IBM Global Technology Services
Thought Leadership White Paper

June 2015

Resiliency in the Cloud

IBM

Contents

Executive summary

In today's mobile and socially connected world, users expect to connect wherever and whenever they want, across multiple access platforms and multiple locations. Any remaining tolerance for downtime and data loss has now evaporated in the face of the always-on world. Caught between the two extremes of low-cost disaster recovery solutions and high-cost replication, many organizations are embracing cloud models as a more affordable way to meet requirements for rapid recovery of systems and data.

Cloud undeniably offers enormous potential as a way to cut costs, increase agility and reduce risk. But failure to properly plan for and implement resiliency can result in an unintentional increase in overall enterprise risk. Many users of cloud solutions assume the existence of a level of continuous availability that is not necessarily engineered into all cloud services. Put another way, just because the potential for improved, more agile resiliency exists, it does not mean that resiliency is automatically attained simply because a workload has been migrated to the cloud. When it comes to resiliency, cloud must be viewed and assessed as another technology domain—with functional and non-functional resiliency requirements to be defined, assessed, measured and monitored.

The right decisions on cloud are critical for organizations to reduce the overall spending and increase the ability to respond to cloud related risks, threats and opportunities. Yet however necessary, identifying requirements, risks, prioritizing them and allocating funds to address them is not always easy. In order to do this, organizations need to gather and analyze the right information to make value-driven decisions regarding the cost effective management of risks related to resiliency.

Whether migrating workloads to cloud-based platforms or pursuing a Disaster Recovery as a Service (DRaaS) model, cloud requires a fundamental shift in thinking about integrated enterprise risk management. While there is a pervasive lack of resiliency planning in most cloud implementations today, better up-front assessment and planning can help organizations realize the enormous potential cloud offers for improved, more agile resiliency and strike the right balance between business service availability requirements and tolerance for risk.

The tipping point of resilience

In past years, ecosystems of internal and external systems used to conduct daily business were more restricted and the risk of exposure was much less. In a typical organization today, users engage with multiple systems leveraging information and services contained internally and externally to perform work. As the number of systems, data and users grows—along with reliance on analytics and system interdependencies—so too does the complexity and necessity of resiliency. For example, when an online order is placed, success for that transaction depends on the website, the ecommerce application, the back-office transaction system, the system that tells the warehouse what to pick and where to send it, and the inventory system that

refreshes the stock in the warehouse. If any one of those goes down, the entire process is disrupted, including the systems of engagement and systems of insight to engage the customer up front to make the purchase, the systems of engagement for the customer to track delivery and the systems of insight for the business to track customer experience and sentiment.

Have you reached the tipping point of cloud resilience**? View the infographic:**

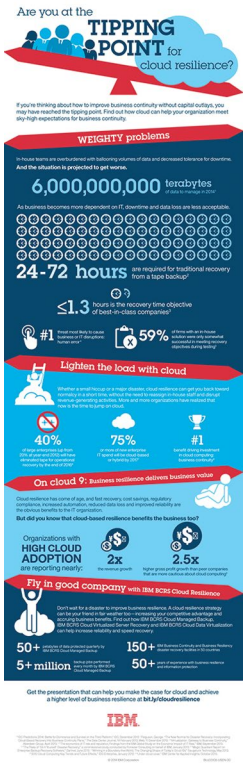

The cost and effort of maintaining resiliency infrastructures tips many organizations beyond traditional disaster recovery technologies into cloud resilience. With the capacity for quick failover, a cloud solution can help ensure near continuous application and data availability no matter how geographically dispersed the organization's systems and users may be. Cloud can also bring unprecedented scalability to business resilience and the ability to rapidly provision new services—features that many commercial resiliency and disaster recovery providers have offered for many years. The difference is that cloud can now bring some of these features under the control of the customer.

