

# **NextGen 360° Advanced Business Continuity™ White Paper**

---

## **The BIA's 10 Fatal Flaws**



---

**The goal of a traditional BIA is sound in theory—to identify and prioritize the business’ critical processes relative to the losses that would result from their interruption, and then, to convince management to fund the most cost- and effective architecture to recover those processes in the aftermath of a disaster.**

However, in practice, the BIA process repeatedly fails to produce the desired results and in virtually every instance we have seen, it delays and complicates the recovery planning process. We believe that there are at least ten fatal flaws in the traditional BIA process.

1. Invented for the wrong reason
2. Losses are not linear
3. Probability is not predictable
4. Criticality is not weighted
5. Criticality is not aggregated
6. The wrong solution curve
7. The business will make up the difference
8. RTO is a dinosaur
9. The devil is in the details
10. The static report

---

## **Invented for the wrong reason**

The BIA was invented for the wrong reason...a self-serving reason. And unfortunately, that heritage has survived until today and still thrives.

In the earliest days of the Disaster Recovery industry, sooner or later, the commercial hotsite providers were asked the same question by every prospective client. “How can we recovery our systems in your hotsite...it is so much smaller than our environment”. Clearly a credible answer was needed. That answer was “because you will obviously not recover everything...only your most critical systems”. The very next question was predictable. “How do we define what is critical?” Hence the BIA was born as the prescribed way to identify the critical system subset so that the

commercial hotsites of the day were “big enough to meet the need”.

30+ years later, BIAs are still being used to define the smallest IT recovery footprint in a mistaken effort to minimize Disaster Recovery costs and complexity. Unfortunately, the smallest DR footprint results in the largest Business Continuity overhead and needlessly complicates the total continuity equation.

Instead of seeking the smallest footprint, a good BIA (assuming there is such a thing) should identify the optimal footprint, which in turn will support the simplest and most efficient BC capability.

---

## **Losses are not linear**

Expecting business leaders to determine the financial losses associated with process interruption within today’s large, distributed organizations with their hugely complicated and intertwined business environments is an exercise in futility. Few business leaders are able to accurately estimate losses associated with process interruption, but when pressed by the BIA process for a “guesstimate”, those leaders will provide an answer. In the best cases, their answers will duplicate many of the same loss elements that other departments included in their guesstimates. In the worst cases, their answers will be totally discounted by management and will cast doubt on the whole analysis. Even if we were to assume that losses could be estimated perfectly, the process would be impractical at best. Let’s say we could estimate exactly when an interruption will irritate our best customer versus when it will cost us an order from that customer versus when it will cause us to permanently lose that customer (using a retail metaphor). Let’s further assume that we can predict exactly when the loss will occur and exactly how much it will cost us...to the penny. The next step would be to estimate the loss for our second best customer and then the third best and so on. For each customer, the answer would be different based on that customer’s particular attitudes, needs, priorities, etc. And even if we estimated perfectly for each and every customer, our analysis

has not begun to address: different durations of interruption; whether the supply chain has been impacted too; whether the event caused larger, regional ramifications; variable impact based on time of occurrence; and a host of other variables that can dramatically change the loss profile based on the timing and nature of the event. Clearly, a different metric and analytic process is needed to define recovery priorities and a different lens needs to be used to garner management's commitment.

---

## Probability is not predictable

The next problem with the traditional BIA process is that even if the loss estimates are assumed to be perfectly accurate and are implicitly accepted by management, the very next question is always: "OK, if these are the losses we can expect, what is the likelihood that they will occur?" As soon as probability enters the picture, the battle has been lost. Predicting probability is an exercise in futility. There are literally thousands of variables with nearly infinite permutations that can affect the probability of the event and its resulting impact. If there weren't, you would simply plan to avoid the situation altogether. Risk analysis is a process to identify risks and prioritize their remediation. And as such, it is a valuable process. However, there has never been a risk analysis that has statistically justified disaster recovery planning...although there have been many that have been used as an excuse for avoiding or delaying planning efforts. But even if the odds of a disaster striking were accurately determined to be exactly 1 in 1,000,000, the laws of probability ensure that the one chance might occur tomorrow as easily as a million years from now. Again, a different way of pragmatically justifying recovery planning must be found!

