Resilience –The Modern Uptime Trinity

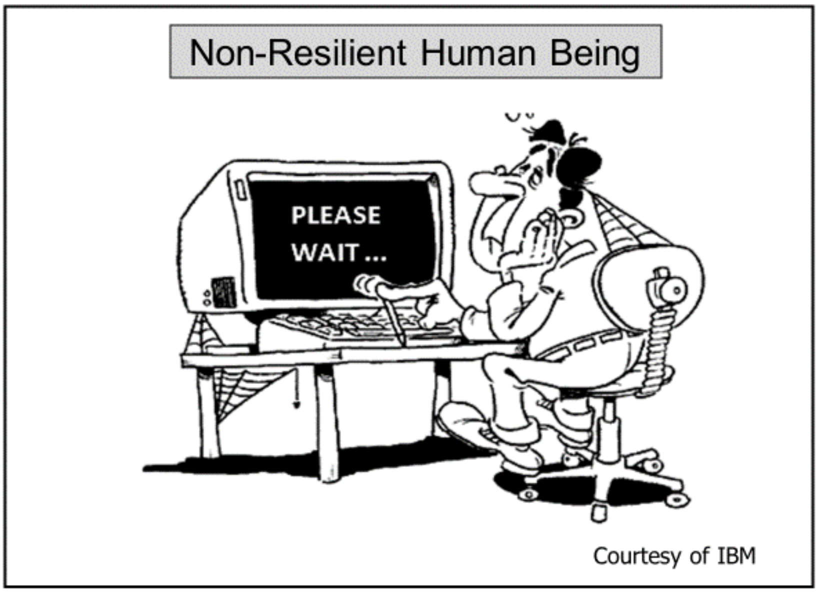

Figure 1: The Poor End User

Introduction

Some years ago, the computer systems' key focus was on PERFORMANCE and many articles, products and efforts were focussed on this area. A few years later, the emphasis moved to HIGH AVAILABILITY (HA) of hardware and software and all the other machinations they entail. Today the focus is on (CYBER)SECURITY. The first two items are of course still important, a fact many forget, but the admixture of them depends on the business in question and the consequences to them of failure.

A lot of water has flowed under the IT technology bridges since then and to a certain extent these foci have merged to form a RESILIENCE spectrum. For example, marketing web sites need good performance since various studies show that poor response times may cause users to switch sites, temporarily or permanently. This is not good for business as you may guess. At the other end of the spectrum, military and many financial and healthcare sites put most emphasis on security, in its widest sense, while the other factors are ancillary. In short, the whole thing is organisation and requirements dependent.

Unfortunately, these key areas are often subsumed in the rush to use 'shiny' new products and ideas such as AI (artificial intelligence), ML (machine learning), big data, mobile access, IoT (Internet of Things devices) and use their functions. Just as unfortunate is the habit of organisations taking their eyes off the 'trinity' ball as described, as several UK banks and the UK NHS will (albeit unwillingly) testify.

The UK NHS (National Health Service) is very keen on these new areas but in my view they fail to look for the possible *gotchas (downsides)* which lurk, iceberg-like beneath them, the 'Titanic Syndrome'. It mirrors Shakespeare's; '*There are more things in heaven and earth than are dreamt of in your philosophy Horatio.*'

It also chimes with my view that you cannot solve a problem fully until you get a full statement of what it is. That is the hard part, otherwise you can get a solution to which there is no problem.

Today

These discrete environments' boundaries have now blurred with the original components above being different sides of same coin. The main components of resilience are:

1. Normal high availability (HA) design, redundancy etc. plus normal recovery from non-critical outages. This applies to hardware and software. Human factors ('fat finger' syndrome and deliberate malice), are extremely common causes of failure.
2. Cybersecurity breaches of all kinds. No hard system failures here but leaving a compromised system online is dangerous. This area has spawned the phrase *cybersecurity resilience*.
3. Disaster Recovery (DR), a discipline not in evidence in May 2017 when *Wannacry* struck the UK NHS (National Health Service). Did they have a tested DR plan? This area of resilience usually sorts out the men form the boys in terms of having a demonstrable, gold-plated DR plan and setup

An Analogy:

In boxing, *resilience* in simple terms means the ability to recover from a punch (normal recovery) or knock down (disaster recovery). However, it has connotations beyond just that, inasmuch as the boxer must prepare himself via tough training, a fight plan and coaching to avoid the knockdown and, should it happen, he should be fit enough to recover and re-join the fray quickly enough to beat the 10 second count[1].