Yet to date, few applications have been designed to take full advantage of the resiliency capabilities cloud can offer. Much of the current focus on cloud-related projects is around identifying and migrating target workloads to the cloud as quickly as possible. Virtualization tools and all cloud solutions are not inherently more resilient. In addition, new cloud-specific services such as brokerage, orchestration capability and the service catalog systems can represent critical single points of failure which if happens, requires stoppage of services to preserve data and transactional consistency. If resiliency is not integrated into initial cloud adoption, organizations are accepting risks whether they realize it or not. While there is consistent reference to leveraging the cloud for disaster recovery, there is little guidance on how to use it effectively, including testing to ensure that cloud-based disaster recovery and resiliency strategies will work as expected.

If organizations are not integrating resiliency into their initial cloud adoption, they are accepting risks, whether they realize it or not.

The reality of cloud resilience

When it comes to resilience, most organizations find these questions difficult to answer for their legacy environments:

- Can you reliably quantify the cost of an hour of downtime?
- Can you provide accurately tested evidence to show how quickly you can resume business operations?
- Do you know where your data corruption concentrations of risk reside?
- Do you know the extent of the impact should something happen?

On average, an infrastructure failure can cost $100,000 an hour and a critical application failure can cost $500,000 to $1 million per hour:[1]

Even fewer organizations can reliably answer these questions for a cloud environment. Cloud deployments typically include unpredictable traffic patterns and sporadic large-scale variations in transaction volume and capacity due to changing business requirements. Architectures are more open, with multiple vendors, ISPs, management systems, connection options and technologies. With changing technologies and emerging standards, cloud can bring significant increase in complexity plus a level of volatility that further compounds complexity. Continuous awareness and real-time monitoring of what runs where is essential.

With any cloud deployment—be it a Disaster Recovery as a Service model or other cloud-based solution, the complexity and impact of outages requires careful involvement and oversight of resilience experts and the application of formalized methods and techniques with adequate testing. For example, moving storage and processing capacity offsite helps reduce the concentration of risk and impact arising from a single site outage. However, shifting to another site without careful planning and design in a resilient strategy and architecture can result in an overall increase in organizational risk rather than reducing it. Concentrations of risk even in multi-site deployments can create bottlenecks and single points of failure. And while cloud assumes continuous access to data and compute power, and replication plays a key role—it also carries the risk of replicating corruption.

Systems recovery across multiple cloud service providers and locations requires high level IT integration skills.

Organizations also need to think about how they will integrate cloud solutions back to the legacy production environment. Most large enterprise cloud-based workloads interact with workloads running on legacy infrastructure environments, which means this interoperability needs to be taken into account for both business as usual and in disruption situations. Painstaking coordination and orchestration of the new hybrid environment is required—especially when it comes to resiliency-related testing and actual disruption failovers when various diverse workloads need to be shifted around and re-synchronized.

Synchronization is more than making sure that the secondary copy of data matches the primary and automation is more than scripting the application restart procedure.

Regulatory pressure for organizations to prove they can continue operations using backup systems continues to increase. Providing evidence that the resilience response meets the needs of the business requires testing that effectively re-creates the user experience—thereby confirming the effectiveness and value of the resiliency program. In the past, regulators have been primarily interested in systems of record. Today, we now see expectations for an always-on experience focused on systems of engagement and increased business reliance on systems of insight to enable real-time processing for in-place marketing, sentiment analytics and unstructured data analysis. The complexity and interactive dependencies between these three processing families need to be included in resiliency strategies and designs. Today, every organization's reputation is built or lost through the eyes of the customer in real time.

Testing in a cloud environment brings its own set of challenges. For example, true disaster recovery testing will require isolating environments and connection paths since (you cannot have your production IP addresses live in two places at the same time - your data center and your cloud resiliency solution). That means figuring out how to have a second set of IP addresses that are masked in the production environment. This will allow comprehensive testing without contaminating the production activity.

IP address management and isolation is more critical than ever in cloud. Getting it wrong can expose the enterprise to huge risk when the internet becomes part of your production processing and resiliency testing environment.

