



Demystifying IT Outages

White Paper



© 2017 Continuity Software Inc. All rights reserved.

Continuity Software, the Continuity Software logo, AvailabilityGuard™ and their respective logos are trademarks of Continuity Software Inc. Other company and brand products are trademarks or registered marks of their respective holders.

Table of Contents

TABLE OF CONTENTS	2
HURRICANE SANDY AND THE IT DILEMMA	3
OUTAGES ARE MORE COMMON THAN PEOPLE THINK	4
MAIN REASONS FOR OUTAGES	5
↳ Drowning in the big "risk waves".....	6
HOW CAN YOU REDUCE THE RISK LEVEL AT ALL TIMES?	7
↳ Eavailabilityguard™ - ensuring availability through automation	8
↳ The benefits of taking a proactive approach to ensuring resiliency.....	10
ABOUT CONTINUITY SOFTWARE	11

Hurricane sandy and the IT dilemma

The evening news on October 27 2012 painted a dramatic picture of a massive storm making its way from the Caribbean. The storm was predicted to hit the east coast shores in three days. Weather forecast analysts who were trying to predict the hurricane's path and potential devastation, presented their "best" and "worst" scenarios by drawing lines on a map, indicating the possible depth of destruction. Although the modelling software used by these forecasters in 2012 was far less accurate than the one used today, it illustrated how the storm would turn west towards land and strike the New York/New Jersey region on October 29, rather than turn east and head out to the open Atlantic as most hurricanes in this position do.



credit: Charles Sykes/AP

Having all this information in front of them, CIOs and IT stakeholders of major corporations faced a major dilemma. What would be riskier for their business--to proactively failover to their disaster recovery site - far away from the path of destruction today? Or to take a chance that the storm would not be as bad as predicted? Although we don't have information about all of the IT systems and data centers, many major outages were reported, suggesting that the IT staff of these organizations decided it would be more risky to failover than to weather the storm. Why would they think that?

Are confidence levels regarding their ability to actually failover so low that companies would rather take the chance on a storm that models predict may have a devastating effect?

The answer is yes. Companies chose to take the chance because they were pretty sure that failing over would not go smoothly -- and even if it did, failing back would also be a nightmare. In fact a [survey from 2014](#) confirmed this uncertainty as it found that 3 out of 4 companies were not properly prepared.

Outages are more common than people think

We hear about outages all the time, it's all over the news, major airlines, leading banks, global payments providers, even stock markets experience IT outages.

We hear about outages all the time, it's all over the news, major airlines, leading banks, global payments providers, even stock markets experience IT outages. Just Google "IT outages" and look at the "news" tab and it's clear that that outages have become a common daily reality. The effects are enormous, with financial damages amounting to billions of dollars, long-term erosion of a company's reputation, along with a lack of stability and low consumer confidence in a time when more and more services are striving to move to the digital age.

The most disturbing part is that the outages that are reported in the news, occur in some of the largest businesses in the world, businesses that are well funded and that probably invest millions in building a redundant IT system that should withstand any catastrophic scenario.

So what is causing all these outages? And how is it that after all this investment, the CIO's and IT managers in these enterprises are not confident that they can avert, or at least smoothly recover from service disruptions?

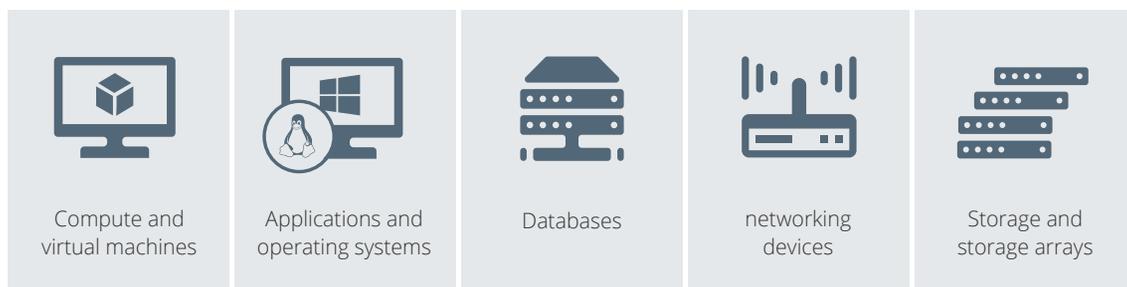


Main reasons for outages

In many of the reported IT outage cases, the companies announced that the outages were caused by a “power switch”, a “network router failure”, an “overload of error messages”, the list goes on. But everyone understands that these faults are just symptoms and not the real root cause of these outages, which happened within environments that were built to have high availability and redundancy. What causes the outages is “statistics of scale”. Let’s look at the following formula to understand the dynamics.

$$\left\{ \begin{array}{l} \textit{Ongoing} \\ \textit{Change} \end{array} \right\} \times \left\{ \begin{array}{l} \textit{Multiple teams /} \\ \textit{vendors /} \\ \textit{solution providers} \end{array} \right\} \times \left\{ \begin{array}{l} \textit{Thousands of} \\ \textit{[ever evolving]} \\ \textit{best practices} \end{array} \right\} = ???$$

IT environments are faced with ongoing cross-domain changes that are introduced on a daily basis. Deployments with thousands of servers usually have to cope with constant and frequent change, such as adding storage or compute capacity and performing updates, upgrades and patches across the entire infrastructure.



