

Evaluation criteria for data de-dupe

How does it work? What are the different implementation methods? And what are the key evaluation criteria?

By Larry Freeman, Rory Bolt, and Tom Sas

SNIA DATA DE-DUPLICATION IS THE PROCESS of eliminating redundant copies of data.

The term “data de-duplication” was coined by database administrators many years ago as a way of describing the process of removing duplicate database records after two databases had been merged.

Today the original definition of de-duplication has been expanded. In the context of storage, de-duplication refers to any algorithm that searches for duplicate data objects (e.g., blocks, chunks, files) and stores only a single copy of those objects. The user benefits are clear:

- Reduces the space needed to store data; and
- Increases the available space to retain data for longer periods of time.

How it works

Regardless of operating system, application, or file-system type, all data objects are written to a storage system using a data reference pointer, without which data could not be located or retrieved. In traditional (non-de-duplicated) file systems, data objects are stored without regard to any similarity with other objects in the same file system.

Identifying duplicate objects and redirecting reference pointers form the basis of the de-duplication algorithm. As shown in the figure, referencing several identical objects with a single “master” object allows the space normally occupied by the duplicate objects to be “given back” to the storage system.

Design considerations

Given the fact that all de-duplication technologies must identify duplicate data and support some form of referencing, there is a surprising variety of implementations, including the use of hashes, indexing, fixed object

length or variable object length de-duplication, local or remote de-duplication, inline or post-processing, and de-duplicated or original format data protection.

Use of hashes

Data de-duplication begins with a comparison of two data objects. It would be impractical (and very arduous) to scan an entire data volume for duplicate objects each time a new object was written to that volume. For that reason, de-duplication systems create relatively small

ensure against hash collisions. Although concern over hash collisions is often raised, depending upon the hash algorithm used and the system design, the probability of a hash collision may actually be orders of magnitude less than the probability of an undetected disk read error returning corrupt data.

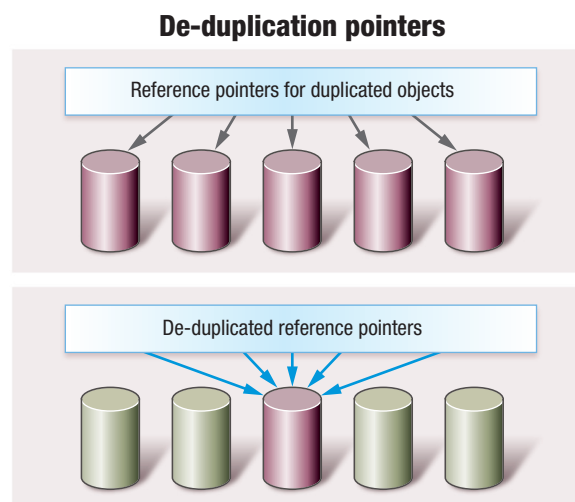
Indexing

Once duplicate objects have been identified (and optionally validated), removal of the duplicate object can commence. There are varying methods that systems employ when modifying their data pointer structures. However, all forms of this indexing fall into four broad categories:

Catalog-based indexing—A catalog of hash values is used only to identify candidates for de-duplication. A separate process modifies the data pointers accordingly. The advantage of catalog-based de-duplication is that the catalog is only utilized to identify duplicate objects and is not accessed during the actual reading or writing of the de-duplicated data objects; that task is handled via the normal file-system data structure.

Lookup table-based indexing—Extends the functionality of the hash catalog to also contain a hash lookup table to index the de-duplicated object’s “parent” data pointer. The advantage of a lookup table is that it can be used on file systems that do not support multiple block referencing; a single data object can be stored and “referenced” many times via the lookup table. Lookup tables may also be used within systems that provide block-level services instead of file systems.

Content-addressable store, or CAS-based, indexing—The hash value, or digital signature of the data object itself, may be used by itself or in combination with additional metadata as the data pointer. In a content-addressable store (CAS), the storage location is determined by the data being stored. Advantages of CAS-based indexing include inherent single instancing/de-duplication, as well



hash values for each new object to identify potential duplicate data.

A hash value, also called a *digital fingerprint* or *digital signature*, is a small number generated from a larger string of data. Hash values are generated by a mathematical formula in such a way that it is extremely unlikely (but not impossible) for two non-identical data objects to produce the same hash value. In the event that two non-identical objects do map to the same hash value, this is termed a “hash collision.”

