Halving Downtime Costs for Mission Critical Apps

by David Floyer | Monday, August 2nd, 2017

https://wikibon.com/halving-mission-critical-downtime-costs/

Premise

The purpose of IT backup and recovery systems is to avoid data loss and recover quickly, thereby minimizing downtime costs. Traditional storage-centric data protection architectures such as Purpose Built Backup Appliances (PBBAs), and the conventional backup and restore processing supporting them, are prone to failure on recovery. This is because the processes, both automated and manual, are too numerous, too complex, and too difficult to test adequately. In turn this leads to unacceptable levels of failure for today's mission critical applications, and a poor foundation for digital transformation initiatives.

Governments are taking notice. Heightened regulatory compliance requirements have implications for data recovery processes and are an unwelcome but timely catalyst for companies to get their recovery houses in order. Onerous malware, such as ransomware and other cyber attacks increase the imperative for organizations to have highly granular recovery mechanisms in place that allow data restoration to the point of infection with as little data loss as possible. Meanwhile, the cost of IT downtime for an average Fortune 1000 enterprise is already high and increasing as digital transformation efforts pressure organizations to deliver always on services.

This perfect storm of heightened regulation, advanced cyber threats and greater business risk creates a mandate for organizations to rethink data recovery for high value applications. The core premise of this research is that deploying an application-led backup and recovery architecture will lead to far less complexity, much easier process improvement over time, and has the potential to **cut the cost of downtime in half.**

Executive Summary

Digital Transformation and an Evolving Regulatory Climate

Virtually every business-oriented conversation Wikibon has with senior IT leaders includes a discussion of digital transformation. The relevance to this research is the increasing interdependency between organizations and the data they use, create, access, share and store. Digital means data and lots of it; and this data must be protected. The data requirements facing organizations today as a direct result of digital initiatives are unprecedented and require new approaches for protecting and enabling recovery for high value data assets.

Governments around the world are trying to keep pace with the digital tsunami and new/evolving regulations will further pressure data protection and recovery requirements. For example, the EU's General Data Protection Regulation (GDPR) states that organizations must take appropriate measures to ensure the ability to restore personal data "in a timely manner in the event of a physical or technical incident." Beginning in May of 2018, penalties for non-compliance to GDPR will be the greater of 4% of turnover or 20M euros.

In September 2016, The U.S. Department of the Treasury's Office of the Comptroller of the Currency issued enforceable guidelines that mandate recovery planning, including addressing threats from destructive cyber attacks.

In January 2016, the Basel Committee on Banking Supervision revised its requirements around market risk which go into effect in 2019. The Fundamental Review of the Trading Book imposes new guidelines that we believe will require major enhancements to data management systems and process methodology enhancements, including those associated with data recovery.

Increasing Complexity and the Escalating Costs of Downtime

The expected financial impact of downtime from IT is about 8% of enterprise revenue for Fortune 1000 enterprises. It is set to grow as digital transformation increases dependence on IT. The biggest downtime exposure is from mission critical systems, which comprise about 16% of all systems, but over 50% of downtime impact (5% of enterprise revenue).

This Wikibon research focuses on the backup and recovery of mission critical systems. It concludes that the overall failure rate for backup and recovery is driven by two components:

1. The complexity of the process (e.g., the number of steps and handoffs between processes and cross-functional groups required to protect data).

2. The probability of correct completion for each step — i.e. the real-world reliability of each step and communication with the step before and after. (See the section "The High Cost of System Downtime" below for more discussion of this complex subject).

Figure 2b below, in the "Key Contributors to Application Downtime" section, illustrates the relationship between complexity and reliability, and the probability of a recovery attempt failing.

Wikibon's research shows that most backup and recovery implementations depend on a mixture of different purchased software packages and hardware platforms, in-house software (including the infamous homegrown scripts), internal documentation and spreadsheets. This complexity requires high levels of expertise and collaboration across database administrators (DBAs), storage specialists and backup/recovery specialists. The result is a large number of steps, held together with error-prone internal software, scripts that nobody dares touch, manual processes, expertise and "tribal knowledge."

This Wikibon research compares Oracle Database backup and recovery using two different backup and recovery architectures for mission critical systems as follows:

1. Traditional installed storage-based backup and restore with Purpose-built Backup Appliances (PBBAs) as the reference architecture. Software and hardware is typically from multiple vendors, and includes native backup tools such as Oracle Recovery Manager (RMAN), backup software, and a backup appliance.
2. An Application-led end-to-end solution architecture using Oracle ZDLRA (RA) as the reference architecture. The database under the application is sending all the data necessary for recovery within the user-defined protection policies to the RA. All the software from production database to database backup and recovery is managed and developed by a single vendor. A formal definition of application-led is in the "Defining Application-Led & Storage-based Architectures" section below. A six month implementation schedule is assumed.

The red column in Figure 1 (storage-centric approach – e.g. PBBAs) shows the 4-year cumulative cost of downtime for mission critical applications, which totals $3.7 Billion ($735 million/year, which is 4.9% of revenue). The blue column shows the 4-year cumulative downtime costs after deploying an application-led architecture. At the end of the period, the annual downtime cost is reduced to $352 million/year, which is 2.3% of revenue. That is less than half of today's figure.

The reason for this dramatic performance is seen in Figure 2b below. It shows the result of implementing an application-led architecture for specifically reducing complexity—the number of steps—from 80 in the current mission critical environment (Year 0) to 14 in year 4. The probability of success for each of the steps is 99.6% in Year 0, improving to 99.9% in year 4.

The net present value (NPV) of the cumulative 4-year project is $1.04 Billion. Year 1 includes the 6-month implementation cost and implementation time. The IRR (Internal rate of return) is 591%, and the breakeven 6 months. For the right applications, this is a no-brainer of a project.

