilience is ultimately about management discipline rather than technology.

Many of the examples in this paper are drawn from our work with the exchanges, utilities, and infrastructure providers that serve the securities industry, which suffered significant and widely-publicized infrastructure damage during the terrorist attacks of 2001. Since then, key players have made significant changes that serve to improve industry-wide resilience. We believe that the lessons learned are of general interest and applicable to other industries and enterprises as well. Other observations and principles are drawn more generally from our long careers in Internetworking and the architecture of resilient distributed systems.

One of our main themes is that assessing or improving business resilience requires disciplined, systematic, and cross-functional thinking about an organization, its mission, its operations, and its technical infrastructure. It should be approached comprehensively not in a piecemeal way. A single department or functional area within an enterprise, or a single enterprise within an industry, can do only so much to foster resilience, which, by its very nature, requires communication and coordination.

A second theme is that resilience is not just about terrorist attacks or natural disasters; it is about dealing with change gracefully on a variety of timescales. Desirable occurrences such as mergers and acquisitions, the introductions of new products or technology, or the emergence of new groups of customers, can be nearly as disruptive as disasters, and the organization that pays attention to flexibility and adaptability is well equipped to capitalize on new opportunities. Taking steps to foster business resilience is not just an “insurance premium” against disaster; it is sound preparation for opportunity as well.

Elaborating and expanding upon these two themes, the paper suggests principles that an organization can evaluate and concrete actions that an organization can take as it assesses its own business resilience and seeks to improve it. While these principles and actions do not present any magic bullet, nor are they a checklist which, if faithfully followed, is guaranteed to improve resilience, they do encourage a comprehensive approach to assessing resilience and practical strategies for improving it. They lead to a broad examination of costs and benefits, ultimately identifying attributes of resilience that are desirable not only because they reduce risk, but because they create opportunity.

An Introduction to Business Resilience

In the aftermath of the September 11, 2001 terrorist attacks, the securities industry, heavily concentrated in lower Manhattan, faced an unprecedented array of challenges. Beyond the immediate and severe loss of life, communication of all types (including radio) had been disrupted. Power was out for an extended period of time. Movement of people, goods, and supplies (including air transport on a national basis) was disrupted. Nearly a quarter of the office space in Lower Manhattan was obliterated. Access to facilities (even undamaged, intact ones) was restricted by police. Emergency services were severely overtaxed. Then, a month later, anthrax emerged as a threat that could shut down entire buildings or city blocks, halt the flow of mail, and disrupt transportation systems. Although 9/11 and the anthrax attacks were caused by terrorist action, a bit of reflection leads to the inescapable conclusion that natural disasters like hurricanes, earthquakes, even bad snow storms can be every bit as disruptive. Even a single failure, such as the regional power outage during the summer of 2003, can cause widespread disruption. Disasters, both natural and man-made, have happened before and will certainly happen again.

The devastating events of 2001, and the subsequent recovery from the damage they wrought, brought intense focus on the issue of *business resilience*. In this paper we look at the idea of resilience as applied to systems, networks, and organizations. We go beyond the lessons of September 11th 2001 and explore some of the general principles that make a system, a network, or an organization resilient in the face of unpredictable events.

To those involved in putting the pieces back together, it quickly became clear that some elements of critical infrastructure had failed and others hadn't, in patterns that were not necessarily predictable. Of the elements that failed, some were quickly repaired and others were not. Of the businesses that depended on failed elements, some were able to work around the failures easily, and others had a great deal more difficulty. In some cases, facts on the ground did not match intended designs or architectural principles: for example, multiple, redundant telecommunications circuits, believed to have been routed in a geographically diverse manner, in fact all transited a single central office building.

By paying attention to what worked and what didn't, many organizations learned a great deal in a short period of time. One thing that became clear is that the robustness of any particular component, system, or network may not be nearly as important as *our ability to deal with failures when they occur*: to put the pieces back together when something breaks, to work around problems, to accommodate disruption. Another thing that became

resilient adj.