Moral: You can't choose which of the three bases you cover; it's all or nothing and in the 'any-2-from-3' choice, disaster beckons. It would be like trying to build then sit on a two-legged stool. Ask any bank, particularly those in the UK.

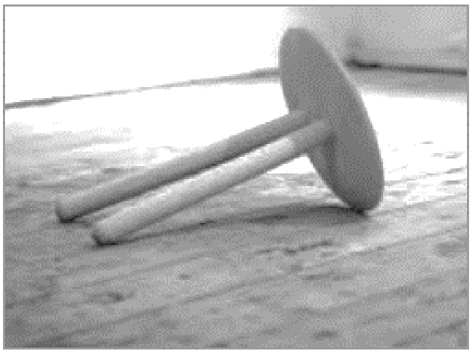

Figure 2: 'Two from Three' Resilience Architecture

UK Finance View[2]

The *UK Finance* organisation defines resilience as:

'A resilient financial system is one that can absorb shocks rather than contribute to them. The financial sector needs an approach to operational risk management that includes preventative measures and the capabilities – in terms of people, processes and organisational culture – to adapt and recover when things go wrong. As recent

[1] Financial penalties in our world.

[2] An organisation which represents and supports over 300 UK financial bodies, ranging in size from small to very large.

high-profile disruptive events have shown, the speed and effectiveness of communications with the people most affected, including customers, is an important part of any firm's or FMI's overall response to an operational disruption.'

Building the UK financial sector's operational resilience https://www.bankofengland.co.uk/-/media/boe/files/prudential-regulation/discussion-paper/2018/dp118.pdf

Gartner Group defines resilience in a broader way, capable of different interpretations, thus:

"Operational *resilience* is a set of techniques that allow people, processes and informational systems to adapt to changing patterns. It is the ability to alter *operations* in the face of changing business conditions.[3]"

Operational Resilience - Gartner IT Glossary https://www.gartner.com/en/information-technology/glossary/operational-resilience

To summarise these views, the key job in hand is '*to keep the show on the road with the same quality of delivery*', in the manner the customers want. You interpret this principle according to your business and clients.

When is an Outage Not an Outage?

This sounds like a Christmas Cracker joke paper but it is a valid question to ask if you understand service level agreements (SLAs). SLAs specify what properties the service should offer aside from a 'system availability clause'. These requirements usually include response times, hours of service schedule (not the same as availability) at various points in the calendar, for example, high volume activity periods such as major holidays, product promotions, year-end processing and so on.

Many people think of a system outage as complete failure – a knockout using our fistic analogy. In reality, a system not performing as expected and defined in a Service Level Agreement (SLA) will often lead users to consider the system as 'down' since it is not doing what it is supposed to do and impedes their work.

This leads to the concept of a *logical outage* (a forced standing count in boxing) where physically everything is in working order but the service provided is not acceptable for some reason. These reasons vary, depending at what stage the applications have reached.

- At initiation of the new service, the user interface (UI) is totally unfriendly and foreign to the application users. This if often the result of a UI designed without user input or knowledge of the business process behind it.
- At initiation, it does not completely map the function of the business function it deals with onto the delivered application, again often the result of user non-involvement.
- During normal running, the performance degrades for some reason, resulting in effects such as loss or productivity through to being totally unusable; loss of performance can have a disastrous effect on web-based sales systems for example.

  Studies have shown that people using a web site to order goods have a mental, often unquantified expectation of interactive response times which, if exceeded, results in their leaving the site. In the worst case, it can mean the buyer obtaining the goods from a competitive site and, worse still never returning.

[3] In boxing parlance change of fight tactics.

