



Microsoft Exchange Server 2007 Best Practices Using 3PAR Utility Storage

Karl L. Swartz
Bill Plein

Document Abstract:

3PAR Utility Storage delivers cost-effective, high-performance, highly available storage that is ideal for Microsoft Exchange. This technical paper shows storage architects and Exchange administrators how to plan and implement a Microsoft Exchange 2007 deployment using 3PAR InServ Storage Servers.

Part Number: 3PAR-mesbp-wp-09.0

Table of Contents

Microsoft Exchange Server 2007 Best Practices Using 3PAR Utility Storage.....	1
1 Introduction	3
2 3PAR Utility Storage for Exchange 2007.....	3
2.1 Email and Messaging Are Mission Critical.....	4
2.2 Platform Scalability and High Performance	5
2.3 3PAR Thin Provisioning	5
3 Exchange 2007 Storage Architecture	6
3.1 Storage Groups	6
3.2 Storage-Related Improvements in Exchange 2007 Compared to Exchange 2003	6
3.3 High Availability.....	7
4 Storage Planning and Deployment.....	7
4.1 Using User I/O Profiles to Determine Performance Requirements	8
4.2 Capacity	9
4.3 Disk Type	9
4.4 RAID Level.....	10
4.5 Storage Layout Considerations	11
4.5.1 Creating Virtual Volumes Using CLI Commands	11
4.5.2 Creating Virtual Volumes Using The 3PAR GUI	12
4.6 Configuring the Exchange Server for 3PAR Storage.....	15
4.6.1 iSCSI Host Connectivity	15
4.6.2 Fibre Channel Host Connectivity.....	16
4.6.3 Partition Alignment	16
5 Testing the Solution.....	16
6 Other Microsoft Recommendations for Storage.....	17
7 Summary.....	17
References.....	18

1 Introduction

Microsoft® Exchange is one of the most popular email and collaboration suites. Because messaging and collaboration are central to an enterprise's operations, intelligent decisions must be made to ensure high availability and performance of the Exchange storage infrastructure. Properly sizing Exchange's storage infrastructure and adhering to Exchange best practices allows you to avoid performance and availability problems as the Exchange environment scales with business growth. This paper is intended to help administrators properly configure and "right-size" their 3PAR® Utility Storage environment in accordance with Microsoft Exchange 2007 best practices.

The 3PAR InServ® Storage Server is a highly virtualized, tightly clustered, dynamically tiered storage array built for utility computing. 3PAR InServ arrays use massive parallelization of all available controller nodes, physical disks, and other system resources to widely distribute volumes for optimal performance and resource utilization. The 3PAR InForm® Operating System automates the process of provisioning these massively striped and scalable volumes, known as Virtual Volumes (VVs). The combination of ease-of-use and automated provisioning with 3PAR Utility Storage yields a highly scalable, extremely adaptable, and easily manageable storage environment that is ideal for Microsoft Exchange deployments.

This paper describes the best practices to apply in the design and configuration of 3PAR Utility Storage as the storage component of a Microsoft Exchange 2007 deployment. Context for this information is provided by a discussion of the benefits of using 3PAR Utility Storage for Microsoft Exchange and a summary of the key features of the Exchange 2007 storage architecture.

2 3PAR Utility Storage for Exchange 2007

3PAR has published Exchange Server Reviewed Program (ESRP) results that demonstrate linear scalability across a wide range of Exchange Server sizes, optionally including CCR and/or Thin Provisioning ^[1,2,3,4,5], allowing IT organizations to easily determine a storage configuration that meets their needs. Additionally, the virtualized nature of 3PAR Utility Storage makes it easy to grow as demand on the messaging service increases.

Microsoft's migration to a 64-bit architecture in Exchange 2007 significantly enhances the scalability and performance of Exchange. This improvement alleviates many of the storage-related performance challenges encountered with earlier 32-bit versions of Exchange. Easing of storage I/O intensity has caused some confusion as to whether direct-attached storage (DAS) or a storage area network (SAN) provides the more appropriate storage environment for Exchange 2007 implementations. For smaller deployments, the performance enhancements and lower storage performance requirements make DAS a potentially viable option for Exchange 2007, but SAN is also suitable as long as the storage array in question is able to scale easily and handle complex, mixed workloads.

For larger deployments, a SAN is more appropriate. However, it is important to choose a SAN storage array for Microsoft Exchange 2007 that can efficiently handle mixed workloads, allowing administrators to leverage the array's ability to consolidate multiple workloads instead of using the array as an expensive DAS subsystem. Deploying a legacy SAN array that does not perform well in mixed workload environments leads to the entire array being committed only to Exchange, as the storage workload from any other application on that same array may impede service to Exchange. Using a SAN-based array as a dedicated device often yields low storage utilization

rates, high acquisition and operational costs, and the misconception that DAS is the only cost-effective option for Microsoft Exchange 2007 in large-scale environments.

Several unique features of the 3PAR InSpire[®] Architecture eliminate the need to use separate arrays to segregate workloads and enable consolidation, which is impossible with DAS. 3PAR Utility Storage delivers high transaction rates along with high throughput by leveraging its unique parallelization of data movement, which is accomplished by the 3PAR Gen3 ASIC and data cache implementation, and its powerful metadata processing, which is powered by the Intel[®] CPUs and control cache.

These 3PAR features enable administrators to meet the performance and availability requirements of their Exchange 2007 environment while concurrently serving file and SharePoint services. Organizations can use 3PAR SAN storage for both small and large Exchange deployments since IT administrators can let the Exchange environment, regardless of size, coexist on a single storage array with other applications. The ability to consolidate applications and deliver high mixed-workload performance eliminates the need to purchase separate storage and allows the IT organization to purchase, manage and maintain fewer physical resources.