1. Marked by the ability to recover readily, as from misfortune.

2. Capable of returning to an original shape or position.¹

¹ The American Heritage® Dictionary of the English Language, Fourth Edition Copyright © 2000 by Houghton Mifflin Company.

clear is that businesses are increasingly interdependent upon each other on a minute-by-minute operational basis: It does little good for a business's own data centers to be up and running and its operations staffed if its customers are off-line and cannot transmit orders, or if the business cannot communicate with suppliers, or trading partners. In the case of the exchanges and financial markets, it does little good for the central infrastructure to be fully operational if large numbers of participants cannot function or communicate.

Resilience isn't just about surviving disasters. Positive events – such as introduction of new products and services or entry into new markets, can also be highly stressful to an organization. Resilience is about preparing for, and dealing effectively with, change of all kinds.

Business leaders began to look at their own operations and ask themselves how they would handle a major disruption, what they would do to get back on their feet, and even whether or not they could survive if their own operation and/or their customers, suppliers, and trading partners were similarly struggling and if the markets and shared infrastructure upon which they depend were damaged.

But business resilience isn't just about surviving a disaster, and achieving business resilience isn't all about preparing for gloom and doom scenarios. In fact, many of the principles that lead to resilience are simply the result of sound business and technical architecture, clear thinking about function, and effective organizational design. At the core,

business resilience is about

solid fundamental structure, agility and the ability to deal with sudden change. While it is impossible to account for and prepare for every possible eventuality, it is nevertheless possible to design for resilience and to train for dealing with disruption.

While it is impossible to account for and prepare for every possible eventuality, it is nevertheless possible to design for resilience and to train for dealing with disruption.

Sudden change comes in many forms. In addition to manmade and natural disasters that disable infrastructure and keep people away from their offices, there are other sorts of occurrences, even apparently positive ones like a sudden inrush of new customers or the departure of a large competitor, that place unexpected, disruptive strain on the systems, networks, and people upon which a business depends. We believe that sound architectural thinking about resilience pays off not only in minimizing the damage done by disruptive events, but also in preparing the organization to react nimbly to opportunity.

The Principles

Our experience, gained not only from dealing with the aftermath of terrorist attacks but also from long careers in Internetworking and in designing and building resilient infrastructure, leads to a number of observations and principles, some of which may be obvious and others less so. As stated, some of them conflict with each other. This is deliberate: reasoning about resilience often requires learning how to make trade-offs among multiple, conflicting objectives as circumstances change and as new possibilities

and technologies become available. We believe that if an organization examines these principles as applied to its own business operation, the examination will lead to a comprehensive assessment of resilience and a practical approach to improving it.

1. Resilience is a set of objectives, not a department

“When you are up to your neck in alligators,” the saying goes, “it’s hard to remember that the original objective was to drain the swamp.” To a telecommunications manager, maintaining the integrity of the link between two of your company’s key facilities may be of the utmost importance. But unless your company is a telecommunications provider, the company as a whole does not intrinsically live and die by its ability to move bits between facility A and facility B. The company as a whole lives and dies by its ability to manufacture and distribute cheese, or execute financial transactions, or deliver packages, or whatever the primary revenue-generating business of the company happens to be.

Moving bits between facility A and facility B may or may not be an essential link in the process. Perhaps order-processing is in facility A and the manufacturing operation is in facility B. The company could not survive for long without a fully functioning communications link between the two, *as things are now configured*.

But might things be configured differently? Could facility A take over the manufacturing function? *Doubtful, all that’s there is an office building*. Could facility B take over order processing? *More conceivable, at least*. Could facility B take over order processing, just on a reduced temporary basis, while awaiting the repair of a damaged telecommunications link? *Probably so, with some investment in cross training and some records and facilities duplication*. Would this cross training and records duplication offer additional benefits, in terms of business resilience and overall agility, which might not be offered simply by a redundant communications link? *Almost certainly*.

