

DATA DRIVEN GLOBAL VISION CLOUD PLATFORM STRATE  
ON POWERFUL RELEVANT PERFORMANCE SOLUTION CLO  
VIRTUAL BIG DATA SOLUTION ROI FLEXIBLE DATA DRIVEN

WHITE PAPER

# Achieve Continuous Computing for Mainframe Batch Processing

By Hitachi Data Systems and 21st Century Software

October 2016

## Contents

<b>Executive Summary</b> .....	<b>2</b>
<b>Introduction</b> .....	<b>3</b>
<b>Difficulties in Recovering Batch Jobs</b> .....	<b>4</b>
<b>The Solution</b> .....	<b>5</b>
<b>Hitachi Distance Replication</b> .....	<b>5</b>
<b>21st Century Software VFI</b> .....	<b>6</b>
<b>Customer Example</b> .....	<b>7</b>
<b>Major Credit Card Processing Service</b> .....	<b>7</b>
<b>Summary</b> .....	<b>7</b>
<b>Appendix A: Collaboration</b> .....	<b>8</b>

## Executive Summary

When downtime is not acceptable and rapid recovery from any failures is essential, fault-tolerant rapid recovery application environments are needed. With mainframes, this includes both transactional and batch processing. Applications and databases have advanced transaction recovery features, but batch jobs present a different problem.

## Introduction

Many corporations and institutions have moved to IT models where continuous processing is an absolute requirement. If systems go down or are taken down, then real revenue is lost, or worse. This is particularly true in the finance, transportation and retail sectors where distributed transactions are the norm, but it also applies wherever mainframes support mission-critical applications. For these businesses, the distinction between disaster and planned outage events is disappearing. Downtime is basically unacceptable, Recovery point and time objectives (RPO and RTO) requirements are well below four hours and often as short as 15 to 30 minutes.

Continuous availability starts with reliable hardware, and Hitachi Virtual Storage Platform G1x00 and F1500 (VSP G1000, VSP G1500, VSP F1500) provide the ideal storage foundation for critical mainframe systems. With full redundancy and superior engineering, VSP G1x00 and F1500 ensure your systems will remain up with performance to meet your most stringent requirements. Hitachi is the only company to provide a 100% uptime warranty to stand behind our storage. But reliable hardware is just the start.

Within mission-critical systems, significant work has been done on online transactional applications to achieve low-loss rapid-recovery goals at both the application and system levels. Organizations can achieve a continuous data protection environment through combinations of application and system enhancements, such as EMC Geographically Dispersed Disaster Restart (GDDR) or IBM Geographically Dispersed Parallel Sysplex™ (GDPS®). However, transactional applications are only part of the picture. What about batch jobs? For continuous computing in mainframe environments, continuous protection of batch processing systems can also be required.



In mainframe systems, critical information is delivered using complicated often-interdependent batch processing jobs, which need to run on very tight schedules to keep systems running. An unexpected failure can take you back to step zero, or worse leave data in an inconsistent state, requiring a more complicated recovery.

Even worse, many major problems happen when a small problem is incorrectly diagnosed and improperly fixed. Small, uncaught errors can turn into major system outage events.

Where is continuous data protection for batch jobs? Hitachi Data Systems and 21st Century Software have the answer.

### **Difficulties in Recovering Batch Jobs**

IT system protection has changed over the past decade. While business users press for secure, smoothly running systems, they're also wanting 24/7 access to data from anywhere, any time, on any device. This is pushing IT architecture away from traditional recovery and towards continuous data protection (CDP). Unlike older backup and recovery data protection architectures, CDP usually involves a strategy centered on frequent snap copies and application log datasets.

In this environment, when a failure occurs there is always the possibility of data or system corruption or problems caused by lack of synchronization. With databases, log datasets are used to roll back transactions to the most recent known good state. However, with running batch jobs there may be multiple open datasets that can be corrupt or in an unknown state. There may also be multiple interdependent batch jobs running. Manually analyzing job status and restoration of appropriate open datasets and job restarts is a laborious task requiring knowledge and time of one or more experts.

#### **What happens to batch jobs when they encounter...**

- Corruption?
- Accidental deletion?
- Out-of-space condition?
- Hardware failure?

To recover from events, application owners have to look at the scheduler and examine relevant job logs and journals to determine which datasets may be corrupt or are in an inconsistent state. If mistakes are made or something is overlooked, the time to recover can be significantly extended. Major outages occur when minor errors accumulate and compound and jobs are not properly restarted.

You should be able to determine where they are in batch processing and which data may be inconsistent and/or corrupt. Consistency points alone are not adequate to recover batch processes that were in flight at the time of disruption.