Evaluation criteria

Understanding a system’s use of hashes is an important criterion when you are evaluating de-duplication. If the technology depends solely on hashes to determine if two objects are identical, then there is the possibility, however remote, that hash collisions could occur and some of the data referencing the object that produced the collision will be corrupt. Certain government regulations may require you to perform secondary data object validation after the hash compare has completed to

Due to the growing interest in data de-duplication and space reduction solutions, the SNIA DMF Data Protection Initiative has recently been tasked with forming a Special Interest Group (SIG) focusing on this topic. This is the first in a series of publications from SNIA on the topic of de-duplication and space reduction. The mission of the DDSR SIG is to bring together a core group of companies that will work together to publicize the benefits of data de-duplication and space savings technologies. Anyone interested in participating can help form the direction and ultimate success of the group. Find out more at www.snia-dmf.org/dpi.

as enhanced data integrity capabilities and the ability to leverage grid-based storage architectures. Although CAS systems are inherently object-based, file-system semantics can be implemented above the CAS.

Application-aware indexing—Differs from other indexing methods in that it looks at data as objects. Unlike hashing or byte-level comparisons, application-aware indexing finds duplication in application-specific byte streams. As the name implies, this approach compares like objects (such as Excel documents to Excel documents) and has awareness of the data structure of these formats.

Evaluation criteria

De-duplication indexing is an important consideration in technology evaluation, particularly when it comes to resiliency of design. When indexing data objects, the index itself could become a single point of failure.

It is important to understand what, if any, single points of failure are present in a de-duplication system. It is equally important to understand what measures are used to protect these single points of failure to minimize the risk of data loss.

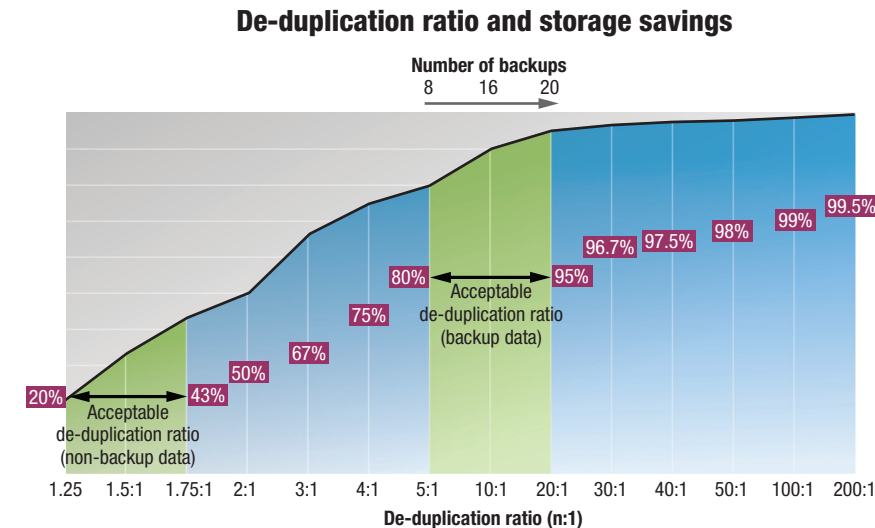
Another indexing consideration is the speed of the index. An inordinate amount of time should not be required to store and retrieve data objects, even when millions of objects are stored in the file system. When evaluating de-duplication, consider both lightly loaded and fully loaded file systems and the potential performance degradation caused by indexing as more and more data is written to the file system.

Fixed object length or variable object length de-duplication

De-duplication may be performed on fixed-size data objects or on variable-size data objects.

With fixed object length de-duplication, data may be de-duplicated on fixed object boundaries such as 4K or 8K blocks. The advantage of fixed block de-duplication is that there is less computational overhead in computing where to delineate objects, less object overhead, and faster seeking to arbitrary offsets.

With variable object length de-duplication, data may be de-duplicated on variable object boundaries. The advantage



of variable object size de-duplication is that it allows duplicate data to be recognized even if it has been logically shifted with respect to physical block boundaries. This can result in much better data de-duplication ratios.

Evaluation criteria

Fixed object length de-duplication offers processing advantages and performs well in both structured data environments (e.g., databases) and in environments where data is only appended to files. In unstructured data environments such as file servers, variable object length de-duplication is able to recognize data that has shifted position as the result of edits to a file. Variable object length de-duplication typically offers greater de-duplication in unstructured data environments. Since fixed object length is a subset case of variable object size, many systems capable of variable object length de-duplication also offer fixed object length.

