Getting to the Bottom of IT Outages:

Ensuring the Resilience of Your Hybrid IT Environment

Table of Contents

IT outages affect millions

In our digital world, unplanned downtime at an organization can be anything from detrimental to disastrous. A recent report from IDC stated that occasional service unavailability, technical failures, unrecoverable data, and malicious attacks have been experienced by close to 100% of enterprises.[1] That is, major industries (banks and financial enterprises, telcos, and airlines to name just a few) that depend on 24x7x365 availability and data, regularly experience sudden downtime. And that's the good news!

Other more severe, longer-lasting IT outage incidents can be termed "disasters," though not natural ones. Such events of late include the unavailability of a major American telecommunications conglomerate's internet service throughout the US caused by a "configuration error"; a "system-wide tech failure" that disabled air traffic monitoring, causing the delay of close to 15,000 flights across Europe; and, the weeks-long meltdown at a major UK bank when their migration of 1.5B records to a new banking platform resulted in customers being closed out of their accounts and led to phishing schemes and theft of funds from accounts.

And, let us not forget natural disasters such as hurricanes, tornadoes or other heavy storms. These often end up cutting the service of organizations located in the affected area, especially when the businesses in question did not failover to their backup sites in time, or if they did, failover failed.

The costs of outages: Outages of such magnitude cause the organization involved to take a beating on several levels: recoverability and restoration of data and return to working systems are difficult and time-consuming processes, customers suffer harm, the company's reputation becomes tarnished, great amounts of time and expense are invested in managing and repairing the negative effects of disastrous events and still, business may be lost and customers may switch over to competitors. A recent estimate of the financial costs associated with downtime put the figure at $700B a year[2] in North America alone!

Major corporations are not achieving IT resilience. A seriously disturbing aspect of these outages is that they occur in some of the largest corporations in the world – businesses that are well funded and which probably invest millions of dollars in building a redundant IT system that should withstand any catastrophic scenario.

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

This whitepaper focuses on uncovering the root cause of IT outages and unplanned downtime and how to proactively prevent them.

[1] https://www.zerto.com/the-state-of-it-resilience-2018/

[2] https://technology.ihs.com/572369/businesses-losing-700-billion-a-year-to-it-downtime-says-ihs

Outages: Why?

IT outages occur for a variety of reasons ranging from human error to extreme weather, as well as technical problems within server environments. When outages occur, organizations generally issue "official" explanations that are vague and mention "power switches," a "technical glitch" or an "overload of error messages," etc. Ambiguity regarding the reasons for the outage are frequently due to the company itself being unsure of where the failure originated. They lack visibility into the development or progression of the problem and consequently, the reason provided to the public typically describes a result or symptom of the root cause - which remains obscure. This is the case even within environments built to have high availability and redundancy.

Finding the cause of an IT outage can be like figuring out clues in an escape room.

Complexity and interconnectedness of IT environments make problems hard to untangle. The reality is that IT environments are becoming increasingly more complex and interconnected/interdependent. Currently, 71% of enterprises use a hybrid-cloud environment that is comprised of up to an average of five public and private clouds.[3] Many others combine cloud environments with an on-premises datacenter.

In a typical, complex, modern IT infrastructure, new systems are introduced all the time, and must be configured correctly in order to achieve resilience. A prodigious task to begin with, it is further complicated by the almost daily stream of updates and changes made by in-house and third-party IT teams, some of which can negatively affect the system in unpredicted ways as well as lead to misconfigurations. Such environments are prone to malfunction resulting from errors and single-points-of-failure that affect entire IT stacks. As a result, maintaining the highest levels of IT resilience is becoming increasingly harder.

In addition, the vendors/providers of these systems regularly, even daily, issue updates to their best practices, which are meant to be implemented by IT teams. And, in general, every environment at every organization is backed up with a twin, redundant environment that must also be changed and configured when vendor updates are issued. However, the reality is that IT teams do not get around to implementing recommended updates because of their sheer volume.

The following formula distills these dynamics:

$$\{\textit{Ongoing change}\} \times \{\textit{Multiple teams/ vendors/ solution providers}\} \times \{\textit{Thousand of [ever evolving] best practices}\} = \textit{Faliures waiting to happen}$$

What can IT teams do to control these unwieldy environments?

