sgi™

Cray SV1™ SuperCluster® Maintenance Guide
(Cray SV1 Series)

108-0246-001

SGI Proprietary

Any shipment to a country outside of the United States requires a letter of assurance from Silicon Graphics, Inc.

Direct requests for copies of service publications to:

Web site: *http://servinfo.csd.sgi.com*

SGI Customer Service - Logistics
1100 Lowater Road
P. O. Box 4000
Chippewa Falls, WI 54729-0078
USA

Direct comments about this publication to:

E-mail: *spt@sgi.com*

SGI Customer Service
Service Publications and Training
890 Industrial Blvd.
P.O. Box 4000
Chippewa Falls, WI 54729-0078
USA

Cray SV1 SuperCluster Maintenance Guide

108-0246-001
Cray SV1 SuperCluster
Last Modified: October 1999

Figures

Tables

Record of Revision

October 1999

Original printing.

Related Publications

The following publications contain more detailed information about the Cray SV1 mainframe, scalable I/O architecture (SIO), and UNICOS system administration.

Publication Number	Title
007-3904-002	*Cray SV1 Hardware Installation Procedures (Cray SV-1 Systems)*
108-0196-001	*System Overview (Cray SV1 Series Systems)*
HMK-396-0	*SIO Maintenance Kit*
004-2253-001	*Cray SV1 SuperCluster Administrator's Guide*
004-2210-002	*UNICOS Basic Administration Guide for CRAY J90, CRAY J90se, and CRAY SV1 GigaRing based Systems*
004-2204-002	*SWS-ION Administration and Operations Guide*
004-5310-002	*UNICOS Installation Guide for Cray SV1 SuperCluster Systems*

Ordering Publications

You may order Cray SV1 maintenance documents from Logistics in Chippewa Falls, WI, or from the Service Publications Web site at: `http://servinfo.csd.sgi.com/cgi-bin/pubsorder.pl`

To order a printed Cray Research document by telephone, call +1-651-683-5907. SGI employees may also order printed Cray Research documents by sending their orders via electronic mail to orderdsk@sgi.com (UNIX system users).

Overview

The Cray SV1 SuperCluster systems (refer to Figure 1) are high-performance supercomputers that are software and hardware compatible with other parallel vector Cray supercomputers. The significant capabilities of the Cray SV1 SuperCluster system include high-speed scalar, dual-pipe vector, and parallel processing; a large, fast central memory; and high-performance input/output.

The Cray SV1 SuperCluster system is a *capacity* cluster in which resources are managed and allocated across the cluster. (A *capability* cluster enables you to run a single job across multiple processors.)

Capacity clusters target jobs to the least-used compute resources. For example, if a new job is launched for execution, it is sized in terms of memory needed, disk needed, and so on, and sent to the most appropriate node in the cluster.

A single Cray SV1 cabinet is called a *node* in a SuperCluster system. Each *node* contains 2 to 32 Gbytes of main memory and 8 to 32 CPUs. Each CPU includes a 256-Kbyte, four-way associative cache for all scalar and vector data. A *building block* comprises four nodes and two I/O cabinets.

The maximum configuration for a Cray SV1 SuperCluster system is formed by combining 32 Cray SV1 nodes into a system, which is 8 building blocks or 32 mainframe cabinets (refer to Figure 1).

A Cray SV1 system node can also be configured with a combination of processors that are referred to as single-streaming processors (SSPs) and multi-streaming processors (MSPs). SSPs have two vector pipes and over 1 Gflop of peak performance.

Note: This document uses the term *single-streaming processor* (SSP) interchangeably with the term *CPU*.

An MSP is configured by combining four SSPs (with each SPP on a different physical processor module) to increase the memory bandwidth of the MSP. Each MSP has eight vector pipes and over 4 Gflops of peak performance. A fully configured Cray SV1 node can contain six MSPs and eight SSPs.

MSPs can be configured only by UNICOS. No hardware configuration changes are required to configure MSPs in the system.

Figure 1. Cray SV1 SuperCluster Systems

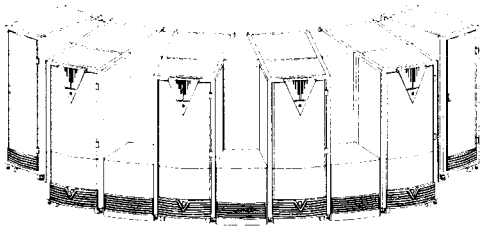

Cray SV1-4 System

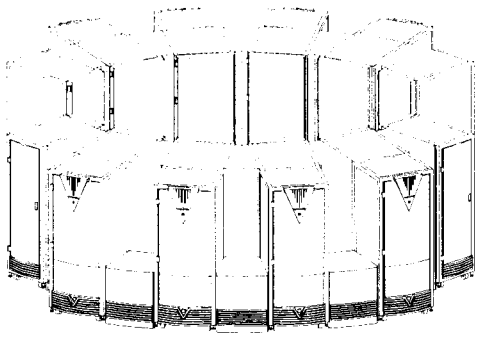

Cray SV1-8 System

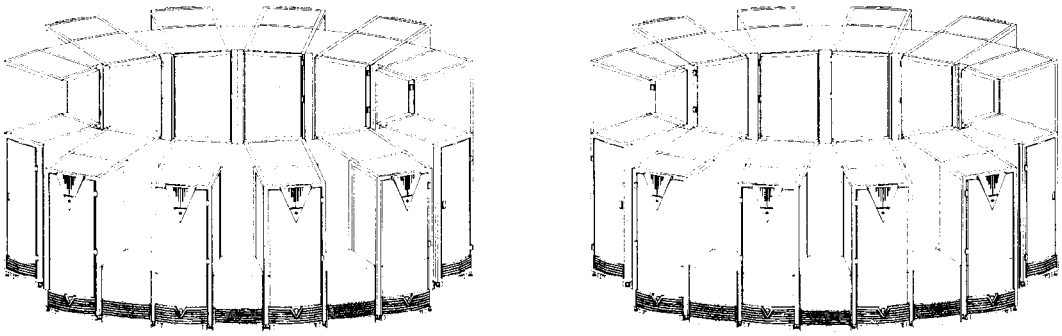

Cray SV1-16 System

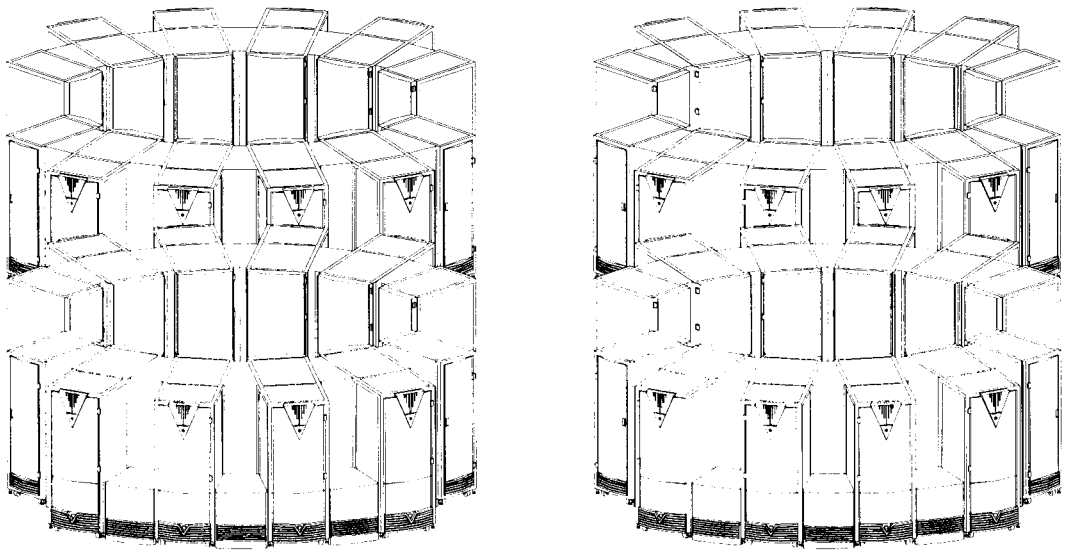

Cray SV1-32 System

System Configurations

The minimum hardware configuration for a Cray SV1 SuperCluster system is a Cray SV1-4 system, or building block. Each building block includes the following components:

- System workstation (SWS)
- Four Cray SV1 mainframe cabinets with two I/O cabinets
- GigaRing channel connection hardware
- One private GigaRing channel per mainframe
- One GigaRing channel per building block for intranode communication

The naming convention for the SuperCluster systems follows:

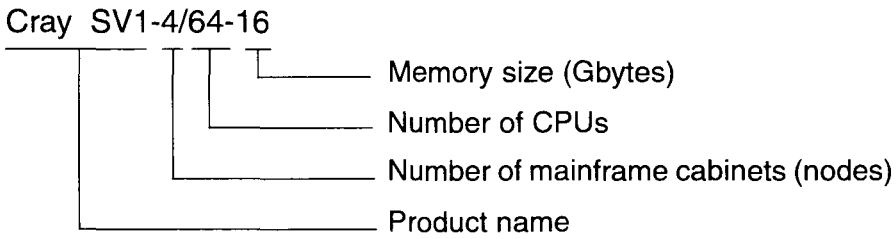

Cray SV1 SuperCluster systems can be expanded to a maximum of 32 nodes and 1,024 CPUs. Table 1 lists the various system CPU configurations. Because of limited space, figures and tables in this document may refer to the Cray SV1 systems as "SV1."

Table 1. Cray SV1 SuperCluster Configurations

SV1-4	SV1-8	SV1-12	SV1-16	SV1-20	SV1-24	SV1-28	SV1-32
Number of SSP Nodes							
4	8	12	16	20	24	28	32
Number of SSPs							
32 – 128	96 – 256	144 – 384	192 – 512	240 – 640	288 – 768	336 – 896	384 – 1024
Maximum Number of MSP at 6 MSPs per node							
24	48	72	96	120	144	168	192
Number of Remaining SSPs at Maximum Number of MSPs							
32	64	96	128	160	192	224	256
Combined Memory Size in Gbytes							
16 – 128	32 – 256	48 – 384	64 – 512	80 – 640	96 – 768	112 – 896	128 – 1024

Power and Cooling

SuperCluster configurations that include four or eight processing cabinets are air cooled. SuperCluster systems that include sixteen or more processing cabinets may be either air cooled or water cooled. All SuperCluster configurations require a raised-floor system.

An optional power conditioning and distribution system is available for Cray SV1 SuperCluster system configurations that include eight or more processing cabinets.

Cabinet Joiners

The Cray SV1 mainframe cabinets are bolted together in a semicircle with cabinet joiners (refer to Figure 2). The cabinet joiners also serve as electromagnetic interference (EMI) shields for the GigaRing and WACs cables that pass between each mainframe cabinet and between the PC-10 and mainframe cabinets.

