

Smart Array Controller technology: drive array expansion and extension

technology brief



Abstract.....	2
Introduction.....	2
Array configuration utility	2
Capacity expansion	2
Logical drive creation	3
Volume extension	4
Volume extension in Windows.....	5
RAID migration.....	7
Stripe size migration	8
Glossary.....	9
Conclusion.....	9
Call to action	10

Abstract

This paper describes four specific functions of the Smart Array Controller family: capacity expansion; volume extension; RAID level migration; and stripe size migration. Capacity Expansion adds physical drives to an array. This is usually done in anticipation of additional changes to the storage configuration, such as volume extension. Volume extension takes this newly available space in the array and allows the OS to use it by increasing the size of an existing volume. If administrators plan to perform volume extension on a Windows 2000 server, they must be aware of specific requirements that must be met to perform the volume extension. Administrators can also optimize their storage by using the Array Configuration Utility to change either stripe sizes or RAID levels within an array.

Introduction

As data storage requirements increase and as computing needs change, customers require flexibility not only within server configurations but also in storage configurations. The Smart Array Controller family includes a standard toolset that can be used to configure array controllers, expand an existing array configuration by adding disk drives, or reconfigure an array by extending volume sizes.

Customers find the abilities of online capacity expansion and volume extension extremely worthwhile. However, these functions are not needed on a daily basis; nor is it easy to remember the differences between the two terms. As a result, customers sometimes find it difficult to distinguish between the two. This paper presents a brief explanation of capacity expansion and volume extension. The paper also briefly discusses other changes that an administrator might want to make, such as RAID level migration and stripe size migration.

This paper assumes that the reader is familiar with basic array controller technology and terms. For those unfamiliar with drive array terms, a glossary is included for reference.

Array configuration utility

The array configuration utility (ACU)¹ is a local, graphical-based application that helps configure Smart Array Controllers. As of this writing, the ACU is a Windows-only utility capable of online configuration changes, that is, changes while the operating system (OS) is running.

The array configuration utility XE (ACU-XE)² is a browser-based version of the same utility. It is available to run online in Linux or Windows. Because ACU-XE is a bootable CD, an administrator using other operating systems³ can run the utility offline by booting the system from the ACU-XE CD.

Customers using NetWare can perform capacity expansion and other configuration changes through the NetWare Online Array Configuration Utility (CPQONLIN), the online tool for array configuration.

Capacity expansion

Capacity expansion is the process of adding physical drives to an array that has already been configured. The logical drives (or volumes) that exist in the array before the expansion takes place are unchanged. Only the amount of free space in the array changes.

¹ For more information about ACU, refer to

<http://h18004.www1.hp.com/products/servers/proliantstorage/software-management/acumatrix/index.html>

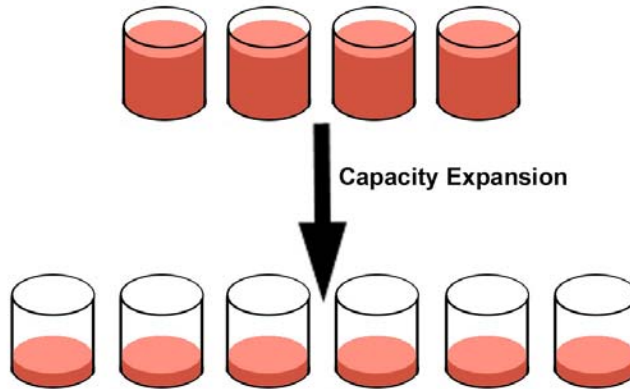
² For more information about ACU-XE, refer to

<http://h20000.www2.hp.com/bizsupport/TechSupport/SoftwareDescription.jsp?locBasepartNum=is-22561-1&lang=English>

³ Customers should check the user documentation to verify that their specific array controller supports their chosen OS.

For example, suppose that an existing array consisted of four physical drives and the administrator wants to expand the capacity to six physical drives. This is like having four glasses full of water, and pouring that water from the original four glasses into six glasses (see figure 1). The amount of water (the volume of the logical drive) has not changed — it has merely been redistributed, or expanded, into a larger number of containers (drives).

Figure 1. Array capacity expansion redistributes an array into a larger number of physical drives. The size of the logical drive does not change.



