StorageTek™
Shared Virtual Array (SVA)
V2X/V2X2 Planning

Part Number : 96218
Revision N

StorageTek™ Shared Virtual Array (SVA)

V2X2

Planning

We welcome your feedback. Please contact the Sun Learning Solutions Feedback System at:

SLSFS@Sun.com

or

Global Learning Solutions
Sun Microsystems, Inc.
One StorageTek Drive
Louisville, CO 80028-3256
USA

Please include the publication name, part number, and edition number in your correspondence if they are available. This expedites our response.

Please Recycle

Adobe PostScript

Table of Contents

List of Figures 11
List of Tables 12
Preface 13
- Notices 13
  - United States FCC Compliance Statement 13
  - Agency Compliance Statement 13
  - CISPR 22 and EN55022 Warning 14
  - Japanese Compliance Statement 14
  - Taiwan Warning Label Statement 14
  - Internal Code License Statement 14
  - Alert Messages 18
  - Mensajes de alerta 18
- Related Documents 19
- Viewing and Printing Web-Based Electronic Documents 21
- History of Changes 21

1 Site Planning 23
- Subsystem Layout 23
- Physical Characteristics 23
- System Cables 25
  - Measuring Cable Lengths 26
- Environmental Specifications 27
  - Floor Loading 27
  - Raised Floors 29
- Power 30
  - Electrical Power 30
  - Electrical Connectors 31
  - Power Source Specifications 31
    - Power Redundancy Levels 31
    - Harmonic Content 32
  - Power Measurements 32
  - Grounding 35
  - Static Electricity 35
  - Power Distribution System 36
    - Power Load 36
    - Power Panel Feeders 36
    - Branch Circuits 36
- Electrical Problem Areas and Issues 36

2 Configuration Planning 39
- Net Capacity Load And Free Space 39
  - Net Capacity Load 40
  - Free Space 40
    - Uncollected Free Space 40
    - Collected Free Space 40
    - How Free-Space Collection Works 41

Storage Capacity Management 41
Selecting a Subsystem Configuration 43
Functional Considerations 44
Functional Configuration 44
Missing-Interrupt Handler (MIH) 45
Disk Array Definition 45
Array Partitions 46
Spares Partition 46
Production Partition 46
Unavailable Partition 46
HCD/IOCP Definition 46
SVA Device Type Support 46
Maximum Functional Devices and Virtual Cylinders 47
Device Support 47
FlexVolume Support (FVS) 47
Restrictions 48
Options 48
SSID 48
Definition 48
Assignment 48
Changing SSIDs 49
Large LUN 49
Physical Capacity Constraints 51
PPRC Considerations 51
Virtual Capacity Cylinder Gauge 52
53
Constraints Imposed by the User Interface 53
3 Channel Configuration 55
Configuration in the Mainframe Environment 55
ESCON Path Definitions 55
Maximum Configuration Values 55
Channel Paths And Mainframe Connectivity 56
Channel Configuration 57
Channel Attachment Options 57
Channel Attachment Considerations 57
Host Path Connection Reduction 59
ESCON Channels 60
4 The SVA In An Open Systems Environment 61
Fibre Channels 61
Fibre Channel Performance And Configuration Options 61
Fabric or PLDA Operation 61
Requested Loop Id. (RLid) 62
Improving Performance 65
Attaching More LUNs Than Normally Allowed To A HBA 65
HP Volume Addressing Method 65
HP Alternate Method Using A Brocade Switch 69
Supported Host Bus Adapters 70
Supported Cluster Configurations 71

Windows 2000 SVA Performance and Block Size 71
Setting I/O Timers 72
HP-UX 72
RS6000 - AIX 72
SUN 73
Windows NT/Windows 2000 74
Linux - Redhat Versions 6.2 through 7.2 74
Netware 5.1 & 6.0 75
Installing Windows (Win2K or NT) Onto An SVA 76
Host Recognition Of New Devices 76
Modifying OS/390 VTOCs for Interoperative Use 77
Record Format Assist Feature (RFA) 77
Read Data Transfers 79
Write Data Transfers 80
Positive Side Effects 80
Track Utilization/Virtual Capacity 80
Intra Track Overhead 80
Cache Bandwidth 80
Compression Ratio 81
Record Format Assist in the Open Systems Environment 81
5 The SVA in a VM Environment 83
Installing the SVA for VM 83
Minimum VM Program Levels 83
SYSGEN Considerations 83
Device Initialization 84
VM Cache Operation 84
Exploiting the V2X2 With SVAA/VM 85
VM Performance Tools 85
Storage Management In VM 86
Special Utilities - Device Support Facilities 86
Defining and Preparing a Functional Volume for VM 86
Data Recovery 87
Data Migration Planning 87
Migration Strategies 88
Mini-disk Migration Strategies 89
Migration Using DIRMAINT CMDISK 90
Migration Using COPYFILE 90
Migration Using DDR 91
Migration Using SPXTAPE 91
Examine Data Access Characteristics 91
Preparing for Data Migration 92
Placing mini-disks On an SVA Subsystem 93
IPL Volumes 94
Common CMS Mini-Disks Linked by Other VMs 94
Page/Spooling Mini-Disks 94
VM NCL Management 94
TDisk Management 94

6 The SVA in a z/OS Environment 97
Installing the SVA 97
Minimum Program Levels 97
Mainframe Environment 97
Configuration Considerations 98
Device Initialization 99
Cache Operation 99
Dual Copy 99
DASD Fast Write And Cache Fast Write 99
IDCAMS 99
LISTDATA 100
Subsystem Counters Report 100
Remote Access Code 100
Pinned Track To Data Set Cross Reference Report 100
Device Status Report 100
SETCACHE 100
Exploiting the SVA With SVAA 101
z/OS Tools 101
Storage Management 102
Special Utilities - Device Support Facilities 102
Data Recovery 103
Data Migration Planning 103
Migration Strategies 103
Examine Data Access Characteristics 104
Preparing for the Movement of Data 105
Placing Data Sets on an SVA Subsystem 105
Space Management 107
Volume De-fragmentation 108
Data Archival Strategies 109
DFSMS, CACHE Console Command 109
Data Class 110
Storage Class 110
Management Class 112
Storage Group 113
Volume Conversion 114
Backups and Disaster Recovery 114
A Device Map And Host Addresses 115
B Alternative Restoration 119
C IOCP/HCD 121
Description 121
Physical Configuration and HCD 122
Clusters And Paths Selections 122
Defining the HCD/IOCP 123
RESOURCE Statement 123
CHPID Statement 124
CNTLUNIT Statement 125
IO DEVICE Statement 126

Device Parameters And Features 128
ESCON Director Definition 128
Esoteric 129
SVA with FICON 131
Channel device addressing support 131
FICON (FCV mode) configuration rule 131
I/O Data Transfers 131
Aggregation of ESCON Channel Paths 132
Example IOCP report: 133
Examples: 135
Example 1 - SVA With ESCON Directors And 8-Path VCU 135
IOCP Statements for Example 1 - 8 Paths To All Devices 138
HCD Example 138
Channel path definition for CHPIDs 40-43: 138
Channel path definition for CHPIDs 44-47: 139
Director entry port definitions: 139
LPAR access list definition: 140
Logical control unit definition: 140
Device definition for host addresses 6000-6FFF: 141
Example 2 - SVA With 2 Processors, ESCON Directors, And 4-Path VCU 141
IOCP Statements 145
HCD Example 146
Channel path definition for CHPID 58: 147
Director (switch) entry port definition: 147
LPAR access list definition: 147
Channel path attribute listing: 148
Logical control unit definition: 149
Device definition for host addresses 1000-10FF: 151
Example 3 – SVA With ESCON Directors and 4-Path VCU 152
IOCP Statements 154
HCD for Example 3 - 16 channels (4-path VCU) 155
Channel path definition for CHPID 58: 155
Director (switch) definition: 156
LPAR access list definition: 156
Channel path attribute listing: 156
Logical control unit definition: 158
Device definition for host addresses 1000-1FFF: 160
Example 4 – SVA With Direct Host Connectivity 160
IOCP Statements 163
HCD for Example 4 163
LPAR access list definition: 164
Logical control unit definition 166
Device definition for host addresses 166
D Device Characteristics 169
3390 Devices 169
3380 Devices 170
Combination Examples 171

E Access Authorization Address Range 173
F ICE/ICF Card Combinations 175
G System Block Diagram 179

List of Figures

Figure 1 Shared Virtual Array Physical Planning Template 25
Figure 2 Shared Virtual Array Weight Distribution and Leveler Locations 28
Figure 3 SVA Axis Reference 28
Figure 4 Power Cable Ends 31
Figure 5 Transient Grounding Plate 37
Figure 6 An V2X2 At 10% Collected Free Space 42
Figure 7 Subsystem Configuration and Status Screen with Cylinder Gauge 52
Figure 8 A Fully Populated SVA with ESCON Connectivity Only 59
Figure 9 A SVA with Two ICF Cards Used For A Fibre Connection 59
Figure 10 HP-UX Fibre Channel Attach For More Than Eight LUNs 70
Figure 11 The Effects Of RFA 78
Figure 12 Channel and ICE Card Configuration 137
Figure 13 Channel And ICE Card Configuration For Example 2 143
Figure 14 Channel And ICE Card Configuration For Example 2 144
Figure 15 Channel and ICE Card Configurations For Example 3 153
Figure 16 Channel and ICE Card Configurations For Example 4 162
Figure 17 System Block Diagram 180
Figure 18 2 GB/sec. SVA Back End 181

List of Tables

Table 1 Operating Environment Requirements 27
Table 2 Superimposed Floor Load 29
Table 3 Horizontal Force Chart 30
Table 4 Dual Line Cord Operation (36 GB Drives) 33
Table 5 Dual Line Cord Operation (73 GB Drives) 33
Table 6 Single Line Cord Operation (36 GB Drives) 34
Table 7 Single Line Cord Operation (73 GB Drives) 34
Table 8 V2X/V2X2 Mixed Drive Size Variations Supported 39
Table 9 An V2X2 At 10% Collected Free Space 42
Table 10 FlexVolume Cylinder Options per Device Type 48
Table 11 SCSI Track Information 50
Table 12 Examples of LLUN Sizes Using 3390 Based Devices 51
Table 13 Requested Loop ID 62
Table 14 Fibre Channel Arbitrated Loop Physical Address Values 63
Table 15 SVA LUN to HP-UX Target Mapping 66
Table 16 Validated Cluster Configurations 71
Table 17 Components and Data Compression 78
Table 18 Data Utilization And Block Size 82
Table 19 Migration Process Tools 90
Table 20 Minimum Version And Release Levels of Software 97
Table 21 ICKDSF Commands Supported by SVA in z/OS 102
Table 22 Dynamic Data Set Caching Classifications 111
Table 23 4096 Device Support Mapping 115
Table 24 Non-sequential Addressing Cross-reference Example 117
Table 25 IOCP Statements 138
Table 26 IOCP Statements 145
Table 27 IOCP Statements 154
Table 28 IOCP Statements 163
Table 29 3390 Device Geometry and Characteristics 169
Table 30 3380 Device Geometry and Characteristics 170
Table 31 3390 Device Combinations 171
Table 32 3380 Access Authorization Address Range 173
Table 33 3390 Access Authorization Address Range 174
Table 34 ICE2, ICE3 and ICF Normal Locations 175

Preface

Notices

Please read the following compliance and warning statements for this product.

Caution: Potential equipment damage: Cables that connect peripherals must be shielded and grounded; refer to cable descriptions in the instruction manuals. Operation of this equipment with cables that are not shielded and not correctly grounded might result in interference to radio and TV reception.

Changes or modifications to this equipment that are not expressly approved in advance by Sun Microsystems Inc. will void the warranty. In addition, changes or modifications to this equipment might cause it to create harmful interference.

United States FCC Compliance Statement

The following compliance statement pertains to Federal Communications Commission Rules 47 CFR 15.105:

Note: This equipment has been tested and found to comply with the limits for a Class A digital device pursuant to part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference in which case the user will be required to correct the interference at his or her own expense.

Agency Compliance Statement

The SVA complies with the following agencies:

UL–Recognized Component by Underwriters Laboratories Inc. to Standard UL 60950, Information Technology Equipment.

CE–Mark to show compliance to European Union Directives (European Union: Safety & EMC).

CISPR 22 and EN55022 Warning

This is a Class A product. In a domestic environment this product may cause radio interference in which case the user may be required to take adequate measures.

Japanese Compliance Statement

The following compliance statement in Japanese pertains to VCCI EMI regulations:

この装置は、情報処理装置等電波障害自主規制協議会（ＶＣＣＩ）の基準
に基づくクラスＡ情報技術装置です。この装置を家庭環境で使用すると電波
妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ず
るよう要求されることがあります。

English translation: This is a Class A product based on the Technical Requirement of the Voluntary Control Council for Interference by Information Technology (VCCI). In a domestic environment, this product may cause radio interference, in which case the user may be required to take corrective actions.

Taiwan Warning Label Statement

The following warning label statement (in Kanji) pertains to BSMI regulations in Taiwan, R.O.C.:

警告使用者：這是甲類的資訊產品，在居住的
環境中使用時，可能會造成射頻干擾，在這種
情況下，使用者會被要求採取某些適當的對策

English translation: This is a Class A product. In a domestic environment, this product may cause radio interference, in which case, the user may be required to take adequate measures.

Internal Code License Statement

The following is the Internal Code License Agreement from Sun Microsystems Inc.:

NOTICE

INTERNAL CODE LICENSE

PLEASE READ THIS NOTICE CAREFULLY BEFORE INSTALLING AND OPERATING THIS EQUIPMENT. THIS NOTICE IS A LEGAL AGREEMENT BETWEEN YOU (EITHER AN INDIVIDUAL OR

ENTITY), THE END USER, AND STORAGE TECHNOLOGY CORPORATION (“Sun Microsystems”), THE MANUFACTURER OF THE EQUIPMENT. BY OPENING THE PACKAGE AND ACCEPTING AND USING ANY UNIT OF EQUIPMENT DESCRIBED IN THIS DOCUMENT, YOU AGREE TO BECOME BOUND BY THE TERMS OF THIS AGREEMENT. IF YOU DO NOT AGREE WITH THE TERMS OF THIS AGREEMENT, DO **NOT** OPEN THE PACKAGE AND USE THE EQUIPMENT. IF YOU DO NOT HAVE THE AUTHORITY TO BIND YOUR COMPANY, DO **NOT** OPEN THE PACKAGE AND USE THE EQUIPMENT. IF YOU HAVE ANY QUESTIONS, CONTACT THE AUTHORIZED SUN MICROSYSTEMS INC. DISTRIBUTOR OR RESELLER FROM WHOM YOU ACQUIRED THIS EQUIPMENT. IF THE EQUIPMENT WAS OBTAINED BY YOU DIRECTLY FROM SUN MICROSYSTEMS INC., CONTACT YOUR SUN MICROSYSTEMS INC. REPRESENTATIVE.

1. **Definitions:** The following terms are defined as follows:

   A. “Derivative works” are defined as works based upon one or more preexisting works, such as a translation or a musical arrangement, or any other form in which a work may be recast, transformed, or adapted. A work consisting of editorial revision, annotations, elaboration, or other modifications which, as a whole, represent an original work of authorship, is a Derivative work.

   B. “Internal Code” is Microcode that (i) is an integral part of Equipment, (ii) is required by such Equipment to perform its data storage and retrieval functions, and (iii) executes below the user interface of such Equipment. Internal code does not include other Microcode or software, including data files, which may reside or execute in or be used by or in connection with such Equipment, including, without limitation, Maintenance Code.

   C. “Maintenance Code” is defined as Microcode and other software, including data files, which may reside or execute in or be used by or in connection with Equipment, and which detects, records, displays, and/or analyzes malfunctions in the Equipment.

   D. “Microcode” is defined as a set of instructions (software) that is either imbedded into or is to be loaded into the Equipment and executes below the external user interface of such Equipment. Microcode includes both Internal Code and Maintenance Code, and may be in magnetic or other storage media, integrated circuitry, or other media.

2. The Equipment you have acquired by purchase or lease is manufactured by or for Sun Microsystems Inc. and contains Microcode. By accepting and operating this Equipment, you acknowledge that Sun Microsystems Inc. or its licensor(s) retain(s) ownership of all Microcode, as well as all copies thereof, that may execute in or be used in the operation or servicing of the Equipment and that such Microcode is copyrighted by Sun Microsystems Inc. or its licensor(s).

3. Sun Microsystems Inc. hereby grants you, the end user of the Equipment, a personal, nontransferable (except as permitted in the transfer terms below), nonexclusive license to use each copy of the Internal Code (or any replacement provided by Sun Microsystems Inc. or your authorized Sun Microsystems Inc. distributor or reseller) which license authorizes you, the end user, to execute the Internal Code solely to enable the specific unit of Equipment for which the copy of Internal Code is provided to perform its data storage and retrieval functions in accordance with Sun Microsystems Inc.'s (or its licensor's) official published specifications.

4. Your license is limited to the use of the Internal Code as set forth. You may not use the Internal Code for any other purpose. You may not, for example, do any of the following:

   (i) access, copy, display, print, adapt, alter, modify, patch, prepare Derivative works of, transfer, or distribute (electronically or otherwise) or otherwise use the Internal Code;

   (ii) reverse assemble, decode, translate, decompile, or otherwise reverse engineer the Internal Code (except as decompilation may be expressly permitted under applicable European law solely for the purpose of gaining information that will allow interoperability when such information is not otherwise readily available); or

   (iii) sublicense, assign, or lease the Internal Code or permit another person to use such Internal Code, or any copy of it.

5. Nothing in the license set forth above or in this entire Notice shall convey, in any manner, to you any license to or title to or other right to use any Maintenance code, or any copy of such Maintenance Code. Maintenance Code and Sun Microsystems Inc.'s service

tools and manuals may be kept at your premises, or they may be supplied with a unit of Equipment sent to you and/or included on the same media as Internal Code, but they are to be used only by Sun Microsystems Inc.'s customer service personnel or those of an entity licensed by Sun Microsystems Inc., all rights in and to such Maintenance Code, service tools and manuals being reserved by Sun Microsystems Inc. or its licensors. You agree that you shall not use or attempt to use the Maintenance Code or permit any other third party to use and access such Maintenance Code.

You, the end user, agree to take all appropriate steps to ensure that all of your obligations set forth in this Notice are extended to any third party having access to the Equipmen

6. You may transfer possession of the Internal Code to another party only with the transfer of the Equipment on which its use is authorized, and your license to use the Internal Code is discontinued when you are no longer an owner or a rightful possessor of the Equipment. You must give such transferee all copies of the Internal Code for the transferred Equipment that are in your possession, along with a copy of all provisions of this Notice.

   Any such transfer by you is automatically (without further action on the part of either party) expressly subject to all the terms and conditions of this Notice passing in full to the party to whom such Equipment is transferred, and such transferee accepts the provisions of this license by initial use of the Internal Code. You cannot pass to the transferee of the Equipment any greater rights than granted under this Notice, and shall hold Sun Microsystems Inc. harmless from any claim to the contrary by your transferee or its successors or assigns. In addition, the terms and conditions of this Notice apply to any copies of Internal Code now in your possession or use or which you hereafter acquire from either Sun Microsystems Inc. or another party.

7. You acknowledge that copies of both Internal Code and Maintenance Code may be installed on the Equipment before shipment or included with the Equipment and other material shipped to you, all for the convenience of Sun Microsystems Inc.'s service personnel or service providers licensed by Sun Microsystems Inc., and that during the warranty period, if any, associated with the Equipment, and during periods in which the Equipment is covered under a maintenance contract with Sun

Microsystems Inc. or service providers licensed by Sun Microsystems Inc., both Internal Code and Maintenance Code may reside and be executed in or used in connection with such Equipment, and you agree that no rights to Maintenance Code are conferred upon you by such facts.

Sun Microsystems Inc. or the licensed service provider may keep Maintenance Code and service tools and manuals on your premises but they are to be used only by Sun Microsystems Inc.'s customer service personnel or those of service providers licensed by Sun Microsystems Inc.. You further agree that upon
(i) any termination of such warranty period or maintenance contract period; or
(ii) transfer of possession of the Equipment to another party, Sun Microsystems Inc. and its authorized service providers shall have the right with respect to the affected Equipment to remove all service tools and manuals and to remove or disable all Maintenance Code and/or replace Microcode which includes both Internal Code and Maintenance Code with Microcode that consists only of Internal Code.

Alert Messages

Alert messages call your attention to information that is especially important or that has a unique relationship to the main text or graphic.

Note: A note provides additional information that is of special interest. A note might point out exceptions to rules or procedures. A note usually, but not always, follows the information to which it pertains.

Caution: *informs you of conditions that might result in damage to hardware, corruption of data, or corruption of application software. A caution always precedes the information to which it pertains.*

Warning:A warning alerts you to conditions that might result in long-term health problems, injury, or death. A warning always precedes the information to which it pertains.

Mensajes de alerta

Los mensajes de alerta llaman la atención hacia información de especial importancia o que tiene una relación específica con el texto principal o los gráficos.

Nota: Una nota expone información adicional que es de interés especial. Una nota puede señalar excepciones a las normas o procedimientos. Por lo general, aunque no siempre, las notas van después de la información a la que hacen referencia.

Precaución: Una precaución informa sobre situaciones que podrían conllevar daños del hardware, de los datos o del software de aplicación. Las precauciones van siempre antes de la información a la que hacen referencia.

Advertencia: Una advertencia llama la atención sobre condiciones que podrían conllevar problemas de salud crónicos, lesiones o muerte. Las advertencias van siempre antes de la información a la que hacen referencia.

Related Documents

The following publications comprise the SVA document set available to Sun Microsystems Inc. customers.

Shared Virtual Array (SVA) Subsystem

Note: The book part numbers changed. The old numbers are shown in parenthesis.