2.1 Email and Messaging Are Mission Critical

The resilience of 3PAR Utility Storage simply cannot be matched by legacy arrays. Even under failure conditions, the InSpire Architecture maintains fast response times to Exchange 2007 databases through advanced error isolation and massive parallelism. InServ arrays are also able to maintain application workloads more consistently since they were designed to eliminate acute component-level dependencies.

3PAR Utility Storage achieves unmatched resilient infrastructure through the following:

- **Mesh-Active Architecture:** Unlike active-passive legacy controller architectures advertised as “active-active,” where each LUN is active on only one controller, 3PAR’s Mesh-Active architecture allows each LUN to be active on every mesh controller in the array. 3PAR’s design automatically delivers more robust, load-balanced performance for Exchange 2007 and greater headroom for mailbox size and quantity to scale.
- **Built-in Redundancy:** InServ arrays offer active hardware components and full redundancy of software. Following an extended power failure, Exchange databases can return to an operational state faster due to the protection provided by the array’s de-staging of cache data to permanent media.
- **Non-Disruptive Upgradability:** Planned email outages are a thing of the past with hardware and software upgrades that are completed online and non-disruptively to minimize impact on production performance levels.
- **Minimal Impact of Drive Failure:** In the event of a failure, rapid RAID rebuild enables reliable, rapid recovery with minimal application performance impact by rebuilding only used chunklets using a many-to-many drive relationship. Rebuilds often occur in just one-quarter the time needed for traditional arrays.
- **RAID Isolation:** By spreading the disks in a RAID group across separate disk chassis, the 3PAR array can lose an entire chassis (up to 40 disk drives) without losing access to data.

2.2 Platform Scalability and High Performance

3PAR InServ Storage Servers scale from 2 to 8 controller nodes and up to 1,280 disk drives with support for hundreds of terabytes of storage capacity. The InServ's clustered architecture uses fine-grained striping of data across massively parallel internal resources for exceptional performance and scalability.

Linear scalability for Exchange Server 2007 within a single 3PAR array is demonstrated by the 3PAR ESRP results summarized in Table 1. These results also show that the 3PAR performance increased uniformly as the mailbox count scaled, with no change in the average database read latency and a nearly imperceptible change in transaction log write latency.

Table 1. Summary of Non-CCR ESRP Results ^[1,2,3]

Mailboxes	Disk Drives	Controller Nodes	Aggregate IOPS	Database Read Latency	Log Write Latency
24,000	160	2	14,753	10 ms	0 ms
48,000	320	2	27,469	10 ms	1 ms
96,000	640	4	55,405	10 ms	1 ms

As can be seen from these results, 3PAR Utility Storage is uniquely designed to deliver cost-effective, high-performance, highly available storage, which is an ideal combination of features for any Exchange deployment.

2.3 3PAR Thin Provisioning

With legacy storage servers, Exchange administrators must routinely overprovision storage to accommodate anticipated future growth. On average, just 25% of physical capacity is used due to this overprovisioning—an enormous waste of capital, power, cooling, and floor space. 3PAR Thin Provisioning solves this problem by allowing virtual volume capacity to be allocated without actually committing physical storage. Physical storage capacity is allocated on a fine-grained basis only when data is actually written. Instead of the traditional approach of tying storage costs to allocation, 3PAR Thin Provisioning aligns storage costs with the amount of data that is actually being written. The Thin Provisioning benefit is discussed in more detail within this paper.

3 Exchange 2007 Storage Architecture

Exchange 2007 stores data in a repository known as an Exchange Store^[6], equivalent to the Information Store used by Exchange 2003^[7]. Microsoft defines an Exchange Store as containing the following components^[8]:

- Storage groups
- Mailbox databases
- Public folder databases

Exchange 2007 Enterprise Edition supports up to 50 storage groups and 50 databases; these limits are reduced to five each in Exchange 2007 Standard Edition^[6,8]. One public folder database is supported per server for either edition.

Storage groups are the containers for the two types of databases and thus are the key components of interest for storage planning. The remainder of this document will therefore focus on storage groups.

3.1 Storage Groups

An Exchange storage group (SG) is a logical container composed of one or more databases and a transaction log. While multiple databases per SG are allowed, Microsoft recommends having only one database per SG. This recommendation becomes a requirement if Cluster Continuous Replication (CCR) is implemented^[9].

Microsoft requires that Exchange databases and their associated transaction logs (i.e., those contained in the same storage group) be segregated on different physical disks. The rationale for this is data availability: a double disk failure, no matter how unlikely, can cause data loss if it affects both a database and its corresponding transaction log. With disk segregation, recovery from a database failure can be accomplished by restoring the database from backups and then applying any changes from the transaction logs. Transaction log failures require an even simpler recovery process: a clean shutdown upon failure of the transaction log will cause all data changes to be written to the database so there is no further need for the transaction log.

3.2 Storage-Related Improvements in Exchange 2007 Compared to Exchange 2003

The most significant improvement Exchange 2007 incorporates over Exchange 2003 is a switch from a 32-bit to a 64-bit version of the Windows operating system. Using 64-bit addresses allows larger data structures which can take full advantage of the larger physical memory, leading to a reduction in storage I/O.

The database cache is one of the most important beneficiaries of the increased 64-bit addressing available with Exchange 2007. Exchange 2003 limited the database cache size to 900 MB by default, with an option to expand to 1.2 GB on systems with 2 GB or more of physical memory^[10,11]. On an Exchange server with 3,000 mailboxes—half the number of mailboxes per simulated server used in 3PAR's initial Exchange 2007 ESRP results—that limit translates to 300-

400 KB of database cache per mailbox^[1,2,3]. This is less than the size of many attachments. In contrast, the database cache size is effectively unlimited on Exchange 2007, allowing it to maximize the use of 32-GB or larger memories commonly found on current servers.

Serving database read requests from cache results in a dramatically lower rate of random disk reads. The improvement depends on a variety of factors—including user profile, mailboxes per server, and physical memory per server—but Microsoft has documented that Exchange 2007 provides 4:1 or better reductions in read I/O with reasonable workloads^[10,12].

