Information and Configuration Guide for Digital's Desktop Storage:

Focus on the RZ Series of SCSI Disk Drives

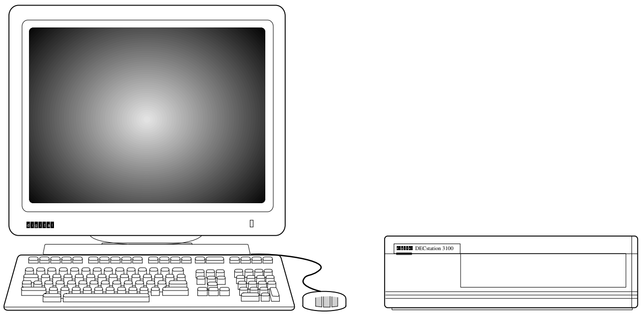

Digital Equipment Corporation
Disk and Subsystems Group

Revision Level: 6.6
August 1991

digital™

Part No: EB-M4451-78

Contents

Figures

Tables

Appendices

1. Introduction

How does one decide on a storage solution for desktop systems? With an ever increasing number of device types, sizes, capacities, and system configurations, you need to know more about the factors that impact your environment so that you can select the best initial and add-on storage devices. Some of the factors that affect these decisions are operating environment, system architecture, applications, operating system, I/O (input/output) interface, disk capacity, backup and media distribution devices, disk form factor, storage expandability, disk performance, and overall application performance. This paper discusses these factors to inform and assist you in configuring the best/most appropriate RZ disks for Digital's SCSI-based, desktop systems.

2. SCSI for the Desktop

Digital selected SCSI as the storage interface for its VAX and RISC workstations and entry-level multi-user systems. SCSI (*Small Computer System Interface)* is the evolving ANSI industry standard I/O bus interface for intelligent peripherals for desktop and deskside systems. Designed to replace the ST506 interface (developed by Seagate), SCSI was originally adopted by workstation vendors for its device independent support of higher capacity devices, greater number of devices, breadth of device types (scanners, printers, communications devices, etc.), and lower cost. SCSI-based devices have embedded, "intelligent" controllers which interpret generic device commands sent by the SCSI host adapter. By using logical addressing, SCSI is able to support the widest range of devices and functions. Due to low command overhead and seek overlap (single command per logical unit), multiple drives on a SCSI bus will demonstrate near linear I/O subsystem increases in small request, random transfers. In general, a SCSI-based subsystem offers the optimal price/performance for the needs of the single-user workstation or smaller multi-user system.

2.1 Why SCSI?

The desktop systems market supports many interfaces - both standard and proprietary. By choosing SCSI for VAX and RISC workstations and entry-level systems in 1989, Digital helped secure SCSI as the most widely accepted open bus for these markets. This is important because the more widely accepted the standard, the wider the array of supported peripherals available. Digital wanted to enable users to access peripherals such as optical readers, scanners, digitizers, and multi-media devices that provide improved productivity.

Compared to the other low-end industry standard interfaces such as ST506 and ESDI *(Enhanced Small Device Interface),* SCSI has the potential for higher performance and speed, more flexibility of features, a wider range of supported device types, and external expansion.

At the higher end of the market, IPI *(Intelligent Peripheral Interface*), which superseded SMD *(Storage Module Device*), is a more complex standard (specified in four levels) that was developed by CDC to be used primarily with larger disks and for very high throughput applications. This complexity and performance restricts this standard in terms of availability and cost.

Digital's DSSI *(Digital Storage Systems Interconnect)* was designed for medium to high-end MicroVAX systems incorporating the high performance and availability features from Digital's high-end SDI *(Storage Disk Interconnect)*. SCSI, on the other hand, was designed from the bottom-up to be fast, simple, and standard for the rapidly growing open desktop environment.

SCSI is best suited to Digital's broad range of desktop systems as it is cost effective for the basic workstation user who, for example, uses the VAXstation 3100 (asynchronous data rate of 1.5 MB/sec., up to 6.6 GB of storage) and yet is scaleable and powerful enough for the high end applications running on, for example, the DECsystem 5500 (synchronous data rate of 4.0 MB/sec., up to 28 GBytes of storage). As the SCSI-2 optional features are further implemented, additional performance enhancements and expandability will be possible.

2.2 SCSI Standards Development

2.2.1 Today

The SCSI interface was originally developed by several independent manufacturers who agreed to form a standard by combining their development efforts. In order to maintain individuality and flexibility for future developments, there are two types of specifications:

SCSI SPECIFICATION

Mandatory Features	Optional Features
- Common Commands Set (CCS)	- Voluntary implementation
- Tighter specifications	- Guidelines only
- Customizable parameters	

The SCSI specification contains guidelines for the electrical, mechanical, and functional aspects of the I/O bus. Interoperability of SCSI devices depends not only on common commands and parameters but also precise timing of signals for device communication on the bus. Because of the customizable parameters and the broad scope of the optional features, not all implementations are compatible. Each vendor may have a different set of parameters and features, such that in some cases, two SCSI devices may not work together or with another vendor's systems. Even if they do initially 'work', there is no guarantee they will do so under all conditions.

The optional features available in today's SCSI are synchronous transmission (versus the standard asynchronous), disconnect/reconnect (which enables overlapped seeks on multiple drives), and differential termination (versus the standard single-ended). Differential SCSI enables maximum cable lengths of 25 meters (versus 6 meters for single-ended termination SCSI) but because it requires extra hardware on the circuit board and different connectors, disk vendors are only beginning to implement this feature. Due to the need for industry-wide consistency, no differential drives are shipping yet.

2.2.2 Tomorrow

Since the original SCSI-1 specification was approved, several revisions of the working draft have moved the standard toward a more enhanced SCSI-2. The original SCSI command set was designed to allow generic commands to run different peripheral devices. However, these commands were too generic to allow advanced disk functions such as disk shadowing, tagged command queuing, and command reordering that users demanded. In addition, there was a need for more commands, more devices, and more features in the SCSI specification to widen the scope of devices, and to enhance performance. SCSI-2 not only makes some SCSI-1 options mandatory, it will also open the door to many additional optional features that will be increasingly required in the years to come: zero latency reads, mirroring, arrays, constant density recording, and scatter gather. The most notable SCSI-2 optional features include "fast" and "wide" SCSI. With fast SCSI, the bus throughput would increase from 4-5 MB/sec. to 10 MB/sec. With wide SCSI, added bandwidth would allow an increase from 8-bit to 16 or 32-bit data paths thereby enabling a maximum fast and wide transmission of 40 MB/sec. on the bus. These optional features are under development today, but not yet shipping.

Thus, SCSI-2 is evolving to offer both tighter specifications and greater functionality than SCSI-1, and as SCSI-3 is defined, more alterations will occur. As SCSI-2 goes through the final approval process, some vendors, including Digital, have chosen to implement some of the commands and features prior to formal approval.

Since SCSI was and is an ANSI standard based on a multi-vendor development effort, it will continue to be designed for flexibility, frequently updated, and consequently, will continue to be broad in scope and open to interpretation. To add to the complexity, SCSI is also gaining acceptance in a broader range of applications. A few implications of this are:

- There will be increasing emphasis on software level architecture: operating system driver and drive microcode enhancements to improve reliability, data integrity, maintainability, and performance (Optional features will require more complex microcode and operating system support).
- Some optional SCSI-2 features will require special cables, connectors, and controller chips which will cost more and require more sophisticated test tools.
- Although SCSI-2 basic commands and wide SCSI (requiring an extra cable) will be backward compatible with SCSI-1, differential termination (recommended with fast SCSI) is not backward compatible. Cables/connectors that support differential SCSI are not interchangeable with single-ended devices or drives.
- As faster CPU and SCSI chips become available, higher bus rates will be less tolerant of timing errors on the bus (requiring more thorough systems integration testing) and performance optimization between devices will become more important.
- With the introduction of multiple device types on the bus, there is greater chance of bus timing conflict and bus failure.

In summary, while SCSI-2 options offer higher performance and expandability, they also bring complexity. Rather than becoming increasingly easy to interchange SCSI devices, it appears there will be greater need to design compatibility up front and test more thoroughly for interoperability in fully configured systems. Due to this and the limited number of applications that take full advantage of the higher performance, system vendors are not yet supporting SCSI-2 optional features and few devices are available in the market. Digital continues to participate actively in the SCSI ANSI committee to refine the standards and to assess demand as it emerges.

2.3 Digital's Implementation and Enhancements

Peripheral and system product developers within Digital have worked together to agree on the interpretations of standards for optimal compatibility. Digital's hardware and software is fully SCSI-1 compatible and adheres to the all the mandatory features of revision 10d of the ANSI SCSI-2 specification draft. In addition, Digital encourages non-storage peripheral vendors to develop fully compatible devices for Digital's systems through a development program. (For more details on Digital's VMS SCSI implementation and developer programs, order the *Small Computer System Interface, An Overview*, Part # EK-SCSIS-OV.)

As discussed in the previous section, consistent SCSI specifications alone do not ensure interoperability. System and devices must undergo extensive, focused testing of all offered configurations. Testing is important to detect bugs in the operating system drivers and drive level microcode (only a portion of which is SCSI specified) as well as sub-optimal performance of peripheral devices. Digital has an extensive suite of custom-designed tests that fully exercise the capabilities of the device, isolate hard to find problems, and ensure full interoperability and optimized performance of peripherals with the systems Digital sells.

Besides complete compatibility assurance, Digital designs its' SCSI drives to exceed industry standards in many aspects. These include software/firmware level enhancements unrelated to the SCSI specifications as well as hardware enhancements:

2.3.1 Hardware Reliability Features

Digital demands high standards for all storage products. Our SCSI devices are no exception. By meeting or exceeding the high corporate standards for environmental tolerance, these disk drives exceed most typical drives in the industry - of particular importance in the desktop environment where systems get knocked, bumped, or moved. For example:

- Operating Shock: 10 G's (vs. industry norm of 2 G's)
- Thermal: 5 to 55 degrees C (vs. industry norm of 10 to 45 degrees C)
- Altitude: 15,000 feet (majority) (vs. industry norm of 10,000 ft.)
- Conformance to all regulatory standards- for all systems in all countries that we sell in. (UL, IEC, FCC, VDE, CSA,VCCI, EN, TUV, etc.)
- Tight electrical specifications to ensure signal integrity and thus data integrity (cables, lengths, number of devices and positions are specified)

2.3.2 Software/Microcode Features

In addition to physical drive quality, Digital designs in performance, data integrity, and maintainability features (some of which are optional SCSI-2 features) through microcode and software which include:

Performance:

- Large track/read-ahead drive buffers (i.e. segmented 256 KB RZ57 buffer)
- Optimal buffer management algorithms for improved throughput

- Disconnect/reconnect support for improved multi-drive performance

Data Integrity and Maintainability of disk and data:

- Full sector write completed on power-down (or reset) prevents data loss (most drives will stop in middle of write)
- Full Write and Read head and media diagnostics on power-up (typical industry drives perform only read test)
- Drive self-diagnostic coverage is greater than 90% (typically about 50%)

2.3.3 Quality, Reliability, and Maintainability Features

Finally, Digital also demands quality in its' drives design, manufacture, and serviceability:

Quality in Design and Manufacture:

- Digital developed sophisticated test tools to detect all microcode problems
- Complete systems integration testing for all configurations of Digital's drives and systems. This includes verification of supported electrical configurations.
- Complete Design and Manufacturing Process Verification testing on every product

Reliability throughout product life:

- MTBF (Mean Time Between Failure) demonstration in lab test beds prior to shipping the first drive (most vendors estimate field MTBF based on statistical inference)
- Customized test beds at vendor to ensure consistent interpretation of specification
- On-going reliability testing throughout product life maintains standards throughout product life. Only the drives that meet our specifications are shipped to our customers

Maintainability of user data:

- One year factory warranty included with list price
- Field replaceable PCB (Printed Circuit Board) enables failed PCB to be replaced at user site to preserve data on the disks
- Bad-block revectoring managed and reported by operating system (typically drive managed revectoring not reported for user tracking) so users can take preventative measures
- Fault reporting and system monitoring of both hard and soft errors, bad blocks, and retries allows for preventative maintenance by user (typically only hard errors reported)

3. Digital's SCSI Solutions

It isn't enough to have high quality devices in the SCSI system market. A complete line of products is also required to complement the broad range of SCSI-based systems offered by Digital. The following tables (Tables 3-1 and 3-2) summarize the mass storage, backup, and load devices that are available for the Digital workstations and entry level multi-user VAX systems: the VAXstation 3100 (Models 30, 38, 40, 48, and 76), the DECstation 3100, 2100, and 5000, the DECsystem 5000, 5100, and 5500, and the MicroVAX/VAXserver 3100. These tables identify the primary use of the device as well as whether they are supported internal or external to the system box. Note: VAXstation 3520 and 3540 are not included in the tables; they support the RZ55 and RZ56 internal to the system box.

