the Availability Digest

www.availabilitydigest.com
@availabilitydig

Why Your Business Continuity Plan May be Inadequate

September 2014
Keith B. Evans, Shadowbase Business Development & Product Management
Paul J. Holenstein, Executive Vice President, Gravic

Disaster Recovery is Not Business Continuity

In today's business world, access to real-time online transactional data is a competitive advantage. To realize the advantage, this data must be available at any time, all the time, from anywhere, and it must be current. The corollary to this advantage is that the inability to access or update this current data carries a significant business cost, possibly measured in many thousands of dollars *per second*. These requirements necessitate an application service that is continuously available, in other words an *IT infrastructure* that is continuously available, and an adequate business continuity plan in place to assure application service continuity with access to data under both planned and unplanned circumstances.

Stuff Happens

Whether it be fire, power failure, software error, malfeasance, or some other cause, the fact is that events will occur that lead to unplanned outages of IT services. It is a matter of when, not if. Studies[1] show that the average business revenue lost per hour of downtime across a range of industry segments is about US$1.4M per hour. The U.S. Bureau of Labor reports that 93% of companies that suffer a significant data

[1] Network Computing, The Meta Group, Contingency Planning Research

loss are out of business within five years. Outages will ultimately happen, and they can be very damaging (even fatal) to the business. Consequently, for those critical IT services necessary for the business to function, steps must be taken in advance to ensure availability of those services no matter what the cause or duration of the outage.

HP NonStop systems – more so than many other platforms – and the mission-critical applications that run on them, must have a business continuity plan in place. NonStop systems are highly fault-tolerant, but they still represent a single point of failure. Hence, there is a need for a business continuity plan to enable operations to survive, despite the loss of a NonStop system or an entire datacenter. Such plans typically include multiple geographically distributed NonStop systems with at least some form of online data replication between them. The question is, are these plans adequate? While you may think so, that belief could be based more on hope than on reality. A recent survey[2] reports some disturbing results:

- Only 32% believe they utilize all best practices in datacenter design and redundancy to maximize availability.
- Only 37% agree there are ample resources to bring their datacenter up and running if there is an unplanned outage.
- 59% agree that the risk of an unplanned outage has increased as a result of cost constraints inside the datacenter.
- 57% believe all or most unplanned outages could have been prevented.

These findings, which illustrate that not enough attention and resources are being applied to outage prevention, are borne out by the fact that 95% of the respondents have experienced a complete datacenter outage, with an average number of 1.24 outages per year and an average duration of 107 minutes.

A study conducted by IBM[3] finds that perceptions of the business continuity plan often differ from reality, with 82% of respondents confident or very confident about their level of outage protection, yet only 65% have 24x7 expert technical support coverage. This same study also found that only 78% perform regular failover testing, and only 67% have a fully documented disaster recovery plan.

While everyone acknowledges that outages do happen, are costly, and need to be protected against, there is substantial evidence that IT departments are not applying sufficient resources to business continuity in practice (even though they might think otherwise). The first lesson is to take a thorough and objective look at your business continuity plans, asking if they are adequate and will they work, or do you just hope they will?

Business Continuity Technologies Are Not All Equal

In implementing a business continuity plan, there are a range of technologies available which provide differing levels of protection, from magnetic tape backup to active/active data replication (Figure 1). Key metrics for defining recovery solutions are how long will recovery take, or the Recovery Time Objective (RTO), and how much data will be lost, or the Recovery Point Objective (RPO).[4]

[2] Ponemon Institute 2010 National Survey on Data Center Outages. Data from 453 IT operations managers across a wide-range of industry segments (Financial, Healthcare, Retail, Communications, Services, etc.).
[3] 2012 IBM Global Reputational Risk and IT Study
[4] See Chapter 6, RPO and RTO, *Breaking the Availability Barrier: Survivable Systems for Enterprise Computing,* AuthorHouse: 2004