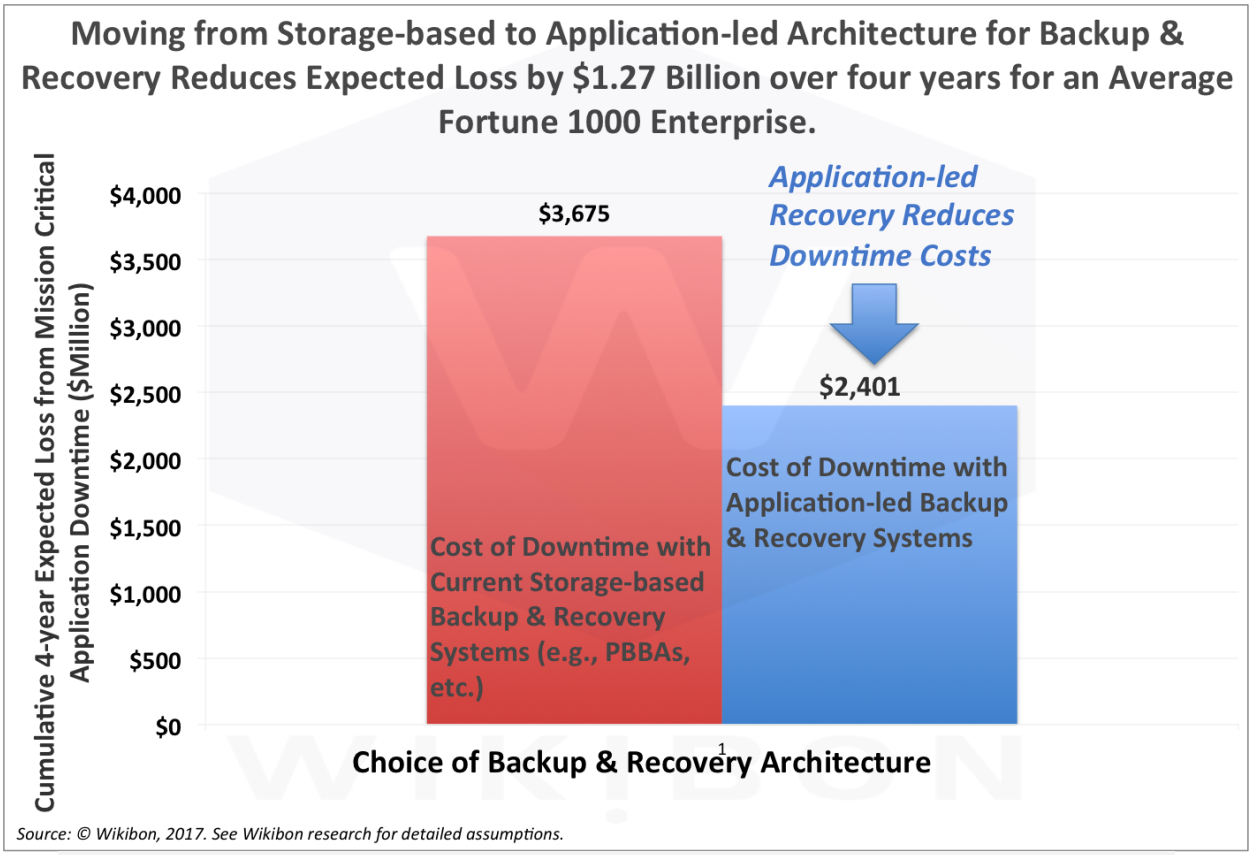

Figure 1 – Benefit of Migration to Application-led Architecture for Mission Critical Backup & Recovery

Source: © Wikibon 2017. See Table 1, Table 2, and Table 3 for Calculations and Assumptions

The research also includes a comparison of a project to migrate fully to an application-led architecture compared with a project to invest in and improve the existing storage-led backup and recovery system. The net NPV for a project to improve the existing systems is projected to be $268 million, with an IRR of 90%, and a breakeven of 11 months. A good project, but not one that puts the enterprise on a strong technical and financial journey. The migration is the better project for the right application portfolio — i.e. high value apps.

The research did not include projects and associated benefits aimed at reducing the number of outages (e.g., DevOps programs, etc.). Wikibon believes that the

fundamental simplification of recovery systems, will make future projects for reducing the number of outages simpler and faster to implement.

Bottom Line: Traditional backup and recovery systems (which are mainly storage-based), and the processes / in-house automation round them, are very complex and error prone, especially for database recovery. Wikibon strongly recommends moving to an application-led backup and recovery architecture for existing applications, and setting the policy that all new application projects should adopt application-led backup and recovery processes.

We believe other database application and file-system suppliers will develop specialized more complete (i.e. end-to-end) data protection solutions. Cloud service providers such as AWS and Microsoft Azure are also investing in end-to-end architectures. Wikibon believe that application-led end-to-end architectures for backup and recovery are a prerequisite for digital transformation projects.

Understanding the Cost of System Downtime

Examples of Spectacular Downtime in Fortune 1000 Enterprises

In just one industry, five airlines had major outages in 2016/2017. Delta, British Airways, JetBlue, Southwest, and United all suffered significant downtime and attracted negative media attention. The impact of that downtime includes:

- Loss of ticket revenues when customers are not able to access sites;
- Cost of compensation for delays and other issues to other airlines, hotels and passengers;
- Loss of trust with flyers, potential customers, and the public, negatively impacting future revenues;
- Damaging the business brand, leading to cancellations, increasing marketing costs and decreasing margins;
- Political fallout due to the high media attention to these types of incidents, resulting in administrative burdens, potential regulatory demands and other business constraints.

Of course, airline outages are publicized in the media. Apple, AWS, Google, Microsoft Office 365, SalesForce, Symantec Cloud, and Twitter can also be added to the list of public outages, because of the large number of customers impacted and the consequent media attention.

Most IT outages will likely fly under the radar of public scrutiny avoiding widespread brand impact, but not the associated IT costs of recovery or post-mortem review of root cause. Downtime for critical business applications such as those impacting manufacturing can add substantial cost to the bottom line estimated at upwards of $10

million dollars for a 15-minute outage of a production line. If customer-facing systems are down or excessively slow, customer frustrations could impact future sales.

Outages that fly under the media don't make them any the less painful and costly to the enterprises that experience them. How often do you hear the words "I'm sorry, my system is down"?

Bottom Line: there are many business-critical applications across all industries in which downtime costs enterprises millions of dollars annually.

What is the Cost of Downtime?

There are many studies on the cost of outages for large (Fortune 1000) enterprises. The following assumptions (see Table 2 below for a full set of assumptions) are used within this report for the annual impact of unplanned downtime on an average Fortune 1000 enterprise:

- Enterprise Annual Revenue for an "Average" Fortune 1000 enterprise = $15 Billion
- Number of Employees = 50,000
- Annual Impact of all IT System Downtime = $1.3 Billion (8.7% of Revenue)
- Annual Impact of Mission Critical System Downtime = $0.7 Billion (4.9% of Revenue)
- Average cost of Downtime for:
    - Infrastructure Failure = $100,000/hour
    - Non-Mission Critical Application Failure = $120,000/hour
    - Mission Critical Application Failure = $750,000/hour

More people, applications and business processes will be impacted by system failure, as enterprise digitization accelerates. The financial and business impact of system failure will grow significantly.

The Root Cause of Downtime

Component vs. System Downtime

If you would believe the IT vendor equipment specifications, nothing would ever go down. In the marketing battles of the "nines", vendors compete with five nines (99.999%), six nines (99.9999%), and now eight nines (99.999999%) availability. Eight nines is an expected downtime of of less that half a second every year. Equipment vendors live in fear of a major outage being blamed on them.

However, equipment and software do not live in isolation, but in systems, together with other equipment and software. Good availability figures of single system components are useful, but not sufficient in achieving high availability of the system as a whole. Systems with redundant lower availability components can perform at high availability

with the right architecture, as Google and AWS have proven. Systems with ultra-high availability components can and do fail catastrophically often.

Recovery from Failure is Key

What is clear is that these airline companies referenced above offered very simple single component explanations of the causes of their downtime. British Airways put the blame on power supplies. Google and AWS pointed to one specific operations individual making a mistake. As any neophyte IT manager knows, these are not honest statements about the true root cause. Investors should be extremely concerned if senior management actually believes these statements. Fixing power supplies and firing an operator will do practically nothing to lower the risk of future outages. **Broken recovery systems are the heart of the problem**.

Investors should demand a detailed and truthful analysis of the root cause of IT failures, and more importantly, ***the inability to recover in a timely manner***. Technology has become more reliable, but will always fail. Humans are error prone. The key to recovering from failure is to have robust and fully tested backup and recovery systems. The root cause is not just the original cause, but the failure of the backup and recovery systems as a whole to recover from that failure.

The most damaging failures occur when mission critical recovery systems fail. These are most always real-time systems, or systems in support of real-time systems. This research focuses on the fundamental architectures and processes of backup and recovery systems for mission critical workloads, and concludes they need to change radically.