The complexity of solving this problem has led some companies to run failover tests during periods when batch jobs were not running. Their hopes were to ensure a "successful" test. However, neglecting to test system failover while batch processing is underway is a major failing that may come back to cost them.

Something better is needed.

Corrupted data, human error, and hardware failure account for 80% of disruptive events. A US\$5000 device failure causes revenue impact, customer goodwill and productivity loss. And the further back in time you have to go to recover or restore data, the more effort it is and the longer it takes to return to normal operations.

## The Solution

How can you achieve “database-like” service level agreements (SLAs) with batch processing? Since individual failures are inevitable, redundancy is required so outages do not affect operations. In addition to fault-tolerant storage, two other elements are central to continuous application solutions. The first is an efficient snapshot and recovery system for application state and data checkpoints and recovery. The second is a remote volume replication system supporting connected multisite data center operations (either synchronous or asynchronous depending on requirements). A good example of this for mainframe storage is Hitachi Virtual Storage Platform G1x00 and F1500 with Hitachi Compatible Mirroring for IBM® FlashCopy® or Hitachi Compatible Software for IBM FlashCopy SE and Hitachi Distance Replication software. When combined with 21st Century Software’s VFI line of batch application protection software products, continuous batch computing becomes possible.

### Hitachi Distance Replication

To provide continuous data availability in the event of site outages the snap volume needs to be mirrored remotely. So a critical component of a continuous-computing, rapid-recovery environment is efficient replication between data centers. Hitachi storage offers a unique solution. Available on VSP G1x00 and F1500, the Hitachi Distance Replication package includes Hitachi Universal Replicator (HUR) and Hitachi TrueCopy to provide versatile storage-system-based replication. Delivering both synchronous and asynchronous, low-latency performance, with key features, such as consistency groups, disk-based journals and multitarget or cascade configuration capabilities. As a result, the system executes replication of small, extremely large or heterogeneous data volumes quickly and efficiently, in 2-data-center and 3-data-center remote configurations for data protection disaster recovery, business continuity or data migration purposes.

#### Features

- Low latency, asynchronous operation requires less network bandwidth, which allows you to maximize network cost savings.
- Disk-based journal enables tight RPO time management with data consistency, which minimizes data-loss exposure.
- Support is provided for any storage connected to Hitachi Virtual Storage Platform G1x00 and F1500.
- Up to 2,064 consistency groups support multiple-volume, multiple-array data sets.

#### Benefits

- Simplifies implementation to meet the most demanding disaster recovery and uptime requirements, regardless of the type of supported storage platform hosting the business-critical data.
- Supports availability of up-to-date copies of data in dispersed locations.
- Maintains the integrity of a replicated copy, minimizing impact to production processing even when replication network outages occur or optimal bandwidth is not available.
- Enhances administrative productivity as well as crisis aversion and response.

To improve performance and simplify operations, it is possible to combine local IBM FlashCopy snapshots with mirrored remote replication. By combining compatible FlashCopy snapshots with Hitachi Universal Replicator or Hitachi TrueCopy,

applications such as 21 Century Software VFI can use compatible FlashCopy to copy updated datasets directly to the Universal Replicator or TrueCopy primary volume. Universal Replicator is used for asynchronous legs, and TrueCopy for synchronous. Then, with Universal Replicator the secondary copy is always consistent, so, in the event of a failover, it can be used directly. This contrasts with other products that cannot directly remotely replicate the target of a dataset-level flash over asynchronous connections.

VSP G1x00 and F1500 with Hitachi Storage Virtualization Operating System (SVOS) counts among many firsts, among which is its status as the first system to include mainframe storage in a virtualized environment. It offers the unique ability to provide native mainframe management of virtualized storage service level policies, replication and tiering. It also provides the unique ability to provide detailed SMF record-based, time-coherent performance reporting from attached mainframe storage, further simplifying mainframe operations.

HDS is the only vendor to provide comprehensive storage management from native mainframe environments. Other vendors can't provide native mainframe configuration, understanding and management of modern storage capabilities, such as virtualized, thin provisioned, protected and policy-controlled automatic tiered storage systems. Optionally integrated with virtualized open systems storage, the Hitachi solution supports the software-defined storage architecture that many companies seeking to implement.



## 21st Century Software VFI

Once in a snapshot and restore environment, traditional database applications by themselves do a good job of recovery. Recovering batch jobs, however, involves identifying the point in time that had been reached when the problem occurred. Research is then required to determine interdependencies before restarting batch processes. The ability to automate this process potentially avoids hours of time recovering from operational problems, particularly corruption.

