TREND 6

Architecting resilience: "Built to survive failure" becomes the mantra of the nonstop business

In the digital era, businesses must support wide-ranging demands for nonstop processes, services, and systems. This has particular resonance in the office of the CIO, where the need for "always-on" IT infrastructure, security, and resilient practices can mean the difference between business as usual and erosion of brand value. The upshot: IT must adopt a new mindset to ensure that systems are dynamic, accessible, and continuous—not just designed to spec but designed for resilience under failure and attack.

Why now?

Digital transformation of enterprises: Transforming to a digital business implicitly increases a company's exposure to risk through IT failures. More business processes are interconnected and automated, all of which become potential points of failure. The average cost of data center downtime by minute has risen by 41 percent since 2010.[1]

Increased cyber threats: It's not just about gaining access to systems; cyber criminals are also trying to bring them down. Denial of service attacks are increasing in frequency and size. The number of attacks has increased by 58 percent in the last year.[2]

Increased IT complexity: More systems are being integrated, and continuous improvement is becoming the IT norm. But constant change to increasingly complex systems is introducing more risk than ever before.

The expectation of "always on": In a digital world, whether your system is under attack, hit by a storm, or just being updated, the expectation is that it always works.

Netflix loves to fail.

Not by delivering movies late, by overbilling customers or in any of the other ways that the video streaming company could fall short. Instead, its engineers try to find fault with their own IT systems—deploying automated testing tools that they refer to as a Simian Army to deliberately wreak havoc in unpredictable but monitored ways.[3]

Why? Because Netflix's engineers know that what doesn't kill their company makes it stronger. Netflix is not alone; these practices were pioneered at Amazon a decade ago and have seen adoption at the likes of Flickr, Yahoo, Facebook, Google, and Etsy.

Those companies' technology chiefs understand something that IT leaders everywhere must grasp: failure is a normal operating condition. It must be anticipated, accommodated, and designed into IT systems. Practitioners of these "game day" strategies—when days are set aside months in advance to perform internal failure testing, with dozens of staff on hand to respond to incidents—regularly find latent defects in their systems, log hundreds of bugs, and continue to test against the repaired defects in future game days.

This continuous improvement strategy involves more than just ensuring that systems have high availability, a condition that still allows for downtime, however minimal. Today, the idea is no longer about designing for "five nines" (99.999 percent) uptime; it's about supporting the nonstop business—literally 24 hours a day, 365 days a year. There can be no exceptions: if systems are to be as nonstop as businesses need them to be, they can no longer be designed just to specification or engineered to handle only particular incidents. They must be designed to work under failure and under attack.

The rationale is simple. As organizations migrate toward digital, every aspect of their business is becoming increasingly interconnected and automated. In natively digital businesses, the digital channel may be the only channel. In this context, resilience—the ability of IT systems to maintain wholly acceptable levels of operational performance during planned and unplanned disturbances—is of growing importance. True resilience is what will help organizations mitigate risks to revenue and brand reputation caused by service disruptions. It's time to architect resilience into all dimensions of the nonstop enterprise, including applications, business processes, infrastructure, and security.

More vulnerable in more ways than ever

As businesses go digital, they are far more susceptible to disruption—vulnerable because IT systems are constantly evolving to do things they were never designed for, because update cycles keep shrinking, and because the intensity and frequency of sophisticated cyber attacks are increasing. Add the impact of natural disasters—seemingly more frequent and more severe than before—and it's easy to sympathize with the challenges being faced by brand managers and risk officers of nonstop businesses. In an always-on world, business leaders have to expect and accommodate the risks posed by internal and external disruptions.

The economic risks associated with business discontinuities can grow incredibly high, incredibly fast. This is especially true for digital companies that rely on Internet-based business models. Take Google's five-minute outage in mid-August 2013 as an example; it's reported to have cost the company $545,000 in revenue.[4]

Not all outages are so costly; a 2013 Ponemon Institute study found that the average cost of data center downtime across industries is approximately $7,000 per minute in losses.[5] The cost of disruption varies by industry and the scale of the compromised infrastructure.

Arguably, the vulnerability that CIOs feel most acutely is from cyber attacks. As transformations to digital multiply, so do the associated risks from cyber criminals. These attacks are increasingly substantial, sinister, and sustained. In 2013, for instance, charges were brought against a group of five hackers based in Russia and Ukraine for stealing more than 160 million credit card numbers over the past eight years. In that same period, they also compromised more than 300,000 accounts from a single banking group.[6]

One of the myriad vulnerabilities highlighted by this group's crimes is the increasing sophistication of brute-force password attacks. Contemporary password cyphers draw from a dictionary with billions of passphrases, route them through rule engines, and use massively parallel graphics-processing units to test trillions of passwords against a single login credential.[7] In short, passwords—even those stored under cryptographic hashes—are vulnerable. Organizations that understand this insist on multifactor authentication policies.