Local or remote de-duplication

Local or remote de-duplication refers to where the de-duplication is performed:

In *local* de-duplication, de-duplication may be performed within a device. This allows transparent operation without the need for APIs or software agents. Local de-duplication is sometimes referred to as target de-duplication in the backup market.

In *remote* de-duplication, for LAN- or WAN-based systems, it is possible to perform de-duplication remotely through the use of agents or APIs without requiring additional hardware. Remote de-duplication extends the benefits of de-duplication from storage efficiency to network efficiency. Remote de-duplication

is sometimes referred to as source de-duplication in the backup market.

Evaluation criteria

The advantage of local de-duplication is total application transparency and interoperability; however it doesn't address remote or distributed systems or bottlenecks in networks. Although it requires a specialized agent or API, remote de-duplication offers tremendous potential for both network bandwidth savings and application performance.

Inline or post-processing

Another design distinction is when to perform de-duplication. Again, there are multiple design options.

With inline de-duplication, de-duplication is performed as the data is written to the storage system. The advantage of inline de-duplication is that it does not require any duplicate data to be written to disk. The duplicate object is hashed, compared, and referenced on-the-fly. A disadvantage is that more system resources may be required to handle the entire de-duplication operation in real time.

With post-processing de-duplication, de-duplication is performed after the data is written to the storage system. The advantage of post-processing de-duplication is that the objects can be compared and removed at a more leisurely pace, and typically without heavy utilization of system resources. The disadvantage of post-processing is that all duplicate data must be first written to the storage system, requiring additional storage capacity.

Evaluation criteria

The decision regarding inline versus post-processing de-duplication has more to do

with the application being de-duplicated rather than any technical advantages/disadvantages.

When performing data backups, the user's primary objective is the completion of backups within an allowed time window. For LAN- and WAN-based backups, remote inline de-duplication may provide the best performance. For direct-attached and SAN-based backup, an assessment should be made to determine which approach works best. Either may be appropriate, depending on data type and volume. If post-processing de-duplication is deployed, users should ensure there is adequate time between backup sessions to complete the de-duplication post-process.

With general applications, the cost of additional storage needed by post-processing needs to be weighed against the cost of system resources and the performance of inline de-duplication to determine the best fit for an environment.

De-duplicated or original format data protection

As is the case with all corporate data systems, de-duplicating storage systems need to be protected against data loss. De-duplicating systems vary with respect to their approach to data protection.

When protecting a de-duplicated system, it is possible to perform backups and replication in the de-duplicated state. The advantages of de-duplicated data protection are faster operations and less resource usage in the form of LAN/WAN bandwidth.

De-duplicated systems can also be backed up and replicated in the original data format. The advantage of protecting data in the original format is that the data can theoretically be restored to a different type of system that may not support data de-duplication. This would be particularly useful with long-term tape retention.

Evaluation criteria

If media usage and LAN/WAN bandwidth are a concern, de-duplicated data protection offers clear cost advantages as well as performance advantages. Note that while original format data protection offers the possibility of cross-platform operation, in practice many data-protection solutions do not allow cross-platform operation. Finally, some

systems offer users a choice of either type of data protection.

De-duplication space savings

So what should you expect in terms of space/capacity savings with data de-duplication?

De-duplication vendors often claim 20:1, 50:1, or even up to 500:1 data-reduction ratios. These claims refer to the “time-based” space savings effect of de-duplication on repetitive data backups. The figure on p. 27 illustrates this theoretical space savings over time. Since these backups contain mostly unchanged data, once the first full backup has been stored, all subsequent full backups will see a very high occurrence of de-duplication. Assuming the user retains 10 to 20 backup images, and the change rate between backups is within the

norm (2% to 5%), this user should expect storage space savings in the range of 5:1 to 20:1. If you retain more backup images, or a reduced rate of change between backups, your ratio will increase. The larger numbers, such as 300:1 or 500:1, tend to refer to data moved and stored for daily full backups of individual systems.

Another area to consider for data de-duplication is non-backup data volumes, such as primary storage or archival data, where the rules of time-based data reduction ratios do not apply. In those environments, volumes do not receive a steady supply of redundant data backups, but may still contain a large amount of duplicate data objects.

