Dramatically Reduce Outage Costs with Advanced Business Continuity Solutions

Keith B. Evans >> **Shadowbase Product Management** >> Gravic, Inc.

Disaster Recovery is Not Business Continuity

In today's business world, consistent 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, or the loss of data, carries a significant business cost, possibly measured in many thousands of dollars per second, or even lives lost. In some cases, absolutely no data loss nor application downtime can be tolerated. 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 current and complete 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 which 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. The U.S. Bureau of Labor reports that 93% of companies that suffer a significant data 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 and the data they depend on no matter the cause or duration of the outage.

HPE 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 36% believe they utilize all best practices in datacenter design and redundancy to maximize availability.
- Only 38% agree there are ample resources to bring their datacenter up and running if there is an unplanned outage.
- 68% agree that availability has been sacrificed to improve efficiency or reduce costs.
- 71% believe at least some 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 all of the respondents have experienced a complete datacenter outage, with an average of one outage per year and an average duration of 91 minutes.[3]

A study conducted by IBM[4] finds that perceptions of the business

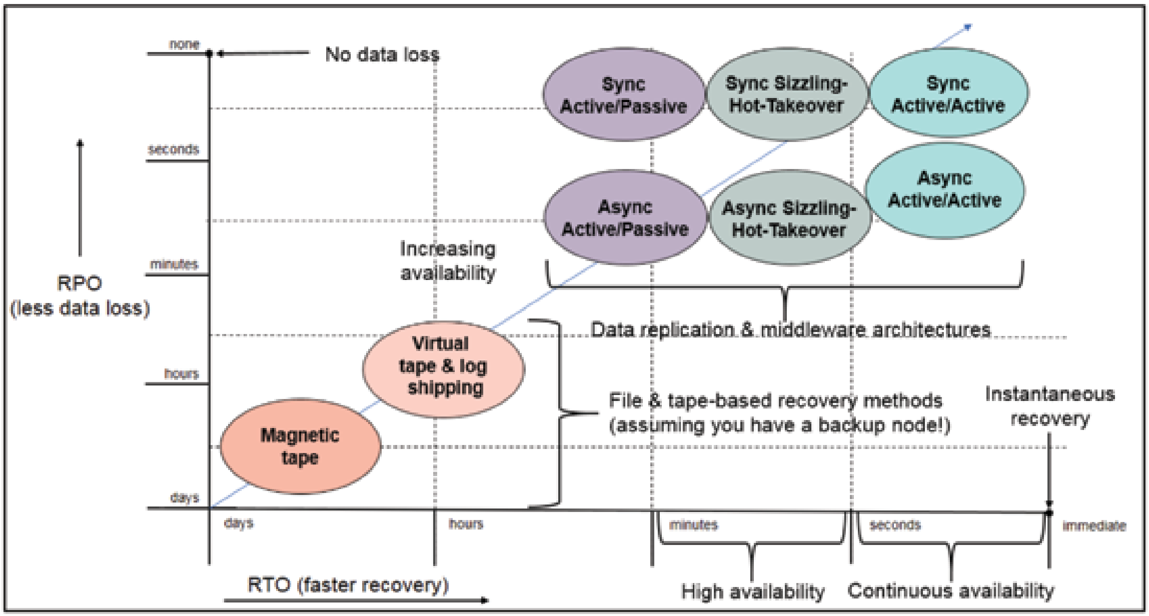

Figure 1 – The Business Continuity Technology Continuum

[1] Network Computing, The Meta Group, Contingency Planning Research
[2] Ponemon Institute, Cost of Data Center Outages
[3] Ponemon Institute, Study on Datacenter Outages
[4] IBM Global Reputational Risk and IT Study

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?

Not All Business Continuity Solutions Are Created Equal

In implementing a business continuity plan, there are a range of solution architectures and technologies available which provide differing levels of protection, from magnetic tape backup to active/active data replication (Figure 1). Key metrics for evaluating recovery solutions are: one, how long will recovery take, or the Recovery Time Objective (RTO), and two, how much data will be lost, or the Recovery Point Objective (RPO).[5]

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. The table also shows that active/passive data replication architectures are inadequate; this inadequacy bears more explanation.