Now, the question on the table has been transformed from “*How can I make this communications link more resilient?*” to “*What are the costs and benefits of making this communications link more resilient, relative to the costs and benefits of the cross-training and facilities duplication that would be necessary to handle order processing in facility B?*” The right answer might very well still be “*Focus on the comms line only*” or “*Do*

1. Resilience is a set of objectives, not a department.
2. It’s about the people.
3. The weakest link breaks the chain, but component analysis is fraught with peril.
4. Context is everything.
5. Resilience is a team sport, not a solo game.
6. Measure the right thing.
7. Great value lies in options.
- 8 Train for the concrete. And for the abstract.
9. Resilience is an approach, not an ingredient
10. It’s not obvious what’s a “non critical” function.

both!” but it’s a difficult question for a telecommunications manager to assess on his or her own. It’s particularly difficult for a functional area manager to answer, “*How could we get by without **my** function being performed?*” even on a temporary, disaster recovery basis. That’s why resilience is not the province of any single department, and must be viewed systemically. The individuals responsible for business strategy, software applications, operations, and infrastructure must all be committed to an explicit, shared, documented set of resilience-related goals and to a common, agreed plan for achieving them.

2. *It’s about the people*

Veterans of disaster recovery report that one of the biggest challenges is **finding and communicating with the right people**. Consider two heavily interdependent organizations, for example a brokerage house and an exchange, or a provider and a user of some electronic service, both of which have been suddenly and forcibly relocated. As they attempt to reestablish data and voice communications, nobody is reachable at his or her usual office telephone number. Some people will be working at disaster recovery sites or other alternative locations, others from home. Many people may be performing other than their usual job functions. Enabling the formation and functioning of ad-hoc teams is essential.

Give people the ability to find each other and communicate effectively and they will find a way to form teams and get the job done.

Experience dictates that **directory services** and **alternative means of contact** become essential components of a resilience strategy, not only within an organization, but between an organization and its customers, suppliers, and trading partners, even industry-wide. This points to the criticality of voice telecommunications, as well as the importance of integrating and using of widespread consumer-grade services (voice, e-mail, web, SMS) that may be external

to your enterprise.

For a contingency directory service to be effective, people must be able to update their own entries, since it can’t be predicted in advance where someone will be working or what function he or she will be taking on. This favors an electronic database over paper. Furthermore, a directory service must be accessible, meaning it cannot reside solely on a corporate LAN. Paper is easily replicated, does not depend upon network connectivity, but is hard to update and works only if you happen to be carrying it. A successful contingency directory service is likely to incorporate a combination of paper, automated systems, and a central telephone number people can call.

Frightening events such as terrorist attacks or major storms raise an additional factor: the safety of their families is going to be the first priority for most people. Attending to this need, perhaps by thoroughly enabling work-from-home support or through various family emergency support services, may play an important role in an overall resilience strategy. If the workplace or disaster recovery site can be counted on as a place of refuge for families, it may make it possible for people to work longer, harder, and more effectively under otherwise stressful conditions.

Reiterating the point raised earlier about the importance of details, experience has shown that even if a facility has backup electrical power, heat and air conditioning, food, drinking water, a place for employees and their families to sleep and change clothes, without water to flush toilets it quickly becomes unsanitary, and ultimately untenable. A demoralized workforce is less effective.

3. The weakest link breaks the chain, but component analysis is fraught with peril.

It's an article of faith among system designers that the most reliability "bang" for the engineering "buck" comes from strengthening the weakest component in a system. For example, if all your vital communications links are redundant except for one, it seems obvious where to spend your next reliability improvement dollar.

What's less obvious, though, is how to measure the "strength" or "weakness" of any given component. Traditionally, we assigned a dollar cost to a component failure, using a mixture of hard and soft criteria, and then multiply that cost by an estimate of probability that the component might fail.

If you make a profit on operations of one million dollars per day, and a failure of your electric power would knock you out and take a day to fix, there's a hard loss of one million dollars, plus the cost of repair, plus whatever (presumably much larger) soft cost you might assign to the loss of reputation that would result. In most cases, the soft costs dominate. For the purposes of this example, assume that the total is \$10 million. Your analysis of statistical data tells you that, for buildings like yours, power fails on average about once every ten years – a one in 10 chance that power will fail in any given year. \$10 million in cost, and a one in 10 chance of occurrence, suggests a risk exposure of about \$1 million per year, which tells you how much you might consider spending to make the problem go away, either by reducing the likelihood (*e.g.*, backup power supply) or by reducing the cost exposure (*e.g.*, distribute operations across multiple sites). Other options include purchasing insurance, or building up inventories or cash reserves to deal with short-term disruptions. In nearly every real-world situation, hybrid strategies tend to be most effective, but only so long as a hybrid strategy doesn't become "do everything."