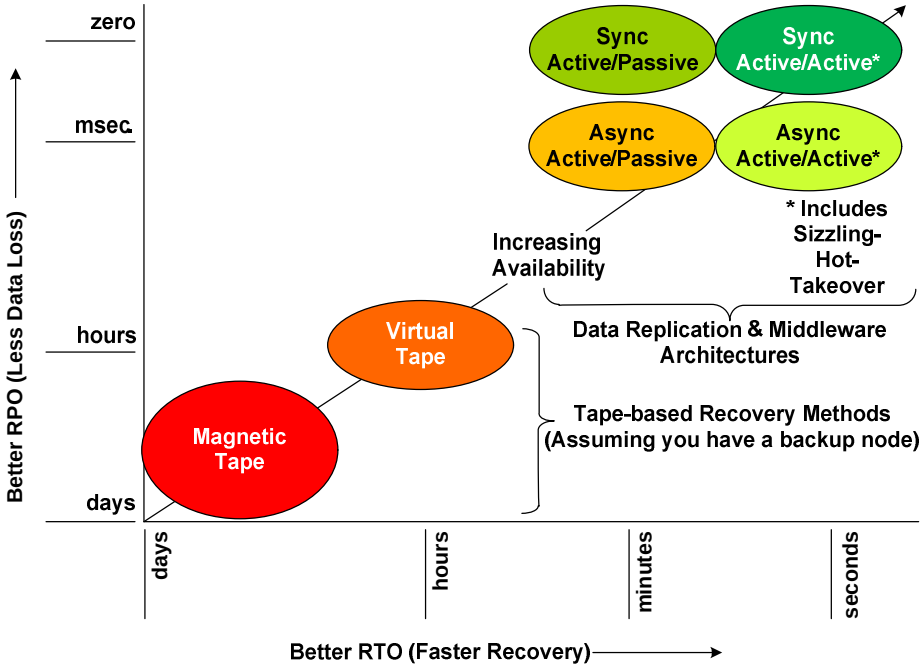

Figure 1: The Business Continuity Technology Continuum

Figure 2 shows some estimated RTO times and costs based on the business continuity technology employed. This table clearly demonstrates that tape-based solutions are insufficient for the purposes of providing adequate availability to mission-critical applications. But what this table also shows is that active/passive style data replication may also be inadequate. This inadequacy bears more explanation.

Technology	RTO	Outage Cost
Magnetic Tape Backup	~ 24 hours (optimistic)	~ $36M
Virtual Tape Backup	~ 12 hours	~ $18M
Active/Passive	~ 3 hours (if at all)[1]	~ $4.5M
Active/Passive	~ 10 minutes[2]	~ $250K
Sizzling-hot	~ 30 seconds[3]	~ $12.5K
Active/Active	~ 30 seconds	~ $6.25K[4]

[1] Worst case: with failover faults, management indecision, etc.
[2] Best case: with no failover faults, prompt management action, etc.
[3] Possibly slightly longer depending on network switching
[4] Half of users see no outage at all (less than half if > 2 replicated nodes)

Figure 2: Estimated Outage Times and Costs by Business Continuity Technology (Financial Application, Average Outage Cost $1.5M/Hour[5])

Active/passive business continuity architectures describe multiple geographically distributed systems, in which one system is active (being used to process online business transactions), and data from that

[5] Network Computing, The Meta Group, Contingency Planning Research

system is replicated to remote standby systems in near real-time. Replication is uni-directional, one-way, from the active to the standby system. The standby systems are not running mission-critical online applications; they may be used for ad-hoc query and other non-update type services. In ideal circumstances, this architecture may seem to provide adequate protection against service outages, but there are many potential issues that make it a less than satisfactory solution:

- **Difficult to test.** In order to test a failover plan, the active system must typically be taken out of service and workload transferred to the standby system (i.e., service to end users is disrupted). Because the standby system is not running the business applications at the time of the takeover (i.e., it is not a known-working system), it is possible it will take several hours before it can be brought into service. Once upon a time there may have been an overnight or weekend maintenance outage window where this length of application outage was acceptable, but in today's always-on world, this outage duration is increasingly not the case. Even if such a window does exist, it is not always possible to complete the testing within that timeframe. When the testing period is over, there is also the risk that the active system may not be able to be brought back online in time. For all these reasons, very often failover plans have not been sufficiently tested, and when they are actually needed, the failover does not go smoothly (so-called "failover faults" occur), and restoring service takes much longer than expected.

- **Management indecision.** Because there is an uncertainty as to whether the failover will be successful, senior management is usually required to authorize the action (as opposed to trying to restore the failed active system, if that is possible). Locating the necessary management personnel and apprising them of the situation and having them reach a decision takes time, further prolonging the outage.