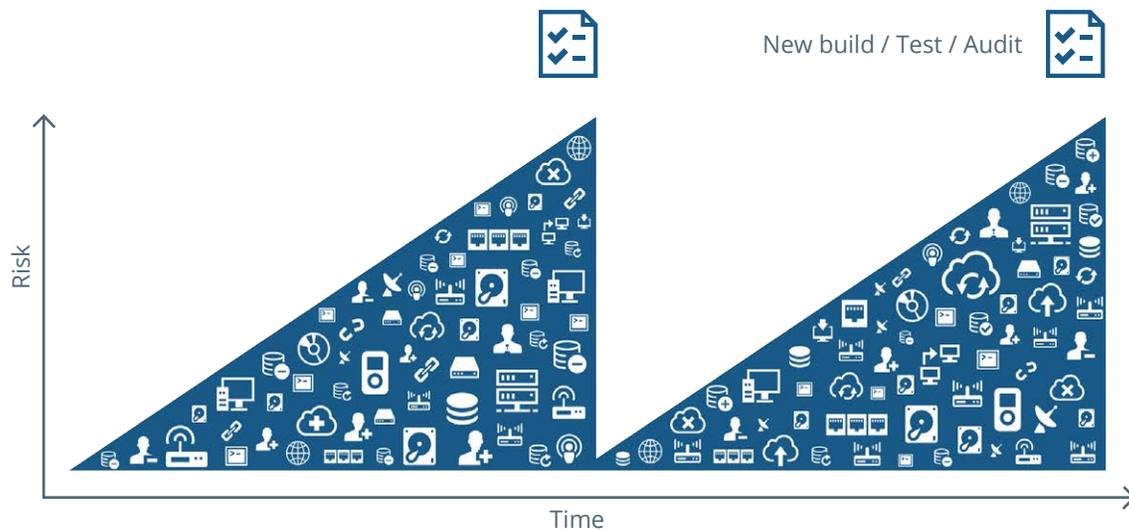
The fact that multiple teams – including external service providers – work on the same environment complicates the process even more as additional problems can take place along the way, due to the challenges in cross-team collaboration and communication. On top of all this, enterprise environments are extremely heterogeneous; containing multiple vendors and combinations of legacy and new systems. When thousands of guidelines are being published and periodically updated for each individual component, it is just impossible to manually keep up with best practices.

With so many moving parts, complexity, and interdependency between the various components, it is no surprise that hidden mistakes happen on a regular basis and the exposure to risk keeps growing.

Drowning in the big “risk waves”

Realizing that it could take a single misconfiguration to paralyze an entire global operation, enterprises must make an effort to detect and contain risk. Running comprehensive tests regularly to review and improve resiliency is a common practice, which is usually held every quarter or on a bi-yearly basis. In the weeks before each test, IT teams work hard to identify and solve all the existing problems. This is important not only in order to improve the likelihood that the test will succeed – but also to mitigate the very realistic risk that the test itself could have an adverse negative impact on normal operations.

So why isn't this enough to raise confidence levels? The following diagram illustrates the problem:



The Y axis represents the risk level and the X axis the time

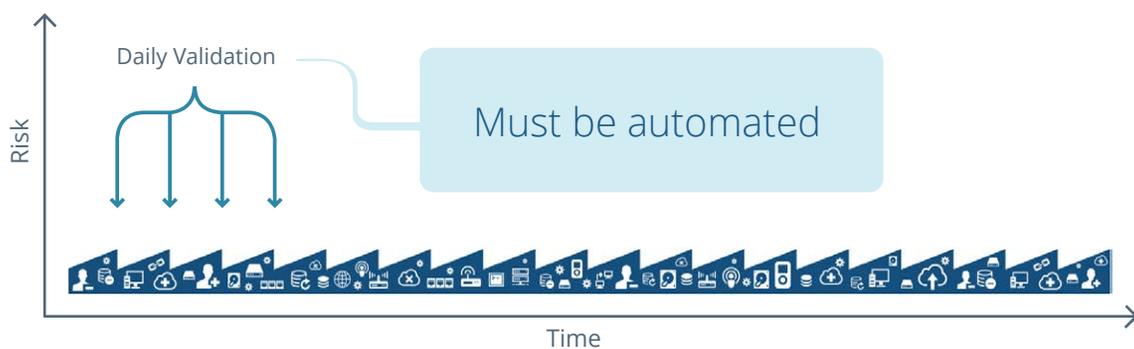
While tests – if correctly planned, exercised, and debriefed – can significantly reduce risk, their effect is extremely short lived.