Unfortunately, this component-by-component analysis fails to take into account a number of significant factors about the overarching business and technical context, and about the way in which the components interact with each other. Often it is tiny, unanticipated details that compound a failure and render a redundant component useless. Consider the case of the company with a backup power generator, fed by diesel. But the tank was on the roof. During a prolonged power outage the company ran out of diesel – and then had no way to get the diesel pumped to the roof to get the backup generator going again. Have you looked at how your supply chain operates in the event of a disaster? Or have you put into place an alternative, disaster-specific supply chain?

A component-by-component "Vulnerability Index", calculated by multiplying the cost of failure by the likelihood of failure may be useful, but, taken by itself, can lead to some very wrong conclusions about resilience

At the same time, component analysis is an “easy sell,” and whole industries have emerged that tend to promote only one “solution” to disaster recovery or improved resilience. It can be difficult to pursue a different path when everyone else has bought into the same component analysis and taken the same mitigating approach. There are even examples of regulators mandating a blanket solution based on component analysis.

4. Context is everything.

Before adopting any proposed resilience “solutions”, and in particular in determining where to invest, how much to invest, and when to stop, it helps to consider a number of broad contextual factors. Three general ones are outlined here; your specific industry may well have others.

The first contextual factor examines the likelihood of a failure. In the real world, **events are not randomly distributed. For example, the presence of a motivated adversary with the means to launch attacks changes everything.** Statistical engineering assumptions about the likelihood of a component failure do not take into account the goals, plans, and actions of an adversary. Locating your data center at a place where power is available from two independent distribution grids and feeding it from both, significantly reduces your chance of an *accidental* power outage disrupting operations. It does hardly anything, though, to your vulnerability to *deliberate* attack. An operative sent to cut one wire can almost as easily cut two. This is a major lesson from the September 11th experience.

The second contextual factor deals with the soft costs of a failure. **Failures that affect only you, damage you worse than widespread ones.** Being offline while everyone else is online (for example, due to an accidental cut to the power or communications cable that serves you) damages your reputation and your competitive position far worse than being offline while everyone else is also offline (for example, due to a widespread power failure that shuts down an entire geographic region.)

“Systems thinking” – attention to business and technical architecture – probably matters more than the robustness or redundancy of any particular system or component.

The third factor examines the value of different repair or mitigation strategies, given that in the real world, system **failures are not statistically independent.** A single failure may affect multiple technology components, or multiple organizational functions, or multiple trading partners. Statistical analysis might lead to the discovery that the most frequent cause of power outages in a specific region is tornado damage and that whenever a tornado damages power lines it also damages the adjacent telephone lines. This is sometimes termed a “fate-sharing assumption.” If continuing operations depends upon data

communications with the outside world, investing in a backup power supply, while it practically eliminates the risk of a power failure, does very little to reduce the total risk exposure. Following a tornado, your systems will be up and running, but out of contact with the outside world and therefore not in business. The backup power supply nearly eliminates the risk of being without power, but it does very little to the risk of being out of business due to a tornado. This analysis suggests a decided advantage for some

mitigation strategies (for example, operation at multiple sites) over others (for example, redundant electrical power at one site)

A related fate-sharing assumption links you and your suppliers, customers, and partners: if, following a disruption, you recover your operations in one hour while it takes six hours for those with whom you do business to get back on line, the additional five hours buys you very little. On the other hand, if your business partners recover in six hours and it takes you seven, the additional hour costs you a great deal.

Similarly, if it takes your staff two hours to relocate to your offsite backup facility and commence operations there, then there is little to be gained by improving your network switchover time from 10 minutes to 50 milliseconds.

5. Resilience is a team sport, not a solo game.

As mentioned above, there's little point being fully operational if your customers, suppliers, and trading partners are not. Businesses are interdependent. Reflecting this point, it is not necessary for every organization to pursue its own independent resilience strategy from the ground up. In fact, in the event of a widespread disruption, your organization's resilience strategy is only effective if it works well with the resilience strategies of your customers, suppliers, and trading partners.

