Getting Started with

SnapShot

In Three days or less…

(Updated for Open Systems Support!)

Introduction

Purpose

The purpose of this guide is to help Information Technology executives and managers' gain an understanding of SnapShot, provide a fast-path implementation process, and recommend specific ideas for achieving its maximum value.

This guide will:

- ✓ Present a high level SnapShot product overview.
- ✓ Outline a pilot project scenario designed for quick implementation.
- ✓ Discuss IS activities that can benefit from SnapShot.
- ✓ Show how SnapShot can shave hours from your daily batch window.
- ✓ Share hints and tips.

Bob Eskenberry, DASD Marketing, Storage Solutions Specialist, Storage Technology Corporation wrote this guide, with additional contributions from Steven Curtis, Jerry Boezel, Chris Saul, Craig Welch and Diane Williams of IBM Corporation. All references to "Virtual Disk Array" include all StorageTek Iceberg, IBM RAMAC Virtual Array (RVA), and StorageTek Shared Virtual Array (SVA) subsystems, including the latest generation known as the SVA 9500. All of these machines harness the leading edge Virtual Disk Architecture pioneered by Storage Technology Corporation.

Preface

Information has quickly become the most valuable asset of many corporations. Emerging technologies, global markets, and a competitive marketplace have resulted in a data explosion. But the problem with this data explosion is not storing it; it's determining how to *leverage* it. There are three major components of leveraging data to produce information. You create data, you use it, and you duplicate it to protect its value. Two of these actions produce value for your business; the other is a painful but necessary overhead. ***SnapShot turns the time and resources used by traditional physical duplication methods into productive time and additional resources for your business***.

Many businesses are beginning to exploit the real potential of E-business, like providing products and services customized for the individual consumer. All the while customers keep demanding key on-line applications be available 24 hours a day, putting even more pressure on the clock. Whether the purpose is to keep overseas factories producing more goods, servicing more customers, or providing individualized products and services on the Web, increased access to information is critical for success.

Most data centers prioritize implementing new technologies based on their potential for solving their most critical challenges. For many, a top priority is reducing the batch window - converting unproductive time into productive business activities.

SnapShot works by creating different "views" of data rather than copying the data itself. By eliminating the need to physically move data for duplication, SnapShot creates copies, almost instantly. While SnapShot can be used in multiple ways to benefit your entire enterprise one key way is to improve application availability. With SnapShot's ultra-fast duplication capability, application backups can be removed from the application critical path, and run concurrently with other processing (fig.1). ***SnapShot significantly reduces CPU costs and channel utilization associated with data movement, and uses no additional physical storage to create the copy***.

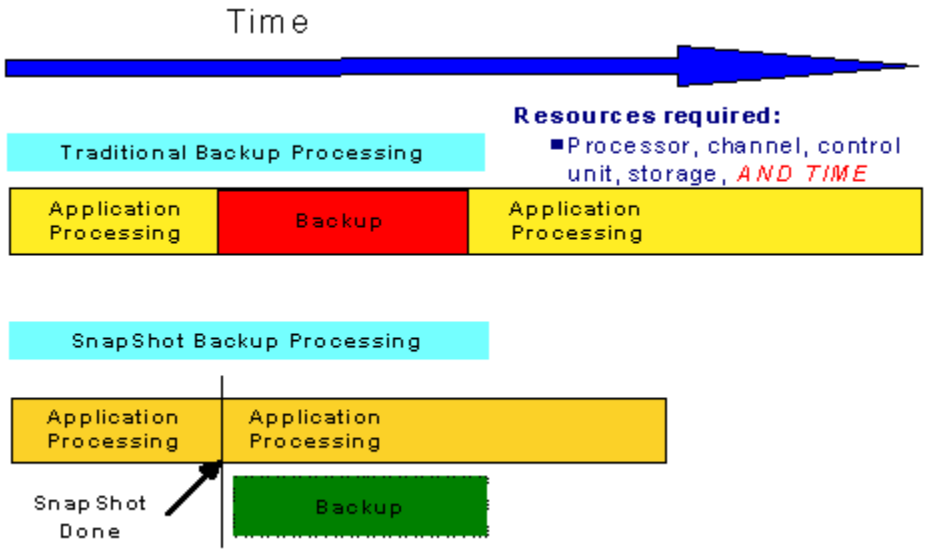

Fig. 1

Many sites worldwide have implemented SnapShot - reducing batch run times, increasing on-line availability and productivity. SnapShot offers a unique extended-lifetime solution, rather than the expensive and temporary relief of physical duplication. Speeding up batch processing by expanding bandwidth and using additional tape and disk systems, offer a short-term remedy that quickly gets eliminated with data growth. SnapShot-based solutions are largely impervious to this data growth.
Experience illustrates dramatic results are possible with SnapShot. An *independent survey* showed that more than 80% of customers gained over 3 hours of availability for business applications every day (fig. 2). The typical savings for a SnapShot customer was 3.8 hours per day, everyday.

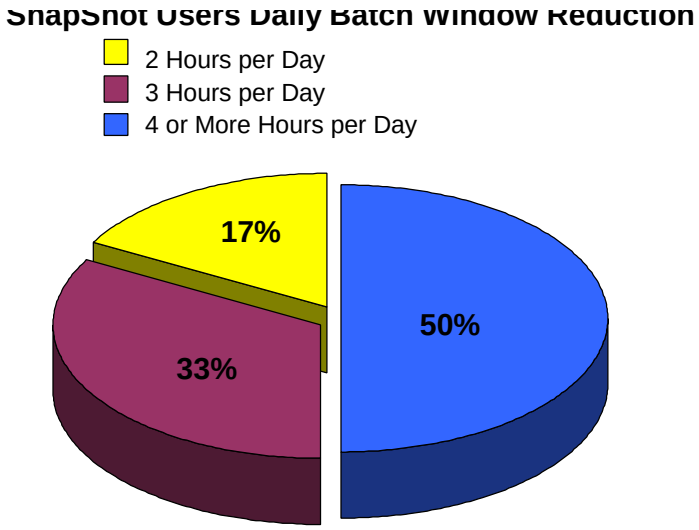

Fig 2

Experience also shows that the sooner SnapShot is implemented into your business applications, the sooner you can convert unproductive time (of physical duplication) into productive time.