- **All users are affected.** When an outage of the active system occurs, all users are denied service until either a failover is effected or the active system is restored.

- **More data loss at failover.** Because all of the updates are being performed on one system, if that system fails, then all of the data in the replication stream that has not been sent to the standby system will be lost (known as the "replication latency").[6]

- **Standby database open read-only.** Even if the business applications are actually up and running on the standby system (but not processing transactions), the database may only be opened read-only. Hence, when the failover occurs, all of the applications must be somehow notified and the database reopened for read-write access. This process complicates application programming, and can be time consuming, extending the outage.

- **Standby database inconsistent.** While replication is occurring, the standby database may be inconsistent ("fuzzy"), which could limit utilization of the standby system for query processing. This inconsistency will happen, for example, if the replication engine does not preserve the source application's transaction boundaries when replaying the data into the standby database.

Due to these issues, recovery times for an active/passive system may be in the order of several hours, costing millions of dollars (Figure 2). Worse, if a serious failover fault occurs, it is possible that the standby system may never be able to be brought into service; the mission-critical application is down and stays down, denying service to users for a prolonged period. This system is insufficient protection for a mission-critical application.

[6] See Chapter 3, Asynchronous Replication, *Breaking the Availability Barrier: Survivable Systems for Enterprise Computing,* AuthorHouse: 2004

Some Technologies Are "More Equal" than Others

There are however alternative business continuity technologies which may be deployed today that do not suffer from these issues. The first of these technologies is known as "sizzling-hot-standby." This technology looks much the same as an active/passive architecture (all transactions are routed to and executed by a primary system, with data replication to a standby system), but it has one big difference – the standby system is "hot." The business applications are all up and running on the standby system with the database open for read-write access; the only difference between it and the active system is that it is not processing online transactions that update the database (it can be processing read-only queries). Sizzling-hot-standby has several important benefits:

- **It greatly reduces RISK.** When a primary outage does occur, failover will be to a known-working standby system with a running application, thereby obviating failover faults. It also removes management indecision issues since the standby system is known to be operational.

- **It greatly improves RTO.** The application is already running, in full read/write mode, on the standby system. It is ready to receive user requests at any time. There is no need to waste the time to bring the application up for processing.

- **It simplifies testing.** A feature of sizzling-hot-standby is that because the applications are hot and the database open for read-write access, it can be tested, end-to-end, at any time even while the production system is in full operation. To verify the end-to-end operation of the standby system, occasionally send it a verification test update transaction. There is no need to take an outage of the active system and worry whether the standby system will come up, or that the testing will cause damage to the production environment.

- **Standby database is consistent.** Replication products that support standby applications opening the database read/write typically maintain transactional database consistency, so there are no issues with using the standby system for query processing.

- **Easier to recover the failed system.** Although all updates are being executed by one system, bi-directional replication is in place between both systems. When the failed system is restored, it is straightforward to recover it and bring the databases back into synchronization.

Overall, a sizzling-hot-standby architecture improves RTO and failover reliability significantly, decreasing recovery times and outage costs substantially (Figure 2). But it does still suffer from the fact that all users are affected when a primary system outage occurs, and causes more data loss than fully active/active architectures. Nevertheless, this architecture represents an excellent solution when the application cannot run in full active/active mode for some reason, and it is no more complex to implement than an active/passive architecture.

Next we turn to active/active architectures. In an active/active configuration there are two or more geographically separated systems, each running online business transactions and updating their local copy of the database, with data replication occurring between each system. Replication is bi-directional, meaning two-way between each active system.

Note that both systems are using replicated copies of the same database, and are running the same applications, with the transaction workload apportioned between them. As shown in Figure 1, active/active solutions provide the absolute fastest takeover times (RTO), with minimal data loss, because only half the data in the replication pipeline is lost in an outage of one system. Recovery times are measured in seconds, and because half of the users see no outage at all, outage costs are half those of the sizzling-hot-standby architecture (Figure 2).

If the sizzling-hot-standby and fully active/active business continuity technologies offer such great benefits versus active/passive architectures, why doesn't everyone use them? Good question. There are really no

additional complexities or limitations with a sizzling-hot-standby architecture. It is just an incremental extension of the active/passive model, which needs a replication product that allows the standby database to be open for read/write access and can be configured for bi-directional replication.