Due to increased caching, a site running Exchange 2003 should be able to transition to Exchange 2007 and expect the same or lower storage I/O, assuming the number of users and mail usage remains the same. Because Exchange 2007 is the underpinning for many new communication services from Microsoft (voice, conferencing, messaging), the opportunity for additional I/O workload will increase with Exchange 2007 as these services are implemented by the customer. Addition of other workloads increases the importance of a storage system that is developed from the ground up to handle mixed workload I/O without sacrificing the performance of applications such as Exchange.

3.3 High Availability

Exchange Server 2007 offers a number of built-in features that can provide quick recovery, high availability, and site resilience for Exchange mailbox servers. These include Local Continuous Replication (LCR), Cluster Continuous Replication (CCR), Standby Continuous Replication (SCR), and Single Copy Clusters (SCC)^[13]. Large sites will most likely be deployed with CCR since it can be deployed in two datacenters for disaster recovery. CCR also provides rapid failover, with failover from the current active node to its passive partner typically taking under two minutes^[14].

With 3PAR InServ Storage Servers, storage is highly available regardless of the HA approach selected for the Exchange servers. With properly configured multipath I/O (MPIO), the failure of any component in an InServ does not impact data availability.

4 Storage Planning and Deployment

Raw capacity is an obvious consideration when planning storage for an Exchange 2007 server. With many storage systems, performance requirements demand more disk drive spindles for I/O performance than are needed for capacity requirements. The 3PAR InSpire Architecture widely stripes volumes to evenly distribute I/O over all available disks within the array, eliminating hotspots or “hot” spindles. As a result, the storage administrator can focus on acquiring the quantity of spindles needed to meet capacity requirements rather than acquiring additional spindles to meet performance requirements.

The remainder of this section discusses User I/O Profiles and Microsoft’s Exchange Calculator and how to use these tools to determine storage performance and capacity requirements. The influence of different disk types and RAID levels on Exchange storage then follows. Moving up the storage hierarchy, the simple steps to provision volumes for Exchange on a 3PAR array are presented, followed by the steps required to configure these volumes on the Exchange server(s).

4.1 Using User I/O Profiles to Determine Performance Requirements

The first step in sizing an Exchange Server 2007 deployment is to determine a User I/O Profile. Combined with the number of users (or equivalently, mailboxes), this profile can be used to determine the number of Exchange IOPS required, which in turn can be used to size the storage from a performance standpoint. Microsoft provides an Exchange 2007 Mailbox Server Role Storage Calculator to assist in this analysis^[15].

Ideally, data from an enterprise's own users would be used to develop the User I/O Profile. Site-specific data may not be readily available, so Microsoft provides guidelines for several types of typical users, based on measurements of live Exchange servers. These profiles are shown in Table 2.

Table 2. Typical User I/O Profiles

Mailbox User Type	Message Profile	Logs Generated per Mailbox per Day	I/O per User per Second
Light	5 sent / 20 received	6	.08
Medium	10 sent / 40 received	12	.16
Heavy	20 sent / 80 received	24	.32
Very Heavy	30 sent / 120 received	36	.48

To determine the overall Exchange I/O rate, the I/O per User per Second (IOPS) metric—as determined through site analysis, Microsoft's calculator, or using the values in Table 1—should be multiplied by the number of mailboxes. For example, selecting the Very Heavy profile from the table results in an estimate of .48 I/Os per user per second. Serving 24,000 Very Heavy users requires .48 times 24,000 or about 11,520 Exchange IOPS.

Keep in mind that not all users are the same, so it may be appropriate to calculate a profile that is a composite of several types of users. Also, use of PDAs like the Blackberry will increase the I/O load for associated mailboxes. Microsoft's discussion of disk I/O causes for Exchange 2007 notes that an increase of two to four times was seen with Exchange 2003^[16].

4.2 Capacity

While it seems that capacity planning can be simply addressed by only considering the mailbox sizes that need to be deployed (i.e., multiplying the number of mailboxes by the average mailbox size), other important considerations must be made. Additional critical considerations include the ever-increasing size of the average size mail message over time, growth in the number of users (and thus mailboxes) over time, and transaction log retention times. Planning for growth must be included in planning storage capacity for an Exchange server. Microsoft's Exchange Calculator includes these and other factors which influence Exchange storage capacity.

Many customers tackle capacity growth needs by substantially oversizing the initial storage purchase and deployment. This is expensive since it requires buying storage that is not initially required, and may never be used. On the flip side, under-provisioning can lead to difficult (and expensive) upgrade challenges in the future.

3PAR Thin Provisioning offers a way to plan for this growth without costly initial overprovisioning. Using Thin Provisioning software, database and log volumes can be provisioned and presented to the Exchange server(s) at their full size, but the underlying physical storage is not committed until data is actually written. For example, if mailboxes are limited to 1 GB and there are 2,000 mailboxes (users) on RAID 1 storage, a traditional storage array would require 2 TB of data space up front (1 GB times 2,000 mailboxes) using 4 TB of physical storage (2 TB times two for the RAID 1 mirror). In this example, if the average mailbox size turns out to be only 250 MB, then only 1 TB of physical storage was actually required (500 GB of data and 500 GB of mirror) and 3 TB of physical storage has gone unused. In this example, if Thin Provisioning is deployed, the server still sees volumes totaling 2 TB of logical space, but only 1 TB of physical storage is allocated and used, which provides a savings of 75% compared to the 4 TB used with non-TP volumes.

A large university with nearly 15,000 mailboxes used 3PAR Thin Provisioning when migrating to Exchange 2007 and was able to increase mailbox size limits fivefold while deploying only 25% of the storage that would have been needed without Thin Provisioning^[17].