- *StorageTek Shared Virtual Array (SVA) V2X/V2X2 Introduction 96216 (MO9135)*
- *StorageTek Shared Virtual Array (SVA) V2X/V2X2 Operation and Recovery 96217 (MO9137)*
- *StorageTek Shared Virtual Array (SVA) V2X/V2X2 Planning 96218 (MO9136)*
- *StorageTek Shared Virtual Array (SVA) V2X/V2X2 Reference 96219 (MO9139)*
- *StorageTek Shared Virtual Array (SVA) V2X/V2X2 System Assurance 96220 (MO9138)*
- *StorageTek Shared Virtual Array (SVA) V2X/V2X2 System Assurance Tables 96223 (MO9169)*
- *StorageTek Shared Virtual Array (SVA) V2X/V2X2 General Information 96221 (MO9134)*
- *StorageTek Shared Virtual Array (SVA) V2X/V2X2 Peer-to-Peer Copy Configuration User's Guide 96225 (MO9211)*

Shared Virtual Array Administrator (SVAA) for OS/390

- *SVAA for OS/390 Configuration and Administration PN 3112905xx*
- *SVAA for OS/390 Reporting PN 3112906xx*
- *SVAA for OS/390 Installation, Customization, and Maintenance PN 3112908xx*
- *SVA SnapShot for OS/390 Installation, Customization, and Maintenance PN 3112913xx*

Shared Virtual Array Administrator (SVAA) for VM

- *SVAA for VM Configuration and Administration PN 3134629xx*
- *SVAA for VM Reporting PN 3134630xx*

- *SVAA for VM Installation, Customization, and Maintenance PN 3134631xx*

Shared Virtual Array Administrator (SVAA) for OS/390 and VM

- *SVAA for OS/390 and VM Messages and Codes PN 3112907xx*

Shared Virtual Array Administrator (SVAA) for Solaris

- *SVAA for Solaris User's Guide PN 3112909xx*
- *SVAA for Solaris Messages PN 3112910xx*
- *SVAA for Solaris Installation PN 3112911xx*
- *SVAA for Solaris Quick Start Guide PN 3134509xx*
- *SVAA for Solaris Command Quick Reference PN 3134119xx*

Shared Virtual Array Administrator (SVAA) for HP-UX

- *SVAA for HP-UX User's Guide PN 3134257xx*
- *SVAA for HP-UX Messages PN 3134244xx*
- *SVAA for HP-UX Installation PN 3134254xx*
- *SVAA for HP-UX Quick Start Guide PN 3134512xx*
- *SVAA for HP-UX Command Quick Reference PN 3134253xx*

Shared Virtual Array Administrator (SVAA) for AIX

- *SVAA for AIX User's Guide PN 3134602xx*
- *SVAA for AIX Messages PN 3134600xx*
- *SVAA for AIX Installation PN 3134599xx*
- *SVAA for AIX Quick Start Guide PN 3134601xx*
- *SVAA for AIX Command Quick Reference PN 3134598xx*

Shared Virtual Array Administrator (SVAA) for Windows 2000 Server and Windows NT Server

- *SVAA for Windows 2000 Server and Windows NT Server User's Guide PN 3134573xx*
- *SVAA for Windows 2000 Server and Windows NT Server Messages PN 3134571xx*
- *SVAA for Windows 2000 Server and Windows NT Server Installation PN 3134570xx*
- *SVAA for Windows 2000 Server and Windows NT Server Quick Start Guide PN 3134572xx*
- *SVAA for Windows 2000 Server and Windows NT Server Command Quick Reference PN 3134569xx*

Shared Virtual Array Console (SVAC) for Windows NT

- *SVAC for Windows NT Quick Start Guide PN 3112993xx*

Other Documents

- *Peer to Peer Remote Copy Configuration Guide MP4007x*
- *Planning For IBM Remote Copy SG24-2595-xx* (IBM document)

- *Remote Copy Administrator's Guide and Reference SC35-0169-xx* (IBM document)

Viewing and Printing Web-Based Electronic Documents

Publications listed in "Related Documents" can be viewed and printed from the Sun Microsystems Inc. Customer Resource Center (CRC) Web site at:

`http://www.support.storagetek.com`

History of Changes

Rev A – Initial release. September, 2002.

Rev B – Second release. December, 2002

Minor changes involved edits and corrections. Major changes include:

- Added "Low Capacity FSC 3E41 Messages" section to chapter four.

Rev C – Third release. March, 2003. Minor changes involving edits and corrections.

Rev D – Fourth release. March, 2003. Minor changes involving edits and corrections.

Rev E – Fifth release. May, 2003. Minor changes involving edits and corrections.

Rev F – Sixth release. December, 2003. Minor changes involving edits and corrections.

Rev G – Seventh release. April, 2004. Minor changes involving edits and corrections.

Rev H – Eighth release. February, 2005. Minor changes involving edits and corrections.

Rev J – Ninth release. December 2005. Minor changes and corrections.

Rev K – Tenth release. May 2006. Minor changes and corrections in addition to:

- Replaced Table 31 on page 174 with new data.
- Removed OS/390 references or replaced them with z/OS except for publication names.

Rev L – Eleventh release. June 2006. Minor changes, additions, and corrections.

Rev M – Twelfth release. Late July 2006. Minor changes and corrections.

Rev N – Thirteenth release. October 2006. Minor changes and corrections in addition to:

- Removed outdated cable part numbers.
- Removed incorrect information about PAV and IFC cards.
- Obsolete feature numbers removed.

Site Planning

1

Subsystem Layout

Floor space and layout requirements differ for each subsystem depend upon the intended applications as well as the physical area available. In general, consider the following points:

- Cable lengths, and service clearance requirements
- Subsystem layout is determined by a combination of factors. Some factors include:
  - equipment components
  - plans for future expansion
  - procedures for using auxiliary I/O equipment
  - operator convenience
  - visibility from the console
  - other physical requirements.
- Storage space might be required within the computer room for temporary storage of media, circuit cards, printouts, and other data processing materials. All materials must be stored in properly-designed and protected storage areas.
- If ServiceTek Plus is to be installed (or is included in future plans), additional space is required for the associated modems and/or Extended Routing Switch (XRS). Additional phone lines are required, as well.
- Space might be required for office equipment, storage cabinets, files, tables, desks, and voice telephone equipment.
- Work-flow between the computer room and other areas must be carefully considered when planning aisles and intermediate storage areas.

Physical Characteristics

Because each customer has different requirements, make an accurate drawing of any proposed layout. The physical planning template, scaled at 1/4 inch to the foot, is provided in Figure 1 on page 25.

The physical characteristics of the SVA, including floor space requirements, are:

Physical AttributesDimensions

Height154.94 cm (61.0 inches)

Width (with two side covers)92.1 cm (36.24 inches)

Depth77.1 cm (30.34 inches)

Service Clearance (front and rear)54.1 cm (21.3 inches)

Side Clearance (left and right)6.4 cm (2.5 inches)

Weight445 kg (982 pounds)

Operators and service representatives require access to both the front and rear of the equipment. Therefore, there must be adequate space (or the ability to make space) for panel openings, test equipment, and for a service representative to work. In addition, there must be outlets available in the room for test equipment and tools.

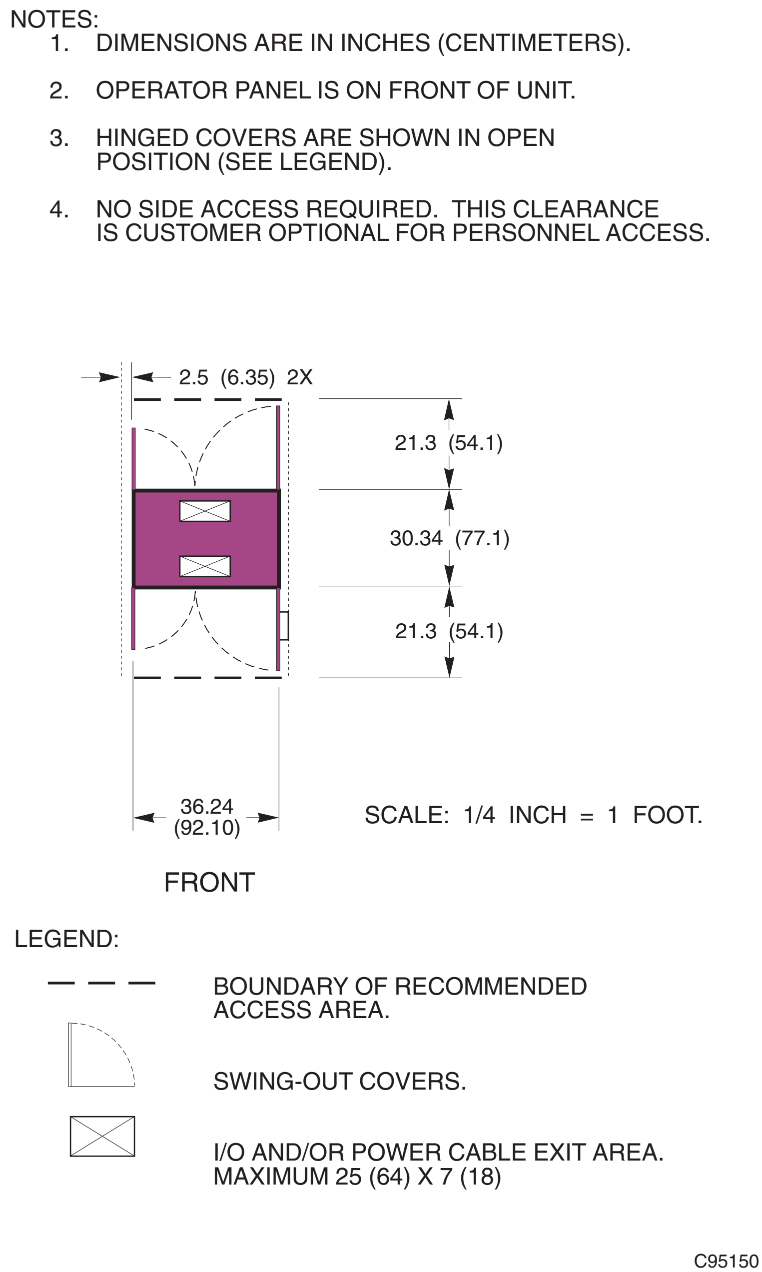

Figure 1 Shared Virtual Array Physical Planning Template

System Cables

Consider the following general issues when laying out the system and preparing the cable order:

- The location and number of control units assigned to each channel
- The number of channel interfaces that are connected to each control unit

- The routing of the I/O cables (i.e. outside/inside the computer room)

Locate subsystem components so the maximum cable lengths given in the equipment specifications are not exceeded.

Planning the cable order:

- Refer to the following section “Measuring Cable Lengths” for the correct cable length.
- Refer to the section “Electrical Power” on page 30 for power connection cable information.

Measuring Cable Lengths

Measure the center-to-center distance between cable entry holes (along the cable route) *plus* the length of cable required to reach the connector located on the device from the sub-floor.

Environmental Specifications

The environmental specifications for shipping, storing, and operating the SVA subsystem are shown in the following table.

Table 1 Operating Environment Requirements

Factors	Requirement
Temperature	60^{o} to 90^{o} F (16^{o} to 32^{o} C)
Thermal Change (max. rate/hr.)	9^{o} F (5^{o} C)
Humidity%	20 to 80 percent
Altitude	0 to 8000 feet (0 to 2439 meters)

Floor Loading

Strict compliance with ratings and weight distribution clearance is critical. The floor loading information provided must be reviewed carefully.

WARNING:
Failure to comply with the floor load rating and service clearance requirements can result in compromised structural integrity or raised floor structure collapse causing injury and/or property damage.

The basic floor load (over the machine footprint alone) for the SVA is 149 lb./sq. ft. (730 kg/sqm), with a superimposed floor load (with the same conditions as in Table 2 on page 29 with 30 inch ‘Z’ service clearance and 0 inches in ‘X’ and ‘Y.’ (also see Figure 3 on page 28)) of 99 lb./sq. ft. (485 kg/sqm).

Caution: The rating and condition of all raised floor panels and pedestals must be evaluated.

As additional clearance is provided around each SVA, the average floor load decreases. This is shown in Table 2 on page 29. Examples of floor loadings are listed in Table 3 on page 30. Figure 2 on page 28 indicates locations and loads for each load point. Equipment loading from adjacent floor area must be considered in the evaluation of the overall floor loading. Consult a structural engineer for a specific loading analysis and determination of floor capacity.

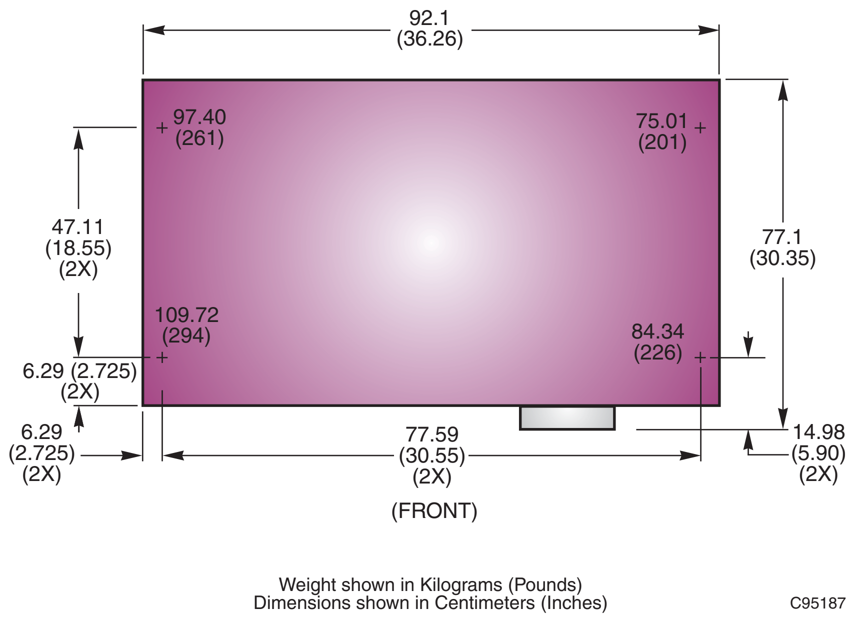

Figure 2 Shared Virtual Array Weight Distribution and Leveler Locations

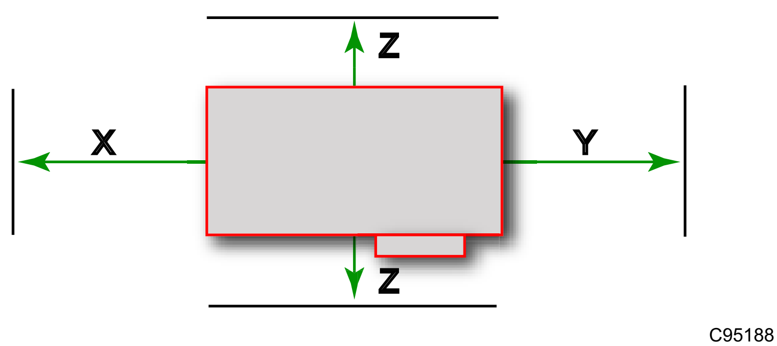

Figure 3 SVA Axis Reference

Table 2 Superimposed Floor Load

		Average End Clearance (X+Y) / 2, Inches (CM)				
		3 (7.6)	24 (61.0)	36 (91.4)	48 (121.9)	60 (152.4)
Service Clearance (Z) Inches (cm)	22 (55.9)	91 lb/sq. ft.^{2} (443 kg/m^2)	70 lb/sq. ft.^{2} (340 kg/m^2)	63 lb/sq. ft.^{2} (308 kg/m^2)	58 lb/sq. ft.^{2} (285 kg/m^2)	55 lb/sq. ft.^{2} (268 kg/m^2)
	38 (96.5)	77 lb/sq. ft.^{2} (374 kg/m^2)	60 lb/sq. ft.^{2} (295 kg/m^2)	55 lb/sq. ft.^{2} (270 kg/m^2)	52 lb/sq. ft.^{2} (252 kg/m^2)	49 lb/sq. ft.^{2} (239 kg/m^2)
	46 (116.8)	72 lb/sq. ft.^{2} (350 kg/m^2)	57 lb/sq. ft.^{2} (279 kg/m^2)	53 lb/sq. ft.^{2} (257 kg/m^2)	49 lb/sq. ft.^{2} (241 kg/m^2)	47 lb/sq. ft.^{2} (229 kg/m^2)
	54 (137.2)	68 lb/sq. ft.^{2} (331 kg/m^2)	55 lb/sq. ft.^{2} (267 kg/m^2)	51 lb/sq. ft.^{2} (247 kg/m^2)	48 lb/sq. ft.^{2} (232 kg/m^2)	45 lb/sq. ft.^{2} (222 kg/m^2)
	62 (157.5)	64 lb/sq. ft.^{2} (315 kg/m^2)	52 lb/sq. ft.^{2} (256 kg/m^2)	49 lb/sq. ft.^{2} (238 kg/m^2)	46 lb/sq. ft.^{2} (225 kg/m^2)	44 lb/sq. ft.^{2} (215 kg/m^2)

Note:

1. Average end clearance calculation formula: (X+Y) / 2.
2. Refer to Figure 2 on page 28. for location and load of each support point.
3. Values in this table assume 15 lb/ft.^{2} (73 kg/m^2) superimposed dead load over entire area for raised floor, cables, etc., and 15 lb/ft.^{2} (73 kg/m^2) live load for personnel and equipment in clearance areas between subsystems.
4. Loading of adjacent floor area must be considered in evaluation of overall floor capacity.
5. Other floor loading configurations are possible. Consult a structural engineer for specific loading analysis/evaluation of floor capacity.

Raised Floors

Locations of all loads at support points of the SVA are shown with physical specifications in Figure 2 on page 28. A representative of the customer's raised floor manufacturer or a structural engineer should be consulted to assure that installations on existing floors meet the following criteria. New floors should be designed and built to this criteria.

The raised floor panels must be capable of resisting a concentrated load of 1000 pounds (454 kilograms) anywhere on the panel with a maximum deflection of 0.08 inches (2 millimeters). In addition, a rolling load of 400 pounds (181 kilograms) anywhere over the panel must be accommodated. Protective floor covering should be used over access floor panels when moving the SVA units to prevent floor surface damage. Where floor tiles are cut to provide service access, additional pedestals might be required to maintain panel capacity. Perforated panels are not required for the SVA units, but if they are used, they must be verified for load rating.

The raised floor pedestals must be able to resist a 5000 pound (2268 kilogram) axial load. Lateral stability of raised floors in areas of high earthquake activity must be considered. The raised floor system shall be capable of resisting the lateral forces shown in the following table. These forces are applied at the top of the pedestal, and are based on the seismic zone in which the installation occurs.

Table 3 Horizontal Force Chart

Seismic Risk Zone	**V (Horizontal Force Applied at Top of Pedestal)**	
	Pounds	Kilograms
1	27.6	12.6
2A	41.4	18.8
2B	55.2	25.1
3	82.8	37.7
4	110.4	50.2

Note: Horizontal forces based on 1991 Uniform Building Code (UBC) Sections 2336 and 2337 assuming minimum operating clearance of multiple V2X2 Shared Virtual Array units. Installations in areas not covered by the UBC should be engineered to meet seismic code provisions of local jurisdiction.

Power

Electrical Power

This equipment operates best with a reliable power source that is free from interference or disturbance. Power supplied by local power companies is generally of sufficient quality.

To take full advantage of the SVA's subsystem AC power redundancy, the two AC inlet/line filter assemblies should receive power from two independent sources that are not likely to fail simultaneously.

Qualified personnel must ensure that the voltage measured at the power receptacle is within the specified tolerance for the machine. In addition to power for the SVA, a separate power feeder is required for such items as lighting, air conditioning and other electrical systems.

Electrical Connectors

The part number of the subsystem power connectors for 60 Hz applications (supplied with the units – see the right side of the following figure) is a Russell Stoll 3750DP

The compatible mating connectors (which are not provided) are:

- Russell Stoll 9R33U0W (rigid/wall mount)
- Russell Stoll 9C33U0 (flexible mount)

Note: There is no Hubbell compatible connector.

Because of the varied connection requirements of 50 Hz power systems, power connectors for 50 Hz applications are not provided. See the left side of the following figure.

Figure 4 Power Cable Ends

Power Source Specifications

The SVA subsystem is designed to operate at the following specifications:

- AC power source – Single Phase, 3 wire.
- Voltage range – 200-240 VAC.
- Frequency range – 50-60 Hz.
- Circuit size – 30 Amperes.

The following paragraphs provide additional specifications for SVA subsystem power.

Power Redundancy Levels

The following are examples of different levels of possible AC redundancy (from low to high):

1. No redundancy. Both AC cords from the same unit connected to different panels and breakers which receive power from the same local utility source.
2. Secondary Selective System. The site is redundantly powered from two sources and is equipped with a transfer mechanism; essentially, a building with “two line cords.”
3. One side of the Shared Virtual Array subsystem powered from the utility; the other side connected to the same utility *with* an uninterruptable power source (UPS) backup.
4. Both sides of the subsystem powered from the same utility; both sides backed up from individual UPS systems.
5. Separate utility sources available at the site; each side of each subsystem unit powered from a different utility.
6. Separate utility sources available at the site; each side of each subsystem powered from a different utility, and with a UPS backup on one source.
7. Separate utility sources available at the site; each side of each subsystem unit powered from a different utility, and with a UPS backup on both utility sources.

Harmonic Content

When the equipment is operating, the maximum total harmonic content of the power system voltage waveforms on the equipment feeder must not exceed 5% with the equipment operating.

Power Measurements

The following tables show power figures for single line cord operation and dual line cord operation and dual line cord operation with one power source failing or one PDU not operational. The value for current shown in those tables is the total current measured into both line cords. All arrays contain sixteen drives.

The SVA dual line cord current is defined as the current in each line cord.

Table 4 Dual Line Cord Operation (36 GB Drives)

No. of 16-Drive Arrays	AC Voltage In	AC Current In	KVA	KWatts	Power Factor	KBTU/Hr
4	264	7.2	3.8	3.6	0.95	12.4
	208	8.9	3.7	3.6	0.98	12.4
	180	10.2	3.7	3.6	0.99	12.4
3	264	6.5	3.4	3.3	0.95	11.1
	208	8.0	3.3	3.3	0.98	11.1
	180	9.1	3.3	3.3	0.99	11.1
2	264	5.8	3.0	2.9	0.95	9.9
	208	7.1	3.0	2.9	0.98	9.9
	180	8.1	2.9	2.9	0.99	9.9
1	264	5.0	2.7	2.5	0.95	8.6
	208	6.2	2.6	2.5	0.98	8.6
	180	7.1	2.6	2.5	0.99	8.6

Table 5 Dual Line Cord Operation (73 GB Drives)

No. of 16-Drive Arrays	AC Voltage In	AC Current In	KVA	KWatts	Power Factor	KBTU/Hr
4	264	6.9	3.7	3.5	0.95	11.9
	208	8.5	3.5	3.5	0.98	11.9
	180	9.8	3.5	3.5	0.99	11.9
3	264	6.3	3.3	3.1	0.95	10.7
	208	7.7	3.2	3.1	0.98	10.7
	180	8.8	3.2	3.1	0.99	10.7
2	264	5.6	3.0	2.8	0.95	9.6
	208	6.9	2.9	2.8	0.98	9.6
	180	7.9	2.8	2.8	0.99	9.6

Table 6 Single Line Cord Operation (36 GB Drives)

No. of 16-Drive Arrays	AC Voltage In	AC Current In	KVA	KWatts	Power Factor	KBTU/ Hr
4	264	14.5	3.8	3.6	0.95	12.4
	208	17.8	3.7	3.6	0.98	12.4
	180	20.3	3.7	3.6	0.99	12.4
3	264	13.0	3.4	3.3	0.95	11.1
	208	16.0	3.3	3.3	0.98	11.1
	180	18.3	3.3	3.3	0.99	11.1
2	264	11.5	3.0	2.9	0.95	9.9
	208	14.2	3.0	2.9	0.98	9.9
	180	16.2	2.9	2.9	0.99	9.9
1	264	10.1	2.7	2.5	0.95	8.6
	208	12.4	2.6	2.5	0.98	8.6
	180	14.2	2.6	2.5	0.99	8.6

Table 7 Single Line Cord Operation (73 GB Drives)

No. of 16-Drive Arrays	AC Voltage In	AC Current In	KVA	KWatts	Power Factor	KBTU/ Hr
4	264	13.9	3.7	3.5	0.95	11.9
	208	17.0	3.5	3.5	0.98	11.9
	180	19.5	3.5	3.5	0.99	11.9
3	264	12.5	3.3	3.1	0.95	10.7
	208	15.4	3.2	3.1	0.98	10.7
	180	17.7	3.2	3.1	0.99	10.7
2	264	11.2	3.0	2.8	0.95	9.6
	208	13.8	2.9	2.8	0.98	9.6
	180	15.8	2.8	2.8	0.99	9.6

Grounding

For personnel safety, as well as electrostatic discharge (ESD) protection, a machine must be properly grounded. An insulated green or green/yellow wire ground, at least the same size as the phase wire, is required between the branch circuit panel and the receptacle. *This is not a neutral line*.