In short, the traditional and narrow scope of validating and restoring IT is no longer sufficient. Nor is running a resiliency environment the same as running a production environment.

Cloud resiliency requires skills and experience to ensure the right design, architecture, orchestration, management, monitoring, reporting and governance are in place to make it all work. It may be simple to stand up a disaster recovery environment in the cloud but much more needs to be done in terms of recovering critical systems of record, engagement and insight to provide proof that the organization can continue to run the business in the event of a disruption.

The right approach to cloud resiliency

The reality of cloud resiliency is that if you don't design, implement and maintain it, you don't get it. But how does an organization ensure they get it right from the beginning? Best practices suggest doing it right takes a structured approach that lets organizations:

- Understand the criticality of business services supported by target cloud workloads, associated business and IT resiliency requirements, and cloud/legacy application/data dependencies
- Identify related risks and required mitigation treatments
- Determine best fit cloud strategies that meet business based resilience needs
- Document linkages, interdependencies and synchronization points between cloud and legacy environments
- Develop, implement, test and sustain corresponding business and technology resiliency plans and procedures
- Create cloud specific and integrated cloud or legacy resiliency validation and testing plans
- Develop resilient cloud transition roadmap

It can be daunting to undertake the design and planning of a resilient cloud infrastructure, let alone implementation and management. However, there are some key tactics and steps organizations can take to plan and manage cloud resiliency. Figure 1 outlines the four steps of the IBM Resilient Enterprise Blueprint methodology that considers cloud resiliency from assessment through management.