Storage reduction of this order might be expected to come at a significant cost in performance, but 3PAR's published Exchange 2007 ESRP results demonstrate an insignificant performance impact from deploying Thin Provisioning^[5,4]. Two 24,000-mailbox CCR runs were published, differing only in the use of fully-provisioned Virtual Volumes (VVs) in one report and Thin Provisioned Virtual Volumes (TPVVs) in the other. Average database disk read latency was the same (13 ms) in both tests, comfortably below the 20 ms limit required to pass a Jetstress 2007 test^[18]. Thin Provisioning increased average transaction log disk write latency from 0 ms (less than 0.5 ms, rounded by Jetstress) to 1 ms, far below the 10 ms limit.

4.3 Disk Type

E-mail has become a critical resource for most organizations and poor e-mail performance can quickly and severely have a negative impact on the organization. Recognizing that performance is paramount, 3PAR's Exchange 2007 ESRP submissions used fast, 15k-RPM, 146-GB Fibre Channel disks in a RAID 1 configuration. Even with a simulated Very Heavy User I/O Profile (Table 1), with an additional 20% I/O to simulate periods of unusually high activity, average disk latency was well below Microsoft's recommended limit (see Section 4.2). Actual disk IOPS for the database disks were slightly over half of their nominal capability. This suggests that half as many 15k-RPM, 300-GB disks (providing approximately the same total capacity) could handle a similar

workload with only a modest increase in response time. Alternately, larger disks could be used for additional capacity for group file servers or other services with lower performance demands while maintaining Exchange performance.

Transaction logs present a sequential workload to the storage system with a very low IOPS requirement. This workload fits well with the performance characteristics of Nearline (enterprise-class SATA) disks. 3PAR InServ arrays readily accommodate mixing of disk technologies and can be configured with both Fibre Channel and Nearline disks. In an array configured with both types of disks, placing the transaction logs on Nearline disks is an option. Additional Nearline capacity can also be used for consolidation of archival data.

Nearline disks may also provide adequate performance for Exchange databases which have a more modest IOPS requirement than that simulated in 3PAR's Exchange 2007 ESRP testing. The Exchange Calculator may be used to determine whether Nearline disks are a viable option for further reducing the cost per gigabyte.

4.4 RAID Level

RAID 1 mirrors the data written to one disk to a second disk. Either disk can fail and the data will still be available, but physical storage exceeds usable storage by 100%. RAID 5 provides the redundancy needed to recover from a failed disk at lower cost by grouping the data from multiple disks and calculating the parity of that group, then writing it to an additional disk. For example, if three disks of data are in the group, four disks are required; this 3+1 configuration requires only 33% more physical storage than usable storage.

Microsoft recommends using RAID 1 due to the lower performance of most vendors' RAID 5 implementations^[19]. This diminished performance is caused by the need to read two or more blocks in order to recompute parity before writing a single data block and its group's parity block. However, 3PAR's Fast RAID 5 combines three technologies to deliver high performance RAID 5: wide striping to reduce spindle contention, a large battery-backed memory cache, and a high-performance ASIC for data movement and parity calculation. Due to this unique combination of features, 3PAR's Fast RAID 5 minimizes the performance impact typical of RAID 5 on legacy storage architectures. As a result, 3PAR's RAID 5 is a viable option for customers to implement with Exchange.

Internal tests performed by 3PAR show that using RAID 5 (7+1)—i.e., seven data blocks per parity block—demonstrated that the same simulated Exchange workload used for Exchange 2007 ESRP testing had disk latencies that were higher than RAID 1 but well within Microsoft's recommendations^[19]. The internal tests demonstrated that disk IOPS for the database disks increased by about 50% to approximately 80% of the nominal capability of the disks. (As noted in Section 4.3, the disks used for 3PAR's submitted Exchange 2007 ESRP results only saw a little over half the IOPS of which they are capable).

The reality is that few Exchange environments consist entirely of users who match the Very Heavy I/O profile as tested in 3PAR's ESRP submissions. With a less demanding User I/O Profile, RAID 5-based response times approaching those demonstrated in the Exchange 2007 ESRP results using RAID 1 can be expected with much less raw disk space.

4.5 Storage Layout Considerations

Microsoft requires that Exchange databases and their associated transaction logs be segregated on different disks to avert data loss in the event of a double disk failure (Section 3.1). A Common Provisioning Group (CPG) is essentially a template or rule for configuring similar volumes within a 3PAR array. To meet Microsoft's requirements, 3PAR recommends creating two CPGs, one for database virtual volumes (VVs) and another for transaction log VVs, using filters to ensure that the database and transaction log CPGs do not share physical disks. Use the storage requirements calculator to determine the required ratio of database and transaction log sizes (Section 4.1).

4.5.1 Creating Virtual Volumes Using CLI Commands

Creating two CPGs on separate disks is simple on 3PAR InServ arrays. For instance, in a 240-spindle system, the administrator might wish to have a 9:1 ratio of disks for databases and for transaction logs. Creating a CPG for databases using the first nine magazines of each disk chassis and a second CPG for transaction logs using the tenth magazine of each disk chassis can be easily and quickly accomplished using the following 3PAR InForm Operating System CLI commands:

```
InServ cli% createcpg -p -mg 0-8 e2007db_cpg
InServ cli% createcpg -p -mg 9 e2007log_cpg
```

After creating the CPGs, if four databases were needed, the database VVs could then be created using the following commands:

```
InServ cli% createvv e2007db_cpg e2007dbvv1 100g
InServ cli% createvv e2007db_cpg e2007dbvv2 100g
InServ cli% createvv e2007db_cpg e2007dbvv3 100g
InServ cli% createvv e2007db_cpg e2007dbvv4 100g
```

Creating Thin Provisioned Virtual Volumes (TPVVs) just requires one additional option to the command:

```
InServ cli% createvv -tpvv e2007db_cpg e2007dbvv1 100g
InServ cli% ...
```