In many cases, it may be highly effective to address resilience on an industry-wide basis, creating shared organizational and technical infrastructure of a scale and quality that few organizations could manage on their own, amortizing the costs over a broader industry group. Such an approach is effective if it provides clear, cost-effective benefits relative to independent strategies, and if it does not alter the competitive landscape. Some industries already have natural candidates to become the leader and focus of shared, industry-wide resilience activities; in other cases such a player can emerge or may need to be created. Although the work is daunting, the prospect of new opportunities and substantial cost savings is attractive.

Shared, industry-wide organizational and technical infrastructure can be a highly effective approach to resilience, and can potentially create additional value by enabling new products and services

6. Measure the right thing.

Which is better, 99% availability or 90% availability?

The answer is: "It depends."

A management team, assessing system reliability, discovers that one of the company's critical systems has an availability of only 90%, and a mean time between failures of only around a minute. Decidedly unimpressed by these numbers, they begin interviewing vendors, one of whom promises a system that fails only once per year, on average, and achieves 99% availability.

The team is on the verge of committing to a replacement system when they calculate that 99% availability and a mean time between failures of one year jointly imply that the system, when it fails, takes almost four days to get back onto its feet. On the other hand, the existing system fails often, but is only out for six seconds on average when it fails. Mean time to repair (MTTR) is at least as important as mean time between failures (MTBF) – either one tells an incomplete story without the other.

What gets measured gets managed.

Selecting appropriate metrics is harder than it looks, and is a critical element of any resilience strategy, worthy of the full attention of the organization's senior strategists.

99% availability and a mean time between failures of one year jointly imply that the system, when it fails, takes almost four days to get back onto its feet. On the other hand, the existing system fails often, but is only out for six seconds on average when it fails. Mean time to repair (MTTR) is at least as important as mean time between failures (MTBF) – either one tells an incomplete story without the other.

The costs of non-availability, both hard and soft, are often highly non-linear. There are many applications in which a six second outage, even one occurring as often as once per minute, would not even be noticed. A four-day outage, on the other hand, could be devastating. Every technical and operational metric must be linked,

thoughtfully, to overarching business goals.

7. Great value lies in options.

A surprising effect of the 2001 terrorist attacks was the wholesale failure of point-to-point telecommunications lines, both voice and data, serving lower Manhattan. Large numbers of circuits passed through a single telecommunications facility, due in part to the accumulated effects of years of field maintenance operations, in particular the practice of “grooming” by which field technicians reorganize and locally optimize the infrastructure by relocating circuits from one physical trunk to another. That facility, adjacent to a destroyed building, continued to operate until its batteries were exhausted, and then failed. Physical access, either to bring in temporary power or to begin restoring circuits, was impossible.

Enterprises who thought they had acquired geographic route diversity across their many circuits found they had not achieved their geographic circuit diversity goal. Even those who had attempted to obtain additional levels of redundancy by buying from multiple vendors found that, because vendors lease infrastructure from each other, the additional redundancy was illusory.

The securities industry, in part because of a history of localized, department-by-department acquisition of infrastructure of the sort described above under Principle 9 (“resilience is not an ingredient”), owned huge numbers of point-to-point telecommunications lines (both voice and data) connecting various participants with the exchanges and utilities upon which the financial markets depend. Many of these lines failed. Many others, however, did not. What was lacking, however, was a comprehensive, systematic way of pressing an intact, working circuit into service to replace one that had failed. The industry had many functioning lines, but no overarching resilience strategy for using them.

In contrast, organizations that had consciously built their network architecture around general-purpose infrastructure (for example, Internet-style IP-based networking) had an

easier time than those relying on purpose-built point-to-point links. This Internet-style approach makes more extensive and explicit use of shared, common infrastructure; furthermore, the fabric of the network itself (e.g. switches and routers) is designed from the ground up to accommodate change. Shared general-purpose networks were (relatively) easily put back on their feet; large numbers of dedicated lines took far longer to repair.