Table 3-1 SCSI Mass Storage Solutions for Digital's Desktop Systems

Fixed Drives:	Function	VS3100	MV3100 DS21/31/5100 DS5000/100	DS5000 Model 200	DS5500
RZ22 (52 MB)	Swapping Disk	Internal	N/A	N/A	N/A
RZ23L (121 MB)	VMS System Disk ULTRIX Swap Disk	Internal	Internal	N/A	N/A
RZ24 (209 MB)	System/data disk	Internal	Internal	N/A	N/A
RZ25 (426 MB)	System/data disk	External	Internal	External	N/A
RZ55 (332 MB)	System/data disk	External	External	External	N/A
RZ56 (665 MB)	System/data disk	External	External	External	Int/External
RZ57 (1.0 GB)	System/data disk	External	External	External	Int/External

The following are Digital's SCSI backup and software distribution devices. The remainder of this paper focuses only on performance aspects of the mass storage devices.

Table 3-2 SCSI Backup and Distribution Solutions for Digital's Desktop Systems

	Function	VS3100	MV3100 DS21/3100 DS5100	DS5000	DS5500
Diskette Drives:					
RX23 (1.44 MB) (3.5" diskette)	Third Party S/W Dist.* Load, Transfer, Backup	Internal	Internal	External	N/A
RX33 (1.2 MB) (5.25" diskette)	Third Party S/W Dist.* Load, Transfer, Backup	External	External	External	N/A
Tape Drives:					
TZ30 (95 MB)	Backup & S/W Dist.*	Internal	Internal	External	N/A
TK50Z (95 MB)	Backup & S/W Dist.*	External	External	External	Internal
TLZ04 (1.2 GB)	Backup	External	External	External	Internal
TZK10 (525 MB)	Backup	External	Int/External	External	N/A
TKZ08 (2.2 GB)	Backup	N/A	External**	External	N/A
TSZ05 (40 MB)	Backup	External	External	External	N/A
TSZ07 (140 MB)	Backup	External	External	External	N/A
CD ROM:					
RRD42 (600 MB)	S/W Distribution*	External	Internal	External	External

*Software Distribution ** N/A on the MV3100

4. *I/O Subsystem Performance- Metrics and Variables*

With so many storage devices available, we're back to the question: How does one decide on a storage solution? As stated in the introduction, key factors that affect the decision include operating environment, system architecture, applications, operating system, I/O interface, disk capacity, disk form factor, storage expandability, disk performance, and overall application performance. Although we can't cover all these topics in this paper, we will provide the basic information to help you understand and optimize your disk performance.

4.1 Traditional Performance Metrics

4.1.1 Drive Components

The basic components of request processing at the drive level are as diagrammed here:

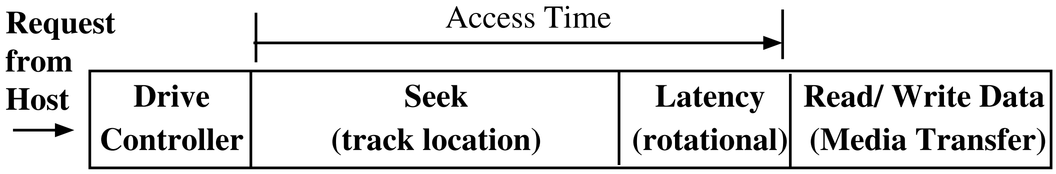

Figure 4-1 Drive I/O Components

A read or write request includes the controller's processing of the command, the seek (the arm/head movement to locate the track, the latency (positioning the head over the correct sector) and then the actual read or write of data (including data transfer back to the controller). Figure 4-2 depicts a disk platter that identifies both physical drive components and basic drive performance terms.

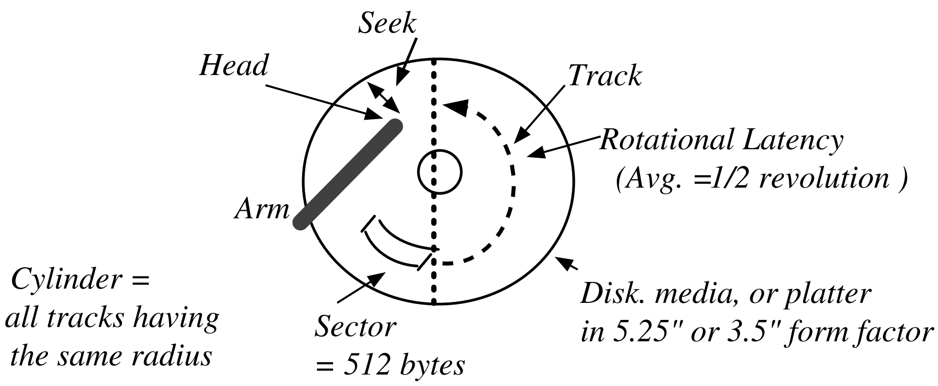

Figure 4-2 Head/Disk Assembly Diagram

Using the above diagrams, the traditional metrics that drive manufacturers use to describe performance are defined as:

Average Seek time: The average time to perform all possible seeks. A seek is the process of positioning the heads over the correct track where target data will be read from or written to. It is usually quoted as the time required to traverse one-third of the data area of the disk. This measure is most meaningful for small, random data transfer situations where seeks are a large portion of the total time and are performed on almost every request. It is not as meaningful for large, sequential file transfers for which seeks are a smaller portion of total time or may not occur at all due to data buffering (described below).

Average Access time: Access time includes both the average seek and average rotational latency and thus the real time required for the heads to be positioned to transfer data. Average rotational latency is quoted as the time in milliseconds required to do 1/2 a disk revolution - at 3600 RPM that means 8.3 ms. This measure is more meaningful than seek time because RPM can differ significantly between drives. When comparing drives from different vendors, be sure that vendors are quoting full access time and not just seek time as these terms are often used interchangeably.

4.1.2 Total Performance Components

The traditional disk drive performance metrics above should not be the only metrics relied on to make configuration decisions. As statistical measures, they do not reflect all operating environments and they only represent a portion of the total I/O system performance. The following diagram demonstrates the components of a complete I/O request:

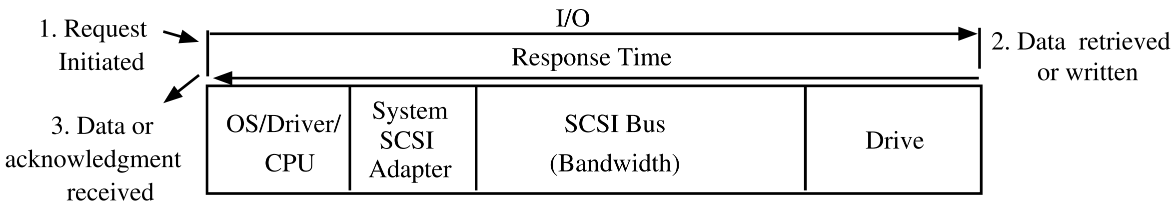

Figure 4-3 Total I/O Components

A request is processed by the operating system and CPU (maybe queued in system driver if bus is busy), sent to the adapter to determine which device to access, transmitted across the bus to the appropriate device, executed by the drive controller as described in the previous section, then passed back through same steps to complete one I/O (request) cycle. (Note: the relative size of each phase of time above does not represent reality nor is it fixed, but is dependent on factors to be described below (generally system overhead is smaller than shown)).

The most comprehensive metrics of I/O performance are, therefore:

Response Time: The time from request initiation by the user to request completion. This is the best measure of single request I/O time and is quoted in milliseconds.

Request Rate: A measure of *throughput* which is the number of requests which are completed in one second (I/O's / second). This measure is very dependent on the system workload (number of users, number of requests, queue depth, and file sizes). This is the most appropriate metric of performance for server/multi-user environments.

On-Track Data Transfer Rate: A measure of *bandwidth* which is the rate at which data is transferred either from the disk media or across the bus (KBytes/second). Physical limiting bandwidth directly off the drive media is calculated using:
[Data blocks/track * B/block (usually 512 bytes) / time for 1 revolution (= 60/RPM)].

Make sure that the on-track transfer rate is distinguished from total data rate (includes spare blocks, headers, servo data, etc.). Actual bandwidth across the bus depends on buffering scheme and workload characteristics. This is the most appropriate metric of performance for large file, sequential access environments where bus saturation is usually the bottleneck to performance. (It isn't meaningful in frequent request (seek), small file environments where the bus is rarely saturated).

4.2 Variables Affecting Total I/O Performance

The three broad categories of variables that impact response time, request rate, or data transfer rate are the *drive,* the *system,* and the *computing/application environment* factors. The following three sections describe the specific variables in each category that influence I/O or bandwidth performance. In subsequent sections, we will provide performance information for two general application environments. Understanding the factors below will allow you to tailor the general performance data to your specific system and application environment.

- **Command Overhead and Mechanical latencies:** Drive command overhead consists of SCSI command decoding and SCSI protocol overhead. This is dependent on the SCSI chip on the drive. This actually is a small portion of the overall transfer time. Mechanical latencies include seek time and rotational latency. During small block transfers, mechanical latencies are the largest portion of total time while in large block transfers, data transfer from the media becomes increasingly important. As the latency and overhead decrease, the overall performance increases.
- **Data transfer from disk or Drive Sector Density (sectors/track):** As more sectors of data are packed into the same track length of a disk, the data transfer rate from the disk media increases. Thus, higher density drives of the same form factor will inherently have higher media transfer rates. (Note: be sure to compare *formatted* capacities (usable storage) of drives since vendors may quote the higher *unformatted* capacities without specifying).
- **Form Factor:** 5.25" disks have longer tracks than 3.5" disks (generally more sectors/track) and more platters and heads. This means 5.25" disks do fewer head switches while reading large files and they have access to more data under the heads (see KB/cylinders below). Therefore, larger form factor disks tend to have inherently higher bandwidth even when access time is the same (Note RZ24 and RZ55 in Table 4-1).
- **Dedicated versus embedded servo:** In dedicated servo drives, a platter and head are dedicated to keeping the head on track while in embedded servo drives, the positioning

data is interspersed with user data. The constant tracking of dedicated servo enables more rapid head positioning during read/write and after head switch (at end of track) and thus improved bandwidth. All Digital's 5.25" SCSI disks and the RZ25 currently have dedicated servo while the lower capacity 3.25" disks have embedded servo.