Figure 2. Cabinet Joiner

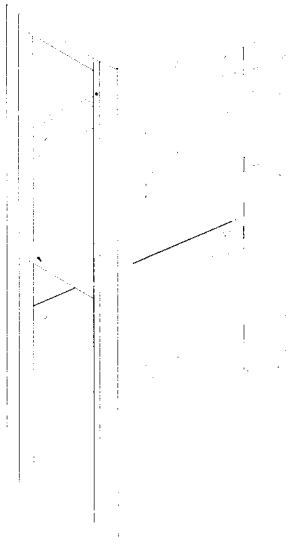

Cabling Bulkheads

SuperCluster cabling bulkheads are designated by their physical position in the cluster. Figure 3 and Figure 4 show the cable bulkheads on a Cray SV1-4 cluster.

Mainframe bulkheads are designated by the mainframe number and whether the bulkhead is on the right side or the left side. Each mainframe bulkhead has four panels of connectors.

Figure 3. SuperCluster Cabling Bulkheads

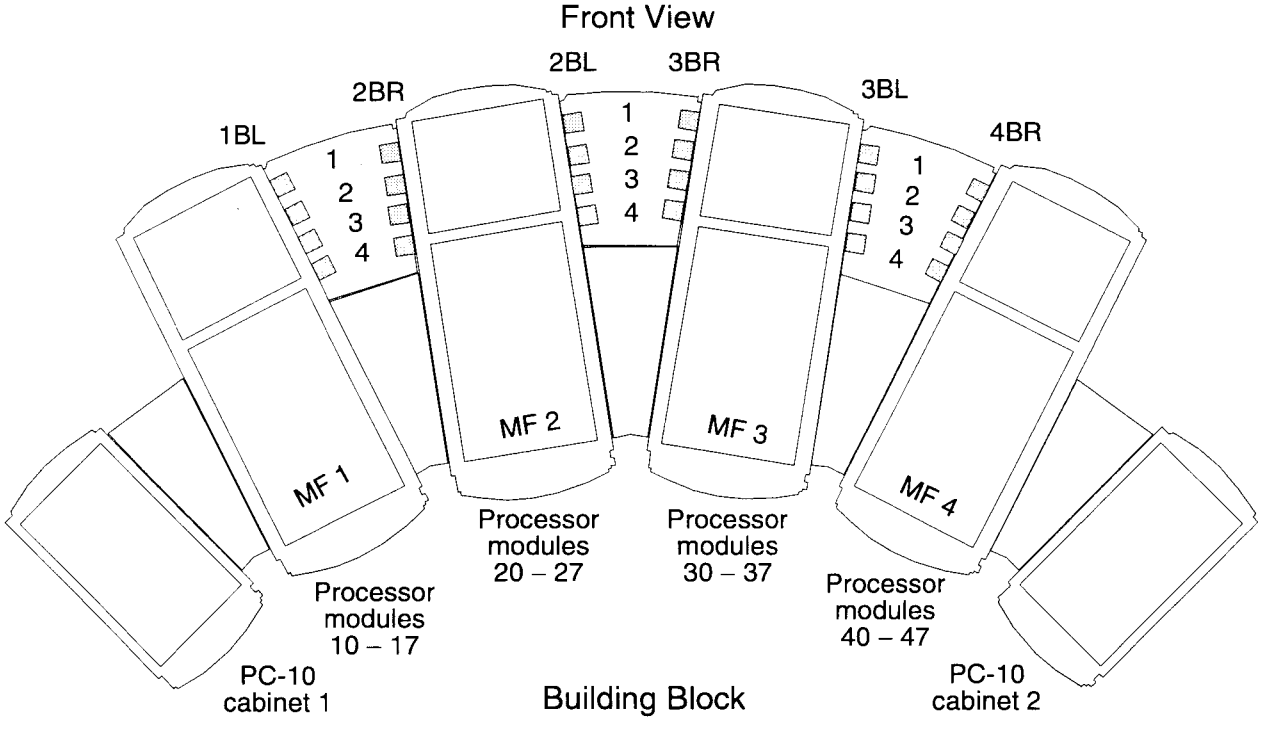

Figure 4. Mainframe Bulkhead Designations

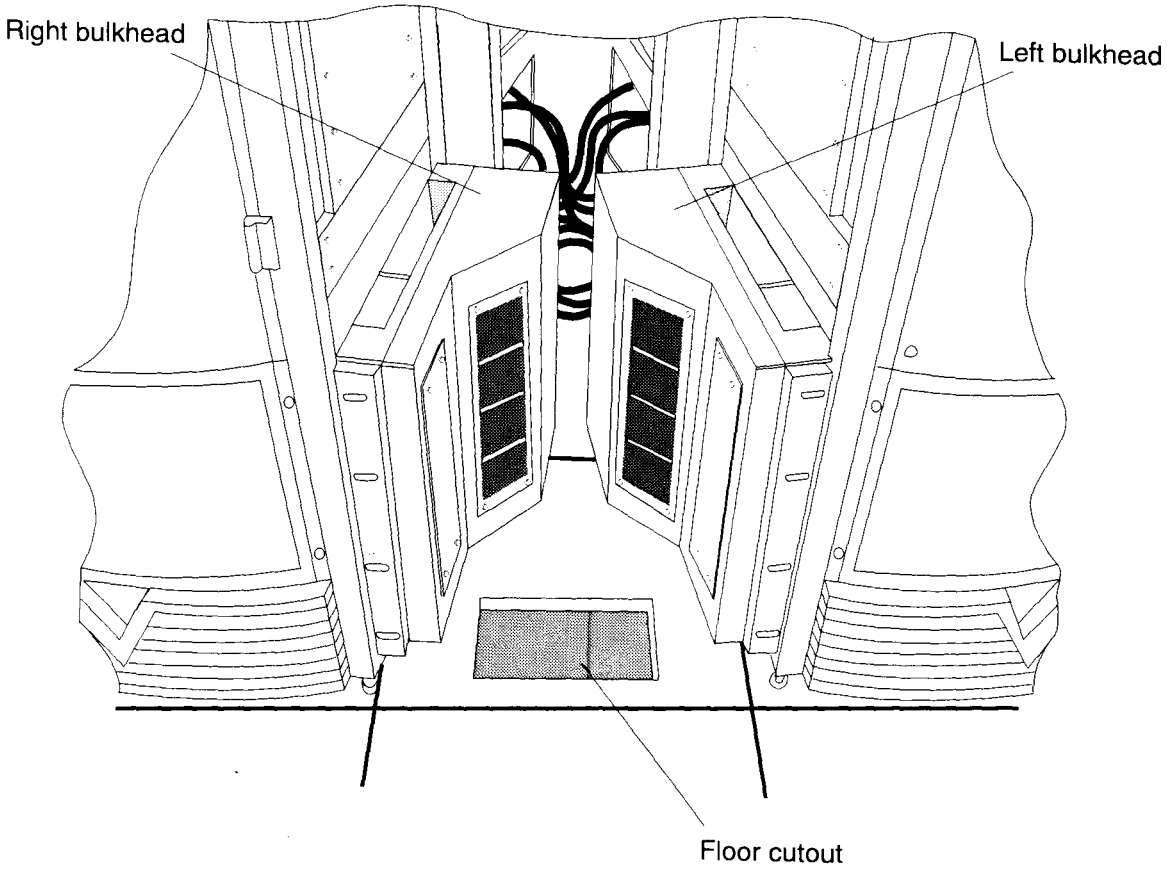

The top panel on the cable bulkhead is used for the warning and control system (RMTCCU) and ALARM connections. The bottom three panels may be used for GigaRing channel connections (refer to Figure 4).

Figure 5. Bulkhead Connections

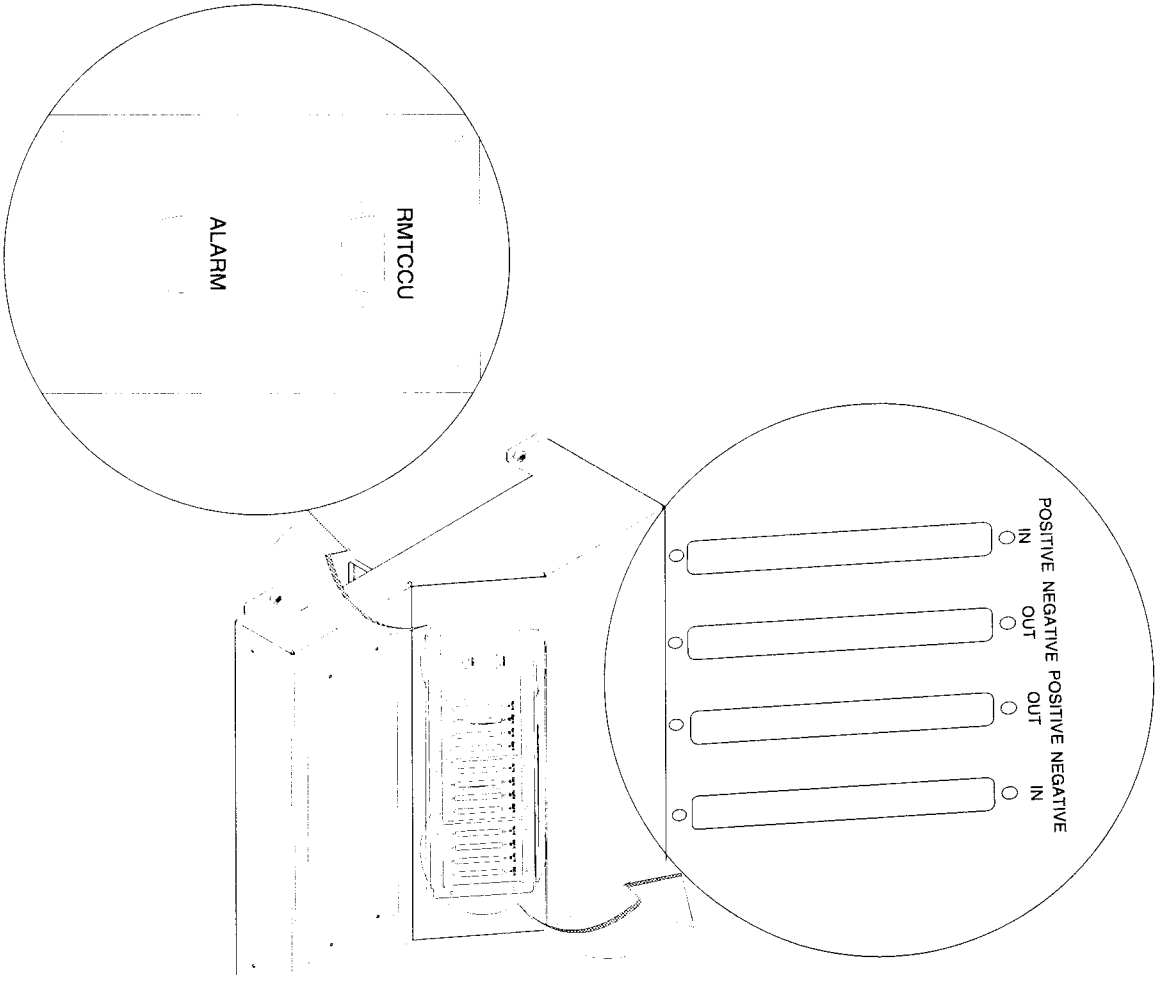

SuperCluster Cabling Configurations

This section shows how the GigaRing channels can be routed for Cray SV1-4 through Cray SV1-32 clusters. Additional I/O rings, which vary according to site-specific needs, are not addressed in this document. Figure 6 shows the Boot, X-ring, and Y-ring connections.

