



YES, YOU DO NEED TO THINK ABOUT RESILIENCE AND DISASTER RECOVERY IN “THE CLOUD”

A FIVE-PART SERIES TO STIMULATE DISCUSSION ON THE TOPIC
OF RESILIENCE IN THE CLOUD

Table of contents

Executive Summary	01
PART ONE	02
Introduction	03
PART TWO	05
Resilience for Workloads Deployed “in The Cloud”	06
PART THREE	07
Using “The Cloud” to Extend Your Resilience Strategy	08
DR as a Service (DRaaS)	08
Technology Synchronization	08
Backup as a Service	09
PART FOUR	10
The Resilience Continuum: Exploring the Impact of the Evolution of Cloud on Resilience Strategies	11
Resilience Continuum	11
PART FIVE	13
DevOps and Resilience	14
What about Re-purposing Non-production Environments for Recovery?	14
Conclusion	16
About the Author	16
About GIS	16

Executive Summary

Cloud changes the conversation about business resilience, but it does not render the consideration of resilience insignificant, not in the least.

- You must still consider resilience and recovery in the design of systems deployed “in The Cloud”.
- The application of cloud-enabled technologies, with careful planning and understanding, can have a net positive impact on an enterprise’s overall resilience and recovery posture.
- To best optimize cloud for application deployment requires collaboration between Development and Operations in the model of **DevOps**, more so as the technologies enabling infrastructure independence continue to advance.

This series puts forth an introductory level discussion on several topics related to resilience and recovery in the context of “Cloud”. It is intended as a primer to outline the key considerations for enterprises moving forward in their journey to Cloud.

The marketplace is demonstrating an increase in the frequency and depth of conversations related to cloud and resilience – both resilience of workloads deployed in The Cloud, and leveraging cloud in overall enterprise resilience strategy. There are a number of reasons for this heightened attention, but two primary drivers appear to be:

- Enterprises who are both the most at risk and have the most resources to invest are, after years of consideration, actively moving business critical workloads to The Cloud, and
- Line of Business leaders who are taking a more significant role in defining IT strategy and have been leading the adoption of cloud are becoming more IT savvy and are asking harder questions about the readiness of cloud for their workloads.

This discussion outlines the key trends and puts forth several questions and topics that should be considered by enterprises moving forward on the journey to cloud.

Cloud can be an ambiguous term. For purposes of this discussion, “**The Cloud**” is defined as IT services that are made available in an on-demand model, on a pay-for-use

basis, specifically owned by somebody else, hosted by somebody else, running somewhere else, accessible over the Internet - essentially, Public Cloud. When the discussion is referring to the general abstraction of cloud as a consumption model for IT service, it is just “**Cloud**”. The focus of this discussion is primarily Infrastructure as a Service (IaaS), though the conversation also applies to Public Cloud Platform as a Service (PaaS). In the case of Software as a Service (SaaS), the responsibility for resilience should lie completely with the service provider, though some of these concepts may be useful in performing due diligence in evaluation of SaaS offerings.

As the conversation about resilience in The Cloud has generally moved from denial – “We don't need to worry about resilience and DR if we're in The Cloud” – to questioning – “Do we need to consider resilience and DR for workloads deployed in The Cloud?” – to acceptance – “Can you help us plan for resilience and DR of our workloads deployed in The Cloud?” – there is a clear trend towards greater appreciation for the importance of resilience planning in cloud deployment. With so much at stake and so much change in the marketplace, it is important to continually stay abreast of the topic. Where does your enterprise fall within this continuum?

While cloud presents a great opportunity to drive increasing value and differentiation from IT, the (often complex) challenges of resilience and recovery must still be considered and addressed. The world is nothing if not a risky place. The people with the skills and disciplines associated with resilience and DR must be actively engaged as we, consumers and providers of IT services, move forward to cloud.

Starting with a look at the importance of including resilience planning in the deployment of workloads in “The Cloud”, then examining various aspects of leveraging cloud resources for resilience and recovery, and finally exploring how the evolution of cloud technologies will impact resilience planning, this series will answer the question, “Why do you need to think about resilience and DR if you have The Cloud?”