Bottom Line: A profound truism is espoused in this simple statement coined by Fred Moore:

> **“*Backup is one thing: Recovery is everything*“**

Exponentials are NOT Intuitive

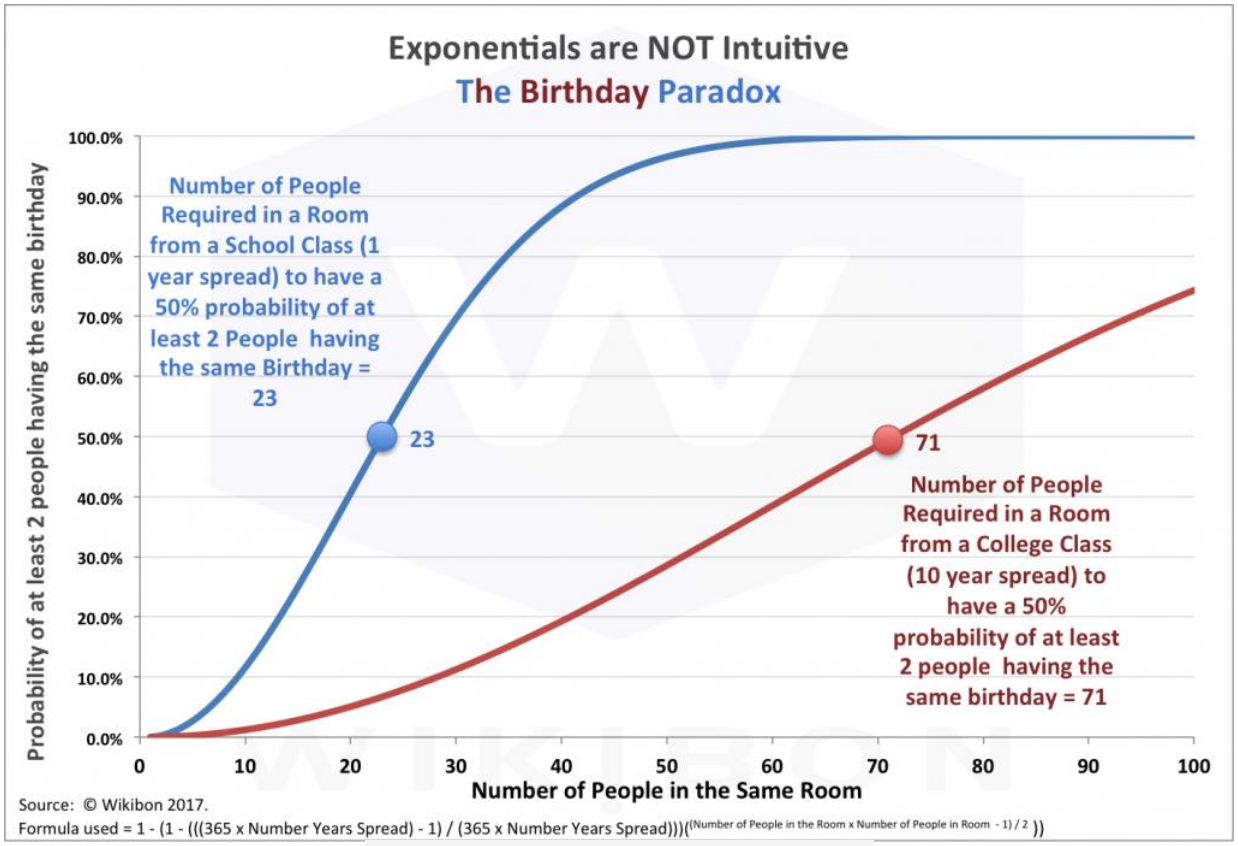

Figure 2a – The Birthday Paradox

Source: © Wikibon 2017

See [Understanding the Birthday Paradox for a full explanation of the math]

Statisticians can skip this section. Remember the “Birthday Paradox”? The question is “How many people do you need in a room to have a 50% probability that 2 people or more have the same birthday?” The answer is 23.

Imagine a school class reunion. The people in the room are typically born the same year. The probability of two people in a room having the same birthday is 1 – (364 ÷ 365) = 0.3%. Yet the blue line in Figure 2a shows that only 23 people are required in the room to have a 50% chance of two people having exactly the same birthday. This is much lower than human intuition would estimate.

Imagine a college class reunion. Let’s assume the birthdays are spread over 10 years. The probability of two people having exactly the same birthday is much smaller. The calculation is 1 – ((365 x 10 – 1) ÷ (365 x 10)) = .03%. Ten times smaller than the high school class reunion. Yet the brown/red line in Figure 2a shows that less than 3 times the number of people are required at the college reunion to have a 50% probability of at

least 2 people have exactly the same birthday, including the year. This is much lower than human intuition would estimate.

Why the paradox? The answer is in the formula at the bottom of chart 2a. The formula has an exponential in it. You can read the "Birthday Paradox" to get the full detail.

Bottom Line: Even if a single failure event has a low probability, the probability with multiple events will usually have an exponential component. This will lead to true probability being much higher than human intuition would estimate.

Key Contributors to Application Downtime

This Wikibon research focuses on the backup and recovery for mission critical systems. Wikibon's working hypothesis is that the overall failure rate for backup and recovery is driven by two major components:

a. The complexity of the backup and recovery process (the number of steps in the process).
    - The steps are defined as when control is passed to another program, process, person or piece of equipment.
    - Included in the step is the hand-over from the preceding step and reporting back to the previous step.
    - In the section "Calibrating the Application Downtime Model in the Real World" below, Wikibon shows the results of analysis of the processes being used in a small number of enterprise installations and concludes that the number of steps is very high in traditional backup systems.

b. The probability of completing each step without error, including communication back to the preceding step and starting the subsequent steps(s), in a stressed environment.
    - In the Wikibon model, all the steps are given an average probability of failure. Some steps, especially those requiring expert judgement, will have much higher probabilities of failure. Others will have lower probabilities.
    - Note: While applying an average probability to all steps is not perfectly precise (and is a simplifying assumption), our research concludes that the impact of using this methodology on the figures below is minimal.

Remember the Birthday Paradox from the last section. The formula that connects a and b above to the application recovery failure probability, illustrated in Figure 2b, is:

$$\text{Application Recovery Failure Probability} = 1 - (1 - \text{Probability of step failure})^{\text{Number of Steps}}$$

This formula has an exponential component. The resulting probability will be much higher than human intuition will calculate.

Figure 2B shows probability of failure to recover (y axis) as a function of:

1. the number of steps (x axis, 1-100)
2. the probability of step failure (different colored curved line, from a probability of step failure of 2% down to 0.1%.