SnapShot is a quantum leap forward in technology, enabling businesses to gain hours of availability for key applications every day.

Critical Success Factor

A key to faster implementation is to engage industry experts to find where SnapShot can provide the largest benefit for your specific enterprise. StorageTek provides tailored SnapShot planning, implementation and knowledge transfer professional services to get you going with SnapShot fast. Our experts will train your personnel on the technical aspects, work with your applications people to determine the best way to use SnapShot, and will test the implementation prior to turning it over to you to begin reaping the savings.

SnapShot Evolution

SnapShot, the cornerstone of the Virtual Power Suite™ of software, is ***the*** industry standard for fast-efficient data duplication. Other disk vendor claim "SnapShot-like" products, but cannot come close to the flexibility and efficiency of SnapShot, such as zero disk capacity for any number of copies you need, created virtually-instantaneously and ready to use ***now***.

Since SnapShot's introduction into the market there have been a number of significant enhancements and third party vendor partnerships that have magnified the value of SnapShot even further.

- ***Support for virtually all types of data set SnapShots.*** SnapShot supports VSAM, including extended format and Alternate Indexes. It also supports multi-volume data sets.

- ***Support via IBM DFSMSdss.*** If you are a user of DFSMSdss, you can gain the benefits of SnapShot and continue to use your same utilities. Furthermore this integration also enables a function called "Virtual Concurrent Copy". Virtual Concurrent Copy (or VCC) enables you to create non-disruptive image copies of DB2 and IMS databases. This means that you can take backups of key databases while users are online! By providing this capability, through the "CONCURRENT" keyword on the DB2 or IMS copy utility, you can now take backups more often and during prime time. By having additional backups, if a recovery is required, the time to get back in business is dramatically reduced…instead of forward-processing 24 hours of log data you will only have to process a fraction.

- ***Support via Innovation Data Processing "Instant Backup".*** The Instant Backup feature for FDR/ABR provides significant benefits. Simplify DR backup and recovery procedures with FDR. Gain the benefit of incremental backups using SnapShot with ABR. Run volume de-fragmentation in a fraction of the time with FastCPK. See information about IDP Instant Backup using SnapShot on the Web.

- ***Support for BMC Software Utilities.*** Utilities that utilize BMC XBM, now exploit the STK SnapShot capability to provide a superior method of getting data base backups and image copies. Read more about the BMC and STK alliance on the Web.

- ***Support for VSE/ESA.*** Using IBM software, VSE/ESA users can now benefit from the power of SnapShot. A new command "IXFP" with the sub-parameter SNAP provides this capability.

- ***Support for Open Systems platforms.*** Maybe the most exciting announcement for SnapShot is its capability to Snap Open Systems data. Whether e-commerce is key to your business today or in the future, SnapShot will help you overcome the dilemma between true 24x7 availability and complete backups. SnapShot is now supported SUN Solaris, and HP/UX platforms, with Windows 2000, AIX and other platforms quickly following.

And it doesn't stop there. There are more partnerships and enhancement to this powerful industry-leading duplication tool.

Understanding SnapShot

Virtual Disk Architecture

SnapShot is a *virtual duplication* product that uses the *unique* virtual disk architecture of the Virtual Disk Array invented by Storage Technology Corporation. To best understand SnapShot, a quick system review is needed. Virtual Disk Architecture automatically maps data to its disk arrays, as it is written, similar to database log files. Virtual Disk Architecture is also referred to as Log Structured File organization (LSF). This automation translates into large savings for the enterprise by reducing storage management costs, and dramatically improving access to information.

Virtual disk architecture manages software-mapping tables in the storage controller and disk arrays, to store and retrieve user information. These tables map host tracks to disk array storage locations where the data is stored. When data is written to the system, it is compressed and compacted, assembled into fixed blocks, and written to the RAID-6 disk arrays (fig 2). All write operations in virtual disk architecture are always directed to a new place in the disk arrays. Traditional disk array systems *reserve* disk array space equal to the size of each host volume, emulating the old big 14" head disk assemblies.

The unique architecture of the Virtual Disk Architecture provides flexibility, while at the same time maintaining a standard look and feel that the storage administrators are familiar with. In fact, Shared Virtual Arrays and RVA are so flexible you can define new host volumes or Open Systems Logical Units (LUNs) in just seconds, right from your desk. It allows you to configure up to 1024 host devices, or up to 2.89TB of "virtual" disk space, no matter what capacity system you have purchased. The capability to define more space than is physically installed creates a *virtual* storage space.