As soon as a test is over, with the rate of changes described above, the risk and uncertainty swells again. In the day following the successful test, CIOs would have no problem to failover the system when facing an eminent storm. But ask them if they would do it two weeks after the test, and they will probably feel some degree of discomfort. Ask them a couple of months after the test, and you will likely hear *“storm? Ohh it will probably quiet down before it reaches the shore..”*

How can you reduce the risk level at all times?

Due to the massive amount of changes between each test cycle, periodical tests and audits do not provide CIOs/IT Managers with a clear understanding of their risk, nor a certainty that their IT environment is truly resilient and recoverable. If this is not a good-enough solution - then how can companies sustain low-risk levels at all times?

To do so, they first and foremost need to adopt a completely different approach. They should look at achieving availability and resiliency objectives as an ongoing process that needs to be continuously planned, executed, and measured. The frequency of validation must match the rate of change, which means that auditing must be performed more frequently.



So instead of big risk waves that build up between each cycle, they should have very small risk waves that will ensure their confidence in the ability to failover at any time.

To be truly prepared, IT organizations would ideally look to continuously:

- Collect and analyze best practices from all vendors, across all layers to learn about added/dropped functionalities and which changes are required within their environment to accommodate it.
- Perform vendor reviews and bring in cross domain experts to assess whether the dependencies between the various components are not compromised.
- Conduct quality checks following every change in the environment.
- Conduct table top simulations.

This sounds like a completely unrealistic process, especially in a time when IT departments are asked to constantly cut their budgets, right? Fortunately all these processes and methodologies can be implemented using automation with Continuity Software's AvailabilityGuard.

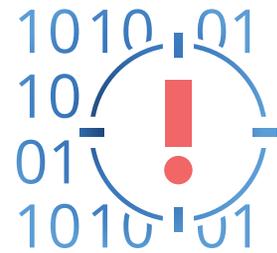
AvailabilityGuard™ - Ensuring availability through automation

Used by leading enterprises worldwide, AvailabilityGuard/Enterprise is a unique IT Resiliency Assurance solution that empowers IT organizations to proactively detect misconfigurations and eliminate outages across all IT infrastructure layers.

The solution is based on 4 key components

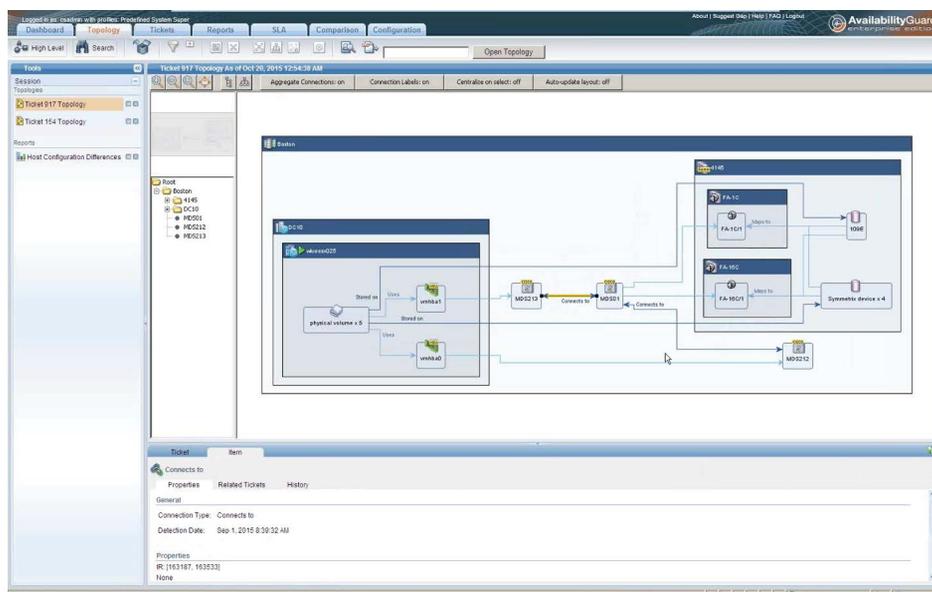
1. Cross-domain, non-intrusive data collection

AvailabilityGuard performs agent-less, non-intrusive, scan of your IT infrastructure configuration. As a vendor-agnostic solution, AvailabilityGuard can collect data from all major vendors and tools, across all the IT layers. Although the solution can collect all the necessary data directly, it can also integrate with leading enterprise CMDBs and vendor management consoles when available.



