

Checklist for Reducing the Impact of Backup and Recovery in Your Organization

DOWNTIME IS MONEY: FIVE WAYS TO MODERNIZE COPY DATA MANAGEMENT

The time it takes to perform a backup, the time to restore data, the time to recreate data that's been lost: This is all time that your organization's operations are impacted, costing you money. As the amount of your data continues to grow, and as your information technology infrastructure becomes more complex and distributed, these costs are likely spiraling out of control.

Reduce complexity, cost and risk with Hitachi Data Instance Director (HDID) and the Hitachi Virtual Storage Platform (VSP) family. Here's how:

❑ 1. ELIMINATE THE BACKUP WINDOW

Backup is performed at night and on weekends because production systems need to be stopped or paused to ensure all data is copied in a complete and consistent manner. For an incremental backup, it can take hours to scan through the system for the files that have changed since the last backup. For full backups, it can take hours or even days to copy all the data, especially large databases. In today's always-on world, how many applications can tolerate this amount of downtime?

HDS advantage: These times are eliminated when you perform a storage-based snapshot or clone on a VSP family system, automated and orchestrated by HDID. With this combination, the application is put into a backup-ready state without impacting performance or availability, and then HDID calls for the snapshot to be executed. Alternatively, use the host-based continuous data protection (CDP) feature in HDID to capture every new block of data as it is written to disk. CDP runs in parallel to the production system with no impact on performance or availability.

❑ 2. IMPROVE RECOVERY POINT OBJECTIVES (RPO) AND RECOVERY TIME OBJECTIVES (RTO) BY 95% OR MORE

Performing a backup once each night supports an RPO of 24 hours. This means that as much as 24 hours of your newest data is at risk of loss. Restoring backup data from disk, tape or cloud services can take hours or even days. This is not acceptable for important or critical applications and data.

HDS advantage: Since the application-consistent snapshot and clone operations of HDID and VSP do not impact production performance, they can be performed much more frequently. For example, creating a snapshot once per hour will improve your RPO by 95.8%; using CDP will drive it to nearly zero. Likewise, restoration from a snapshot, clone or off-site replica can be completed within seconds or minutes. Also, recovery with HDID is a single-step process, unlike restoring the last full backup and then each of the subsequent incremental backups. This approach greatly reduces the time to restore, as well as most of the risk of (human) error.

❑ 3. REDUCE THE AMOUNT OF DATA TO PROTECT BY 60% OR MORE

It is likely that most of the data in your production systems has not been accessed in a while. You probably want to keep that data for any number of reasons, but why are you continuing to create additional copies of it with each full backup cycle?

HDS advantage: HDID automatically moves inactive data, based on policies that you set, to Hitachi Content Platform (HCP), a self-protected, cost-effective object storage system. HCP applies the retention policies set in HDID, and fully indexes the files and emails sent by HDID, providing easy discovery. For example, an automated policy can move files that have not been accessed in 90 days. This capability could easily reduce the amount of capacity used in the production storage system by 60%. As new data is created, older data is moved to the archive. The less-expensive archive storage will grow instead of the more expensive primary storage.

❑ 4. REDUCE BACKUP STORAGE CAPACITY REQUIREMENTS BY MORE THAN 70%

When you perform a full (or synthetic full) backup, most of the data being copied is identical to the data that was backed up last time. If you perform an incremental backup each day and a full backup once per week, and retain each for 12 weeks, you could need more than 20TB of backup capacity for every 1TB of production data. Deduplication appliances were developed to correct this problem, but why should you pay extra to fix a problem caused by your backup software?

HDS advantage: HDID employs an incremental-forever data capture model, so each new or changed block of data is copied only once. Instead of duplicating most of the data every week, 1TB of production data will only need 3TB to 5TB of protection storage. And HDID also includes deduplication at no extra charge. Depending on the daily data change rate and the amount of duplicate data blocks across source systems, HDID can reduce backup storage requirements by 70% to 90% or more. Also, data captured by HDID can be re-used by other applications and departments: There is no need to create additional copies.

❑ 5. REDUCE ADMINISTRATION COSTS BY AT LEAST 50%

Most organizations have more than one data management tool to cover the range of operating systems, applications and locations in their IT environment. For example, do you have separate tools for databases, virtualized servers, remote offices and employee workstations? Or, do you keep separate tools to cover operational recovery, disaster recovery and long-term retention requirements? All of the technologies represented by these tools are important to ensure the availability of your data and the continuity of your business, but multiple point solutions lead to silos of secondary systems and administrative skill sets.

HDS advantage: HDID brings all of these technologies together in a single, fully integrated platform. And its unique, whiteboard-like user interface makes creating and managing complex data protection, retention and recovery workflows a breeze. You can quickly and easily assign a variety of capture, retention, replication and expiration policies to the same data set, and tie them together in graphical workflows. What would take days to script in legacy environments, if that were even possible, takes just a few minutes in HDID.

To learn more about data protection solutions from Hitachi Data Systems, please visit [HDS.com/go/protect](https://hds.com/go/protect).



¹ Hitachi Data Instance Director supports the hardware-based snapshot, clone and remote replication capabilities of VSP, VSP G series, VSP F series and Hitachi NAS Platform, with Microsoft® Windows®, Exchange and SQL Server®, as well as Oracle and SAP HANA platforms. HDID CDP works on Windows platforms, and its archiving to HCP supports Windows file systems and Exchange. Host-based batch backup with HDID can be performed on Windows, Linux, Oracle Solaris and IBM® AIX® file systems.

Hitachi Data Systems

Corporate Headquarters
2845 Lafayette Street
Santa Clara, CA 96050-2639 USA
www.HDS.com community.HDS.com

Regional Contact Information
Americas: +1 866 374 5822 or info@hds.com
Europe, Middle East and Africa: +44 (0) 1753 618000 or info.emea@hds.com
Asia Pacific: +852 3189 7900 or hds.marketing.apac@hds.com