Some advantages of Virtual Disk Architecture include:

- Lowest cost of ownership
- Reduced storage management requirements
- SnapShot, which capitalizes on this unique virtual disk architecture

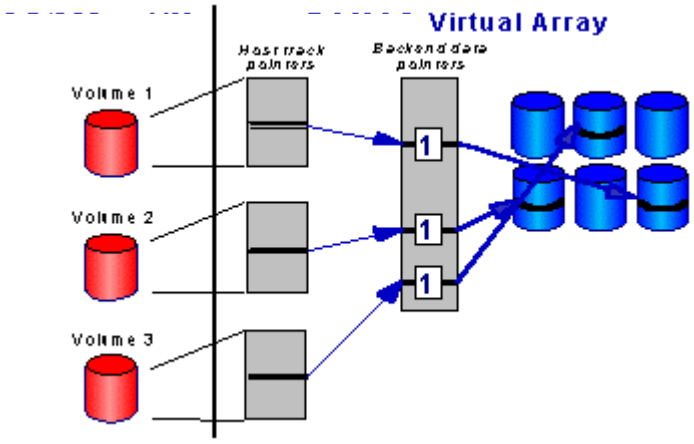

Fig. 3

Traditional Physical Duplication vs. SnapShot

A traditional physical disk copy reads the source data from the disk into the host, and writes a duplicate physical copy back to the receiving disk. This method uses computing, I/O, disk storage, *and time*. In fact, the amount of time and resources consumed are directly proportional to the amount of data being copied. The larger the size of the data, the more resources, *and time*, used. Any physical duplication method also requires *time* to complete, complicating and lengthening operations and scheduling as well as recovery processes. ***SnapShot provides unsurpassed benefits for a duplication product.***

SnapShot simply updates pointers (within the Virtual Disk Array mapping tables) for the data views being duplicated at electronic memory speeds. These updated pointers reference the same disk array locations as the original source data - one physical copy of the data with multiple host "views" (fig. 3). Snaps use no disk storage initially since only one physical copy of the data exists at the point in time when the snap is created. Only changes to either the source data or target(s) use additional storage, and any parts of the data that remain common are shared. Space savings are just the icing on the cake; the *true benefit* from SnapShot is the time redirected to productive business activity.

Unlike traditional physical duplication, SnapShot does not move any user data to create a "snap" copy, and therefore uses virtually no time or resources. The output of a SnapShot is virtually the same as that of a traditional copy, the difference is the **time** required to create the copy is a mere fraction of traditional physical duplication products. ***With SnapShot, data duplication can now be viewed as opportunity instead of a necessity.***

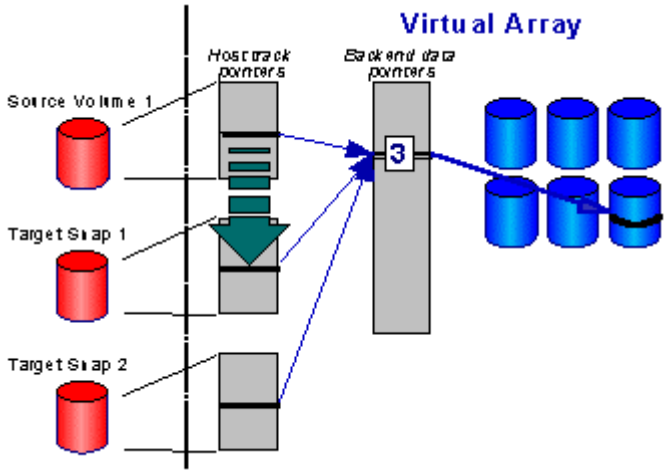

Fig. 4

SnapShot is VIRTUALLY INSTANTENOUS

Eliminating the need to physically move user data enables SnapShot's *virtually instantaneous* operation. With SnapShot a typical volume "snap" takes just 1-5 seconds for an full 3390-9 (9 GB) volume, and most of that time is host software overhead. Because SnapShot copies can be created almost instantaneously, off-site backups can now be created in parallel, or concurrently, with batch processing allowing key applications to be made available much sooner. SnapShot copies can be created at any time without the need to synchronize (or re-synchronize) source and target volumes, as other "time saving" duplication products, which require physical data movement and massive resources to create copies. The comparative speed of SnapShot versus other duplication products is exponential to the number of copies being created, only SnapShot can almost instantaneously create multiple copies without physical data movement.

True Point in Time

SnapShot commands are directed to the Virtual Disk Array using standard I/O channel commands (CCWs). Since these commands are processed serially and extremely fast with SnapShot, a true point in time “copy” results. No matter how many copies you need or how quickly you need them, SnapShot will produce them with point in time integrity to meet you needs. To maintain the integrity of your data, save to disk any file buffers in the CPU's host memory prior to creating the "snap".

Functions

There are three commands that perform the functions of SnapShot.

SNAP VOLUME will snap an entire disk volume to another disk volume, in MVS, VSE/ESA or VM/ESA operating environments. Additional keyword parameters enable options for flexibility and ease of use. For example the target of a Snap Volume operation can retain its original volume serial number (or name).

SNAP DATASET will snap a data set, creating a fully cataloged ready to use data set. Additional keyword parameters enable options for flexibility. For example the target of a Snap Dataset operation can have specific DFSMS (Data Facility System Managed Storage) characteristics set. The ability to easily and quickly snap individual data sets is required to gain the maximum benefit from any duplication product. Many customers need to target individual applications for improved availability. SnapShot with its data set copy capability fits this requirement perfectly. Because application files are spread out among multiple volumes, and intermixed with other application and user files, the need to "pick-out" specific files is not negotiable, it's required. Other competitive duplication products will claim this capability, but require you to first duplication every volume with at least one data set for the application volume first, and then spend time and effort to rename the data sets for use.

SNAP MINIDISK enables Virtual Machine (VM) users to snap minidisks.

SNAP enables open systems platforms to execute SnapShots. Depending on the sub-parameters for source and target determines the object to be snapped (i.e. SUN partition, HP “disk”).

SnapShot can be run in several ways, using a batch job, command line (CLI), or using a REXX or CLIST programming interface. Additional keyword parameters enable options for flexibility.