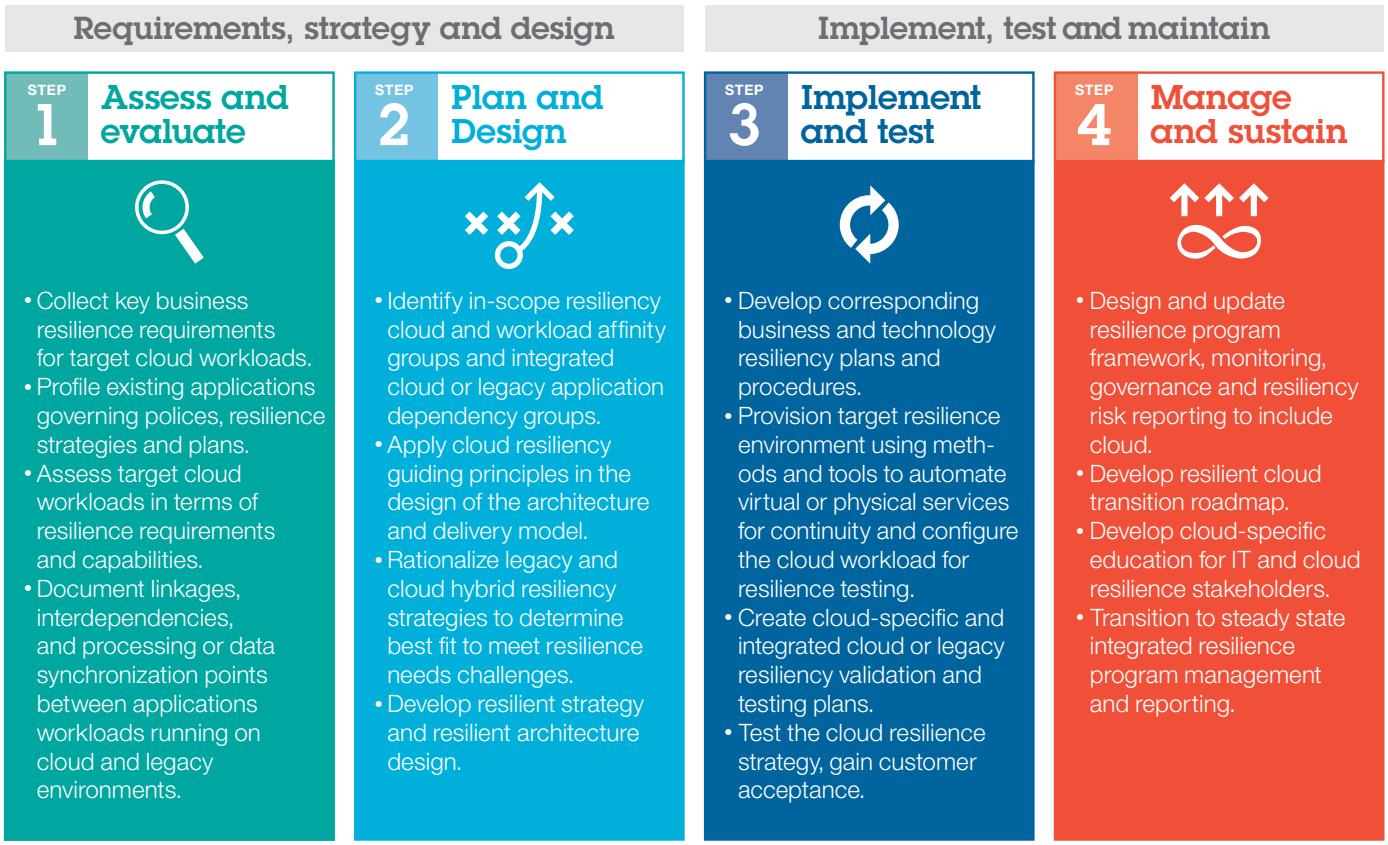

Figure 1. The four steps of the IBM Resilient Enterprise Blueprint methodology.

Step 1 – Assess and Evaluate

Few, if any, organizations can afford to make their entire IT environment "always-on." Instead, resiliency should be linked to business value. A structured approach can help organizations carefully identify business requirements from the resiliency perspective that will drive resiliency objectives and metrics (e.g., key performance and key risk indicators). You can start by identifying and documenting tiers of resiliency and getting agreement from your business users as to what the recovery time and recovery point objectives are for each tier and what business functions will be mapped to which tiers. This will give you some yardsticks to use when you are evaluating different cloud features and functions.

With this information, organizations need to be able to understand the services and infrastructure interdependencies and impact they have on attainment of these business objectives—while expecting them to shift. After you understand the risks,

you can develop an appropriate resilience strategy and architecture that not only meets the needs of the business—but enables you to prove it. Key activities in this process include:

- Collecting key business resilience requirements for target cloud workloads
- Profiling existing application chains with associated infrastructure components and data, governing policies, resilience strategies and plans
- Assessing target cloud workloads in terms of resilience requirements and capabilities
- Documenting linkages, interdependencies and processing/data synchronization points between applications workloads running on cloud or legacy IT environments

Organizations cannot afford to make their entire IT environment "always on." Resiliency needs to be linked to business value.

Step 2 – Plan and Design

Cloud design must include the fully integrated manner in which organizations "operate the business," including service dependencies. However, high levels of abstraction in the cloud can create challenges in identifying and documenting service dependency blueprints. That's where the Cloud Computing Reference Architecture can help with management tool sets that include provisioning, replication, automation, monitoring and reporting.