For the purposes of this discussion the scope of resilience and DR is contained to the IT environment. The full scope of

everything needed to ensure business continuity would be a series without end, more like a book, and the basic message of it would be lost. Further, there is distinction between “resilience” and disaster recovery. Resilience implies a broader context and a more proactive approach to creating an environment that can absorb the realization of any number of risks and still keep serving the business. Disaster Recovery implies a reactive approach to being able to recover after the fact.

The next topic in the series, Part Two, will consider the topic of Resilience for Workloads Deployed “in The Cloud”. The discussion will, hopefully, dispel the notion that you don't need to worry about resilience or recovery if you have your workloads deployed “in The Cloud”.



THE PEOPLE WITH THE SKILLS AND DISCIPLINES ASSOCIATED WITH RESILIENCE AND DR MUST BE ACTIVELY ENGAGED AS WE, CONSUMERS AND PROVIDERS OF IT SERVICES, MOVE FORWARD TO CLOUD.





PART TWO

Resilience For Workloads Deployed “In The Cloud”

This part of the series explores the often expressed sentiment that if you go to The Cloud, you don't need to worry about resilience or recovery, or at the very least, it is somebody else's problem. That is very rarely the case, and there are definitely many questions to ask and issues to consider.

Start by looking carefully at the services agreement of your cloud provider(s). On the surface, the SLAs are generally quite impressive, with numbers like 99.9 availability being typical. A closer inspection will very likely uncover that this is for availability of "the service". Meaning – you can get to the portal and provision a server or storage, somewhere in the vendor's cloud any time, but it says nothing about your workloads, or where you can put them. In all likelihood, your "content" (meaning your applications and your data) will be specifically excluded from the SLA.

What does that mean? For starters, it means **you still have to plan to recover your workloads**. If the cloud data center goes off-line for whatever reason, or your specific rack of servers takes a failure, or the vendor does a rolling emergency outage, or a security breach shuts down some or all of the servers you're running on – what are you going to do to keep your applications running?

Most leading commercial public cloud service providers have facilities you can use to add resilience to your environment. Things like disk replication, image cloning, backup utilities and load balancing can be applied to your systems so that your workloads can stand a better chance of surviving through outages that are likely beyond your control.

However, implementing these facilities requires careful planning, design and regular testing. You'll need to have a plan in advance. While these services are generally available on an on-demand basis, making them work effectively in your

environment can take weeks, or even months. You shouldn't wait till you have a disaster and it becomes an emergency. Also, you shouldn't expect to order these services and turn them up without smart people, who understand how to design and implement resilience, in the environment you're using.

An important consideration while implementing the cloud service provider's resilience technologies will be the degree of vendor tie-in this creates. On the one hand, if the facilities are easy to use and work well, you'll want to apply them to save yourself extra work and improve your resilience posture. On the other hand, if you use their integrated features, will you lose flexibility in moving your workloads to other cloud service providers, or back in-house, should the need arise to do so? This is a key consideration, and illustrates a dilemma that long preceded cloud. That is - if resilience is custom-built into the infrastructure (cloud or otherwise) it may be well tuned to the needs of respective workloads, but it may limit the options available for recovery. Large enterprises have historically solved this problem by creating their own dedicated recovery environments, or outsourcing it to have somebody do it for them. Smaller businesses or smaller entities within large enterprises don't often have this luxury or the economies of scale to create dedicated resilience environments, so it may be more apt for them to rely on the Cloud. The point here is –to understand the implications.

We will discuss more on this later as the series explores the developing technologies that allow a higher level of independence from the infrastructure specifics. The next part, Part Three: Using The Cloud to Extend Your Resilience Strategy, will examine aspects of using cloud to extend your enterprise resilience and recovery strategy.



YOU SHOULDN'T WAIT TILL YOU HAVE A DISASTER AND IT BECOMES AN EMERGENCY.