The unconscious response time expectation of buyers varies from a few hundred milliseconds to a few seconds, depending on the study involved. Incidentally, poor design of the user interface is another *logical outage* but this and other aspects of outages are too detailed (and gory) to cover here.

Resilience Areas

Resilience in bare terms means the ability to recover from a knock down, to use the boxing analogy once more. However, it has connotations beyond just that inasmuch as the boxer must prepare himself by tough training and coaching to avoid the knockdown and, should it happen, he should be fit enough to recover, get to his feet and continue fighting. Referring to our fistic analogy, the information technology (IT) scenario this involves, amongst other things:

- 'Fitness' through rigorous system design, implementation and monitoring plus staff training in risk and crisis management
- Normal backup and recovery after outages or data loss not due to criminal activities or severe system damage.
- Cybersecurity tools and techniques to counter attacks by malware, both external and internal. In my view, the current internet architecture, intrinsically open, is unsuitable to stem the tide of malware, particularly in the rapidly growing areas of mobile computing and the internet of things (IoT) devices.
  The latter is expected to number 22 bn. by 2021, across many industries, and, unfortunately, security was not the primary feature in their design and manufacture and still isn't in many cases. Security is notoriously difficult to retrofit.
- Disaster Recovery (DR) when the primary system(s) is totally unable to function for whatever reason and workload must be located and accessed from facilities – system and accommodation (often forgotten) - elsewhere.
  The location of a secondary DR site is important since some natural disasters affect large areas and an on-site or near-site secondary system can be rendered useless by the same disaster that struck the primary. On the other hand, cybersecurity attacks are geography neutral.

I WAS JUST WONDERING, DO WE HAVE A DISASTER RECOVERY PLAN?
IT DEPT
© D.Fletcher for CloudTweaks.com

Figure 3: The Need for a Disaster Recovery (DR) Plan: *Courtesy Cloudtweaks*

These factors involve experience, planning and common sense. I remember the case of the UK data centre plan where the location was decided by AI-like software which decided the optimum location for the centre was in the middle of the Bristol Channel, although now I think Microsoft are planning on under the sea. Perhaps this ‘AI’ was cleverer than we thought.

- Spanning the resilience ecosphere are the monitoring, management and analysis methods to turn data into information to support the resilience aims of a company and improve it. If you can’t measure it, you can’t manage it.

Figure 4 is a simple representation of resilience and the main thing to remember is that it is not a pick and choose exercise; you have to do them all to close the loop between the three contributing areas of resilience planning and recovery activities.

In view of the series of failures of financial and other service’s computer systems over the past few years, it is evident that *Performance* and *HA lessons* learned (if indeed they ever were) have largely been forgotten. *Security* (cybersecurity) is a new threat which the business world has to be aware of and take action on, not following the Mark Twain dictum; ‘*Everybody is talking about the weather, nobody is doing anything about it’*.

The key factor is covering all the ‘resilience’ bases at a level matching the business’s needs. It is not a ‘*chose any n from M*’ menu type of choice; it is all or nothing for optimum resilience. Failures make regulatory bodies and stakeholders angry and ‘you wouldn’t like them when they are angry’.