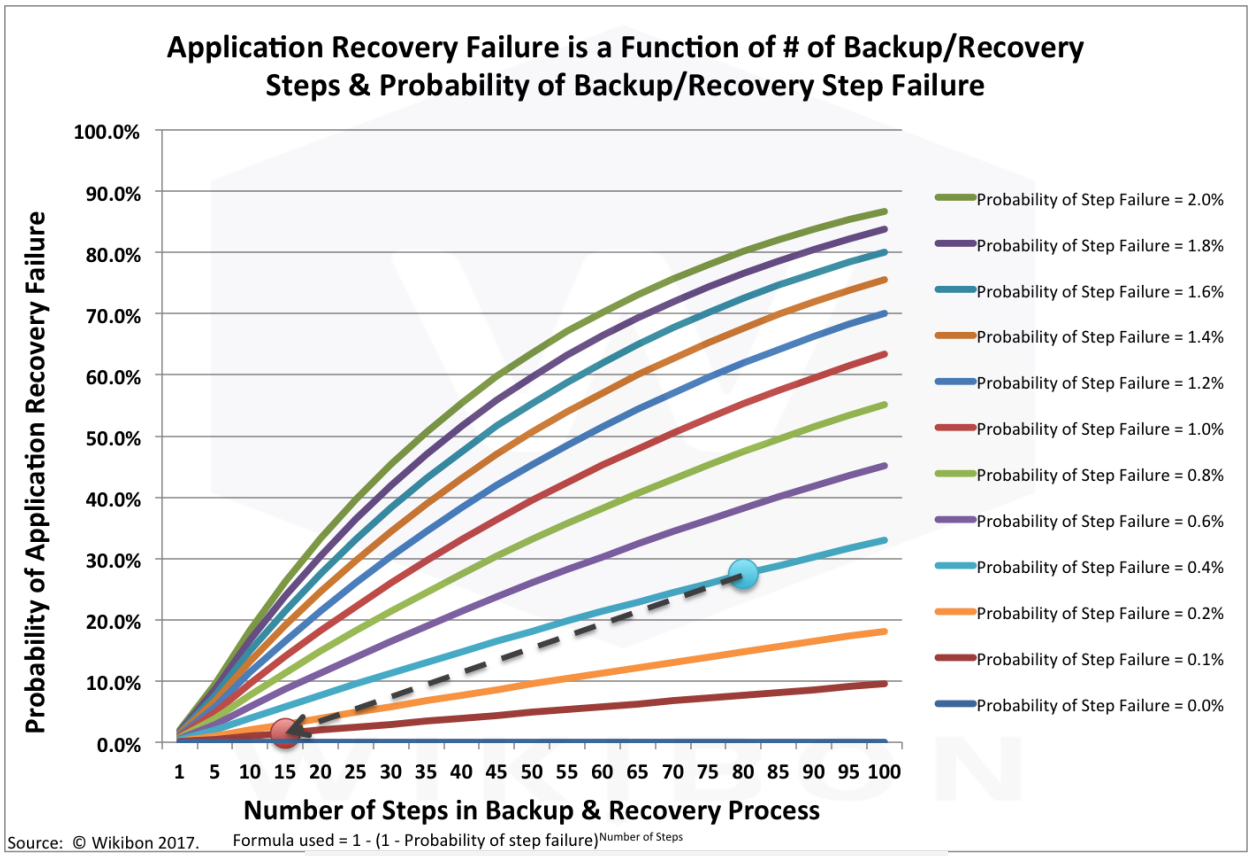

Figure 2b – Key Contributors to Downtime

Source: © Wikibon 2017. Formula used = 1 – (1 – Probability of step failure)Number of Steps

The big blue dot in Figure 2b shows the result of improving both factors. If the starting point is 80 steps, with a probability of step failure of 0.4%, the probability of failure is 27.4%. If we reduce the number of steps to 14, and the probability of failure to 0.1%, the probability of failure reduces to 1.5%. These figures will be used in the detailed case study below, and are shown in Table 2.

Bottom Line: The key focus items for reducing the downtime impact of an outage is to reduce the number of steps and increase automation and testing. This will decrease the probability of Step Failure.

Calibrating the Application Downtime Model in the Real World

Wikibon's research into real world customers using traditional storage-based backup and recovery systems confirmed the theoretical findings of the model illustrated in Figure 2b. Wikibon found that failover recovery works comparatively well in Oracle

database environments, especially in Oracle Database environments where Oracle RAC and/or Active Data Guard features are used.

However, in these same environments application and database recovery from human or software failure (e.g., corruption of a table by an application bug) fails 30-40% of the time when traditional storage-based backup and recovery systems are used. Wikibon's research also confirms that these recovery procedures are tested infrequently, and rarely tested in high-load situations or from backups made in high-load situations. Due to the usual day-to-day focus on meeting backup windows, the recovery processes are frequently compromised in ways that are not discovered or understood until a recovery actually fails.

This research is the basis for using 80 steps and 0.4% failure rate in the Wikibon business cases in the following section. The 0.4% failure rate is the light-blue line in Figure 2b above. The formula calculates a 27% probability of recovery failure, which is slightly conservative, but in line, with the real-world analysis done.

Of course, each failure does not result in total failure, but an extended downtime while the failure is analyzed and resolved, and the applications and database are brought back online. For example, an earlier backup point can be used, and more log-files applied. Alternatively, a manual intervention/correction may cure the problem. However, there is usually some significant impact on recovery time, and an increased risk of other failures. The financial models below assume that in this case the average time to resolve is increased by a factor of 4 times. If data is lost, the average time to resolve is increased by a factor of five (See the assumptions in Table 2 below for details and the calculations of failure impact).

Defining Application-Led & Storage-based Architectures

Many Oracle environments use storage-based systems for backup. EMC was the first to introduce SRDF synchronous and asynchronous failover functionality, with multi-site support. EMC followed that up with a Purpose-built Backup Appliance (PBBA), acquired from Data Domain. All these systems operate, and only operate, at the storage level. They interface with Oracle via Oracle Recovery Manager (RMAN) and Oracle Enterprise Manager (OEM) software. They are not using application-led architectures, with knowledge of transactions or recovery status; rather they are storage-based systems mainly focused on performing backup within the backup window.

A high percentage of the mission critical applications utilize databases as critical middleware for data management and for real-time failover without data loss. Oracle dominates this space with Real Application Clusters (RAC) and Active Data Guard; with Microsoft SQL Server with AlwaysOn and IBM's DB2 with Q Replication also present. All of these databases use an application-led architecture for failover.

For this research, Oracle environments have been chosen as the reference high-availability database, because of its widespread usage, and because it offers an integrated recovery appliance. Wikibon expects Microsoft and IBM to deliver similar recovery capabilities in the future.

The classic models of backup and recovery, from tape to PBBA, are storage-based. An application consistent copy of the database and/or other files at a point in time is created by flushing the database and IO buffers to disk, and a backup copy is created.

In contrast, An application-led architecture sends transactions from memory buffers (before even flushing to disk), thereby reducing data loss exposure to the sub-second level, as ongoing transactions are stored on the production database and Recovery Appliance. As an example, the application-led reference system (The Oracle ZDLRA, referred to as Oracle Recovery Appliance – RA). The RA:

- Transmits database redo logs, using the same technology as Data Guard for redo transport;
- Receives redo and upon a log switch, the archived log is backed up by and on the Recovery Appliance.
- Creates “virtual fulls” in its catalog for each incremental backup taken (only an initial full backup is needed and incremental backups consisting of changed blocks only from that point forward). The advantage of virtual full backups is that during a backup, only an incremental is needed yet during a restore, the Recovery Appliance can send a full database backup or sub-set (e.g. datafile) to the point of an incremental backup,
- Schedules for incremental backups (virtual fulls) that can be set individually for each database. This can be daily or more or less frequent depending on SLAs. Retention of “Virtual fulls” is based on the “Recovery Window” parameter user-defined within a RA Protection Policy which can be days, months or years. As only changed blocks are stored, there is no storage impact for maintaining longer term retention on disk.
- Ensures all database block retention and purging is automatically managed, based on Protection Policy parameters without the need for user-intervention.
- Ensures all database blocks are validated end-to-end, so that recovery will not fail because of corrupted blocks.
- Is designed, architected and tested as fully integrated software, including Oracle’s RMAN, RAC, Active Data Guard, Far Sync, Enterprise Manager (OEM), and from one vendor, with one throat to choke. This end-to-end architecture gives the best probability that databases backup recovery is complete, error free, and consistent.