PART THREE

Using “The Cloud” To Extend Your Resilience Strategy



This part of the series explores key concepts and considerations for using The Cloud to extend the resilience of your existing systems. This is a particularly attractive use case for cloud, but again, requires careful planning and testing to be successful.

DR as a Service (DRaaS)

The notion of using The Cloud for DR has gained momentum of late, creating a sort of new sub-category of cloud - DRaaS. Based on the concept that Cloud lets you acquire services when you need them and not until you need them, DR seems like a perfect workload. Who actually wants to pay for insurance until you need it? To capitalize on that idea, many of the companies who started as facilitators of migration to cloud are now offering their migration technologies (and expertise) - usually with higher levels of scripting and/or automation - as DR as a Service. While this concept seems to be getting most traction in on-premise/legacy to Cloud recovery, it can also apply in a Cloud-to-Cloud resilience strategy.

Sounds good, doesn't it? But, like anything else in this business, you need to look carefully at the details. This area is getting a lot of focus and will continue to improve rapidly, but consider the following:

Many of these technologies were invented to facilitate migration and they have been extended to be purposed for DR. Something you do once, in a carefully planned manner, usually at a leisurely pace, is not the same as something you need to do by instinct, quickly, during an emergency. You should look carefully at how much you have to build around the solution to make it highly automated, virtually fool-proof, and meet your recovery requirements (Recovery Point Objective and Recovery Time Objective). Whether you're using DRaaS to protect your cloud environment or your on-premise environment, getting this sort of capability working will require time, testing and skilled resources. It's highly recommended to plan accordingly.

Another consideration for these DRaaS offerings is whether or not the capacity of the target cloud will be there when you need it, possibly when a lot of other businesses also need it. This would be the case, for example, if the event that caused your outage is widespread. Commercial cloud vendors try to maintain excess capacity, but subject to peaks of demand, how their business is doing, or when they last brought online additional capacity. There may or may not be available capacity exactly when you need it and where you want it. You may want to look again at the services agreement with your cloud provider. Unless it is a cloud solution specifically built (and priced and contracted) for DR, or you specifically negotiated reserved capacity (which has a cost to it), then you may not be able to count on being able to provision a large number of servers, and a lot of disk capacity, all at once. Smaller businesses and/or smaller workloads may have less risk of insufficient capacity, but if the workload is critical, you should think about having an alternate plan in place.

One way to mitigate the risk of not having capacity available when you need it is to provision your recovery environment to be present and/or active all the time. If your design point is high availability (or very low RPO/RTO), which is often the case in the world of "always on" cloud applications, then this is very likely what you've done anyway. But, if this practice is just a hedge against capacity not being available when you need it, then did you just negate one of the anticipated benefits of using the cloud for DR?

Technology Synchronization

If you're using DRaaS, can you insure that technology upgrades in the target cloud environment don't cause breakage in your ability to recover? If you don't control both the source and the target, and/or if they are different vendors or of different cloud ecosystems, there is a pretty good probability that they will get out of sync, if they ever were in sync. You'll need to test and remediate frequently. Ask yourself, is annual testing going to be sufficient? (If it ever was!)



SOMETHING YOU DO ONCE, IN A CAREFULLY PLANNED MANNER, USUALLY AT A LEISURELY PACE, IS NOT THE SAME AS SOMETHING YOU NEED TO DO BY INSTINCT, QUICKLY, DURING AN EMERGENCY.