The ability to reduce space in these volumes through de-duplication is measured in “spatial” terms. In other words, if a

500GB data archival volume can be reduced to 400GB through de-duplication, the spatial (volume) reduction is 100GB, or 20%. Think of it as receiving a “storage rebate” through de-duplication. In these applications, space savings of 20% to 40% may justify the cost and time you spend implementing de-duplication.

Summary

Data de-duplication is an important new technology that is quickly being embraced by users as they struggle to control data proliferation. By eliminating redundant data objects, an immediate benefit is obtained through space efficiencies.

When evaluating de-duplication technologies, it is important to consider major design aspects, including use of hashes, indexing, fixed or variable object length,

local or remote operation, inline or post-processing, data protection, and of course, the space savings reduction that you will receive. By considering these items and understanding how they impact your data storage environment, informed decisions can be reached that benefit your environment. Watch for more on this topic at www.snia.org. □

This article was written on behalf of the SNIA Data Management Forum. **Larry Freeman** is a senior product manager at Network Appliance; **Rory Bolt** is chief technology officer, Avamar Products, EMC; and **Tom Sas** is a product marketing manager at Hewlett-Packard.

VENDORS MENTIONED
EMC, Hewlett-Packard, Network Appliance

Cutting the costs of remote disaster recovery continued from p. 25

replicate both primary production data and backup data to the company’s warm disaster-recovery site. Although currently still backing up to tape, Shapelow says the company plans to move toward using CommVault software with a virtual tape library (VTL) with data de-duplication functionality. “We’ll put a de-duplication device on-site and replicate to the off-site location,” he explains. Shapelow says the company is also currently in the process of switching its VMware virtual servers to use VMware Consolidated Backup (VCB) in conjunction with

CommVault’s Galaxy software to direct virtual server images to the Symmetrix array where data can be replicated to the warm site via SRDF.

Don’t forget virtualization

Server virtualization and its ability to help reduce the cost of remote DR has gained traction over the past several months. “The biggest expense, after personnel, on the secondary site tends to be hardware,” says Lamorena. “If you use server virtualization tools, you can reduce the number of boxes you have at the second site.”

The growing popularity of server virtualization is putting greater emphasis on virtual storage architectures as well when it comes to cost savings in remote DR. When it comes to replication, companies can choose from a wide range of solutions. Software applications from vendors such as DoubleTake and Neverfail are examples of server-based replication solutions. In addition to server-based approaches, replication choices include appliance-, array-, and database-level replication.

Among this range of current choices for remote replication is also a growing mix of storage virtualization software vendors such as DataCore and FalconStor, as well as combined hardware/software virtualization vendors such as EMC, Hitachi Data Systems, and IBM. Users also have the choice to implement local and remote replication services on top of the vendors’ respective virtualization platforms.

The Taneja Group’s Norall sees three primary levels of technology at which data is moved across the wire from a primary to secondary site: host-based replication products installed as software on the host agent, network virtualization products that support both asynchronous and synchronous replication, and array-level replication and mirroring.

Datalink’s Robinson says companies have a lot more choice and cost-savings opportunities in the type of storage sys-

tem they can now deploy at their secondary sites, noting that there are now a variety of heterogeneous replication solutions available to help drive down the cost of remote DR.

Although Robinson doesn’t see large companies wanting to replicate between a Symmetrix and a Nexsan system, for example, he has seen the desire to replicate to lower-cost systems in the same vendor’s line. “We may see customers going from a Symmetrix to a Clariion. The same thing with Hitachi and its replication technology on the USP [Universal Storage Platform],” he says. “Technically, while I still need a USP to replicate on the other end, I can use cheap disk behind it. Although it’s still frame-to-frame, it’s a way to reduce the overall cost.”

Regardless of which technology you choose, there’s no doubt that there are more choices than ever when it comes to cost savings for remote DR. □

Michele Hope is a freelance writer covering enterprise storage and networking. She can be reached at mhope@thestoragewriter.com.

VENDORS MENTIONED
Brocade, Cisco, CommVault, Data Domain, DataCore, Datalink, DoubleTake, EMC, FalconStor, Hitachi Data Systems, IBM, Juniper Networks, Neverfail, Nexsan, Overland Storage, Riverbed, Silver Peak Systems, Sun, Symantec, VMware

Case study: Tiered data protection – Large manufacturing company
(Deferring purchases of Tier-1 storage for two years)