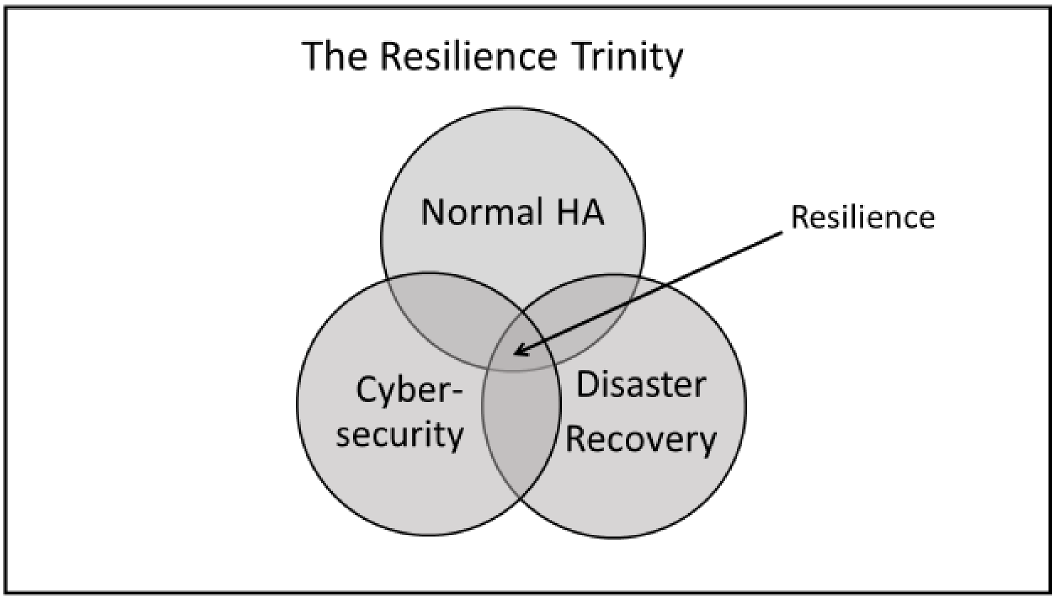

Figure 4: Resilience Components

In view of the series of failures of financial and other service’s computer systems over the past few years, it is evident that *Performance* and *HA lessons* learned (if indeed they ever were) have largely been forgotten. *Security* (cybersecurity) is a new threat which the business world has to be aware of and take action on, not following the Mark Twain dictum; ‘*Everybody is talking about the weather, nobody is doing anything about it’*.

The key factor is covering all the ‘resilience’ bases at a level matching the business’s needs. It is not a ‘*chose any n from M*’ menu type of choice; it is all or nothing for optimum resilience, otherwise you have the *2-legged stool*. Failures make regulatory bodies and stakeholders angry and ‘you wouldn’t like them when they are angry’.

What Needs to be Resilient?

A good question to which the answer is probably '*anything which can stop or hinder the show*'. Here are several of them:

- Server and peripheral devices, especially storage devices (*)
- Software – application and system (*)
- Physical facilities failure – rooms, power supplies, sprinklers etc.
- Networks and its 'peripherals – mobile device, IoT etc. (*)
- People – numbers, skills, finger trouble, commitment (*)
- Backup and recovery methods and systems (resilience management) (*)
- DR facilities (complete primary site failure) (*)
- (*) Security facilities across the entities above marked (*) and including encryption at rest and in flight.

Lussers' Law[4] shows that entities operating in parallel (unison) give higher availability than those in series.

To stretch a point a little, I think that resilience will be enhanced by recognising the 'trinity' aspect of the factors affecting resilience and should operate as such, even in virtual team mode across the individual teams involved. This needs some thought but a 'war room' mentality might be appropriate.