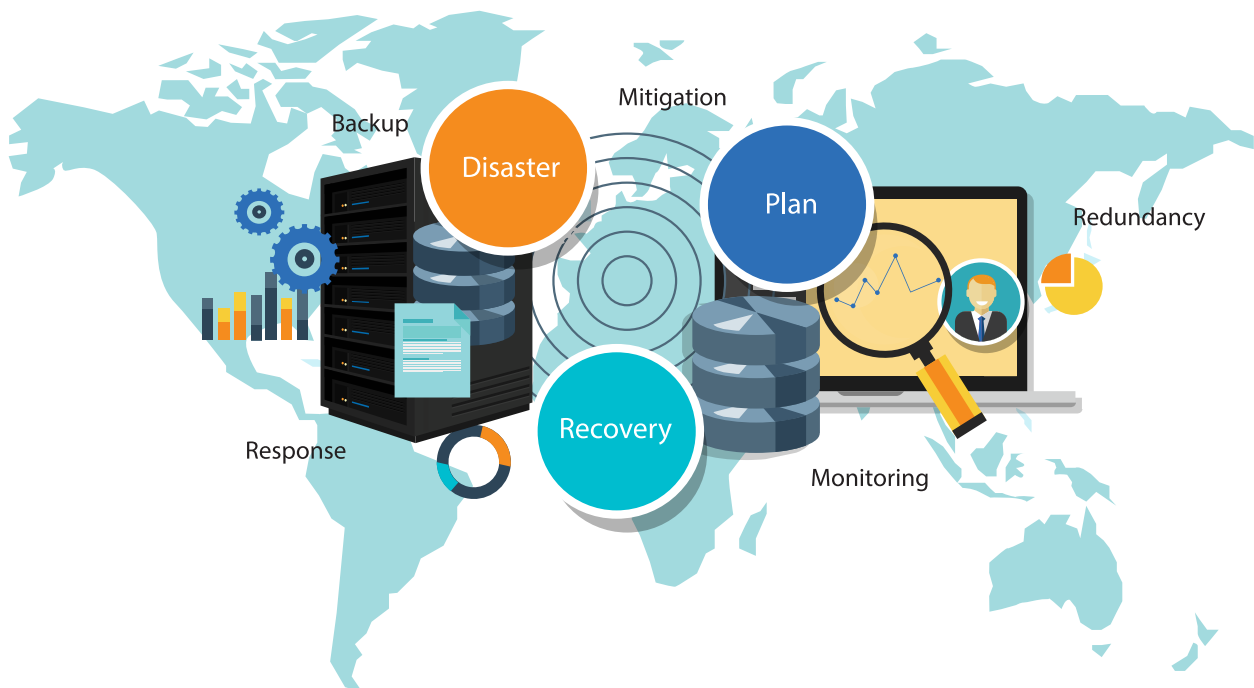
Backup as a Service


Closely related to, and sometimes used interchangeably with DRaaS, is Backup as a Service, or Cloud Backup. Backup is a wonderful thing, and cloud for backup is a great way to get your data encrypted, off-site, and still be fairly easily accessible. Backup is absolutely critical, even if you've built the best, most resilient, and high-availability system the cloud has ever seen. (Who hasn't heard somebody else's nightmare story of a job-gone-bad or a replication-gone-awry or a particularly gnarly malware that completely trashed both sides of a replicated database?)

But, backup is not the whole story of resilience. You also have to be able to recover. If backup is in addition to high-availability, you may be fine without DR, depending on the scope of risk you want to manage to and how you've designed your High-Availability (HA) deployment (locally and/or geographically disperse). However, if you're thinking that backup combined with DRaaS (as described above) is going to be your approach, then all those considerations for DRaaS still apply.

It is important to make sure the resources will be there when you need them and the technical environments will be consistent enough to not inhibit your recovery. In addition, if your strategy is to use your cloud backup for large scale data restoration, you will need to consider the proximity of your backup to your recovery resources, how much bandwidth it will require, and how long it will take. Will you be able to meet your recovery objectives if a full-scale restoration of all the data associated with even a relatively small system could take days? To meet your recovery time objectives you may need significant bandwidth and you may be forced to have at least partially provisioned infrastructure on standby, using periodic replication and/or continuous, incremental backup/restore.

In Part Four of the series, *The Resilience Continuum: Exploring the Impact of the Evolution of Cloud on Resilience Strategies*, we will explore how the evolution of cloud technologies will impact resilience planning.