A command similar to the one used to create the database VVs is used to create the transaction log VVs:

```
InServ cli% createvv e2007log_cpg e2007logvv1 11g
InServ cli% ...
```

Because of 3PAR's simple provisioning feature, the above steps to provision the CPGs and VVs can be completed in little more time than it takes to enter the commands. Once completed, all of the VVs can then be presented to the Exchange servers by creating appropriate Virtual LUNs (VLUNs) as in the following examples:

```
InServ cli% createvlun e2007dbvv1 1 e2007server
InServ cli% createvlun e2007logvv1 2 e2007server
InServ cli% createvlun e2007dbvv2 3 e2007server
InServ cli% createvlun e2007logvv2 4 e2007server
```

For larger deployments it should be noted that a single CPG may contain a maximum of 64 VVs, with a recommended limit of 32. While the CPG creation examples above are quite simple, a more complex disk configuration than homogenous, full disk chassis or a ratio of database space to transaction log space greater than 9:1 might require several additional arguments to the CPG-creation command. In these cases, to simplify the task and to reduce the opportunity for errors, a CPG template should be used. The following example illustrates how to create a template for database CPGs and then create two CPGs using that template:

```
InServ cli% createtemplate cpg -p -mg 0-8 e2007db_cpg_t
InServ cli% createcpg -templ e2007db_cpg_t e2007db_cpg1
InServ cli% createcpg -templ e2007db_cpg_t e2007db_cpg2
```

4.5.2 Creating Virtual Volumes Using The 3PAR GUI

Configuration of the storage may also be performed using the 3PAR InForm Management Console (IMC) graphical interface. Figures 1 through 3 show the steps required using the 3PAR GUI to create the database CPG on the first nine disk magazines (equivalent to the command line **createcpg -p -mg 0-8 e2007db_cpg**). Limiting the disk magazines used for the database CPG is done to ensure that the database and transaction log CPGs do not share physical disks, as described at the start of this section.

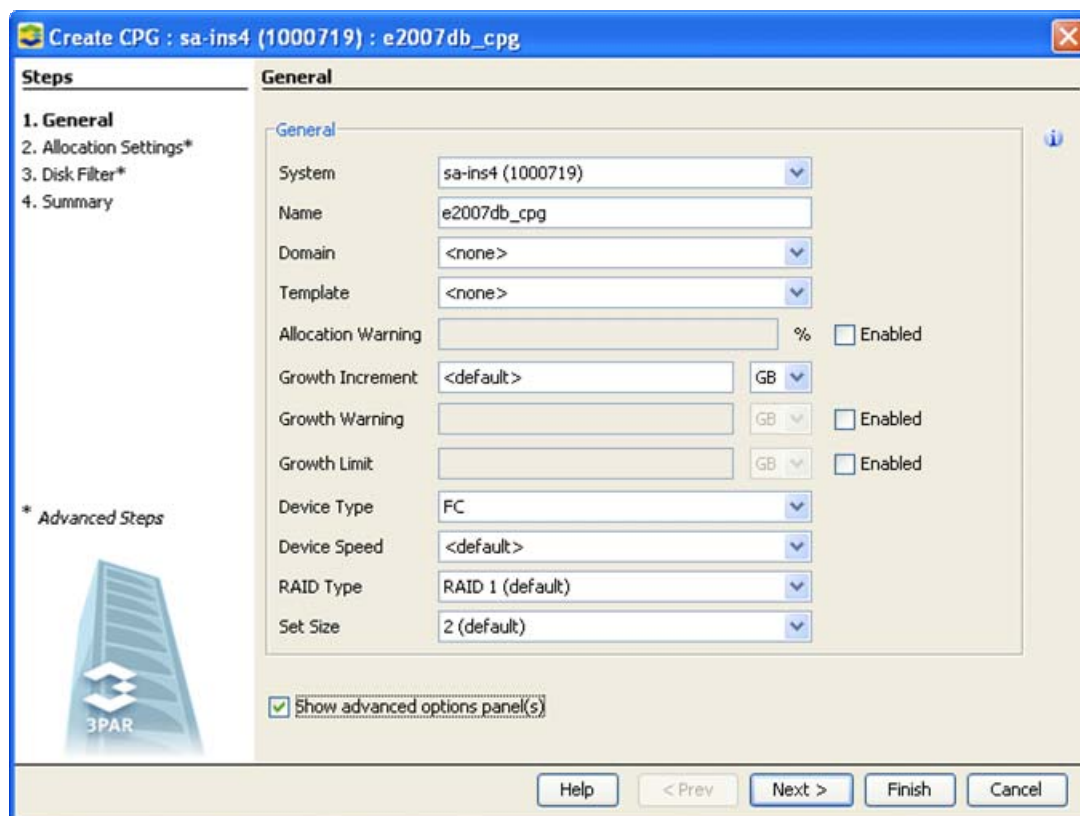


Figure 1. Create a CPG for Databases, Step 1: Define General Parameters for CPG

Note that “Show advanced options panel(s)” has been selected. This includes the Disk Filter step in the workflow so disk selection may be filtered to include only those disks on magazines 0 through 8 (Figure 2), leaving disks on magazines numbered 9 for use by transaction logs.