Note: Local codes and standards might establish more stringent requirements. Always comply with local codes.

For personnel safety, the ground must have sufficient low impedance to limit the voltage to ground and to facilitate the operation of circuit protective devices. For example, the ground path shall not exceed one ohm for 120 volt, 20 ampere branch circuit devices. This yields a maximum machine voltage of 20 volts in the event of internal shorts.

Sun recommends good bonding techniques such as bared metal and star washers. It is not sufficient to provide ground paths through anodized material or hinges. Sheet metal screws must never be used as a means of attaching a ground.

A low impedance grounding and lightning protection must be considered when planning and installing an electrical system. The electrical contractor must meet local and national code requirements.

Sun equipment power cords contain an insulated equipment grounding wire (green or green with a yellow stripe) that connects the machine frame to the ground terminal at the power outlet.

Static Electricity

Static electricity can be caused by the movement of people, furniture, and equipment. In addition to causing discomfort to personnel, an electrostatic discharge (ESD) can cause a machine malfunction. It is possible for electrostatic discharge (ESD) to destroy discrete components on logic cards, change information on magnetic media, or cause other problems.

Special attention must be given to floor paneling and floor covering. Panels must be made with nonconducting cores for high resistance. Anti-static floor coverings are required.

Other sources of static electricity include furniture, casters, vacuum cleaners, and plastic access panels. The resistance between furniture and the floor must be at least 10 *u*ohms.

Every possible precaution must be taken to minimize static discharge. This can be accomplished by providing a conductive path from raised floors to ground and maintaining the computer room humidity within the recommended control parameters.

Power Distribution System

The following defines the electrical power distribution system required for proper operation of your SVA.

Power Load

By calculating the power load of all the equipment, you should determine the number of amps required for your power service. Large inrush conditions, such as motor start times, need to be considered. When multiple subsystems/units are to be installed, your electrician should rotate the phases when connecting power (to prevent overloading any single phase).

Power Panel Feeders

Ensure that the feeder wires to the branch-circuit distribution panel have adequate current handling capacity for the SVA power load.

Branch Circuits

Note: Electricians are to review and adhere to local codes and regulations at all times when planning and installing branch circuits.

The individual branch circuits on the panel must be protected by suitable circuit breakers properly rated according to manufacturer specifications and applicable codes. The branch circuit breaker must not exceed the local code requirements to protect branch circuit conductors. Breakers must have delays for motor start inrush. Each circuit breaker must be labeled to indicate which branch circuit it is protecting.

Branch circuits installed under a raised floor must be within 10 feet (three meters) of the machines they supply. If the branch circuits are required to run in a metal conduit, either rigid or non-rigid, the conduit system must be continuous and uninterrupted from the machine power cord connection to the building or transformer ground.

Electrical Problem Areas and Issues

Equipment errors can be caused by outside (radiated or conducted) transient electrical noise signals being superimposed on the power provided to the equipment. The following figure depicts a typical transient grounding plate, which is not to be confused with the code approved grounding electrode.

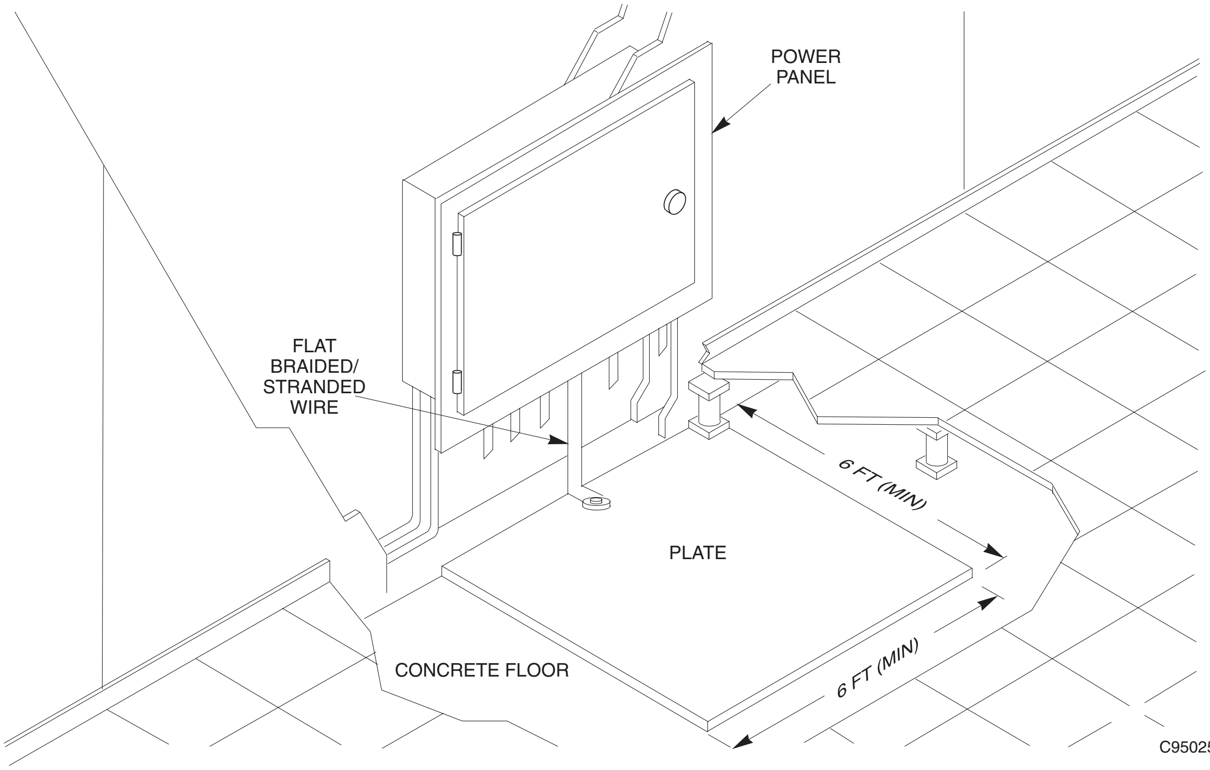

Figure 5 Transient Grounding Plate

A means for disconnecting power from all equipment in the computer room must be provided. This means of disconnecting must be in locations readily accessible to the operator at the principle exit doors. Consult local and national codes to determine the requirements for this disconnect.

The SVA is designed to operate within most common power line disturbances. The quality of incoming power, however, can make a substantial difference in the performance of the equipment. Power disturbances or transients can cause equipment power failures or errors. Transients can come into the site on the power company lines, but (with the exception of lightning) are more often caused by electrical equipment installed within the building. For example, transients can be caused by welders, cranes, motors, induction heaters, elevators, copy machines and other office equipment.

Failures caused by the power source are basically of three types:

1. Power outages include short duration dips in voltage as well as prolonged outages (brownouts). To take advantage of the SVA subsystem redundant AC power distribution function, one subsystem input should be connected to a UPS (uninterruptable power source) system.

2. Transient electrical noise superimposed on power lines can be caused by a variety of industrial, medical, communication, or other equipment:

   - Within or adjacent to the computing facility
   - Near the power company's distribution lines.

3. Switching large electrical loads can cause problems, even though the source is on a different branch circuit. If such a condition is suspected, Sun recommends a separate dedicated feeder or transformer for your SVA directly from your power source.

   Motor-operated devices on the same power source as an SVA can, under certain conditions, cause intermittent electrical disturbances. To minimize the effects of high frequency noises, the branch circuit power panel servicing the equipment must be mounted in contact with bare building steel or connected to it by a short length of cable.

   If this is not possible, a metal area (power panel plus conduit, plus plate) of at least 10 square feet (approximately 1 square meter) in contact with masonry can be used. The plate must be connected to the green wire common. The connection must not be more than 5 feet (1.5 meters) long and must consist of No. 12 AWG (0.0051 square inches or 3.3 square millimeters) or larger wire.

   If transient-producing devices have been removed from the power bus and the computer room power panel and power line disturbances still exist, it might be necessary to install power line isolation equipment. This equipment includes:

   - Transformers
   - Motor generators
   - Uninterruptable power sources (UPS)
   - Other power conditioning equipment

Configuration Planning

2

The SVA supports the use of different drive sizes in the same frame. However, all trays must be fully populated and all drives within a drive tray must be the same size. All drive sizes are shown in giga-bytes.

Table 8 V2X/V2X2 Mixed Drive Size Variations Supported

Drive Tray #1[a] (bottom)	Drive Tray #2	Drive Tray #3	Drive Tray #4 (top)
36	36	36	36
36	36	36	
36	36		
36	36	73	73
73	73	73	73
73	73	73	
73	73		
73	73		

a. All drive trays contain 16 drives.

Net Capacity Load And Free Space

Net Capacity Load (NCL), Collected Free Space (CFS), and Uncollected Free Space (UFS) comprise the total capacity of the V2X2. The physical capacity of the V2X2 is the Disk Array Capacity (DA Capacity).

The Physical Capacity Control (PCAP) is the effective capacity that the customer purchased and includes an assumed compression ratio. The

UFS might be displayed as 0% until the difference between the actual physical capacity and the PCAP is exceeded.

The NCL, CFS, and UFS capacity categories support the virtual architecture, where storage is assigned to functional volumes as needed and data is not updated in the same location.

Net Capacity Load

Net Capacity Load (NCL) represents the percentage of the SVA's storage capacity used to store the compressed functional tracks. The NCL percentage fluctuates during the processing of the data, including writes and deletions. An increase or decrease of NCL is detected in any environment other than read-only. When the disk utilization (i.e. NCL) exceeds 90%[1], the customer should be aware that they might experience some performance degradation. The customer then needs to determine if the performance level is acceptable for their requirements. If the levels of performance are not acceptable, the customer should contact Sun to order additional capacity. (Also see "Storage Capacity Management" on page 41.)

Free Space

There are two forms of free space within the V2X2; uncollected free space, and collected free space.

Uncollected Free Space

When a record is updated in a V2X2, the space in which the original data resides is not over-written. The data is written on a new cylinder and the pointer to that data is changed. The old data, while no longer valid, still resides on its track. That space is made available for use when the SVA collects free space as a background task within the SVA. See "How Free-Space Collection Works" on page 41.

Collected Free Space

In a SVA, data written to an array is always written to a free array cylinder[2]. the SVA keeps a list of free cylinders and uses those cylinders for writing new data or updating data in the array. Collected free space is a short term indicator of capacity consumption.

The disk array controller only schedules write operations to free array cylinders. Data written to these array cylinders as records is updated or created on functional volumes.

---

1. This figure could be closer to 75% or 80% in subsystems not running PCAP.
2. Functional volumes and functional tracks are the host view of the V2X2. Arrays, array tracks, and array cylinders are the V2X2 view of the physical storage devices.

How Free-Space Collection Works

The SVA performs free space collection (FSC) as a background task to maintain the availability of collected free space (CFS). As allocated array cylinders age with unused space identified as UFS, they become potential candidates for FSC. Thresholds and algorithms provide the mechanisms for free space collection to operate efficiently.

The free-space list identifies the uncollected free space on the allocated array cylinders. At a high priority, FSC processes cylinders with a high percentage of free space. When the utilization of the SVA is less, the FSC processes cylinders with less free space. After the valid data is consolidated and written to free array cylinders, the processed cylinders append to the free-cylinder list and the CFS.

Storage Capacity Management

To manage the capacity effectively, implement the Deleted Data Space Release (DDSR) function provided by Shared Virtual Array Administrator (SVAA). DDSR identifies the storage occupied by the host-deleted data to be reclaimed by free space collection (FSC). Otherwise, this storage remains assigned to a functional volume until allocated and updated. The space reclaimed by FSC is available for use by any functional volume. For further information concerning DDSR, refer to the manual titled Shared Virtual Array Administrator for z/OS or VM Configuration and Administration.

The amount of actual data stored and the NCL percentage varies based on the V2X2 compression ratios and the use of the DDSR function. If the host system or application compresses any data, then review the rationale. Host compression are not reflected in the compression ratios displayed on the SVA's space utilization reports. Monitor the NCL and the CFS during data migration, database reorganizations, and changes to workload activity. If full volume migrations or restores are performed and host-deleted data is copied, then execute interval DDSR on these volumes. Prior to initializing or restoring volumes with existing data, delete and define the volume. Periodic evaluation of the data growth, calendar cycles, new applications, and workload attributes, especially high write activity, is recommended.

For information concerning the LISTCFG command and the space utilization reports to obtain statistics for NCL and free space, refer to the manual titled Shared Virtual Array Administrator for z/OS or VM Reporting.

The following figure show a graphical representation of the total capacity of an SVA and the smaller segments of the capacity taken up by customer date, the uncollected free space, the boundary where the SVA sends a message to the host if more information (data) is stored on the SVA.

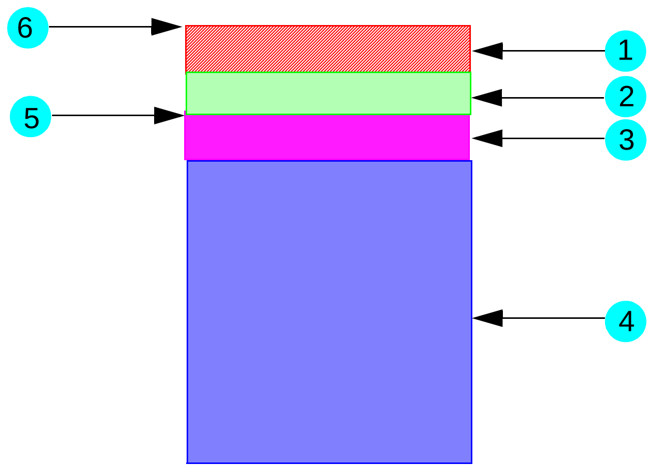

Figure 6 An V2X2 At 10% Collected Free Space

Table 9 An V2X2 At 10% Collected Free Space

Figure 6 Reference	Meaning
1	Physical storage beyond the PCAP limit that has been used by the SVA to temporarily store Uncollected Free Space
2	Uncollected Free Space.
3	If Collected Free Space falls to 10% of NCL, the SVA sends a SIM message with an FSC 3E41 to the MVS host console
4	Customer Data
5	PCAP Limit
6	Physical capacity

Caution: Performance Impact: as the NCL increases and the management of the storage by the SVA intensifies, an impact to performance could be detected. Low collected free space percentages might degrade performance and require immediate response to warning messages. When workload activity is high, CFS might decrease and UFS might increase.

Recognize warnings that the available storage capacity is low or exhausted:

- Less than 10% collected free space, SIB1930W, an SVA message, or a service information message (SIM) identified as IEA480E with the serial number of the SVA and a REFCODE=3E41, the fault symptom code, indicates low capacity.
- Due to 0% collected free space and the consequence of no more writes, SIB1931S, a V2X2A message, or a SIM identified as IEA480E with the serial number of the SVA and a REFCODE = 3E40, the fault symptom code, indicates exhausted capacity.
- Respond to the warnings by using the DDSR subcommand to identify candidates for free space or one of the methods to manage capacity. Further action might require the deletion of functional devices if the data is not required or available from backups or copies.

For further information concerning the SIM messages, refer to the *V2X2 Shared Virtual Array Operations and Recovery*.

As the NCL number begins to approach the 90% figure, several things can be done:

1. Use SVAA NCL Management Utilities to reclaim deleted data space.
2. Add more physical capacity to the subsystem by configuring another disk array, or for the SVAs that use PCAP, increase the PCAP logical limit to the next supported increment.
3. Move some data to another SVA with a lower NCL percentage.
4. Reduce the update write and write/delete content of the user workload by swapping data to another SVA.
5. Install an additional SVA.

Selecting a Subsystem Configuration

The process of selecting a disk hardware configuration has always been complex. Factors such as I/O rates, cache hit rates, the peak capacity of data, and response times influence the decision-making process. Typically, storage managers make decisions that optimize a specific configuration to meet either capacity requirements or performance requirements.

Considerations for performance and capacity requirements might be contingent upon:

- redundancy and availability, since the SVA consists of two distinct logical redundancy parts (clusters 0 and 1),

- hardware specifications and configuration - Cache capacity, storage capacity, channel connectivity,
- host processor capability and configuration, sysplex, geoplex,
- data characteristics - workload and application attributes, host-compressed data, data compressibility, reads versus writes ratio, skewed versus balanced workload, online and batch processing, modeling, growth rate, quantity,
- peak capacity required to process applications and respective calendar cycles,
- distribution of temporary and sort-work space,
- production, test, and backup applications,
- placement of primary and alternate critical system data sets on separate subsystems, i.e., COUPLE, JES, RACF,
- SnapShot and PPRC capacity,
- service level agreements, performance commitments,
- device types and configuration,
- physical or logical data migration and the source of the data,
- and multiple subsystems.

Functional Considerations

The SVA's architecture requires some functional configuration considerations, which are described in the following sections.

Note: You can perform the configuration tasks and sub-tasks from the local operator panel or from the host using the Shared Virtual Array Administrator (SVAA). Because of its flexibility, StorageTek recommends performing most of the configuration sub-tasks using SVAA configuration services.

Functional Configuration

Before configuring an SVA subsystem from the host operating system, you must first establish a minimum functional configuration from the local operator panel.

The initial SVA subsystem configuration is identified during the system assurance process. Then, as part of the subsystem installation, service representative sets up the minimum subsystem configuration. (Refer to the *V2X2 Shared Virtual Array Operation and Recovery* for more information about the initial subsystem configuration.) Next, a trained service representative modifies the minimum configuration, which includes defining the functional configuration of the subsystem.

The only restriction is that the physical capacity of the attached disk arrays supports the actual Net Capacity Load (NCL) of the subsystem. The NCL is the physical disk array capacity required to store all of the compressed functional tracks actually occupied by data in a functional configuration.

Functional 3380 or 3390 volumes might be intermixed freely in a subsystem configuration; however:

- 3380 track compatibility mode (TCM) for 3390 volumes is not supported,
- and HCD/IOCP, and other host system constraints do apply.

Note: Implementing 3380 and 3390 functional volumes requires the same operating system maintenance level required to implement actual 3380 and 3390 volumes, and 4-digit addressing.

Missing-Interrupt Handler (MIH)

Although SVAA adjusts the Missing-Interrupt Handler (MIH) for SVA devices automatically, it is important that you know the minimum value to which you can safely set the MIH manually.

The SVA's internal error recovery procedures require that the MIH time-out value be set to four minutes and ten seconds (the default MIH time-out value is 15 seconds). Known I/O timer or MIH values are:

- **For z/OS or MVS:** Set the MIH to four minutes ten seconds, and If you are installing SVAA, you can elect (by default) to have SVAA set the MIH time-out value automatically.
- **For Solaris:** If installing SVAA, set the I/O timer to 270 seconds.
- ***For VM*:** You must set the MITIME value manually to four minutes ten seconds; you cannot elect to have SVAA set the MIH time-out value automatically. Refer to the CP command "Set MITIME."

Note: For systems not listed, set the I/O timer or MIH to 270 seconds.

This change forestalls the invocation of the ERPs for the missing interrupt handler in normal error recovery situations.

Note: Adjust applications running on your system that detect missing interrupts and that are independent of the system MIH setting to allow a 270 second missing interrupt time-out value.

Also see "Setting I/O Timers" on page 72 for HP-UX, Windows NT etc.

Disk Array Definition

In an SVA, fifteen physical devices are organized into a logical group. Within the group, user data is recorded on identically addressed tracks on all but two of the devices. The identically addressed tracks on the other two devices are used to store the two levels of redundancy data generated by the subsystem.

The redundancy data is dynamically placed across the devices in the array, so that all of the devices in the array contain user data and redundancy data.

In a 15-device array, which is referred to as a 13+2+1 array, one physical device in the array is reserved as a spare device and is globally available to the subsystem. This spare device is referred to as a “hot spare.” This configuration provides the maximum data capacity with an adequate number of spare drives for data recovery.

Array Partitions

Part of the SVA configuration process includes assigning physical devices to an array partition. A physical device might be associated with one of the following partitions:

Spares Partition

Drives in the “Spares Partition” are used to form test and production arrays, to reconstruct failed disks, and to receive data from a draining device.

Production Partition

Physical devices in the “Production Partition” comprise the arrays that store user data. These arrays are formed from the physical devices in the “Spares Partition.”

Unavailable Partition

The “Unavailable Partition” includes any device slots that have not had devices installed, any powered off devices, or any broken devices. The subsystem manages the “Unavailable Partition.”

HCD/IOCP Definition

Refer to “Device Map And Host Addresses” on page 115 and “IOCP/HCD” on page 121 for planning the HCD/IOCP definition. This definition is a prerequisite to the mainframe host recognition of the V2X2 subsystem.

Caution: Performance Impact: Unit = ***3990*** must be coded in the control unit definition (no model number).

SVA Device Type Support

Functional 3380 or 3390 devices might be intermixed freely in a subsystem configuration. However, storage administration might be tedious. Grouping devices of the same type is recommended. The following device types are supported: 3380-J, 3380-K, 3390-1, 3390-2, 3390-3, and 3390-9.

The device type, 3390 or 3380, as defined on the subsystem ***must match*** the ***device type*** specified in the IOCP/HCD. Functional devices might be defined with SVAA screens or the subsystem operator panels.

Maximum Functional Devices and Virtual Cylinders

Neither of the following limits can be exceeded during device definition:

- The **maximum number** of **functional devices** is **4,096** defined as 3390-3 or smaller device types.
- The **maximum number** of **total cylinders** for all devices is **13,676,544.** (See "Device Characteristics" on page 169 for more information on 3380 and 3390 devices.)

A maximum configuration is 4,096 devices defined as a 3390-3 or smaller device type.