PART FOUR

The Resilience Continuum: Exploring the Impact of the Evolution of Cloud on Resilience Strategies

This part of the series introduces the concept of the “Resilience Continuum” as a way to explain how changes in technology are impacting the way IT service consumers and providers think about, and design for resilience of cloud-deployed workloads. (Author’s note: this is not an official thing; just the name I came up with for the concepts I was attempting to describe.) As the technologies associated with cloud continue to advance, the roles and responsibilities related to designing and implementing resilience and recovery are changing.

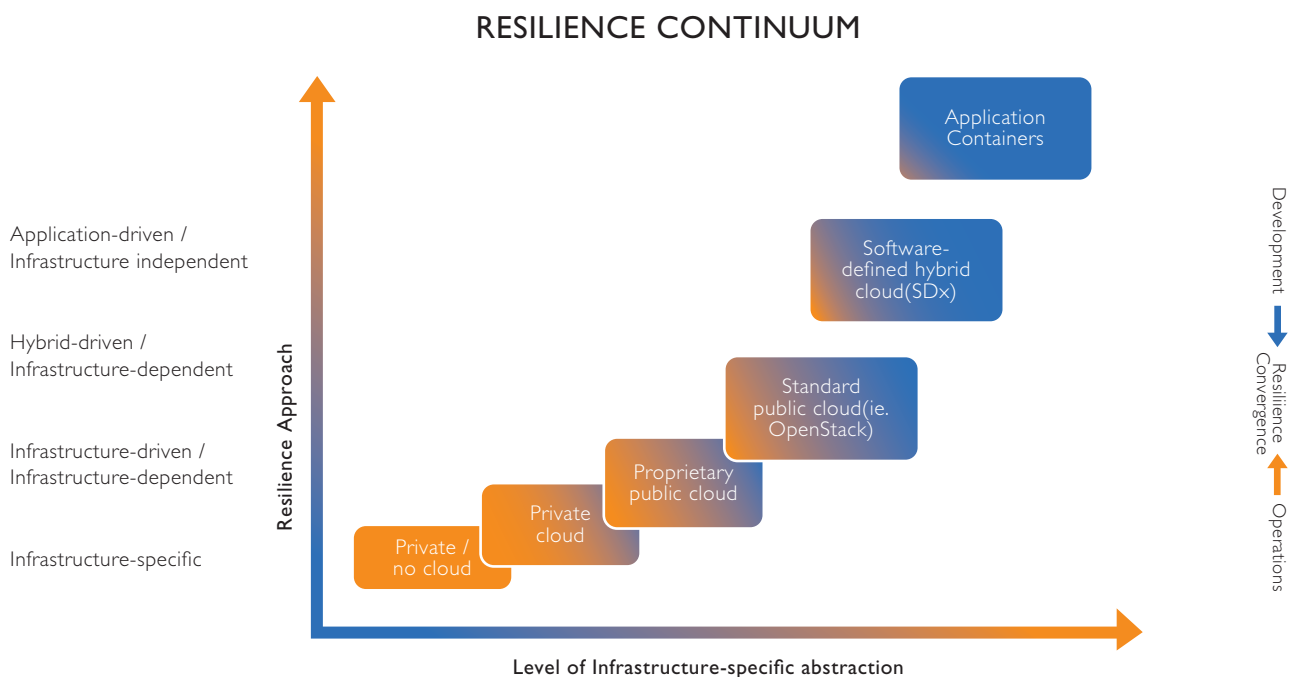
Resilience Continuum

At this point in the discussion, hopefully you’re sensing that “going to The Cloud” doesn’t mean you can leave the concerns about resilience and recovery in the past. With cloud (except maybe for SaaS) you will still require a fairly high level of infrastructure awareness, especially in the IaaS service model. While this is true, you should stay the course on your journey

to cloud. Cloud is a good thing. It will help your business grow in ways you hadn’t before thought possible, but it’ll take some time and planning to get it right. The idealistic view of things – that Cloud is Cloud and it’s all the same, like the electrical grid and indoor plumbing – is still out there a ways, but this is a really fast moving business. The future will be here before you know it.