- **Buffer size and type**: Buffers dynamically store data being read from the disk to enable the system bus to access the data faster (in large bursts) rather than at the rate that it streams off the disk (bus rate is generally 2-3 times faster than media rate). Buffer algorithms can be simple (holding data only until it is transferred to the system) or more advanced such as "read ahead" (reads and stores data beyond the requested location in the event that the next request will be sequential and "hit" or match the data in the buffer). In the case of the RZ57, the 256 KB buffer is segmented such that four different read ahead requests can be stored, thereby improving chances of a buffer hit in a multi-user environment.

Refer to the chart of technical specifications below to see how the variables described above differ by RZ disk drive. Note that the RZ24 3.5" drive has slightly higher data density and media transfer rate than the lowest capacity 5.25" drive (RZ55). However, the RZ55 will still outperform an RZ24 in most situations primarily because of its dedicated servo but also due to the higher number of heads and platters (see KB/cylinder calculation).

Table 4-1 RZ Series Metrics

	RZ22	RZ23L	RZ24	RZ25	RZ55	RZ56	RZ57
Capacity (Formatted)	52	121	209	426	332	665	1.0
Form Factor	3.5"	3.5"	3.5"	3.5"	5.25"	5.25"	5.25"
Ave seek time (ms)	25	19	16	14	16	16	14.5
Ave rotational latency (ms)	8.4	8.8	8.6	6.8	8.3	8.3	8.3
Rotational speed (RPM)	3600	3400	3500	4412	3600	3600	3600
Ave access time (ms)	33.4	26.8	24.6	20.8	24.3	24.3	22.8
Num. of R/W (data) heads	4	4	8	9	15	15	15
Num. of cylinders	776	1523	1248	1474	1218	1626	1928
Blocks per track*	33	39	38	48-73	36	54	71
KB/cylinder	67	79	155	221-336	276	468	539
Areal Density (MB/Sq.")	26.96	67.06	54.06	68.28	22.82	45.86	70.88
Max. Transfer rates: (MB/sec. burst)							
to/from Media*	1.25	1.5	1.5	2.0-3.13	1.25	1.875	2.2
data to/from Media	1.0	1.13	1.13	2.33	1.11	1.66	2.19
across Bus (async)**	1.6	3.0	3.0	3.0	1.5	1.6	1.6
across Bus (sync)**	n/a	4.0	4.0	4.0	4.0	4.0	4.0
Servo(Embedded/Dedicated)	Emb	Emb	Emb	Ded	Ded	Ded	Ded
Buffer Size (KB)***	16	64	64	60	64	64	256
Power - max. seek (Watts)	8.3	3.8	6.6	14	25.2	32	32

*RZ25 uses Variable Density Recording, with bands of denser data radiating out from the disk center.
Burst rate on the bus (sustained varies with system and application) * Read ahead buffer

4.2.1 System Variables

When the application environment and drive configurations are held constant, SCSI drives will perform differently on different systems offered by Digital. The sources of differentiation are:

- Central Processor architecture/speed (CPU)
- SCSI system level controller/adapter architecture and speed
- Operating system features
- I/O bus architecture (potential number of SCSI buses supported)

To summarize differences by system:

	CPU Speed*	SCSI Adapter	Operating Sys.	I/O Buses
VAXstation 3100 Models 30/40	4 MIP	asynchronous	VMS/Ultrix	one/two
VAXstation 3100 Models 38/48	5 MIP	asynchronous	VMS/Ultrix	two
VAXstation 3100 Model 76	9 MIP	asynchronous	VMS/Ultrix	two
VAXstation 3520/3540	6/10 MIP	synchronous	VMS/Ultrix	one
MicroVAX 3100	5.5 MIP	asynchronous	VMS/Ultrix	two
DECstation 2100	11 MIP	synchronous	Ultrix	two
DECstation 3100	15 MIP	synchronous	Ultrix	two
DECstation/system 5000 Mod 100	18 MIP	sync (2nd gen.)	Ultrix	four
DECstation 5000 Model 200	24 MIP	sync (2nd gen.)	Ultrix	four
DECsystem 5100	19.4 MIP	sync (2nd gen.)	Ultrix	one
DECsystem 5500	28.2 MIP	sync (2nd gen.)	Ultrix	four

* MIP = Million instructions per second

4.2.1.1 Central Processor Factors

Faster central processors affect the performance of the I/O subsystem just as they do any other component, however this effect is usually minor. There are two cases where CPU performance does make a difference. First, when data is read directly from the file system cache, performance generally scales with CPU speed. Second, for processors with first generation SCSI adapters (all but DS5000), CPU cycle availability may limit transfer rate to or from the disk.

4.2.1.2 SCSI Controller/Adapter Factors

- **SCSI overhead** is associated mainly with the SCSI chip on the system board or adapter card (DS5000) and consists of command setup, SCSI protocol overhead, and data processing. Second generation SCSI chips (available on some systems) have implemented DMA (direct memory access) that reduces the amount of required host intervention and thereby significantly improves performance. (Note: SCSI *controller* applies to 'busless' systems whereas SCSI *adapter* applies to systems (i.e. DS5000) that have an internal system bus that is adapted to SCSI).

- **Type of bus transfer** supported also differs by system and SCSI adapter. Some first generation chips (i.e. on VAXstation 3100) support only *asynchronous* bus transfer (1.5 MB/sec) while others (i.e. on DECstation 3100) also support *synchronous* transfer (4-5 MB/sec). Synchronous transfers enable higher bandwidth than asynchronous due to the lower protocol overhead. Synchronous is therefore optimal for large data transfers. All RZ drives except the RZ22 support either synchronous or asynchronous transfer depending on the system SCSI chip and will therefore perform faster on the DECstation 3100 than the VAXstation 3100 under the same conditions. Faster data transmission becomes increasingly important as more devices are added to a bus and multiple requests are initiated.

4.2.1.3 Operating System Factors

- **File System (host-based) buffer cache:** the most important difference between VMS and Ultrix is that Ultrix has an intelligent system buffer that stores the most frequently accessed files in memory. This reduces the frequency of accessing the disks and improves total system response time. The hit rate to this cache is highest for single users where the entire cache is dedicated to one user. For multiple users, the hit rate to this cache would be lower, yet some software development applications still demonstrate over 90% hit rate. The implication is that for applications that make use of this system buffer cache, total Ultrix system performance is less a function of drive capability than of cache capability. VMS I/O performance is more dependent on raw drive access time and bandwidth.

Cache capability also depends on amount of memory dedicated to it. Since memory is a higher cost/MB storage medium, the user needs to determine the optimal memory to disk ratio (price/performance willing to pay) for the application environment. In Ultrix, the user can adjust the percent of memory dedicated to buffer cache (default is 10%) and the basic file system request (segment) size (default is 8 KB).

- **Device driver:** VMS and Ultrix SCSI device driver buffers are also implemented differently. As requests are placed, the OS must segment them into the host buffer, which manages all of the peripheral devices at once.

  - VMS (on VAX) allocates part of the 128 KB memory dynamically, as a request is issued, to a maximum of 16 KBytes (32 sectors of 512 bytes each).

  - Ultrix (on VAX) allocates a static portion of the 128 KB memory to tapes and CD ROM on spin-up, with the remainder divided evenly between the disk drives on the bus. The buffer size per disk can vary from a minimum of 16 KB (32 sectors) to a maximum of 64 KB (128 sectors).

  - Ultrix (on RISC systems) allocates a 16 K "double pipe" buffer per SCSI device that keeps the I/O reads and writes flowing without segmentation.

The less a request is segmented, the less overhead time is spent processing the request in several segments. Thus for multi-user systems with small requests, VMS would be more beneficial while ULTRIX for RISC systems is better for large requests. Some very technical and large file applications bypass the system buffer cache (above) to take advantage of the "raw" data access of the driver for maximum bandwidth. For these applications, Ultrix's driver buffer scheme can optimize the I/O performance.

4.2.1.4 I/O Bus Factor

- **Number of Buses:** Multi-user (or multi-process) environments will see performance improvements as drives are separated among buses. Especially with large data transfers keeping the I/O bus busy, splitting devices among several buses can significantly improve I/O subsystem response. While only one transmission can occur across a single bus at any time, multiple buses enable multiple parallel transmissions simultaneously from different disks that are on different buses. Such multi-channel configurations can increase total system response time and bandwidth, as long as the system processor can efficiently handle the load.

4.2.2 User Environment Variables: (Configuration and Application)

You may have noticed that many of the drive and system variables influencing performance are dependent on characteristics of the user environment, application or system configuration (single-user or multi-user). The user environment determines the workload characteristics that are seen by drives. The primary mass storage dependent workload variables are:

- **Queue depth:** The number of requests on average being stored in host buffers awaiting processing by the SCSI controller. Generally this is a function of the number of users and the rate that they read or write to disk. Total system throughput generally decreases with higher queue depths. In the data displayed later, a queue depth of 0 was used. If you believe your system receives, on average, more requests than the number of disks on the system, you should expect to see a lower request rate (see multiple disks below).
- **Transfer size** (small vs. large requests): The average request size and the variance of sizes for a specific user environment is important to understand because a direct/indirect relationship exists between transfer size and disk performance (depending on which metric is being used). In other words, larger transfer sizes demonstrate higher bandwidth but lower throughput than smaller transfer sizes. Large requests take full advantage of high bandwidth drives while small requests benefit from fast access from the disk.
- **Sequential vs. Random access:** Applications/user environments tend to access drives either predominately sequentially or randomly. For example, for single user applications that use very large files, sequential access mode is more common - large files are written or read in a single command and any subsequent requests for data will tend to follow where the previous request left off. In this case, drives with high sector density and large buffers will have the best performance. In multi-user applications where the I/O subsystem sees many small, random requests, multiple small drives outperform single large drives.
- **Reads vs. writes:** In sequential environments, read access is faster than writing to the disk due to read ahead buffers as well as the lower seek overhead associated with reads. With random transfers, writes are faster than reads since the seek time can be slightly overlapped with the initial data transfer. All data provided later is based on reads only and should be tempered to reflect the predominance of reads or writes to your system.
- **Multiple small disks vs. single large disks:** Due to the near linear improvements in drive performance at smaller transfer sizes, multiple small drives can often out-perform

a single large disk in a multi-user environment. As long as the load on the system is distributed evenly across spindles, overlapped seeks (with the disconnect/reconnect drive feature) can improve the overall request rate. However, single large disks perform better in a single user environment, characterized by large requests, due to high disk data rate.

- **Distribution of files:** Applications/files must be evenly distributed across multiple disks to see the maximum performance advantage in multi-user environments. This includes balancing the location of swap files as well as hot (frequently used) files across two (or more) disks.
- **Disk fragmentation:** With frequent deleting and writing to a disk, data becomes increasingly fragmented and less sequentially located on the disk. This decreases the data rate from the disk due to increased seeks and rotational latency. Performing disk defragmentation will maintain maximum I/O performance (back-up to tape and re-load onto disk).

5. Drive Performance by Application and Computing Environment

How does one take into account all the interdependent environmental, system, and drive variables described above to optimize one's storage configuration? We recommend the following approach:

1. Understand your particular *user/application environment* (request size, access type, queue depth)
2. Determine current or future *system attributes* (CPU, operating system, SCSI adapter, buses)
3. Understand how the above affect or are dependent on the different *drive technologies* and *combinations* of drives for performance (form factor, capacity, access time, buffer size)

Every system is a unique environment based on the workload and configuration variables we've discussed in this paper. With an ever-growing number of Ultrix and VMS applications, we cannot characterize every specific environment to make recommendations on the optimal storage configuration for each individual user/situation. However, we can try to generalize all environments to two representative situations (single-user and multi-user environments) which could apply the largest set of users. We will then provide the maximum relative I/O performance metrics under those two situations. Users should attempt to characterize their own environment as discussed in Section 4.2 (performance variables) and then customize this generalized data to their unique environment.