SnapShot operates using the Virtual Disk Array mapping tables, therefore source and target data elements must reside on the same Virtual Disk Array system. The optional DATAMOVERNAME parameter is available to insure that your SnapShot jobs will function when SnapShot may not available, such as in a disaster recovery test. With this parameter specified SnapShot will call DFSMSdss to physically move data to and from any 3990 compatible disk system.

SnapShot Installation

SnapShot installs quickly and easily using standard system installation utilities. SnapShot takes advantage of virtual disk architecture, available only in the RAMAC Virtual Array, or Shared Virtual Array (or Iceberg). Therefore any of these machines are required for SnapShot. SnapShot currently runs in OS/390, SUN, HP, VSE/ESA and VM/ESA environments, requires the Shared Virtual Array Administrator (SVAA/SVAC) (or Iceberg eXtended Facilities Product (IXFP) Version 2 Release 1) software, and Licensed Internal Code. The service engineer installs a non-disruptive diskette, enabling SnapShot within the Virtual Disk Array system.

Now you are ready to begin your pilot-project using SnapShot. For detailed installation instructions reference the SnapShot installation guides for your specific platform.

Testing SnapShot

Once SnapShot has been fully installed, you should test it. To confirm a successful installation, test all of the methods and commands, for example, batch job, command line, and the REXX / CLIST programming interfaces.

The installation documents referenced in the above installation section contain an installation verification checklist. And directions for executing the installation verification process.

Demonstrating SnapShots ultra fast operation to others in the organization will amaze and impress them, and clearly show SnapShots potential value. Developing a simple, easy to run demo will be useful for this purpose.

The Virtual Power Suite

SnapShot is the cornerstone product in the Virtual Power Suite. The virtual power suite is a set of unique software products that harness the power of the Virtual Disk Architecture of the Shared Virtual Array, IBM RAMAC Virtual Array, and Iceberg disk subsystems, invented and manufactured by StorageTek. Other components of the Virtual Power Suite include:

- Peer-to-Peer Remote Copy. PPRC is a fast disk mirroring solution that harnesses the built-in compression capability of the SVA to transmit critical data to a secondary site faster. A typical data compression rate of 75% means that only ¼ of the data is actually transmitted, slashing response times. It is fully supported by Globally Dispersed Parallel Sysplex software that prevents a rolling-disaster from corrupting your secondary site. PPRC can also be used as a high-speed data mover that does not use CPU resources. By combining PPRC with SnapShot you can get a powerful data mover that operates asynchronously to your applications.

- High Speed Data Mover. HSDM like PPRC harnesses the built-in compression capability of the SVA, except that you can use it for backups. By dumping compressed data, you can cut your backup times in half or better. But since you have SnapShot to save your backup times, the best benefit of HSDM is the restore. When restoring HSDM data to an SVA (with HSDM feature), the restore times will also be cut in half, and that's when you really need the speed of HSDM. And if you choose to restore to a non-virtual disk subsystem, the host software will automatically decompress you data and restore to any 3990 compatible storage device so it's a totally open solution.

- Transparent Data Migration Facility. TDMF allows you to non-disruptively migrate data for any disk to SVA. Optional features also allow you to migrate from any disk to any other disk while applications are fully active. Furthermore TDMF does not actually move your data until you give the word so you never have to worry about where your data really resides.

Pilot Project Outline

In this section we will discuss a suggested process for quickly implementing SnapShot into the enterprise. Many SnapShot customers have used this process successfully, and StorageTek has used this process as a guideline to their current SnapShot service offering. The process outlined below will produce a fully tested and documented SnapShot implementation, including a plan to go into production in just three days of dedicated effort.

Day 1: Select an application for SnapShot

The initial selection of where to target SnapShot is by far the most important. Before selecting an application you should understand where other customers have found value for SnapShot in their environments. Listed below are some areas where you may wish to focus your attention.

Potential Uses for SnapShot:

- ✓ Batch Window Reduction (Online window elongation)
- ✓ Disaster Recovery/Business Continuance
- ✓ Decision Support/Data Mining and Data Warehousing
- ✓ Application Development and Testing
- ✓ Duplication Utility

Since over 70% of customers polled want to target the batch window, this is the example that we will use in describing this pilot project.

The first activity is to hold a "brainstorming" meeting for the purpose of identifying which application would best benefit from the timesaving possible with SnapShot. A key to the success of this meeting is to have the appropriate people attend. These people must have a good understanding of the business challenges facing the company, and how the IS applications serve the business. Another key to the success of this meeting will be insuring the group understands the basics of SnapShot and what it can do. The demo mentioned earlier in this guide, would be a good way of accomplishing the task of training. This process will typically reveal several applications as potential targets. The next step will be to prioritize this list of applications in order to pick the best one for the pilot project.

Application Selection Criteria:

1. **Business value** - which application would provide the largest business value if availability were significantly improved. For example:
   Customer Service - serve customers better by staying "open" longer.
   Competitive Advantage - increasing availability may allow better production for global operations across multiple time zones.
   e-Business - duplicate legacy databases for electronic business using the Internet.

2. **Ease of implementation** - for a pilot project, selecting a simply structured application would help speed the process.
3. **Application dependencies** - many applications have some amount of dependence on other applications. Selecting one with few or no dependencies would also help speed the process.

In some cases the selection process will produce multiple applications because of dependencies.

Use this checklist to document the activities for Day 1's activities:

Complete	Activity
	SnapShot Project Manager appointed.
	Brainstorming Meeting Scheduled.
	Meeting attendees selected (including system support, operations, and scheduling).
	Meeting held.
	Top 3 applications identified and prioritized.
	Initial application selected from prioritized list for Pilot Project.