Planning for the future starts with having a vision for where the technology is going. In the interest of getting ready for the future, consider the illustration below.

This diagram is intended to illustrate the evolution away from infrastructure dependence in how resilience strategies are implemented to a situation where the infrastructure is controllable by software - first in the form of API-controlled infrastructure, then to standardized interfaces (APIs) to the infrastructure, and over time in the form of standard and open application runtime environments.



In the “old” way of doing things, application developers created the business functionality and turned it over to the infrastructure team to ensure the technical service management capabilities are implemented, availability included. The infrastructure team would divide the applications into “tiers”, and then build infrastructure solutions to meet the needs of each tier. In those environments, we saw very specific dependencies emerge in terms of things like disk-based replication being tightly integrated with, and controlled by, the database engine, for example. How it all worked, and how well it worked was **very specific to the infrastructure**, often right down to the patch levels of the microcode in the hardware.

Virtualization enabled the removal some of the infrastructure dependencies, but it was still the role of the infrastructure team to make resilience and recovery work. Proprietary public clouds are really not that much different. Largely, it is still expected that the infrastructure team will use the facilities available to them to design the environment to meet the resilience needs of each application. Once designed and implemented, it remains pretty much statically defined.

APIs to the Cloud infrastructure offer the first glimpse into being able to control the infrastructure from the application to meet the performance, availability and security needs of the application. The ability to use an API to detect a lost service, then restart it, or deploy another environment automatically is common in The Cloud. With proprietary clouds (and proprietary APIs) you're pretty much limited to recover in that

particular cloud, but with pervasive or open APIs, the options are extended, and the **dependence on the infrastructure is lessened**. If we get to a place where all of the infrastructure is software defined and all the APIs are the same (or at least pervasive and well-known), then in theory we have complete infrastructure independence, and the vision of a cloud that manages its own resilience comes closer to reality.

Software defined infrastructure will open up the possibility of being able to deploy your applications pretty much anywhere (“in The Cloud” – public/private/hybrid), and then control the environment very dynamically to meet the needs of the application. But, the developers who write the applications have been focused almost exclusively on business functionality for the last 20+ years. The expertise for resiliency design has been the domain of the infrastructure team. Developers need to learn new interfaces (infrastructure APIs) and design resilience into the application. As containers (i.e. Docker) evolve to provide facilities for enterprise Qualities of Service and become integrated with tooling for Software Defined Infrastructure (OpenStack and VMware are already moving in this direction), the environment will at least be familiar to developers. But, **understanding how to apply these facilities will require developers to work closely with the infrastructure team, in the mode of “DevOps”**.

Part Five, DevOps and Resilience, will tie up the discussion with emphasis on the importance of including resilience in your collaborative development process.



PLANNING FOR THE FUTURE STARTS WITH HAVING A VISION FOR WHERE THE TECHNOLOGY IS GOING.



PART FIVE

DevOps and Resilience



This final part of the series summarizes the role of resilience in the DevOps life cycle – which prescribes a partnership between those who develop the systems and those who understand the implications of designing for resilience - and summarize the discussion overall.

While the basic premise of DevOps is all about getting new capability deployed faster, that's only part of the story. DevOps is the continuous transformation of the whole lifecycle to remove bottlenecks and inconsistencies, end-to-end, including Operations. DevOps implies (requires!) collaboration between Development and Operations so that new functionality is deployed both quickly and flawlessly, and that it meets the needs of the business the first time. Technology is helping lower the barriers between Application and Infrastructure, but it's going to require people collaborating across the IT supply chain to realize this vision. This highlights another aspect of the transformation needed to enable effective use of cloud. Those skilled resources who used to be applied in silos, including resilience expertise, are going to need to work together, continuously, with Design and Development, and Deployment teams to realize the full value of The Cloud. **The idea that resilience gets designed in after the application has been developed simply won't work.**

But, the good news is it should get easier, not harder, as the realization of Software Defined Infrastructure moves closer and system design transforms from monolithic systems with marked division between Application and Infrastructure, to "thinner", vertically-integrated services, as illustrated in the graphic below. While there will be more moving parts to coordinate, each component should be smaller, less complex, infrastructure-independent, and able to control its own resilience.