All I/O subsystem performance data presented here was measured using Digital's own drive exerciser test programs (IOX for VMS testing and UIOX for ULTRIX testing). The programs vary the request size, the queue depth, and the seek length to obtain measures of drive throughput and bandwidth for sequential and random reads or writes. These measurements allow us to make comparisons between RZ disk drives and amongst groups of drives to optimize performance and to assist in selection of drives for your system.

In order to isolate the performance of the drives from the many other system variables, these programs bypass the Ultrix file system and host-based cache. They saturate the bus to demonstrate raw drive and adapter performance rather than attempt to portray maximum system performance. This data is useful for comparison within the disks used, but should not be used as a competitive benchmark tool against other drives on other systems.

5.1 Single-user Environments

We will define a simple 'single-user' environment as a stand-alone system or networked satellite workstation that has its operating system and data files on a local disk(s). Data access of the disk is typically sequential. We differentiate between two groups of single-user applications based on the typical file sizes and whether or not the application utilizes the I/O file system cache. For single-user workstation applications that need to rapidly read or write large sequential files (such as image processing and seismic analysis), the file system is usually bypassed to maximize the data rate of the drives. Most references to 'single-user' environments in this paper (and the performance data presented) refer to this category of high bandwidth, single-user applications.

On the other hand, some workstation applications may be limited by the local disk performance and benefit from the use of the file system cache. When the cache is effective, the user's perceived performance will be dominated by the performance of the cache and not the I/O subsystem behind it. These two classifications are broadly defined and leave it to the user to understand how their specific application utilizes the system cache before applying the associated performance metrics and data.

Applications: 1) Low bandwidth commercial and technical applications (Examples: desktop publishing, software development and some CAD packages, etc.)
2) High bandwidth scientific and technical and commercial applications (Examples: image processing, finite element analysis, seismic analysis, solids modeling, fluid dynamics computation, animation, meteorology, some CAD packages, multi-media, etc.)

Information type: 1) Text, charts, 2D graphics, etc.
2) Images, sequences of images, 3D graphics, sound, etc.

Drive access: Sequential access from the disk (or buffers)
Typical request size: 1) < 16 sectors (8 KB) (applications w/ file system)
2) > 16 sectors (8 KB) for applications bypassing the file system cache (i.e. imaging)
Queue depth: 0,1 (unless window, multi-process environment)
Request mix: mostly reads

Primary Performance Metric: 1) Response time, Request Rate, Bandwidth
2) Bandwidth (Data rate or KB/sec) See section 5.1.1.

Given these assumptions, optimal attributes for a storage subsystem would include:

Disk features: Large, intelligent drive buffers (to take advantage of sequential access)
Synchronous bus transfer support for high bus data rate

High data rate off the disk to maximize bandwidth (primarily category 2)
Fast drive access time (more important for applications in category 1)

System features: 2) Fastest SCSI adapter, maximum amount of system memory

Operating System: 2) File system buffer cache (customizable for Ultrix and reduces dependence on drive performance)

5.1.1 Bandwidth metric for the Single-user

The following performance charts use bandwidth as the primary metric comparison for the large file single-user environment that features sequential reads and writes and transfer sizes larger than 16 sectors. Although response time is also an appropriate measure, the limiting factor in very large file applications that bypass the file system cache (i.e. imaging and seismic analysis) will tend to be bandwidth which should therefore be optimized when choosing mass storage disks.

Figure 5-1 shows the maximum transfer rate for each of the RZ drives. Refer to section 4.2.1 for the MIPs ratings of the various Digital systems.

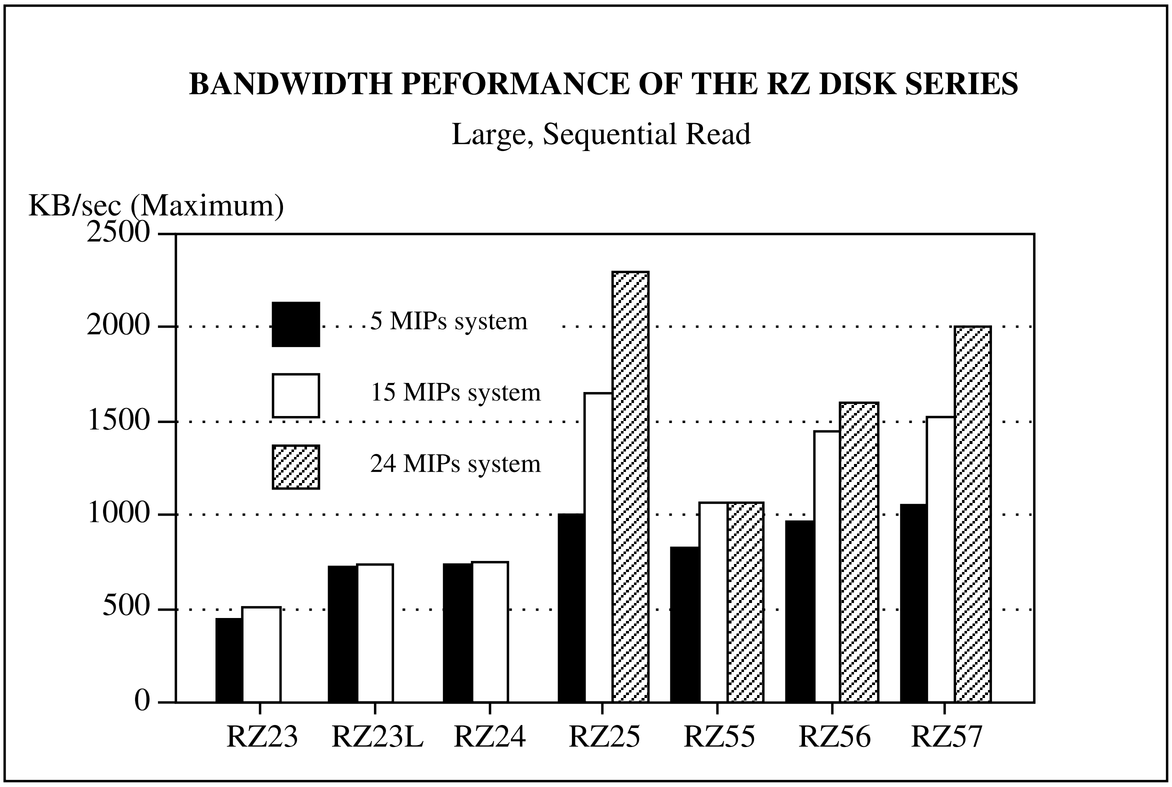

Figure 5-1 Maximum Bandwidth of RZ Disk Series

There are many variables influencing the data in Figure 5-1: First note that larger and newer drives appear to perform better than smaller drives (on the same system) which is due primarily to the higher sector density, dedicated servo (lower head switch time), and the buffer management algorithm.

The second point to note is that drives will perform better on systems with more processor power. That is, the limiting factor for high I/O rates is actually more than just the drives.

The largest improvement in Bandwidth would be to add drives to different buses. As we now add drives and buses to the DECstation/system 5000 and compare random and sequential reads, we can see dramatic improvements in raw data transfer rates (See Figure 5-2). This maximum bandwidth would only be experienced on the DECstation 5000 if all four available channels each had one RZ57 that was reading four separate files and transmitting simultaneously to the CPU. (Note that tests bypass the file system. For applications that utilize the file system, lower total system bandwidth would be experienced).

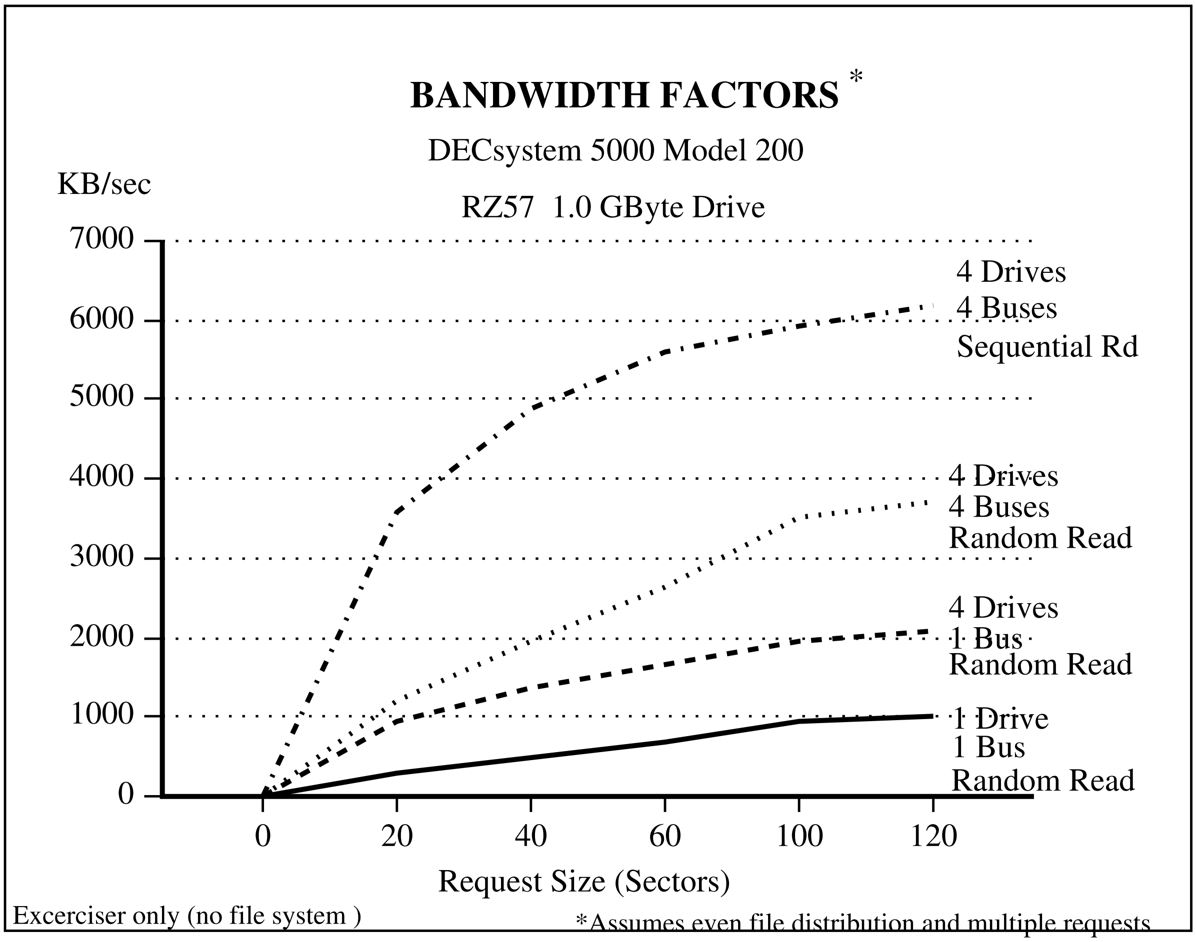

Figure 5-2 Multiple Device and Bus Bandwidth Improvements

5.2 Multi-user Environments

In this environment, the system could be a device server, dedicated application server, or general purpose system for multiple interactive users and/or windows. The environment may be characterized by multiple, short, random disk accesses, and the following general scenarios:

Applications: Commercial applications (Examples: office automation, PC integration, multi-user, time-share, transaction processing)
Departmental application server (Examples: computation, database, distributed bank branch server, retail point of sale automation)

Information type: Text, documents, charts, low end graphics

Drive access: Transfer sizes: range 1 KB - 1 MB (See Figure 5-3 for VMS sizes)
Request mix: ~80% reads for time-share
~70% writes for transaction processing
Random data accesses
Average queue depth: greater than 1 (dependent on user request rate)

Primary Performance Metric: Request rate (I/O requests/sec)**.** See next section.