Thus, if an existing array is nearly filled with logical drives, an administrator can add new physical drives and initiate a capacity expansion through ACU or ACU-XE. The program automatically checks the drive hardware configuration. Then, the existing logical drive is distributed across all physical drives in the expanded array without affecting any existing data. If the array being expanded contains more than one logical drive, data will be redistributed one logical drive at a time.

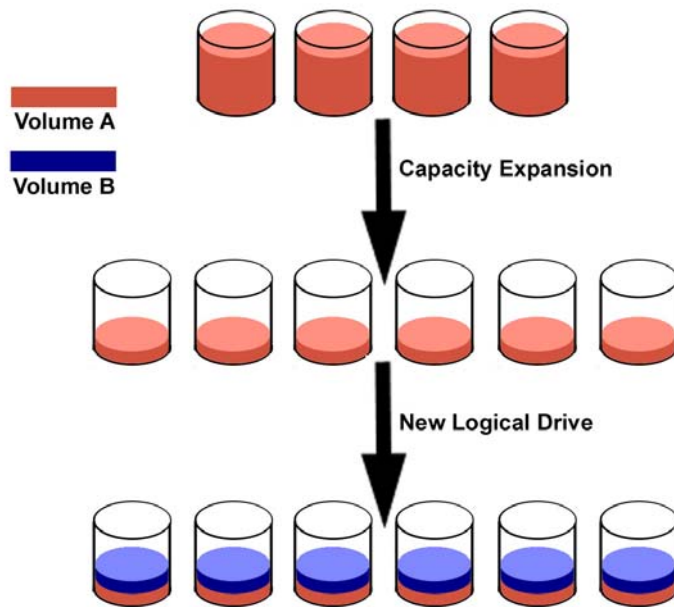
The expansion process is entirely independent of the system OS. For example, if a 10-gigabyte (GB) logical volume is expanded from four drives into six drives, the OS is unaware of this change.

The amount of time it takes to perform an online capacity expansion depends on several parameters: drive speed, the server microprocessor speed, the amount of other I/O work the server is doing, and the priority level of the capacity expansion. This priority level can be changed from low (the default in ACU) to medium or high to expand the volumes as quickly as possible. Depending on these factors, the expansion process takes somewhere between 10-15 minutes per GB.

Logical drive creation

Once the drive capacity has been expanded, the added capacity can be used to create a new logical drive (Figure 2) or to extend the size of an existing logical drive (as described in the section titled volume extension).

Figure 2. After capacity expansion, an administrator can use this free space by creating an entirely new logical drive.



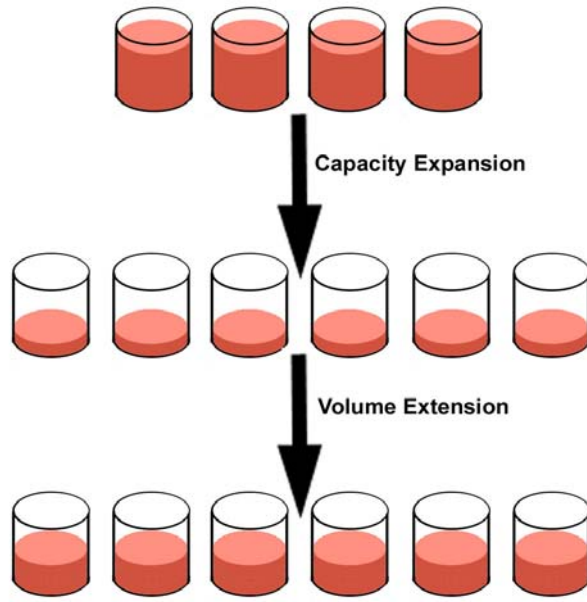
Some people might want to add a new logical drive, for example, if every department in a company had its own logical storage volumes. If a new department were created, the administrator might need to create an entirely new logical drive for that department, or need to improve performance of an existing volume by increasing the number of spindles (drives).

However, most customers will choose to extend the size of their volume after capacity expansion. The volume extension allows customers to increase the logical volume size if they are running out of storage space on the existing volume.