The principle is that, in an environment of unpredictable failure (or unpredictably changing business requirements), **general-purpose tools and components increase one's options for recovery.**

Anybody who is responsible for the care and feeding of a complex mechanical environment, whether that system is a battleship or a television studio, soon learns that, no matter how wonderfully stocked the spare parts locker may be, there is also great value in having on hand a few boards of wood, screws and nails, some hand tools, and, of course, duct tape! Such a repair kit buys a lot of options in dealing with unexpected circumstances. The system architect can work in concert with the repair technician, standardizing parts and making the system more amenable to general-purpose tools and techniques and to a smaller spare parts kit. The legendary WWII-era Jeep, which could be disassembled and reassembled in the field with little more than a screwdriver and an adjustable wrench, is an example of such resilience-oriented engineering.

Identical, interchangeable components, multiple sites, cross-trained staff, and a general bias towards the use of general-purpose "commodity" technology all increase your options.

Another, related factor is that **interchangeability contributes to resilience.** For a business dependent upon networked computing, are your routers and servers all on one hardware and software platform, thereby increasing your ability to substitute one for another in a pinch, and increasing the chance that the staff on hand will have the right training and spare parts to repair damage or deploy new configurations? Of course, there is a countervailing risk commonly termed the *monoculture vulnerability*: if all your servers are the same hardware/software configuration, your entire infrastructure is vulnerable to a common-mode failure such as a virus or programming bug. The resilience-increasing value of interchangeability must be weighed against the resilience-decreasing risk of common-mode failures. Common mode failures may be of greater concern where some threats come from motivated adversaries.

Typically, using a single vendor or single technology platform raises the monoculture threat; but offers the most reusability of components, most bang for the buck with spares and parts, and the most interchangeability of staff functions. Use of two technology platforms or vendors raises the cost and lowers the utility of spares and staff training, but addresses the monoculture problem. Going to a higher number (three or more) technology platforms or vendors further exacerbates the cost and complexity of stocking spares and training staff, and may arguably not improve the monoculture situation much beyond what is achieved by having two. Detailed analysis is called for.

A number of factors tend to increase one's options: multiple sites, cross-trained staff, repurposeable and interchangeable components. Each of these comes with its own costs and risks; the art continues to be finding a balanced approach that is appropriate to the specific environment, business conditions and risk profile.

8 Train for the concrete. And for the abstract.

There's a reason why sports teams scrimmage, and why military units participate in war games. Concrete scenarios are extremely useful in thinking about resilience. An organization ought to be able to work through a number of scenarios or simulated events, both broad and specific. As with many other complex endeavors, working through concrete scenarios, running extended, realistic simulations, and conducting drills will often reveal details that require attention.

By their very nature, unanticipated critical details are nearly impossible to identify in an abstract plan. In one "war-game" style exercise involving a large number of participants, it was determined that coordination and information-sharing was a major objective, and that a twice-daily conference call would be effective. It's one thing to write "stage conference call" on a plan and check it off as complete when the arrangements have been made. It's another thing altogether to run the exercise, and to discover that, in actual practice, key participants don't have the call-in number because they're working away from their usual office and files (part of the scenario denied access to a couple of buildings), and that the number of people needing or wanting to participate in the calls vastly exceeds the technical capacity of the conference bridge, resulting in people being unable to hear and some participants being closed out. "The devil," to paraphrase Ludwig

Well-formulated plans for dealing with a variety of abstract conditions (loss of transportation, area denial, loss of data communications, failure of a major supplier or customer, and so forth) become a rich library from which planners can piece together a solution to deal with whatever concrete situation presents itself in the real world.

Mies van der Rohe, "is in the details."

As useful as detailed, specific scenarios and complex, high-fidelity simulation exercises are, it's essential to remember that *the scenario in the simulation exercise is probably not the one you will encounter in real life*. The purpose of the exercise is not to become adept at handling **that** particular scenario; it's to uncover details that need to be addressed, and to become adept at thinking about and handling **a range** of actual situations. There's a tradeoff here between breadth and depth of training, because there isn't enough time to train in great depth for large number of divergent scenarios. A highly-detailed, extended exercise involving a single scenario explores the specifics: Will there be enough slots on the conference call bridge? Will it be possible to deliver more diesel fuel to the backup generator once the initial supply is exhausted? On the other hand, exploring a large number of very different scenarios sharpens a team's general skills for reasoning about and dealing with adverse conditions. It's as much about the lessons learned and skills acquired during the *process* of designing and conducting drills and exercises as it is the specific knowledge gained from the exercise scenario itself.