A majority of the storage interactions are in multiple, short, random disk accesses. On most VMS systems, 95% of all transfers are 5 sectors or less, almost 50% of them are only 1 sector as shown on the following chart.

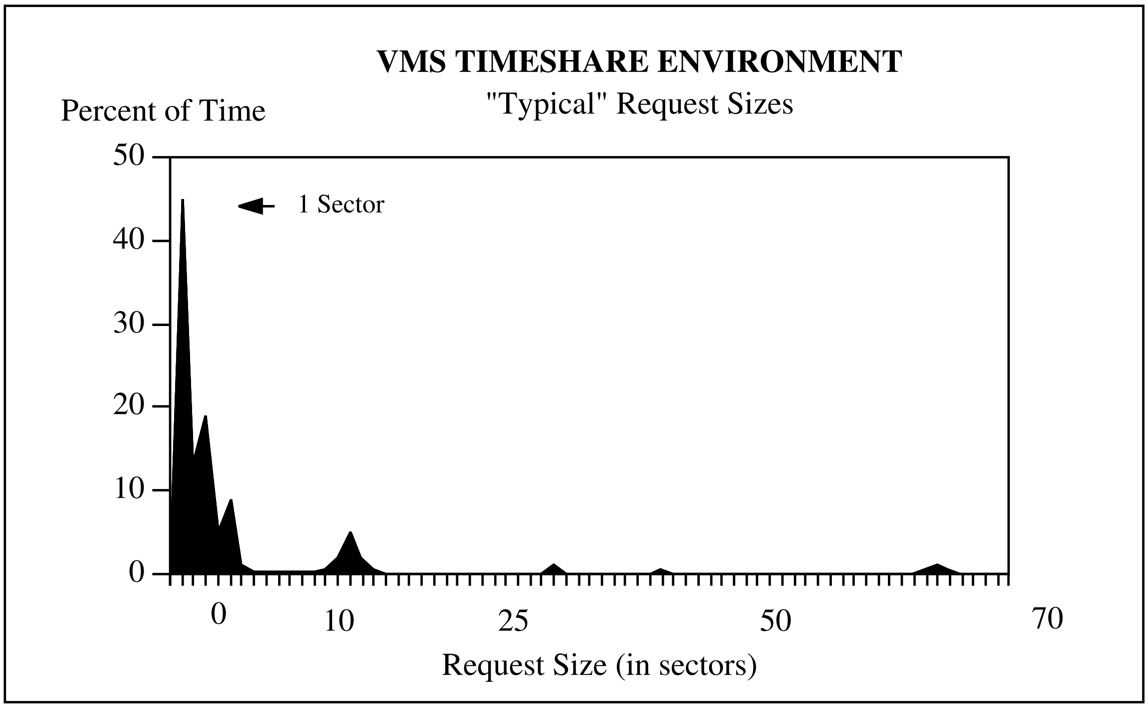

Figure 5-3 VMS Time-share Request Sizes

Given these assumptions, optimal attributes for a storage subsystem would include:

Disk features: Since buffer hit rate is reduced, access time becomes a larger proportion of overall response time.
For request rate optimization, multiple drives enable overlapped seeking (disconnect/reconnect at the drive level must be supported and data should be evenly distributed over the drives)

System features: Fast processor, optimal queue and file management (host buffer segmentation is more important than drive buffer capability)
Increased number of buses

5.2.1 Throughput for Multi-user Systems

As with the single-user environment, we will first compare the typical single drive multi-user performance for VMS and Ultrix systems (1 sector, random read environment) in the key multi-user metric: throughput. In Figure 5-4, we notice that there is less of a difference between the 3.5" disks and the 5.25" disks and that single disk throughput seems to be relatively independent of SCSI adapter, CPU speed, and operating system. Refer to section 4.2.1 for the MIPs ratings of the various Digital systems.

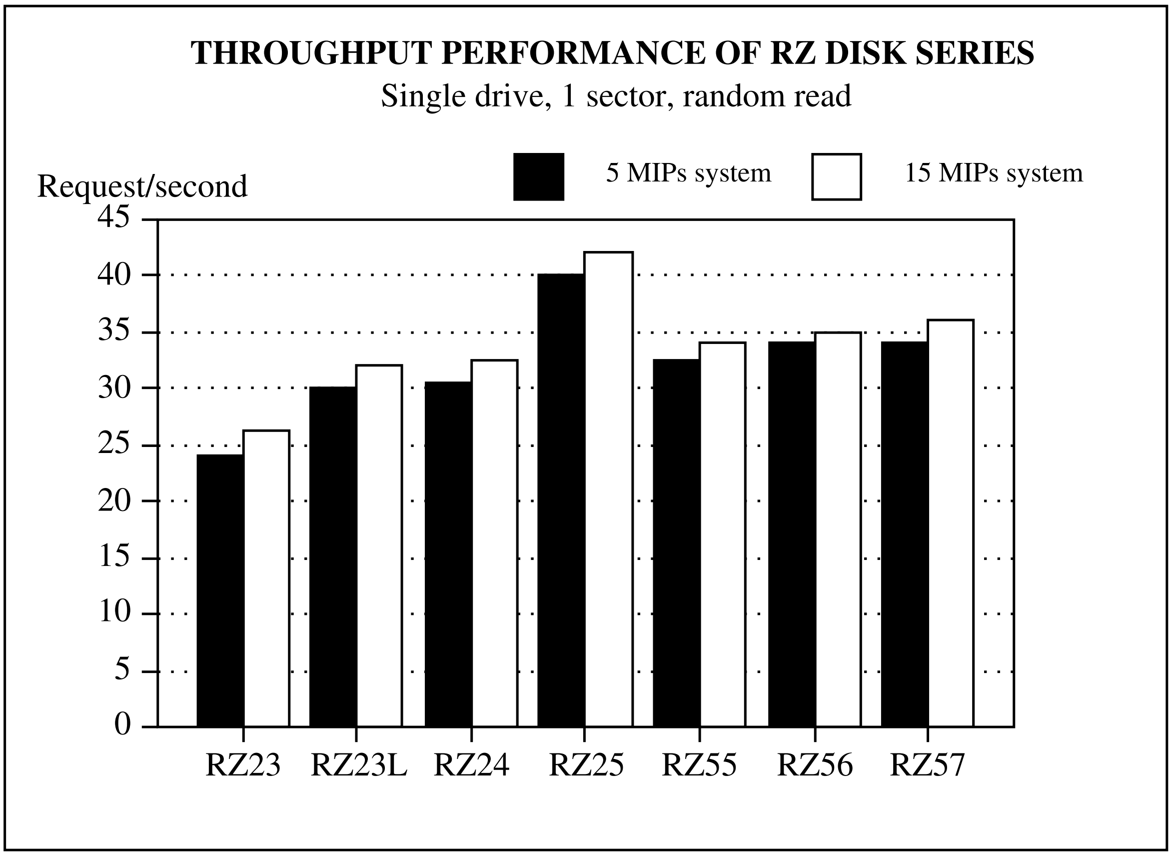

Figure 5-4 Throughput (I/O) Performance of RZ Disk Series

Now, by configuring multiple drives in a multi-user system, dramatic I/O improvements can be demonstrated. Multiple drives improve the I/O rate primarily because of the overlapped seeks as long as drive disconnect/reconnect is supported. The following chart displays the improvements in multiple drives with a VAXstation 3100 VMS system with 1 sector random reads.

Notice that two RZ23L drives together outperform any other single drive. Keep in mind that this incremental performance applies not only to system and data disks but also to multiple swapping disks in a networked system where a workstation user may be running multiple processes on the network (see operating system documentation for implementation).

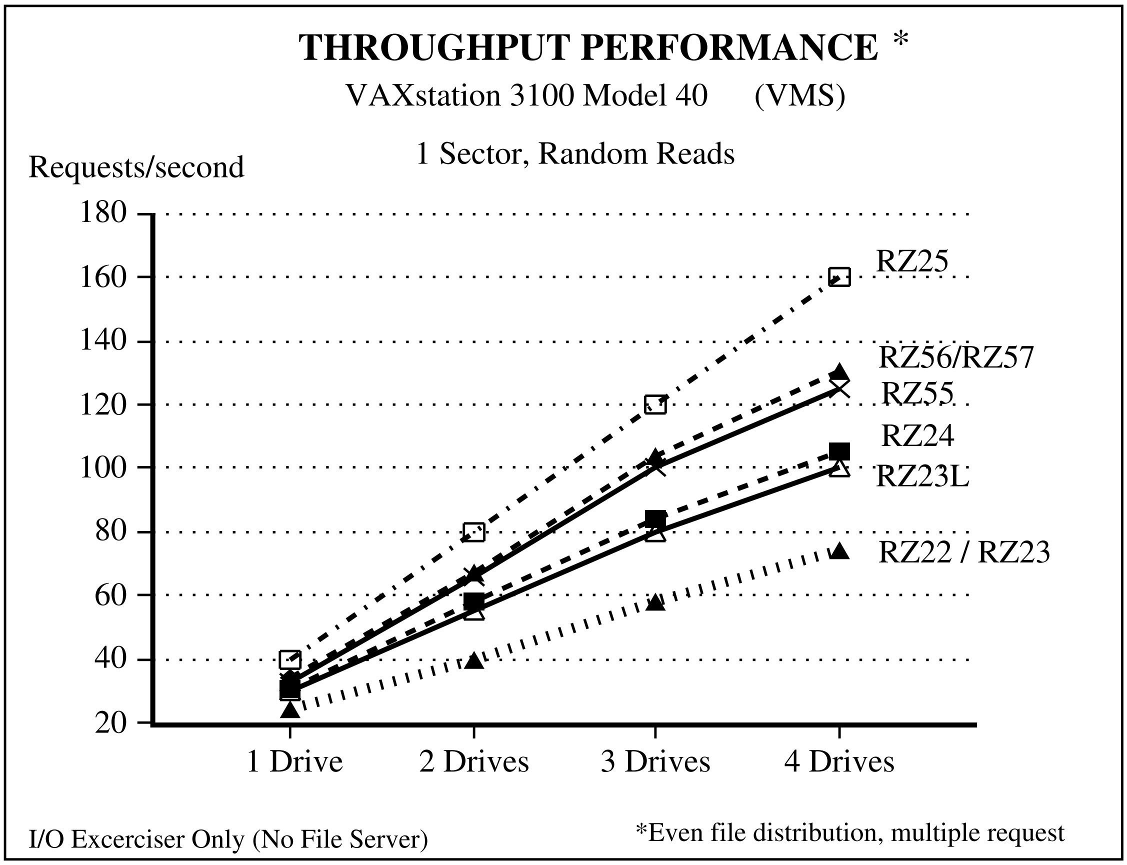

Figure 5-5 Multi-drive I/O Performance

Similarly, if we take the same DECstation/system 5000 scenario as we used in Figure 5-2 four RZ57 drives (either on 1 bus or across 4 buses), we see nearly four times the throughput compared to one drive on one bus during random reads. The four bus configuration performs better than single bus configurations in large request sizes since the single bus becomes saturated more frequently. For sequential read environments, the four drive/four bus configuration demonstrates even greater performance improvements: for large requests environments, sequential reads can show nearly two times higher request rates than random read requests whereas for single sector request sizes, this configuration can demonstrate up to nine times the request rate (notice graph is in log scale). Improvements come from the increased use of drive buffer as a cache as the request size increases.