The key benefit of this application-led end-to-end architecture is that there are fewer components and many fully automated steps. In addition, the backup and recovery software is specifically designed and integrates with full knowledge of the database software architecture. The result is a system where all the updating and upgrading of all the software components is coordinated and tested as a single system.

Building a Model of Application Downtime

The purpose of building a model is to help understand the interaction between technical choices and the financial implications. As is usual practice for Wikibon, the full assumptions and calculations of the model are disclosed in the tables below. This allows interactive discussions and discoveries to take place, and adoption of the model for specific purposes.

This also allows our clients to build their own model themselves, or develop it with Wikibon as part of the Wikibon service.

Key Components of Wikibon Downtime Model

Table 1 shows the key elements of the Wikibon model, and includes the key downtime financial impact assumptions in the first three data rows (taken from the “Cost of Data Loss” section above).

1. The first column details the formula used, in spreadsheet format. The calculations refer to other rows in the model, and to rows in the assumptions table detailed in Table 2 below.
2. The second column describes each component.
3. The third column gives the modeling assumptions and calculation results for a traditional storage-based backup & recovery system. It reflects the current environment, referred to as “Year 0”. It uses the assumptions from the section above entitled “What is the Cost of Downtime” above and the Table 2 assumptions below to calculate the IT and financial impact of downtime.
4. The fourth column gives the modeling results for Application-led backup & recovery architecture in its 4th year after its implementation in year 1. The reference model is the Oracle Recovery Appliance using Active Data Guard and Far Sync. More details of the architecture and solution are shown in “Business Case Conclusions” section below.

Business & IT Impact of Failure and Recovery for Average Fortune 1000 Customer			
Formula (Spreadsheet Format)	Business & IT Impact of IT Failure and Recovery	Traditional Storage-based Database Backup & Recovery	Application-led Database Backup & Recovery Architecture (after 4-years)
a	Infrastructure Failure Cost/Hour	$100,000	$100,000
b	Non-Mission Critical Application Failure Cost/Hour	$120,000	$120,000
c	Mission Critical Application Failure Cost/Hour	$750,000	$750,000
d=a×p + b×q + c×r	Average Cost of Failure/Hour	$213,000	$213,000
e	Average Number of Outages	700	700
f=x×(1 - y) + x×y×z	Average Hours per Outage	8.75	4.19
g=e×f	Average Hours of Failure	6,125	2,930
h=d×g	**Business Cost of System Failure and Recovery**	**$1,304,649,125**	**$624,049,546**
i=c×g×r	**Business Cost of Mission Critical System Failure and Recovery**	**$735,013,591**	**$351,577,209**
j=s×z	Maximum Outage Time (hours)	20	20
k=w×y	Probability of Maximum Outage	1.37%	0.001%
l=g×aa×ab÷ac÷ad	Staff Cost IT for Failure	$6,562,621	$3,139,082
m=ae×af	Cost of Staff Operations for Backup and Recovery	$6,600,000	$2,400,000
ah	Cost of Hardware & Software for Backup & Recovery	$13,162,621	$13,162,621
n=l + m + ah	**Total Operational Cost of Recovery**	**$26,325,243**	**$18,701,704**
o=n×ag	**Total Operational Cost of Mission Critical Backup & Recovery**	**$9,572,816**	**$6,800,619**
Source: © Wikibon, 2017			

Table1 – Wikibon Backup and Recovery Model for Large Organizations

Source: © Wikibon 2017. See Table 2 for detailed assumptions

Table 2 below lists the assumptions below used to calculate Table 1. It uses the same format as Table 1.

Formula (Spreadsheet Format)	Assumptions	Traditional Storage-based Database Backup & Recovery	Application-led Database Backup & Recovery Architecture (after 4-years)
p	Percentage Infrastructure Failures	39%	39%
q	Percentage Non-Mission Critical Application Failures	45%	45%
r	Percentage Mission Critical Database Application Failures	16%	16%
s	Time to recover when Recovery works 1st time (hours)	4	4
t	Time to recover when Recovery does not work 1st time (hours)	16	16
u	**Number of Steps in Backup & Recovery**	**80**	**14**
v	**Probability of Step Success**	**99.60%**	**99.89%**
w=1-v^u	**Probability of Recovery Error**	**27.4%**	**1.4%**
x=s×(1-w) + t×w	Unplanned Outage Time to Recover	7.3	4.2
y	Probability of Data Loss or Manual Recovery given a Recovery Failure	5%	0.1%
z	Impact on Data Recovery time if Data is Lost/Manually Recovered (times)	5	5
aa	Additional Hours of IT for every hour of outage	10	10
ab	Average IT Recovery Salary	$150,000	$150,000
ac	Average Hours/day for IT	7	7
ad	Average Days/year for IT (less training, vacation, etc.)	200	200
ae	Average IT Salary for Operational Staff in Backup & Recovery	$120,000	$120,000
af	Number of Operational Staff in Backup & Recovery	55	20
ag=3×r÷(p + q + 3×r)	Percentage of IT Staff & Hardware for Mission Critical Backup & Recovery	36%	36%
ah=l + m (for Trad)	Cost of Hardware & Software for Backup & Recovery for Traditional = App-led	$13,162,621	$13,162,621
ai	Annual Revenue for Average Fortune 1000 Enterprise	$15,000,000,000	$15,000,000,000
aj	Number of Employees	50,000	50,000
ak=ai÷aj	Revenue/Employee	$300,000	$300,000
al=h÷ai	Percentage Revenue Impact of all Downtime	8.7%	4.2%
al=i÷ai	Percentage Revenue Impact of Mission Critical Downtime	4.9%	2.3%
am	Time to Implement Application-led Backup & Recovery System (months)	n/a	6
Source: © Wikibon, 2017			

Table 2 – Assumptions for Wikibon Backup and Recovery Model for Large Organizations

Source: © Wikibon 2017.

There are four inputs in Table 2 which are different between the third and fourth columns. They are:

1. Number of steps in recovery (u), which declines in steps from 80 in year 0 to 14 in year 4
2. Probability of step success (v), which declines in steps from 99.6% in year 0 to 99.9% in year 4
3. Probability of data loss (y), which declines with the number steps from 5% in year 0 to 0.1% in year 4.
4. Number of Operational Staff in Backup & Recovery (af), which declines from 55 in year 0 to 20 in year 4.

As a calculation from 1. and 2. above, the probability of recovery error (w) declines from 27.4% to 1.4%.

The full details of assumptions and calculations by year for the alternatives of improving the current storage-based backup and recovery system and deploying an application-led system are shown in Table 5 in the Footnotes below.

Business Case for Migrating to Application-led Architecture

Table 3 below shows business case for migrating to from a traditional storage-based backup and recovery to an application-led end-to-end backup and recovery system. The analysis shows two business scenarios, one with no change to the storage-based systems, and the other migration to an application-led architecture.