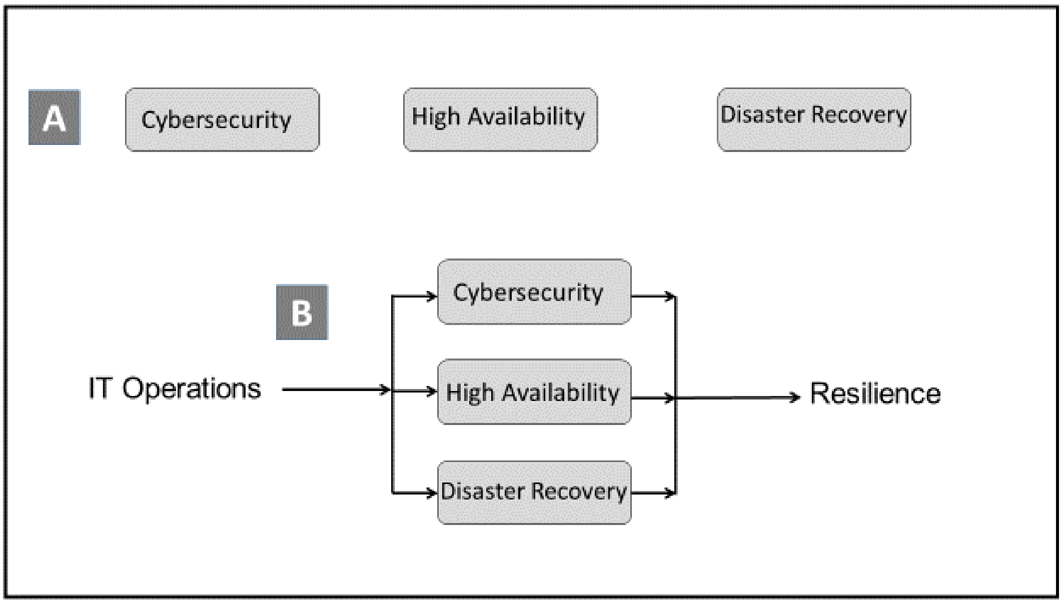

Figure 5: Resilience *a la* Lusser's Law

The three areas considered in *parallel* (**B**) make for a more resilient systems than different teams treating them in isolation as *serial* or siloed activities (**A**). Another downside of A is the three sets of change management activities - a possible 'triple whammy'.

Tools & Methods

'There are scores of methods and technologies to 'harden' areas where vulnerabilities lurk before, during and beyond the three areas shown in Figure 2. Some resilience topics are to be found in the links below:

[4] See https://en.wikipedia.org/wiki/Lusser%27s_law for the simple mathematics of this.

HA Design

There are many ways of 'hardening' systems to make them more resilient but they are to detailed to cover here. Nonetheless, a few ideas and sources are outlined in what follows.

High Availability Spectrum; Factors
http://www.availabilitydigest.com/private/0201/availability_best_practices.pdf

The other parts of the triumvirate, disaster recovery and cybersecurity, depend on an organisation's requirements and is essentially a 'knitting exercise'.

Cybersecurity

There are a great number of ways to protect one's data, am attack on which can cause and outage or make the owner take it down for safety. These methods are too numerous to outline here but I'll mention one – encryption – which according to IBM represents a key safety area. Only 2.8% of attacks are successful against systems with encrypted data.

Encryption Article
http://www.ibmsystemsmagmainframedigital.com/nxtbooks/ibmsystemsmag/2020mfse/index.php#/34

Disaster Recovery

Again, a large area to cover but once more a key message. How far one goes in configuring disaster recovery site depends on how many applications are important enough to require elaborate systems. This is decided by a failure business impact analysis (BIA) between business and IT, part of BCP (business continuity planning).

Conclusion

Review

Like any major activity, the results of any resilience plan need review and corrective action take. This requires an environment where parameters relating to resilience are measurable, recorded, reviewed and acted upon; it is not simply a monitoring activity since monitoring is *passive*, management is *active* and *proactive*.

Management = Monitoring + Analysis + Review + Action

Why Resilience is IT's Responsibility

I find that vendors and some organisations often split these areas under different departments or functions, especially cybersecurity or HA. Unfortunately, systems cannot live by any of these alone but only in synergistic development and operation.

They need to be considered together, especially as there will be areas where changes to one of these entities can adversely affect another, thereby *mandating change management* as a corollary discipline[5].

Resilience, like a bespoke garment, is a personal thing and it is unlikely that the vendor of tools will know more than you about your business and its vulnerabilities. If you think that tools will do all the heavy lifting for you, you are deluding yourself.

Resilience is Hard

Your impression of resilience may now be; 'That is a big job'. Yes, it is, but a big problem usually involves hard work to solve or at least mitigate it. If you think it is expensive to ensure resilience, try frequent failure instead and consider the penalties for failure levied by regulatory bodies, particularly in finance.

[5] This should be in place anyway in a well-run site.

In computer *security*, a vulnerability is a weakness which can be exploited by a threat actor, such as an attacker, to perform unauthorized actions within a computer system. However, in our quest for resilience, of which vulnerability is its antithesis, we need to expand this definition to cover the other two elements of resilience.

In the case of *HA,* the vulnerabilities include the inherent failure rate of the hardware and software components, their configuration (series if parallel), the environment and many other factors, including human error and vindictiveness.