The maximum number of 3390-9 devices is 1,365. This restriction is due to the total maximum number of cylinders for devices defined on the subsystem.

Device Support

See "Device Characteristics" on page 169 for information regarding 3380 and 3390 device characteristics. Note the cylinders per device if the total number of cylinders for functional devices in the subsystem needs to be calculated.

In MVS, a minimal initialization with ICKDSF is supported to define the volser, VTOC, and VTOC Index. Performance could be degraded on MVS operating systems if a VTOC Index is not defined. The VTOC, VTOCIX, and VVDS might be placed at the beginning of the volume.

- Consider the physical capacity of the subsystem to support the actual NCL of the subsystem. Compression rates vary depending upon the type of data.
- Programs restricted to 64k tracks do not support 3390-9 device types.
- Allocation of large databases or files on 3390-9 device types is recommended. Numerous, small, active files could degrade performance due to VTOC enqueues.

FlexVolume Support (FVS)

FlexVolume Support allows for the definition of variable sized volumes of 3380 and 3390 device types. These FlexVolume definitions allow you to define a lower number of cylinders than the device type would normally have.

Note: FlexVolume requires a code level or B01.09 or higher.

Restrictions

- No FlexVolumes on SCSI devices - only on CKD device types.
- No FlexVolumes on PPRC bridge devices.
- No 3390-1 or 3390-2 FlexVolumes. This is in keeping with how IBM does it on the 2105 ESS.

Options

The following table lists the allowed range of cylinders you can define for each device type.

Table 10 FlexVolume Cylinder Options per Device Type

Device Type	Cylinder Definition Limits
3380-J	1-885
3380-KE	886-1770
3380-K	1771-2655
3390-3	1-3339
3390-9	3340-10017 (without LVS)
	3340-32760 (with LVS)

SSID

Definition

The subsystem identifier (SSID), a user-assigned value, uniquely identifies each functional storage control unit in the customer configuration.

The subsystem identifiers (SSID) ascertain the identity of the functional storage control units (the functional 3990 storage controls). A maximum of sixteen SSIDs can be assigned to a single SVA.

The functional storage control unit is also known as the virtual control unit, logical control unit, or control unit image.

Assignment

To assign the SSIDs:

- Define a SSID in the hexadecimal range 0001 to FFFF.

Note: When entering SSID numbers on the CD21 screen, leading zeros are not accepted by the SVA – which excludes the possibility of using an SSID of zero.

- Designate 16 sequential SSIDs, the first one being an even numbered SSID, to the 16 functional control units.

Caution: Performance Impact: each SSID number must be unique within the customer configuration, including PPRC primary and secondary subsystems. At initial program load (IPL) and vary on time, the MVS system builds control blocks that associate a device to the storage control subsystem. A storage control with a duplicate SSID cannot be initialized as online and operational. MVS issues error message "IEC334E DUPLICATE SUBSYSTEM SSID DEVICE addr NOT BROUGHT ONLINE" to the master console.

Changing SSIDs

Changing the SSIDs after the subsystem is operational can be disruptive. PPRC must be disabled with an options disk provided by Sun. All attached devices must be varied offline. After the SSIDs are changed, vary all devices online to permit the connected systems to access the devices.

If PPRC is disabled with the options disk, the SSIDs can be changed dynamically. However, a problem might not be detected until the next IPL or the devices are varied offline and online.

Large LUN

Larger Logical Unit Number (LLUN), also called a disk in the open systems environment, functionality provides a mechanism for grouping small virtual disk devices into a single larger virtual "disk" for presentation to the host operating system.

These LLUNs can be created and managed through the Local Operator Panel (LOP) or by SVAA and is supported by SnapShot 3.1 or higher.

In Order to understand LLUN better, be aware that:

- The management method for a grouping of functional devices (FDID's), called LLUN, establishes a relationship between virtual devices.
- A LLUN is configured with either 3390 model 3 or model 9 type devices.
- A LLUN must be composed of all 3390 model 3 or 3390 model 9 virtual devices. Mixed configurations are not allowed.

- A LLUN is made of a number of devices; by definition, more than one.
- Any given device can only belong to one LLUN.
- Changes made on any configuration flags or attributes of a LLUN affects all devices of that LLUN even if they are maintained at “index 0” device level only.
- Maximum size of a LLUN is 2 TB. This is a SCSI limitation and not a function of the V2X2.

Note: Sun does not recommend LLUNs beyond 1 TB due to many host operating system limitations.

- The maximum number of devices per LUN is 1,024. This number might be lower due to subsystem capacity and configuration constraints.

The SVA always presents itself to the host as a 512 block size device with out regard to the block sizes used in the back end as shown the following table.

Table 11 SCSI Track Information

Internal Block Size	Internal Records/ Track	Bytes/Track	SCSI Blocks/Track
512	48	24576	48
2048	20	40960	80
4096	12	49152	96
8192	6	49152	96
16384	3	49152	96

The following table is offered for reference only. Various host limitations could reduce the number of devices per LLUN and the LLUN size.

The dark shaded LLUN sizes in the following table exceed 2 TB. While LLUNs of those sized could be built, the host system would not be able to access storage beyond 2 TB with the current SCSI 10 CDB limitations.

Table 12 Examples of LLUN Sizes Using 3390 Based Devices

Block size	1 vDev=	Number of vDev / LUN					
		32	64	128	256	512	1024
3390-3							
512	1.146	36.661	73.323	146.646	293.291	586.582	1,173.164
2048	1.909	61.102	122.205	244.409	488.818	977.637	1,955.273
4096	2.291	73.323	146.646	293.291	586.582	1,173.164	2,346.328
8192	2.291	73.323	146.646	293.291	586.582	1,173.164	2,346.328
16384	2.291	73.323	146.646	293.291	586.582	1,173.164	2,346.328
3390-9							
512	3.438	110.028	220.056	440.112	880.225	1,760.449	n/a[a]
2048	5.731	183.380	366.760	733.521	1,467.041	2,934.082	5,868.164
4096	6.877	220.056	440.112	880.225	1,760.449	3,520.898	7,041.797
8192	6.877	220.056	440.112	880.225	1,760.449	3,520.898	7,041.797
16384	6.877	220.056	440.112	880.225	1,760.449	3,520.898	7,041.797

a. The maximum capacity of a LUN made of 3390-9 devices formatted in 512 bytes blocks is 2,045.835 GB (595 vDev).

Physical Capacity Constraints

You can configure each SVA with any combination of devices *provided*:

- You do not exceed the total number of 13,676,544 cylinders, given the following cylinder usage by device type:

  3390-3 = 3,339 cylinders

  3390-9 = 10,017 cylinders

  3390-27 = 32,760 cylinders

  FlexVolume = User definable cylinder

- You do not exceed the total number of 4096 physical devices

Contact Sun for help determining the maximum number of volumes and their mix.

PPRC Considerations

Primary volumes must still be paired with like size or larger Secondary volumes.

Note: Even though you can PPRC to a larger secondary volume, you shouldn't for disaster recovery purposes. See the SVAA Configuration and Administration Guide for your operating system for SVAA language on the subject.

Virtual Capacity Cylinder Gauge

The Virtual Capacity Cylinder Gauge (VCCG) monitors and displays the availability and usage of SVA's virtual capacity for ease of management and configuration.

The VCCG feature provides a usage display of the following data fields in the Subsystem Configuration and Status screen (see the following figure). This gauge, in the Configuration Status area (center of the screen) shows:

- Total Cylinders - Total number of virtual cylinders supported by this subsystem
- Configured Cylinders - Number of virtual cylinders in use by this subsystem.
- Unconfigured Cylinders - Available number of virtual cylinders available on this subsystem

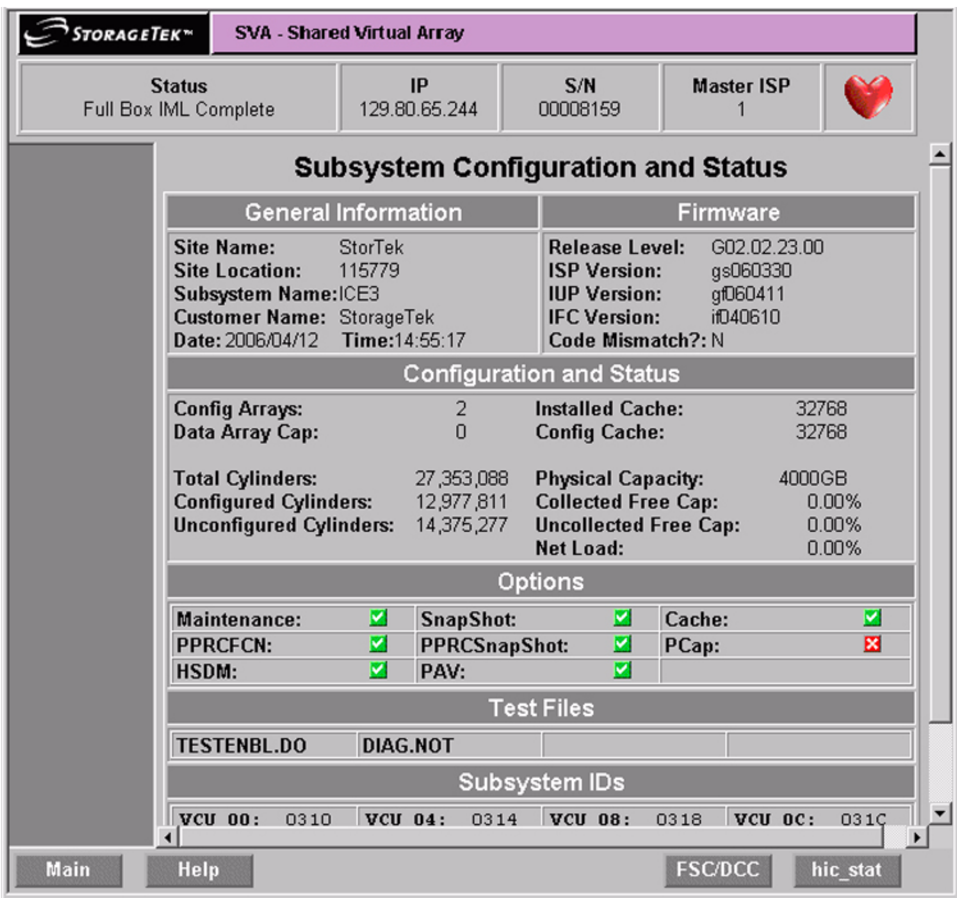

Figure 7 Subsystem Configuration and Status Screen with Cylinder Gauge

The implementation of Large Volume Support (LVS) involves support for virtual devices with more than 10,017 virtual cylinders and support for devices with non-standard numbers of virtual cylinders. Sun provides these two capabilities together because the only technique that can be used for specifying a device with a cylinder count greater than 10,017 is to provide a means of selecting the cylinder count rather than just selecting the device type. This is because a device with more than 10,017 cylinders is still a 3390 Model-9 device. The host does not recognize device types of Model-27 despite the widespread use of these terms to describe volumes with 32,760 cylinders.

Constraints Imposed by the User Interface

There are two high level external interfaces that allow a user to define virtual devices: SVAA and the Detached Operator Panel (DOP). These user interfaces impose the following constraints on the virtual devices that are defined through these interfaces:

1. The user may select a device type from the following list:
   - 3380 J (885 cylinders)
   - 3380 K (2655 cylinders)
   - 3380 6EA_K (1770 cylinders)
   - 3390 Model-1 (1113 cylinders)
   - 3390 Model-2 (2226 cylinders)
   - 3390 Model-3 (3339 cylinders)
   - 3390 Model-9 (10,017 cylinders)
2. The user may select a track format and a cylinder count. The track format may be either 3380 or 3390. The cylinder count may be any value from 1 through 65,520. When the user specifies a cylinder count, the user will either:
   - Not have the option to specify the device type
   - Specify the valid virtual device type for the specified cylinder count. The user interface (SVAA or DOP) selects the valid virtual device type based on the specified track format and the cylinder count.

   A. For a 3380 track format, the valid virtual device types are:

   - 3380 J for cylinder counts 1 through 885
   - 3380 KE for cylinder counts 886 through 1770
   - 3380 K for cylinder counts 1771 through 2655

B. For a 3390 track format, the valid virtual device types are:

-3390 Model-3 for cylinder counts 1 through 3339

-3390 Model-9 for cylinder counts 3340 through 32,760

The ability to select a non-standard cylinder count that is greater than the largest standard device is required to enable the user to define a device with more cylinders than a 3390 Model-9. The ability to select a non-standard cylinder count that is less than the cylinder count for a standard device is useful to customers who want to place a frequently accessed dataset on a dedicated volume in order to eliminate the possibility that access to that dataset will be delayed due to an access to another dataset residing on the same volume. In the SVA architecture, when a single, small dataset is placed on a dedicated volume the remainder of the volume does not consume physical disk space, as it does on all competitors' machines. However, it does consume virtual address space. This reduces the number of virtual cylinders that the customer can use for other virtual devices. This is an important issue for customers who need a much larger virtual capacity than the machine's physical capacity either due to storing highly compressible data, defining volumes with unused space that is available for dataset expansion or using SnapShot copy.

When defining a custom volume on the IBM ESS, the only two device types it will create are a 3390 Model-3 (for cylinder counts up to 3339) and a 3390 Model-9 (for cylinder counts greater than 3339). If a customer replaces an IBM ESS with an SVA, it should be possible to duplicate the IBM ESS's device configuration on the SVA. To do this, the SVA uses the same strategy as the IBM ESS for imposing a device type when the user has specified a cylinder count. Note that the IBM ESS documentation explicitly states that even when a user selects a cylinder count that is equal to the size of some other, smaller 3390 model, the ESS will still create a 3390 Model-3 device. The IBM ESS does not support 3380 devices. It only supports 3390 Model-2 and 3390 Model-3 devices in 3380 track format. Thus 3380 device support is a capability that the SVA architecture provides above and beyond the capabilities of the IBM ESS. When a user specifies the virtual cylinder count for 3380 device, the subsystem selects the largest device type that does not exceed the specified virtual cylinder count. For a cylinder count value that is smaller than the smallest standard 3380 device type, the subsystem sets the device type to the smallest standard 3380 device type.

Channel Configuration

3

Fibre channels provide connectivity to open systems. ESCON channels and ICE cards provide connectivity to z/OS, MVS, VM, or OS/390 systems.

Configuration in the Mainframe Environment

ESCON Path Definitions

- **Host Path** – defined as the communication path composed of 4 logical paths between a host and any 4 of the 16 logical control units (LCUs) or virtual control units (VCUs) to the functional devices.
- **Logical Path** – defined as the communication path between a channel and a Virtual Control Unit (VCU[3]) with capability to access any of its 256 devices.
- **Physical Path** – defined as the ESCON cable between the channel and the control unit interface (ICE3 card).

Maximum Configuration Values

- Physical paths - 32
  - Links per ICE3 card - 42
  - ICE3 cards - 8
- Host paths - 512
  - Host paths per Link - 16
  - Physical paths - 32
- Logical paths - 2048
  - Host paths - 512
  - Logical paths per host path - 4
- Logical paths per Link - 64

3. A Virtual Control Unit is also know on the host end as a Logical Control Unit. For all intents and purposes, they are the same thing.

- Host paths per Link - 16
- Logical paths per host path - 4

- LCUs or VCUs - 16 (256 devices per LCU)
  - LCUs per ESCON channel - 4 (host limit - 1024 device addresses)
  - Host paths to access all 16 LCUs or 4096 devices - 4 (minimum)
- ESCON channels per LCU - 8 (host limit)

  A maximum of 512 host paths to the SVA can be established by utilizing ESCON and FICON directors or shared channels on an MIF (multiple image facility) - capable CPC (central processor complex). A shared channel uses 1 host path for each LPAR defined to the channel.

Channel Paths And Mainframe Connectivity

A maximum of eight channel paths can be configured to a specific device or a LCU from one host operating system. This limit is imposed by the Central Processor (CP) and the Channel Subsystem (CSS) – not the SVA.

The SVA could be configured with the following method to access all 4096 devices (16 LCUs):

- Up to a collection of eight paths accesses four sets of 256 devices, four logical control units. Each collection could be different CHPIDs, allowing up to 32 CHPIDs.

Note: The 512 host path total can NOT be exceeded. The SVA eight ICE3 cards, quad-link, configuration mandates a maximum of sixteen host paths per link. THIS IS A DIFFERENT RESTRICTION THAN THE 8-PATH GROUP TO A DEVICE.

To elaborate on the 8-path connection to a specific device:

- The 8-path group to a device from a system image, either in LPAR Mode or a CP in Basic Mode, is a host constraint. The 8-path group is valid for recent processor models. Earlier models support only 4-path groups.
- The channel subsystem control blocks and pathing masks for each device control the number of channel definitions. Without major architectural modifications in the mainframe hardware and software platforms, the 8-path group limit reigns.
- From a system image to a logical control unit, define up to 8 CHPIDs (channel path identifiers). If a processor running in LPAR mode has ESCON multiple image facility (EMIF) capability, its

logical partitions can share channel paths. EMIF and/or ESCON directors is used in the configuration to reduce the number of physical connections to the SVA subsystem. However, be aware of the host 8-path group limit from each operating system and the maximum of 512 logical host paths from all systems to the SVA.

- The machine limits and rules is found in Appendix F of the IBM manual *IOCP User's Guide and ESCON Channel-to-Channel Reference* Document Number GC38-0401.

Note: A maximum of 1024 devices (four logical control units) is configured to an ESCON channel.

Channel Configuration

Dependent upon the number of ICE and ICF cards installed,

- A maximum of **32** physical **ESCON and fibre channels** is directly connected to a subsystem.
- Physical channel connectivity is not required to each ICE or ICF card.
- Refer to “ICE/ICF Card Combinations” on page 175 for the supported configuration of card placement in the slots.
- Ensure that ALL channels on a single ICE card link are configured for the same BASE, RANGE, and BFDID parameters on the SVA.
- The default parameters (even if all devices are defined) are:
  - Base 00
  - Range 4096
  - BFDID 00

Channel Attachment Options

The SVA supports:

- ESCON-compatible FX channels with data transfer rates of 20 megabytes per second. An SVA attaches to all IBM S/390 - equivalent data streaming channel architectures.
- Fibre channel.

Channel transfers are never synchronized with array transfers. Therefore, the timing problems associated with channel extenders for traditional disk storage systems are not a consideration for a subsystem.

Channel Attachment Considerations

An ICE3 card facilitates ESCON channel attachment to an SVA subsystem. Each ICE3 card is functionally partitioned into two ESCON

Channel Adapters (ECAs). Each ECA provides two ESCON channel attachments for a total of four attachments per ICE3 card. The top two links are controlled by ECA0 and the bottom two links are controlled by ECA1.

Caution: Performance Impact

-ICE2 or ICE3 cards cannot have 2 paths from the same LPAR connected to the same physical link (via ESCON director) if both paths are using the same address range.

- **ICE3 cards cannot have 2 paths from the same LPAR connected to a single ECA if both paths are using the same address range.**
- **The ICE3 card does NOT support PPRC primary or secondary connections. If an Establish PPRC Path (CESTPATH) command is attempted, the command are rejected with a fault symptom code of 0B3D. ICE2 cards must be used for Power PPRC connectivity. ICE2 and ICE3 cards is installed in the same V2X2.**

For optimal performance:

1. *Balancing the ICE/ICF cards between clusters is always a good idea*. If there's an odd number of these interface cards:

   A. **For new SVA installations**, all you can do is pick a cluster and plug in the ICE card; at a later time, the SE and/or service representative might have look at various channel activity reports to determine which channels/cards are least busy, consult with the customer, and make a cluster change with the odd ICE/ICF card.

   B. **For SVA upgrades or SVA replacements**, consult the earlier SVA's channel loads to get an indication of where the odd ICE/ICF card should be placed.

2. Split the channels for a logical control unit (LCU) between cluster 0 and cluster 1 from a specific operating system. Assigning ICE card ports to an LCU is an IOGEN configuration item.

3. An ICE3 card is logically a 4x2 switch where one Channel Interface Processor (CIP) controlling two interfaces of a link pair. If only two interfaces are being used, for optimal performance the two cables should be plugged into port 0 and port 2, or port 1 and port 3[4] (counting from the top).

4. Ports are also known as links, so this would be link 0 and link 2 or link 1 and link 3.

Host Path Connection Reduction

The first consideration when operating the SVA in the co-located environment is the reduction in the number of ICE cards available to a mainframe host. With the allocation of ICF cards to the open systems host, there is a reduction of the number of ICE cards available to the mainframe.

The ICE3 card allows 16 host paths per ICE card port.

The following figure depicts a fully populated SVA with ESCON attach to an OS/390 or z/OS host. It depicts eight host paths per link on the ICE card (there are four links on an ICE card) and all eight ice cards installed (8 ICE cards X 4 links per card X 16 hosts paths per ICE card link = 512 host paths)

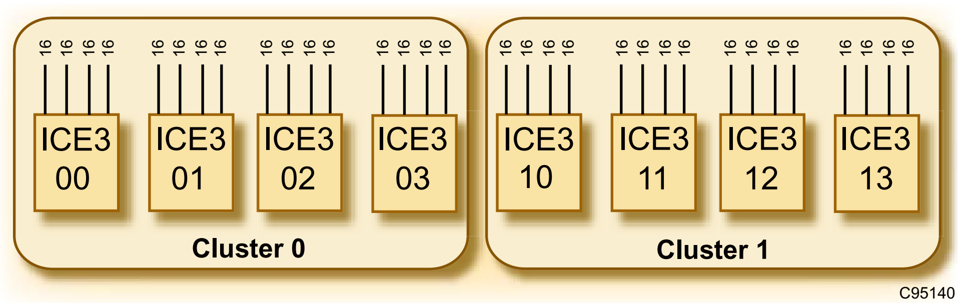

Figure 8 A Fully Populated SVA with ESCON Connectivity Only

The following figure depicts a fully populated SVA with ESCON attach to an OS/390 or z/OS on six of the ICE cards and a fibre host attach with two ICF cards. The ESCON host paths now available would be the number of ICE cards, multiplied by the number of ports used (four), multiplied by the number of host paths available on each link (sixteen), or 6 X 4 X 16 = 384.

Note: The ICF cards are NOT in the same cluster as a performance and access reliability consideration.

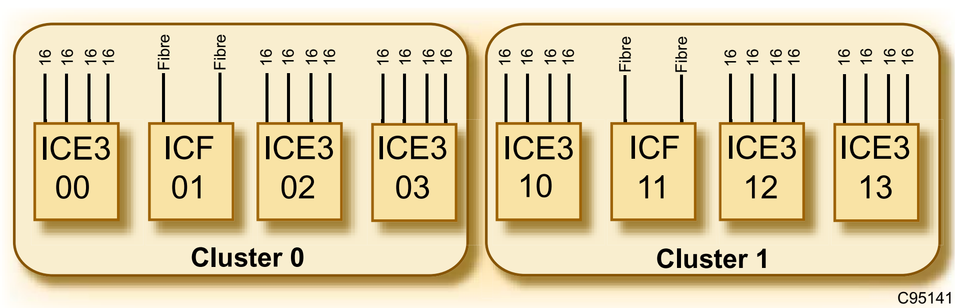

Figure 9 A SVA with Two ICF Cards Used For A Fibre Connection

ESCON Channels

The ICE3 cards provide ESCON interfaces on the V2X2. A maximum of eight ICE3 cards is installed. The installation of eight ICE3 cards allows up to 32 physical connections.