Figure 6. Boot Ring, X-ring, and Y-ring Cabling

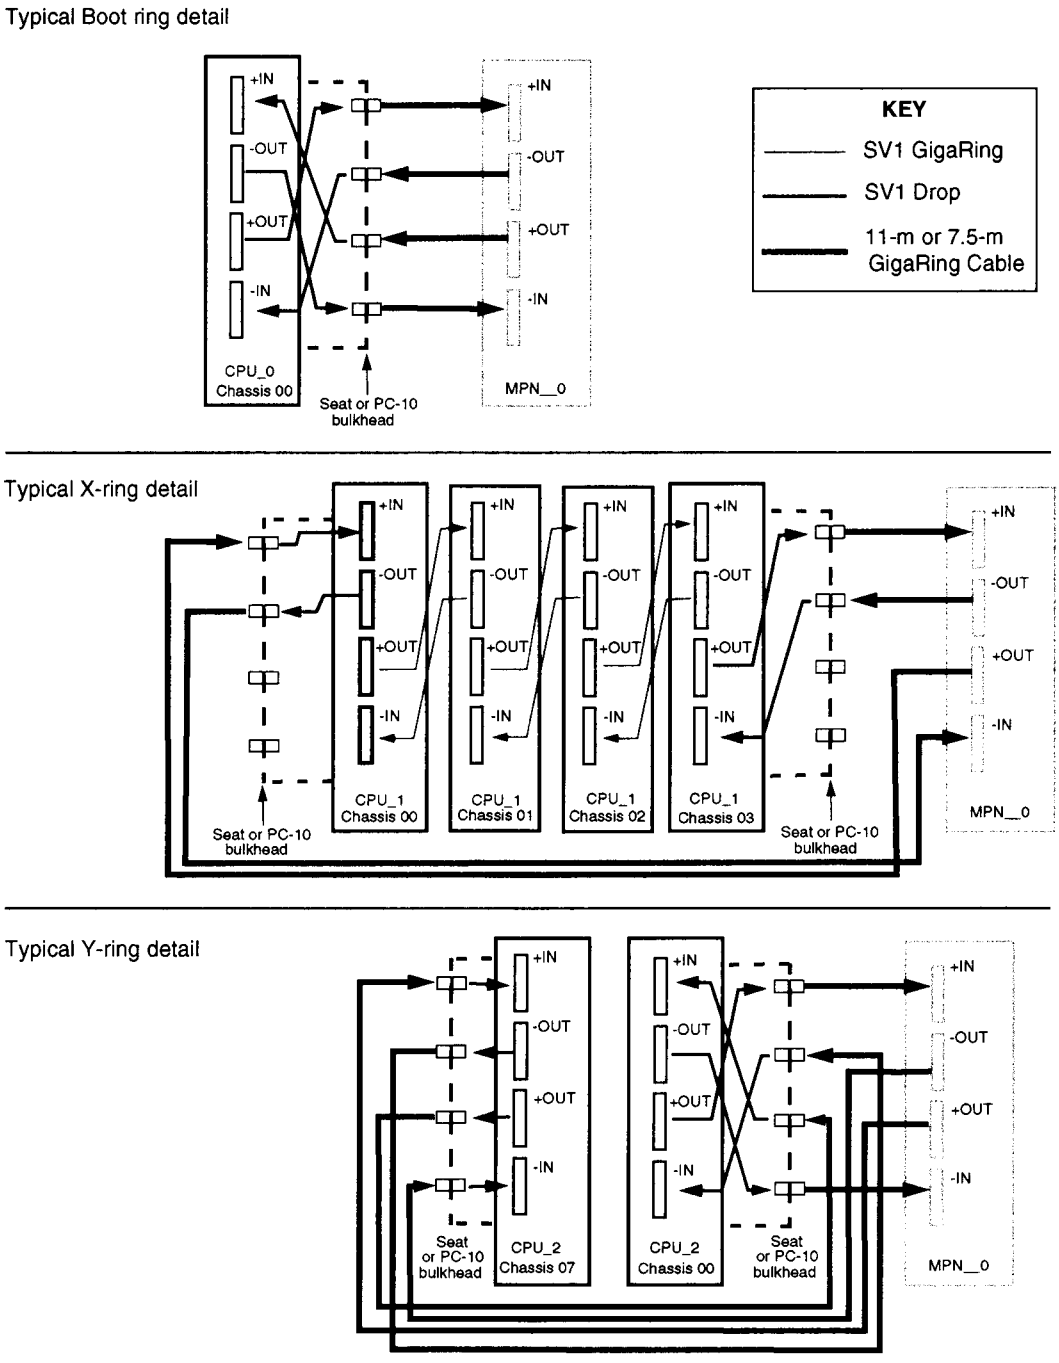

Cluster Connection Matrices

Table 2 through Table 6 list the chassis routing of Y rings for Cray SV1-8 through Cray SV1-32 clusters. The term *chassis* in this discussion refers to all of the mainframe cabinets and PC-10 cabinets on each Y ring. Each Y ring

ends at its starting point, although these connections are not listed in the following tables. For example, in Table 2 the Y 01 ring connection from M 08 returns to M 01 to close the ring.

The chassis, PC-10 cabinet, and mainframe cabinet numbers are listed in the following tables for each cluster configuration. PC-10 cabinets are preceded with a "PC-10" designation; mainframe cabinets are preceded with an "M" designation. More detailed cable routing diagrams begin on page 16.

Table 2. SV1-8 Y-ring Matrix

Y ring	Chassis 1	Chassis 2	Chassis 3
Y 01	M 01	PC-10 02	M 08
Y 02	M 02	PC-10 02	M 07
Y 03	M 03	PC-10 03	M 05
Y 04	M 04	PC-10 03	M 05

Table 3. SV1-12 Y-ring Matrix

Y ring	Chassis 1	Chassis 2	Chassis 3	Chassis 4
Y 01	M 01	M 09	PC-10 02	M 08
Y 02	M 02	M 10	PC-10 02	M 07
Y 03	M 03	M 11	PC-10 03	M 06
Y 04	M 04	M 12	PC-10 03	M 05

Table 4. SV1-16 Y-ring Matrix

Y ring	Chassis 1	Chassis 2	Chassis 3	Chassis 4	Chassis 4
Y 01	M 01	M 09	M 16	PC-10 02	M 08
Y 02	M 02	M 10	M 15	PC-10 02	M 07
Y 03	M 03	M 11	M 14	PC-10 03	M 06
Y 04	M 04	M 12	M 13	PC-10 03	M 05

Table 5. SV1-24 Y-ring Matrix

Y ring	Chassis 1	Chassis 2	Chassis 3	Chassis 4	Chassis 5	Chassis 6	Chassis 7
Y 01	M 01	M 09	M 17	M 24	M 16	PC-10 02	M 08
Y 02	M 02	M 10	M 18	M 23	M 15	PC-10 02	M 07

Table 5. SV1-24 Y-ring Matrix (continued)

Y ring	Chassis 1	Chassis 2	Chassis 3	Chassis 4	Chassis 5	Chassis 6	Chassis 7
Y 03	M 03	M 11	M 19	M 22	M 14	PC-10 03	M 06
Y 04	M 04	M 12	M 20	M 21	M 13	PC-10 03	M 05

Table 6. SV1-32 Y-ring Matrix

Y ring	Chassis 1	Chassis 2	Chassis 3	Chassis 4	Chassis 5	Chassis 6	Chassis 7	Chassis 8	Chassis 9
Y 01	M 01	M 09	M 16	M 17	M 24	M 32	M 25	PC-10 02	M 08
Y 02	M 02	M 10	M 17	M 18	M 23	M 31	M 26	PC-10 02	M 07
Y 03	M 03	M 11	M 18	M 19	M 22	M 30	M 27	PC-10 01	M 06
Y 04	M 04	M 12	M 19	M 20	M 21	M 29	M 28	PC-10 01	M 05

Upgrades to Y Rings

When you upgrade a cluster, you must reroute some Y-ring links. The layout of the ring connection matrix is designed to minimize this effort, but some rerouting of existing Y rings is inevitable. Remove the Y-ring cables for each configuration upgrade as follows:

Cray SV1-8 to Cray SV1-12 M 01 to PC-10 02, M 02 to PC-10 02, M 03 to PC-10 01, M 04 to PC-10 01

Cray SV1-12 to Cray SV1-16 M 09 to PC-10 02, M 10 to PC-10 02, M 11 to PC-10 03, M 12 to PC-10 03

Cray SV1-16 to Cray SV1-24 M 09 to M 16, M 10 to M 15, M 11 to M 14, M 12 to M 13

Cray SV1-24 to Cray SV1-32 M 09 to M 17, M 10 to M 18, M 11 to M 19, M 12 to M 20
M 24 to M 16, M 23 to M 15, M 22 to M 14, M 21 to M 13
M 16 to PC-10 02, M 15 to PC-10 02, M 14 to PC-10 03, M 13 to PC-10 03

GigaRing Cabling Guidelines

Use the following guidelines to determine how to route GigaRing cables on Cray SV1 SuperCluster systems:

- Use CPU slot 0 as the boot slot
- Use CPU slot 1 for X-ring connections.

  Note: When more than 4 CPU slots are used for GigaRing connections, use CPU slots 4 through 7 for easier X-ring cable routing. This is the case with most cluster configurations.

- Use CPU slot 2 for Y-ring connections
- Route the X ring to the even-numbered PC-10 cabinet in the building block.

GigaRing Cabling Diagrams

Figure 7 through Figure 16 show the cabling diagrams for each Cray SV1 SuperCluster configuration. Each line represents two cables, unless otherwise indicated with a number inside the box superimposed on top of the ring segment. This number applies to both elements in the link: the GigaRing cable and the drop cable. Because of limited space, figures and tables in this document use "GR" to represent GigaRing cables.

