



The Economic Value of Axxana's Zero Data Loss Solution

Authors: David Floyer & Dave Vellante

Wikibon.org

Table of Contents

Introduction.....	1
Model Overview	2
Base Information Required by Model.....	2
DR Alternatives Assessed by Model	2
Base Case of Offsite Tape Only.....	2
Two Data Center Synchronous	3
Two Data Center Asynchronous.....	3
Three Data Center DR solutions.....	3
Enterprise Data Recording.....	3
General Findings.....	4
Specific Findings	4
Summary and Conclusions from Model	6
Caveats.....	7
Action Item	7
Figure 1 - Benefit of Alternative DR Solutions for Bank.....	5
Figure 2 - Cost of Alternative DR Solutions for Bank	6
Table 1 - Inputs and Assumptions	8

Introduction

In late 2008, Israeli-based startup Axxana approached Wikibon to conduct an independent assessment of its zero-data-loss solution. This was a paid consulting engagement that involved interviewing beta customers and building out a cost-and-value model to compare the cost and expected data loss of Axxana's solution to likely alternative approaches and, consequently to each other.

The economic model Wikibon developed leveraged an existing tool set created by Wikibon in collaboration with several large banks. The model uses off-site tape backup as the baseline to

compare the costs and value of alternative disaster recovery solutions and was extended to add Axxana's Enterprise Data Recording (EDR) technology.

Tape recovery was used as a baseline because it represents a lowest common denominator for comparison. Virtually all organizations can relate to, conceptualize and roughly quantify the costs and business process impacts associated with a tape-based recovery solution. As such this technique allows Wikibon to make mathematically consistent and defensible comparisons across any use case scenario. The comparisons described here are based exclusively on open systems markets and do not include an analysis of mainframe environments.

The purpose of this note is to describe the findings of the model and share our economic conclusions with the Wikibon community. It was developed without Axxana's input, although the company is free, as is any Wikibon member, to provide edits and comments.

Model Overview

The economic value model we developed captures several information sources necessary to evaluate the costs and benefits associated with virtually any disaster recovery (DR) solution. They include information necessary to determine both cost and expected loss in real value terms, as well as drive key assumptions relative to various DR scenarios which are described below.

The information is then fed to a back-end calculation engine that determines the expected loss for each scenario and the cost to meet RTO and RPO requirements. These inputs and assumptions drive a series of results that allow IT management to perform cost/benefit comparisons across various solutions.

Base Information Required by Model

The basic data include:

- The revenue contribution of the applications supported by the DR infrastructure;
- The storage capacity being protected;
- The distance between the primary and backup sites (as the line lurches);
- I/O rate for applications being protected;
- The business impact of an outage;
- The probability of an outage;
- Line costs for the specific locations;
- Cost of primary storage;
- Data necessary to calculate NPV and other financials.

DR Alternatives Assessed by Model

The model addresses a spectrum of five (5) main scenarios:

1. Base Case of Offsite Tape Only

This is a well understood, low cost and time honored method of managing disaster recovery. A coherent set of data is written to tape, and transported off-site. The amount of data that would be lost in a disaster (Recovery Point Objective, RPO) is determined by how often the system is backed up and how quickly it is transported off-site. If a system is backed up daily, and it takes six hours to move the data off-site, then the average amount of data lost will be

that created an 18 hours production window. There are new technologies for improving the time to transfer the data (snapshots, de-duplication, and transmission over the wire) but the fundamentals have not changed significantly.

The Recovery Time Objective (RTO) to bring up the system again on another system, reconcile the data lost and bring it back into full production is usually measured in days.

2. Two Data Center Synchronous

Creating a second copy of the data in a second site at the same time as the original data is created means that if the primary site is lost, the second site can recover without any data being lost. This RPO-zero solution is well understood and used extensively in the financial sector. Usually the second data center is fully equipped with servers, and because no data is lost, recovery time (RTO) is usually quick and can be a matter of hours. However, the distance that the two data centers can be apart is limited to less than twenty miles for most applications (see Asynchronous below). There is a significant chance that both sites will be taken out by a rolling disaster (disaster being defined in a broad sense that could, for example include local unrest, union action etc.) Then recovery would have to be made from a normal tape backup. Financial watchdogs such as the SEC strongly recommend distances of 200 miles or more between data centers to eliminate this risk.