These days, cyber criminals are highly sophisticated and strategic in their approaches—and rarely brought to justice. Three of the five hackers in the aforementioned example are still at large. Individuals are not the only offenders: organized crime, nation states, and sometimes unscrupulous competitors are also guilty of cyber crimes.

Cyber threats are not just about gaining access to systems. In the case of distributed denial of service (DDoS) attacks, it's also about shutting down or disabling services—or at least causing enough secondary discomfort to damage a company's brand. Security company Prolexic reports that in the third quarter of 2013, its clients experienced a 58 percent increase in the total number of DDoS attacks compared with the year-earlier quarter.[8]

More advanced threats are not aimed at entire systems; they target specific products and services that may be beyond the protection of a conventional security perimeter and may include physical assets. The "black

hats" now have ready access to many helpful tools: for example, the Shodan search engine—labeled the "Google for hackers"—quite easily finds infrastructure components that can be probed quickly for insecure authentication and authorization.[9] Today, a botnet that can do millions of dollars of damage within minutes can be rented for $7 per hour.[10]

A surprisingly large proportion of companies concede that they are unprepared for the scope, severity, and sophistication of today's attacks. Nearly 45 percent of CIOs surveyed in Accenture's 2013 High Performance IT Research admit that they have been underinvesting in cyber security.[11] Many feel overwhelmed about where to begin; their chances of catching up seem daunting and expensive.

"Arguably, the vulnerability that CIOs feel most acutely is from cyber attacks. These attacks are increasingly substantial, sinister, and sustained."

Engineering to be a nonstop business, even under attack

The more professional and prolific cyber attacks become, the greater the role that cyber security plays in business continuity. CIOs must use a business-driven strategy to managing risk across the enterprise by understanding which assets are critical and then prioritizing resilience and active defense measures accordingly. These investments should be proportional to the downside risk in the event of a disruption.

The time to start architecting for resilience is right now—not when customers expect it or when losses in trade secrets, revenue or brand value have reached painful levels. After the necessary discussions about risk with the organization's most senior executives, IT leaders must begin to map out the threat models specific to their businesses. With this information in hand, they can use business process economics to identify the services most critical to the organization's strategic direction and thus those most in need of resilience. This might mean giving different tiers of service to different users.

After that, it's necessary to look for investments that provide security "bang for the buck," leveraging existing investments and going beyond compliance. Once these steps are complete, organizations can start to look at advanced detection and external threat intelligence capabilities to better orient their investments toward the areas most in need. This process will provide the CIO with an immense amount of data necessary to move from a compliance-focused stance to one that is more threat-centric and tied to strategic risk. Resilience is far ahead of compliance and best practices.

Security experts must also architect for a diversity of economic conditions, business risk factors, and a multitude of entry points—including their own security fabric. Can their own control systems trust the information they're receiving? Is their white listing (identifying known entities that are trusted) really working? Has their end-point protection been deactivated by trojan malware?

Ensuring trust among all components of a system—through attestation—is the next security frontier. One of the best examples of exploits that could have been mitigated through proper attestation was the targeted

remote attack of Iran's centrifuge control systems at a uranium enrichment facility. The trojan malware deployed against Iran's nuclear refining capacity caused centrifuges to spin beyond their designed operating parameters while reporting normal operating conditions back to the control systems.[12]

In response to this new class of attack, companies as diverse as HP and Siege Technologies are innovating attestation solutions at the hypervisor level, while others such as Mocana are concentrating on the machine and embedded device level.[13] Putting it another way, the former are focusing on ways to verify and trust the operating conditions of systems while the latter are securing end points so that they're less likely to fall prey to an attack.

Once an organization has the technical solutions in place (DDoS appliances, highly skilled security personnel, applications and infrastructure designed to detect early warning signs, security analytics feeding into proactive quarantining, and automated traffic swings and sink-holing), the most effective response is coordination among peers. This practice has been adopted by the financial services community as a response to a repeating pattern of prolonged, serial attacks against its members. The victims later in the attack chain learn from earlier victims, share architecture recommendations and IP reputation scoring, and provide for continuity in relationships with law enforcement. This has proven to be an effective countermeasure and is being mimicked in other industries and by regulatory bodies as a result of the successes in the financial services sector.