The ICE2 card is only used to provide PPRC physical connectivity between SVAs with two links per card. The recommended maximum number of ICE2 cards for PPRC is four (bi-directional mode).

In addition to ESCON channel connectivity, consider availability, redundancy, and performance to determine the number of ICE3 cards. A minimal configuration of 4 ICE3 cards is recommended. Up to 8 ICE3 cards is recommended in a multiple LPAR environment with high performance requirements to access up to sixteen internal data paths and exploit the capability of the SVA.

The SVA In An Open Systems Environment

4

Fibre Channels

Fibre channel cards include two ports per card.

Note: The SVA allows for attachment of up to 256 LUNs (LUN numbers 0 through 255) on each fibre connection. However, the actual number is often much less than that as per restrictions imposed by the host system and the host bus adapter. Some software work-arounds might overcome some of the before mentioned restrictions.

Fibre Channel Performance And Configuration Options

Fabric or PLDA Operation

The SVA supports fabric operation in addition to Private Loop Direct Attach (PLDA). This feature can be selected on a per Fibre Channel (FC) port basis. Refer to the section on “Requested Loop Id. (RLid)” on page 62 for details on how to set the operating mode.

PLDA is a Fibre Channel specification that controls which FC features are allowed, required, or prohibited. PLDA does not allow for operation through an FC switch, which is why it is called “Private”. Most switches provide a translation function that allows PLDA devices to operate through a switch. The switch also has to support a Fabric Loop port (FL port) for this to work since PLDA is for loop devices only. There are restrictions on what can be done when using the PLDA translative function of these switches.

Fabric operation requires support for several additional link services beyond those required in PLDA. This includes the SVA doing Fabric Logon. While not required for fabric, selecting fabric operation on the SVA also makes the ICF2 port operate in point-point topology (N port) instead of Arbitrated Loop topology (Nl port). The HBA should also be set for fabric operation. This is typically done by setting the HBA topology to point-point (PT-PT). The SVA does not support the ACK0, so this option should not be set for the HBA. Also the SVA does not support point-point direct connection with the HBA. You must have a switch when the SVA is set for Fabric operation.

The use of HBA binding is strongly recommended when operating the SVA with a switch. In most circumstances WWNN or Port Address Binding should be used. WWPN binding is typically not recommended since the WWPN is stored on the ICF2 Card. If WWPN binding is used and the ICF2 card needs to be replaced, the binding options needs to be changed and a host reboot is required. Refer to the appropriate HBA documentation for specifics on setting binding options. Note that HP-UX currently uses a hard address model for fabric operation and does not use binding options.

Requested Loop Id. (RLid)

The Requested Loop ID (RLid) is used to select the operating mode. The loop address on a FC port is called the RLid because the address is actually requested or negotiated with the host system. See the following table.

Table 13 Requested Loop ID

Rlid	FC Topology	Loop ID	Operation	Address Method
00 - 7D	Aribtrated Loop	00 - 7D	PLDA	PDA
7E	PT to PT	NA	Fabric	PDA
80 - FD	Aribtrated Loop	00 - 7D	PLDA	VSA
FE	PT to PT	NA	Fabric	VSA
FF	Aribtrated Loop	Soft	PLDA	PDA

Note:

1. Volume Set Addressing (VSA) is currently only used on HP-UX, all other hosts should use Peripheral Device Addressing (PDA)
2. Rlids 7E and FE are only valid for ICF2 cards and require E02.02.13.00 microcode or higher.
3. 7F is not used.

If you set the RLid to a value (the recommended method) then you are setting the SVA to specifying “Hard Addressing.” Now you have assigned the SVA a specific loop address and when the ICF card in the SVA negotiates with the host system for this address, if it cannot get it, it does not participate. If you leave the RLid set to FF then “soft addressing" is used - the ICF card in the SVA negotiates and uses whatever address is it is assigned by the host.

The ALid field on the Local Operator Panel (LOP) screen CD14 indicates the Actual Loop Id. If LOP screen CD14 indicates FF, no address has been negotiated - loop initialization has not occurred or the address requested could not be obtained. It indicates FF if the cable is not plugged in or some other device has the requested address. ALid should not be confused with the Arbitrated Loop Physical Address.

"Hard addressing" is recommended as soft addressing allows the device address to change when a device or controller joins or leaves the loop. This can cause problems with some hosts.

The RLid is a human translated version of the Arbitrated Loop Physical Address (ALPA) - and should not be confused with it. The RLid can be either decimal or hexadecimal. Sun Microsystems uses the hexadecimal version of this address when entering numbers on the Local Operator Panel. While the RLid numbers are sequential, the ALPA numbers are not (see Table 14).

Sun Microsystems recommends using a low RLid number. One is the suggested number. Zero or 7D is often used by host bus adapters and should be avoided.

Table 14 Fibre Channel Arbitrated Loop Physical Address Values

ALPA	RLid (decimal)	RLid (Hex)	ALPA	RLid (decimal)	RLid (Hex)	ALPA	RLid (decimal)	RLid (Hex)
Lowest Priority								
EF	0	0	A3	43	2B	4D	86	56
E8	1	1	9F	44	2C	4C	87	57
E4	2	2	9E	45	2D	4B	88	58
E2	3	3	9D	46	2E	4A	89	59
E1	4	4	9B	47	2F	49	90	5A
E0	5	5	98	48	30	47	91	5B
DC	6	6	97	49	31	46	92	5C
DA	7	7	90	50	32	45	93	5D
D9	8	8	8F	51	33	43	94	5E
D6	9	9	88	52	34	3C	95	5F

Table 14 Fibre Channel Arbitrated Loop Physical Address Values (Continued)

ALPA	RLid (decimal)	RLid (Hex)	ALPA	RLid (decimal)	RLid (Hex)	ALPA	RLid (decimal)	RLid (Hex)
D5	10	A	84	53	35	3A	96	60
D4	11	B	82	54	36	39	97	61
D3	12	C	81	55	37	36	98	62
D2	13	D	80	56	38	35	99	63
D1	14	E	7C	57	39	34	100	64
CE	15	F	7A	58	3A	33	101	65
CD	16	10	79	59	3B	32	102	66
CC	17	11	76	60	3C	31	103	67
CB	18	12	75	61	3D	2E	104	68
CA	19	13	74	62	3E	2D	105	69
C9	20	14	73	63	3F	2C	106	6A
C7	21	15	72	64	40	2B	107	6B
C6	22	16	71	65	41	2A	108	6C
C5	23	17	6E	66	42	29	109	6D
C3	24	18	6D	67	43	27	110	6E
BC	25	19	6C	68	44	26	111	6F
BA	26	1A	6B	69	45	25	112	70
B9	27	1B	6A	70	46	23	113	71
B6	28	1C	69	71	47	1F	114	72
B5	29	1D	67	72	48	1E	115	73
B4	30	1E	66	73	49	1D	116	74
B3	31	1F	65	74	4A	1B	117	75
B2	32	20	63	75	4B	18	118	76
B1	33	21	5C	76	4C	17	119	77
AE	34	22	5A	77	4D	10	120	78
AD	35	23	59	78	4E	0F	121	79
AC	36	24	56	79	4F	08	122	7A

Table 14 Fibre Channel Arbitrated Loop Physical Address Values (Continued)

ALPA	RLid (decimal)	RLid (Hex)	ALPA	RLid (decimal)	RLid (Hex)	ALPA	RLid (decimal)	RLid (Hex)
AB	37	25	55	80	50	04	123	7B
AA	38	26	54	81	51	02	124	7C
A9	39	27	53	82	52	01	125	7D
A7	40	28	52	83	53	N/A	FABRIC	7E
A6	41	29	51	84	54	Highest Priority		
A5	42	2A	4E	85	55			

Improving Performance

Where there is a mismatch in the I/Os per second capability of the HBA and the fibre connection port in the SVA, there is a tactical advantage to using a fibre channel hub as shown in Figure 10 on page 70. This allows the host processor to initiate a larger number of I/Os per second against the SVA, resulting in a higher data transfer rate between the HBA and the SVA.

In tests using two connections as shown in Figure 10, performance doubled. Using three connections did not triple the performance, but still produced a significant performance improvement. More than three connections from the SVA to the fibre channel hub showed no increase in performance - depending on the circumstances. Actual performance varies with actual configuration and the jobs being run against the SVA.

Attaching More LUNs Than Normally Allowed To A HBA

HP Volume Addressing Method

HP-UX provides two LUN address methods for Fibre Channel attached devices:

- Peripheral Device Addressing (PDA)
- Volume Set Addressing (VSA)

The SVA Fibre Channel Interface Code (FIP) currently only supports the “Peripheral Device Address” method. “Peripheral Device Addressing” is supported on Solaris, AIX, NT4.0, Win2K, and Novell as well as by HP-UX.

HP-UX limits the number of LUNs per fiber channel port to eight when using the “Peripheral Device Addressing” method. Using the HP

"Volume Addressing" method, the SVA supports up to 256 LUNs per fiber channel port when used with HP-UX 11.0.

Since the "Volume Address" method is only supported on HP-UX, "Peripheral Device Addressing" support is still required. A method of selecting which addressing method to use is required. This is accomplished using previously restricted values (8D – FE) for the Requested Loop ID of the existing "Configure Interface" menu on the LOP.

Note: Changes to the Configure Interface menu requires an IML of the SVA.

When using VSA, HP-UX maps each group of 8 SVA LUNs as a different target as seen in the special character and block device files using the cxtydz nomenclature, where x is the controller instance, y is the target instance, and z is the device instance. These special files are actual just links to the hardware path as shown by the "ioscan" command. The following table lists the mapping of the SVA defined LUN to the HP-UX target (ty) and Device (dz).

***Caution: Potential Data Loss - I*f changing the address method on HP-UX for devices that have existing data on them that needs to be preserved, the user should read and understand chapter 4 section titled "Moving Disks to a Different Address" HPUX for Peripherals manual B2355-90698 before changing the address method. Failure to do this will result in the loss of user data on these devices.**

Table 15 SVA LUN to HP-UX Target Mapping

SVA LUN	HP Target	HP Device	SVA LUN	HP Target	HP Device	SVA LUN	HP Target	HP Device	SVA LUN	HP Target	HP Device
0	0	0	64	8	0	128	0	0	192	8	0
1	0	1	65	8	1	129	0	1	193	8	1
2	0	2	66	8	2	130	0	2	194	8	2
3	0	3	67	8	3	131	0	3	195	8	3
4	0	4	68	8	4	132	0	4	196	8	4
5	0	5	69	8	5	133	0	5	197	8	5
6	0	6	70	8	6	134	0	6	198	8	6
7	0	7	71	8	7	135	0	7	199	8	7
8	1	0	72	9	0	136	1	0	200	9	0
9	1	1	73	9	1	137	1	1	201	9	1
10	1	2	74	9	2	138	1	2	202	9	2

Table 15 SVA LUN to HP-UX Target Mapping (Continued)

SVA LUN	HP Target	HP Device	SVA LUN	HP Target	HP Device	SVA LUN	HP Target	HP Device	SVA LUN	HP Target	HP Device
11	1	3	75	9	3	139	1	3	203	9	3
12	1	4	76	9	4	140	1	4	204	9	4
13	1	5	77	9	5	141	1	5	205	9	5
14	1	6	78	9	6	142	1	6	206	9	6
15	1	7	79	9	7	143	1	7	207	9	7
16	2	0	80	10	0	144	2	0	208	10	0
17	2	1	81	10	1	145	2	1	209	10	1
18	2	2	82	10	2	146	2	2	210	10	2
19	2	3	83	10	3	147	2	3	211	10	3
20	2	4	84	10	4	148	2	4	212	10	4
21	2	5	85	10	5	149	2	5	213	10	5
22	2	6	86	10	6	150	2	6	214	10	6
23	2	7	87	10	7	151	2	7	215	10	7
24	3	0	88	11	0	152	3	0	216	11	0
25	3	1	89	11	1	153	3	1	217	11	1
26	3	2	90	11	2	154	3	2	218	11	2
27	3	3	91	11	3	155	3	3	219	11	3
28	3	4	92	11	4	156	3	4	220	11	4
29	3	5	93	11	5	157	3	5	221	11	5
30	3	6	94	11	6	158	3	6	222	11	6
31	3	7	95	11	7	159	3	7	223	11	7
32	4	0	96	12	0	160	4	0	224	12	0
33	4	1	97	12	1	161	4	1	225	12	1
34	4	2	98	12	2	162	4	2	226	12	2
35	4	3	99	12	3	163	4	3	227	12	3
36	4	4	100	12	4	164	4	4	228	12	4
37	4	5	101	12	5	165	4	5	229	12	5
38	4	6	102	12	6	166	4	6	230	12	6
39	4	7	103	12	7	167	4	7	231	12	7
40	5	0	104	13	0	168	5	0	232	13	0

Table 15 SVA LUN to HP-UX Target Mapping (Continued)

SVA LUN	HP Target	HP Device	SVA LUN	HP Target	HP Device	SVA LUN	HP Target	HP Device	SVA LUN	HP Target	HP Device
41	5	1	105	13	1	169	5	1	233	13	1
42	5	2	106	13	2	170	5	2	234	13	2
43	5	3	107	13	3	171	5	3	235	13	3
44	5	4	108	13	4	172	5	4	236	13	4
45	5	5	109	13	5	173	5	5	237	13	5
46	5	6	110	13	6	174	5	6	238	13	6
47	5	7	111	13	7	175	5	7	239	13	7
48	6	0	112	14	0	176	6	0	240	14	0
49	6	1	113	14	1	177	6	1	241	14	1
50	6	2	114	14	2	178	6	2	242	14	2
51	6	3	115	14	3	179	6	3	243	14	3
52	6	4	116	14	4	180	6	4	244	14	4
53	6	5	117	14	5	181	6	5	245	14	5
54	6	6	118	14	6	182	6	6	246	14	6
55	6	7	119	14	7	183	6	7	247	14	7
56	7	0	120	15	0	184	7	0	248	15	0
57	7	1	121	15	1	185	7	1	249	15	1
58	7	2	122	15	2	186	7	2	250	15	2
59	7	3	123	15	3	187	7	3	251	15	3
60	7	4	124	15	4	188	7	4	252	15	4
61	7	5	125	15	5	189	7	5	253	15	5
62	7	6	126	15	6	190	7	6	254	15	6
63	7	7	127	15	7	191	7	7	255	15	7

Note:

1. The controller instance changes at SVA LUN 128 and the target and device numbers restarts at t0d0.
2. The 8 SVA Luns within the HP target must be contiguous starting with HP device 0 in order to be recognized by HP-UX “e.g. If you defined SVA LUNs 8, 9, 11,12 and 16, HP-UX would see t1d0 (LUN 8), t1d1 (LUN 9), and t2d0 (LUN 16).” LUNs 11 and 12 would not be seen since LUN 10 is missing. LUN 16 is seen because it is in the next target.
3. In E02.01.xx code you must also have a LUN 0 defined to see any LUNS when using VSA on HP-UX This restriction is removed in the E02.02.xx Code release.

HP Alternate Method Using A Brocade Switch

Figure 10 on page 70, indicates more LUNs than usually allowed attached to a single HP-UX fibre channel host bus adapter[5]. Figure 10 includes a Brocade switch, but this is accomplished with a fibre channel hub just as easily. This larger than usual number of LUNs is accomplished by:

1. Using two distinct domains (1 and 2) in the SVA. The target is always 0 and the LUN numbers are 0-7.

Notes:

   1. For HP-UX HBAs, the loop Id number can NOT be zero.
   2. While the following figure indicates eight LUNs in each domain, having an equal number of LUNs in each domain is not a requirement.

2. This arrangement uses two SVA fibre channel ports assigned to these domains (one per port).

The HP-UX HBAs and some Solaris HBAs take the arbitrated loop Id as the target Id.

The LUN numbers are the same in each of the domains AND in the HP as shown in Figure 10.

5. For HP-UX, eight is the maximum number of LUNs. The maximum number of LUNs varies for Sun Solaris, Windows NT, and AIX depending on the HBA and the software version being used.

***Caution: Performance Impact* - Due to problems with most HP host bus adapters, you cannot have more than one host bus adapter attached on a "loop."**

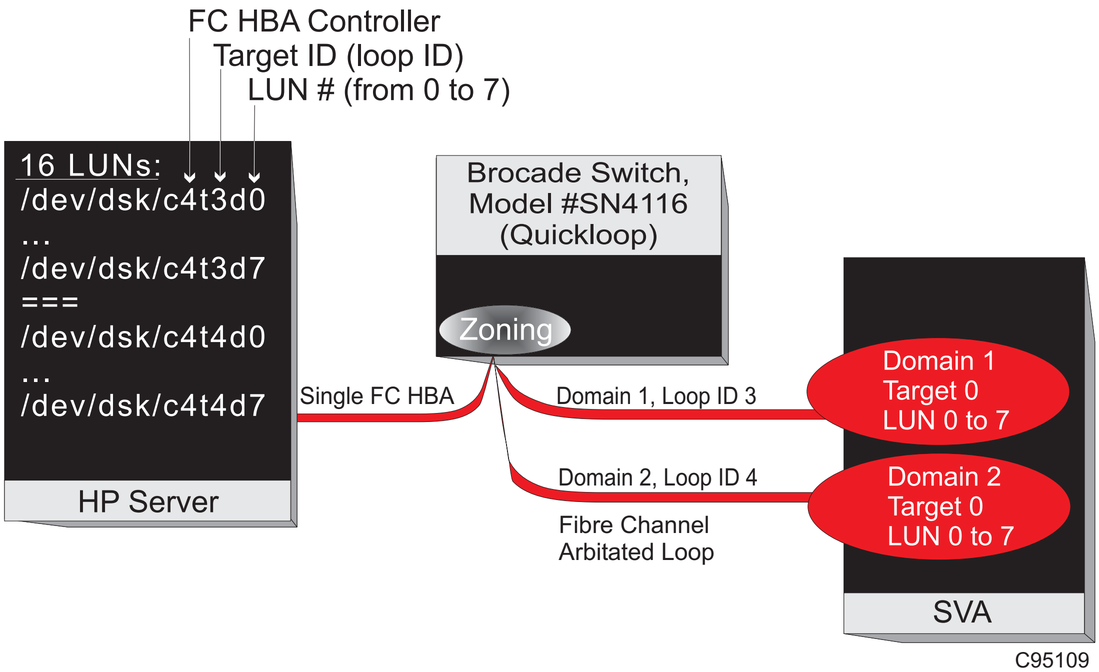

Figure 10 HP-UX Fibre Channel Attach For More Than Eight LUNs

Note: HP host bus adapter A5158 supports "fabric" so a fibre channel hub or Quick loop is no longer necessary.

Supported Host Bus Adapters

For a complete list of the host bus adapters supported by the V2X2, visit the Sun Microsystems Web site and look in the Customer Resource Center.

Supported Cluster Configurations

Table 16 Validated Cluster Configurations

Host	Software	Details	Min. Microcode Level
RS6000	HACMP	Emulex HBAs/IBM FC6227 AIX 4.2.4 & 4.3.2	e02.01.14
HP	MC/ServiceGuard	HP HBAs HPUX 11.0	e02.00.33
NT	MSCS	Qlogic HBAs/Emulex HBAs NT 4.0 SP6	e02.01.14
Windows 2000	MSCS	Qlogic HBAs/Emulex HBAs Advanced Server 5.00.2195	e02.01.14

Windows 2000 SVA Performance and Block Size

While the host operating system can address disk storage in block sizes as low as 512 bytes, it predominately addresses the data on larger boundaries such as 4 KB or 8 KB to better match the file system or data base in use. Accommodating this, the SVA allows the selection of various block sizes for internal data storage from 512 bytes through 16,384 bytes. (Also see Table 18, “Data Utilization And Block Size,” on page 82.)

Note: Larger block sizes are not only more efficient for storage and compression, but increase performance.

While Windows 2000 and NTFS allocates space in 4KB units, these units does not start or end on a 4KB boundary. Windows 2000 uses the first 32,256 bytes (blocks 0 through 62) of the disk for metadata. This causes the first user partition to start at block 63 on the disk which is not a 4KB boundary. Misaligned accesses are significantly slower than aligned accesses due to the internal overhead to process them. This is most noticeable on small write operations.

Note: Even in the worst case situation, a misaligned 4KB write to a 4KB block is faster than an aligned 4KB write to a 512 block size device.

This misalignment problem did not occur on NT4.0 and devices that were created on NT and then moved to Windows 2000 does not exhibit the problem, unless all the partitions are removed and recreated on

Windows 2000. NT4.0 metadata used 16,384 bytes (blocks 0 through 31) on the disk, so the first user partition would start at block 32 which is aligned with all SVA block sizes.

Setting I/O Timers

Use the following information to change the time out values (to 270 seconds) on the supported interoperative system platforms. These changes enhances the recovery of processing jobs and or data in the event of a check 0 recovery on the SVA.

You need a system administrator with root login access. All commands to change the time out values to a minimum value of 270 seconds must be made from the root. If any value already exists which exceeds the 270 second minimum limit do not change these values.

HP-UX

There are two commands that can be used to view and change IO time-out parameters on a per device basis. These commands are "pvchange" and "pvdisplay". See the system manual pages for more information on each of these commands. An example for "pvdisplay" would be:

```
pvdisplay /dev/dsk/c1t3d0
```

This returns an IO_time value in seconds which could also say default. To change this value to 270 seconds use the command:

```
pvchange t 270 /dev/dsk/c1t3d0
```

The device needs to be part of a volume group before "pvdisplay" and "pvchange" works. Once the devices are changed to the 270 time out value the host system needs to have a REBOOT to download the changes to the drivers.

RS6000 - AIX

The changes can be made via "smit", via a command line entry, or script file using "chdev" to change the time out value and "lsattr" to display the attributes of the devices. The changes can be made upon initial set up to the raw devices or with a system REBOOT to devices already configured to a volume group.

If queuing is not turned on request that queuing be turned on to at lease the minimum levels in the following list.

SMIT

Devices

Fixed Disk

Change / Show Characteristics of a Disk

Select the hdisk # to change:

Queue Depth[1]

Queuing TYPE[simple]

Use QERR bit[no]

Device CLEARS its Queue on error[no]

READ/WRITE time out value[270]

START UNIT time out value[270]

REASSIGN time out value[270]

Note: If the time-out values are already greater than 270 do not change these values. Also if queuing values are set higher than those listed here no queue value change is needed.