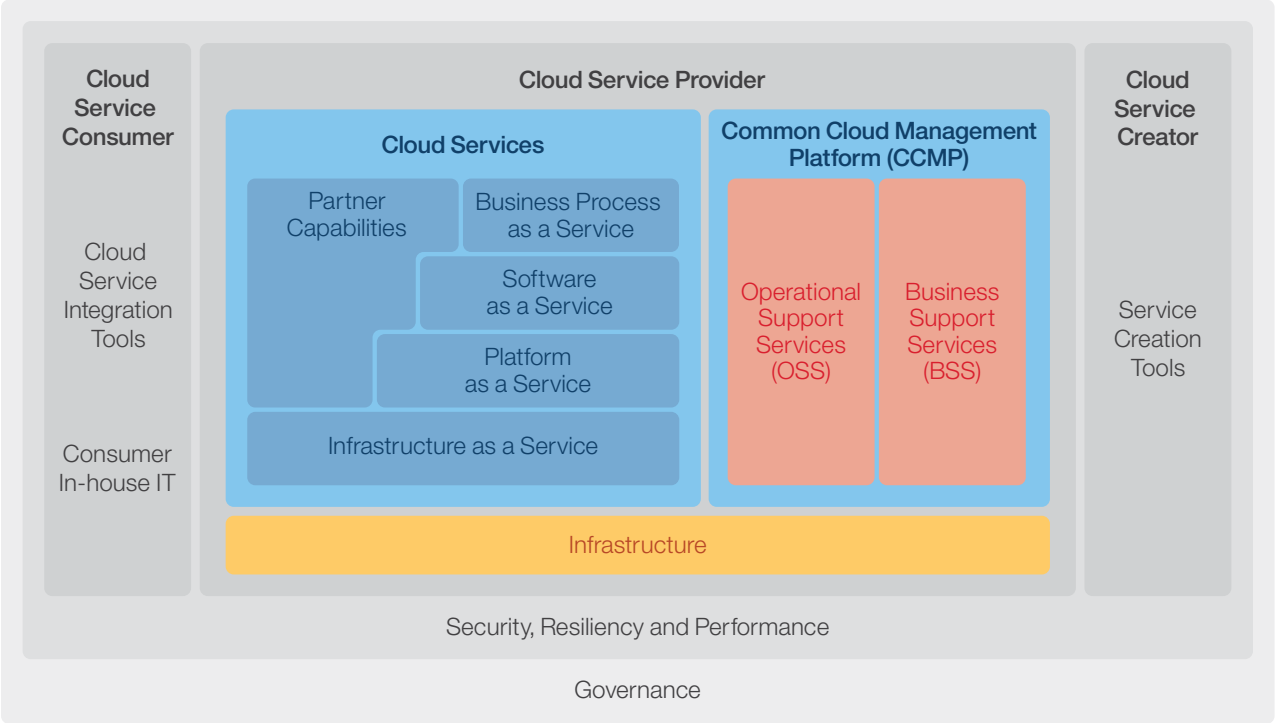

Figure 2. The Cloud Computing Reference Architecture can help identify interdependencies.

To help leverage these tools effectively, you can combine the reference architecture with IBM's Resilient Enterprise Blueprint methodology to provide a clear understanding of application level interdependencies, shutdown and restart sequencing as well as processing and data synchronization points for workloads across both cloud and traditional IT environments. Levels of concentration of risk should also be documented whether due to potential capacity bottlenecks or single points of failure.

The Resilient Enterprise Blueprint methodology can help you understand whether you have a risk that you can accept or whether you have a risk that you want to avoid and mitigate. In other words, you may choose to do nothing about a risk, or you may improve your infrastructure to help ensure that you can handle events if they occur. Only when the complete set of functional and non-functional requirements (including resiliency) are understood can organizations move into the planning and design phase, which includes:

- Identifying in-scope resiliency requirements for cloud and workload affinity groups and integrated cloud/legacy application dependency groups
- Understanding the impact of provisioning lead times on recovery time and recovery point objectives
- Applying cloud resiliency guiding principles in the design of the architecture and delivery model
- Rationalizing legacy and cloud hybrid resiliency strategies to determine best fit to meet resilience needs
- Developing resilient strategy and resilient architecture design

Step 3 – Implement and Test

False perceptions about cloud have perpetuated the idea that testing is simply a matter of provisioning what you need, performing the testing and then de-provisioning to go back to where you were. Properly creating and designing a testing environment requires much more than simply activating an infrastructure or dumping a copy of data over to the cloud. Organizations need to be able to maintain stringent IP address management, jurisdictional boundaries over data copies and an evidence trail that they not only performed the testing but can produce expected to actual results.