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 system is replicated to remote standby (passive) 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 reporting, or for 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 an unsatisfactory 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., application services to the end users are 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 as operations fail-back to the original system. 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 fail-over action (as opposed to trying to restore the failed active system, if that is possible). Locating the necessary management personnel, 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.** Along with the unavailability of services, data loss accounts for the majority of the costs associated with unplanned downtime. In an active/passive architecture, 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 successfully delivered to the standby system will be lost (known as the replication latency).[7] This amount of data loss is far more than will occur in the most advanced architectures.
- **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

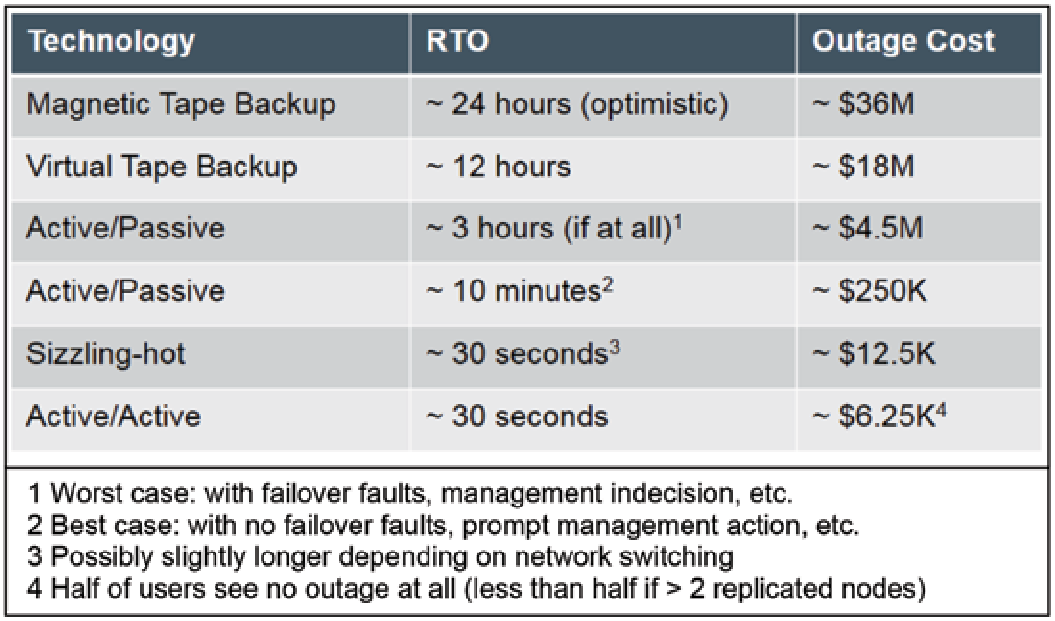

Technology	RTO	Outage Cost
Magnetic Tape Backup	~ 24 hours (optimistic)	~ $36M
Virtual Tape Backup	~ 12 hours	~ $18M
Active/Passive	~ 3 hours (if at all)[1]	~ $4.5M
Active/Passive	~ 10 minutes[2]	~ $250K
Sizzling-hot	~ 30 seconds[3]	~ $12.5K
Active/Active	~ 30 seconds	~ $6.25K[4]

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] See Chapter 6, RPO and RTO, Breaking the Availability Barrier: Survivable Systems for Enterprise Computing, AuthorHouse: 2004

of the applications must be somehow notified (or restarted) 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 will often be on the order of several hours, and data loss may be significant, resulting in outage costs of 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. An active/passive architecture is therefore insufficient protection for a mission-critical application.

But Some Business Continuity Solutions Are "More Equal" than Others

However, there are alternative business continuity solution architectures and technologies which may be deployed today that do not suffer from these issues; the first is known as sizzling-hot-takeover (SZT). This architecture 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 (passive) 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). An SZT architecture 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. No delay is required to bring the application up for processing.
- **It simplifies testing.** A feature of SZT 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. Taking an outage of the active system is not needed, so there is no concern whether the standby system will come up or the testing will cause damage to the production environment.
- **The standby database is consistent.** Replication products that support standby applications opening the database read/write typically maintain transactional database consistency, so there are no data consistency issues with using the standby system for query processing.
- **It is 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, an SZT 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 incurs 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 not more complex to implement than an active/passive architecture.

Application Availability – It Doesn't Get Any Better Than This

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 (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 (RPO), because only half the data in the replication pipeline is lost in an outage of one system. Recovery times are measured in seconds to sub-seconds, and because half of the users see no impact at all, outage costs are half those of the active/passive and SZT architectures (Figure 2).

If the SZT and fully active/active business continuity technologies offer such great benefits versus active/passive architectures, why doesn't everyone use them? Good question. Compared with active/passive, there are really no additional complexities or limitations with an SZT 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. An SZT architecture should be considered the absolute minimum configuration for mission-critical applications.