If this is a raw disk you do not need to make any further changes. If this disk is already established in a volume group, make the following changes and note that the changes are not applied until the host system is rebooted.

```
Apply change to DATABASE only        yes
```

Via Command Line or you can build a script to include several command lines:

Raw Disk:

```
chdev l 'hdisk#' -a queue_depth='16' -a q_type='simple' -a
rw_timeout='270' -a start_timeout='270' -a reassign_to='270' -a
q_err='no'
```

Established volume group (REBOOT required):

```
chdev l 'hdisk#' -a queue_depth='16' -a q_type='simple' -a
rw_timeout='270' -a start_timeout='270' -a reassign_to='270' -a
q_err='no' P
```

SUN

The changes on the Sun are made by editing the '/etc/system' file. Add or modify the following line within this file. Again the value of 270 is a minimum value. If a larger value already exists use the larger value.

```
set sd:sd_io_time=270
```

A system REBOOT must occur for this value to take effect. It is suggested that you reboot your system with the following parameters:

```
boot r
```

Also see “Missing-Interrupt Handler (MIH)” on page 45.

Windows NT/Windows 2000

To increase the disk I/O time-out value on NT perform the following:

1. Click on Start on the task bar. Select run and enter regedt32 in the popup window. Click OK to invoke registry editor.
2. Select HKEY_LOCAL_MACHINE->SYSTEM->CurrentControlSet->Services->Disk.
3. Under Edit, select Add Value. Enter a Value Name of “TimeOutValue” and select Data Type of REG_DWORD.
4. The DWORD editor window pops up. Select a Radix of Dec and enter a value of 270. This value is in seconds. The registry entry should be TimeOutValue: REG_DWORD: 0x10e.

Linux - Redhat Versions 6.2 through 7.2

To change timeout value for SCSI R/W operations:

1. **For Redhat version 7.1,** either install the patch that came with the SVAA installation disk, or in module sd.c change the value for SD_TIMEOUT from 30 (default) to 270. It should look like this:

   SD_TIMEOUT (270 * HZ)

   Dir (/usr/src/linux-2.2.16/drivers/scsi)
2. **For Redhat version 6.2 only,** in scsi.c or scsi_scan.c find #ifdef CONFIG_SCSI_MULTI_LUN. Under this statement max_scsi_luns has a default value of 8 (this could change in a later release). Change this to a value of 128. The #else value for max_scsi_luns is left at 1.
3. **Change Makfile in linux 2.2.xx** directory so that the extended revision value is a new or added value. This effects the library name that is stored in the modules /lib/modules directory. If you don’t change the make file the old lib is overwritten.
4. To rebuild the kernel:

   A. Form dir usr/src/linux

   B. Make menuconfig

   C. Change multiple scsi luns to * and scsi low level driver qlogic or the type you have to * and remove others.

   D. Make dep

   E. Make clean

   F. Make –j2 bzImage

G. Make modules

H. Make modules_install

I. Copy the bzImage file and the System.map files to the /boot directory.

J. Save a copy of the original lilo.conf file by doing:

```
cp /etc/lilo.conf /etc/lilo.conf.save
```

K. Edit lilo.conf and add new image (bzImage) by copying the 1st image group of text and changing the 1st line (image=) to point to the bzImage file. Then change the 2nd image group label line to linux-smp.

L. Remove linux. It is a link to the System.mapxxxx file for the default kernel. Replace it with a link file to your map (ex. ln –s System.map-2.2.14-5.fc).

Netware 5.1 & 6.0

1. Go to the Web site[6]:

   ```
   http://support.novell.com/servlet/filefinder
   ```

2. Enter *SCSIHD.CDM* in the search box.
3. Click on “SEARCH.”
4. Find *SCSIHD.CDM Version 3.01* associated with package SCSIHDFT.exe in the list of search matches.
5. Download SCSIHDFT.exe.
6. Run SCSIHDFT.EXE to unzip the package.
7. Find and open the file SCSIHDFT.TXT. (It gives installation instructions for SCSIHD.CDM.)
8. Install SCSIHD.CDM as per the previous instructions.
9. After installing SCSIHD.CDM, edit the STARTUP.NCF file. This file is located in the directory that contains your new SCSIHD.CDM file (usually C:\NWSERVER).
10. On the “LOAD SCSIHD.CDM” line, add the parameter:

    ```
    /retrydelay=270
    ```

Note: Make sure there is a space before the /

11. Save the STARTUP.NCF file.
12. Restart your Server.

6. This Web site link was valid at the time of publication – however, Web sites change frequently.

Installing Windows (Win2K or NT) Onto An SVA

First, make sure the host bus adapter, driver, firmware, bios, etc. supports booting over fibre. If all that is the case, then:

1. Create an SVA LUN and set it's LUN address to X.0.0, where X is the domain. Windows ONLY installs to LUN 0 - it does not give you an option.

Notes:

   1. Make sure to have the appropriate HBA driver on a floppy before you go through this procedure because sometimes the installation CD won't have the ones you need.
   2. Make sure you have the drivers for any internal RAID devices or motherboards
   3. Windows only boots from the 0 LUN. Therefore, there can only be one boot LUN per domain.

2. Boot from the Windows installation CD.

Caution: Performance Impact (Optional) - **As a precaution, disable the host's hard drive for this process. Sometimes the Windows installation process writes undesired data on it with disastrous consequences.**

3. When given the option to choose the drive you want to install Windows on, choose the SVA LUN you created for this purpose.
4. When the installation is complete, reboot the host and go into the system setup to configure the HBA as your primary choice over the local hard drive for boot.
5. Configure the HBA settings to support booting. Configuring the HBA settings to support booting is done through the firmware, driver or utility of the HBA. Each HBA manufacturer bundles these operating instructions with the firmware/driver/utility. Nothing has to be done to the physical SVA itself.

Host Recognition Of New Devices

It is beyond the scope of this document to explain all of the different host, host bus adapter, and software command lines required to get a host to recognize newly attached devices. See the host system manuals for those commands.

Modifying OS/390 VTOCs for Interoperative Use

Earlier versions of the Iceberg product line modified OS/390 VTOCs for interoperative or open systems use by means of a JCL job. This operation is currently handled in the V2X2 by the SVAA product. This VTOC modification is performed when the storage device is initialized and is a prerequisite to loading the data. See the *SVAA for OS/390 Configuration and Administration* or *SVAA for VM Configuration and Administration*. Located in chapter eight, see the heading **Initialize Device subcommand** or the **Define Device subcommand**.

Record Format Assist Feature (RFA)

Fibre channel data transfers use the Record Format Assist Feature (RFA). This standard feature provides the following features:

- Improved compression due to longer run lengths:
  - compression done on larger blocks
  - from +53% to 325% increase depending on the data type
  - physical space increased
- Dramatically improves performance for virtual disk in open systems environments (see “Record Format Assist in the Open Systems Environment” on page 81).
- User tailored parameter; a choice of record block sizes of 512 bytes, 2KB, 4KB, 8KB, and 16KB
- Increases usable space with larger block sizes (also see “Track Utilization/Virtual Capacity” on page 80):
  - mod-3 user capacity rises from 1.2GB to 2.65GB
  - mod-9 user capacity rises from 3.7GB to 7.4GB
  - Usable virtual space is increased
- Increases the maximum amount of user cache available to the customer because fewer directory extents are required.
- Supports multiple virtual devices per LUN.
- A two times performance improvement for write and read hit accesses.

The effect of RFA can be seen in the following figure. Without RFA, there are a lot of inter-block gaps (1), compression run on small block which does not optimize physical space (2), and small records staged one at a time in cache (3). The following figure also indicates that with

RFA, there is no lavish waste of virtual space, increased use of physical space, and performance is dramatically improved.

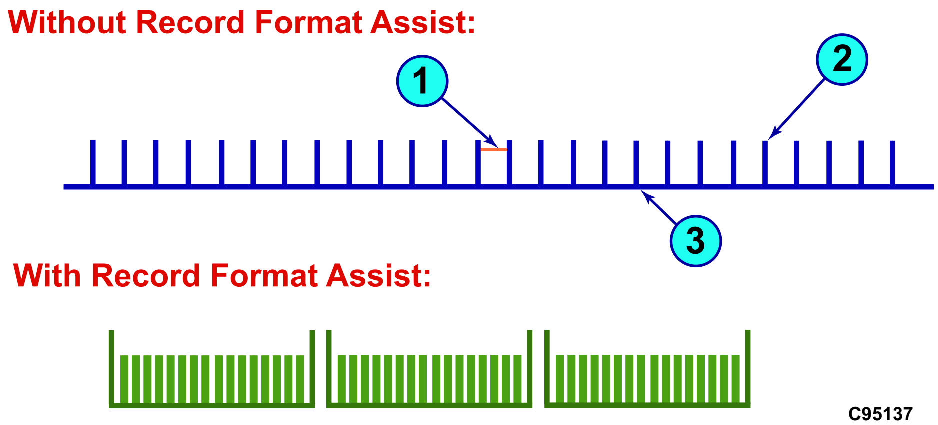

Figure 11 The Effects Of RFA

The following table identifies the components and how the compressability of the data affects the component.

The net effect of placing a high percentage of un-compressible data on an SVA is negligible until the various resources such as cache, NVS, and bandwidth are saturated. When these points are reached, for cache and NVS, performance is impacted as if you had less cache and for bandwidth it appears that the bandwidth, e.g., back-end performance is affected. The major concern is from a physical capacity perspective where the logical capacity is reduced as a function of the compression ratio.

Table 17 Components and Data Compression

Component	Description of Effect
Channels	Compression does not affect channel performance except as described in the following sections[a] because the SVA always sends and receives un-compressed data. The Interface Card takes ESCON data at 18 MB/s and feeds an internal data path that runs at 12.5 MB/s, requiring a 1.5:1 compression ratio to maintain channel speed. This is usually masked by the full track buffer.
Cache and NVS	Un-compressible data uses more cache and NVS. If the cache and NVS limits were exceeded, the subsystem would perform as if it had less cache. As long as the cache and NVS limits are not exceeded, performance would be similar to a system with compressible data.

Table 17 Components and Data Compression (Continued)

Component	Description of Effect
Back End Storage	Un-compressible data consumes more bandwidth between the controller and the back end (disk arrays), which could limit the maximum stage and destage rates. Performance should be close to nominal.
Disk Arrays	Un-compressible data consumes more space. However, the compression rates seen by the customer reflect both data compression and compaction, which eliminates inter record gaps (IRG). The smaller the blocks the more IRGs and the more benefit from compaction. On mainframe data, most customers still enjoy an approximate 2 to 1 net effect from compression and compaction on data that is otherwise un-compressible. Customers who want to selectively place data on subsystems might experience significantly greater than 4 to 1 compression ratios by choosing more compressible data for the SVA. Compression information can be provided by SVAA, a utility program that can be used to mange the SVA.

a. There are currently two exceptions to this method of operation. First, for all PPRC activity between two SVA subsystems, the data is transferred in a compressed format. Second, when using HSDM, the data is transferred between the host and the SVA in a compressed format.

Read Data Transfers[7]

For read operations that are aligned to a valid block size, the data transfer proceeds in a normal fashion. The transfer is handled by the CIP decoding the logical block address sent by the host and mapping it to a configured block size. For unaligned accesses (512 byte accesses that don't align on a block boundary), a new command is issued by the CIP to the IUP. This command has additional information that includes number of 512 blocks to be transferred and the starting block address. The record that contains the unaligned data is read into the ICE card's full track buffer. Please note that the compressed data Cyclic Redundancy Check (CRC) is checked at the ICE card in the decompression logic. The data is then written un-compressed to temporary space in cache. Accumulated CRC values are saved for each block, including the original CRC bytes. The accumulated CRC is used to verify the data has been stored correctly in cache. The host

7. While fibre channel attachmet supports RFA, fibre channel operation for read and write data transfers is slightly different.

data is then sent to the ICE card with CRC appended. Data integrity is ensured by this CRC as data is sent to the host. After the successful transmission of data to the host the temporary cache space is freed up.

Write Data Transfers

For write data transfers that are aligned to a block size, the data transfer proceeds in a normal fashion. As with the reads, the record that contains the unaligned data is read into the ICE card's full track buffer. The data is then written un-compressed to temporary space in cache. Accumulated CRC values are saved, including the original CRC bytes. The accumulated CRC is used to verify the data has been stored correctly in Cache. Host data is transferred to cache with CRC generated from the ICE card. The updated host data along with the unchanged data is transferred from cache to the ICE card buffer. Virtual field CRC is generated and appended to the data by the sector CRC logic on the CMI card. The PERC logic on the CMI card verifies the integrity of the data that is being transmitted to the ICE card. The last step is for all of the data to be sent through the compression logic on the ICE card and written to cache as a compressed record. After the successful transmission of data to cache, the temporary write space is reclaimed.

Positive Side Effects

The new track format with larger block sizes presents a number of positive side effects:

Track Utilization/Virtual Capacity

The current implementation of storing data in 512 byte fixed sized records limits the number of records on a track. This limits a track to holding only 24 KB of customer data. The new format with 4 KB records doubles the amount of customer data per track to 48 KB. This allows base device sizes to go from 1.2 GB to 2.4 GB for mod 3 devices and from 3.5 GB to 7 GB for mod 9 devices.

Intra Track Overhead

Storing twice as much data per track reduces by half the number of track boundary crossings for large transfers. Current measurements with a channel monitor show that this intra track time can be as high as 900 μsec. The new format allows a 1MB I/O operation to save 19 msec.

Cache Bandwidth

The software transfers each record in cache as an individual entity. In addition, the cache hardware breaks up large transfers into 512 byte segments to allow the available bandwidth to be distributed among the active cache ports. Records that are less than 512 bytes compressed reduce the effective cache bandwidth because of the additional arbitration, control and status cycles. 512 byte records with a 2:1 compression ratio reduce the cache bandwidth down to a number between 120 MB/sec. and 140 MB/sec. Tracks formatted with 4K byte records with a 4:1 compression ratio allow the cache bandwidth to maintain its current throughput of 160 MB/sec.

Compression Ratio

The current data compression design within the SVA makes use of a history algorithm[8]. The greater the data set to generate history over, the better the compression ratio. At 512 byte block size it is predicted that the compression ratio is approximately 2:1. A 4 KB block size should allow the compression ratio to approach the 4:1 value that is observed with OS/390 data.

Record Format Assist in the Open Systems Environment

RFA offers several significant enhancements to the SVA implementation in the Open Systems environment.

- **Dramatically improves performance of virtual disk.** Multiple records are staged at one time in cache, thus eliminating back end overhead. Customers can tune their disk block sizes to best suit a given application or operating system. One customer performance test went from 2 MB/sec. to 15 MB/sec.
- **Increases usable space with larger block sizes.** Stores more records per track than before, and reduces inter record gaps. With a mod-3, user capacity rises from 1.2 GB to 2.56 GB, and with a mod-9, from 3.7 GB to 7.4 GB. The table on the following page demonstrates how data utilization increases with block size.
- **Improves data compression.** As compression is done on larger blocks of data, then compression effectiveness increases from 53% to 325% depending on the data type.

The SVA in the Open Systems environment initially offered a maximum of 1.26 TB of addressable space (virtual space), and about 1.1 TB of user capacity (data). As most users would use SnapShot, about half of the virtual space would be dedicated for that purpose and the

8. The History Algorithm uses the Advanced Data Compression algorighm. The Advanced Data Compression algorithm is a dynamically adaptive, lossless, and string oriented data compression algorithm.

remaining virtual space for the user would be reduced to half of the 1.26 TB or just 630 GB.

With RFA, the addressable space (virtual space) is now about 2.5 TB and the user capacity is about 1.56 TB. Now, if a user wants to use SnapShot, they have a 1.2 TB user capacity.

Table 18 Data Utilization And Block Size

Block Length (rec size byte)	Records per Track	Bytes per Track (used)	Bytes per Cylinder (used)	Mod 3 Bytes/ Device	Mod 9 Bytes/ Device
512	48	24,576	368,640	1,230	3,691
2048	20	40,960	614,400	2,050	6,153
4096	12	49,152	737,280	2,460	7,384
8192	6	49,152	737,280	2,460	7,384
16384	2	49,152	737,280	2,460	7,384

The SVA in a VM Environment

5

The SVA subsystem is fully compatible with the z/VM operating systems. However, there are some special considerations in defining and allocating a subsystem under VM.

Installing the SVA for VM

After installation, as with any other I/O device, an SVA subsystem must be defined to the VM operating system in the same manner as other subsystems. Special software considerations for defining the SVA subsystem to the VM operating system, along with a list of supported releases of VM, are provided in this section. Functional disk initialization using ICKDSF is also discussed.

Minimum VM Program Levels

The VM operating system release level required for the SVA implementation is dependent on the functional device types used.

Before installing the initial the SVA subsystem, install and verify all of the current maintenance recommended in support of the 3990-3, 3380 and 3390 volumes. Check with the support center for the current PSP buckets.

***Caution: Performance Impact* - Host software must be supported by the vendor. IBM is the vendor for z/VM and z/OS.**

SYSGEN Considerations

The SVA subsystem does not have physical disk volumes in the traditional sense; however, the operating system does access disk-resident data as though it were on traditional disk subsystem. Data volumes, as seen by the operating system, are functional volumes in SVA terms. (Refer to the *V2X2 Shared Virtual Array Introduction* for a more detailed discussion of functional volumes).

From an operating system perspective, functional volumes are defined just as traditional disk volumes are defined in the Real I/O (RIO) definitions. In the VM environment, the RDEV macro is used to define a real device block for a functional volume.

You can define a maximum of 4096 functional devices. The primary advantage to defining 4096 functional volumes to VM is that if fewer than 4096 functional volumes are initially configured on the SVA subsystem, more can be added (i.e., defined) at a later date without an IPL. New volumes can be varied online after they are defined to the SVA subsystem.

The primary consideration when defining 4096 functional devices to VM is the added amount of real storage required for real control blocks for each functional device. The amount of storage required to support the control blocks for 4096 real device definitions is:

z/VM 409,600 (X’64000’) bytes

You should understand these implications for your specific environment. The industry trend has been toward systems with increasing amounts of memory, and the benefits of improved I/O response probably outweigh the cost of used host storage.

Device Initialization

The use of the ICKDSF format command to initialize SVA functional volumes in VM is identical to the process used for traditional disk volumes. Due to differences in the SVA architecture over conventional disk subsystems, there are some differences in the recommendations for formatting. Refer to “Preparing a Functional Volume” for a full discussion on initializing functional volumes using ICKDSF.

VM Cache Operation

DASD Fast Write (DFW) and Cache Fast Write (CFW) are both fully supported by the SVA subsystem. All writes in an SVA subsystem are DASD fast writes. All data transfers from the cache to the disk array are performed asynchronously in the background.

The implementation of cache fast write in the SVA subsystem is slightly different from the 3990. the SVA subsystem acknowledges and records cache fast write commands. All data transfers are performed exactly the same as for DASD fast write. the SVA subsystem indicates a successful write before the data is transferred to the disk array. Unlike the 3990, data is written to both cache and nonvolatile storage. This ensures that there is always a secure copy of the data under any circumstance.

The CP SET CACHE command is used to activate or deactivate caching by functional device address or by entire subsystem. From the

VM perspective, the CP QUERY CACHE command is used to display the caching status of the V2X2 subsystem.

Note: The SVA subsystem always operates in a caching mode, even if the response from the QUERY CACHE command indicates that caching is disabled.

Exploiting the V2X2 With SVAA/VM

SVAA is a host software package that provides the ability to configure, manage, and provides virtual disk administration within the SVA subsystem. Included in these capabilities are functions such as cache control and enabling DASD Fast Write (DFW).

SVAA maximizes the benefits of Sun's virtual storage architecture by providing the following capabilities:

- CMS command line interface for configuration and administration.
- ISPF full screen interface for configuration and administration
- VM NCL Management Utilities.
- SnapShot Copy Utility.
- Instant Format Utility.
- Data Collection and Reporting.

The V2X2A can be used to produce cache performance reports, including information on cache hit ratios and cache occupancy. The information from these reports can be used to effectively manage cache utilization.

Note: Refer to the Shared Virtual Array Administrator manual for VM for further information.

VM Performance Tools

In any VM environment, the standard VM/Monitor can be used to collect resource utilization and capacity planning data at user-specified intervals over desired time periods. Different packages, depending on the level of the VM system in use, are available for the reduction and reporting of this data.

The Virtual Machine Performance Reporting Facility (VM/ PRF) can be used to analyze and produce standard reports on the collected data.

The Virtual Machine Performance Analysis Facility (VM/PAF), which is a combination of VM/PRF with a capacity planning capability, can be used for longer term analysis.

The Realtime Monitor (RTM) can be used for short-term, dynamic monitoring and analysis.

Storage Management In VM

In VM, a Shared Virtual Array functional device is initialized and prepared for CP use in either system stand-alone mode (when the software tools used for volume initialization and preparation run on the CPU without benefit of an operating system) or under the control of a VM.

Note: In VM, a prerequisite to SVA functional volume initialization is the definition of the real addresses. These are defined in System Configuration, HCPRIO, or dynamically. Refer to the IBM VM Planning and Administration document.

Special Utilities - Device Support Facilities

Device Support Facilities (ICKDSF) is an IBM software utility that performs functional volume initialization and disk media maintenance. Volume initialization prepares a functional volume for VM operating system usage. However, the use of ICKDSF for media maintenance on SVA functional volumes does not perform the desired service.

Disk errors that cause an interruption in data availability for traditional disk storage systems are minimized with an SVA subsystem. a V2X2 can automatically correct media errors and dynamically reconstruct a failed disk drive. Refer to the *V2X2 Shared Virtual Array Reference* for further information on SVA drive reconstruction and automatic media maintenance.

Defining and Preparing a Functional Volume for VM

To bring a new functional device online to VM one of the following definitions must occur:

- SYSTEM CONFIG – RDEVICE statement
- HCPRIO – RDEVICE Macro (requires reassemble and CP Nucleus Generation)
- Dynamically sensed/defined during VM IPL
- Dynamically sensed/defined after SVA functional device define (attention interrupt)
- Dynamically defined by CP Command

In order to use an SVA functional device as a VM Volume, it must be defined to CP and initialized as a CP Volume (CPVOL). The ICKDSF utility is used to initialize these volumes into 4096 byte (4K) blocks, either from execution of the ICKDSF or CPFMTXA command. When a CP Volume is initialized, the CP allocation table (CYL 0, REC 4) must reflect the allocation type (DRCT, PAGE, PARM, PERM, SPOL, or

TDSK) for each 4K block (page) of contiguous disk space on the device.