Table 4 below shows the business case for improving the current storage-based system over 4 years.

Figure 3 below compares the two business cases, the case for migration to application-led, and the case for improving the existing storage system. Both are valid cases; the case for application-led migration is overwhelmingly better, from both strategic and financial perspectives.

The Business Case for IT

The business case for IT is the first case in Table 3. The main savings is in the decrease in staff from 50 in year 0 to 20 in year 4 and resulting decrease in operational costs. From a budget perspective, it almost draws even after 4 years. When the cost of money in included, it shows a net return of -$393K.

The purpose for showing this business case is to emphasize that the return on investment is not great just for IT, but the business benefits in the section below are phenomenal. IT will not take a major hit on its four-year budget, but will need the buy-in from the business to justify the investment over other projects. With $1 Billion in business savings shown in Table 3, it should not be difficult.

Business & IT Impact of Migrating from Traditional Storage-based Recovery to an Application-led Architecture for an average Fortune 1000 Customer							
IT Migration and Annual IT Costs		**Initial Annual Costs**	**Year 1**	**Year 2**	**Year 3**	**Year 4**	**Total**
	Staff Cost IT for Failure	$2,386,408	$1,716,127	$1,351,121	$1,198,836	$1,141,484	**$5,407,569**
	Cost of Staff Operations for Backup and Recovery	$2,400,000	$1,963,636	$1,527,273	$1,090,909	$872,727	**$5,454,545**
	Cost of Hardware & Software for Backup & Recovery	$4,786,408	$4,786,408	$4,786,408	$4,786,408	$4,786,408	**$19,145,631**
	Project Cost for Reducing steps and probability of failure/step		$5,000,000	$1,100,000	$1,100,000	$1,100,000	**$8,300,000**
Total 4-year Operational Cost of Mission Critical Backup & Recovery		**$9,572,816**	**$13,466,171**	**$8,764,802**	**$8,176,153**	**$7,900,619**	**$38,307,746**
	Saving for IT		-$3,893,356	$808,014	$1,396,663	$1,672,196	**-$16,483**
	Net Cumulative Savings for IT		-$3,893,356	-$3,085,342	-$1,688,680	-$16,483	**-$16,483**
IT Financial Analysis for Migrating to an Application-led Architecture for Backup & Recovery (4-year)							
	Net Present Value for IT Project (5%)		**-$392,856**				
	IRR for IT Project		**0%**				
	Breakeven for IT Project		**>48 Months**				
Total 4-year Business Impact of Migrating to an Application-led Architecture for Backup & Recove		Year 0	Year 1	Year 2	Year 3	Year 4	4-year Total
	Business Cost of Mission Critical System Failure and Recovery	$735,013,591	$631,790,390	$416,145,409	$369,241,465	$351,577,209	$1,768,754,473
	Total Savings for the Business		$99,329,846	$319,676,196	$367,168,789	$385,108,578	$1,171,283,409
	Net Cumulative Savings for the Business		$99,329,846	$419,006,042	$786,174,831	$1,171,283,409	**$1,171,283,409**
Financial Analysis for Migrating to an Application-led Architecture for Backup & Recovery (4-year)							
	Net Present Value For Total Project (5%)		**$1,037,581,414**				
	IRR for IT Project		**591%**				
	Breakeven for Total Project		**6 Months**				

Source: © Wikibon, 2017

Table 3 – Business Case for Deployment of an Application-led architecture for Backup & Recovery for Large Organizations

Source: © Wikibon 2017.

The Business Case for the Business as a Whole

The second of the business cases in Table 3 shows an overall case for migrating to an application-led architecture. The net savings are projected to be $1.2 Billion over the four years. After taking into account the time value of money, the NPV (5%) for the project is $1.0 Billion. The breakeven is just six months (the time allowed for the project). The Internal Rate of Return (IRR) is 591%, enough to make even a finance director smile.

Migrating to Application-led vs. Improving Storage-based

The last issue to consider is an alternative to migrating to an Application-led system, which is to upgrade the existing storage-based systems and processes. These systems (PBBAs, etc.) have been around for a number of years, as have the manual procedures and scripts around them. The low-hanging fruit for improvement is long-gone.

Figure 4 below shows the business case for improving the current storage-led system. The savings again come from the business, and are driven by the same key metrics as discussed above. The detailed assumptions and improvements can be seen in the red section of Table 5 in the footnotes.

In normal circumstances, the business case for improvement would be good. $268K in NPV savings, 90% IRR, and an 11-month breakeven certainly pass muster. However,

both strategically and financially, this is not an optimum project when compared with the migration to application-led project.

Business & IT Impact of Improving the Installed Traditional Storage-based Backup & Recovery System							
Total 4-year Business Impact for Improving Existing Traditional Storage-based Backup & Recovery		Year 0	Year 1	Year 2	Year 3	Year 4	4-year Total
	Business Cost of Mission Critical System Failure and Recovery	$735,013,591	$721,148,985	$670,144,103	$635,679,446	$603,720,463	$2,630,692,996
	Total Savings for the Business		$13,864,607	$64,869,489	$99,334,146	$131,293,128	$309,361,370
	Net Cumulative Savings for the Business		$13,864,607	$78,734,095	$178,068,241	$309,361,370	**$309,361,370**
Financial Analysis for Improving Existing Backup & Recovery Infrastructure (4-year)							
	Net Present Value (5%)		**$267,517,284**				
	IRR for IT Project		**90%**				
	Breakeven for Total Project		**11 Months**				

Source: © Wikibon, 2017

Table 4 – Business Case for Upgrading the Existing Storage-Based architecture for Backup & Recovery for Large Organizations

Source: © Wikibon 2017.

Business Case Conclusions

Figure 3 below is a graphical summary of the two business cases.

1. The blue columns show the cumulative benefit of migration to an application-led backup and recovery architecture, expressed as a net present value assuming a 5% value of money per year.
    - The rightmost blue column shows a cumulative NPV value of $1.0 billion dollars at the end of four years.
    - The blue line shows cost of downtime as a percentage of annual revenue for mission critical applications. It starts at the current value of 4.9% for year 0 and the first six months of year 1. It reduces to 2.3% by year 4.
    - By year three, the cost of downtime for mission critical systems has been halved
2. The red columns show the cumulative benefit of improving the installed storage-based backup and recovery system, expressed as a net present value assuming a 5% value of money per year.
    - The rightmost red column shows a cumulative NPV value of $268 million dollars at the end of four years.
    - The red line shows cost of downtime as a percentage of annual revenue for mission critical applications. It starts at the current value of 4.9% for year 0 and the first six months of year 1. It reduces to 4.0% by year 4.

The overall conclusion is clear and unambiguous. The NPV of selecting the migration project is $770 million over four years. In addition, it positions the enterprise to reduce the number of outages, mainly through development and DevOps initiatives.