What about re-purposing non-production environments for Recovery?

As part of their DevOps transformation, many organizations have been using cloud for development and testing for a while, so there is practice in non-production environments being built and taken down and refreshed over and over. Some are thinking to repurpose parts of their development/test environment for recovery, thinking that would solve the problem of making sure the environments are in sync between production and non-production, that the automation would be

tested, and that incompatibilities would be ironed out.

Over the years many organizations have implemented a similar approach, way before cloud. This approach can work well, especially when the environment to be re-purposed is the one most closely aligned with production. Assuming you can afford to stop testing and deploying new functions while you're testing or executing recovery, and that the QA environment and process isn't critical to running the business, it can work. But, even this seemingly safe concept has its considerations.

If you've leveraged The Cloud for development and testing, including pre-production, you may need to consider whether your cloud environment can meet the same stringent security, audit, and compliance requirements as your production environment. If the whole system, all environments, is in The Cloud, did you design them to the same specifications, or did you scrimp a little in non/pre-production to save costs?

You should plan that whatever facilities you need in production you'll very likely need in your recovery environment. While you might be able to get by with reduced scale, you're not likely to get away with cutting back in areas like security and compliance. For example, if your non-production environment relies on data redaction or obfuscation to be considered "safe", is that same environment also ready to support production workloads, or are there other limitations that you overlooked because the data wasn't a risk? If you use encryption and/or dedicated infrastructure to support your database servers in production, for example, do you have those same facilities available in non-production if you were planning to use that environment for recovery?

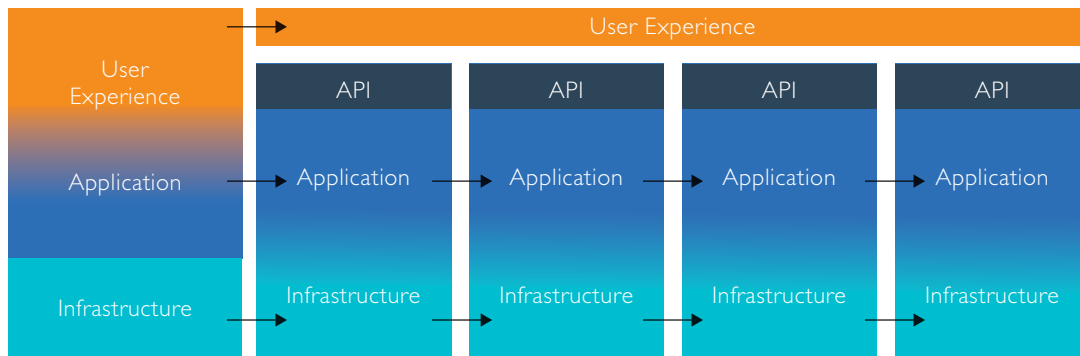
Further, if part of your non-production environment (i.e. QA) is temporarily unavailable because you're using it for recovery, can you really afford to stop testing and deploying updates? What if one of those updates is some critical patch or mandated fix, perhaps even the one that caused the outage? What are the risks of not being able to test it before you drop it into what would then be production? If you break the repurposed non-production environment that you are now using for recovery during some sort of an emergency, then what? Have you ever seen what happens when you put one emergency on top of another? Chaos ensues. Might be worth the risk, but know what you're getting into.

The great thing about Cloud is that you should be able to adjust things quickly and easily, and with stipulations sited

previously, get additional resources when you need them to implement your resilience strategy. With planning and **automation** you should be able to create fast and efficient recovery in The Cloud. Many do. Automation is one of the key

ingredients in DevOps and in taking advantage of Cloud. Implementing resilience is no different. **Do you have a spot on your roster for automation specialists? You're definitely going to need them.**

EVOLUTION OF APPLICATIONS, FROM MONOLITHS TO INTEGRATED SERVICES



There's an old saying - "There's no such thing as a free lunch". This diagram also illustrates an important consideration that arises in recovery and/or resilience testing in this new application paradigm – how do you perform complete testing of one of the vertically integrated services if all the other slices that make up the 'system' of the user experience aren't available for testing at the same time? It defeats the agility purpose of this application design pattern if you have to lump the whole group together for testing.