VFI is an application that inventories your batch environment. It monitors and creates a database of job and database activity. It provides an inventory of all datasets in batch and virtual storage access method (VSAM) and when and how they are used. All information about file usage and status is stored in the VFI Unified Recovery Architecture database, and is used by VFI reporting, simulation and restore modules to ensure visibility and integration and allow you to take action with knowledge no matter what the event might be. This job-specific information in the database can be mirrored across sites to enable rapid remote recovery in the event of a site failover event.

The VFI TimeLiner component adds information into the database on when all job, steps, and datasets are opened and closed. Hardware consistency points don't help with application itself. TimeLiner monitors jobs through the supervisor call (SVC) activity and system exits and tracks and records by automatically taking snapshots of both VSAM and non-VSAM data at key points in time with the needed extra information for job step and dataset level recovery.

VFI TimeLiner and VFI InstaSnap give you point-in-time (PIT) copies at all job steps to recover from, as well as automating the location and recovery from those points in the event of a failure event. VFI includes a dependency mapping component, which, based on prior job runs, maps dependencies between jobs. This lets you avoid potential "toe stepping" if jobs are started without this consideration.

Online TimeLiner panels are used to identify PITs (rollback with TimeLiner, key by job name), get reports of all cascading jobs, and tell you what you need to restore, including which jobs and which datasets. Criteria includes where data is input or output. TimeLiner lets you look at existing batch jobs to see how to do it. It uses ISPF panels, these panels show you which jobs and datasets are open or running. The VFI InstaRestore feature then allows you to recover any job or dataset as required. You can then restart the jobs from that point.

## Customer Example

### Major Credit Card Processing Service

A major International credit card processing service faced the challenges of continuous data processing and availability. For their business, improving customer transaction times for customer facing applications is of the utmost importance. If a phone cannot be purchased, groceries bought, accounts accessed quickly, then customers walk out of the figurative and literal door to the competition. It only takes a minute to lose a customer and read about it later on social media. Also, because of its online nature, the credit card industry has many interlocking dependencies, so there are shared Payment Card Industry (PCI) mandated service level standards. Not meeting them can result in substantial fines and added processing fees.

With data centers in multiple countries, the credit card processing company had recently embarked on a major internal initiative to set a 15-minute recovery time objective (RTO) goal for all customer-facing systems. A major trouble spot was identified in the mainframe batch processing area. When a failure occurred, the existing processes involved assembling a team of specialists, including application owners to identify at what stage each job had failed and what states needed to be recovered and restored. It typically took more than 15 minutes just to assemble and activate the team.

To recover, application owners had to look at the scheduler to see what jobs were running when the failure event happened. They then needed to figure out interdependencies. Did all jobs fail at the same time? Over what time period? What datasets were open? The next problem was determining which datasets to recover and to restore them.

This led them to the automation of batch job recovery with VFI: Instead of recovery being a manual people intensive process, it was automated. Installation and setup was an easy process. And now the production control system programmer is able to simply make the determinations without relying on application teams or time-consuming analysis processes.

The solution enabled the company to meet its 15-minute RTO objective for their critical mainframe batch processing jobs. It simplified and improved their operations, and reduced dependencies on individual application owners. It also increased savings by avoiding downtime caused by simple mistakes cascading into larger problems.

## Summary

When downtime and continuous computing is required, fault-tolerant rapid recovery application environments are needed. With mainframes, this includes both transactional and batch processing. Applications and databases have advanced transaction recovery features, but batch jobs present a different problem. For continuous computing in mainframe environments continuous protection of batch processing systems can also be required.

Continuous availability starts with reliable hardware, and Hitachi Virtual Storage Platform G1x00 and F1500 provide the ideal storage foundation for critical mainframe systems. With full redundancy and protection while supporting advanced storage technologies, such as policy-based automated tiering and flash, VSP G1x00 and F1500 ensure your systems will remain up and with performance to meet your most stringent requirements.

For rapid recovery and restart batch jobs the 21<sup>st</sup> Century Software VFI product can provide advanced disaster recovery capabilities to complement those used for transactional systems. Combining the Hitachi and 21<sup>st</sup> Century solutions helps you achieve continuous computing with your business-critical mainframe systems, which translates into improved business competitiveness and increased bottom-line dollars.



## Appendix A: Collaboration

Hitachi Data Systems has published this white paper in association with:



21st Century Software, Inc.

940 West Valley Road

Wayne, Pennsylvania 19087-1823

800.555.6845

Office: 610.971.9946

Fax: 610.964.9072

Sales: [sales@21stcenturysoftware.com](mailto:sales@21stcenturysoftware.com)