3. Two Data Center Asynchronous

The ideal location for a second data center is hundreds of miles away from the primary site. However, for normal applications it is not possible to keep the data exactly consistent in both locations (if you wait for an acknowledgment that the data is safe at the second site, the delays in the transmission, even at near the speed of light, mean that wait time for I/O is poor and throughput slows to a crawl). An asynchronous DR solution keeps a small buffer of information at the primary end, and ensures that a coherent set of data is transmitted to the other end. The RPO is much better than a tape backup solution, but the fact that some data is lost means that the RTO is longer than Synchronous solutions because the databases need to be reconciled with other business records.

4. Three Data Center DR solutions

Three data center solutions are a hybrid between synchronous and asynchronous solutions. A two data synchronous solution is set up between the primary and the “B” site less than twenty miles away, and a second asynchronous connection is set up with the remote site (“C” site). Either the B site is connected to the C site or the primary A site is connected to the C site or both. This cascaded or multi-hop approach ensures that most of the time, failovers can occur to the B site without data loss, and in the case that both the primary A and B sites are taken out, the C site can recover with less data loss, much more quickly than a tape recovery approach. However the transmission line, data center infrastructure and storage costs of such solutions are very high and only used by a relatively small number of organizations (mostly financial).

5. Enterprise Data Recording

The Enterprise Data Recording solution from Axxana is logically the same as a three data center solution. The difference is that a synchronous copy of the data that has not been sent to

the remote site is held on the primary site. It is protected from a disaster not by distance, but by Axxana's "Black-box" technology that provides physical protection from fire, water or earthquake. In the event of a disaster the data in the black-box can be recovered by cellular transmission and transmitted to the remote location. This enables a zero-data loss solution with two data centers at extended distances.

General Findings

Our conclusions focus on two main areas:

- What is the impact of Axxana on benefits?
- What does the model say about the impact of Axxana on costs?

At the highest level, the Axxana technology brings the probability of losing data very close to zero at asynchronous distances. It provides the same level of business protection from data loss as a three data center solution and a higher level of protection than either synchronous or asynchronous solutions. Because no data is lost, RTO time should also be better than asynchronous solutions. RTO will be slower than synchronous solutions if the second site is unaffected, but much faster in the case that both synchronous sites are affected by a disaster. Conceptually, compared to alternatives, the Axxana approach simplifies implementation and testing of near zero data loss solutions.

From a cost perspective, Axxana's solution reduces the expense of lines from the remote site and lowers the cost of storage and other equipment. The Axxana solution should also decrease staff/management costs, primarily because the environment is simpler. It also allows much easier testing and fail-over testing and provides IT and Business Managers a more solid basis on which to judge business resilience.

The main impacts of an Axxana solution on cost are:

- It reduces line costs by decreasing the peak threshold required for a desired service level;
- It reduces the cost of storage because less redundancy is needed to meet the same recovery objectives;
- It simplifies the set up and operating environment;
- It allows much easier testing of DR function.

In theory, the Axxana approach will allow organizations to eliminate or avoid building an entire data center (e.g. the B site) in a three data center solution. However the solution must be proven in the market before this strategy is widely adopted.

The Axxana solution is not appropriate for very small systems (e.g. below about 20TB) where the cost of the Axxana solution is higher than a simple replicated solution and the cost of Axxana would, therefore, exceed the benefits.

Specific Findings

In order to illustrate the economics of Axxana's solution, we have run the following case example through the model. This was based on an engagement with a US bank. The customer profile is:

- A mid-sized bank with assets of \$20B, growing through regional acquisition;
- Core retail banking applications with some commercial and high-net worth applications;
- Two locations about 350 miles apart;
- As the bank has grown, the bank is under increasing pressure from auditors and the SEC to improve recovery time objectives (RTO) to under four hours.