Volume extension

Volume extension grows the storage space of a logical drive (see Figure 3). During this process, an administrator adds new storage space to an existing logical drive on the same array, usually after the array has been expanded. An administrator may have gained this new storage space by either capacity expansion, as previously discussed, or by deleting another logical drive on the same array.

Figure 3. Volume extension grows the size of a logical drive.



Unlike capacity expansion, the OS must be aware of changes to the logical drive size. The following operating systems support online volume extension:

- Microsoft Windows 2000, NT 4.0
- Novell NetWare 4.2, 5.x
- SCO UnixWare 7.1.x

Offline volume extension is possible for some operating systems by backing up data, reconfiguring the array, and restoring the data from backup. Refer to the OS documentation for additional information.

Volume extension in Windows

Operating systems within the Windows family support online drive extension under specific conditions. To understand these conditions, an administrator must first recognize the differences between Windows basic and dynamic disks.

A basic disk uses Microsoft's classic partition system in which each volume uses only one partition. This system is the default in Windows NT, Windows 9x, and DOS. An administrator can perform volume extensions using ACU or ACU-XE if the system uses basic disks. The administrator can use the new space on the logical drive for an additional partition, or can grow the existing partition by using a third-party software utility such as PowerQuest's Partition Magic.

Windows 2000 supports basic and dynamic disks, with dynamic disks being the default partition system. A Windows 2000 dynamic disk is a physical disk that doesn't use static partitions or logical drives. Instead, it contains only dynamic volumes that can span multiple partitions. Dynamic disks are created in the Disk Management console of Windows 2000 and contain database tables at the beginning and the end of the dynamic volume to recognize the partition boundaries. Because this database may still exist even after a logical drive is deleted, an administrator must be careful when performing volume extension on a dynamic drive.

The best way to grow a volume size in Windows 2000 is to take advantage of this dynamic disk functionality, rather than performing a volume extension in ACU or ACU-XE. Because a dynamic disk can span multiple partitions on separate logical drives, an administrator can create a completely new logical drive when free space exists. Then the administrator can create a new partition on the logical drive and add it as a segment to the existing volume (this is what Microsoft calls expansion⁴). This functionally is the same as extending the volume size in ACU or ACU-XE.

If an administrator desires to use ACU or ACU-XE to extend a volume, the administrator should determine whether the drives are pristine – that is, the physical drives have never been used for any other dynamic disk. If so, the administrator should download and apply the Microsoft Service Pack 3. This ensures proper compatibility between Microsoft dynamic disks and logical volume extension using Smart Array Controllers.

If the administrator is trying to extend a volume of dynamic disks on previously used physical drives, the database used for a dynamic disk may still exist even after the logical drive is deleted. In this case, when an administrator performs volume extension, the OS may see this old database. In this case, the partition on the logical drive may be marked as a foreign disk and become unusable.

⁴ Microsoft expansion refers to adding partitions to a volume using dynamic disks and growing the size of that volume.

RAID migration

Administrators may perform volume extension in order to increase raw data storage capacity, to improve performance by increasing the number of spindles in a logical drive volume, or to change fault-tolerance (RAID) configurations. With ACU or ACU-XE, an administrator can reconfigure a logical drive to a new RAID level online, without disrupting system operation or causing data loss. The following situations would require additional drive space: changing RAID 0 to RAID 1, 5, or RAID ADG; or changing RAID 5 or RAID ADG to RAID 1. Table 1 summarizes the RAID levels and the amount of space required for each type of fault tolerance.

Table 1. Summary of RAID methods.

	RAID 0 (striping)	RAID 1 (mirroring)	RAID 5 (distributed data guarding)	RAID ADG (advanced data guarding)
Usable drive space*	100%	50%	67% to 93%	50% to 96%
Usable drive space formula	n	n/2	(n-1)/n	(n-2)/n
Minimum number of drives	1	2	3	4
Tolerant of single drive failure?	No	Yes	Yes	Yes
Tolerant of multiple simultaneous drive failures?	No	If failed drives not mirrored to each other	No	yes
Read performance	High	High	High	High
Write performance	High	Medium	Low	Low
Relative cost	Low	High	Medium	medium

* The values for usable drive space are calculated assuming a maximum of 14 physical drives of the same capacity (or a maximum of 56 for RAID ADG) with no online spares. HP recommends that these maximum figures (excluding any allowable online spares) are not exceeded when configuring a drive array, due to the increased likelihood of logical drive failure with more hard drives.

