

4.6.5—HP StorageWorks Disk System SC10

Note: Effective June 1, 2002, the SC10 will be discontinued. Customers can continue purchasing add-on disk drives until July 1, 2003 or while supplies last. The HP Disk System 2300 is the replacement product for the SC10.

Description

Hewlett-Packard's Disk System SC10 (referred to in this guide as the disk system or SC10) is a high-availability, AC or DC-powered Ultra2 SCSI storage product. Dual SCSI ports on dual bus controllers provide Ultra2 SCSI 80 MB/s LVD connections to the host. Ten slots accept high-speed, high-capacity LVD/SE SCSI disks connected to an SE backplane. Data throughput is 40 MB/sec per SCSI bus, with each enclosure capable of two separate SCSI bus connections to the host for an aggregate throughput of 80 MB/sec. Eleven disk systems fill a 2-meter System/E rack, yielding up to 8.0 TB of storage.

Modular and redundant components are easy to upgrade and maintain. Disks, fans, power supplies, and Bus Control Cards (BCCs) are replaceable parts that plug into individual slots in the front and back of the disk system. Redundant fans, AC or DC power supplies, and BCCs, when in full bus mode with redundant host connections, can be removed and replaced without interrupting storage operations. Disks also can be replaced with the system on and with only the affected file systems taken offline. Hewlett-Packard technical support is optional for these procedures.

Special electronics and HP-UX software enable remote monitoring and diagnostics. Sensors on the BCCs monitor the disk system environment, including temperature, voltage, fan speed, and component status. Hewlett-Packard's Event Monitoring Service (EMS) reports any changes in environmental status to user-defined locations, including email accounts and a system console. Standard HP-UX diagnostic utilities also report environmental data for enhanced troubleshooting.

Figure 4.6.5.1 HP Disk System SC10



Table 4.6.5.1 HP Disk System SC10—Field Rackable

A5272A		Disk System SC 10, AC-powered
	001	Add-on Bus Controller Card
	150	18.2-GB 10K RPM Ultra3 SCSI Drive (Note: Order quantity 1 to 10 per enclosure)
	250	36.4-GB 10K RPM Ultra3 SCSI Drive (Note: Order quantity 1 to 10 per enclosure)
	350	73.4-GB 10K RPM Ultra3 SCSI Drive (Note: Order quantity 1 to 10 per enclosure)
	450	18.2-GB 15K RPM Ultra3 SCSI Drive (Note: Order quantity 1 to 10 per enclosure)
	550	36.4-GB 15K RPM Ultra3 SCSI Drive (Note: Order quantity 1 to 10 per enclosure)
Z7536A		Disk System SC10, DC-Powered

Table 4.6.5.2 HP Disk System SC10—Factory Racked

A5272AZ		Disk System SC 10
	001	Add-on Bus Controller Card
	150	18.2-GB 10K RPM Ultra3 SCSI Drive (Note: Order quantity 1 to 10 per enclosure)
	250	36.4-GB 10K RPM Ultra3 SCSI Drive (Note: Order quantity 1 to 10 per enclosure)
	350	73.4-GB 10K RPM Ultra3 SCSI Drive (Note: Order quantity 1 to 10 per enclosure)
	450	18.2-GB 15K RPM Ultra3 SCSI Drive (Note: Order quantity 1 to 10 per enclosure)
	550	36.4-GB 15K RPM Ultra3 SCSI Drive (Note: Order quantity 1 to 10 per enclosure)

Table 4.6.5.3 Upgrade Kits

A6272A		Add-on 18.2-GB 15K RPM Ultra3 SCSI Disk Module
A6273A		Add-on 18.2-GB 15K RPM Ultra3 SCSI Disk Module
A6274A		Add-on 36.4-GB 10K RPM Ultra2 SCSI Disk Module
A6275A		Add-on 36.4-GB 15K RPM Ultra3 SCSI Disk Module
A6276A		Add-on 73.4-GB 10K RPM Ultra3 SCSI Disk Module
A5273A		Add-on Bus Control Card
A5250A		Rack Rail Accessory Kit
A5251A		HP Rack System/E Rack Rail Access Kit
A5612A		Field Install Swing Out Bracket for PDU
A5612AZ		Factory Install Swing Out Bracket for PDU
A5629A		Add-on Blower Module
A5630A		Add-on Power Supply

