

Designing High to Low Cost Solution for Crash Recovery of Servers through Virtualization using Cloud Computing

G. Rama Rao, J. Purna Prakash and M. Rama Raju

Abstract--- Crash recovery is a process of recovering the data, plus all the hardware lost in the crash, and returning it to a reasonably functional state so that the business organization can run. It is important to understand the proper way to use the cloud for crash recovery and how to further low costs and how to shorten the time it takes to get back to business work after the crash. The key is to choose one from the plenty of options for crash recovery for specific business need.

Keywords--- RPO, RTO, Virtualized

I. INTRODUCTION

Fast forward 50 years to today's "always-on" world. The flow of information and commerce in our global business environment never sleeps. With the demands of an around-the-clock world, organizations need to start thinking in terms of application continuity rather than infrequent crashes, and crash recovery service providers need to enable more seamless, nearly instantaneous failover and failback of critical business applications. Yet given the reality that most IT budgets are flat or even reduced, these services must be provided without incurring significant upfront or ongoing expenditures. Cloud-based business resilience can provide an attractive alternative to traditional crash recovery, offering both the more rapid recovery time associated with a dedicated infrastructure and the Reduced costs that are

consistent with a shared recovery model. With pay-as-you-go pricing and the ability to scale up as conditions change, cloud computing can help organizations meet the expectations of today's frenetic, fast paced environment where IT demands continue to increase but budgets do not. This white paper discusses traditional approaches to crash recovery and describes how organizations can use cloud computing to help plan for both the mundane interruptions to service—cut power lines, server hardware failures and security breaches—as well as more-infrequent crashes. The paper provides key considerations when planning for the transition to cloud-based business resilience and in selecting your cloud partner. Traditional crash recovery-A choice between cost and speed When choosing a crash recovery approach, organizations have traditionally relied on the level of service required, as measured by two recovery objectives:

- Recovery time objective (RTO)—the amount of time between an outage and the restoration of operations
- Recovery point objective (RPO)—the point in time where data is restored and reflects the amount of data that will be ultimately lost during the recovery process.




Fig. 1: Measuring Level of Service Required by RPO and RTO

G. Rama Rao, Asst. Professor, IT Department, Christu Jyothi Institute of Technology & Science, Jangaon, India. E-mail: itsmerams@gmail.com

J. Purna Prakash, Asst. Professor, CSE Department, Christu Jyothi Institute of Technology & Science, Jangaon, India. E-mail: jpurnaprakash@gmail.com

M. Rama Raju, Asst. Professor, CSE Department, Christu Jyothi Institute of Technology & Science, Jangaon, India. E-mail: mrraju808@gmail.com

In traditional crash recovery models—dedicated and shared—organizations are forced to make the tradeoff between cost and speed to recovery, as illustrated in Fig 2

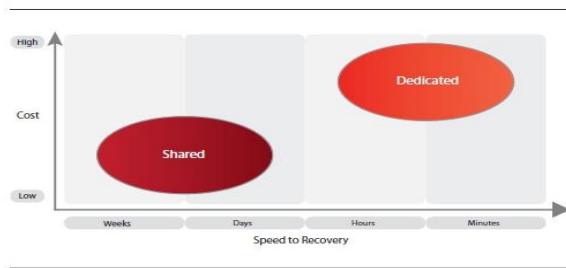


Fig. 2: Traditional Crash Recovery Approaches Include Shared and Dedicated Models

In a dedicated model, the infrastructure is dedicated to a single organization. This type of crash recovery can offer a faster time to recovery compared to other traditional models because the IT infrastructure is mirrored at the crash recovery site and is ready to be called upon in the event of a crash. While this model can reduce RTO because the hardware and software are preconfigured, it does not eliminate all delays. The process is still dependent on receiving a current data image, which involves transporting physical tapes and a data restoration process. This approach is also costly because the hardware sits idle when not being used for crash recovery. Some organizations use the backup infrastructure for development and test to mitigate the cost, but that introduces additional risk into the equation. Finally, the data restoration process adds variability into the process. As illustrated in Fig 3, data restoration can take up to 72 hours including the tape retrieval, travel and loading process.