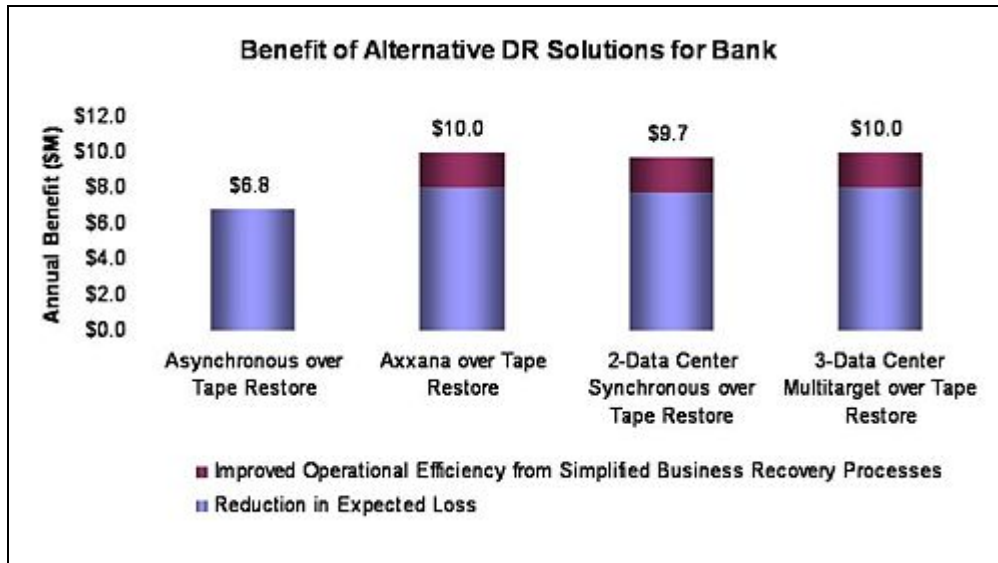


Figure 1 - Benefit of Alternative DR Solutions for Bank

The impact of Axxana's solution on the benefit side is notable and essentially identical to non-tape alternatives. Specifically, Figure 1 shows these benefits relative to alternative DR approaches. The primary benefit calculated is the reduction in expected loss (i.e. the lower probability of losing data) as a result of putting in place a disk-based recovery solution (synchronous, asynchronous or three data center). In each scenario, the benefits of the target DR solution are based on a comparison to tape-based recovery. As such, relative to tape-based recovery, all solutions show substantial benefits from a reduction in expected loss.

This factor is due primarily to the following points:

- The RTO of all alternative scenarios is dramatically improved over tape's 96 hours;
- The RPO in all alternative cases is dramatically improved relative to tape's 18 hours of data loss;
- The expected loss of the asynchronous solution is greater than the solutions with a synchronous component;
- The simplification of business recovery processes from zero data loss did not accrue to the asynchronous solution.

Overall, all disk based solutions were significantly better than the current tape-based solution and Axxana's EDR demonstrates benefits that are equal to or greater than alternatives.

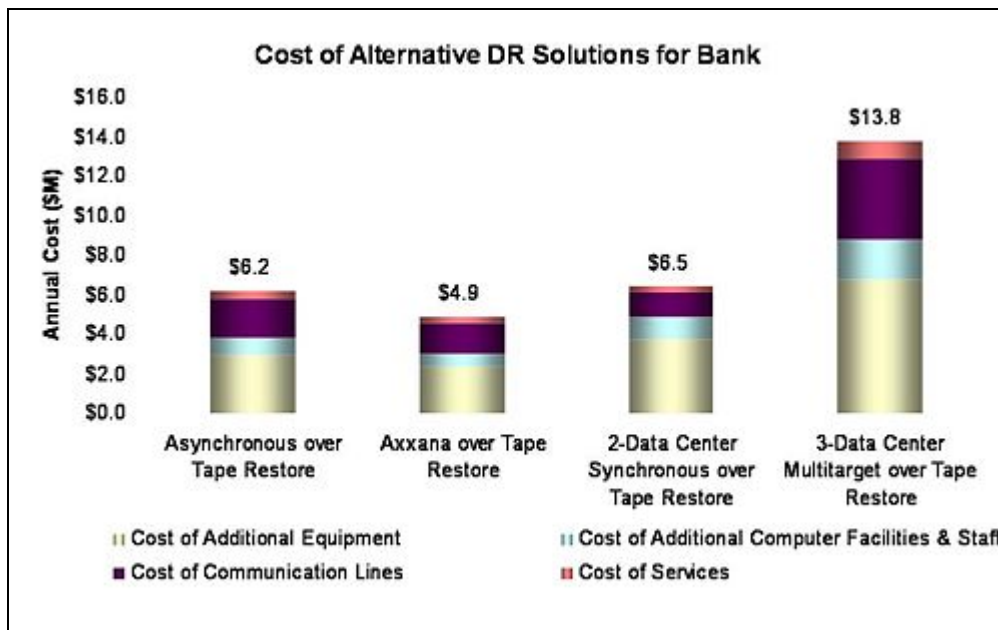


Figure 2 - Cost of Alternative DR Solutions for Bank

As seen in Figure 2, Axxana's solution has a lower cost of ownership than alternative disk-based DR solutions. Our analysis for this specific example shows the following:

- Costs for Axxana's solution are approximately \$1.3M lower than those required to run asynchronous or synchronous data protection;
- Costs for Axxana are nearly \$9M lower than those required to run a 3-node data center solution.