Technologies to improve resilience

Cyber attacks aside, businesses that are striving to become digital are racing to keep up with always-on expectations. It is no longer acceptable to simply post notices about planned downtime. There is less and less tolerance for service interruptions in any form. Whether systems are brand-new or state-of-the-art digital systems from the likes of Google and Facebook, or conventional legacy systems, there are many tools available to help systems administrators provide always-on delivery of digital services.

To a large extent, CIOs already understand that annual release cycles are a thing of the past. Facebook and Yammer are among the leading organizations showing the way forward: they answer the call to be always on by deploying updates in staged releases and using quantifiable metrics and statistical modeling to measure their effectiveness. Only if the features reach predetermined performance metrics are they rolled out to a broader spectrum of users.[14]

Technology companies are not the only ones moving in this direction; high performers in IT are beginning to embrace agile development practices and are adopting related methodologies for operations—that's six times the rate at which other IT departments do it, according to Accenture's latest High Performance IT study.[15] The challenge of transitioning to agile at scale is being met by a suite of operational tactics and technologies, including DevOps, performance monitoring and failure tracing, workload management, and software-defined networking (SDN). Combined, these practices and technologies pave the road to resilience by making it possible to build always-on software and hardware systems.

DevOps is the business-driven integration of software development and IT operations. DevOps tools such as Chef and Puppet allow for highly automated deployment of entire systems from version control. This enables the rapid deployment of new or extended systems throughout the compute fabric of the enterprise without disrupting the nonstop business.

The agile practice of automated unit testing has transitioned to operations as well, where newly committed code automatically goes through thousands of test cases before being deployed; once deployed, best practice calls for it being deployed on a "canary" server first. If there are any issues, the canary discovers them and stops the cascading of flawed code or configuration to the rest of the production environment. Amazon, Facebook, and Google all use Chef to manage the continuous integration of new hardware and software on their cloud infrastructures—while staying always on.[16]

Performance monitoring and failure tracing tools such as Nagios and New Relic provide data center managers with real-time insights so that they can inspect and troubleshoot their systems, from source code to hardware components.

And workload management tools help to make applications more portable across heterogeneous infrastructure—a factor that is increasingly important with cloud-first infrastructure strategies. With tools such as Akka and Docker, developers can now go beyond agile and leverage their cloud infrastructure investments to build more distributed and concurrent applications and services, adding resilience to the organization while decreasing deployment timelines. Gilt, the flash sales site, uses Akka to build a concurrent, distributed, and fault-tolerant event-driven application that handles the daily burst in traffic when flash sales go live.[17]

Traditional content delivery networks (from vendors such as Akamai, CDNetworks, CloudFlare, Cisco, and F5) are providing businesses with integrated workload management technologies that allow them to stay agile all the way to their consumer-facing activities. In many cases, these CDNs also give businesses access to innovation they may not have otherwise. For instance, CloudFlare's proprietary technology was used to reduce the severity of the DDoS attacks on Eurovision's annual Song Contest that reaches 170 million viewers. By moving to CloudFlare after the site experienced crippling DDoS attacks during the semifinal round, service disruptions were eliminated—something that Eurovision could not have done on its own.[18]

For enterprises using private cloud solutions such as OpenStack, CloudStack, and Eucalyptus, SDN enables seamless bursting to public cloud infrastructure when business demands on compute capacity overwhelm internal capabilities. SDN is also invaluable for helping manage the transition to agility at scale. When data centers (or clouds) fail, SDN-enabled organizations can instantly transfer operations to other online assets, often in automated ways and without meaningful service interruptions. SDN showed its ability to contribute to resilience during Super Storm Sandy in late 2012. CurrenEx used a Vello SDN solution to dynamically reconfigure routes, service providers, and hybrid cloud infrastructure. As a result, the company was the only currency exchange in New York City that was able to maintain connectivity throughout the storm and the ensuing cleanup.[19]

These types of services make IT systems better able to withstand failure, notifying administrators of dysfunction, increasing portability, and providing self-healing capabilities—features that circumvent the deficiencies of the highly available, state-of-the-art systems of just a few years ago. Those earlier systems were about hardware; now they're about instances and processes. Rather than trying to design resilience into every component, it is now best to take a systemic approach where the service delivery architecture should be able to survive the loss of any component—including that of entire data centers. And when components or data centers do fail in a resilient architecture, it's no longer a disaster recovery event; it is a high-availability event.

A mindset for resilience in the digital business

Resilience is the new high ground for CIOs who take their strategic business roles seriously.