Before implementation and testing, you'll want to think about the end results: realistic testing and being able to "run the business" on an alternate environment when faced with a major service disruption. With this information, you can develop corresponding business and technology resiliency plans and procedures by:

- Provisioning target resilience environment using methods and tools to automate virtual or physical servers for continuity and configure the cloud workload for resilience testing
- Creating cloud specific and integrated cloud or legacy resiliency validation and testing plan (may include regular introduction of planned failures to test real time component/service resiliency)
- Testing the cloud resilience strategy and gaining customer acceptance
- Retaining evidence of test scope and outcomes for audit and reporting purposes

Partial or component level testing combined with blind faith that things will "work themselves out" at time of disruption is no longer acceptable.

Step 4 – Manage and Sustain

The speed and flexibility of cloud can bring great agility to users, but this is often accompanied by a high degree of volatility as well. Robust governance, management and control processes are required to keep resiliency capabilities synchronized with the production environment. Monitoring and reporting are key, as resiliency is not a "once and done" project. Instead, it's an operational process and executives need to know the current state of enterprise at any time—and especially in the event of an unplanned outage. Steps to manage and sustain resiliency capabilities include:

- Designing/updating resilience program framework, monitoring, governance and resiliency risk reporting to include cloud
- Developing resilient cloud transition roadmap
- Developing cloud specific education for IT/Cloud resilience stakeholders
- Transitioning to steady state integrated resilience program management and reporting
- Maintaining appropriate checks and balances for:
  - Assurance (including third-party) and attestation
  - Application readiness
  - Process readiness, vendor stability and reputation, mobility to migrate to another location/vendor
  - Continuity requirements
  - Network
  - Data—location, protection, segregation
  - Governance, risk and compliance (GRC)

Conclusion

In an always-on world, the complexity and impact of outages requires more attention to the design and management of resilience solutions. Service-based resiliency in complex multi-provider systems means that you need to think differently. The evolution of cloud brought expectations of significant efficiencies and economies for disaster recovery that have largely been unrealized. Today, there are still relatively few properly designed, implemented and tested cloud resiliency solutions.

The reality is that the combination of traditional and cloud-based IT environments isn't going away any time soon, nor do the assumed "built-in" disaster recovery capabilities of cloud simplify resilience—in fact, they make it more complex. Risk can be significantly increased when organizations acquire and deploy cloud solutions without the involvement and oversight of resilience experts and the application of formalized methods and techniques with adequate testing.

From innovative cloud services to full scale compute, data and applications resiliency solutions, IBM has the proven expertise, knowledge and technology to help enable built-in—not bolted-on—resilience of your business in the face of threats as well as of opportunity across all layers of your enterprise and in any IT environment, including public, private or hybrid clouds.

Learn more about how the four-phase resilient enterprise blueprint methodology has been applied to developing continuous availability using cloud:

Your journey to always-on in four steps

Our business resiliency services and time-tested solutions can help you plan and design for the implementation and management of cloud resilience, with a strong commitment to understanding ever-changing business requirements.

For more information

To learn more about IBM Resiliency Services, please contact your IBM representative or IBM Business Partner, or visit the following website: **ibm.com**/services/continuity

Additionally, IBM Global Financing can help you acquire the IT solutions that your business needs in the most cost-effective and strategic way possible. For credit-qualified clients we can customize an IT financing solution to suit your business requirements, enable effective cash management, and improve your total cost of ownership. IBM Global Financing is your smartest choice to fund critical IT investments and propel your business forward. For more information, visit: **ibm.com**/financing

IBM Global Technology Services
Route 100
Somers, NY 10589

Produced in the United States of America
June 2015

[1] IDC, "DevOps and the Cost of Downtime: Fortune 1000 Best Practice Metrics Quantified." Stephen Elliot. December 2014, IDC #253155.

Please Recycle

BUW03048-USEN-00