Table 4.6.5.4 Cables

C2373A	0D1	2-meter VHDTS68 (M/M) Multimode
C2374A	0D1	5-meter VHDTS68 (M/M) Multimode
C2375A	0D1	10-meter VHDTS68 (M/M) Multimode
A5668A	0D1	2-meter VHTDS68 LVD/SE ILT (M/M)
A5669A	0D1	5-meter VHTDS68 LVD/SE ILT (M/M)
A5670A	0D1	10-meter VHTDS68 LVD/SE ILT (M/M)

Features

The disk system occupies 3.5 EIA units in a standard 19-inch rack. Disk drives mount in the front of the system. Redundant power supplies, fans, and BCCs mount in the back. A lockable front door shields the environment from RFI and provides access to the disk drives and power button.

High Availability

High availability is a general term describing computer systems that are designed to minimize planned and unplanned downtime. Hewlett-Packard's Five Nines Initiative promises systems with 99.999% availability (less than 5 minutes downtime per year) by 2001. The SC10 provides high availability through the following features:

- Hot-pluggable, high-capacity, high-speed disks
- Redundant, hot-pluggable, and user-replaceable fans, power supplies, and BCCs
- Support for mirrored disks in the HP-UX environment
- Online BCC firmware upgrades.
- Hardware event monitoring and real-time error reporting
- The DC-powered version of the Disk System SC10 has been NEBS Level-3 certified.

Note: Since disks require exclusive access, data is unavailable during the disk firmware download.

Environmental Services

Environmental services circuitry monitors the following elements:

- Fan rotation
- Fan status
- Power supply outputs
- Power supply status
- Disk drive status, except fault condition
- BCC status
- Temperature
- Self-test results

Each BCC reports the status of all elements in the disk system, even if the BCC does not have direct access to the element. Additionally, the EEPROM on each BCC stores 2 KB of configuration information and user-defined data, including the manufacturer serial number, and product number.

Hardware Event Monitoring

A hardware event monitor specifically for the Disk System SC10 is part of Hewlett-Packard's Event Monitoring System (EMS) for HP-UX. This special software monitors the disk system and reports changes in environmental status, which are called events. Hardware event monitoring is an important tool for implementing high availability. Using hardware event monitors, you can virtually eliminate undetected hardware failures that interrupt system operation or cause data loss.

The EMS Hardware Monitors User's Guide is included in Adobe® Acrobat® format on the IPR Support Media.

Status Indicators

LEDs on the disk system enable you to detect and replace failed components and thereby prevent or minimize downtime for users.

On the front of the disk system, a pair of LEDs indicate the status of the disk system, and an LED for each slot shows disk I/O activity:

- The system power LED indicates that power is on or off.
- The system fault LED indicates whether or not a fault has occurred in the BCCs of the disk system.
- Above each disk slot, the left LED indicates the presence of I/O activity on the disk.
- The second LED above each disk slot can be flashed to help a service engineer locate the disk for physical inspection or removal.
- LEDs on the back of the disk system indicate the status of replaceable components and the configuration of the SCSI bus: See chapter 4, Troubleshooting, of the User and Service Guide, for specific LED information.

Upgradability

You can increase disk system storage capacity by:

- Replacing disk drives with higher-capacity disk drives
- Adding disks in unused slots
- Adding another disk system (to fill a bus)

None of these actions require shutting down the product, but some may require the use of system utilities to manage file systems. Upgrade BCC and disk firmware using a download function.

Capacity

The Disk System SC10 has ten (10) slots for storage mechanisms. Low-profile devices (1 inch high) as well as half-height devices (1.6 inches high) require 1 slot each, therefore a maximum of ten low-profile devices or a maximum of ten half-height devices can be inserted. The enclosure can accept a mixture of low-profile and half-height mechanisms, or, stated another way, the enclosure can accept multiple capacity points such as 9 and 18-GB HDDs simultaneously, subject to file system (LVM) restrictions.

Power Switch

Located in the upper right corner behind the disk system door, the power interrupts power from the power supplies to the BCCs and other internal components. Power to the power supplies is controlled by the power cords and the AC source.