---

## Criticality is not weighted

The fourth flaw comes from the inability of traditional BIA methodologies to weight and differentiate process criticality. Most methodologies prioritize processes according to a scale (1 –

5, low – medium – high, critical – important – deferrable, etc.). Regardless of the chosen scale, the problem arises when multiple impact categories are considered (legal, financial, image, health and safety, upstream, downstream, etc.) and the BIA methodology fails to provide a technique to weight criticality (compare the criticality of one process to another's) across different impact categories and/or across different priority tiers. Assume the 1 – 5 scale is used...is a process with a level 4 financial impact, a level 4 legal impact, and a level 3 image impact more or less important than one with a level 4 financial impact, a level 4 legal impact, a level 2 image impact and a level 2 downstream impact? Few methodologies help you make this decision (fyi...the second process is the most critical). Conversely, is a process with a level 3 financial impact, a level 3 legal impact, a level 3 image impact and a level 2 downstream impact more or less important than a process with a level 4 operational impact? (the level 4 operational impact is more important) Few BIA methodologies address this challenge and most default to labeling the process according to its highest criticality, disregarding same-tier or cross-tier weighting. This overly simple approach sacrifices granularity and as a result, it sacrifices the ability to optimally distinguish and adjust real priorities on-the-fly at-time-of-disaster based on the actual impacts of the moment. In reality, criticality should not be weighted...until it should be!

---

## Criticality is not aggregated

The next flaw is a variation of the previous flaw. BIAs do not aggregate process criticality across various interdependencies. That is, they do not provide a methodology to dynamically adjust a process' criticality based on unplanned, but nevertheless present interdependencies that change the assumptions that were used when determining the process' inherent criticality.

One such interdependency is location. Interruption of a process at a single location is one thing. Loss

of that same process across all locations is a different thing entirely. Few events would interrupt a process across all of its locations, and those events are usually substantially mitigated. As such, the traditional BIA will typically determine criticality assuming that only one (or a few) locations have been impacted. However, what if despite best efforts, all locations have been interrupted? A traditional BIA will not provide the ability to adjust the process' inherent criticality on-the-fly to accommodate the unlikely, but nevertheless manifested increased impact, nor will it offer any guidance on how the increased criticality effects other processes and resource allocation.

Another example of interdependency is the shared dependency on a common resource. Say for example, three processes all depend on Application ABC as a critical resource. Planning efforts have "ensured" that Application ABC will not be interrupted...it has been implemented with continuous availability and synchronous replication. As such, the three processes will never be interrupted concurrently, an assumption taken into account when defining their inherent criticality. However, now assume that Application ABC has been interrupted...despite best efforts its data has been corrupted across all instances. As a result, all three processes have been interrupted...a theoretical impossibility. The traditional BIA will do nothing to help reprioritize the three processes' unplanned concurrent interruption, nothing to help you decide how their priority must change relative to other processes, and nothing to recommend how their resource needs should be reprioritized given their "new criticality". In the traditional BIA, criticality is not aggregated...but it must be!

---

## The wrong solution curve

This flaw manifests in sub-optimal and/or overly expensive recovery architectures. Most BIAs gather information based on the "Sweet Spot" premise. The "Sweet Spot" premise was formulated in the DR/BC industry's early days and was illustrated as two opposite curves plotted