Day 2: Form Cross-Functional Team

The next step will be to assemble a project manager and team. Team members should consist of people who have technical knowledge of the target application, SnapShot, system operation, and production change processes.

The team will accomplish these tasks:

- ✓ Analyze selected application.
- ✓ Identify logical points in the application for SnapShot.
- ✓ Develop test plan.
- ✓ Identify project metrics.
- ✓ Develop production plan.
- ✓ Produce final project report.

Application Analysis

The team should review the application's design and current operation, and look for where to use SnapShot. Application flow charts are very useful for this task, and a technical presentation by a system analyst would be wise. There are two major areas to review for improvement.

1. Application backups. Traditional backup programs can be replaced with SnapShot, allowing batch updates to begin immediately, and the physical backup to run

concurrently with batch updates. For example, an application takes backups prior to and following batch updates, taking 60 minutes total. With SnapShot, Snap the files associated with this backup in just a few seconds, then begin the batch process while backing up the "snap" copies to tape concurrently, reducing the applications elapsed time by almost 1 hour. Analyzing your batch process and identifying the critical path is important to effectively reducing the overall batch run time.

2. New opportunities for duplication. Until SnapShot, duplication was viewed as a painful necessity, to be done as little as possible. With SnapShot the pain of duplication has been eliminated, presenting new opportunities for data duplication. For example, adding additional checkpoint disk backups to be used in the event of an application failure reducing recovery times to a minimum. In fact, many SnapShot customers have pointed out that many times the reason for missing the online window start time is due to this very problem.

Use this checklist to document the activities for Day 2's activities:

Complete	Activity
	Cross-Functional team organized (based on Pilot Project application).
	Selected application's flow charts and/or other documentation available.
	SnapShot opportunities identified within selected application.
	Testing strategy decided and planned.
	Operations and Support training strategy decided and planned.
	Operational procedures modified for SnapShot changes.

Day 2 (Night): Copy selected application files to Virtual Disk Array

The first thing to do is to replicate the application files necessary for testing. If all of the files are currently on the SnapShot capable or Iceberg Virtual array, then use SnapShot to make the copies. The Job Control Language (JCL) libraries containing the applications jobs should also be copied.

Day 3: Test the new application, and plan for production

Apply SnapShot Changes

Sufficient time should be given to thinking about exactly the best way to implement SnapShot. The team should decide at which points in the applications' batch process where SnapShot will provide timesaving. Team members should also decide which SnapShot command is appropriate. Typically, when targeting a production application for batch window reduction, the Snap Dataset command will be the choice. This is because application files tend to be spread across multiple volumes, for which the Snap Volume

command would not provide the level of granularity necessary. However if all of the application files are alone on host volumes, Snap Volume may be chosen.

It is important to distinguish some differences between Volume and data set Snaps. The Volume function will "snap" every host track from the source volume, to the selected target volume of same device type and model. When the operation is complete, every data set on the original source volume will have a duplicate data set on the target volume. MVS system catalogs do not recognize data sets with the same name. In this example you could still backup the "snapped" data sets by backing up the entire disk volume. Volume snaps typically run, depending on system activity, between 3-5 seconds each. One reason for this speed is the fact that no systems catalog, or data set allocation services are used.

The Snap Dataset command operates differently. It can "snap" into either a pre-existing data set, one created in the SnapShot job with attributes defined by the user, or *dynamically* allocated by SnapShot with attributes copied from the original source. The Snap Data set uses system catalog, data set allocation, and system managed storage (DFSMS) services. These services add additional time to the Snap Data set command, which typically runs between 5-25 seconds, depending on system activity and data set type. ***Experience shows that dividing up the data sets being snapped into a maximum of eight jobs will significantly improve the overall elapsed time.***

An important decision should be which parameters to use for SnapShot. A careful study of the Using SnapShot manual and Implementing SnapShot red book will fully explain all of the keyword parameters and their use.

Another key decision is to determine exactly how long to keep the SnapShot created data sets or volumes accessible on the Virtual Disk Array. This can be important because changes that occur on either the source or target use additional disk array storage. The usefulness of the data, and the reason for creating the copy, will determine when to delete it. For example, if the only reason a copy is created is to create an offsite backup concurrently with batch updates, then deleting the file following a successful backup is wise. However, if this same copy adds value for running ad-hoc reports during the day, alleviating contention on the production copy of the file, then keep it online. In this case only updates to the production data will occupy new space, until the copy is recreated in the next cycle. Any data that remains common between the source and target will share disk array storage. The worst possible case is the same amount of space will be used as if traditional duplication were used. The output of this meeting is an implementation plan, including estimated benefits.

Testing

A primary activity in the project will be proving the application's new flow, and verifying the recovery procedures. Test the new SnapShot steps, by running the entire application in your test environment. Verifying recovery procedures are critical in the testing phase,

since this is the purpose of backups. ***Most shops have their own testing criteria for moving changes into production, use yours***.

Project Metrics

As with any project, or process improvement, measuring key information before and after change is necessary to fully determine the project's level of success. Sometimes IS professionals get "caught-up" in computer metrics like CPU time and I/O, while missing an opportunity to measure a metric like amount of increased company revenues, or level of customer satisfaction. The key point of this section is to suggest that along with metrics like CPU time, I/O, and disk space consumed, to also measure the SnapShot improvements from a business perspective. In many cases, these business metrics are difficult to quantify, but even a conservative "ballpark" estimate would go a long way in communicating the true value of your SnapShot investment. ***Don't just measure the savings received from SnapShot, but also measure how you've leveraged these savings into increased business value***.