Active/active solutions on the other hand can suffer from complexities which do not arise in active/passive or SZT 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 each copy of 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 the 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 subsequently 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, which are resolved using pre-defined rules (e.g., the transaction update with

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

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

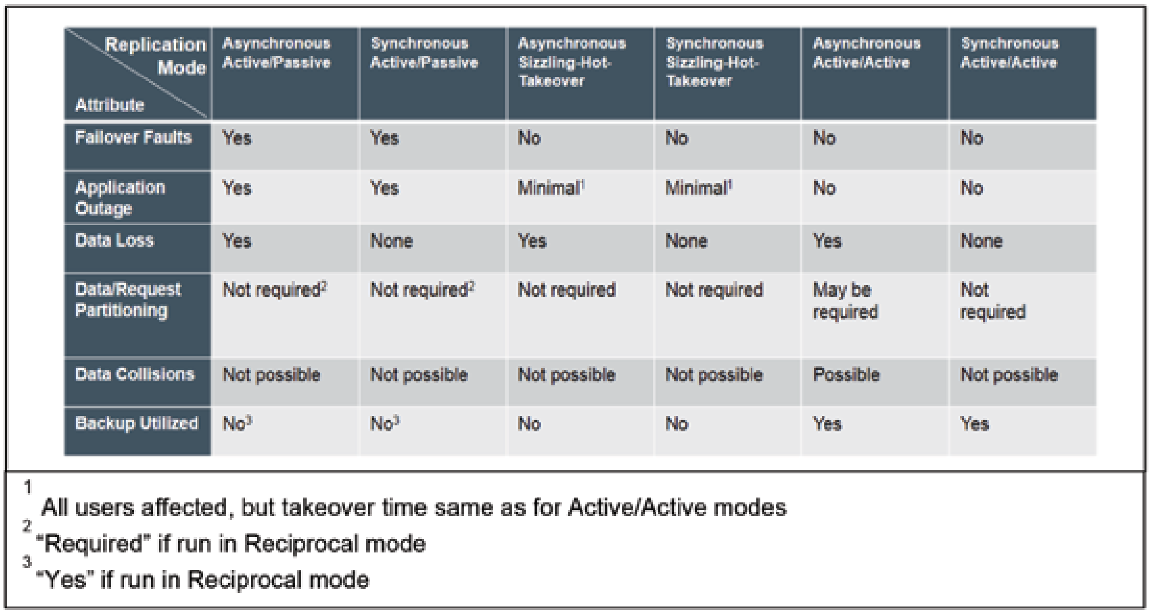

Replication Mode / Attribute	Asynchronous Active/Passive	Synchronous Active/Passive	Asynchronous Sizzling-Hot-Takeover	Synchronous Sizzling-Hot-Takeover	Asynchronous Active/Active	Synchronous Active/Active
Failover Faults	Yes	Yes	No	No	No	No
Application Outage	Yes	Yes	Minimal[1]	Minimal[1]	No	No
Data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

[1] All users affected, but takeover time same as for Active/Active modes

[2] "Required" if run in Reciprocal mode

[3] "Yes" if run in Reciprocal mode

Figure 3 – Pros and Cons of Replication Technologies and Architectures

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).

One Business Continuity Solution to Rule Them All – Synchronous Replication!

But what is necessary for 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 database updates made by the application. In this mode, when a failure occurs data can be lost and data collisions can occur in active/active route anywhere architectures, during the replication latency interval mentioned above. Synchronous replication resolves all of these issues. It is the business continuity solution which provides the greatest protection against the many impacts and costs of unplanned outages.

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 replication 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). Hence the costs arising from data loss simply do not occur with synchronous replication.

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 (caused by a distributed lock collision), which is easily resolved by the data replication engine (one lock/transaction wins, the other loses and should be resubmitted by the application similar to any other error that the application receives that requires request resubmission). There is never any visible data inconsistency.

In summary 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.

Time for 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 solutions readily available, such as SZT 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 (a new release, ***HPE Shadowbase ZDL,*** is now available which supports synchronous replication and zero data loss). If your business is relying on an active/passive or asynchronous solution for business continuity, take another look at whether or not it really provides a sufficient guarantee of protection against the impacts and costs of downtime and data loss. Chances are that it doesn't, and now is the time to consider moving to one of the other higher levels of business continuity solution.

Mr. Evans earned a BSc (Honors) in Combined Sciences from DeMontfort University, England. He began his professional life as a software engineer at IBM UK Laboratories, developing the CICS application server. He then moved to Digital Equipment Corporation as a pre-sales specialist. In 1988, he emigrated to the U.S. and took a position at Amdahl in Silicon Valley as a software architect, working on transaction processing middleware. 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. In 2006, he moved into a Product Manager position at NED, responsible for middleware and business continuity software. Mr. Evans joined the Shadowbase Products Group in 2012, working to develop the HPE and Gravic partnership, internal processes, marketing communications, and the Shadowbase product roadmap (in response to business and customer requirements). A particular area of focus is the newly patented Shadowbase synchronous replication technology for zero data loss (ZDL) and data collision avoidance in active/active architectures.