For more information about RAID levels, refer to the Smart Array Controller User Guide, or the document titled Assessing RAID ADG vs. RAID 5 vs. RAID 1+0. Both are available from the [smart array controller webpage](#).⁵

⁵ The Smart Array controller webpage is available at the following URL:
<http://h18000.www1.hp.com/products/servers/proliantstorage/arraycontrollers/documentation.html>

Stripe size migration

When an administrator expands the volume capacity of an array, this does nothing to change the stripe size of the data stored within the array. Each RAID level has a default value designed to provide good performance across many types of applications (see table 2). In addition, table 2 shows the range of stripe sizes available per RAID level. If an administrator changes the RAID level, and this new RAID level does not support the existing stripe size, the ACU or ACU-XE changes the stripe size automatically.

Table 2. Stripe sizes available per RAID level.

Fault tolerance level	Default (KB)	Available stripe sizes (KB)
RAID 0	128	8, 16, 32, 64, 128, 256
RAID 1 or 0+1	128	8, 16, 32, 64, 128, 256
RAID 5	16	8, 16, 32, 64
RAID ADG	16	8, 16, 32, 64

Certain applications, especially those that perform a majority of one type of transaction (for example, primarily write transactions), may require tuning of the stripe size to increase performance. Table 3 recommends stripe sizes for general types of server applications. An administrator can take these general recommendations and try progressively larger or smaller stripe sizes to determine the best overall performance for a particular application. Through ACU or ACU-XE, the administrator can make these changes online without disruption or data loss.

Table 3. Recommended stripe sizes

Type of server application	Suggested stripe size change
Mixed read/write	Accept default value
Mainly sequential read (such as audio/video applications)	Larger stripe sizes
Mainly write (such as image manipulation applications)	Smaller stripes for RAID 5, RAID ADG Larger stripes for RAID 0, RAID 0+1, RAID 1

Glossary

Term	Definition
Array	A group of physical drives. The ACU groups physical drives together for logical drives to be created upon them. A single physical drive can be part of only one array (excluding spares), but an array can be made up of multiple physical drives. Logical drives can span only one array.
Basic Disk	Microsoft's classic partition system on which each volume uses only one partition. A basic disk cannot grow in size without using a third party application. This system is the default in NT 4.0.
Dynamic Disk	Microsoft's new default partition system for Windows 2000. In this system, a volume can span multiple partitions. Thus an administrator can grow a volume by adding additional partitions.
Expansion	The process of adding a physical drive to an existing array. The logical drives that exist on the array before the expansion takes place are unchanged. Only the amount of free space on the array changes
Extension	The process of growing the size of a logical drive. In this case the size of the logical drive reported to the OS is increased.
Free space	Blocks on an array that are currently not assigned to a logical drive. Free space can be used for creating new logical drives, changing the RAID level of an existing logical drive, or growing the size of an existing logical drive.
Logical drive	A logical partitioning of space on physical drives on an array controller that is presented to the operating system as though it were an actual single physical drive attached to a SCSI controller. This entity can make use of differing RAID levels for fault tolerance.
MS Expansion	Microsoft's use of the term expansion refers to adding partitions to a volume using dynamic disks and growing the size of that volume.
Partition	Logical entity created on a logical drive by the operating system that can be used by the operating system.
RAID	Redundant Array of Independent Disks.
Volume	In Windows, this is what is given a drive letter. It is created on a partition or group of partitions.

Conclusion

Capacity expansion simply adds drives to an array. This may improve performance by increasing the number of spindles available. However, most customers will want to take advantage of this newly available space by either creating a new logical drive or extending the size of an existing one. An administrator can also change RAID levels or stripe sizes as needed for performance or fault-tolerance improvements. The management software included with the Smart Array Controllers makes all these functions available online to eliminate the need for administrators to shut down the server to reconfigure their storage.

Call to action

To help us better understand and meet your needs for ISS technology information, please send comments about this paper to: TechCom@HP.com.

© 2002 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft, Windows, Win32, Win32s, Win64, Windows NT, and Visual C++ are trademarks of Microsoft Corporation in the U.S. and/or other countries.

TC021003TB, 10/2002