CP Volumes are usually defined in the SYSTEM CONFIG file (or in HCPSYS) as a CP-Owned DASD or DASD Used for Minidisks. This is done by:

- SYSTEM CONFIG – CP_OWNED list
- SYSTEM CONFIG – USER_VOLUME_LIST

– OR –

- HCPSYS – SYSCPVOL Macro (requires reassemble and CP Nucleus Generation)
- HCPSYS – SYSUVOL Macro (requires reassemble and CP Nucleus Generation)

CP Volumes are then initialized (formatted/allocated) using ICKDSF or CPFMTXA as follows:

CP-Owned – A CP system residence volume must be formatted/ allocated by the System Administrator. These volumes can contain one or more of the following allocation types: DRCT, PAGE, PARM, PERM, SPOL, or TDSK. Care must be taken to allocate the space appropriately.

User Volumes – A CP volume intended for user minidisk use is formatted or allocated using of allocation type PERM (default). To initialize a user volume on a V2X2 it is only necessary to format or allocate Cylinder 0 (“CPFMTXA vdev volser 0-0”).

Note: Cylinder 0 of a CP Volume must be formatted before any other cylinders on that functional volume.

Refer to the IBM VM Planning and Administration Documentation for specific details on DASD (Auxiliary) storage for your particular version of VM.

Data Recovery

The SVA recovers from most hardware failures automatically. However, as with any other disk data storage subsystem, data corruption caused by application programs or user errors must be corrected by the customer.

Data Migration Planning

Migrating data to an SVA subsystem is similar to migrating data to any other form of disk storage. the SVA subsystem is designed for reliability and you should encounter no problems during the data migration process. To enhance the integrity and speed of the migration process, create and document a detailed migration plan prior to any movement of data. Some key components of a complete VM data migration process include:

- A report documenting the current allocations residing on a disk storage system
- A report documenting the utilization rates of the allocated mini-disks
- A complete tape backup of data from the disk storage system
- A report documenting the converted allocations from one device type to a different device type emulated by the SVA subsystem (if required)
- A plan to archive/off load any data determined to be not in use
- A description of the impact of data migration from the end-user's perspective
- A report documenting the I/O workload skew across the disk storage system devices containing data to be migrated to an SVA subsystem
- A plan to recreate the data residing on the original disk in the event that a process back-out is required.

Migration Strategies

The simplest migration method is a direct volume copy from the source volumes to the target functional devices. The usual IBM utility for this operation is Disk Dump Restore (DDR). This strategy is also the most expedient and similar method to existing migration plans. However, this method does not attempt to exploit any of the special features of the V2X2 subsystem.

The V2X2 subsystem features that are automatically available, regardless of the migration strategy used, include:

- **Fault tolerance** - the data is protected from hardware failures.
- **Data compression** - the data on an SVA subsystem is compressed upon entering the SVA Controller. This strategy reduces the amount of cache and physical space required to store the data.
- **Data compaction** - inter-record gaps are removed, further decreasing the amount of space required to store the data. Additionally, RPS miss is eliminated.

Mini-disk Migration Strategies

Several software tools are used to perform the actual data migration process in a VM environment:

- **CMDISK -** A DIRMAINT command that moves the data residing on a mini-disk from one device type to the same or another device type.

  DIRMAINT is an IBM licensed program product.

Note: VMSECURE has comparable commands, and CMDISK is called CHGMDISK.

- **COPYFILE -** A CMS command used to copy files residing on a mini-disk from one device type to the same or another device type.
- **SPXTAPE -** A CP command used to dump spool files to tape, and load spool files from tape to a functional device.

While each of these software tools has different characteristics, the following steps can be identified as required for a complete migration of a disk storage system volume to an SVA functional volume:

- Confirm that a valid VOLSER has been written to each functional volume
- MDISK statements in the VM directory reflect allocations on the functional volume
- Actual data is migrated to the mini-disks on the functional volume(s)

While these software tools (with the exception of SPTAPE) use a disk-to-disk copy process, it is also possible to use a disk-to-tape-to-disk migration process. Regardless of which process is used for data migration, it is recommended that a reliable backup copy of the data to be migrated be maintained until the migration process is complete and mini-disk integrity has been validated.

The selection of the software tool to use for data migration depends on the environment in which SVA is being implemented. Consult the following table to determine which tools optimizes the migration process in your environment:

Table 19 Migration Process Tools

ACTION	CMDISK	COPY FILE	DDR	SPTAPE
Used to migrate CMS formatted mini-disk data	Yes	Yes	Yes	No
Allows for altered mini-disk allocation sizes between DASD and a V2X2	Yes	Yes	No	Yes
Automatically updates VM directory with functional volume mini-disk volumes	Yes	No	No	No
Copies a physical device volser to functional volume	No	No	Yes	No
Migrates data between unlike device types (such as 3380 to 3390)	Yes	Yes	No	Yes
Migrates more than one mini-disk per command input	Yes	No	Yes	Yes

Migration Using DIRMAINT CMDISK

DIRMAINT CMDISK (change mini-disk) is a DIRMAINT command that moves the data residing on a mini-disk from one device type to the same or another device type. (DIRMAINT is an IBM licensed program product.) When the DIRM CMDISK command is issued, the following tasks are performed:

- A new mini-disk allocation, on a functional volume, for the CMDISK cylinder specifications is automatically created using a dummy virtual address for the DATAMOVE VM.
- The user's existing mini-disk is automatically reassigned to the DATAMOVE VM at another temporary address.
- The new mini-disk allocation on the functional volume is automatically formatted with the CMS FORMAT command.
- All the files on the existing mini-disk allocation are copied to the new mini-disk allocation on the functional volume.
- The space allocated for the existing mini-disk allocation is returned to the general pool of available mini-disk space, and the new mini-disk allocation on the functional volume is reassigned as the existing mini-disk for the target user id.

Migration Using COPYFILE

The CMS COPYFILE command can migrate a copy of the CMS files residing on a single mini-disk allocation to another mini-disk allocation.

With the COPYFILE command, the input and target mini-disks can reside on different device types (such as 3380 to 3390). For further information on the CMS COPYFILE command format, consult the appropriate CMS command reference document for your operating system environment.

Migration Using DDR

The DASD Dump Restore program can be used to perform a physical copy of a DASD volume to a functional volume. The DDR program can be executed stand-alone on the CPU (without benefit of the VM operating system) or under the control of a VM.

DDR requires the input and target volumes to be the same device type (such as 3380 to 3380). For this reason, an SVA functional volume must be defined as a 3380 to receive migrated data from a 3380 volume via DDR. Correspondingly, an SVA functional volume must be defined as a 3390 to receive migrated data from a 3390 device via DDR. No other device configurations of data migration are supported with DDR.

Note: The target volume of a DDR move must be the same size or larger than the input volume.

A series of DDR commands used to perform the physical migration of a volume can be placed in a control file. The control file name can then be specified in the DDR command to automatically execute the DDR statements in the control file. If a disk-to-tape-to-disk migration strategy is desired, DDR can facilitate such a strategy, limited only by the device type constraints previously described.

Migration Using SPXTAPE

The SPXTAPE command is used to move spool files from disk to tape, and load spool files from tape to disk. The input and target device types might differ. Spool files migrated to another device are assigned a new spooled to avoid a duplication conflict within the VM spooling system. For further information on the CP SPTAPE command format, consult the appropriate CP command reference document for your operating system environment.

Examine Data Access Characteristics

A more comprehensive process for migrating data to an SVA subsystem involves studying and understanding several activity and utilization characteristics of the current disk storage system configuration. These characteristics include:

- **I/O workload skew -** Minimizing the workload skew provides a performance benefit by reducing the I/O contention for a functional volume. Determine the percentage of I/O activity submitted to each volume over several time intervals. I/O data saved from the online VM monitor can be used to analyze workload skew by device. Volumes should then be ordered by the amount of activity they currently support, with the most active volumes given the highest priority for spreading across multiple functional devices.
- **Seek skew -** Understanding the seek skew assists in determining appropriate volumes for caching. Minimal seek skew offers the highest probability of a cache hit and therefore better performance. Determine the percentage of zero seek skew on a volume and mini-disk basis. One strategy to consider is migrating mini-disks with high zero seek skew rates to SVA functional volumes designated as cache volumes.
- **Capacity utilization -** By evenly distributing mini-disk allocations across functional volumes, greater environmental flexibility is offered to the process of adding new mini-disks to an SVA subsystem. Determine the percentage of the existing disk storage system capacity that is currently allocated, as well as the aggregate utilizations of the existing allocations. Data migration should be considered as an opportunity to review the current utilization of existing allocations and to modify these allocations based upon under- or over-utilization.

Using more functional volumes allows the reduction of access delay due to device queueing and improves the performance of the I/O subsystem. Understanding the I/O activity of the volumes to be migrated to a V2X2 subsystem maximizes the benefits of data migration.

Preparing for Data Migration

The major considerations for selecting data for migration involve the trade-off between the desired performance level and the data capacity of the subsystem. Remember that the fault tolerant design of an SVA subsystem protects data stored on it. This is the most important consideration when selecting data for migration.

Data movement should be planned to minimize the risk of data corruption or loss. The highest chance for data loss to occur is through human error or misunderstanding.

The first stage of migration should emphasize data that benefits significantly when placed on an SVA subsystem, but for which impact is low if unexpected problems are encountered during the migration process. Data in this category includes general user mini-disks and relatively inactive system data, such as CP and CMS source and aux file mini-disks. These mini-disks are most easily managed by systems programming personnel.

The next group of data to be moved is application mini-disks. It is particularly important to understand the I/O activity of these mini-disks to select candidates that realizes the greatest benefit from migrating to an SVA subsystem (refer to "Examine Data Access Characteristics" on page 91).

Application data can be the most sensitive data from the users' perspective. It realizes the greatest performance benefit from movement to an SVA subsystem.

Due to the disruptive nature of migrating system mini-disks, these migrations should be the most carefully planned. Such added consideration minimizes the time required to perform the migration and minimizes the risk of an error occurring during the migration process.

Some of the most active system mini-disks benefit the most from migration to SVA. The MAINT mini-disks linked to by VMs (190, 19D, 19E, etc.) show marked performance improvement from the fault tolerant aspects of the SVA subsystem.

Placing mini-disks On an SVA Subsystem

The placement of mini-disks on functional volumes under SVA is less critical than data placement on traditional disk storage systems. This is because traditional DASD has a single actuator between the stored data and the control unit, creating data transfer bottlenecks that adversely impact performance. Spreading active mini-disks across volumes is the usual method for addressing this problem with traditional DASD.

With an SVA subsystem, functional volumes are associated with multiple paths and actuators. The data transfer from the disk arrays to the Controller occurs on several paths simultaneously. The data for any functional volume can potentially be on any or all physical volumes on the array. This design minimizes the I/O bottlenecks within the subsystem, and reduces the concern for mini-disk placement on particular volumes.

The I/O queueing study performed in the planning stage (refer to "Examine Data Access Characteristics") has indicated a need to separate certain mini-disks because of concurrent activity to a volume. A problem condition is demonstrated by a repeated occurrence of device busy for I/O targeting specific mini-disks. Additional investigation of the mini-disk activity on the volumes to be migrated indicates which (if any) mini-disks should be allocated to different functional volumes. the SVA subsystem’s ability to define and support 4096 functional volumes allows for a high level of flexibility in mini-disk distribution.

IPL Volumes

Mini-disks containing IPL-able data, such as MAINT 190 have good performance characteristics because of SVA’s sequential detection and pre-staging. System areas such as the CP directory and saved system areas also benefit in this manner.

Common CMS Mini-Disks Linked by Other VMs

Mini-disks that are linked at VM logon by more than one VM, such as MAINT 19D, 19E, PROFS, or language compilers, are excellent candidates for disk caching.

Page/Spooling Mini-Disks

Page/spooling mini-disks are not good candidates for caching. In fact, overall VM system performance is degraded if these mini-disks are cached.

VM NCL Management

The SVA's disk array space is a limited resource that must be effectively managed. Reclaiming deleted data space is a critical part of space management. When a VM user deletes a file, the CMS file system is updated and the space becomes a free extent on the volume. In the SVA, as well as in traditional DASD, this deleted data still occupies the physical tracks. For traditional disk storage systems, physical tracks reside on a physical volume and cannot be used by any other volume. In the SVA subsystem, physical tracks (disk array space) are a shared resource and can be used by all functional devices.

The SVA Administrator software provides a facility for reclaiming the tracks in a deleted extent. This facility, called “VM NCL Management”, is documented in the *V2X2A for VM Configuration and Administration* manual.

TDisk Management

TDISK usage proves to be very efficient on SVA. The TDISK CLEAR option causes all TDISK space to be rewritten when the user releases the TDISKS. This results in good space management by the SVA.

TDISK formats also take less time on an SVA than on traditional disk storage systems. The TDISK clear option has some overhead since it causes the rewriting of this space. For shops that have large amounts of TDISK, this might be a consideration before enabling this option. A TDISK Clear Utility enables users who do not wish to have this happen automatically but would rather submit a utility at an off-peak time to do so.

The TDISK Clear option causes all TDISK space to be rewritten when the users release the TDISKs. The following screen is how you set the TDISK clear option in effect. The option is specified in the SYSTEM CONFIG file and invoked at initial program load (IPL). TDISK must reside on a CP-owned volume and the TDISK space is defined in the directory on MAINT and allocated to the CP-owned disk when it is formatted with CPFMTXA.

The SVA in a z/OS Environment

6

Installing the SVA

After installation, as with any other I/O device, an SVA subsystem must be defined to z/OS, MVS, or VM and the channel subsystem. The process for defining devices might vary, depending on the level of the operating system. This chapter describes the special software considerations for defining the SVA subsystem to an operating system. Functional disk initialization using ICKDSF is also discussed.

Minimum Program Levels

Minimum prerequisite host software levels are listed in the following table. This provides a *baseline* for additional PTFs and levels of software applications required for implementation of an SVA system.

SVAA available maintenance is recommended to be applied to the host operating systems.

Mainframe Environment

Table 20 Minimum Version And Release Levels of Software

Operating System	Required Minimum Level
z/OS	Version 1, Release 4
OS/390	Version 2, Release 10

Note: At the time of this writing, the information in this table was considered current. If this does not meet the conditions of your environment, contact Sun Microsystems for a possible update.

Caution: Performance Impact - **Host software must be supported by the vendor. IBM is the vendor for z/VM and z/OS.**

The host operating system must be able to support a 3990 storage control units and the functional device types being defined. Before installing the initial SVA subsystem, install and verify all of the current maintenance and release levels. Preventive software management strategy is highly recommended.

Configuration Considerations

The subsystem does not have physical disk volumes in the traditional sense; however, the operating system does access disk-resident data as though it were on a traditional disk subsystem. Data volumes, as seen by the operating system, are functional volumes in SVA terms.

From an operating system perspective, functional volumes are defined just as traditional disk volumes are defined: IODEVICE macro definitions describe all of the functional devices that are to be accessed by the operating system.

The subsystem can be defined with up to sixteen logical control units with a full complement of 3380 and/or 3390 devices.

If properly implemented, defining the maximum number of volumes allowable on the subsystem reduces the I/O delay at the host due to UCB queuing.

The disadvantage to defining 4096 devices in the host definition for a subsystem is the amount of *host virtual storage* required to support 4096 UCBs. This is controlled by defining only the devices needed in the IODEVICE definition. Usually you would not define more devices in the SVA that you have on the host. Adding more logical devices later to an SVA is a non-disruptive process and this is *usually* a non-disruptive process on the host system as well.

The implications of defining 4096 devices should be evaluated for each specific site. The industry trend has been toward systems with increasing amounts of memory, and the benefits of improved I/O response probably outweigh the penalty paid in virtual host storage used.

Using RMF to study and understand the UCB queuing in the current storage subsystem helps to determine the need for additional volumes. Similarly, storage utilization reports can help to determine possible storage constraints.

SVA allows you to configure functional volumes from the local operator panel or at the host console through SVAA.

Device Initialization

The use of the ICKDSF format command to initialize the V2X2 functional volumes in MVS, VM, or z/OS is identical to the minimal process used for traditional disk volumes. A minimum of one privileged ECAM device must be defined on the V2X2's detached operator panel (LOP). The prerequisite is to define the devices within the V2X2 and define the HCD/IOCP. Using SVAA, define functional devices within each VCU prior to varying the channels online.

In an open systems environment, use the Initialize Device subcommand described in the manual *SVAA for OS/390 Configuration and Administration*.

Cache Operation

Although the cache operation on an SVA subsystem is essentially the same as that of a 3990, the SVA's response to some of the 3990 commands varies from the 3990 response. The differences are described in the following sections:

Dual Copy

The SVA does not support the 3990 dual copy feature because the feature is not necessary in SVA. In response to dual copy commands such as Establish, Terminate, or Suspend Duplex Pair, the SVA presents status and sense data indicating that the device does not support dual copy.

DASD Fast Write And Cache Fast Write

DASD fast write (DFW) and cache fast write (CFW) are both fully supported by the subsystem. All data transfers from the cache to the disk array are an asynchronous background task.

The implementation of cache fast write in the subsystem is slightly different from the 3990. the SVA acknowledges and records cache fast write commands, but all writes and data transfers in a subsystem are DASD fast writes. Unlike the 3990, the data is written to both cache and to nonvolatile storage. This process ensures that in any circumstance, there is a secure copy of the data.

IDCAMS

The *IBM 3990 Storage Control Planning, Installation, and Storage Administration Guide* documents the cache functions of IDCAMS. The similarities and differences in SVA's responses to IDCAMS cache control commands are documented here.

LISTDATA

The LISTDATA command provides activity and status reports about cache, including:

Subsystem Counters Report

This report generates a report on subsystem counters at the device, subsystem, or system level by specifying DEVICE, SUBSYSTEM, or ALL, respectively. Device counters are for the functional devices defined with SVA subsystems.

Remote Access Code

IDCAMS can be used to generate a remote support access code for the V2X2 subsystem.

Pinned Track To Data Set Cross Reference Report

Because the SVA does not experience pinned data, this report always indicates that no data is pinned. the SVA's fault-tolerant design allows it to recover from both cache and NVS failures.

Device Status Report

This report describes the device status with the channel connection address and the director-to-device connection address for each device. SVA supports this function, and "cache inactive" is reported for cache status requests for devices that have caching turned off, even though SVA subsystems continue to cache all data.

SETCACHE

The recommendation is to allow the cache settings to be enabled at all time on the SVA.

The SETCACHE command has several functions. In the SVA, all directives for volumes operate on functional volumes.

The SETCACHE command makes functional volumes eligible for caching operations.

In a subsystem, all data is cached. Setting cache off for a device minimizes the cache storage available for that volume by changing the cache management algorithms. Modified tracks are immediately queued for de-staging, and unmodified tracks are discarded when other tracks need the space they occupy.

All other SETCACHE commands, except those directed to a device, are reflected in the reporting only. No action is taken by the SVA. Problems might occur when the host detects disabled cache settings.

Exploiting the SVA With SVAA

SVAA provides an additional level of control for an SVA subsystem beyond that provided by the local operator panel. In addition, the SVAA maximizes the benefits of the SVA's virtual storage architecture by allowing interactive control of activities such as subsystem administration and reporting.

The SVAA has additional reporting capability for the SVA. the SVAA can be used to produce cache effectiveness, space utilization, and device performance reports. The information from these reports can be used to effectively manage the SVA.

There might be some differences between the data produced by the SVAA reports and the data produced by other host software. These differences might be explained by the host perspective versus the subsystem perspective. Additional overhead might be recorded in the host statistics.

***Caution: Performance Impact* - Devices defined in the HCD must match the SVAA device definitions.**

z/OS Tools

In a z/OS environment, the following standard tools are normally available to collect key planning statistics about data usage:

- **Resource Measurement Facility (RMF) –** The most appropriate RMF reports are the Cache and the Direct Access Device Activity Reports. For this report, statistics are gathered for one or more peak processing periods (e.g., prime shift). With the peak I/O rate value and knowledge of the existing DASD configuration, peak access density can be determined. The potential or need for functional device re-configuration to exploit SVA's unique architecture might be determined using the same report.
- **System Management Facility (SMF) –** Numerous data reduction and analysis packages are available to aid in the analysis of SMF data. SMF data is often most useful as a long-term (i.e., days or weeks) resource measurement tool, but it can also be used to support specific data set or data pool analysis, where RMF cannot.
- **IDCAMS Listdata –** Can provide real time cache status and statistics for the subsystem or device.

Storage Management

Special Utilities - Device Support Facilities

Device Support Facilities (ICKDSF) is a software utility that performs volume initialization and disk media maintenance. Volume initialization prepares a disk volume for use by the operating system, while disk media maintenance provides a means to detect and correct problems with disk surface defects.

Use ICKDSF to initialize the SVA functional volumes, either as 3380s or 3390s. With traditional disk subsystems, incorrect placement of the volume table of contents (VTOC) can cause excessive disk head movement. Excessive disk head movement causes I/O performance degradation. Because a subsystem spreads data across multiple volumes in the disk array, excessive head movement is minimized. With a subsystem, there are no performance penalties for VTOC placement; however, the VTOC should be placed at the beginning of the functional volume to preserve contiguous free space. Only an ICKDSF MINIMAL INIT is necessary and recommended for the SVA.

Disk errors that cause an interruption in data availability for traditional disk subsystems are eliminated with a subsystem. the SVA can automatically correct media errors and dynamically reconstruct a failed disk drive. It is never necessary to use ICKDSF for media maintenance on SVA functional volumes. Refer to the *V2X2 Shared Virtual Array Reference* for further information on the SVA drive reconstruction and automatic media maintenance.

The following table provides a list of ICKDSF commands that are supported by a subsystem.

Table 21 ICKDSF Commands Supported by SVA in z/OS

ICKDSF Commands	Restrictions
AIXVOL with no surface checking functions	READCHECK is not supported
ANALYZE SCAN function only	DRIVETEST is not supported
BUILDIX	None
CPVOLUME with no surface checking functions	READCHECK is not supported
INIT minimal initialization	CHECK and VALIDATE are not supported
REFORMAT	None

Data Recovery

The SVA recovers from most hardware failures automatically. However, as with any other disk subsystem, data corruption caused by application programs or user errors must be corrected by the customer.

As a precaution when correcting VTOC errors on a SVA functional volume, suspend Deleted Data Space Release (DDSR) until the VTOC errors have been corrected.

Data Migration Planning

Migrating data to a subsystem is similar to migrating data to any other form of disk storage. the SVA subsystem is designed for 3990 compatibility. As insurance against accidental data loss and to expedite the migration process, develop and document a detailed migration plan prior to moving data. Make backups prior to moving any data to guard against accidental loss during migration. Also, secure a backup and recovery plan as part of the complete migration plan.

Migration Strategies

The simplest migration method is a direct volume copy from the source volumes to the target functional volumes. This strategy is also the most expedient. Some data migration software tools can be used to concurrently copy data on the target volume during production processing. Data migration strategies require precise planning and scheduling.

The subsystem features that are automatically available, regardless of the migration strategy used, include:

- **Fault tolerance –** the data is protected from hardware failures.
- **Data compression –** the data on a subsystem is compressed upon entering the SVA controller. This strategy reduces the amount of cache and physical disk space required to store the data.
- **Data compaction –** inter-record gaps are removed, further decreasing the amount of space required to store the data. In addition, RPS miss is eliminated.