How do IT teams generally reduce risk to IT envirnments? In the effort to detect and contain risk, IT departments engage in accepted practices to test the resilience of their IT environments. That is, on a quarterly or half-yearly basis, they run comprehensive tests to check if they can safely move to their secondary environment and if they need to improve resilience. In the weeks leading up to each test, IT teams attempt to identify and resolve all the existing problems including the real risk that the test itself could have an adverse impact on normal operations. After all this intense preparation, many environments do not pass IT resilience tests. And, even for those that do, following the test, with every day that goes by, because of the vast number of changes that are a normal part of maintaining any IT environment, the risk to resilience grows and quickly becomes an unknown quantity.

Does such testing raise confidence levels? The following diagram illustrates why it does not:

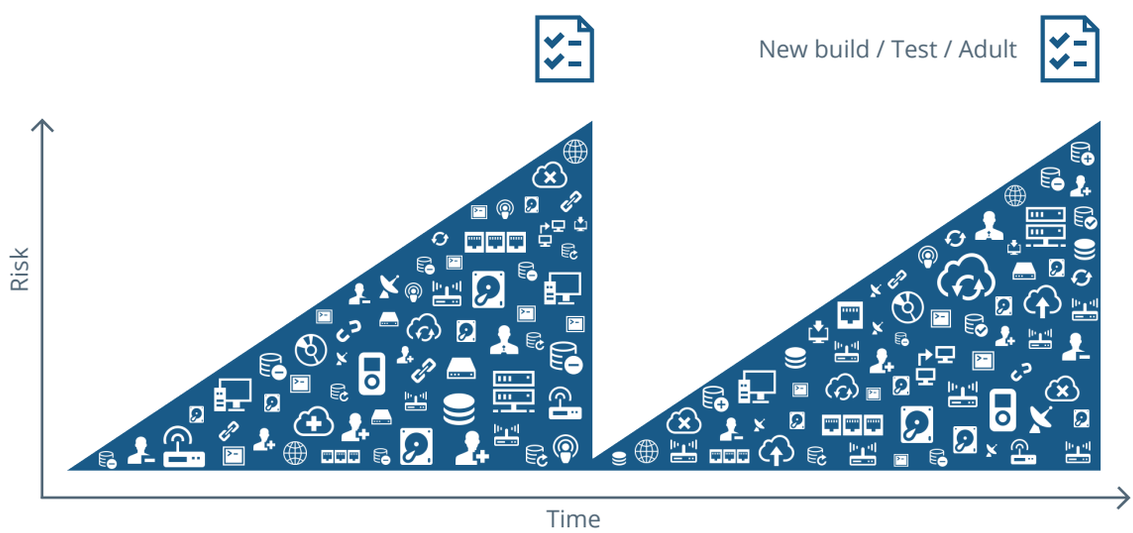

As soon as IT issues are under control and tests are successfully passed, new issues immediately pile up.

The irony is that even correctly planned and executed tests, while significantly reducing risk, unfortunately have an extremely short-lived positive effect on resilience. The rate of change in IT environments, described above, guarantees that risk to resilience continues and is ever-present.

Making a conceptual paradigm-shift in order to reduce risk level

Since IT teams cannot keep up with the vast number of changes to practice that are recommended by vendors and because caring for IT infrastructure is subject to human error and affected by other factors as well, periodic 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. As a result, current test methodology is inadequate. *So, how can companies attain continuously low-risk levels?*

First and foremost, they must realize that achieving availability and resilience objectives must be an ongoing process that is continuously planned, executed, and measured. The frequency of validation must match the rate of change, which means - all the time and in a continuous manner.

A new approach is required which assures that on going validation would prevent the accumulation of large risk waves between each test cycle. This will increase confidence in the ability to recover at any time.