Table 7. Number and Type of GigaRing Cables for Each Configuration

Cable Type	Cray SV1-4	Cray SV1-8	Cray SV1-12	Cray SV1-16	Cray SV1-24	Cray SV1-32
SV1 Drop Cables P/N 90514700	20	72	108	144	216	288
SV1 4-m GR Internal P/N 90517900	6	12	18	24	36	48
2-m GR Internal P/N 13334602	2	4	6	8	12	16
1-m GR Internal P/N 13334601	8	16	24	32	48	64
7.5-m GR Subfloor P/N 35746902	8	20	28	36	52	76
11-m GR Subfloor P/N 35746900	2	24	34	44	64	76

Figure 7. Cray SV1-4 System Cabling

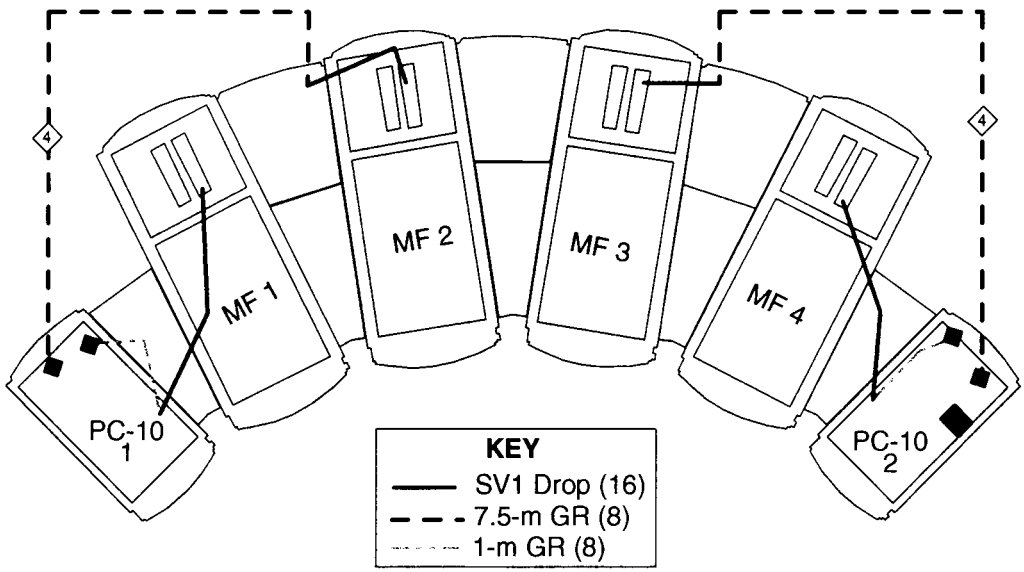

Boot rings

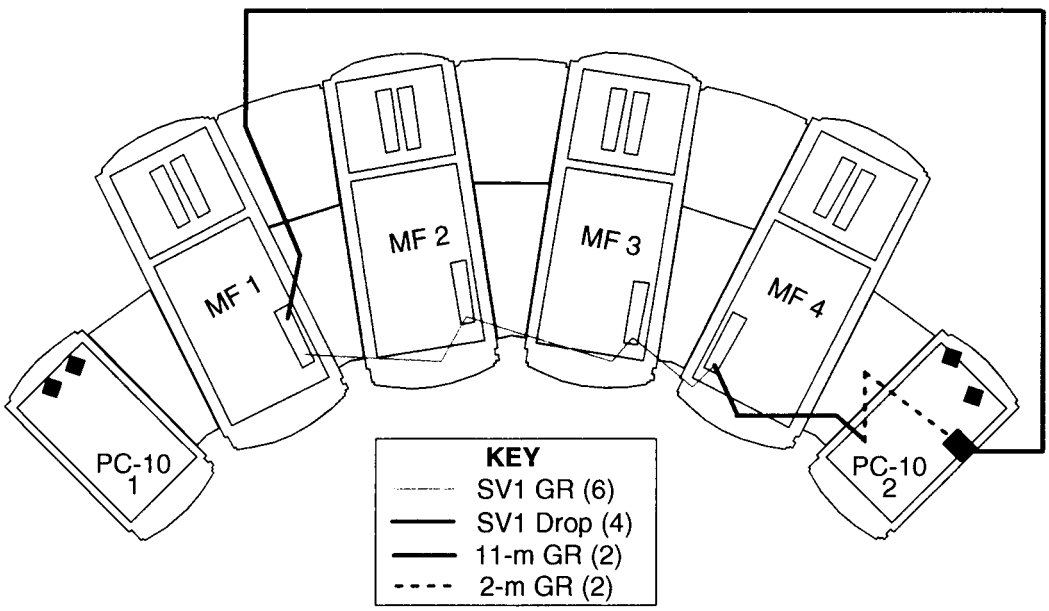

X rings

Cray SV1-4 (128 CPUs)

Figure 8. Cray SV1-8 System Cabling

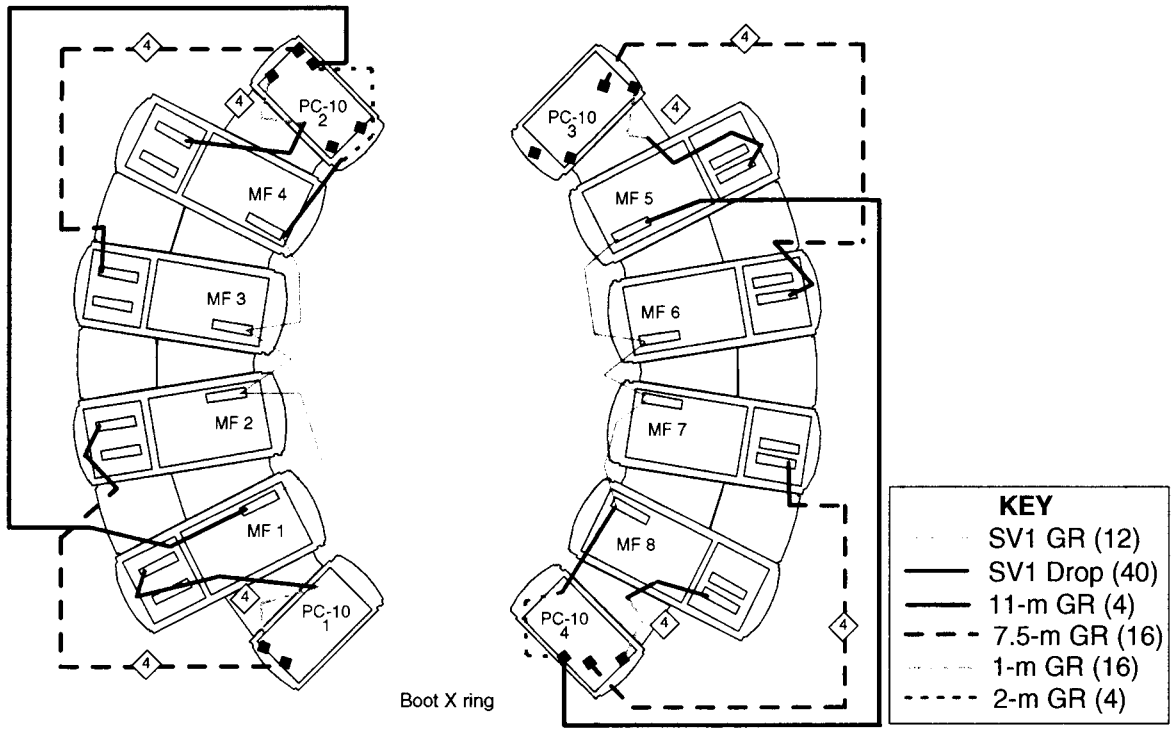

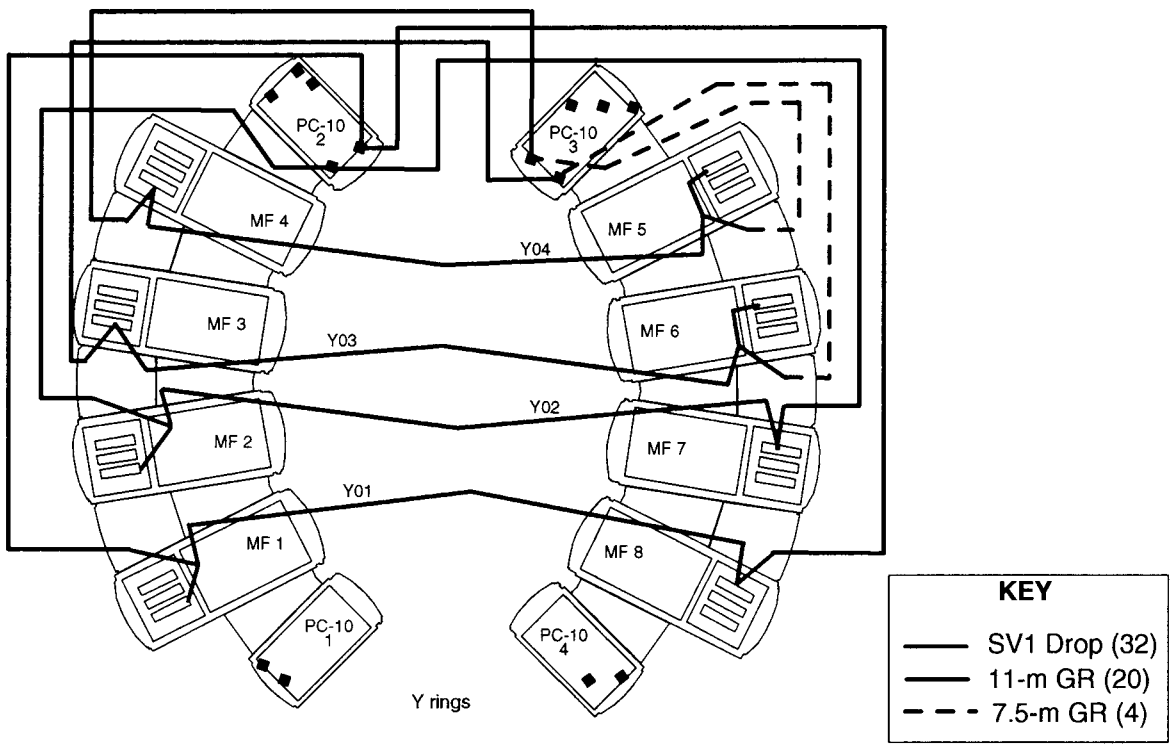

Cray SV1-8 (256 CPUs)