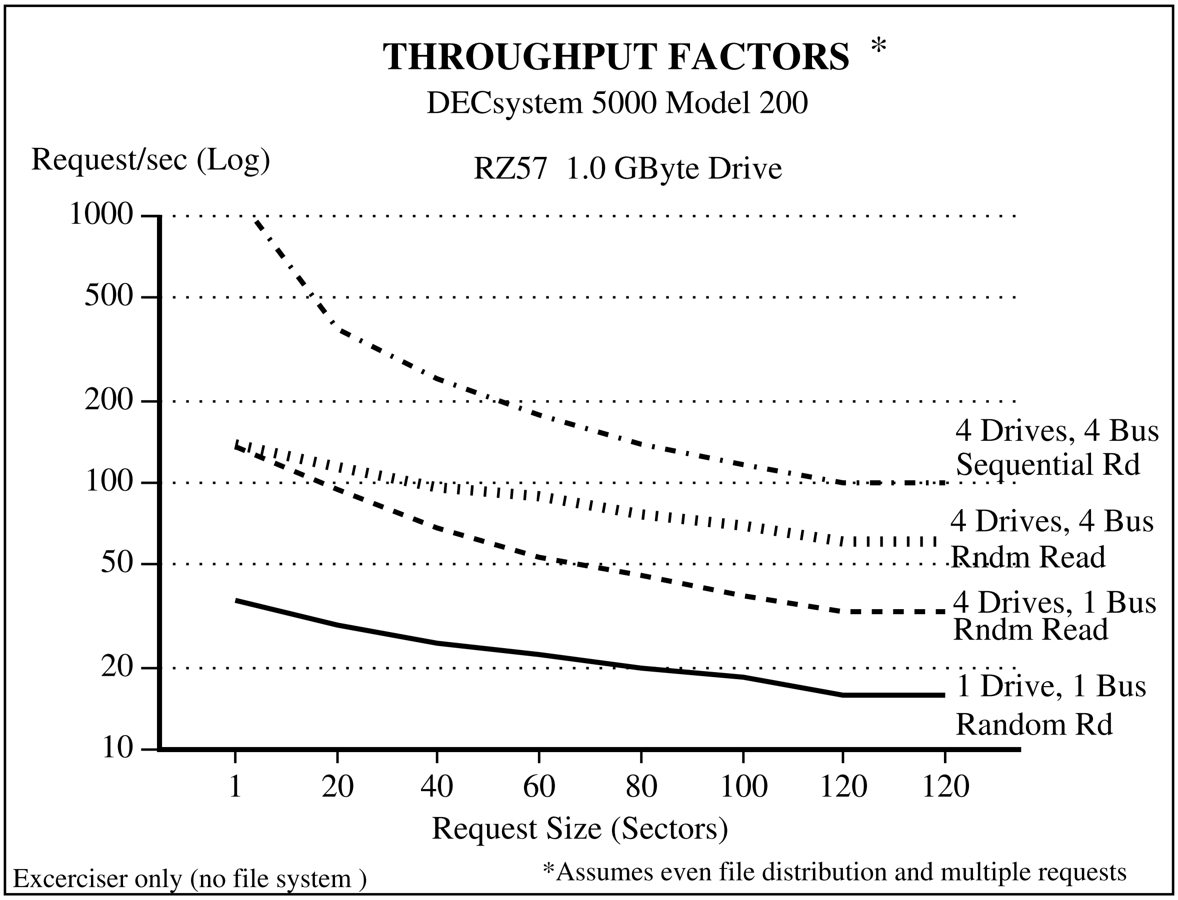

Figure 5-6 Throughput Factors, DECsystem 5000 and RZ57

5.2.2 Configuring Systems for Multi-user Throughput

Since multi-user environments can benefit from multiple drives, configuration decisions are more complex. The next two figures provide an overview of possible drive combinations for the MicroVAX /VAXstation 3100 and the DECstation 3100. One can draw two general conclusions from the graphs:

- multiple drives deliver more throughput than single drives of the same aggregate capacity
- the higher the total system capacity, the higher overall I/O performance

Both of these observations imply that optimal performance solutions come at greater expense given that multiple drives tend to have higher cost/MB than single drives of the same aggregate capacity. However, if we normalize the cost per unit of performance derived, there is in fact a more favorable price/performance for multiple drives:

300 MB Configurations		**I/O/sec**	**Example $**	**$/MB**	**$/I/O/sec**
(3) RZ23L	**(363 MB)**	80	$3600	10	45
(1) RZ24&(1) RZ23L	**(330 MB)**	58	$2700	8.2	46
(1) RZ55	**(332 MB)**	33	$3500	10.5	106
(1) RZ25	**(426 MB)**	40	$3500	8.2	87.5

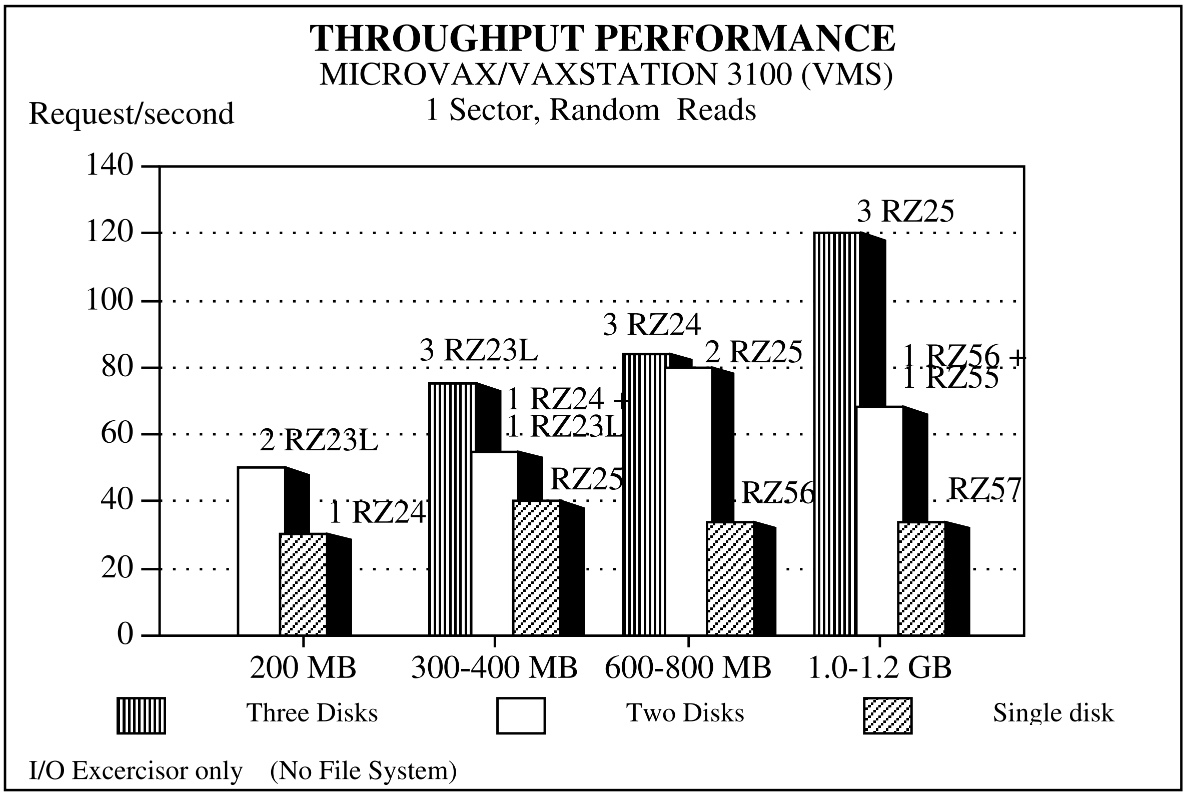

Figure 5-7 Multi-User Disk Configurations, MicroVAX and VAXstation 3100

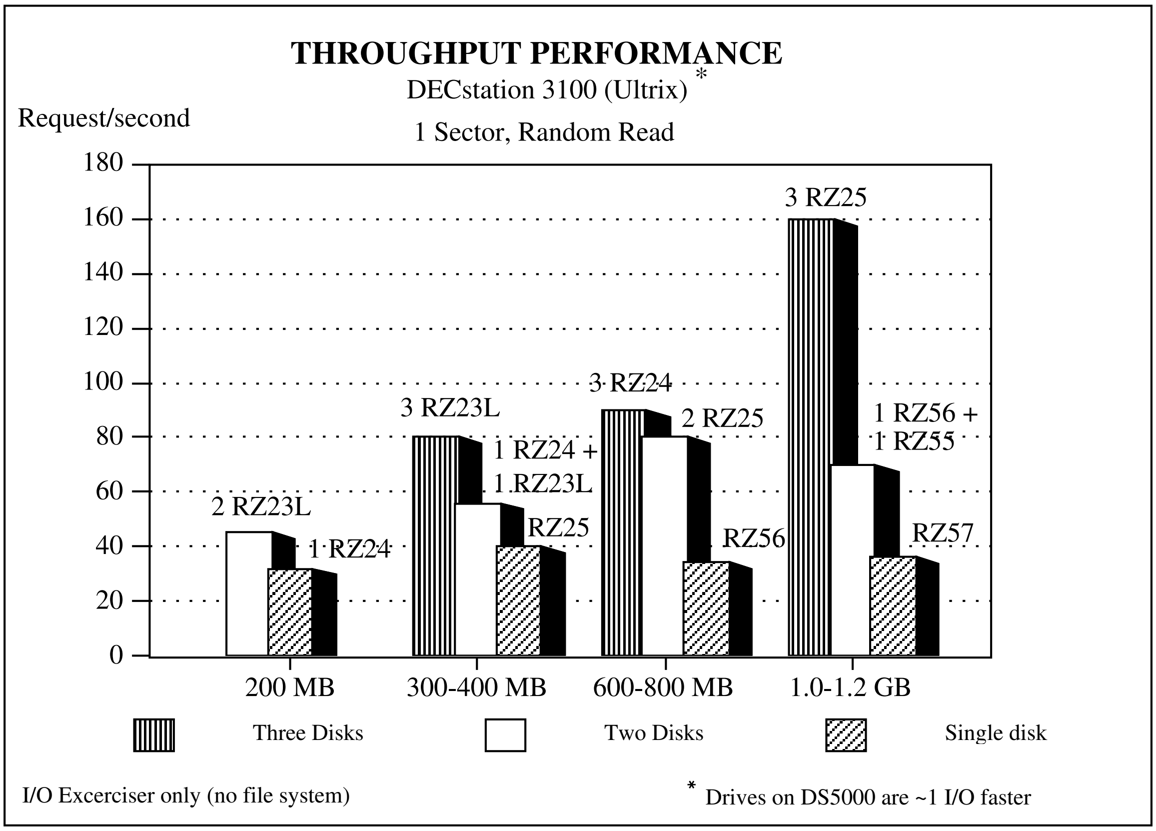

Figure 5-8 Multi-User Disk Configurations, DECstation 2100/3100

6. Summary and Other Considerations

This guide has attempted to explain the I/O interface and performance issues which you need to understand to select a more balanced I/O storage solution for your existing or future open desktop system.

In the open desktop environment, users demand exceptional performance. Our emphasis in this paper has been to help you optimize your I/O performance based on understanding how your application environment, system architecture, and operating systems interact with I/O subsystems.

