the Availability Digest

www.availabilitydigest.com
@availabilitydig

The High Availability Design Spectrum – Part 3

Dr. Terry Critchley
February 2017

[Editor's Note: In his book "High Availability IT Services," Dr. Terry Critchley lists twenty-three areas that can have an impact on the availability of IT business services. In this multipart series, and with his permission, we publish his observations. In Part 1 of this series, we reviewed his first four reflections - his Parts A through D. In Part 2, we examined his next nine considerations – his parts E through M. In this Part 3, we publish his six observations N through S]

Dr. Terry Critchley: Most of the documentation on HA/DR I have come across majors on hardware, mainly redundant or fault-tolerant, and to some extent software. My thesis is that the spectrum of activity needed to design, implement and maintain a high availability business IT system and recover from failures small and large (DR) is much, much greater. Below, I have listed 23 areas (A to W) that can have an impact on the availability of business services that are IT-based. I am sure it will be evident that these areas can have a significant impact on the availability and non-availability of any service or system.

Remember, focusing on availability and focusing on avoidance of non-availability are not the same thing if you think about it.

The book and chapter references following refer to 'High Availability IT Services': https://www.crcpress.com/High-Availability-IT-Services/Critchley/9781482255904.

N. Availability by Partnerships

It is usually the case when systems experience faults or crash that there will be more than one vendor who may need to help in problem determination and resolution. If these vendors do not cooperate, there is often 'finger pointing' or 'not our product' activity. Cooperating vendors can aid in designing high availability systems and reduce problem determination and recovery time in an installation. They should also be on hand to expedite recovery from failures involving their products.

It may also be beneficial to consult other organizations on IT matters like HA/DR where there is a common interest (and no business competition).

There is synergy and cross fertilization of ideas in this sort of cooperation, but your company may have to initiate it.

O. Availability by Change Management[1]

'*Bank loses data'*. This was all over the UK newspapers, radio and television in June 2012 when a major bank had an outage that blocked updates, transfers and withdrawals from the bank. It transpired that there had been a *scheduling software upgrade* which resulted in the non-availability issue. The outage, which had wide-ranging effects, lasted several days and, although the bank said it would underwrite any consequential loss, it was a case of 'egg all over their faces'.

.. and lo and behold, there was another outage at the same bank in early March 2013, this time attributed to *hardware*. Although the fix apparently took effect after 3 hours, there were ripple effects on other institutions connected with the bank as well as the bank itself.

Of course, the hardware was probably fine in the first instance, but this was a typical 'logical outage' that hardware at 99.999% availability would not have solved – only cost the customer a lot of money. As it was later reported, the 'caper' cost the bank £125m and an unknown amount in customer defections.

Change Management is mentioned elsewhere in this document, but suffice it to say that one of the prime considerations in changes is that they should be capable of being *'backed out'* if they fail. This doesn't mean that any outages caused by the original change will not be serious. 'Risk-laden' changes ought to be pre-tested if possible [2].

Though it is known that a piece of scheduling software was involved, it is not known who 'goofed' - the software or operations liveware.

P. Availability by Performance/Capacity Management

We have seen that performance issues can masquerade as outages when user work is seriously interrupted or degraded by poor responses. These topics are covered in Chapter 7, but suffice it to say what the fundamental differences between performance issues and outages are as follows.

Performance Management is a proactive *operational* exercise, mainly concerned with the 'here and now' of resource consumption, response and turnaround times and the provision of data for Capacity Management. On the other hand, Performance Monitoring is more of a passive exercise whereas Capacity Management is the prediction of, and planning for, increased resource utilization in systems with a view to proactive solutions. It is *tactical* and, in some cases, *strategic*.

Capacity management has two major legs:

- The use of performance data and its intelligent extrapolation based on business volumes and their resource requirements.

- The estimation of resource requirements for new business applications and services, often using operational data from similar workloads already running. It is important to understand the difference between a *business transaction* and the subsequent multiple *system* transactions when assessing resource requirements.

The two disciplines, Performance and Capacity, should not be confused in their objectives, methodology, conclusions or output. One is tactical and operational, the other planning and strategic.

---

[1] Excellent document and checklist set on Change Management at**;** http://www.cisco.com/warp/public/126/chmgmt.*pdf*

[2] Amazon web site's 'caper' in October 2012

Q. Availability by Monitoring

It is now fairly common to have a pseudo-client accessing services, such as web sites and business applications. The 'client' acts as a real user submitting transactions or reading data from a live system. It is plain to see that if this 'client' is to monitor availability as well as other items, it must survive outages of the system being monitored. Common sense dictates that this 'client' be duplicated. For example, one client should be located at the center and one located in an end user location or locations. A complete center outage can be compensated for by the continuing activity of a remote 'client'. This could not happen if a component and its redundant component are side by side in a flooded room.

An adjunct to monitoring availability will be tools for measuring KPIs (key performance indicators), for example, for Performance and Capacity Management. They will assist in situations where the users consider a system 'down' based on its performance characteristics.

See Chapter 7, alongside 'SLAs, Management and Methods'.

R. Availability by Cleanliness

The machine room should resemble a hospital operating theatre (the better ones anyway) in its cleanliness. Coffee, hamburgers and sundry other comestibles have no place in a machine room for fear of contamination. I have seen a console keyboard ruined when coffee was spilt on it and the miscreant tried to dry it with a hair dryer - the keyboard just buckled and died. The system it was controlling wasn't pleased, and the end users didn't think much of it either. There was no duplicate console facility.

Anything exposed that carries data, such as a tape reel, should not be handled with dirty hands. Another possible *gotcha* can be introduced by extraneous electronic equipment in the machine room, since its radiations can be picked up by the circuitry in other systems, possibly causing bits to change at random.

S. Availability by Anticipation

One of the benefits of experience is that a person can judge where problems might arise and concentrate more effort in those areas. Where feasible, such experience ought to be enshrined in Operation Runbooks so that less experienced personnel can learn. For example, if operations staff have found by experience that when A happens, F and G are likely culprits and should be checked out first, it provides valuable insight into problems and should be recorded in Operation Runbooks. This allows other less experienced operators to minimize problem determination time and hence recovery time.

In the future, this might be done by software which is capable of learning and making judgements. In the meantime though, it is a human's task to learn from experience and pass this learning on.

Predictive Maintenance

Predictive maintenance, as the name implies, is the maintenance or replacement of parts of systems based on certain criteria which might indicate imminent or near-future failure. The technique is very often applied to mechanical systems such as oil rigs, large machines and so on; but it is a little harder to fit into an IT environment. It is still worth investigating for applicability to your IT environment, an example being the detection, logging and analysis of 'soft' errors in system parts. An IBM Redbook gives an overview of the topic in the first chapter of:

http://www.redbooks.ibm.com/redpapers/pdfs/redp5035.pdf

There is also an Aberdeen Report ' *Asset Management: Using Analytics to Drive Predictive Maintenance.*' Although mainly applicable to manufacturing, this report has uses in *'telecoms and IT management'*, according to the report abstract:

http://aberdeen.com/Aberdeen-Library/8380/AI-predictive-asset-analytics.aspx.aspx

- It should now be glaringly obvious that there are many bases to be covered in the search for high availability, and I have outlined quite a few here. You may think of others. If you do, elaborate on them and tell others (not your competitors though).