Figure 9. Cray SV1-12 System Cabling

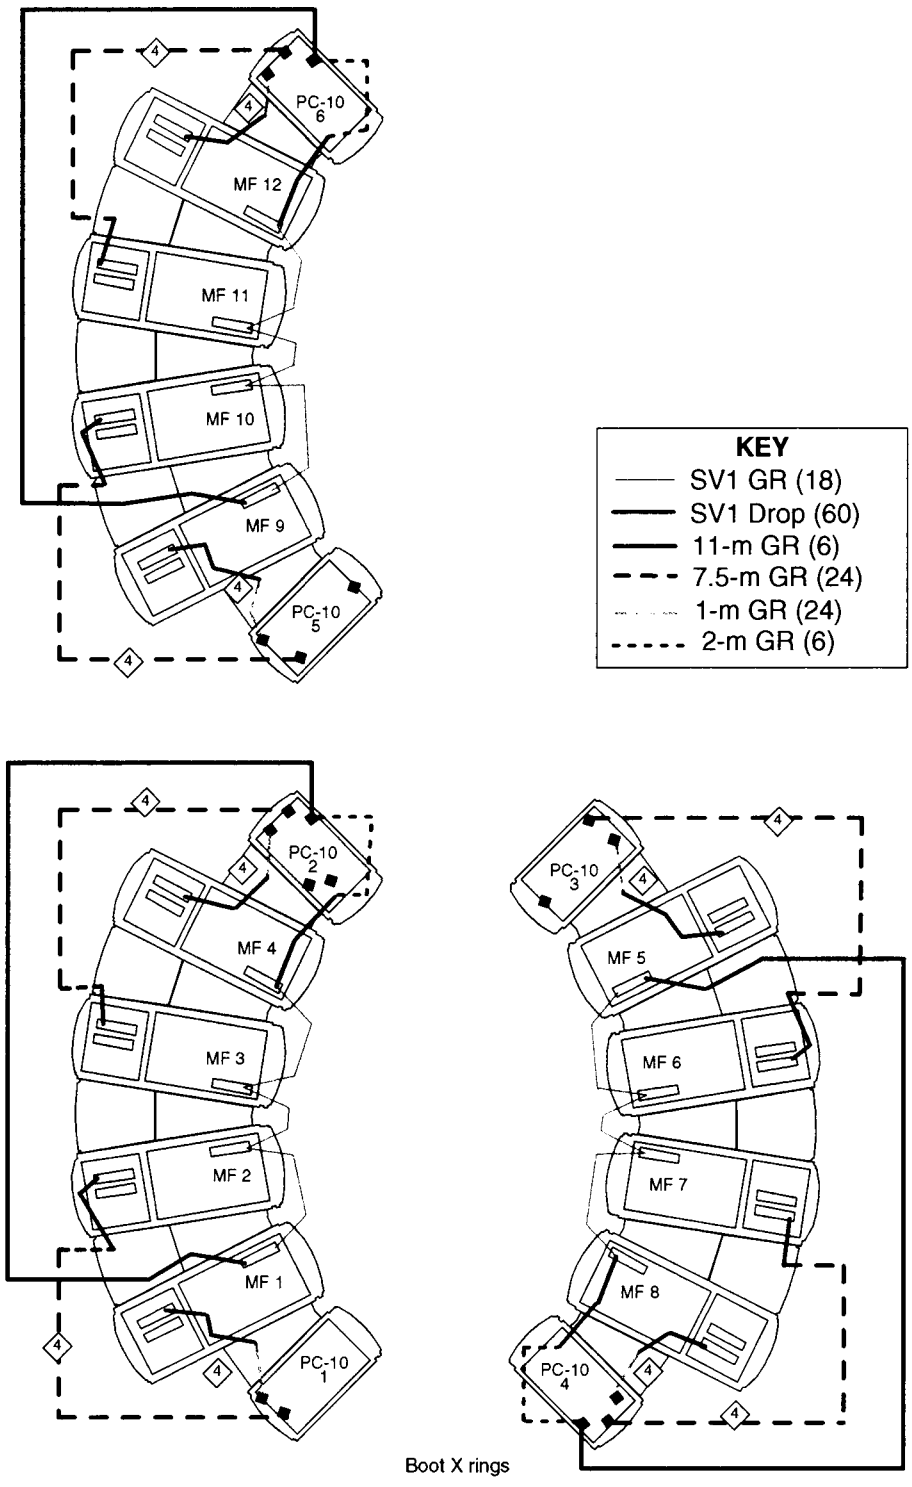

Cray SV1-12 (384 CPUs)

Figure 10. Cray SV1-12 System Cabling (Y Rings)

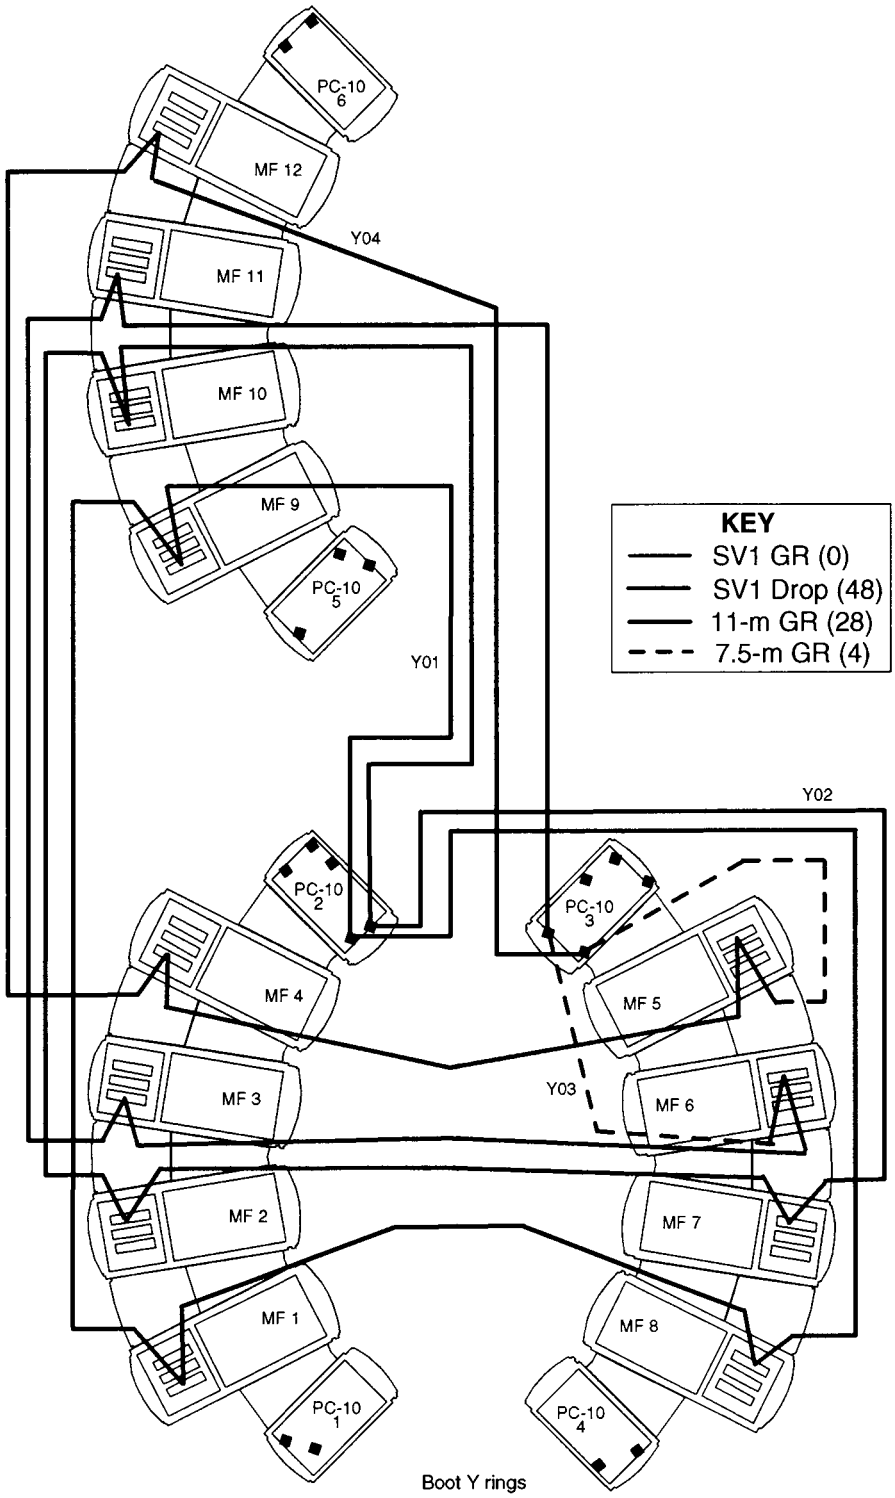

Cray SV1-12 (384 CPUs)