In our assessment, Axxana will have the lowest cost of staff because the solution is simpler to install, test and manage. For example, an Axxana approach reduces the amount of equipment needed to be managed. In a 2-Data center solution, exact copies of servers are required in two sites and in a 3-Data center approach, three sets of servers are needed. In addition, the cost of lines is lower for Axxana over asynchronous distances because of a reduced peak bandwidth requirement. Two data center solutions require dark fibre over shorter distances, increasing costs.

EMC RecoverPoint was assumed for the Axxana solution and EMC SRDF or HDS TrueCopy was assumed for alternatives, meaning fewer copies of data were required for Axxana. The footnotes provide additional detail about the inputs and assumptions used for the model.

Summary and Conclusions from Model

The question the model attempts to address is: **Relative to advanced disk-based DR solutions, how does Axxana fare?** From the case study above and other analysis using the model, the following key points are highlighted:

- Axxana's EDR approach decreases the expected loss relative to asynchronous solutions;
- Axxana's solution provides risk reduction substantially similar to both synchronous and 3-node data center approaches at a much lower cost;
- As a result, despite the higher costs for disk-based DR solutions, for environments with high data value the ROI of all these solutions is evident.
- The incremental CAPEX and OPEX of Axxana's solution is much lower than alternatives.

The bottom line is Axxana's approach appears to substantially cut the cost of achieving near-zero data loss and can do so at asynchronous distances, dramatically decreasing infrastructure costs relative to 3-node data centers and reducing proximity risk as commonly seen in synchronous operations.

In configurations above 20 TB's, our findings indicate that Axxana's solution almost always provides a more attractive cost and benefit relative to alternative disk-based platforms. Newness of the platform notwithstanding (Axxana went into Beta in early 2009), data center managers should, we believe, investigate this technology and identify applications that are candidates.

Caveats

Wikibon analysts have extensive experience in assessing the economic value of disaster recovery solutions. Our experts have studied this issue for more than a decade and have constructed dozens of models to support large financial institutions and a variety of cross-industry organizations. We have done so in both mainframe and non-mainframe environments and studied virtually every vendor's solution in this space.

While we are excited and encouraged by Axxana's innovative approach, the EDR solution only began Beta shipments in January, 2009. We do not have the field data to confirm that our model's expected results will actually be achieved.

Axxana is a startup and must prove to us and the world that it can execute on its vision of providing high quality disaster recovery solutions at substantially reduced operational costs. There are several hurdles Axxana faces in this regard including product stability, channel uptake, the ability to evolve its product and customers willingness to fit the solution into their business processes, or potentially alter processes to fit the solution.

Nonetheless, on balance we are impressed with the Axxana management team, we feel they are capable of securing the continued funding necessary to execute and have the wherewithal to deliver on the company's vision.

Action Item

The EDR technology is potentially a game-changer for high-value applications sensitive to data loss. IT and Business executives responsible for business continuity should make themselves familiar with Axxana's technology, and implement a pilot. Wikibon is happy to assist users in evaluating the technical and business case for change. As always, don't hesitate to [contact us](#).

Footnotes: The following table shows the main inputs and assumptions made in the model:

Inputs and Assumptions	Value
Number of application terabytes at primary site (excluding Replication)	250
Average IOPS for core application system	20,000
Peak IOPS for core application system	50,000
Peak to average IO rate	2.5
Average blocksize for application Tbytes	16.4
Degradation of transmission speed/100 miles distance	5%
OC-3 Lines required for asynchronous connection	23
Number of fibre links for synchronous connections	7
Time buffer for remote transmission errors at main site (mins)	0.5
Distance between primary site and local replication site (synchronous)	30
Distance between primary site and remote replication site (asynchronous)	350
Telecommunication Cost (Price per DSO/mile/month)	\$0.010
Probability of disaster causing a systems outage in a year	5%
Organizational Revenue (\$M) for core applications	\$ 400
Initial revenue loss if data is lost (% revenue)	2%
Revenue loss per hours if data is lost (% revenue)	1%
Loss/ hour RTO	\$ 10,000

Table 1 - Inputs and Assumptions

Retrieved from http://wikibon.org/The_economic_value_of_Axxana%27s_zero_data_loss_solution