Production Plan

Most IS operations do not put new changes into production until they are very comfortable with them, and they have gone through the normal change process. This process that we have just taken you through recognizes this fact, and that is the purpose of producing a production plan. The production plan should include training for operations and programming staff. This plan should include a date for moving the changes into production, along with an ongoing measurement strategy to verify the long-term effect from SnapShot. These measurements will help justify resources for future SnapShot projects.

Final Report

It is important to understand the value received from investments, and some investments return much more than others. SnapShot customers worldwide are continuously proving that SnapShot turns unproductive hours, every day, into productive revenue-generating hours. Executives depend on information as the lifeblood of business. Reporting excellent investments is always good news. This report should effectively communicate in business term the savings achieved, and how effectively they were translated into business success. This report should be completed as part of the three-day exercise, but the metrics will reflect *estimated* benefits calculated from the test phase of the project. Once the changes have been migrated into production and measured for a reasonable time period, the report should be updated with actual benefits achieved. When presenting this report, include other application managers, analysts and end users. You will find that many of these people are not aware of new technology, and will be able to supply you with additional ideas concerning their applications. ***A final report template is included in the implementation plan.***

Use this checklist to document the activities for Day 3's activities:

Complete	Activity
	SnapShot commands and keyword parameters identified.
	Job Control modified in test environment.
	Retention of SnapShot copies determined.
	All aspects of implementation (including recovery procedures) fully tested.
	Application metrics identified and documented.
	Production plan decided and documented.
	Final report outline created and ready for final production metrics.

What's Next?

Now that you have seen how to get SnapShot up and running in short order, let's take a look at what other things can be done using SnapShot. The categories listed are just a few major operational activities that can benefit from SnapShot.

Business Continuance

Surprisingly many IT shops either have no business continuance plan, or an inadequate one. SnapShot can help you achieve a recovery plan that best fits your needs without the pitfalls associated with using traditional physical duplication utilities. For example a very large customer stops their entire shop, Snaps *ALL* of their volumes in just a few minutes, then restarts the applications and dumps the Snapped volumes to be taken offsite. This procedure is run daily, insuring that they will be able to continue operations in the event of a disaster. Another customer choose to improve their recovery process by quiescing their database three times per day instead of just once, dramatically reducing their recovery time in the event of a disaster. Using SnapShot to enhance your business continuance process depends on your particular recovery requirements, SnapShot allow you the luxury to re-think this very important and often overlooked business requirement. ***SnapShot uses minimal time and resources allowing you to maximize your recovery capability with minimum disruption.***

Shadow Database

There are several possibilities in creating a Shadow database. You can create a copy of your critical databases prior to the batch update process, then make the database available in read only mode. Once the batch process has finished simply shutdown the read only version and restart the updated database. Another option is that following the batch process you could make additional copies of the updated database to be used for generating reports or adhoc queries, reducing contention on the production copy of the database improving overall performance.

System Software

Duplicating System Software is another way that SnapShot can save you time and resources. With today's mergers and acquisitions, it may become necessary to have the ability to run different levels of system software until the systems merge. With SnapShot you have the flexibility to quickly and easily make copies of system volumes and data sets. Once you have created these copies you can easily apply changes such as maintenance and new releases. After you have successfully tested the changes, you can use SnapShot to copy the new volumes into production. If problems are encountered simply vary the new volumes offline and vary online the original copies and restart your system.

Application Testing

The problems associated with application testing include the amount of capacity required to fully test applications, the *time* it takes to duplicate the data, to reset the test environment for re-testing. SnapShot can be an invaluable tool for overcoming these issues. The primary advantage of using SnapShot for application testing is the time saved resetting (or restoring) the application files following a test iteration. Instead of restoring files from tape or disk, you can keep a base-copy of the test data on the system and simply re-snap for each iteration of the test. Also, you may choose to make several copies of this base-data for individual testers private use, effectively increasing the concurrency of your testing effort. All of these suggestions will help you reduce the *TIME* required to achieve maximum application quality and on-time delivery. By reducing the time required to replicate data for application testing, you can move them into production quicker and do additional testing to improve quality. All of the benefits of improved testing - without the expense of duplicate capacity and valuable resources.

Valuable Tool

Special Uses are limited only by ones imagination. For example, using SnapShot to pre-format data base volumes or VM minidisks. With traditional tools, adding volumes to a database pool would require formatting each volume, adding them to the pool, and consuming system resources and precious time. With SnapShot, you simply snap one pre-formatted volume to each target volume, then add them to the pool. Another use for SnapShot could be as an "instant initialize" utility. Define a volume on the Virtual Disk Array and leave it empty. Then simply snap the empty source to those target volumes being initialized using COPYVOLID (NO).

Identifying a SnapShot candidate is as simple as knowing which of your applications are the most critical to your business. Which applications require maximum availability? Minimizing down time for key applications could translate directly into increased productivity, higher net income, and increased earnings per share and revenues. Once these simple approaches have been used, you will be familiar enough with the product to want to find those other areas that help your business.

Technical Hints and Tips

This section of the guide will provide some hints and tips learned from other customer experiences. These tips are not always applicable to all scenarios, so careful examination of the appropriate technical documentation is suggested to augment this information. To improve this document in the future, your suggestions and experiences are very valuable to other customers who may be starting or expanding their SnapShot implementation. Please provide any feedback or additions to this section to your STK storage representative for possible inclusion in a future update of this document.

Tip 1:
Review your storage management architecture before SnapShot implementation.
As part of your implementation, give consideration to your storage management configuration. For a SnapShot to be successful the target must reside on the same Virtual Disk Array, so defining your storage pools (SMS or otherwise) to insure that the snap will always be successful is recommended. For example, define a storage class on the Virtual Disk Array that will always contain the target, and specify the storage class name in the SnapShot parameters. Since every IS shop has different storage requirements, specific suggestions are not given, other than analysis of the storage architecture.