Figure 11. Cray SV1-16 System Cabling

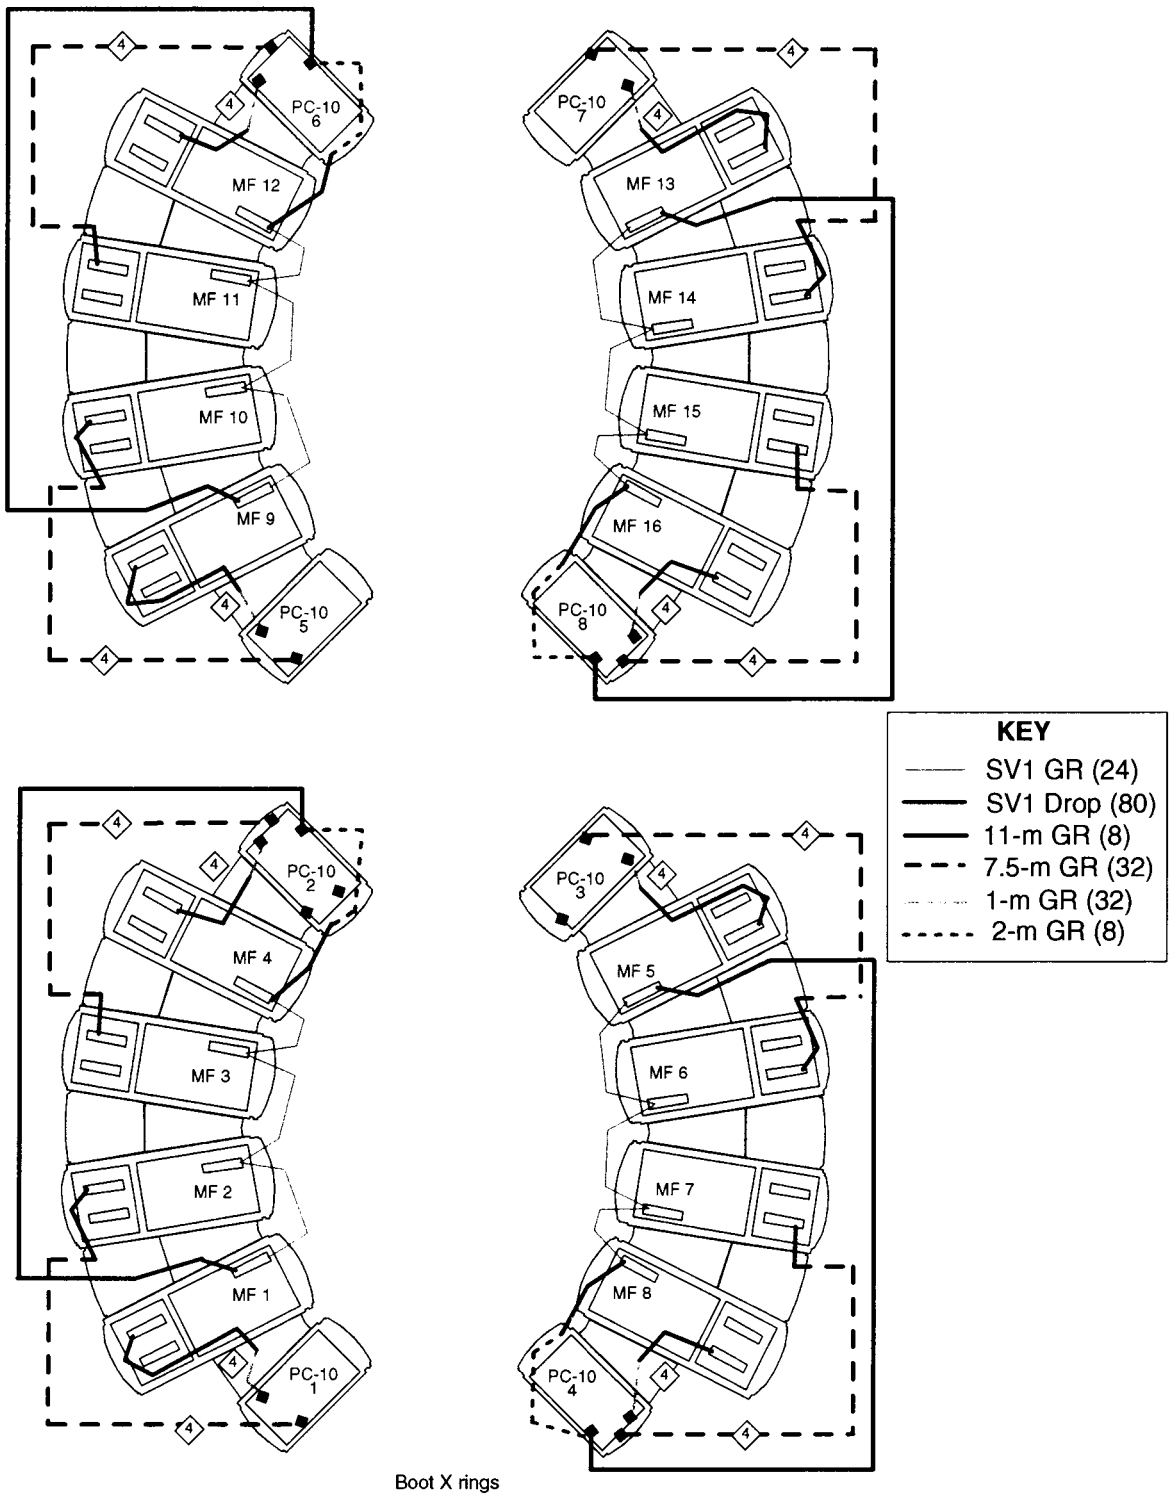

Cray SV1-16 (512 CPUs)

Figure 12. Cray SV1-16 System Cabling (Y Rings)

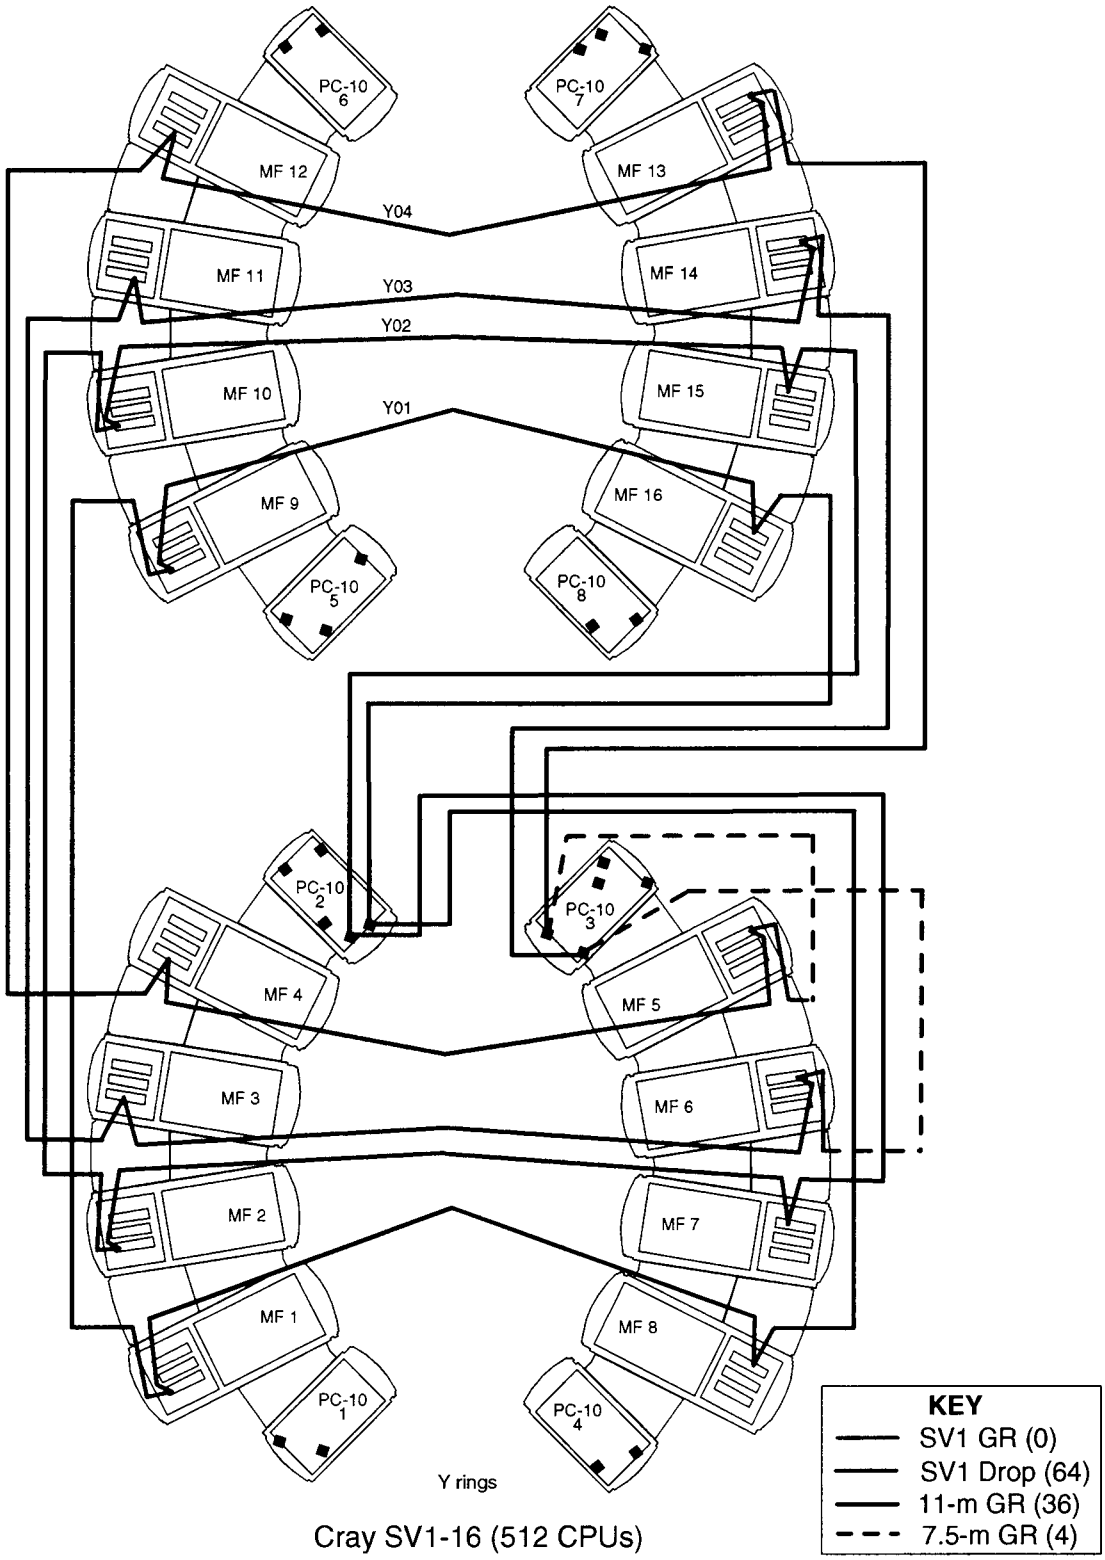

Cray SV1-16 (512 CPUs)