Table 4.6.5.5 HP Disk System SC10 Summary

Functionality	HP Disk System SC10	
	Rackmount	
Number of 3.5 inch mechanisms	10 LP ¹ – 10 HH ¹	
Disk Drive Support		
9 GB 10K Ultra2 SCSI	Yes	
18.2 GB 10K Ultra2 SCSI	Yes	
18.2 GB 15K Ultra2 SCSI	Yes	
36 GB 10K Ultra2 SCSI	Yes	
73 GB 10K Ultra2 SCSI	Yes	
Enclosures Per Rack		
	HP Standard Rack	HP Rack System/E
1.1-meter rack	4	4
1.6-meter rack	8	8
2-meter rack	10	11
High Availability Functionality		
Hot Plug	Yes	
Redundant Fans	Yes	
Redundant Power Supplies	Yes	
Number of power cords	2 with 2nd PS	
Number of Ultra SCSI Busses	2	

¹ LP - Low-Profile, 1 inch high mechanisms HH - Half-Height, 1.6 inch high mechanisms

Powerfail Protection

A UPS or user provided power protection is required for powerfail protection for the Disk System SC10. The 5.2-kVA UPS is recommended for use with the Disk System SC10 in AC-powered configurations.

Cables

The cables must be explicitly ordered for both factory-racked and field-racked configurations. The cables may be ordered via options, 860 (2-meter) or 861 (5-meter), or ordered separately by product, A5306A (2-meter) or A5307A (5-meter).

Configuration

Step 1: Determine the required number and type of Ultra2 SCSI storage mechanisms. For example, 1 TB of protected storage requires 220 9.1-GB 10K Disks, or 110 18.2 GB 10K Disks. It is recommended that the SC10 be configured with at least one disk per enclosure, regardless of the number of power supplies. Each solid state device is counted as one disk when calculating the number of disks in an enclosure. (In order for customers to purchase an SC10 with less than four drives, they should order A5272A and select the number of drives needed.)

Step 2: Determine the number of SC10 enclosures required. Each SC10 has a spindle capacity of ten. Ten 9.1-GB LP or 10 18.2-GB HH. Make your selection based upon customer data requirements. For example, a higher spindle count usually means higher performance in OLTP, or transaction based environments. Data warehousing, DSS, and scientific applications tend to favor larger spindle capacities, however this is not true in all cases. Also consider that protected storage configurations require twice the capacity of unprotected JBOD. For example, a 1 TB data warehouse will require 2 TBs of disk space in a mirroring configuration.

Consider the following ways to achieve 1 TB of protected storage with the SC10:

Table 4.6.5.6

	Spindle Count	SC10 Chassis
9.1 GB	220	22
18.2 GB	110	11

The number of enclosures required for a given capacity is equal to the spindle count/divided by ten plus an additional enclosure for any remaining spindles.

Step 3: Determine the required overall configuration of the SC10. The SC10 supports high availability through a variety of configurations. Each redundant bus in the storage system can be connected to a different adapter in the same or different hosts.

Disk Addressing

Dip switches on BCC A (in the top slot) determine bus architecture, bus addressing, and some bus behavior. If BCC A is removed, then BCC B switch settings define the bus. Switch positions are described

Table 4.6.5.7 Dip Switch Settings

Switch	ON = 1	OFF = 0
1 Full Bus	Forms a single bus of 10 addresses	Defines split buses, 5 addresses each
2 JBOD	Sets disk-to-host configuration for JBOD mode	Sets disk configuration for array mode
3 Bus Reset: Power Fail	Automatically resets the bus when power fails	No reset on power fail
4 Address High	Assigns high SCSI addresses (8-12) to disks on split bus	Assigns low SCSI addresses (0-3, 13 when switch 2 is set to "1") to disks on split bus
5 Bus Reset: Hot Swap	Automatically resets the bus whenever a disk is removed or inserted	No reset on disk removal or insertion

A few rules govern the setting of these switches.

Note: *The switch settings on BCC A and B must match.*

The BCC self-test detects any discrepancy in the dip switch settings between BCC A and BCC B., The buzzer sounds 5 times, and the BCC's fault LED and system fault LEDs flash. If the mismatch is discovered at enclosure powered on, the disks do not spin up. If the system is operating, the disks and second BCC continue to operate.

1. When switch 1 (Full Bus) is "1", switch 4 (Address High) has no effect.
2. When switch 1 (Full Bus) is "0" and split buses are daisy-chained to another disk system, switch 4 (Address High) must be "1" for one of the disk systems and "0" for the other.