against a vertical axis of solution costs and a horizontal axis of recovery time. The first curve sloped smoothly from the high left (greater cost for faster recovery) to the low right (lower cost for slower recovery). The second curve rose from the low left (little loss from short outages) to the high right (great cost from long outages). The "Sweet Spot" was where the two curves intersected and was intended to represent the optimal balance of the cost of prevention relative to the cost of outage. At first, this concept appears logical. However, in reality, neither the cost of prevention or the cost of outage are curves at all. They are stair cases and very uneven stair cases at that. The difference is that in the real world, as you move up the cost scale or along the time scale those moves are not smooth and gradual. They are abrupt and dramatic. For example, going from 3-day recovery to sub 24-hour recovery is not an incremental cost of a few percentage points. It is a significant jump that might represent 3, 4 or 5 times more cost and complexity. Unless the BIA methodology specifically recognizes and gathers this "outlier" data, the warnings necessary for optimal solution modeling will not be available. For example, there will be no warning that a process should be revisited to determine if its RTO could be delayed 24 hours to reduce costs and/or simplify recovery. Conversely, there will be no warning that a process' RTO should be shortened to "guarantee" that it is achievable so that the unacceptable losses of a missed RTO are avoided.

---

## The business will make up the difference

The next problem comes from a fundamental flaw in thinking that the BIA process perpetuates. By focusing on business processes in an attempt to place responsibility for defining criticality and budget with the business owners (as opposed to IT), the BIA infers that "process recovery" is the goal. In fact, there are few if any critical business processes in today's organizations which stand alone and which can be recovered or "continued" manually. In virtually all cases, the automated

applications which support the critical processes must be recovered in order for the process itself to continue. This is the BIA's seventh fatal flaw. Business owners rarely understand the applications that they use in sufficient detail to provide the information that is necessary for disaster recovery planning. They do not understand application inter-dependencies, they do not know which physical assets their applications run on or which platforms must cross-communicate, they do not understand the network requirements for minimal connectivity and they usually do not have knowledge of the physical nature of the files that contain their data or of how those files are backed-up or rotated. The result is that the BIA process cannot be truly complete without Applications and Operations involvement, which is usually outside the original project's scope and which usually, requires a second project once the primary BIA is completed—which in turn further delays proactive planning efforts. To avoid this outcome, there is a tendency to assume that through a superhuman manual effort the business will make-up for the technology shortfall. This is seldom the case. More often, well-intentioned but naive assumptions about the viability of manual alternatives will actually exacerbate the situation and complicate recovery.

---

## RTO is a dinosaur

Who would have thought when the terms RTO (Recovery Time Objective) and RPO (Recovery Point Objective) were coined, that 30+ years later, the industry as a whole would still be arguing about what they meant, the “certifying bodies” would still be debating about when their clocks start and end, and the standards organizations would still be on a mission to invent new terms to better explain the concepts, while actually exacerbating the confusion (i.e. MTPOD).

Now enter the BIA. A good BIA (once again, if there is such a thing) will attempt to define both Application RTOs and Process RTOs. A valid effort as far as it goes. However, in today's business environment of ever-increasing competition and

cost-containment; continually more demanding availability requirements; increasing natural and man-made disruptions; and never-ending automation challenges...the effort doesn't go far enough. The traditional terms do not answer today's questions. They do not contribute to optimal planning and prevention efforts. And they do not facilitate the critical decision making and trade-offs that must be brokered at-time-of-disaster to enable the optimal recovery effort given the specific event's impact.

In addition to the traditional RTO, there are at least eight additional “RTO Flavors” that are a mandatory part of any worthwhile BIA...or BIA Alternative.

- Process – Application RTO
- Minimum Application RTO
- Cumulative Application Criticality
- Application Recovery Group
- Application Recovery Group RTO
- Application Recovery Sequence
- Process RPC
- Application RPC

The key is that your BIA methodology must identify and define these “RTO Flavors” without any more effort and with indisputable accuracy. Unfortunately, most BIA's do not even acknowledge these concepts. If your BIA doesn't provide these “flavors”, it's probably time to conduct a new BIA.