Figure 13. Cray SV1-24 System Cabling

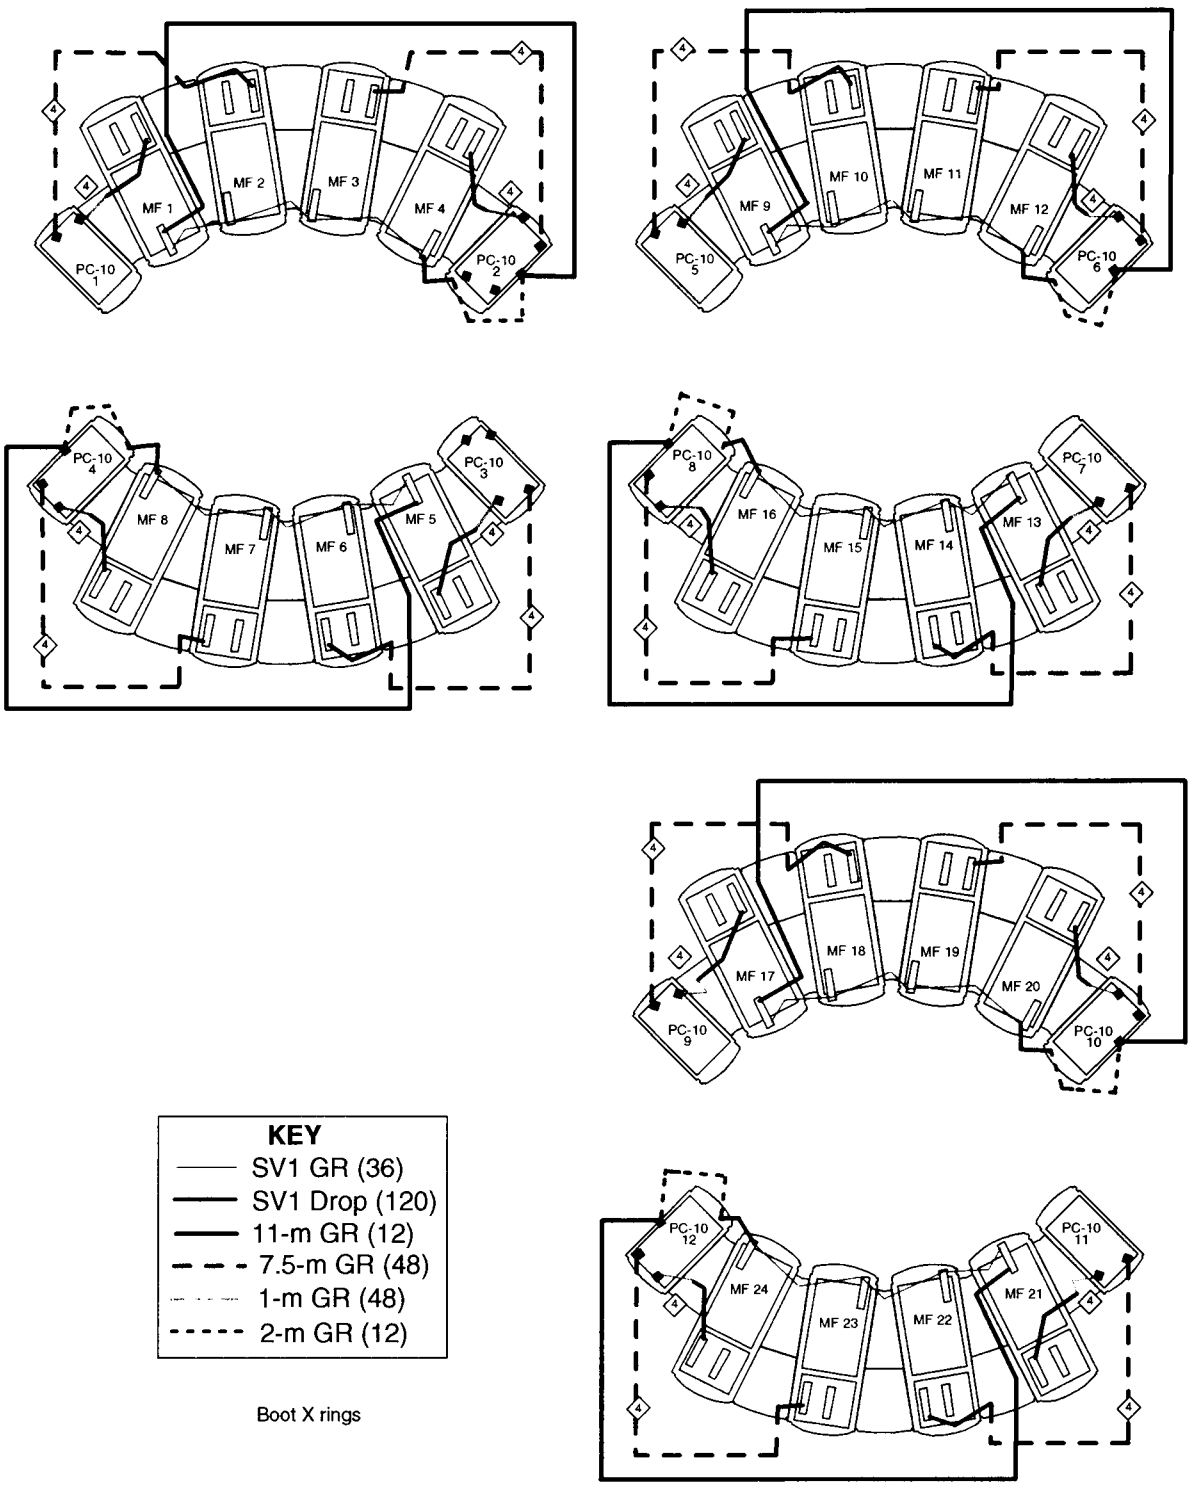

Cray SV1-24 (768 CPUs)

Figure 14. Cray SV1-24 System Cabling (Y Rings)

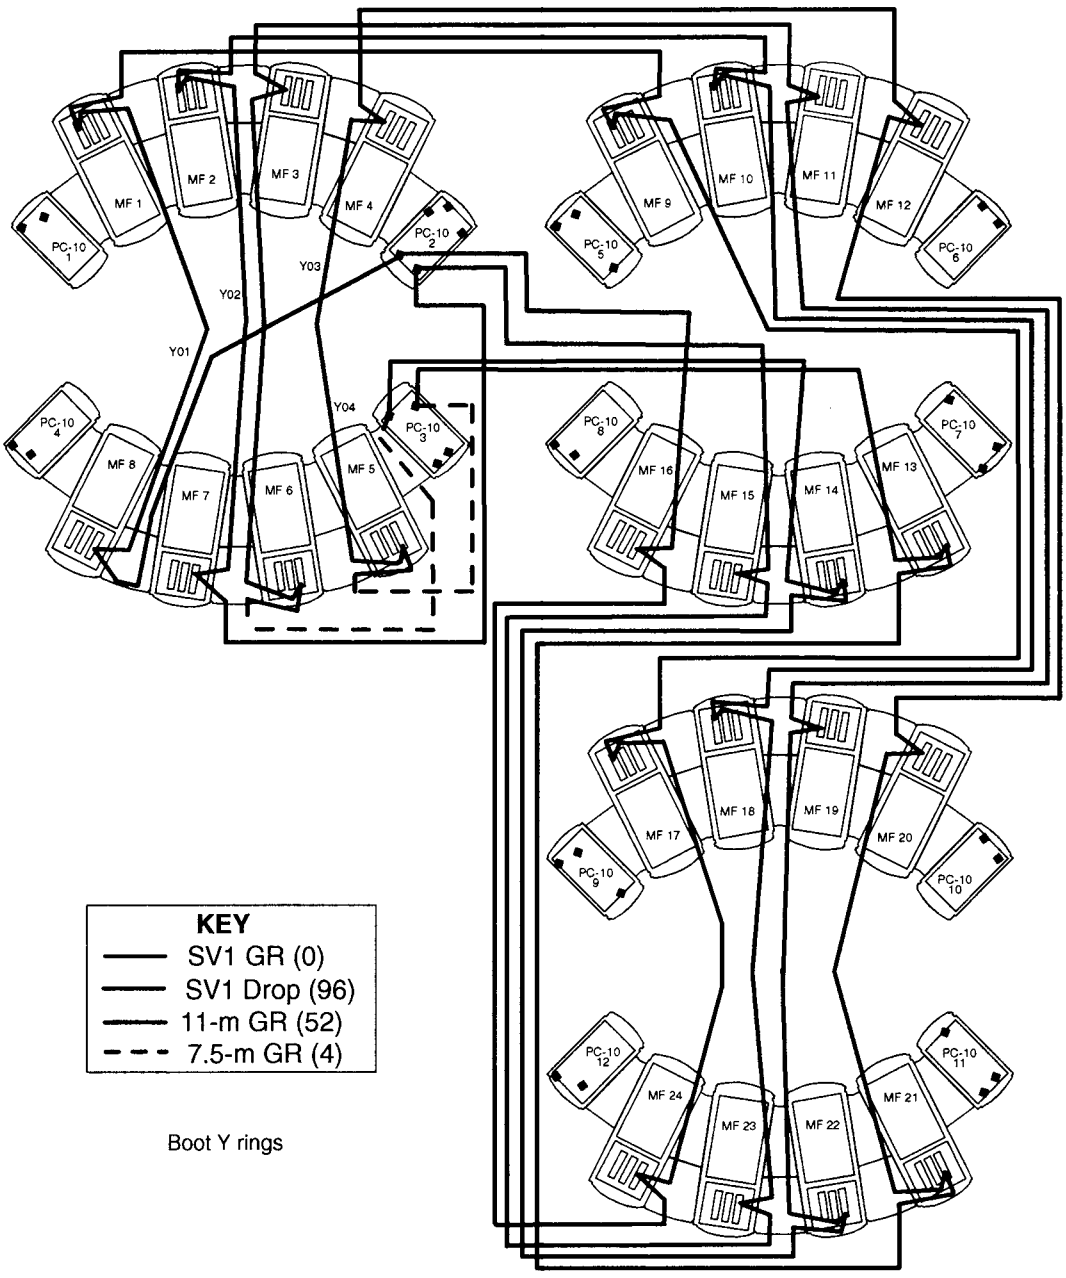

Cray SV1-24 (768 CPUs)