Note: *High and low addresses are not supported in the same disk system.*

Each disk in the SC10 disk system occupies a separate address (SCSI ID) on the SCSI bus. Disk addresses range from 0 to 3 and 8 to 13 in full-bus mode. In split-bus mode, disks are addressed according to the High Address switch setting: 0, 1, 2, 3, and 13 when High Address is "0"; 8, 9, 10, 11, and 12 when High Address is "1". Addresses 4, 5, 6, and 7 are reserved for host bus adapters. The BCCs take addresses 14 and 15. This is for JBOD mode. Dip switch #2 has to be "1" for this to be true.

SCSI IDs

BCC A generates all bus addresses, in full-bus and split-bus modes. If BCC A is removed, then BCC B generates the bus addresses. BCC A resumes the addressing function as soon as it is reinstalled.

Table 4.6.5.8 shows all target SCSI IDs for full-bus and split-bus modes and for low and high addressing.

Table 4.6.5.8 Disk and BCC SCSI Addresses for Full and Split Bus Modes*

Slot	Full Bus	Split Bus	
		High Addr = 0	High Addr = 1
0	0	0	8
2	1	1	9
4	2	2	10
6	3	3	11
8	13	13	12
BCC A	14	15	14
1	8	0	8
3	9	1	9
5	10	2	10
7	11	3	11
9	12	13	12
BCC B	15	15	14

* Values assume JBOD mode, i.e., dip switch 2 is set to "1".

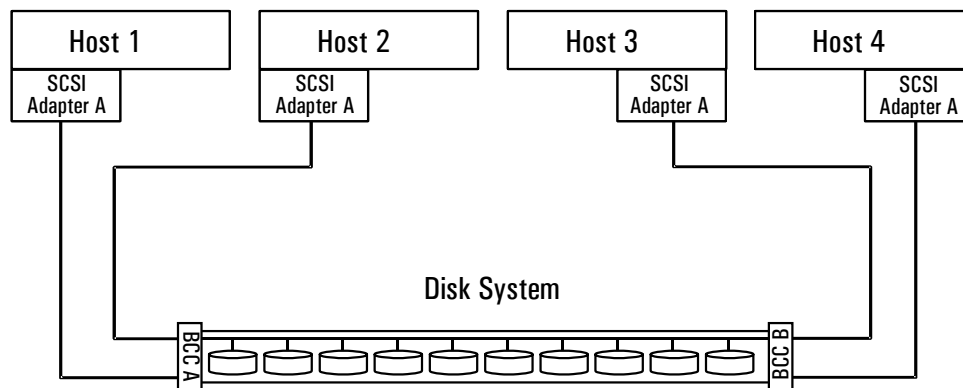
Topologies

The SC10 supports high availability through redundant components and redundant connections to redundant hosts. Each SCSI port on a BCC can be connected to a different host bus adapter in the same or different hosts.

Single Host, One Disk System

Connecting one disk system to one host will not achieve high availability. The SC10 is supported in single host; single enclosure environments for those customers who desire a cost reduced implementation and are willing to trade off high availability features. The SC10 supports data mirroring between the two internal busses of the same enclosure. It is expected that customers will use this configuration for performance reasons, and should not expect any level of data availability from a single host, one disk system configuration.

Figure 4.6.5.2 Four Initiators and One Disk System




Redundant Hosts, One Disk System

Connecting one disk system to redundant hosts achieves system level high availability. In **Figure 4.6.5.2**, a single host bus adapter (HBA) in each host is connected to a different port in the disk system. With the disk system in full bus mode (switch 1 on), each host can reach all ten disks. If BCC A fails in this topology, there are still two paths to the disks through BCC B. With the disk system in split bus mode (two internal busses), the SC10 supports data mirroring between the two internal busses within the same enclosure in MC/ServiceGuard configurations. All connections from the host to the disk system are SCSI LVD cables.

Redundant Hosts, Mirrored Disk Systems

Another high availability topology connects mirrored disk systems to redundant hosts. In **Figure 4.6.5.3**, dual host bus adapters (HBA) in each host are connected to mirrored disk systems. With the disk systems in full bus mode (switch 1 on), each host can reach all ten disks in both disk systems. If one of the disk systems fails in this topology, all four hosts still have access to the data on the mirrored system. All connections from the host to the disk system are SCSI LVD cables.