Tip 2:
Use Volume Preferencing in an SMS managed shop.
SnapShot has an optional service built-in, which will enhance storage management control in an SMS environment. Volume preferencing, an installation option for SMS shops, receives the SMS generated volume list from the ACS (automatic class selection) routines, and limits the list to volumes in the same Virtual Disk Array.
For more detailed information see "Implementing SnapShot" Red Book SG24-2241-00 section 2.2.3 on pg. 18.

Tip 3:
Use DATAMOVERNAME optional parameter in every SnapShot job.
Always use the DATAMOVERNAME(dss) parameter in all SnapShot jobs. If SnapShot cannot allocate the target on the same Virtual Disk Array, this parameter will help insure that SnapShot jobs will not fail for that reason. Although the duplication will require additional resources and time to complete, it will work. For example, if you are testing your Disaster Recovery capability at a location that does not have SnapShot, then your test will still be successful without having to make changes to your production JCL, of course the benefits of SnapShot will not be realized. Discussing SnapShot with your disaster recovery vendor is highly recommended so the service level your customers have come to expect can be achieved even in a disaster scenario.
Note: IBM's DFSMS/dss program product is required.
For more detailed information see "Implementing SnapShot" Red Book SG24-2241-00 section 2.7 on pg. 27.

Tip 4:
Divide Snap Data set commands between a maximum of 8 jobs.
When using the Snap Dataset function you must specify every data set being duplicated in a SnapShot command. Many applications have a large number of data sets associated with them. If all of the data sets for an application were put into a single job, each SnapShot request would be executed sequentially or serially. Virtual Disk Array has the ability to process up to eight concurrent SnapShot commands, dividing these commands in up to eight jobs will provide a significant savings in wall clock execution time of the SnapShot requests.

Tip 5:
Use SnapShots dynamic allocation for target data sets.
When using the Data Set SnapShot function, let SnapShot use dynamic allocation. As stated earlier, there are three methods for allocation target (output) data sets for SnapShot.

1. Pre-allocated or pre-existing
2. Allocated in a DD (data definition) JCL card in the SnapShot job. All data set DCB (data set control block) characteristics are defined by the JCL programmer.
3. Dynamic allocation. All data set DCB characteristics are copied from the source by SnapShot automatically.

By allowing SnapShot to copy the data set DCB information from the source data set, there will be less chance for error. Also, SnapShot will attempt to combine source data set extents into the fewest possible (typically one)-target extents.

Tip 6:
Review security configuration before SnapShot production implementation.
As SnapShot gets more familiar in your shop as an ad-hoc duplication tool, many users will want to use it. If allowed to run uncontrolled SnapShot copies can significantly increase storage capacity used in the Virtual Disk Array system, if not deleted. SnapShot has support built-in for SAC (security access and control) in which the system administrator can assign security for the SnapShot commands, thereby limiting it's use. Your shop has unique security requirements, and they should be included in the SnapShot implementation plan.
For more detailed information see "Implementing SnapShot" Red Book SG24-2241-00 section 2.10.3 on pg. 31.

Tip 7:
Develop a naming convention for SnapShot output.
Experience shows that developing a naming convention for data created by SnapShot is highly beneficial. Naming conventions allow you to easily identify data elements for which action may be needed, such as a deleting, or overlaying.
For example, if a data "snapped" multiple times during an application run a naming convention such as *hlq.applx.history.snap1* would help identify this file as the first one created in the cycle.

Tip 8:
Use caution when specifying the REPLACE (YES) keyword parameter.
The REPACE (YES) parameter will overlay existing data. If it is specified on a volume snap, any and all data on the target (output) volume will be replaced with the source volume data sets. If specified on a data set snap, any and all data in the target data set will be replace with the source data. This function should not be a large concern in a SnapShot implementation where specific cycles and processing steps are the same every time. It is more important when using SnapShot as an adhoc tool, where a human error can potentially cause a problem.

Tip 9:
Document new process definitions, and train operations / programming personnel.
Many customer change control procedures include this critical step, but many other do not. Because SnapShot can dramatically alter the run characteristics of an application or process, such as shaving off several hours of run time, it is critical that the operations staff is fully trained on the new procedures. Particular focus should be spent on the changes to documented recovery procedures. For example, if an application's recovery procedure has been modified from a tape restore, to a SnapShot "snap-back" (reversing the source and target on a SnapShot operation) the time difference can be a few seconds versus hours or minutes using the old procedure. Having these new procedures fully documented, combined with training will minimize those annoying midnight calls.

Tip 10:
Use COPYVOLID (NO) CONDVOL (LBL) for full-volume Disaster Recovery volumes for OS390.
When using the Snap Volume command for the purpose of creating full volume disaster recovery backups, or other full volume backups, the use of these parameters in recommended. The default of a volume snap is that every track from the source volume will be duplicated on the target volume. These parameters give you additional flexibility to enhance usage of these volumes. The following list shows the effects of a volume snap using these specific keywords.

1. COPYVOLID (YES) *Default* - All source tracks are duplicated on the target volume.
   - ✓ Volume serial number (VOLID) duplicated
   - ✓ Volume forced offline in MVS (remains online in VM)
2. COPYVOLID (NO) - all source tracks copied PLUS
   - ✓ Target VOLID changed back to its original name.
   - ✓ Target VTOC data set name changed back to its original name.
   - ✓ Target VTOC Index data set name changed back to its original name.
   - ✓ Target VVDS entries changed back to original data set names on the volume.
   - ✓ Volume varied online by SnapShot.
3. COPYVOLID (NO) CONDVOL(LBL) - all source tracks copied PLUS
   - ✓ Target VOLID changed back to its original name.
   - ✓ Volume varied online by SnapShot.