In the *DR world*, the vulnerabilities in primary systems again are many and in theory apply to the DR site as well but the probability of both sites suffering catastrophic outage simultaneously are vanishingly small, assuming the DR location and other factors are chosen well.

For example, if you chose to site both on a geological fault line you are asking for trouble; similarly in flood or hurricane areas.

To ensure optimum resilience of systems is not a 'one solution fits all' situation but a rather time-consuming search for vulnerabilities which can affect your systems(s) in each area nor is it 'let us give this a try and see how it goes' – the 'burn your boats' technique if it fails. No, it needs a **CFIA** (component failure impact analysis) where each impact is traced to which business aspects are compromised or fail if a particular component ceases to function.

If you think that throwing suitable, trendy products at the resilience design is the answer, you are deluding yourself. As Sir Winston Churchill said, in paraphrase; '*All I can offer is blood, sweat and tears.*'

End Note

NIST Cyber Resilience Definition

'Cyber resiliency is defined as "the ability to anticipate, withstand, recover from, and adapt to adverse conditions, stresses, attacks, or compromises on systems that use or are enabled by cyber resources." ……… Systems with this property are characterized by having security measures or safeguards "built in" as a foundational part of the architecture and design and, moreover, can withstand cyber-attacks, faults, and failures and continue to operate even in a degraded or debilitated state, carrying out mission-essential functions and ensuring that the other aspects of trustworthiness (in particular, safety and information security) are preserved.'

[*Draft NIST Special Publication 800-160 Volume 2*]

Developing Cyber Resilient Systems: A Systems Security Engineering Approach
https://csrc.nist.gov/publications/detail/sp/800-160/vol-2/final

The Author

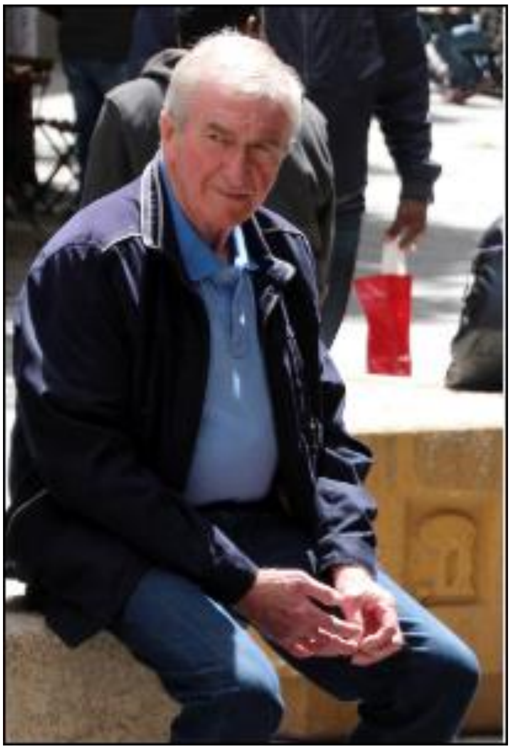

Meditating on Resilience in Madrid

Dr. Terry Critchley is a retired IT consultant living near Manchester in the UK. He joined IBM as a Systems Engineer after gaining a Ph D and spent 24 years there in a variety of accounts and specialisations, latterly joining Oracle for 3 years. He joined his last company, Sun Microsystems in 1996 and left there in 2002 and then spent a year at a major UK bank. He is now an author of numerous IT articles and books.

Publications

Modern IT Concepts and Technology: An IT Study Guide for Beginners and Practitioners Kindle Edition [2019]
https://www.amazon.co.uk/dp/B0826XN9L2/ref=sr_1_1

Open Systems: The Reality
BCS Practitioner Series ISBN 0-13-030735-1

High Availability IT Services
https://www.crcpress.com/High-Availability-IT-Services/Critchley/9781482255904

High Performance IT Services
https://www.crcpress.com/High-Performance-IT-Services/Critchley/9781498769198

Making It in IT
https://www.crcpress.com/Making-It-in-IT/Critchley/9781498782760