Fig. 3: Time to Recovery Using a Dedicated Infrastructure

In a shared crash recovery model, the infrastructure is shared among multiple organizations. Shared crash recovery is designed to be more cost effective, since the off-site backup infrastructure is shared between multiple organizations. After a crash is declared, the hardware, operating system and application software at the crash site must be configured from the ground up to match the IT site that has declared a crash, and this process can take hours or even days. On top of that, the data restoration process must be completed as shown in Fig 4, resulting in an average of 48 to 72 hours to recovery.

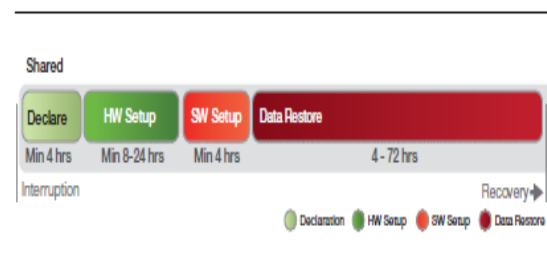


Fig. 4: Time to Recovery Using a Shared Infrastructure

A. *The Pressure for Continuous Availability*

According to organizations are being challenged to keep up with the growing demands on their IT departments while keeping their operations up and running a making them as efficient as possible. Their users and customers are becoming more sophisticated users of technology. Research shows that usage of Internet-connected devices is growing about 42 percent annually, giving clients and employees the ability to quickly access huge amounts of storage. In spite of the pressure to do more, they are spending a large percentage of their funds to maintain the infrastructure that they have today. They are also not getting many significant budget increases; budgets are essentially flat. With dedicated and shared crash recovery models, organizations have traditionally been forced to make tradeoffs between cost and speed. As the pressure to achieve continuous availability and reduce costs continues to increase, organizations can no longer accept tradeoffs. While crash recovery was originally intended for critical batch “back-

office” processes, many organizations are now dependent on real-time applications and their online presence as the primary interface to their customers. A downtime reflects directly on their brand image and interruption of key applications such as e-commerce, online banking and customer self service is viewed as unacceptable by customers. The cost of a minute of downtime may be thousands of dollars shown in Fig 5.

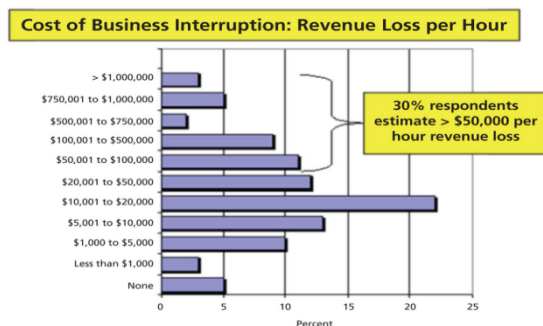


Fig. 5: Cost of Business Interruption: Revenue Loss Per Hour

B. Thinking in Terms of Interruptions and not Crashes

Traditional crash recovery methods also rely on “declaring a crash” in order to leverage the backup infrastructure during events such as hurricanes, tsunamis, floods or fires. However, most application availability interruptions are due to more mundane everyday occurrences. While organizations need to plan for the worst, they also must plan for the more likely—cut power lines, server hardware failures and security breaches. Fig 6 shows the kinds of disruptions has helped its customers respond to over the past few years. While weather is the root cause of just over half of the crashes declared, note that almost 50 percent of the declarations are due to other causes. These statistics are from clients who actually declared a crash. Think about all of the interruptions where a crash was not declared. In an around-the clock world, organizations must move beyond crash recovery and think in terms of

- More predictable monthly operating expenses can help you avoid the unexpected and hidden costs of do-it-yourself approaches

application continuity. You must plan for the recovery of critical business applications rather than infrequent, momentous crashes, and build resiliency plans accordingly.