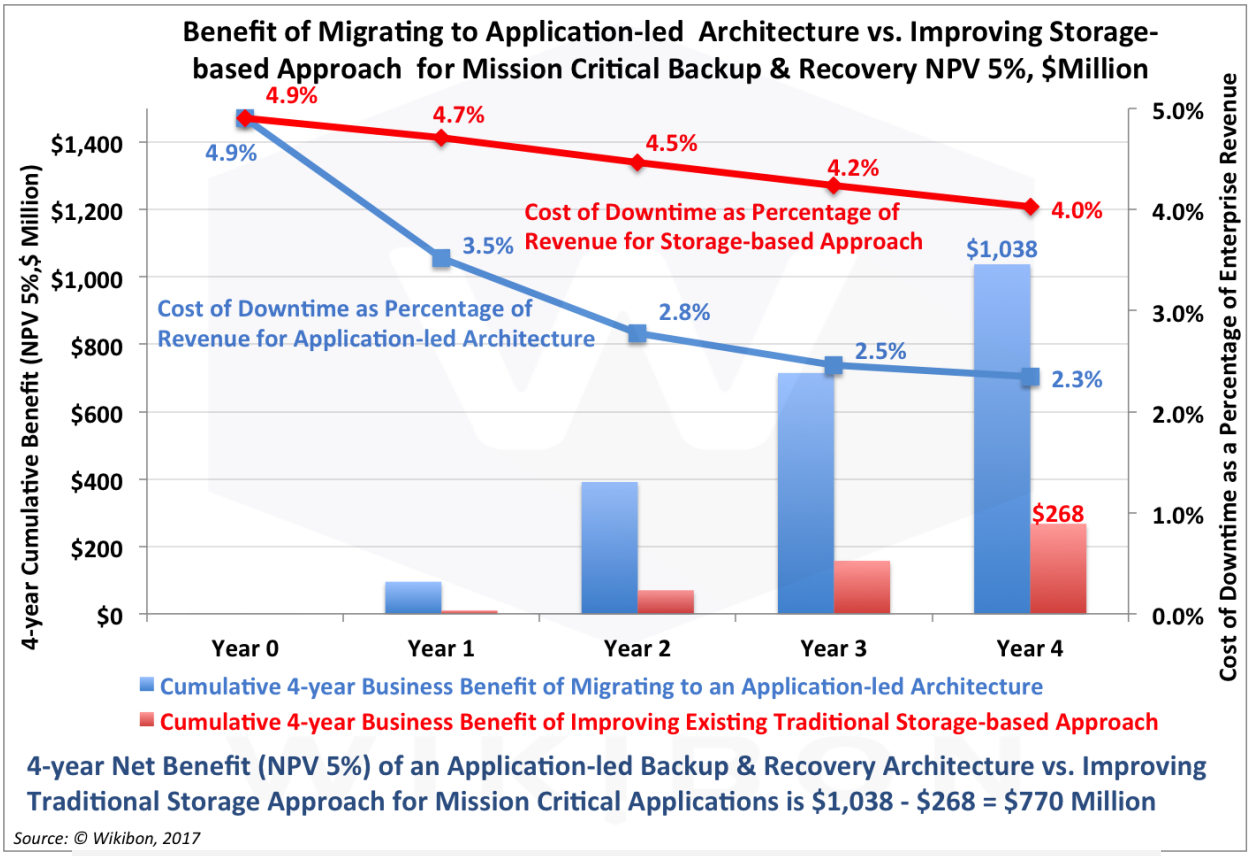

Figure 3 – Comparison of Migration to Application-led Architecture with Upgrading existing Storage-based for Mission Critical Backup & Recovery

Source: © Wikibon 2017. See Table 1, Table 2 and Table 5 for detailed assumptions.

Strategically, for Oracle users, implementing an application-led backup and recovery strategy for the mission critical applications is the superior decision.

- The application-led strategy is by far-and-away the best strategy for reducing data loss. A Recovery Point Objective (RPO) of close to zero is made possible with the Oracle Recovery Appliance. In addition, an Oracle Far Sync implementation can improve asynchronous distance recovery, and reduce this to close to zero.
- The addition of Axxana and its amazing crush/flood/fire resistance capabilities can make a local Far Sync server break the speed of light constraints for asynchronous replications, and reach a theoretical RPO of very close to zero. Reduction of data loss is an important contribution to reducing the recovery time, and in the work required by the lines of business to recover from lost data.
- Remote fail-over and recovery services and management services are available from cloud service providers such as Secure-24 and Oracle Public Cloud, with the same advantages of an integrated solution.

- The relational Data Services from AWS are taking a similar architectural route by providing availability and recovery IaaS services on top of MySQL. Although AWS is nowhere near the completeness of solution, or has the experience and capabilities of Oracle, their approach may be useful for applications with less stringent SLAs that do not share data with Oracle applications.
- Wikibon expects Microsoft to take a similar application-led approach with the Azure stack for Microsoft SQL Server.

Overall Downtime Conclusions

The Business Importance of Recovery

The business importance of effective recovery is illustrated forcefully by the average cost of mission critical downtime at 5% of revenue. Digitalization is increasing the dependence of all departments on IT systems. The airline outages illustrated that the ground staff could not help customers in anyway. All the systems were linked.

Adding to the challenges of recovery is the threat of ransomware and other malware. The ability to recover from these attacks is essential. In addition multiple backups with air-gaps are important. The ability of transmit real-time backup data to multiple independent sources is important (e.g., Moreover, onerous malware, such as ransomware, and heightened regulatory compliance requirements (e.g. GDPR & FRTB) further underscore the need to re-think data recovery approaches.

Benefits of Application-led Architectures

The benefits of moving to an Application-led backup and recovery architecture are:

- Reduced complexity of the environment (fewer steps)
- The end-to-end architecture improves the probability that the steps and the communication to/from the adjacent steps before and afterwards will complete correctly, as measured from a systems perspective.

The benefits, whatever the application-led architecture chosen, will not come immediately. It is hard work to refresh the processes and procedures needed, then identify and reduce the steps required, increase automation and testing with the integrated orchestration tools, and perform the analysis to implement improvements that will eliminate weak links.

The vendor is an integral part of the journey, as it is far better to improve the product for many customers together with integrated testing, than allow each customer to use sealing wax and scripts to patch the product.

Overall Downtime Impact Conclusions

From an architectural and business case points of view, Wikibon believes Oracle has hit a home run with the fundamental architecture of the Oracle Recovery Appliance. Clearly the IT and Business financial savings and IRR are superior for the Recovery Appliance compared to a storage-based backup appliance. Migrating to the Recovery Appliance is a no-brainer—even the CFO will smile.

The users that were interviewed by Wikibon were all very happy with the performance of the appliance for backup and recovery. Most of the praise is aimed at the end-to-end management support of the Recovery Appliance. They have a clear idea of the recoverability status of every protected database, meeting/not meeting SLAs, and the confidence that recoveries just work. User management viewed this as a major improvement over the management experience with the previous storage-based solutions.

There are still areas where the Recovery Appliance needs additional functionality, such as improved role-based access, ability to store application code levels with the data, Far Synch fail-over within Axxana, remote cloud services, non-Oracle data and application management, improved automated recovery testing and much more. The integration with automation and orchestration procedures needs to be improved. The ZDLRA functionality ensures end-to-end checking between RMAN and the ZDLRA virtual fulls; this needs to be extended to automated recovery checking with different scenarios. And the customers we talked to would love less internecine warfare between Oracle representatives, and for Oracle to be an easier company to do business with.

Strategically, Wikibon recommends moving to an application-led architecture, and allowing the PaaS/IaaS providers to deliver the necessary services to implement high availability with high recoverability.