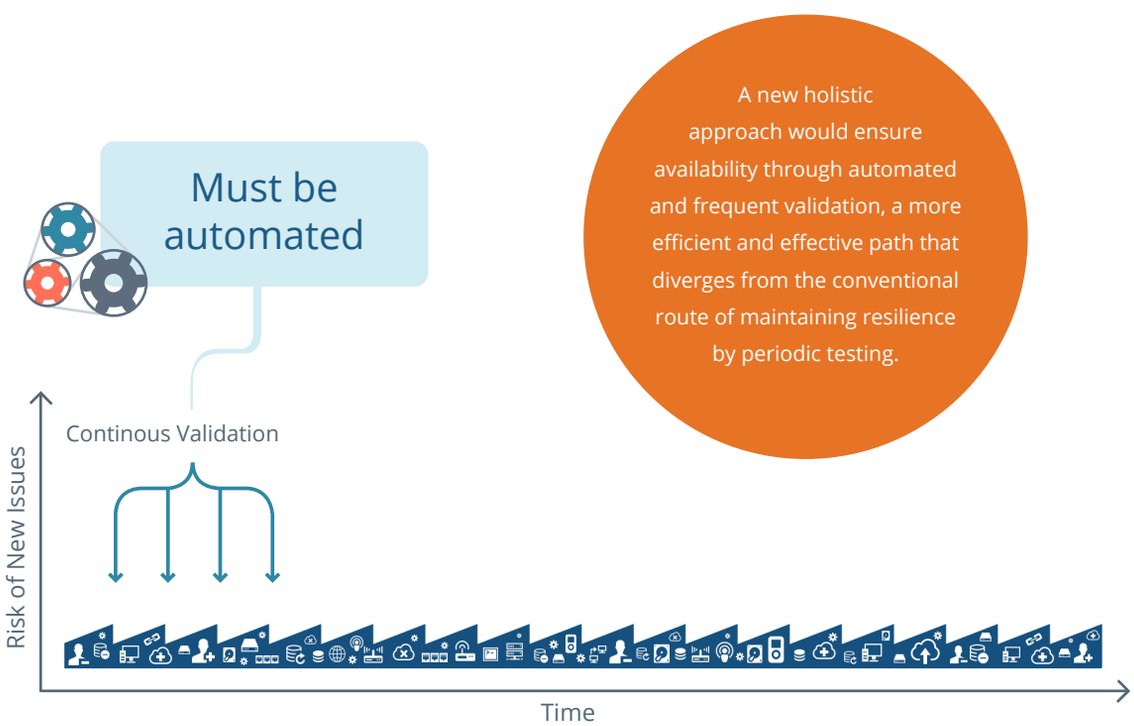

AvailabilityGuard™ - putting theory into practice

Continuity Software applies this new holistic approach, as exemplified by the AvailabilityGuard™ resilience assurance solution. Its design and operation reflects the understanding that IT environments change all the time and resilience must be continually checked, analyzed and verified to ensure that they are healthy and available.

To facilitate this goal, AvailabilityGuard is:

Automatic

Uses a deep knowledgebase of vendor best practices

Evaluates detected errors with respect to potential risk to the business and the technologies used

Provides IT with a detailed protocol for repairing the errors

Reports on the status of the secondary environment

The vendor-agnostic solution empowers IT organizations to proactively detect misconfigurations and eliminate outages across all IT infrastructure layers. Operationally, it relies on regular and frequent non-intrusive scans of all IT infrastructure configurations, collecting data from all major vendors and tools across all the IT layers. It compares scan results to its enormous proprietary knowledgebase of vendors' best practices which also includes input from actual users of the relevant technologies. The solution was created to be used by IT teams as often as every day to examine the health of their infrastructure environment and to then make immediate repairs to problems uncovered.

The AvailabilityGuard approach is based on four pillars that support resilience assurance of hybrid IT environments:

1. Continuous & proactive resilience assurance

The solution proactively scans the entire IT infrastructure to detect risks to resilience and single points of failure. This continuous process leads to improved and more efficient IT infrastructure since risks are resolved in a timely manner, before they escalate into costly service incidents. Scans can be scheduled at the IT team's desired frequency and/or be event-driven. They are agentless and carried out in read-only mode, ensuring that enterprise data is undisturbed.

2. Support for hybrid IT

The solution meets the challenges of ever-more-complex hybrid IT environments that can include a mix of physical, virtual, on-premises and cloud infrastructures. It automatically scans and inspects for cross-domain and cross-layer resilience risks and checks for misconfigurations that could affect recoverability. It automatically discovers the connections and interdependencies between the various components across the layers that play a role in IT resilience, while correlating all these resources to the business applications.

The various configurations, dependencies, and resilience schemes and how they are related to the various business applications are mapped for IT teams to see. Mapping is continuous and dynamic, providing up-to-date views of the scanned systems.

3. Deep knowledge base

The information gathered in scans is compared against a vast knowledge base of technology vendors' best practices augmented by user community input. Currently, more than 7,000 guidelines exist on how misconfiguration issues, large and small, should be corrected.

The knowledge base - probably the largest repository of its kind – goes back more than a decade to include knowledge on just about every system used in all types of IT environments anywhere in the world. This library continues to grow and accumulate more information so that the very latest know-how available can be utilized to assure IT resilience.