Fig. 6 : Types of Different Business Interruptions

C. Cloud-Based Business Resilience—A Welcome, New Approach

Cloud computing offers an attractive alternative to traditional crash recovery. “The Cloud” is inherently a shared infrastructure: a pooled set of resources with the infrastructure cost distributed across everyone who contracts for the cloud service. This shared nature makes cloud an ideal model for crash recovery. Even when we broaden the definition of crash recovery to include more mundane service interruptions, the need for crash recovery resources is sporadic. Since all of the organizations relying on the cloud for backup and recovery are very unlikely to need the infrastructure at the same time, costs can be reduced and the cloud can speed recovery time.

Cloud-based business resilience managed services like Virtualized Server Recovery are designed to provide a balance of economical shared physical recovery with the speed of dedicated infrastructure. Because the server images and data are continuously replicated, recovery time can be reduced dramatically to less than an hour, and, in many cases, to minutes—or even seconds. However, the costs are more consistent with shared recovery

D. Cloud-Based Business Resilience Offers Several other Benefits over Traditional Crash Recovery Models

- Reduced up-front capital expenditure requirements, because the crash recovery infrastructure exists in the cloud

- Cloud-based business resilience managed services can more easily scale up based on changing conditions
- Portal access reduces the need to travel to the recovery site which can help save time and money.

While the cloud offers multiple benefits as a crash recovery platform, there are several key considerations when planning for the transition to cloud-based business resilience and in selecting your cloud partner. These include:

- Portal access with failover and failback capability
- Support for crash recovery testing
- Tiered service levels



- Support for mixed and virtualized server environments
- Global reach and local presence
- Migration from and coexistence with traditional crash recovery The next few sections describe these considerations in greater detail.

E. Facilitating Improved Control with Portal Access

Crash recovery has traditionally been an insurance policy that organizations hope not to use. In contrast, cloud-based business resilience can actually increase IT's ability to provide service continuity for key business applications. Since the cloud-based business resilience service can be accessed through a web portal, IT management and administrators gain a dashboard view to their organization's infrastructure. For example, clients can access the Smart Cloud Virtualized Server Recovery

portal via the Internet and identify their servers to be protected and replicated. Through this portal, customers can download the Smart Cloud Virtualized Server Recovery client software to install on their covered servers. Once the environment is defined through the portal, the client can view the protection status of their servers, generate reports and conduct other administrative tasks. While this administrative view through the portal is useful, it is critical to evaluate cloud-based business resilience services to help ensure that the portal is not merely an administrative configuration tool but that it also provides the opportunity to initiate a failover and failback. With Smart Cloud Virtualized Server Recovery, clients can use the portal to failover in near real time (for the "Always Available" service-level protected servers described later), reducing the need to contact the cloud service provider to "declare as crash" or to initiate the failover. Without the need for a formal declaration and the ability to fail over from the portal, IT can be much more responsive to the more mundane outages and interruptions described above

F. Building Confidence and Refining Crash Recovery Plans with more Frequent Testing

One traditional challenge of crash recovery is the lack of certainty that the planned solution will work when the time comes. Typically, organizations only test their failover and recovery on average once or twice per year, which is hardly sufficient, given the pace of change experienced by most IT departments. This lost sense of control has caused some organizations to bring crash recovery "in house," diverting critical IT focus for mainline application development. Cloud-based business resilience provides the opportunity for more control and more frequent and granular testing of crash recovery plans, even at the server or application level. Smart Cloud Virtualized Server recovery provides a DR Testing view in the portal so that IT can test the failover and failback process more frequently.