Active/active solutions on the other hand can suffer from complexities which do not arise in active/passive or sizzling-hot-takeover modes. Principal among these complexities is the possibility of data collisions. Because the same logical database is being updated on multiple nodes, and the same business applications are executing on those nodes, it is possible for a transaction to be executed simultaneously on each system, which updates the same record in the database. When that change is replicated to the other system, each will overwrite its update with that from the other system, and consequently both databases will be incorrect.

There are two potential solutions to this problem. The first is to avoid the possibility of data collisions altogether, which can be done by partitioning either the data or the applications, with transactions routed to the appropriate system, such that the same record will never be updated on both systems at the same time. For example, transactions for customer data records with names A-M are executed by one system, and those for names N-Z by the other system. One downside of this approach is that not all business services are amenable to partitioning in this way; the other is that workload may not be evenly distributed between each system, under-utilizing capacity and affecting response times.

The second solution is to route the requests to either system based on load (the so-called “route anywhere” model) and detect and reconcile any data collisions which do occur. Data replication solutions which support active/active modes generally include automated mechanisms for detecting data collisions, and resolving them using pre-defined rules (e.g., the transaction update with the more recent timestamp wins). This approach does not suffer from the workload distribution issue, but may not be feasible where there is no easy way to automatically resolve the collision (or where collisions cannot be tolerated by the application at all).

But what of those business services where application or data partitioning is not possible, and data collisions and/or loss of any data cannot be tolerated? Up until now this discussion has been all about asynchronous replication, where the replication engine sends data to the standby system asynchronously from the updates made by the application. In this mode, data can be lost (the replication latency mentioned above), and data collisions can occur (in active/active route anywhere architectures).

Synchronous replication resolves all of these issues. With synchronous replication, application data updates are not committed (made visible and permanent) by either system unless the updated data has been replicated to the standby system. This technology guarantees that no data is lost in the event of an outage of the system performing the update (known as “zero data loss,” or ZDL).

Additionally, in an active/active environment, it is not possible for data collisions to occur because the updated data records are locked on both systems before any changes are committed on either system. The same simultaneous update situation is instead manifested as a transaction deadlock, which is easily resolved via application programming using timeout and retry/resubmission logic similar to any other application processing error requiring a request resubmission. There is never any visible data inconsistency. Therefore, synchronous replication further reduces outage costs by avoiding any data loss, and by eliminating data collisions, opening up the benefits of active/active architectures to any application. It is the pinnacle of business continuity replication solutions.

For comparison, Figure 3 gives a summary of the most significant characteristics of each of the various replication architectures discussed.

Attribute	Asynchronous Active/Passive	Synchronous Active/Passive	Sizzling-Hot-Takeover	Asynchronous Active/Active	Synchronous Active/Active
Failover Faults	Yes	Yes	No	No	No
Application Outage	Yes	Yes	Minimal[1]	No	No
Data Loss	Yes	None	Yes[2]	Yes	None
Data/Request Partitioning	Not required[3]	Not required[3]	Not required[3]	May be required	Not required
Data Collisions	Not possible	Not possible	Not possible	Possible	Not possible
Backup Utilized	No[4]	No[4]	No[4]	Yes	Yes

[1] All users affected, but takeover time same as for Active/Active modes
[2] "None" if run in Synchronous mode
[3] "Required" if run in Reciprocal mode

Figure 3: Replication Technologies – Pros and Cons

We have discussed the sizzling-hot-standby and fully active/active architectures that offer greater protection than an asynchronous active/passive architecture, but are they cost effective? Yes, because the recovery times are generally so much better. With sizzling-hot-standby and active/active configurations, the outage costs and application risk are greatly reduced (Figure 2). While there may be some incremental cost increases with these solutions (e.g., the price of replication software licenses used in these higher availability modes), any such increases are more than offset by the significant reductions in outage costs, resulting in lower overall total cost of ownership (TCO) for these solutions (Figure 4).