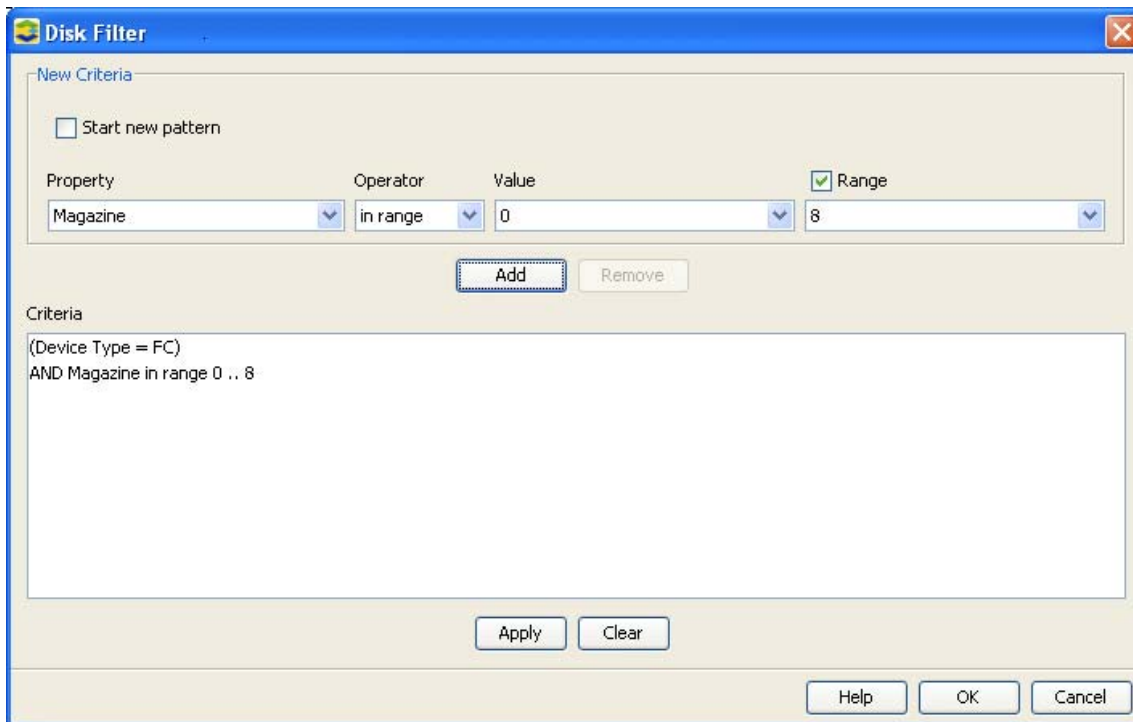


Figure 2. Create a CPG for Databases, Step 2: Filter Disks to Select Magazines

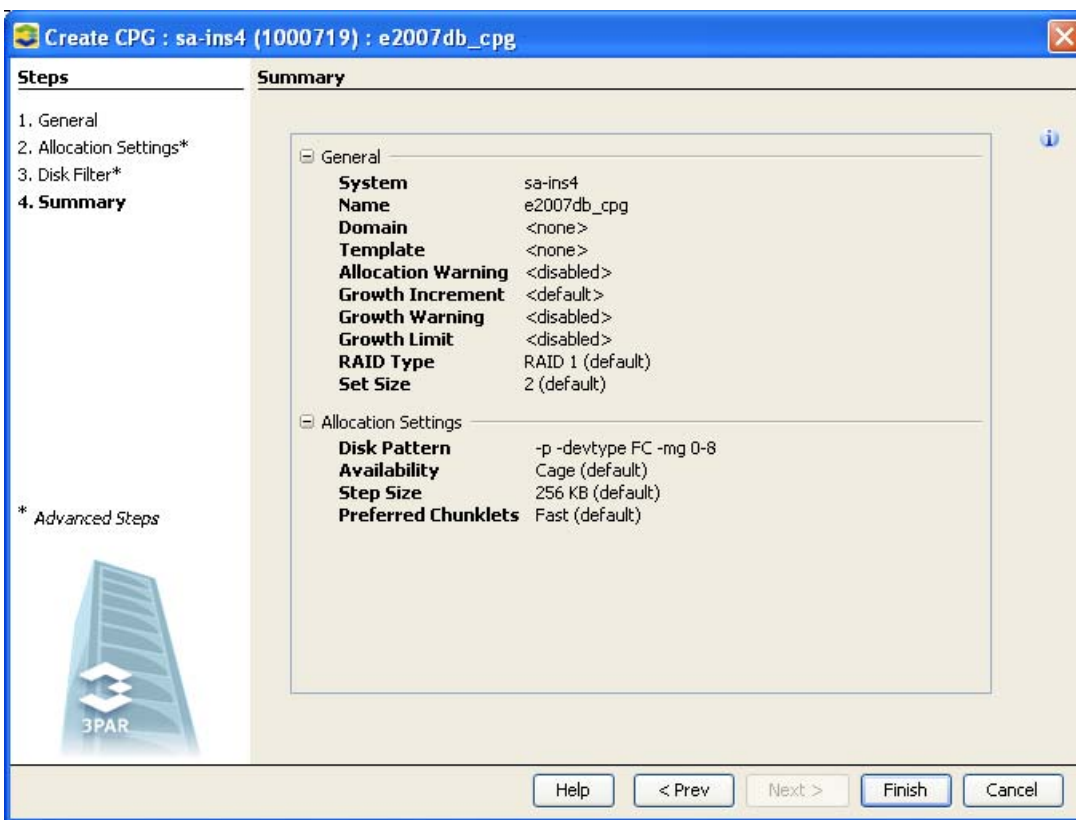


Figure 3. Create a CPG for Databases, Step 3: Review and Finish

Once the database CPG has been defined, a database virtual volume (VV) can be created in a single step as illustrated in Figure 4 (equivalent to `createvv e2007db_cpg e2007dbvv1 100g`).