Clients can generally tailor testing to their schedule. For example, a critical e-Commerce application can be tested prior to a peak online shopping period such as

Cyber Monday. Or an online banking system can be tested after a version upgrade in order to ensure that the failover and failback process still works seamlessly

Table 1: Virtualized Server Recovery Service Levels

Smart Cloud virtualized server recovery service level	Recovery time Objective(until system boot start)	Description
Fast Recovery Always Available Virtual Machine	1 to 5 minute failover	For mission- critical applications that require near zero RTO/RPO and that need a recovery infrastructure with near-continuous availability
Medium Recovery Crash and Test Virtual machine	30-40 minutes failover	For applications that need repaid recovery infrastructure that is remotely accessible at the time of crash
Slow Recovery Imported Media Virtual machine	6 to 24 Hours or best effort basis, assisted failover, and assisted failback	For clients that want to import server images and data from tape/disc into our cloud infrastructure at time of crash for failover and failback provides migration path for infrastructure recovery service clients to cloud recovery service

G. Supporting Optimized Application Recovery Times with Tiered Service Levels

Cloud-based business resilience offers the opportunity for tiered service levels that enable you to differentiate applications based on their importance to the organization and the associated tolerance for downtime. For example, Smart Cloud Virtualized Server Recovery provides three service level options: Gold, Silver and Bronze. These tiers enable organizations to optimize their bud-get, paying more for their mission-critical applications to have nearly continuous availability and paying less for non-critical applications. With Smart Cloud Virtualized Server Recovery, the frequency of the data replication and the resulting RPO/RTO is based upon the service level assigned to the server. Multiple servers supporting the same application and business process can be collectively assigned the same group and service level to help ensure consistency and synchronization for failover and failback operations.

H. Efficiently Supporting Mixed Environments with Virtualized Crash Recovery

The notion of a “server image” is an important part of traditional crash recovery. As the complexity of IT departments has increased, including multiple server farms with possibly different operating systems and operating

system (OS) levels, the ability to respond to a crash or outage becomes more complex. Organizations are often forced to recover on different hardware, which can take longer and increase the possibility for errors and data loss. Organizations are implementing virtualization technologies in their data centers to help remove some of the underlying complexity and optimize infrastructure utilization. The number of virtual machines installed has been growing exponentially over the past several years. According to a recent survey of Chief Information Officers, 98 percent of respondents either had already implemented virtualization or had plans to implement it within the next 12 months. Cloud-based business resilience solutions must offer both physical-to-virtual (P2V) and virtual-to-virtual (V2V) recovery in order to support these types of environments. Smart Cloud Virtualized Server Recovery supports virtualized, non virtualized and mixed environments, including those with multiple operating systems.

I. Enabling Bandwidth Savings with a Local Presence

Cloud-based business resilience requires ongoing server replication, making network bandwidth an important consideration when adopting this approach. A global provider like offers the opportunity for a local presence, thereby reducing the distance that data must travel across the network. With SmartCloud Virtualized Server Recovery, the client’s server configuration, operating

system, application software and associated data are replicated to the Resiliency Center across the Internet or designated network connection. While data will be replicated to the closet Resiliency Center running SmartCloud Virtualized Server Recovery, added resiliency and backup is provided within the network of secure centers. SmartCloud Virtualized Server offers Recovery Synchronization and Bandwidth Estimator to assist with the assessment of network bandwidth requirements. While many of our clients will not need to increase capacity, the Estimator can confirm your capacity needs. Clients should identify all servers that support a single business application and include those servers in a single Virtualized Server Recovery plan. The solution assures cross-server consistency across those servers for failover and failback, helping to enhance security and reduce risk

J. Coexisting More Effectively with Traditional Crash Recovery

While cloud-based business resilience offers many advantages for mission-critical and customer-facing applications, an efficient enterprise-wide crash recovery plan will likely include a blend of traditional and cloud-based approaches. The Bronze level service of SmartCloud Virtualized Server Recovery can help ease the transition from traditional methods. Clients can also use this approach when integrating their cloud-based business resilience with data back-up solutions like SmartCloud Managed Backup. In a recent study, respondents indicated that minimizing data loss was the most important objective of a successful crash recovery solution. With coordinated crash recovery and data back-up, data loss can be reduced and reliability of data integrity improved