4. Pursuit & realization of operational excellence

The solution filters, forwards and assigns incident tickets to the relevant stakeholders and includes a protocol for resolution. All information is presented on the solution's dashboard which includes visual representation of data infrastructure status, the tickets to be resolved, and the steps for resolution. Issues are marked by their level of severity and IT resilience quality and risk are analyzed with respect to their effects on the technologies used and on the enterprise. Existing incident management and ticketing systems such as ServiceNow, IBM Tivoli, BMC Remedy, and others, can be used to manage incident tickets.

AvailabilityGuard dashboard of issues discovered .

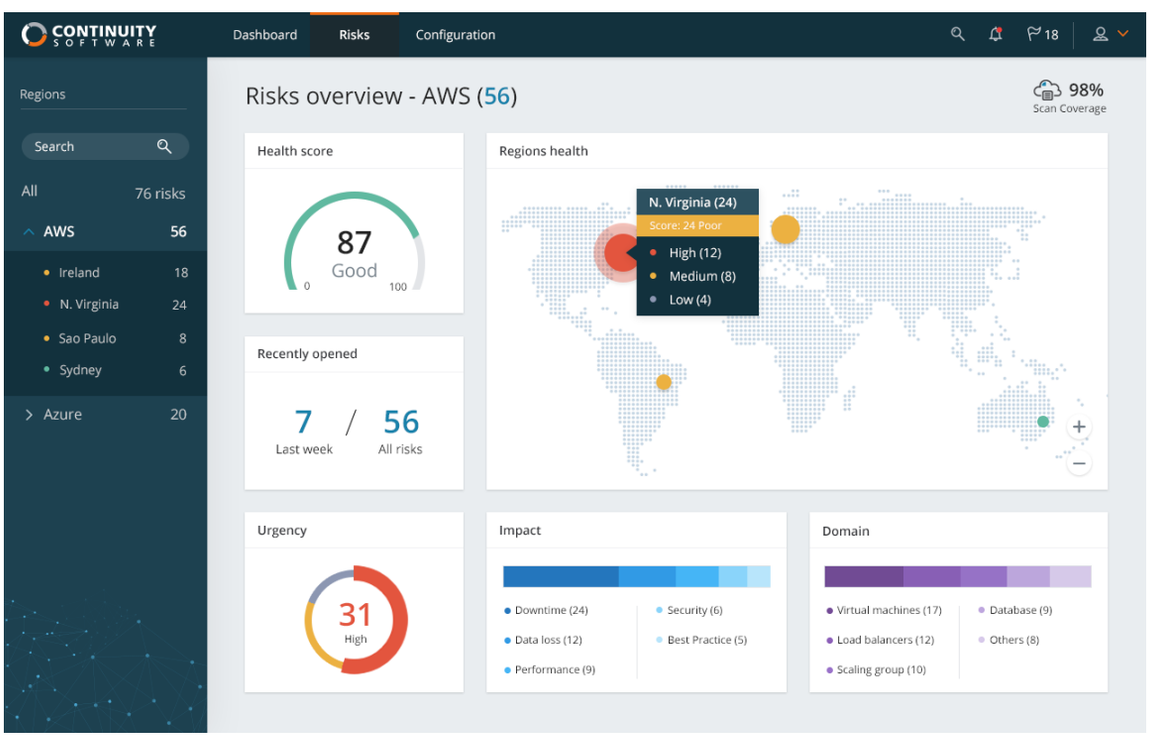

Why implementing AvailabilityGuard makes sense

It establishes an ongoing and simple process that enables IT teams to gain more control over assuring resilience and availability – in the most complex of new and existing environments.

It quickly validates and improves operational resilience in all environments

It enables organizations to attain and maintain high availability, recoverability and stability of all environments

It reduces operational costs due to significantly less downtime, and other related and acute incidents

Ever-increasing IT environment complexity is the new normal. AvailabilityGuard provides the tools and the roadmap for achieving resilience for modern IT environments.

About Continuity Software

Continuity Software helps the world's leading organizations, including 6 of the top 10 US banks, to achieve resilience for their hybrid IT environments.
As a global leader in IT resilience assurance, our solutions proactively prevent outages and data loss incidents on critical IT infrastructure.

As a result, unplanned infrastructure outages are reduced by an average of 80% and configuration errors are resolved before they turn into costly service incidents. Continuity Software was founded in 2005 by a team of experienced IT and data protection professionals, with a passion for building innovative enterprise software solutions.

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