Different strategies are evolving to address this challenge. The age old technique of writing code to "stub out" the interfaces is one way. Another approach is to have always-on versions of each service for testing against, so components can be rotated in/out of test cycles. This requires a high level of coordination to make sure your tests aren't colliding. Perhaps the most scalable approach is to implement a test harness with service virtualization to provide a ready, but virtualized, test environment to drop your application components into when needed. Whichever approach, you will likely need to have a plan as your applications move to this new, cloud-enabled, model.

Conclusion

Cloud and resilience do go together, but not without purpose and a plan. Designing for resilience requires specific skills and collaboration between Development and Operations. While today the common approaches are still very-much infrastructure dependent, over time the approach will become

less infrastructure dependent and more application driven, as the divide between application and infrastructure is blurred. As you begin to leverage cloud for more workloads, including those that require high levels of resilience, consider engaging experienced resources to help guide your journey.

About the Author



Derek is a Principal Consultant and Enterprise Architect on the Transformation Services team at Wipro, in the Global Infrastructure Services practice. In that role he works to develop and deliver the next generation of technology enabled solutions for Wipro's customers who


want to transform the dynamics of the value of Information Technology for greater efficiencies and strategic advantage.

Derek has 31 years of experience in information technology, supporting customers and engagements across many industries. His areas of expertise include e-Business solutions design, ERP, I/T architecture and governance, Cloud, Smarter Commerce, Analytics, and a thorough understanding of the application of best practices in helping clients solve business and technology problems. Derek has worked with clients in many industries, including Distribution, Retail, Travel and Transportation, Healthcare, and Professional Services.

About GIS

Global Infrastructure Services (GIS), a unit of Wipro Limited, is an end to end IT infrastructure & outsourcing services provider to global customers across 57 countries. Its suite of Technology Infrastructure services spanning Data Center, End User Computing, Networks, Managed Services, Business Advisory and Global System Integration. Wipro, is a pioneer in Infrastructure Management services and is amongst the fastest-growing providers across the world. GIS enables customers to do business better by enabling innovation via

standardization and automation, so that businesses can be more agile & scalable, so that they can _nd growth and succeed in their global business. Backed by our strong network of Integrated ServiceNXT™ Operation Centers and 11 owned data centres spread across US, Europe and APAC, this unit serves more than 500+ clients across with a global team of 23,800 professionals and contributes to over 30% of Wipro's IT Services revenues of Wipro Limited.



About Wipro Limited

Wipro Ltd. (NYSE:WIT) is a leading Information Technology, Consulting and Business Process Services company that delivers solutions to enable its clients do business better. Wipro delivers winning business outcomes through its deep industry experience and a 360 degree view of "Business through Technology" - helping clients create successful and adaptive businesses. A company recognized globally for its comprehensive portfolio of services, a practitioner's approach to delivering innovation, and an organization wide commitment to sustainability, Wipro has a workforce of over 150,000, serving clients in 175+ cities across 6 continents.

For more information, please visit www.wipro.com

DO BUSINESS BETTER

WWW.WIPRO.COM

CONSULTING | SYSTEM INTEGRATION | BUSINESS PROCESS SERVICES

WIPRO LIMITED, DODDAKANNELLI, SARJAPUR ROAD, BANGALORE - 560 035, INDIA TEL : +91 (80) 2844 0011, FAX : +91 (80) 2844 0256, email : info@wipro.com

© WIPRO LIMITED 2015

"No part of this booklet may be reproduced in any form by any electronic or mechanical means (including photocopying, recording and printing) without permission in writing from the publisher, except for reading and browsing via the world wide web. Users are not permitted to mount this booklet on any network server."