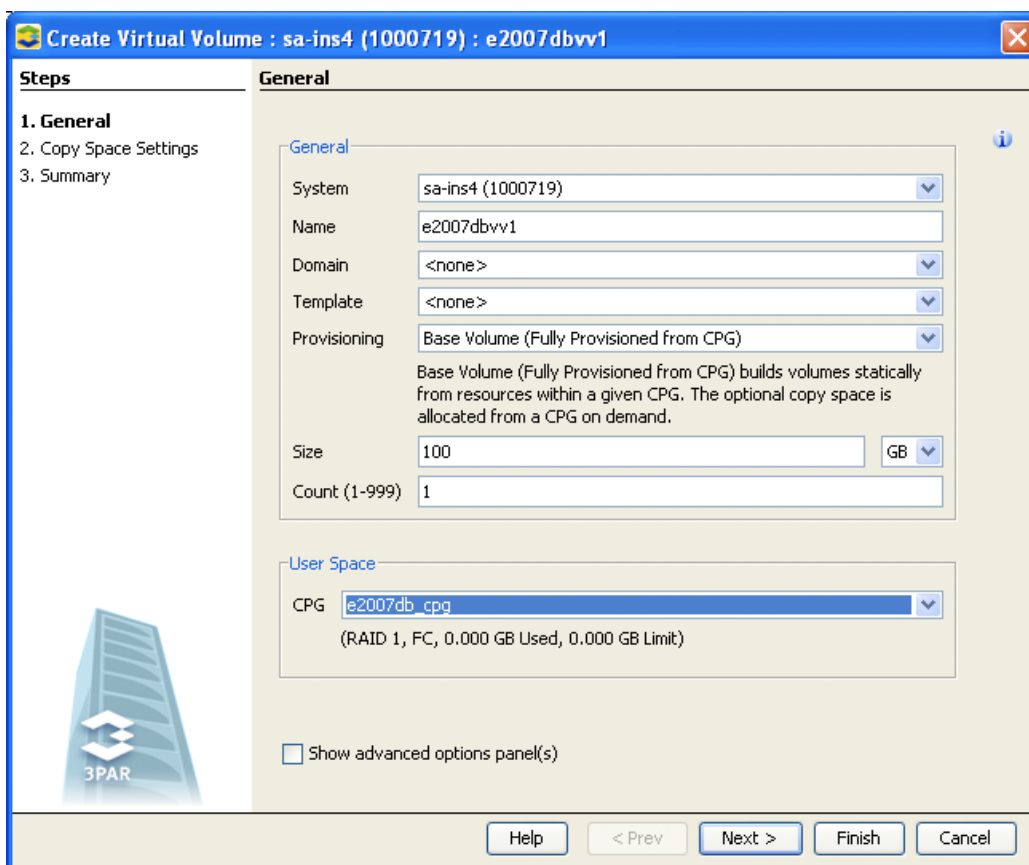


Figure 4. Create a Database Virtual Volume (VV)

Using the GUI, similar steps can be used to create the CPG and VV(s) for the transaction log.

4.6 Configuring the Exchange Server for 3PAR Storage

There are several storage-specific configurations that the administrator needs to consider when configuring each Exchange Server. Important items are selecting and configuring host connectivity (iSCSI or Fibre Channel) and setting partition alignment.

4.6.1 iSCSI Host Connectivity

iSCSI provides a cost-effective host connectivity option for environments with relatively low I/O requirements and no Fibre Channel SAN infrastructure. If iSCSI is selected to attach the Exchange Server to the 3PAR array, follow the procedures in the 3PAR iSCSI Implementation Guide for Windows Server 2003^[20] or for Windows Server 2008^[21] to properly configure the Microsoft iSCSI Initiator software. It is important to make sure that active-active multipath I/O has been enabled so a single Ethernet connectivity failure will not cause an outage.

4.6.2 Fibre Channel Host Connectivity

Fibre Channel host connectivity provides high I/O performance for busy Exchange servers. If Fibre Channel is selected to attach the Exchange Server to the 3PAR array, multipath I/O (MPIO) software should be installed and configured:

- Windows Server 2003 requires installation of 3PAR Multipath I/O software ^[22].
- Windows Server 2008 includes Microsoft Multipath I/O software, which should be installed and configured using the instructions in the 3PAR Implementation Guide for the HBA type installed in the server.

MPIO provides redundancy so a single Fibre Channel connectivity failure will not cause an outage. In active-active mode, MPIO also offers higher data throughput because it supports dynamic load balancing.

Microsoft recommends increasing the queue depth—the number of concurrent Fibre Channel I/Os which may be outstanding—for Exchange servers to assure adequate performance. However, testing by 3PAR has demonstrated that the default value that Microsoft provides for queue depth does not need to be changed. For further information consult the 3PAR Implementation Guide for the HBA brand and models implemented on the Windows Server hosting the Exchange Server.

4.6.3 Partition Alignment

To ensure good performance on Windows Server 2003, the Windows Diskpart utility ^[23] should be used to align all Exchange-related volumes before formatting. Specify a value of 16 to align the NTFS partitions with 3PAR's 16-KB cache page. The specific command within Diskpart is `create partition primary align=16`. (Windows Server 2008 automatically aligns partitions on a 1024-KB boundary.)

5 Testing the Solution

Microsoft provides Jetstress as “a tool that helps administrators verify the performance and stability of the disk subsystem before putting their Exchange server in a production environment” ^[24]. For the results to be meaningful, the mailbox count and size should be configured in Jetstress to match the planned Exchange deployment. Default settings should be used elsewhere except for thread count, which may require tuning to determine maximum performance of the system.

The key storage-related metrics in the Jetstress output are those related to average database read and transaction log write latency. Microsoft specifies that these latencies must not exceed 20 ms and 10 ms respectively ^[18] and Jetstress will fail if these values are exceeded. In 3PAR's published Exchange 2007 ESRP results, average database read latencies are 10-13 ms while average write latencies (both database and transaction log) are 0-2 ms ^[1,2,3,4,5]. Configurations using slower disks and/or RAID 5 will tend to have higher latencies, while a lighter workload may result in lower latencies.

6 Other Microsoft Recommendations for Storage

This document details storage considerations and Microsoft's storage recommendations for Exchange Server 2007. In addition to this paper, review Microsoft's detailed planning documents carefully when implementing Exchange Server 2007 ^[25].