Figure 15. Cray SV1-32 System Cabling

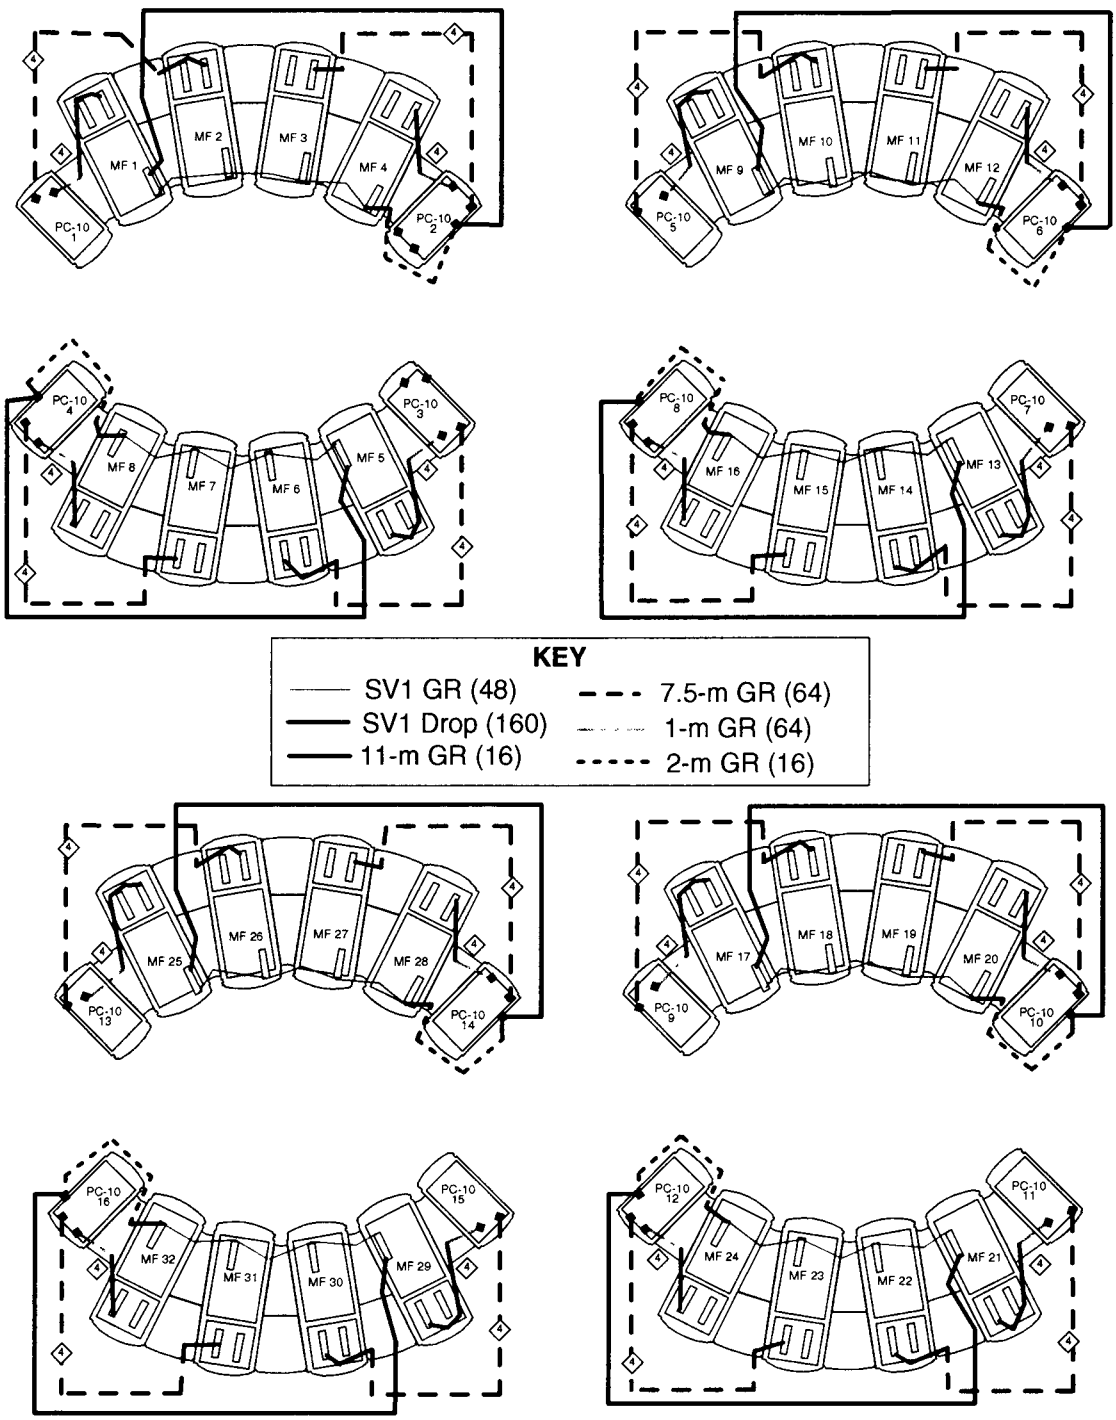

Boot X ring

Cray SV1-32 (1,024 CPUs)

Figure 16. Cray SV1-32 System Cabling (Y Rings)

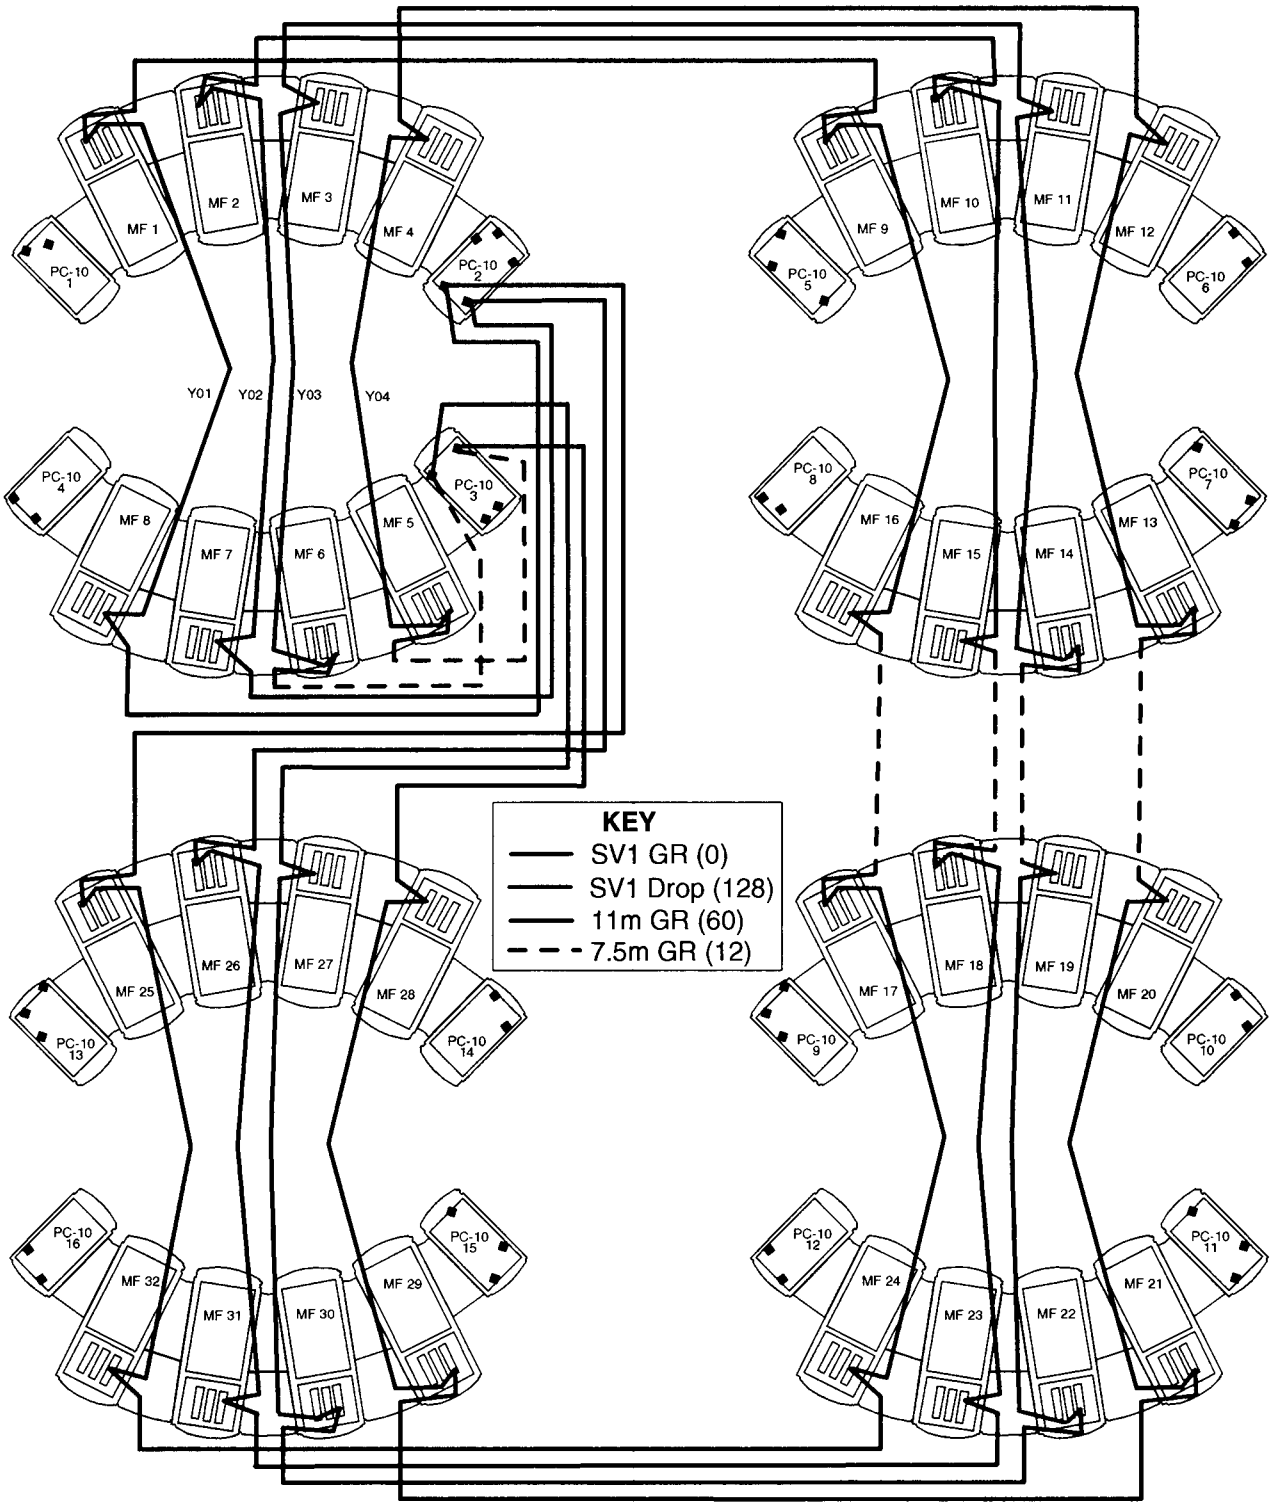

Boot Y rings

Cray SV1-32 (1,024 CPUs)

Reader Comment Form

Title: **Cray SV1™ SuperCluster® Maintenance Guide (Cray SV1 Series)**

Number: **108-0246-001**

Your feedback on this publication will help us provide better documentation in the future. Please take a moment to answer the few questions below.

For what purpose did you primarily use this document?

_____Troubleshooting
_____Tutorial or introduction
_____Reference information
_____Classroom use
_____Other - please explain

_____

Using a scale from 1 (poor) to 10 (excellent), please rate this document on the following criteria and explain your ratings:

_____Accuracy _____
_____Organization _____
_____Readability _____
_____Physical qualities (binding, printing, page layout) _____
_____Amount of diagrams and photos _____
_____Quality of diagrams and photos _____

Completeness (Check one and explain your answer)

_____Too much information _____ Too little information _____Correct amount

_____

_____

_____

You may write additional comments in the space below. Mail your comments to the address below, fax them to us at +1 715 726 4353, or e-mail them to us at *spt@sgi.com.* When possible, please give specific page and paragraph references. We will respond to your comments in writing within 48 hours.

_____

_____

_____

NAME _____
JOB TITLE _____
E-MAIL ADDRESS _____
SITE/LOCATION _____
TELEPHONE _____
DATE _____

[or attach your business card]

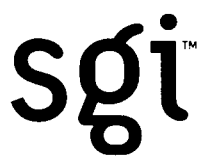

Attn: Service Publications and Training
890 Industrial Boulevard
P.O. Box 4000
Chippewa Falls, WI 54729-0078
USA