---

## The devil is in the details

The DR/BC industry is on a constant crusade to Simplify! The traditional BIA is never far behind. Unfortunately, the crusade is founded on a false belief...that DR/BC planning can be simplified by reducing or eliminating the details. We believe this grail search is naïve at best, and potentially catastrophic at worst. We believe that the “devil is in the details” and that artificially simplifying those details puts your recoverability at risk. A thorough, traditional BIA might define RTO, RPO, recovery tier, recovery group, recovery impact and, if you are lucky, maybe a couple of other data points. Unfortunately, 70 or 80 data points are required for

an optimal recovery! (That's required, not nice-to-have.) Even if only 20 or 30 of those data points were required, where does the traditional BIA leave you? What do you do when your BIA doesn't define internal and external / upstream and downstream dependencies? Or, what happens when your BIA overlooked the priority changes and additional resources needed to address your once-per-year worst-case priority. And how would you conduct your process when your third-party inputs were interrupted if your BIA didn't take the time to identify viable alternative sources?

Instead of ignoring or avoiding the details, why not find a way to define them and manage them that takes no more effort but which produces a far superior result. Why not use a methodology that defines 10 or 12 times the data points of a traditional BIA in less than half the time with twice the accuracy?

Why not really simplify the BIA and concurrently simplify your recovery planning and your actual recovery?

---

## A static report

The final flaw may be the most critical. A BIA typically is a static report whose data reflects the business' needs at a specific point in time. Once the report is presented to management, it is usually put on the shelf to gather dust until the next refresh...two or three years from now. But gaining management's commitment and prioritizing recovery requirements for planning purposes is possibly the least important use for this data. Instead, this data must be kept evergreen and should be immediately available at-time-of-disaster to dynamically model cross-process dependencies to determine the ripple-effect of the event-specific impact; and, to "calculate" in real-time, the dynamic recovery plan to address the specific loss profile for the specific event given its specific damage and impact relative to the specific timing of the event while taking into account the then available tools and resources available for

mitigation and recovery efforts. And, if you are unlucky enough to have another disaster on the next day, at a different location, with a different impact profile, you should be able to model a completely different plan for a completely different mitigation and recovery effort...again in real-time.

The days of "all or nothing" plans and "smoking hole" disasters are long past. This dictates that the BIA must be a tool not a static report—regardless of how good that report might be. The tool must contain all of the recovery objectives, needs and resources correlated by location, department and process. It must be able to interactively illustrate the unique recovery priorities and sequences based on the actual loss from the specific event. It must be able to support decision making at time of disaster by clearly illustrating increasing impact relative to time and the corresponding degrading effectiveness of manual procedures. And it must be able to facilitate dynamic reassignment of resources based on current needs and priorities.

---

## **Only by understanding the true needs of the business and how applications and processes interact at the detail level can proportionate and cost-effective recovery strategies be designed.**

A traditional BIA is too often an artificial project that is intended simply to convince management to invest in business continuity planning by painting a picture of abstract risk and losses. The BIA needs to be redesigned to reveal more detailed process and application information that is mandatory to craft the most cost-effective and workable final solution. It must produce data that will enable the design of much more finely tuned recovery strategies which in turn will offer much better recoverability at a much lower price point...strategies that will take maximum advantage of existing resources and infrastructure and support the business process requirements in the most functional and cost-effective manner.

---

The William Travis Group is a dedicated disaster recovery, business continuity and all-risk incident management consulting firm that has been in business for over 25 years. The founder of WTG has been in the disaster recovery industry since its inception and developed many of the technologies and methodologies that represent the standard in the industry and the baseline for today's practitioner certification. Today, WTG's NextGen 360<sup>0</sup> ABC™ methodology offers a holistic All-Risk Incident Management approach that combines best practices in disaster recovery and business recovery planning with leading edge all-risk initiatives such as management succession planning, supply chain continuity, pandemic operations, manufacturing resource planning, production line continuity and other advanced continuity solution. WTG works with across all industries with organizations of all sizes, both public and private and guarantees its clients 100% satisfaction.

The William Travis Group can be contacted at 1827 Walden Office Square, Suite 220, Schaumburg, IL 60173  
■ Phone 847-303-0055 ■ fax 847-303-0378 ■ [www.williamtravisgroup.com](http://www.williamtravisgroup.com) ■