7 Summary

Many organizations use Microsoft Exchange Server 2007 as their mission-critical communications tool, as it offers enormous productivity gains by enabling workplace collaboration. Given the importance of Exchange, organizations devote significant amounts of storage capacity and management resources to this critical application. Even so, many IT groups still struggle with Exchange scalability, complex mechanisms to provide high-availability, infrequent or extended backups, and unrealized disaster recovery plans.

3PAR InServ Storage Servers provide an ideal storage foundation upon which to build a consolidated Microsoft Server environment that addresses the needs of email, collaboration, and file services. 3PAR's massively-parallel architecture and support for mixed workloads provides benefits that include

- High, predictable, and linearly scalable performance of 3PAR Utility Storage with Exchange Server 2007.
- Improved capacity utilization without performance degradation using 3PAR Thin Provisioning.

The ability to scale Exchange Server 2007 databases simply and confidently within a single, massively scalable storage platform enables any organization to reduce the costs and complexities associated with the maintenance and administration of multiple systems. The InServ's scalability allows the consolidation of additional services such as group shares on the same storage platform, delivering substantial cost savings. In addition, 3PAR Thin Provisioning software helps achieve further cost savings and maximize efficiency without sacrificing performance. Avoiding the costs and complexities of purchasing and maintaining multiple storage systems can mean the difference between struggling and surviving in these lean economic times.

References

- [1] 3PAR InServ T800 24,000 Mailbox Exchange 2007 Storage Solution, January 2009.
- [2] 3PAR InServ T800 48,000 Mailbox Exchange 2007 Storage Solution, January 2009.
- [3] 3PAR InServ T800 96,000 Mailbox Exchange 2007 Storage Solution, January 2009.
- [4] 3PAR InServ T800 24,000 Mailbox Clustered Continuous Replication Exchange 2007 Storage Solution, February 2009.
- [5] 3PAR InServ T800 Thin-Provisioned 24,000 Mailbox Clustered Continuous Replication Exchange 2007 Storage Solution, February 2009.
- [6] Understanding the Exchange 2007 Store, <http://technet.microsoft.com/en-us/library/bb331958.aspx>, February 2007.
- [7] Exchange Information Store Service Architecture (Exchange 2003), <http://technet.microsoft.com/en-us/library/bb123916.aspx>, May 2005.
- [8] Logical Components of the Exchange Store, <http://technet.microsoft.com/en-us/library/bb331956.aspx>, March 2007.
- [9] Planning for Cluster Continuous Replication, <http://technet.microsoft.com/en-us/library/bb123996.aspx>, July 23, 2008.
- [10] Schnoll, Scott: Exchange Server 2007 Storage Changes and Design Considerations (UCM309), Microsoft Corporation, 2007.
- [11] The Microsoft Exchange Team Blog: Understanding Exchange 2007 Memory Usage and its use of the Paging File, <http://msexchangeteam.com/archive/2008/08/06/449484.aspx>, August 6, 2008.
- [12] The Microsoft Exchange Team Blog: Understanding Exchange 2007 I/O improvements from 64 bit, <http://msexchangeteam.com/archive/2006/09/08/428860.aspx>, September 8, 2006.
- [13] High Availability, <http://technet.microsoft.com/en-us/library/bb124721.aspx>, March 21, 2009.
- [14] Cluster Continuous Replication, <http://technet.microsoft.com/en-us/library/bb124521.aspx>, March 21, 2008.
- [15] The Microsoft Exchange Team Blog: Exchange 2007 Mailbox Server Role Storage Requirements Calculator, <http://msexchangeteam.com/archive/2007/01/15/432207.aspx>, January 15, 2007.
- [16] What Causes Exchange Disk I/O, <http://technet.microsoft.com/en-us/library/bb738154.aspx>, March 18, 2009.
- [17] University Upgrades Messaging Software to Raise E-Mail Performance and Mailbox Sizes (Microsoft Windows Server Systems Customer Solution Case Study), April 2009.

- [18] Reading the Test Reports, <http://technet.microsoft.com/en-us/library/bb643098.aspx>, February 15, 2008.
- [19] Planning Storage Configurations, <http://technet.microsoft.com/en-us/library/bb124518.aspx>, October 22, 2007.
- [20] InForm OS 2.2.3 Windows 2000/2003 x86, x64, IA64 iSCSI Implementation Guide, 2008.
- [21] InForm OS 2.2.3/2.2.4 Windows 2008 x86, x64 iSCSI Implementation Guide, 2008.
- [22] 3PAR Multipath I/O 1.0.22 for Microsoft Windows User's Guide for (IA64, x64, x86), 2009.
- [23] A Description of the Diskpart Command-Line Utility, <http://support.microsoft.com/kb/300415>, December 1, 2007.
- [24] Microsoft Exchange Server Jetstress, <http://technet.microsoft.com/en-us/library/bb643093.aspx>, July 16, 2007.
- [25] Planning Your Server and Storage Architecture, <http://technet.microsoft.com/en-us/library/bb738142.aspx>, March 21, 2008.

About 3PAR

3PAR[®] (NYSE: PAR) is the leading global provider of utility storage, a category of [highly virtualized](#), tightly clustered, and dynamically tiered storage arrays built for utility computing. Organizations use utility computing to build cost-effective virtualized IT infrastructures for flexible workload consolidation. 3PAR Utility Storage gives customers an alternative to traditional arrays by delivering resilient infrastructure with increased agility at a lower total cost to meet their rapidly changing business needs. As a pioneer of thin provisioning—a green technology developed to address storage underutilization and inefficiencies—3PAR offers products designed to minimize power consumption and promote environmental responsibility. With 3PAR, customers have reduced the costs of allocated storage capacity, administration, and SAN infrastructure while increasing adaptability and resiliency. 3PAR Utility Storage is built to meet the demands of open systems consolidation, integrated data lifecycle management, and performance-intensive applications. For more information, visit the 3PAR Website at: www.3PAR.com.