That does not simply mean putting in place the right cyber security structures and deploying best-of-breed highly available systems. It calls for a wholesale shift in mind-set to the idea of 100 percent uptime. It is a mindset rooted firmly in the context of business risk and a deep understanding of the constant threats of business disruptions—from hurricanes, hackers, or internal upgrades—and the risks that those threats pose to maintaining operational continuity and brand value.

Above all, the resilience mindset is categorically not about compliance. Compliance means complacency; in an always-on world, it is not enough to simply check the Sarbanes-Oxley boxes to confirm that this or that risk management process is being followed. To be clear, leaders don't follow compliance frameworks; they set them.

It's important to know that many of the tools and methods to engineer for resilience—to design for always-on operation—are available and improving all the time. It is not necessary to wait for the maturation or proliferation of a particular technology. As noted, agile development methodologies are already in use, and

Your 100-day plan

In 100 days, consider where you can make the most impact in building a more resilient company.

- Shift conversations with senior executives about security to conversations about mitigating business risks. Talk about the benefits of designing for failure.
- Map and prioritize security, operational, and failure scenario threat models to existing and planned business operations.
- Develop a strategy to handle elastic business demand for IT services.
- Reaffirm a force-ranking alignment of IT systems and their dependent components with business-driven KPIs for success and downside revenue risk. Evaluate the top five for resilience.
- Test your resilience by planning a "game day" exercise for IT operations.
- Consider hiring an outside security firm to attack your infrastructure, monitor the events internally, and reconcile with logs from the security firm to see where your defenses are deficient.
- Perform a data security review. Determine from a business risk perspective where data can reside; consider using the public cloud as a disaster recovery solution.
- If not already doing it, plan a pilot for software-defined networks and a software-defined data center. Start small and scale over time.
- Create a governance model for auditing and testing the entire ecosystem of IT system and process dependencies—both internally and externally. Be sure it includes policies for managing capacity utilization and using hybrid infrastructure.

This time next year

In 365 days, be ready to embark on projects that will build resilience and reduce the operational risks of your digital business.

- During the budgeting process, look for security- and infrastructure-related investments that maximize business process resilience per dollar spent.
- Publish a plan to transition IT operations to a DevOps-based agile organization.
- Mitigate business downtime risks by aiming to shift compute loads to public cloud infrastructure—either during peak times or while under attack.
- Consider piloting automated root-cause analysis tools in the data center.
- Use results from game day exercises to create a prioritized list for operational upgrades.
- Test your system against agile software outputs. Verify that deploying faulty code leads to safe environment fallbacks.
- Create a security road map to build advanced detection and external-threat intelligence capabilities.

they can be used to even greater advantage in building resilient operations and infrastructure. Even some of the hackers' most useful tools, such as Shodan, can be used by the security community as tools to actively defend infrastructure.

The CIOs who truly get the concept of resilience have begun transitioning their organizations to an always-on state. Knowing that it is neither simple nor cheap to provide real resilience, they are taking a pragmatic approach, phasing in resilience over time as business risk and process economics dictate. And some are already thinking ahead to the time when their entire business is digital, cloud-based and always on.

SIDEBAR

A framework for a resilient future

How can IT leaders start to design for failure? In Michael Mehaffy's and Nikos A. Salingaros' studies of resilience in the natural world, they uncovered four key principles that can be adapted for IT.[20] Any truly resilient IT system should demonstrate the following:

Interconnectedness. The evolution of networks, from point-to-point, to hub-and-spoke, and now to mesh, embodies the benefits that interconnectedness brings. When there are more connections at the edge and throughout a network, aggregate decision-making improves, happens more quickly, and has a greater tolerance for the failure of any one node. Many of these same features appear as part of the sharing economy and expanded workforce as well, which further underscores the disruptive power of interconnectedness.

Diversity and redundancy. There should be no reliance on singular data sources; embracing redundancy, IT systems should demonstrate diversity and be designed for failure. The Hadoop Distributed File System is a prime example of these concepts in action; it has data redundancy at the document, file, and system levels. This redundancy allows analytic jobs to be broken into smaller parts, distributed across the cluster, and run in parallel to achieve results in a highly scalable way. Similarly, high availability is a primary benefit.

Modular scalability. Modular systems can be replaced easily and they enable rapid scalability. They find uses across solution architectures and they work well in large and small deployments. Furthermore, when modular systems are also decentralized, each cluster of nodes becomes less and less significant to the functioning of the whole and more independent of centralized control systems.

Adaptation. Sensors that are able to make localized decisions based on quantified measurements, domain experience, and collaboration with peer nodes can have a significant impact on the physical world around them. These decisions on the edge are informed by shared knowledge and, over time, can gain decision making characteristics akin to biological intelligence.