Simply summarized, single-user systems tend to perform best with higher capacity 5.25" form factor drives. Multi-user systems will demonstrate nearly linear increases in performance with additional drives on the bus thereby encouraging multiple spindle configurations.

There are, of course, other factors than performance that we have not discussed which may also affect your I/O subsystem decision such as:

- Total storage capacity/system constraints (see Appendix A)
- Internal storage vs. expansion box requirements
- Removability
- Back up strategy (match disk capacity with unattended backup capability)
- Cost/performance

Although beyond the scope of this paper, you should be aware of other aspects of storage management that will ensure continued high performance of your storage systems:

- Ensure that devices are configured according to SCSI specifications and Digital's implementation including proper ID assignment and priority (0 = lowest priority, 7= highest priority on the bus) and supported/recommended configurations (correct cable quality and length, proper distance between nodes on the bus, etc.). Contact customer services.
- If you are adding non-Digital devices to the bus, contact VES (Vendor Equipment Services) and perform thorough testing, checking error rate in maximum load to ensure interoperability and data integrity
- Develop understanding of workload characteristics and monitor for changes over time
- Customize file system allocation defaults as needed (see Ultrix user documentation)
- Select applications that manage I/O intelligently (use of file systems and buffers) and /or program with I/O performance in mind
- Balance the file load and hot files across multiple disks
- Perform disk defragmentation when needed (backup to tape and reload disk)
- Monitor drive error rates for excessive incidences (call customer service as needed)

A. SCSI Storage Devices Ordering Information

The following is a guide to ordering storage products (excluding memory) for SCSI-based systems. It is sorted by system and includes specific model numbers and maximum configuration information by system. The following storage products are included:

BASE DISK AND TAPE DRIVE OPTIONS

RX23	1.4 MB 3.5" diskette drive
RX33	1.2 MB 5.25" floppy drive
RZ22	52 MB paging disk drive for VMS
RZ23L	121 MB sys or paging disk for VMS, paging for Ultrix
RZ24	209 MB system disk or paging disk
RZ25	426 MB system disk or paging disk
RZ55	332 MB system disk drive
RZ56	665 MB system disk drive
RZ57	1.0 GB system disk drive
TZ30	95 MB cartridge streaming tape drive (internal)
TK50Z	95 MB streaming tape in expansion box
TLZ04	1.2 GB Digital Audio Tape drive in expansion box
TZK10	525 MB QIC streaming tape drive
TKZ08	2.2 GB Helical scan exabyte drive
TSZ05	40 MB streaming tape drive
TSZ07	140 MB 9 track tape drive
RRD40	600 MB compact disk drive
RRD42	600 MB compact disk drive

SZ12 EXPANSION BOX OPTIONS Note: device upgrade kits also available

	Disk Drives				Diskette Drives		Tape Drives		CD Drive
	RZ25	RZ55	RZ56	RZ57	RX23	RX33	TZ30	TZK10	RRD42
	3.5"	5.25"	5.25"	5.25"	3.5"	5.25"	Cart.	QIC	CD
Part #	426	332	665	1.0	1.44	1.2	95	525	600
	Mbyte	Mbyte	Mbyte	Gbyte	Mbyte	Mbyte	Mbyte	Mbyte	Mbyte
SZ12A-XA/XB		1							
SZ12A-AA/AB		2							
SZ12A-LA/LB		1			1				
SZ12A-MA/MB		1				1			
SZ12A-HA/HB		1					1		
SZ12A-EA/EB		1						1	
SZ12A-FA/FB		1							1
SZ12B-XA/XB			1						
SZ12B-BA/BB			2						
SZ12B-HA/HB			1				1		
SZ12B-EA/EB			1					1	
SZ12B-FA/FB			1						1
SZ12C-XA/XB				1					
SZ12C-CA/CB				2					
SZ12C-EA/EB				1				1	
SZ12C-FA/FB				1					1
SZ12G-XA/XB	1								
SZ12G-GA/GB	2								
SZ12G-HA/HB	1						1		
SZ12G-EA/EB	1							1	
SZ12G-FA/FB	1								1
SZ12X-LA/LB					1				
SZ12X-MA/MB						1			
SZ12X-HA/HB							1		
SZ12X-EA/EB								1	
SZ12X-FA/FB									1

Note: For any technical or configuration questions on these devices, contact the Technical Consulting Center (TCC) at 1-800-DIGITAL (1-800-344-4825).

SYSTEM OPTIONS

* Please refer to the Systems and Options Catalog for full system information.

VAXstation 3100

Configuration rules:

- Max. 7 storage devices
- Max. 3 internal devices (Models 30, 38, 76); Max. 5 internal devices (Models 40,48)
- Max. 3 external expansion boxes (6 external devices)
- Max. 5 external devices (Models 40,48); Max. 6 external devices (Models 38,38,76) when system box holds TZ30, RRD40,or 3 RZ23L.

	Internal to System Box					External to System Box			
	RZ22	RZ23L	RZ24	RX23	TZ30	TK50Z	TLZ04	RRD40	SZ12
Model 30									
Field installed	-EF	-EF	-EF	-EF	-EF	-GA/G3	-MA	-FA/F3	All
Factory installed	-EJ	-EJ	-EJ	-EJ	-EJ	-GA/G3	-MA	-FA/F3	All
Model 38									
Field installed	-EH	-EH	-EH	-EH	-EH	-GA/G3	-MB	-FA/F3	All
Factory installed	-EJ	-EJ	-EJ	-EJ	-EJ	-GA/G3	-MB	-FA/F3	All
Model 76									
Field installed	--	-EH	-EH	--	--	-GA/G3	-FA	-FA/F3	All
Factory installed	--	-EJ	-EJ	-ER	--	-GA/G3	-FA	-FA/F3	All

	Internal to System Box					External to System Box			
	RZ22	RZ23L	RZ24	RRD40	TZ30	TK50Z	TLZ04	RRD40	SZ12
Model 40									
Field installed	-EF	-EF	-EF	-EF	-EF	-GA/G3	-MA	-FA/F3	All
Factory installed	-EG	-EG	-EG	-EG	-EG	-GA/G3	-MA	-FA/F3	All
Model 48									
Field installed	-EF	-EF	-EF	-EF	-EF	-GA/G3	-MB	-FA/F3	All
Factory installed	-EG	-EG	-EG	-EG	-EG	-GA/G3	-MB	-FA/F3	All

VAXstation 3520/3540

Configuration rules:

- Max. 4 drives

	Internal to System Box		External to System Box
	RZ55	RZ56	RRD40
Model 3520			
Field installed	A-SF	-SF	-GA/G3
Factory installed	A-SA	-SA	-GA/G3
Model 3540			
Field installed	A-SF	-SF	-GA/G3
Factory installed	A-SA	-SA	-GA/G3

DECstation 2100/3100

Configuration rules:

- Max. 7 devices
- Max. 2 internal hard disks; 1 internal diskette drive
- Max. 6 external devices (4 external disks)

	Internal to System Box			External to System Box				
	RZ23L	RZ24	RX23	TLZ04	TKZ08	TSZ07	RRD42	SZ12
Field installed	-FF	-FF	-EZ	-FA	-AA	-CA	-FA/DG	All
Factory installed	-FG	-FG	-EJ	-FA	-AA	-CA	-FA/DG	All

DECstation 5000 Model 100 Series

Configuration rules:

- Max. 2 internal devices
- Max. 3 expansion boxes (6 devices) per SCSI bus
- Up to 3 other Turbochannel options to expand number of SCSI buses & devices (6 per bus)

	Internal to System Box						External to System Box			
	RZ23L	RZ24	RZ25	RX23	RRD42	TZK10	TLZ04	TKZ08	RRD42	SZ12
Field installed	-FL	-FL	-FL	-FL	-FL	-FL	-FA	-AA	-FA/DG	All
Factory installed	-FM	-FM	-FM	-FM	-FM	-FM	-FA	-AA	-FA/DG	All

DECstation/DECsystem 5000 Model 200 Series

Configuration rules:

- No internal devices; Max. 6 external devices
- Max. 3 expansion boxes (6 devices) per SCSI bus
- Up to 3 other Turbochannel options to expand number of SCSI buses & devices (6 per bus)

	External to System Box				
	TLZ04	TKZ08	TSZ07	RRD42	SZ12
Field installed	-FA	-AA	-CA	-FA/DG	All
Factory installed	-FA	-AA	-CA	-FA/DG	All

DECsystem 5100

Configuration rules:

- Max. 5 internal devices
- Max. 3 external enclosures
- Max. 7 total devices

	Internal to System Box						
	RZ23L	RZ24	RZ25	RX23	TZ30	TZK10	RRD42
Field installed	-EF	-EF	-EF	-EK	-EF	-FG	-EL
Factory installed	-EG	-EG	-EG	-EM	-EG	-FF	-EM

	External to System Box						
	TK50Z	TKZ08	TSZ05	TSZ07	TLZ04	RRD42	SZ12
Field installed	-GA/G3	-AA	-AA/BA	-CA	-FA	-FA/DG	All
Factory installed	-GA/G3	-AA	-AA/BA	-CA	-FA	-FA/DG	All

DECsystem 5500

Configuration rules:

- Max. 4 internal devices
- Max. 2 external enclosures
- Max. 7 total devices per SCSI adapter card

Note: Only SCSI options are shown here. See the System and Options Catalog for all devices.

	Internal to System Box			External to System Box			
	RZ56	RZ57	TLZ04	TSZ07	TLZ04	RRD42	SZ12
Field installed	E-AF	E-AF	-JF	-CA	-FA	-FA/DG	All
Factory installed	E-AA	E-AA	-JA	-CA	-FA	-FA/DG	All

MicroVAX/VAXserver 3100

Configuration rules:

- Max. 3 internal devices (Model 10e); Max. 5 internal devices (Model 20e)
- TZ30 tape drive uses the space of two devices in Model 10e systems
- Max. 4 external devices

	Internal to System Box					External to System Box					
	RZ23L	RZ24	RZ25	RX23	TZ30	TZK10	RRD42	TLZ04	TSZ05	RRD40	SZ12
Model 10e											
Field installed	-EF	-EF	-EF	-EP	-EF	-FG	-GL	-FA	-AA	-FA/F3	All
Factory installed	-EG	-EG	-EG	-ET	-EG	-FF	-GM	-FA	-AA	-FA/F3	All
Model 20e											
Field installed	-EF	-EF	-EF	-EK	-EF	-FG	-GL	-FA	-AA	-FA/F3	All
Factory installed	-EG	-EG	-EG	-EM	-EG	-FF	-GM	-FA	-AA	-FA/F3	All