II. CONCLUSION

Cloud computing offers a compelling opportunity to realize the recovery time of dedicated crash recovery with the cost structure of shared crash recovery. However, crash recovery planning is not something that is taken lightly; security and resiliency of the cloud are critical

considerations. SmartCloud Virtualized Server Recovery is hosted within the network of Resiliency Centers—so clients can feel confident that is helping to protect their sensitive data. Second, there is no need to rush in—clients can start to work with SmartCloud Virtualized Server Recovery with as few as five virtual machines under managed contract—so getting started is easier and relatively risk free. With more than 1,800 dedicated business continuity professionals and more than 160 business resilience centers located around the world, respected industry analysts recognize as a leader in business continuity and resilience. Our virtually unparalleled experience is based on more than 50 years of business resilience and crash recovery experience and more than 9,000 crash recovery clients. Further, it has been in the systems business for 60 years, and just about no other company understands systems like does. Using our vast business process and technology expertise, we can help you design and implement a business resilience solution that meets your organization's needs

REFERENCE

- [1] Kresimir Popvoic and Zeljko Hocenski, "Cloud Computing Security Issues and Challenges", MIPRO, Opatijia, Croatia, May 24 -28, 2010
- [2] M. Armbrust, A. Fox, R. Griffith, A. Joseph, R. Katz, A. Konwinski, G. Lee, D. Patterson, A. Rabkin, I. Stoica, and M. Zaharia, "A View of Cloud Computing", Communication of the ACM, Vol. 53, No. 4, April 2010
- [3] CloudComputing vs.Virtualization <http://www.learncomputer.com/cloud>
- [4] W. Tsai, X. Sun, J. Balasooriya, "Service-Oriented Cloud Computing Architecture", 7th IEEE International Conference on Information Technology, 2010
- [5] V. Sarathy, P. Narayan, and R. Mikkilineni, "Next generation Cloud Computing Architecture", 2nd International IEEE Workshop On collaboration & Cloud Computing, 2010.
- [6] Y. Chen, X. Li and F. Chen, "Overview and Analysis of Cloud Computing Research and Application", International Conference on E-Business and E-Government (ICEE), May 2011



M.Rama Raju3. Jangaon, 19-04-1983. B.Tech in Computer Science and Engineering at Sreenidhi Institute of Science and Technology Completed in the year 2006. At Hyderabad, Andhra Pradesh, India. His M. Tech in Computer Science and Engineering at Vathsalya Institute of Science and Technology, Bhongir, Nalgonda, Andhra Pradesh, India. Currently Working as a Asst. Professor in Christu Jyothi Institute of Technology and Science 6 years Experience in Engineering Colleges. Previous publications with the book titles National conference on Advance Research in computing. New Delhi, Allied publishers in the year 2013. The Titles are 1. Proximity: A privacy measure on micro data for data publishing



J. Purna Prakash2. Jangaon, 15-01-1988. B.Tech in Computer Science and Engineering at Christu Jyothi Institute of Technology and Science Completed in the year 2009. At Jangaon, Warangal, Andhra Pradesh, India. His M. Tech in Computer Science and Engineering at Indur Institute of Engineering College, Siddipet, Medak ,Andhra Pradesh, India. Currently Working as a Asst. Professor in Christu Jyothi Institute of Technology and Science, 2 years Experience in Engineering Colleges. Previous publications with the book titles National conference on Advance Research in computing. New Delhi, Allied publishers in the year 2013. The Titles are 1. Novel two-stage Queuing Architecture in Heterogeneous Flow Networks.2. Network modification techniques on QOS Stability failures



G. Rama Rao1. Jangaon, B.Tech in Information Technology at SRKR Institute of Technology .At West Godavari, Andhra Pradesh, India. His M. Tech in Software Engineering at Andhra University, Vizag ,Andhra Pradesh, India. Currently Working as a Asst. Professor in Christu Jyothi Institute of Technology and Science, 6 years Experience in Engg Colleges. Previous publications with the book titles National conference on Advance Research in computing. New Delhi, Allied publishers in the year 2013. Title is Multipath switching system with load balancing flow slice finer granularity