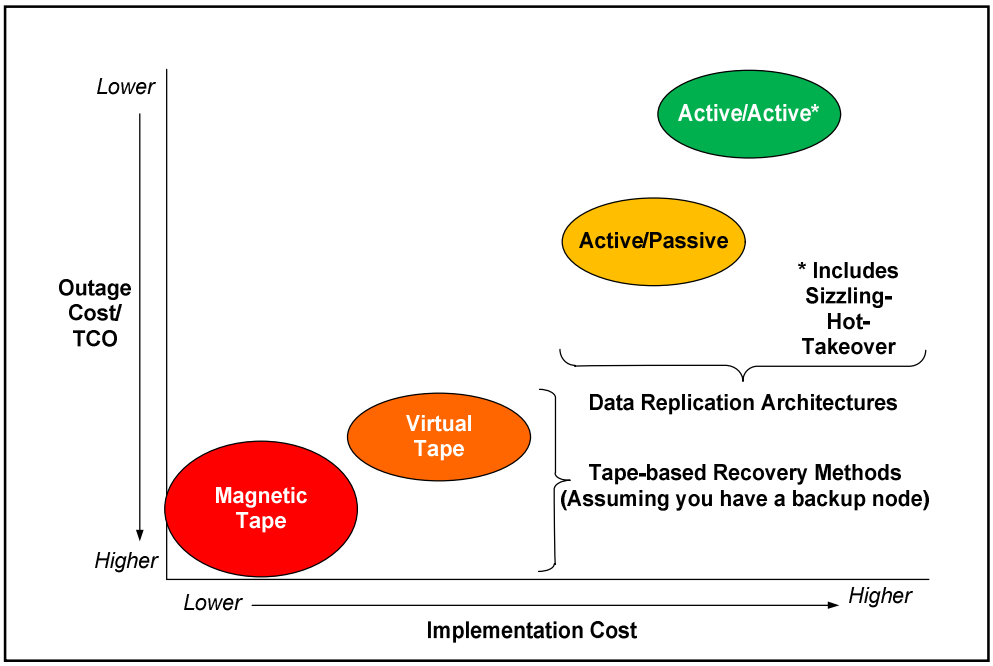

Figure 4: Business Continuity Technology and TCO

Time for a Reassessment?

Even though you may already have a business continuity plan in place, it may not be adequate, well-tested, or well-supported. Worse, it may be providing you with a false sense of security, and will fail when called upon. If this plan relies on an active/passive replication architecture, there are significant issues with this approach which could hamper a fast and successful takeover in the event of an outage. The key point is that you can avoid this risk, since there are other replication technologies readily available, such as sizzling-hot-standby and active/active architectures, which mitigate the issues with active/passive, and with better TCO. Further, for the highest levels of availability with no data collisions and zero data loss, synchronous replication may be utilized. If your business is relying on an active/passive architecture for service continuity, take another look at whether or not it really provides a sufficient guarantee of availability. It may now be time to consider moving to one of the other higher level replication architectures.

Keith B. Evans began his professional life in 1981 as a software engineer at IBM UK Laboratories, developing the CICS application server. He then moved to Digital Equipment Corp. as a pre-sales specialist. In 1988 he took a position at Amdahl in Silicon Valley as a software architect, working on transaction processing middleware for its UTS (Unix) O/S. In 1992 Mr. Evans joined Tandem and was the lead architect for its open TP application server program (NonStop Tuxedo). After the Tandem mergers he became a Distinguished Technologist with HP NonStop Enterprise Division (NED) and was involved with the continuing development of middleware application infrastructures (Pathway, J2EE, SOA, et al). In 2006 he moved into a Product Manager position at NED, responsible for middleware products. In 2010 he also became Product Manager for the NonStop Business Continuity product suite (TMF, RDF, AutoTMF, AutoSync, et al). Mr. Evans joined Gravic in 2012 and focuses on Shadowbase business development and product management for the Shadowbase synchronous replication products. He earned a BSc (Honors) in Combined Sciences from DeMontfort University, Leicester, England.

Paul J. Holenstein *is Executive Vice President of Gravic, Inc. He is responsible for the Shadowbase suite of products. The Shadowbase replication engine is a high-speed, unidirectional and bidirectional, homogeneous and heterogeneous data replication engine that moves data updates between enterprise systems in fractions of a second. It also provides capabilities to integrate disparate operational application information into real-time business intelligence systems. Shadowbase Total Replication Solutions® provides products to leverage this technology with proven implementations. For further information regarding Shadowbase data integration and application integration capabilities that can assist in solving big data integration problems, please refer to the companion documents Shadowbase Streams for Data Integration and Shadowbase Streams for Application Integration, or visit www.Gravic.com/Shadowbase for more information. To contact the author, please email: SBProductManagement@gravic.com.*