Figure 4.6.5.3 Four Initiators and Mirrored Disk Systems



Racking Considerations

Currently the SC10 is available only in a racked configuration. When installing the SC10 in an existing customer rack, be sure to determine the rack model. The SC10 ships with rack rail kits for either existing HP rack series or the new Rack System/E. In addition, determine whether or not the SC10 is to be racked at the factory. If so, assume the A4900A, A4901A, or A4902A rack models. If installing the SC10 in a Rittal 9000 series rack, order A5656A. Also, order A5612A or A5612AZ when installing the SC10 in a rack with a 30-amp PDU.

- The SC10 takes 4U per enclosure plus rail set in the A1883A/A1896A, A1884A/A1897A, A3656A.
- The SC10 takes 3.5U per enclosure and rail set in the A4900A, A4901A, or A4902A rack models.
- Redundant PDU configurations as shown are defined as including redundant power cords to the wall. Non-redundant PDU (power) configurations can be built but are not recommended for high availability environments.

The following table details the available configurations for the HP Rack Model A1883A/A1896A, A1884A/A1897A, A3656A.

Table 4.6.5.9 Maximum Number of SC10 Enclosures per HP Rack (A1883A/A1896A, A1884A/A1897A, A3656A)

Required PDU/PDRU Type	1.1-meter (21 U)	1.6-meter (32 U)	2.0-meter (41U)	Rack Switch Support
2× 3 ft., 16 Amp	4	—	—	Yes
2× 5 ft., 16 Amp	—	4	4	Yes
2× 19 in., 16 Amp	4	4	4	—
4× 19 in., 16 Amp	—	—	8	—
2× 19 in., 30 Amp	—	8	8	—
2× 19 in., 30 Amp and 2× 19 in., 16 Amp	—	—	10	—
2× 19 in., 60 Amp	—	—	10	—

Table 4.6.5.10 Maximum Number of SC10 Enclosures per HP Rack System/E (A4900A/J1502B, A4901A/J1501B, A4902A/J1500B)

Required PDU/PDRU Type	Power Cords to Wall	1.25-meter (25 U)	1.6-meter (33 U)	2.0-meter (41U)	Rack Switch Support
2× 16 Amp or 2× 30 Amp	2	4	4	4	Yes, add 1x dual 16-amp PDRU
4× 16 Amp or 2× 30 Amp	4	–	8	8	Yes, add 2× 30-amp PDRU
	2	–	8	8	Yes
6× 16 Amp or 2× 30 Amp and 2× 16 Amp or 4 30 Amp	6	–	–	11	Add 4× 30-amp PDRU
	4	–	–	11	–
	4	–	–	11	Yes

1. SC10 takes up 4U per box plus rail set in Rack System, and 3.5U per box and its rail set in Rack System/E
2. Redundant PDU configurations shown are defined as including redundant power cords to the wall. Non-redundant PDU configurations can also be constructed per User needs.
3. Groupings show solution options to achieve 4, 8, or 11 SC10 in the racks. Options within the groups vary w.r.t. the number of power cords going the wall AC.
4. (*) These are traditional (existing for HP Rack System) PDUs.
5. (**) To achieve the highest rack densities, large PDRUs must be placed behind the storage systems. Installed with hinged brackets, these PDUs can rotate up or down for convenient servicing of rear FRUs.
6. The 60-amp PDRU will not be available at first release.
7. The on/off switch option for HP Rack System/E will work only with the 2×16A PDUs (need to add the dual 16A PRU), and with the large PDRUs. The 4×16A and the 6×16A combinations do not support the on/off switching. Combinations of 16-amp plus 30- or 60-amp PDU/PDRUs can be configured to always support on/off switching.
8. The 4 and 6×16A configurations will be susceptible to high inrush currents during start-up scenarios. The 30- and 60-amp PDRUs are designed to minimize inrush compared to the 16A PDUs.

HBA Considerations – Ultra2 SCSI Low Voltage Differential Host Bus Adapters

- PCI Single-port Ultra2/Wide/Low Voltage Differential SCSI (A5149A).
- PCI Dual-port Ultra2/Wide/Low Voltage Differential SCSI (A5150A).
- 80 MB/s burst speed.
- Supports up to 15 Ultra2 or SE peripherals.
- A5149A used with A-Class (180A/180C), V-Class (V2500) and N-Class (N4000) Enterprise servers.
- A5150A used with N-Class only.
- Do not connect FWD SCSI-2 peripherals to Ultra2 interface cards. They are not compatible.
- Ultra2 interface cards and FWD SCSI-2 interface cards may coexist in the same system.