Additionally, it's important not to let the vividness of a specific recent example (*e.g.*, a terrorist strike against lower Manhattan) blind an organization to a more likely but less dramatic threat (*e.g.*, a bad ice storm that keeps half the workforce at home). In addition to planning for concrete scenarios, it may also be useful to plan for generalized, ones. For example, "how would my organization deal with a situation in which nobody could work at his or her usual location?" Well-formulated plans for dealing with a variety of abstract conditions (loss of transportation, area denial, loss of data communications, failure of a major supplier or customer, and so forth) become a rich library from which planners can piece together a solution to deal with whatever concrete situation presents itself in the real world.

9. Resilience is an approach, not an ingredient

In general, resilience is ***designed in*** at a fundamental level and is seldom improved solely by ***adding something*** to an existing environment. Resilience is the result of successfully applying a coherent mindset to strategic, architectural, engineering, operational, and organizational issues.

Note carefully that "taking a broad, systematic view" does not equate to "go back to the drawing board and re-engineer everything." Far from it.

In some cases, re-examining an organization's technical and business architecture will result in a decision to add computing hosts, communications links, facilities, or staff. In many cases it will not. We know of one organization that deployed a new network design largely motivated by resilience concerns, but then, reaped a windfall because the network enabled the elimination of hundreds of leased telecommunications circuits, with a concomitant reduction in direct costs and indirect support costs.

In fact, a common source of vulnerabilities that can reduce resilience is the gradual accretion of many layers of technology, processes, and organizational structures over many years. Like trash that accumulates until it becomes a fire hazard, many businesses have tended to add new systems, components, communications links, and organizational functions without planning for the obsolescence and removal of the old or otherwise rationalizing the overall architecture.

So, while there are specific changes that a business can implement that will lead to greater resilience, it must be recognized that any vulnerability can diminish resilience. Hence, it is important to approach the challenge of improving resilience by considering what to add, and what to take away, while recognizing that the mere addition of patches or workarounds to mitigate vulnerabilities is, at best, a short-term tactic.

This observation is closely tied to the need for constant cross-functional and even cross-enterprise analysis. Consider, for example, two departments or business units, both of which need a similar capability (for example: an Internet connection, or backup electrical power, or a file server) If each department pursues its need independently, it's possible to end up having ***duplication***, without ***added resilience***. Each department for example, would acquire its own Internet connection. On the face, the overall enterprise having two Internet connections seems more resilient than having only one, but without explicit coordination:

1. It might not be possible to press one into service to cover for the other in the event of a failure, due to technical incompatibilities, physical layout, or simple lack of knowledge. Rather than redundancy, the two connections might represent two single points of failure rather than one.
2. The two connections might be obtained from a single vendor using a single geographic circuit path, when, for no additional cost, they could have been obtained from two vendors over diverse geographic routing. An opportunity to increase resilience would have been lost.
3. Each department, pursuing its own contingency strategy, might purchase more capacity than needed or a redundant link, resulting in more cost and complexity that would have been necessary to achieve the same result if the two had been coordinated

More generally, *more* (circuits, components, facilities, etc.) is not necessarily *more resilient*. It's the architecture that creates resilience, not the components.

10. It's not obvious what's a "non critical" function.

A fairly obvious tactic for dealing with unusual, irregular operations in the aftermath of a disaster or other disruptive event is to reallocate (previously cross-trained) staff from "non-critical" to "critical" functions. It's not always clear, however, what constitutes a "non-critical" function.

It might be tempting for a network provider to say, for example, "Provisioning new service is the last thing we'll be worrying about in the aftermath of a disastrous, widespread outage. Therefore we'll shut down our order-taking capability and put all those people to work on repair..." Further analysis, though, reveals that industry-wide recovery may depend upon precisely that capability: businesses forced from their usual locations may be reconstructing, reconfiguring, and operating from new, unexpected places. They won't be able to get back on line until new telecommunications service can be established.