2. Automated resiliency blueprint discovery

AvailabilityGuard automatically discovers the connections and interdependencies between the various components across the layers that play a role in IT resiliency, while correlating all these resources to the business applications. This creates an elaborate mapping/modelling of various configurations, dependencies, and resiliency schemes and how they are related to the various business applications. This mapping is continuous and dynamic, providing IT managers with continuous and up-to-date view of the resiliency deployment.



AvailabilityGuard™ - Ensuring availability through automation

3. Risk detection

Continuity Software patented technology constantly analyzes the discovered and updated resilience blueprint in order to identify risks. The technology relies on two key features:

Risk Detection Engine

The engine analyzes the deployment's configurations dependencies and resiliency schemes to establish if it was properly configured/deployed based on industry standards. The analysis relies on a vast knowledgebase, that has been accumulated and developed for over 10 years. Detected risks are captured in detailed tickets that includes the root cause, full description of the problem, the impact of the issue, history of the issue, and step by step instructions on how to solve it.

Knowledgebase

The knowledgebase includes industry best practices from both vendors and end-users, and is based on vendor literature, user forums, online publications and constant inputs from hundreds of large enterprises.

In addition to the risk signature itself, the vast knowledgebase - which contains well over 7,000 risk signatures, and constantly growing - also includes detailed information about the business impact, and the recommended remediation

4. Proactive resilience management

Based on sophisticated prioritization schemes AvailabilityGuard can filter, forward and assign risk tickets to the relevant stakeholders, empowering IT teams to resolve issues before they impact the business.

In addition, tickets can be also generated in your existing IT management system and ticketing systems such as ServiceNow, IBM Tivoli, BMC Remedy and more.

The AvailabilityGuard Dashboard



The Benefits of Taking a Proactive Approach to Ensuring Resiliency

AvailabilityGuard allows companies to move from a “reactive mode” to proactively detecting resiliency risks. This allow them to mitigate risk and correct misconfigurations before they actually impact business services and reputation which often leads to costly corrective actions, or even customer compensation and fines.

AvailabilityGuard helps IT organizations achieve the following:

1. Improve resiliency

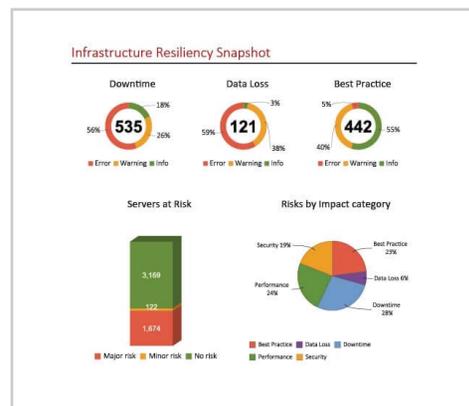
Increasing availability, recoverability and stability of the organization’s environments through regular validation of their infrastructure configuration and resiliency metrics.

2. Optimize IT Operations

Increase agility by allowing shorter time from design to implementation while reducing the risk in software updates and rolling out new systems.

Significantly ,reduce operational cost:

- » Reduced firefighting - due to the drastic drop in downtime or data loss incidents, as well as the reduction in root cause analysis and remediation when issues occur.
- » Dramatically reduced testing costs - as a result of the minimal manual labor and intrusive processes that are required for successfully passing tests and audits.



Using AvailabilityGuard, for the first time, companies have the ability to properly plan their efforts, implement them, and constantly measure, using a process of continuous improvement and operational excellence. Continuity Software’s solution provides automation of regulatory and audit reports, so tests and audits are no longer stressful periods filled with sleepless nights and nerve wracking frustrations. They are predictable and smooth, like wind surfing across a flat sea.



For more information about AvailabilityGuard visit: www.continuitysoftware.com

About Continuity Software

Continuity Software is a leader in IT Resilience & Service Availability Assurance. Continuity Software's award-winning solution enables IT teams to proactively prevent infrastructure outages and data loss incidents. Taking a proactive approach to early detection of IT resiliency risks, allows our customers to mitigate risk by remediation of configuration errors and deviation from policies before they turn into actual costly service incidents.



Founded in 2005 by a team of IT infrastructure and data protection experts, the company is focused on a single mission – helping the world's leading organizations prevent unplanned IT outages across the entire IT infrastructure – including High Availability, Cloud, and Disaster Recovery (DR) environments.

Today, AvailabilityGuard is trusted by many of the largest and most successful enterprises around the globe, ensuring that service availability, data protection, and business continuity goals are met or exceeded 24/7/365.

Additional Resources



Product Sheet

AvailabilityGuard/
Enterprise



e-Book

The Agile IT Operations
Approach to Resiliency



Video

4 Principals to
Preventing Your Next
Outage