The SCSI address of a device dictates the device’s priority when arbitrating for the SCSI bus. SCSI address “7” is the address for the highest priority device (and is usually reserved for the host). Address “7” is followed in priority (from highest to lowest) by the subsequent addresses 6, 5, 4, 3, 2, 1, 0, 15, 14, 13, 12, 11, 10, 9, and 8. When used with the SC10, only addresses 4, 5, 6 and 7 are available for Host Bus Adapters.

Under extreme I/O workload conditions, it is possible for lower priority devices on an Ultra2/Wide/Low Voltage Differential bus to be “starved” for data throughput and time out. A retry will be automatic.

Devices on the Ultra2 SCSI bus are connected to each other in a “daisy chain.” The first and last devices on the SCSI bus must provide proper termination on the bus. A terminator, which fits on a SCSI connector, is shipped with the adapter card and can be used to terminate the last device on the SCSI bus.

Table 4.6.5.11 Ultra2/Wide/Low Voltage Differential SCSI Specifications

A5149A	<ul style="list-style-type: none"> • PCI Ultra2/Wide/Low Voltage Differential SCSI • SCSI-3 compliant • 16-bit SCSI bus • 80 MB/s burst speed • Parity support • Maximum cable length is 12 meters • Connects up to 15 devices per adapter
--------	---

HP 9000 Enterprise Server and Visualize Workstation support for Ultra2/Wide/Low Voltage Differential SCSI

Please refer to the NSSO Configuration Matrix at <http://turbo.rose.hp.com/spock/> for the most up to date configuration information.

Table 4.6.5.12 summarizes the number of Ultra2 cards supported on the various HP 9000 Enterprise Servers.

Table 4.6.5.12

	V-Class
PCI slots	28
Number of PCI Ultra2 cards supported	20

Step 5: Determine the cabling requirements. One cable per bus is required to interface to the HBA of the host, and additional optical cables are required for redundant connections.

Step 6: Determine the count and type of PDUs required. See above PDRU table.

Miscellaneous

Operating System Support

Functionality	HP Disk System SC10
HP-UX 9.04	No
HP-UX 10.0	No
HP-UX 10.01	No
HP-UX 10.10	No
HP-UX 10.20	Yes
HP-UX 11.00	Yes
HP-UX 11.01	Yes
Windows NT Server v4.0 SP3	No
Windows NT Server Enterprise Edition v4.0 SP3	No
MPE	No
Boot - Single-host	Yes ¹
Boot - Multi-host	No
Root	Yes
Swap	Yes
Install	Yes
LVM	Yes
SwitchOver/UX	No
MirrorDisk/UX	Yes
MC/ServiceGuard	Yes
Cluster Lock Disk	Yes
MC/LockManager with OPS environments up to 4 nodes	No
Hot Swap	Yes

¹ The SC10 is supported on Superdome as a boot device.

Electrical Requirements

All electrical wiring to the service point (plug) must be sized to carry the following inrush and steady state currents:

Table 4.6.5.13 Inrush (Surge) Current and Duration

Number of Disk Systems on Circuit	Inrush Currents and Duration
1	Up to 20 amps for 10 to 12 cycles
2	Up to 40 amps for 10 to 12 cycles
3	Up to 60 amps for 10 to 12 cycles
4	Up to 80 amps for 10 to 12 cycles

Table 4.6.5.14 Maximum Operating Current

Incoming Voltage AC RMS	Maximum RMS Current Drawn by One Disk System
100-120 volts	7.7 amps
200-240 volts	3.6 amps

Circuit breakers must be adequately rated for inrush and operating currents. Hewlett-Packard recommends magnetic-type circuit breakers, which are capable of handling large inrush currents for short durations (10 to 12 cycles) and are rated adequately for steady state currents. In Europe, install the following breaker types:

Table 4.6.5.15 Recommended European Circuit Breakers*

Number of Disk Systems	Breaker Rating	Breaker Type
1 to 3	16 amps	Type C per IEC 898 or Type D per IEC 898
4	16 amps	Type D per IEC 898 or Type K per IEC 947-2

* Data assumes no other devices share the circuit breaker.