B. *Detailed SCSI 3.5" Disk Product Specification sheets*

3.5" SCSI PRODUCT SPECIFICATIONS

PERFORMANCE	**RZ22**	**RZ23**	**RZ23L**	**RZ24**	**RZ25**
Capacity					
Unformatted (MB) [1]	N/A	N/A	N/A	N/A	493
Formatted (MB)	52	104	121.65	209.7	426
Per Surface (MB)	13	13	30.41	26.2	47.3
Per Track (Bytes) [8]	16,896	16,896	19,968	19,456	24,576-37,376
Per Block (Bytes)	512	512	512	512	512
BlocksTrack/Spare [9]	33,1	33,1	39,1	38,1	48,1-74,1
Blocks/Drive	102,432	204,864	237,588	409,792	832,031
Spare Blocks/Cylinder	4	8	4	8	9
Spare Blocks/Drive	3104	6208	6092	10,944	14,148
Spare Cylinders	0	0	0	0	2
Buffer Size (Kbytes)	16	16	64	64	60
Transfer Rate					
To/From Media MB/Sec [2]	1.25	1.25	1.5	1.5	2.1-3.2
Data Transfer Rate MB/Sec [3]	1.00	1.00	1.13	1.13	2.33
Bus Asynchronous MB/Sec	1.25	1.25	3.0	3.0	3.0
Bus Synchronous MB/Sec	N/A	N/A	4.0	4.0	4.0
Latencies					
Track to track seek (MS)	<8	<8	8	5	2.5
Average Seek (MS) [4]	<25	<25	<19	16	14
Maximum (full stroke) seek (MS)	<45	<45	<35	<35	28
Average Rotational Latency (MS)	8.4	8.4	8.8	8.6	6.8
Average Access (MS) [5]	33.4	33.4	26.8	24.6	20.8
Rotational speed (RPM)	3575 ± 0.1%	3575 ± 0.1%	3409 ± 0.5%	3497 ± 0.5%	4412 ± 0.5%
Start Time (Sec Max)	20	20	20	20	20
Stop Time (Sec Max)	20	20	20	20	30
Interleave	1:1	1:1	1:1	1:1	1:1
Functional					
Interface	SCSI	SCSI	SCSI	SCSI	SCSI-2
Recording Density at ID (BPI)	23,441	23,441	36250	31,800	38,843
Flux Density at ID (FCI)	15,627	15,627	24170	21,200	28,669
Track Density (TPI)	1150	1150	1850	1700	1760
Areal Density (MB/Sq.In)	26.96	26.96	67.06	54.06	68.28
Tracks per surface unformatted	N/A	N/A	N/A	N/A	N/A
Tracks per surface formatted	776	776	1523	1248	1474
Servo (Embedded/Dedicated)	EMB	EMB	EMB	EMB	DED
R/W Heads (Data)	4	8	4	8	9
Disks	2	4	2	4	5
Recording Mode [10]	CF	CF	CF	CF	ZBR

SCSI PRODUCT SPECIFICATIONS

PERFORMANCE	RZ22	RZ23	RZ23L	RZ24	RZ25
Software Support					
VMS Operating System Support	5.1-b	5.1-b	5.4-1	5.3	5.4-1
Ultrix (UWS)	2.2	2.2	4.1	2.2D	4.1
Physical					
Height	1.66"	1.66"	1.00"	1.6254	1.63
Width	4.00"	4.00"	4.0"	4.0	4.0
Depth	5.75"	5.75"	5.75"	5.75	5.75
Weight	1.72lb	1.82lb	1.2lb	1.7lb	1.8lb
Power					
Max Seeking (w)	8.3	8.3	3.8	6.6	14
Max Startup (w)	23.9	23.9	14.5	27.5	34.5
Environmental Limits					
Operating					
Ambient Temperature [6]	10°C - 60°C	10°C - 60°C	10°C - 55°C	10°C - 55°C	10°C - 55°C
Relative Humidity	8% - 80%	8% - 80%	8% - 80%	8% - 80%	8% - 80%
Max Wetbulb (Noncondensing)	25.6°C	25.6°C	25.6°C	25.6°C	25.6°C
Altitude	-1000 - 10,000FT	-1000 - 10,000FT	-1000 - 15,000FT	-1000 - 15,000FT	-1000 - 10,000FT
Heat Dissipation					
Typical	8W	8W	3.6W	6.6	10W
Random Seek	9W	9W	3.8	6.8	12W
Temperature Gradient	11°C (20°F)per hr	11°C (20°F)per hr			
Nonoperating					
Ambient Temperature	-40°C - 66°C				
Relative Humidity	8% - 95%				
Max Wet Bulb (noncondensing)	46°C				
Temperature Gradient	20°C (36°F)per hr				
Shock [7]	10 Gpk 10 ± 3ms				
Vibration	5-22 HZ, 0.01 in DA, 22-500 HZ, 0.25Gpk				
Acoustic Noise					
Idle (bels)	4.3	4.3	4.3	4.3	4.6
Seeking (bels)	5.0	5.0	5.0	5.0	5.0
Error Rates					
Soft Errors	<1E10				
Hard Errors	<1E13				
Seek Errors	<1E6				
Date Available	Feb 89	Feb 89	Apr 91	Apr 90	Jul 91

FOOTNOTES

1. Drives with dedicated servo systems differ in unformatted vs formatted capacity. It is not uncommon for vendors to quote unformatted capacity since it makes the drive appear to have a higher capacity. Digital quotes usable capacity.

2. Specified media Transfer rates include spare blocks, headers, servo data etc. See note 10.

3. The real (user measurable) data rate (this is sometimes called the on-track data rate) can be calculated as:

$$\frac{512 \text{ X Data Sectors per track}}{\text{Time required for 1 revolution}} \quad \text{or for the RZ57} \quad \frac{512 \text{ X } 71}{16.6} = 2.19 \text{ MB/Sec}$$

(A 3600 RPM drive requires 16.6MS to make 1 revolution)

4. The average seek is the time required to traverse 1/3 of the data area of the disk. This is the average time required to position the heads over the correct track.

5. Average access is the real time required for the heads to be positioned to transfer data. It is the sum of the average rotational latency and average seek times. It is a more realistic measurement and should be the quoted figure of data location time.

6. Most vendors specify 10°C - 45°C. Digital specifies a minimum of 55°C. This higher temperature limit is critical in noncomputer room environments where ambient temperatures are higher.

7. Most vendors specify 2 G's operating shock vs 10G's for Digital products. This is critical in desktop environments.
8. ZBR Drives consist of a range of bytes per track. See note 10.
9. ZBR Drives consist of a range of sectors per track. See note 10.
10. Two recording modes are employed in the RZ products. They are continuous frequency (CF) and Zone Bit Recording (ZBR). In continuous frequency recording, the same number of sectors are recorded on all tracks. In Zone bit recording, the disk is divided into logical bands or zones like the rings of a bulls eye. The minimum number of sectors per track is found in the inner most zone and the maximum number of sectors per track in the outermost zone. ZBR trends to maximize drive performance.

NOTE: All performance specifications are maximum averages. Individual drives may perform better than specified.

C. Detailed SCSI 5.25" Disk Product Specification sheets

SCSI PRODUCT SPECIFICATIONS

PERFORMANCE	RZ55	RZ56	RZ57
Capacity			
Unformatted (MB) [1]	380	760	1200
Formatted (MB)	332.3	665.17	1000
Per Surface (MB)	22.48	50.90	66.70
Per Track (Bytes)	18.432	31,248	35,976
Per Block (Bytes)	512	512	512
BlocksTrack/Spare	36,0	54,0	71,0
Blocks/Drive	649,040	1,299,174	1,954,050
Spare Blocks/Cylinder	8	11	11
Spare Blocks/Drive	10,300	20,316	96,009
Spare Cylinders	3	3	71
Buffer Size (Kbytes)	64	64	256
Transfer Rate			
To/From Media MB/Sec [2]	1.25	1.875	2.2
Data Transfer Rate MB/Sec [3]	1.11	1.66	2.19
Bus Asynchronous MB/Sec	1.5	1.6	1.6
Bus Synchronous MB/Sec	4.0	4.0	4.0
Latencies			
Track to track seek (MS)	<4	<4	<4
Average Seek (MS) [4]	<16	<16	<14.5
Maximum (full stroke) seek (MS)	<35	<35	<33
Average Rotational Latency (MS)	8.3	8.3	8.3
Average Access (MS) [5]	24.3	24.3	22.8
Rotational speed (RPM)	3600 ± 0.5%	3600 ± 0.5%	3600 ± 0.5%
Start Time (Sec Max)	20	20	20
Stop Time (Sec Max)	20	20	20
Interleave	1:1	1:1	1:1
Functional			
Interface	SCSI	SCSI	SCSI
Recording Density at ID (BPI)	21,231	31,846	42,441
Flux Density at ID (BPI)	14,154	21,231	31,830
Track Density (TPI)	1075	1440	1670
Areal Density (MB/Sq.In)	22.82	45.86	70.88
Tracks per surface unformatted	1224	1632	1928
Tracks per surface formatted	1218	1626	1854
Servo (Embedded/Dedicated)	DED	DED	DED
R/W Heads (Data)	15	15	15
Disks	8	8	8
Recording Mode [10]	CF	CF	CF

SCSI PRODUCT SPECIFICATIONS

PERFORMANCE	RZ55	RZ56	RZ57
Software Support			
VMS Operating System Support	5.1-b	5.2-1	5.3
Ultrix (UWS)	2.2	2.2D	2.2D
Physical			
Height	3.25"	3.25"	3.25"
Width	5.75"	5.75"	5.75"
Depth	8.2"	8.2"	8.2"
Weight	8.4lb	8.4lb	8.4lb
Power			
Max Seeking (w)	25.2	32	32
Max Startup (w)	52.9	60.7	68.2
Environmental Limits			
Operating			
Ambient Temperature [6]	10°C - 55°C		
Relative Humidity	8% - 80%		
Max Wetbulb (Noncondensing)	25.6°C		
Altitude	-1000 - 15,000FT		
Heet Dissipation			
Typical	28	28	28
Random Seek	32	32	32
Temperature Gradient	11°C (20°F)per hr		
Nonoperating			
Ambient Temperature	-40°C - 66°C		
Relative Humidity	8% - 95%		
Max Wet Bulb (noncondensing)	46°C		
Temperature Gradient	20°C (36°F)per hr		
Shock [7]	10 Gpk 10 ± 3ms		
Vibration	5-22 HZ, 0.01 in DA, 22-500 HZ, 0.25Gpk		
Acoustic Noise			
Idle (bels)	5.0	5.0	5.0
Seeking (bels)	5.7	5.7	5.7
Error Rates			
Soft Errors	<1E10		
Hard Errors	<1E13		
Seek Errors	<1E6		
Date Available	Feb 89	Nov 89	Apr 90

FOOTNOTES

1. Drives with dedicated servo systems differ in unformatted vs formatted capacity. It is not uncommon for vendors to quote unformatted capacity since it makes the drive appear to have a higher capacity. Digital quotes usable capacity.

2. Specified media Transfer rates include spare blocks, headers, servo data etc. See note 10.

3. The real (user measurable) data rate (this is sometimes called the on-track data rate) can be calculated as:

$$\frac{512 \text{ X Data Sectors per track}}{\text{Time required for 1 revolution}} \quad \text{or for the RZ57} \quad \frac{512 \text{ X } 71}{16.6} = 2.19 \text{ MB/Sec}$$

(A 3600 RPM drive requires 16.6MS to make 1 revolution)

4. The average seek is the time required to traverse 1/3 of the data area of the disk. This is the average time required to position the heads over the correct track.
5. Average access is the real time required for the heads to be positioned to transfer data. It is the sum of the average rotational latency and average seek times. It is a more realistic measurement and should be the quoted figure of data location time.
6. Most vendors specify 10°C - 45°C. Digital specifies a minimum of 55°C. This higher temperature limit is critical in noncomputer room environments where ambient temperatures are higher.
7. Most vendors specify 2 G's operating shock vs 10G's for Digital products. This is critical in desktop environments.
8. ZBR Drives consist of a range of bytes per track. See note 10.
9. ZBR Drives consist of a range of sectors per track. See note 10.
10. Two recording modes are employed in the RZ products. They are continuous frequency (CF) and Zone Bit Recording (ZBR). In continuous frequency recording, the same number of sectors are recorded on all tracks. In Zone bit recording, the disk is divided into logical bands or zones like the rings of a bulls eye. The minimum number of sectors per track is found in the inner most zone and the maximum number of sectors per track in the outermost zone. ZBR trends to maximize drive performance.

NOTE: All performance specifications are maximum averages. Individual drives may perform better than specified.