Action Item

Senior business and IT executives should set a goal to measure the cost of downtime, and halve it within 4 years. Senior IT executives should ensure that the measurement and testing of recoverability is front and center in reducing the high cost of downtime. As a strategic imperative in support of this initiative, enterprise IT needs to move from a Smörgåsbord of backup & recovery solutions held together with in-house scripts, and simplify the recovery system by implementing an Application-Led end-to-end architecture for backup and recovery of mission critical systems.

Footnotes

Table 5 below shows the Year 1 to year 4 calculations and assumptions for application-led systems migration project (green) and the traditional systems improvement project (red).

	Year 1	Year 2	Year 3	Year 4	Year 1	Year 2	Year 3	Year 4
Business & IT Impact of IT Failure and Recovery	**Years 1-4 Improvement of Existing Storage-based Infrastructure**				**Years 1-4 of Migration Project from Storage-based backup/Recovery to an Application-led Architecture**			
Infrastructure Failure Cost/Hour	$100,000	$100,000	$100,000	$100,000	$100,000	$100,000	$100,000	$100,000
Non-Mission Critical Application Failure Cost/Hour	$120,000	$120,000	$120,000	$120,000	$120,000	$120,000	$120,000	$120,000
Mission Critical Application Failure Cost/Hour	$750,000	$750,000	$750,000	$750,000	$750,000	$750,000	$750,000	$750,000
Average Cost of Failure/Hour	$213,000	$213,000	$213,000	$213,000	$213,000	$213,000	$213,000	$213,000
Average Number of Outages	700	700	700	700	700	700	700	700
Average Hours per Outage	8.42	7.98	7.57	7.19	6.29	4.95	4.40	4.19
Average Hours of Failure	5,894	5,585	5,297	5,031	4,405	3,468	3,077	2,930
Business Cost of System Failure and Recovery	**$1,255,429,771**	**$1,189,505,782**	**$1,128,331,016**	**$1,071,603,822**	**$938,206,758**	**$738,658,100**	**$655,403,601**	**$624,049,546**
Business Cost of Mission Critical System Failure and Recovery	**$707,284,378**	**$670,144,103**	**$635,679,446**	**$603,720,463**	**$528,567,188**	**$416,145,409**	**$369,241,465**	**$351,577,209**
Maximum Outage Time (hours)	20	20	20	20	20	20	20	20
Probability of Maximum Outage	1.257%	1.035%	0.841%	0.672%	0.055%	0.015%	0.003%	0.001%
Staff Cost IT for Failure	$6,315,039	$5,983,429	$5,675,709	$5,390,361	$4,719,350	$3,715,584	$3,296,799	$3,139,082
Cost of Staff Operations for Backup and Recovery	$6,600,000	$6,000,000	$5,400,000	$4,800,000	$5,400,000	$4,200,000	$3,000,000	$2,400,000
Cost of Hardware & Software for Backup & Recovery	$13,162,621	$13,162,621	$13,162,621	$13,162,621	$13,162,621	$13,162,621	$13,162,621	$13,162,621
Total Operational Cost of Recovery	**$26,077,660**	**$25,146,051**	**$24,238,331**	**$23,352,983**	**$23,281,971**	**$21,078,205**	**$19,459,420**	**$18,701,704**
Total Operational Cost of Mission Critical Backup & Recovery	**$9,482,786**	**$9,144,018**	**$8,813,938**	**$8,491,994**	**$8,466,171**	**$7,664,802**	**$7,076,153**	**$6,800,619**
	Assumptions				**Assumptions**			
Assumptions	**Application-led Database Backup and Recovery Architecture**	**Application-led Database Backup and Recovery Architecture**	**Application-led Database Backup and Recovery Architecture**	**Application-led Database Backup and Recovery Architecture**	**Application-led Database Backup and Recovery Architecture**	**Application-led Database Backup and Recovery Architecture**	**Application-led Database Backup and Recovery Architecture**	**Application-led Database Backup and Recovery Architecture**
Percentage Infrastructure Failures	39%	39%	39%	39%	39%	39%	39%	39%
Percentage Non-Mission Critical Application Failures	45%	45%	45%	45%	45%	45%	45%	45%
Percentage Mission Critical Database Application Failures	16%	16%	16%	16%	16%	16%	16%	16%
Time to recover when Recovery works 1st time (hours)	4	4	4	4	4	4	4	4
Time to recover when Recovery does not work 1st time (hours)	16	16	16	16	16	16	16	16
Number of Steps in Backup & Recovery	**76**	**72**	**69**	**65**	**60**	**33**	**20**	**14**
Probability of Step Success	**99.62%**	**99.64%**	**99.66%**	**99.67%**	**99.66%**	**99.76%**	**99.84%**	**99.89%**
Probability of Recovery Error	**25.1%**	**23.0%**	**21.0%**	**19.2%**	**18.5%**	**7.6%**	**3.2%**	**1.4%**
Unplanned Outage Time to Recover	7.0	6.8	6.5	6.3	6.2	4.9	4.4	4.2
Probability of Data Loss or Manual Recovery given a Recovery Failure	5.0%	4.5%	4.0%	3.50%	0.3%	0.2%	0.1%	0.1%
Impact on Data Recovery time if Data is Lost/Manually Recovered (times)	5	5	5	5	5	5	5	5
Additional Hours of IT for every hour of outage	10	10	10	10	10	10	10	10
Average IT Recovery Salary	$150,000	$150,000	$150,000	$150,000	$150,000	$150,000	$150,000	$150,000
Average Hours/day for IT	7	7	7	7	7	7	7	7
Average Days/year for IT (less training, vacation, etc.)	200	200	200	200	200	200	200	200
Average IT Salary for Operational Staff in Backup & Recovery	$120,000	$120,000	$120,000	$120,000	$120,000	$120,000	$120,000	$120,000
Number of Operational Staff in Backup & Recovery	55	50	45	40	45	35	25	20
Percentage of IT Staff & Hardware for Mission Critical Backup & Recovery	36%	36%	36%	36%	36%	36%	36%	36%
Cost of Hardware & Software for Backup & Recovery for Traditional = App-led	$13,162,621	$13,162,621	$13,162,621	$13,162,621	$13,162,621	$13,162,621	$13,162,621	$13,162,621
Annual Revenue for Average Fortune 1000 Enterprise	$15,000,000,000	$15,000,000,000	$15,000,000,000	$15,000,000,000	$15,000,000,000	$15,000,000,000	$15,000,000,000	$15,000,000,000
Number of Employees	50,000	50,000	50,000	50,000	50,000	50,000	50,000	50,000
Revenue/Employee	$300,000	$300,000	$300,000	$300,000	$300,000	$300,000	$300,000	$300,000
Percentage Revenue Impact of all Downtime	8.4%	7.9%	7.5%	7.1%	6.3%	4.9%	4.4%	4.2%
Percentage Revenue Impact of Mission Critical Downtime	4.7%	4.5%	4.2%	4.0%	3.5%	2.8%	2.5%	2.3%
Time to Implement Application-led Backup & Recovery System (months)	n/a	n/a	n/a	n/a	6	6	6	6
	Source: © Wikibon 2017				*Source: © Wikibon 2017*			

Table 5 – Year 1 to Year 4 Detailed Calculations and Assumptions for Wikibon Backup and Recovery Model for Large Organizations

Source: © Wikibon 2017.

The calculation details for each line in Table 5 are shown in Tables 1 and 2 above.