Similarly, a number of functions generally labeled "customer support" become critical during disaster recovery or other non-standard operations. Close coordination with customers, particularly as the customers themselves struggle to resume operations, becomes an essential, even critical function. 9/11 taught us that HR, normally considered non-critical, was one of the most vital departments for firms trying to recover, as it became the repository of contact data, and the clearinghouse for staff reallocation. Marketing, communications or public relations may also become essential functions, explaining the progress made towards recovery and maintaining the confidence of customers, suppliers, and partners.

The unexpected criticality of certain staff functions also raises the question as to whether those functions themselves are resilient: Does each of them have backup technology infrastructure? Alternative operating procedures? A contingency plan of its own?

Conclusion: Making resilience a reality

Reviewing these principles and observations, we claim that resilience requires systematic, architectural, organization-wide thinking: that it cannot be addressed by a single department or business function and that it is not in general improved simply by adding more of anything. Sophisticated analysis is called for, not only of technical and organizational components themselves but in the context of likely threats and risks in the real world in which the organization operates. Metrics, while essential, can be deceiving and must be developed with care. In setting goals and establishing priorities, a useful frame of reference is to place a high value on changes that increase the range of your options. Training, drills, and extended, high-fidelity simulations serve not only to develop skills, but also to identify defects in your plans and details that need to be addressed. We've learned from real-world experience that, at the end of the day, it's the people that matter, in particularly their ability to find and communicate with one another. In a related vein, it's very difficult to predict which organizational functions will be critical or non-critical in the case of any particular contingency. Finally, we've suggested that firms not go it alone: that organizations capitalize on any available regional or industry-wide opportunities to cooperate and obtain leverage from the actions of other parties.

Examining these principles and observations, it should be clear that “resilience” shares much with “performance”, “security”, and “quality”:

- It isn't something one adds; it's a result of applying a comprehensive engineering perspective to the entire system, and a comprehensive business perspective to the entire context in which the system operates.
- It's almost a side effect of good overall design.
- It is often very difficult to assess quantitatively; facile and/or poorly-chosen metrics lead to misapplied resources and illusory gains.
- Oftentimes, small, low-cost and low-tech steps, for example backing up a sophisticated electronic directory by placing a photocopied list of home or mobile telephone numbers in every key executive's wallet, can make a big difference in a crisis.
- An approach to resilience must be forward looking and dynamic—not backward looking and reactive.

With this in mind, an organization looking to assess and improve resilience should ask itself the following questions:

1. *Is there substantial, board and senior executive level support for the project?*
2. *Have we got the right team assembled?* For planning, does the group represent all business functions? Is there enough seniority? Is there enough capacity for “systematic thinking”? For execution, is there redundancy and flexibility built into the team? When a real-world disruption occurs, not everyone will be available (e.g., on vacation, away from the area on business, unable to discharge

their duties due to injury or death – or distracted by other issues, for example, worrying about family, friends or loved ones).

3. *Have we posed the right questions?* Do we have a clear hierarchy of overarching business goals? Are our resilience objectives clearly defined? Are our metrics sensible?
4. *Have we considered context?* In particular, how does our resilience thinking deal with the boundaries and points of interface between our organization and our suppliers, customers, and trading partners? Have we looked outside our own organization for help, for partnership? Are we able to use or create elements of an industry-wide solution?
5. *Have we got an appropriate set of knowledge tools?* Being able to gather and reason about statistical data is a start. A clear and well documented picture of what currently exists, (systems, infrastructure, assets, processes, organization) is absolutely essential. Fine-grained, high-fidelity modeling and simulation may or may not be appropriate depending upon the nature of planning being undertaken. Organizational and industry expertise, in particular the ability to tell and learn from “war stories” can be of value.
6. *Have we thought things through in enough **breadth** and **depth**?* Is there a realistic mix of concrete and abstract threat scenarios? Are our scenarios broad enough to obtain coverage of the space of possibilities, and concrete enough to allow us to work out the details?
7. *Have we **documented** and **communicated** our plans?*
8. *Have we practiced? Regularly?*
9. *Have we learned?* From our own exercises and drills? From past disasters and other disruptions? From our peers?
10. *Is the process ongoing and dynamic?* Or have we made a plan once and then relegated it to a dusty shelf?

Putting these principles into effect requires a dedicated, high-performance team, careful attention to objectives, and a rigorous, ongoing, and dynamic process.