Using option three allows you to do a full volume backup of the target; you then rename the volume to the original source name during restore and immediately begin using it.

Tip 11:
Monitor cycles of Net Capacity Load for SnapShot applications.
It is difficult to determine the exact impact on NCL from a SnapShot implementation. A good estimate can be determined from the testing phase prior to production, but since application access typically varies widely an exact determination is impossible to calculate. Monitoring your NCL variance often immediately following a SnapShot implementation is recommended. Determining the upper and lower variances of NCL usage is valuable in understanding the storage cost of the SnapShot data, and also may help determine how long to preserve it. In our example, assume your current NCL is 45% prior to moving your SnapShot application into production. Measure the NCL prior to the batch cycle beginning, again during the cycle (maybe multiple times), following the batch cycle, and then at the beginning of the online window. Plot the usage, and repeat this for several application cycles. Depending on the number of updates, and how long copies are retained on the system, the NCL will vary. Once you have the data you can better understand NCL usage, and better prepare for the next SnapShot implementation.

Tip 12:
Refreshing SnapShots saves physical storage space.
When a snap is created and subsequent updates are made to either the source or target copies, additional storage space is used for those changes, and any common data shares the same physical space. As more updates occur, the source and target copies will diverge using more and more physical space until (at most) a duplicate amount of space is used for the source and target. However, if the SnapShot operation is repeated, the space required for unique data is immediately freed. For example, application X has a master file that is duplicated by SnapShot for adhoc reports throughout the day. The file also accumulates approximately 30% track updates throughout the day. If snapped only once per day you would use approximately 30% of the original data set's space for these updates. However, if the file were quiesced three times per day, this number would drop to approximately 10%.

Tip 14:
Use SnapShot for all application data set files in OS390.
The timesaving associated with Snap Dataset implementations are directly proportional to the size of the data set. For example, if the traditional duplication procedure (i.e. DFSMS/dss) backs up a 5Gb file and takes 30 minutes, replacing this with SnapShot would save over 29 minutes. If the file was only 1Mb, and the traditional copy took 20 seconds, replacing this with SnapShot may take the same or possibly a bit more time. However to be consistent, it is recommended that SnapShot be used even for these smaller data sets where the time benefit may not be significant. ***The size of the data set has negligible impact to the time required for the SnapShot; it's always FAST***.

Tip 15:
Upgrade your Virtual Disk Array to the latest generation.
SnapShot enables you to do more work concurrently, such as application backups during batch processing. This means that additional work or I/O requests need to be processed. The latest generation Virtual Disk Array subsystems have improvements specifically targeted at getting this work accomplished, and fits perfectly into a SnapShot environment. These machines not only perform faster, but also have the additional bandwidth to pump the data out for backups, while concurrently updating the batch copies.

Tip 16:
For complete true-point-in-time copies, quiesce the source files before snapping.
Depending on what you will do with the copies created by SnapShot, quiescing the source (input) files may be required. SnapShot will duplicate disk track pointers associated with host tracks for a file or volume. SnapShot has no knowledge of user data that may be stored in the host memory of the CPU. If any data stored in the host memory is required for the copy, the file must be closed (or data flushed) and all data written to the source disk file before the SnapShot operation is run. If the file will be used for testing, and a complete copy is not needed, then snapping the file while open is permitted. Care should be taken to correctly specify the keyword parameters that control serialization to achieve the desired result.
For more detailed information see "Implementing SnapShot" Red Book SG24-2241-00 section 2.10.2 on pg. 30.

Tip 17:
Use SnapShot for application testing.
The very process of testing is time consuming, and involves massive amounts of data duplication. These are the very ingredients that make SnapShot the tool of choice of speeding-up this critical task. Testing involves repetitive cycles, which include recovering to specific test points and re-testing. SnapShot can shave minutes and maybe hours from each iteration of the testing cycle. And because data duplication is no longer painful with SnapShot, you could do more of this testing work in parallel by replicating the test environment for each individual tester. ***What value do you put on launching new applications on time and with high quality?***

Tip 18:
Insure that multi-volume data set extents remain on the same Virtual Disk Array in OS390.
If you will be snapping multi-volume data sets, you should make every effort to keep all of the data sets extents on the same Virtual Disk Array system. If they are not on the same system, SnapShot will determine which extent(s) are largest, use SnapShot for them, and use the data mover (DFSMS/dss) program to physically copy the remaining extents to the same Virtual Disk Array to complete the copy.

For more detailed information see "Implementing SnapShot" Red Book SG24-2241-00 section 2.2.2 on pg. 18.

Tip 20:
Improve application performance with SnapShot.
One way you can use SnapShot is by duplicating high access read-only files across multiple volumes and/or Virtual Disk Array systems. This will spread the I/O workload across these copies, greatly reducing the queue time versus one copy. This technique can also be used for batch files that may be read by multiple batch programs concurrently.

Tip 21:
Share your successes with SnapShot
An excellent way of finding new uses for SnapShot is to share your success stories with others. This may be with other groups with the IT department, but should also include your customers and users. Application users and analysts may have new ideas and suggestions to further improve business operations related to their specific environments.

Final Tip:
Keep your ears and eyes open.
SnapShot is a very successful product with thousands of licenses worldwide. StorageTek recognizes the importance of this key storage product for its customers, and is committed to its continued advancement directly and through partnerships.

Trademarks