When using the volume copy method, be aware that some volume copy utilities copy track images, which includes residual data in unused and un-allocated space on the track. SVA treats the images as valid data, compressing and storing un-allocated space during the copy process. Therefore, use utilities (or copy options) that recognize the end of data, and copy only active data. This reduces the net capacity load on the subsystem following a data migration. If storing the unused portion of track images is unavoidable, the net capacity load can be reduced to actual data allocation after migration by using the DDSR feature of SVAA.

DFSMSdss is a commonly-used data migration tool.

Using this package is exemplified by the following full-volume copy statement:

```
COPY FUL1 INDD(ddn1) OUTDD(ddn2) COPYVOLID
```

The values for ddn1 and ddn2 identify DD statements that define the source and target DASD volumes. DFSMSdss defaults to copying only the used space for allocated data sets and nothing for empty data sets. COPYVOLID can be used to copy the VOLSER from the input volume to the output volume (it is required for a full-volume copy of an SMS-managed input volume). At the end of the copy operation, before the volume serial number on the receiving volume is altered, the operator is notified and the operating system initiates a demount of the volume.

This type of migration strategy requires halting applications and updates to the source volumes.

Examine Data Access Characteristics

A more comprehensive process for migrating data to a subsystem involves studying and understanding the I/O activity of the data prior to selection and movement. RMF reports can be examined to find the I/O activity by volume. Volumes should then be ranked by the amount of activity they support, giving the most active volumes preference for movement to the subsystem.

The same method can be used to select volumes for caching, if it has been determined that only selected volumes should be specified as cached volumes. Specify volumes with high I/O rates over volumes with low I/O rates to take advantage of improved performance from caching.

The I/O activity profile indicates if specific volumes have excessive I/O activity and queuing or reserves in a shared DASD environment. Those volumes can be analyzed further to determine the most effective way to spread the data across two or more functional volumes. Using more functional volumes to separate concurrently-accessed data sets minimizes UCB queuing and volume reserves and improves the performance of the I/O subsystem. Single host queuing is reported as the IOSQ in RMF. Reserve delay is reported as PEND in RMF.

Minimizing SVA’s net capacity load is an additional benefit of this method of data migration. Data movement tools that work at the data set level do not move unused allocations. In fact, they can reduce the size of over-allocated data sets, if desired.

Preparing for the Movement of Data

The major considerations for selecting data for migration involve the trade-offs between the desired performance level and the data capacity of the subsystem. Remember that the fault-tolerant design of the subsystem protects the data stored on it. This might be the most important consideration when selecting data for migration.

It is particularly important to understand the I/O activity of application data to select candidates that realizes the greatest benefit from migration to an SVA (refer to "Examine Data Access Characteristics" on page 104). Give consideration to the potential benefit of spreading the data across more functional volumes.

Due to the disruptive nature of migrating system data sets, they should be the most carefully planned. Careful planning minimizes the time required to perform the migration and minimizes the risk of an error occurring during the migration process.

Placing Data Sets on an SVA Subsystem

Data Set Segregation

Placement of data on functional volumes is less critical than data placement has been on traditional disk volumes. Traditional volumes have one actuator to transfer data to and from a volume. This creates data transfer bottlenecks, which adversely impact performance. Spreading active data sets across volumes has been the usual method for addressing this problem.

In a subsystem, functional volumes are associated with multiple physical actuators. The data for any functional volume can potentially be on any or all physical devices in the subsystem. Data transfers from the disk arrays to the Controller occur on several paths simultaneously. This design minimizes the I/O bottlenecks within the subsystem and reduces (but does not eliminate) the concern for data set placement on particular functional volumes.

An I/O queuing study performed in the planning stage (refer to "Examine Data Access Characteristics" on page 104) might indicate that there is a need to separate certain data sets because of concurrent activity to a volume. Such problems are indicated by excessive IOSQ time or PEND time on a device or UCB. Further investigation of the data set activity on that volume indicates which data sets should be spread across multiple UCBs. the subsystem's ability to support up to 4096 functional volumes allows for much more data set separation than has previously been feasible.

System Data Sets

Data sets in this category include operating system data sets, catalogs, paging data sets and spool data sets. All of these data sets benefit from the fault-tolerant design of the subsystem. Dual redundancy provides system data with protection from hardware errors that are a serious concern with traditional disk media.

IPL volumes gain little in performance, except during the initial program load process where sequential loading of the control program dominates the I/O activity. SVA performs sequential read ahead when sequential processing is detected, ensuring that the next data requested is available in cache; therefore, the operation occurs at channel rates, not device rates.

Catalog placement on an SVA is less critical than on traditional disk subsystems. To minimize contention for volume access, catalogs are commonly isolated onto volumes, which have very little other activity. Functional volumes in an SVA have multiple access paths because they are mapped onto multiple physical disks. Given that, it is still important to monitor and control the amount of activity to the functional volume because of the possibility of UCB queuing. UCB queuing should be kept to a minimum on catalog volumes to avoid a system-wide bottleneck, but higher overall activity can be supported on SVA functional volumes than can be supported on traditional volumes.

JES2 and JES3 spool volumes benefit greatly when they are placed on a subsystem. Much of the data that is read and written to spool volumes is processed sequentially, benefitting from cache access, especially for read processing. The longer channel command word (CCW) chains that are used to read and write data to the spool data sets benefit from DASD fast write. The checkpoint data sets, in particular, benefit from DASD fast write because of good locality of reference.

Page and Swap data sets perform on a subsystem much the same way that they currently do on a 3990 subsystem. The data access on paging data sets is random and involves relatively small quantities of data (single 4K blocks). This type of data access does not benefit greatly from the sequential pre-fetch, DASD fast write, and caching features of SVA.

Access to high activity load libraries is much improved (unless Virtual Fetch is used) on SVA. The directory of high activity load libraries are likely to remain in cache.

Application Data Sets

All application data sets benefit from residing on a subsystem, some more so than others. Dual redundancy provides a high degree of continuous availability for all data sets without the overhead and cost of dual copy. In addition, most data set activity benefits from the advanced caching features.

Work and temporary data sets, if active enough, might never be written to the functional volumes to which they are directed. This can have a strong positive effect on throughput of tasks that use temporary data sets. Work data sets are usually accessed sequentially. Data that is written to the functional volumes are retrieved very quickly because of sequential pre-fetch.

VSAM data set definitions and attributes should be reviewed when moving VSAM data to a subsystem. Some of the performance options that are used on traditional disk subsystems are no longer needed in an SVA environment, though they do still work. These options specifically include REPLICATE and IMBED.

The REPLICATE option is a VSAM performance solution directed at the traditional disk architecture. It attempts to minimize the effects of rotational delay on data transfer. Data transfer from an SVA disk begins immediately once the correct track is located, eliminating RPS miss, and minimizing rotational delay. In addition, a read request for a particular record causes the entire functional track containing that record to be staged into cache. Although there is little to no advantage to using REPLICATE and IMBED on a subsystem, there might be some advantage to avoiding it. Maintaining the index as a discrete data component increases the probability of a cache hit. If the index is imbedded in the data or is replicated, the probability of a cache hit is reduced for both the data and the index. In addition, REPLICATE increases the net capacity load of the subsystem.

Space Management

With a subsystem, the physical disk array space is managed automatically. However, DASD space management tasks used to manage traditional disk space should also be used to manage SVA's functional volumes. The necessity to perform space management tasks does not differ with a subsystem. Tasks such as volume de-fragmentation, idle space release, data set reorganizations, etc. are still important to manage a subsystem’s functional space. For additional information on how data is stored in a subsystem, refer to the *V2X2 Shared Virtual Array Reference.*

Volume De-fragmentation

As data sets are allocated and deleted, the DASD space available for data set allocations becomes fragmented. Even though there might be adequate space available for new allocations on a fragmented DASD volume, that space is often not available in sufficient contiguous amounts. In a traditional disk subsystem, management of space fragmentation often means the physical movement of the data to consolidate the available free space. Typically, the storage administrator executes a special utility that performs the consolidation of the fragmented space.

With a subsystem, the functional volume VTOC does not reflect the space available in the disk array but rather the space available on the functional volume. As data sets are allocated and deleted, the functional volume becomes fragmented. This fragmentation can cause “X37” type abends. Functional volume fragmentation can be minimized by defining maximum size functional volumes. This allows data set allocations to occur in larger contiguous areas, thereby reducing the necessity of de-fragmentation operations. However, functional volume fragmentation can occur and a de-fragmentation utility should be executed to consolidate the fragmented functional volume.

The following guidelines apply with or without the utilization of Snapshot capabilities:

- Schedule defragmentation infrequently during low system activity.
- Target application volumes with changeable data.
- Set parameters to achieve maximum results with limited data moves.

These data and functional volume attributes influence the elapsed execution time of defragmentation tasks:

- Extent sizes of the data sets and the free space reflected by the VTOC
- Distribution of the free space on the functional volume
- Percentage of free space on the functional volume
- Device type

For a discussion on how data is stored on a V2X2 subsystem, refer to the *V2X2 Shared Virtual Array Reference.*

Data Archival Strategies

There are many storage management utilities that allow for the migration of inactive data from primary DASD storage to a secondary storage medium. Typically, there are two levels of storage for less active data. Commonly, the first level of storage is DASD and the second is tape. These levels are often referred to as ML1 and ML2. Data migrations often use a host resident data compression routine to provide a more effective utilization of the secondary storage medium.

Small Data Set Packing (SDSP) is a method employed by DFSMShsm that increases the effective utilization of ML1 volumes. The SDSP function stores user data sets as logical records within VSAM key-sequenced data sets (KSDS), thereby allowing more data sets to be stored per track than with the standard DASD allocation unit. Those data sets that are not migrated into an SDSP data set are stored on ML1 DASD space allocated in tracks. Migrated data sets stored in an SDSP data set are stored in space increments of 2 KB. a V2X2 subsystem automatically stores all data in 2 KB space increments. Using SDSP on a subsystem is inefficient and costly in terms of the DFSMShsm overhead (compression, I/O to search for candidate small data sets, maintenance of the SDSP data set, I/O to the DFSMShsm control data sets, etc.). *The use of SDSP with a subsystem has no resultant benefit and is not recommended.* The elimination of SDSP provides a reduction in CPU cycles and DASD I/O.

The SVA provides increased data availability and data compression/ compaction outboard of the host. The following recommendations establish an inactive data set migration strategy that takes advantage of SVA's characteristics:

- Do not use SDSP for data migrations to a subsystem ML1 pool.
- Do not migrate inactive data sets residing on SVA ML0 volumes to an ML1 pool.
- You must continue to migrate inactive data sets to an ML1 pool contained in an SVA - turn “Off” host compression. There is a reduction in CPU cycles.

DFSMS, CACHE Console Command

Automated storage management eliminates the requirement to manage data at the volume level. With DFSMS, management of data can be accomplished at the data set level to achieve availability and performance objectives. the V2X2 subsystem presents an alternative to traditional disk storage subsystems that exceeds today's availability and performance objectives in a cost-effective manner. This portion of the manual addresses some considerations for a V2X2 subsystem in a DFSMS environment and explores ways in which DFSMS can exploit a subsystem.

Data Class

The data class construct defines allocation attributes for a data set. a subsystem can emulate most 3380- and/or 3390- type DASD devices. With a subsystem, data class can be used as though the data set allocation were to traditional disk storage. However, fewer data class constructs are necessary, because SVA allows larger allocations for all files.

Storage Class

The storage class construct describes the service levels for a data set. For the data set to be managed by DFSMS, a storage class must be assigned to the data set. The storage class construct identifies the data set requirements for availability, guaranteed space and performance. Critical data that requires continuous availability should be placed on a subsystem.

SVA does not support dual copy commands. The use of dual copy is not recommended with SVA subsystems because it uses storage resources inefficiently with no resultant benefit. Careful consideration should be given to designing storage class and storage group automatic class selection (ACS) logic to prevent dual copy allocations on a subsystem. Data sets with the Availability parameter set to Continuous should be directed to a storage group containing a 3990 with dual copy enabled. Failure to properly structure the storage group ACS logic causes the data set allocation to fail.

When the DFSMS is active, caching at the data set level can occur dynamically with the 3990. There are three classifications a data set might have for dynamic caching:

- NEVER-CACHE
- MAY-CACHE
- MUST-CACHE

Data sets are classified based on their millisecond response (MSR) values defined in the storage class. The responses for cache reads and DASD-fast writes are shown in the following table.

Table 22 Dynamic Data Set Caching Classifications

Operation	**MSR=999 BIAS=Blank R,W**	**MSR=blank BIAS=Blank R,W**	**25=<MSR<999 BIAS=Blank R,W**	**MSR=blank BIAS=W**	**MSR<25 BIAS=Blank R**
Cache Read	NEVER-CACHE	MAY-CACHE	MAY-CACHE	MUST-CACHE	MUST-CACHE
DASD-Fast Write	NEVER-CACHE	MAY-CACHE	MAY-CACHE	MUST-CACHE	MUST-CACHE

At data set open time, cache usage is determined by the MSR value defined in the storage class. Data sets classified as MAY-CACHE use 3990 cache, if cache is not over-utilized. MUST-CACHE data sets always use cache. Data sets classified as NEVER-CACHE never use cache in a 3990 subsystem. DFSMSdfp periodically collects 3990 subsystem information (DINTERVAL default value is 150 seconds) to ascertain the utilization of cache and NVS. If the 3990 cache or NVS is over-utilized, DFSMSdfp progressively restricts MAY-CACHE data sets from cache. Conversely, if cache and NVS are under-utilized, DFSMSdfp allows previously restricted MAY-CACHE data sets to use cache.

To monitor cache utilization in a DFSMS environment, the DFSMS, CACHE console command can be used. This command provides information for each 3990 cache controller that has at least one DFSMS-managed volume attached to it. The information provided includes:

- The number of DFSMS-managed volumes attached to cache
- The percent of MAY-CACHE reads allowed to use cache
- The percent of MAY-CACHE writes that are allowed to use the DASD fast write feature
- The read hit ratio expressed as a percentage
- The number of DASD fast write waits per minute.

Because the subsystem is a cache subsystem and all writes are DASD fast writes, the storage class construct can be simplified for data residing on a subsystem. For the majority of data residing on a subsystem, the MSR values and bias values should be left blank. For data sets that are MUST-CACHE, the MSR value should be less than

25 milliseconds. The caching algorithms of the subsystem are modified to avoid unnecessary over-usage of cache memory for those data sets that are classified as NEVER-CACHE. When caching is restricted by DFSMSdfp, the MAY-CACHE data sets inhibit CACHE LOAD SET in the DEFINE EXTENT CCW SET to prevent the data set from being loaded into cache. With a subsystem, the restriction of cache by DFSMSdfp results in the modification of the caching algorithms for the MAY-CACHE data sets to provide performance similar to non-cache. To exploit SVA subsystem caching, all functional volumes should have cache turned on. This allows those data sets that are MAY CACHE to utilize the subsystem cache more effectively. Using the DFSMS CACHE command provides insight into the utilization of the subsystem cache in a DFSMS environment.

Management Class

The management class construct identifies a data set's expiration, retention and migration criteria. DFSMShsm is a storage management software utility that performs data migration. The use of DFHSM can be significantly enhanced by a subsystem. The following discussion refers to DFSMS, but the same concepts are applicable to other data management products.

With traditional disk subsystems, DFHSM uses a significant amount of processing resources during the data migration and recall process. When DFHSM data compression is active, CPU cycles are expended to compress and decompress data sets. DASD space consumption increases due to the addition of ML1 volumes. I/O rates increase due to the data migration/recall process. Lastly, access to migrated data is slowed as a result of the recall process.

With an SVA, there is a significant potential to reduce the amount of resources that DFHSM requires. Because a V2X2 subsystem compresses and compacts data outboard of the host, an ML1 pool should be defined to a subsystem. Establishing a DFHSM ML1 pool comprised exclusively of SVA functional volumes allows DFSMShsm data compression to be disabled on the host. Data migrations from traditional disk subsystems to an SVA ML1 pool can occur without the associated DFSMShsm compression overhead. DFSMShsm data compression for ML2 can be eliminated. The elimination of DFSMShsm data compression for ML1 and ML2 minimizes the CPU cycles consumed by DFSMShsm data migration/recall and backup/recovery operations.

The DFSMShsm command to disable host compression for ML1 and ML2 is:

```
SETSYS COMPACT(NONE)
```

To further exploit an SVA, DFSMShsm data migrations to ML1 should be eliminated for data sets residing on an SVA subsystem. CPU cycles consumed by DFSMShsm are further reduced by the elimination of the ML1 migration process for data sets residing on a V2X2 subsystem. To implement this strategy, set the LEVEL 1 DAYS NON-USAGE parameter in the management class to 0 for the data sets residing in the storage group consisting of the subsystem. Because of SVA's compaction and compression, it is unnecessary to migrate data to ML1 to increase primary DASD storage capacity.

Storage Group

In typical DFSMS implementations, storage groups are mapped across various DASD devices to optimize availability and device performance. The fault tolerance and redundancy of all data on a subsystem makes dual copy unnecessary. With an SVA subsystem, the process can be reduced to two steps:

1. Define a storage group with SVA functional volumes.
2. Design ACS routines that direct both critical and noncritical data to the storage group containing the SVA functional volumes.

When defining storage groups for a subsystem, segregation of data based on availability requirements is no longer necessary.

There are two common approaches to configuring storage groups for performance: horizontal and vertical mapping.

The horizontal approach simply maps a storage group to a number of different device types with varying performance characteristics.

The vertical approach maps a storage group to devices that are similar and have similar performance characteristics. The horizontal approach is common when the DASD is expected to be under-utilized or where a major portion of the DASD is attached to cache controllers. The vertical approach is used mainly where there is a limited number of cache controllers available.

DFSMS reduces the number of storage groups or DASD pools that a storage administrator manages. a subsystem with extended capacity and Dynamic configuration should be defined with maximum size and number of functional volumes. With data dispersed across all functional volumes, UCB queueing is reduced and out-of-space conditions are be minimized. Extended capacity and Dynamic configuration allows for the establishment of a single storage group comprised of SVA functional volumes. The creation of a single storage group for data sets residing on a subsystem provides increased data availability and performance while reducing the number of storage groups the storage administrator manages. When multiple storage groups must be defined to a subsystem, the horizontal approach should be used to take advantage of the subsystem's parallelism in channel operations and to maximize dispersal of data across functional volumes. Volumes that are not typically managed by DFSMS (IPL volumes, ML1 volumes, etc.) can co-exist with a storage group defined to a V2X2 subsystem. Because a subsystem provides fault tolerance and redundancy, as well as exceptional performance characteristics, the need to segregate data based on availability and performance requirements is virtually eliminated.

Volume Conversion

Volume conversion with a subsystem is accomplished in the same manner as with a traditional disk subsystem.

Backups and Disaster Recovery

Implementing a backup and disaster recovery strategy is highly recommended. This process might include external physical and logical backups on-site and off-site on a timely basis. Incremental backups of modified data and catalog backups are suggested to be included in this process. SNAPSHOT processes are internal to the subsystem and might not provide all the backup and disaster recovery criteria. The methods might use PPRC and/or tape solutions. The value of the data, resources to backup the data, and the resources to recreate the data (both production and development) are some of the factors to consider. For example, evaluate the potential cost if the only copy of the data is accidentally deleted by initializing or overwriting a volume. Do not bypass the essential process of external backups.

Device Map And Host Addresses

A

To allow accessibility to the devices, a new host IODF must be activated with the SVA hardware definitions.

During a new installation, functional devices are mapped as illustrated in the 4096-Device Support example in the following table. This mapping is mandated by the SVA microcode.

Table 23 4096 Device Support Mapping

Sequential Device Numbers	Parameters	Sequential Device Numbers	Parameters
0 - 255	VCU0 DEV 0-255 FDID 0-FF CUADD=0	2048 -2303	VCU8 DEV 0-255 FDID 800-8FF CUADD=8
256 - 511	VCU1 DEV 0-255 FDID 100-1FF CUADD=1	2304 -2559	VCU9 DEV 0-255 FDID 900-9FF CUADD=9
512 - 767	VCU2 DEV 0-255 FDID 200-2FF CUADD=2	2560 -2815	VCUA DEV 0-255 FDID A00-AFF CUADD=A
768 - 1023	VCU3 DEV 0-255 FDID 300-3FF CUADD=3	2816 -3071	VCUB DEV 0-255 FDID B00-BFF CUADD=B

Table 23 4096 Device Support Mapping (Continued)

Sequential Device Numbers	Parameters	Sequential Device Numbers	Parameters
1024 -1279	VCU4 DEV 0-255 FDID 400-4FF CUADD=4	3072 -3327	VCUC DEV 0-255 FDID C00-CFF CUADD=C
1280 -1535	VCU5 DEV 0-255 FDID 500-5FF CUADD=5	3328 -3583	VCUD DEV 0-255 FDID D00-DFF CUADD=D
1536 -1791	VCU6 DEV 0-255 FDID 600-6FF CUADD=6	3584 -3839	VCUE DEV 0-255 FDID E00-EFF CUADD=E
1792 -2047	VCU7 DEV 0-255 FDID 700-7FF CUADD=7	3840 -4095	VCUF DEV 0-255 FDID F00-FFF CUADD=F

Up to 4096 functional devices might be configured in a subsystem. The functional device IDs are mapped from 0 through FFF. Each of the sixteen Virtual Control Units (VCU) supports 256 devices. Whether the devices are defined or undefined, the mapping as illustrated in Table 23 is applicable. The IOCP/HCD must correlate with the mapping for 4096-device support. CUADD (logical control unit address) is an IOCP/HCD CNTLUNIT parameter. Refer to "IOCP/HCD" on page 121.

Host 4-digit addressing is not required for 4096 virtual device address ability. Devices might be addressed with host 3-digit UCBs (unit control blocks).

Sequential host addressing is highly recommended, but not required in 256 or 1024 ranges. A simple addressing scheme eases disk administration.

Non-sequential addressing could complicate identifying the subsystem for a particular device or increase the exposure for potential problems. If sequential addressing ranges are unavailable, consider readdressing other devices or using host 4-digit ucbs. However, if the addresses must be non-sequential, review the following steps carefully:

1. Identify the available host addresses.
2. Identify the host address availability in sequential ranges. Note 256 and/or 64 ranges, if any.
3. Identify the functional device ids to be defined.
4. Cross-reference the functional device ids, the IOCP/HCD IODEVICE UNITADD parameters, and the host addresses. Do not overlap host addresses. See Table 24 on page 117 for an example.

Table 24 Non-sequential Addressing Cross-reference Example

Virtual Control Unit 0 (VCU0) CUADD=0		
FDID 0-FF	UNITADD=00	ADDRESS=(100,256)
Virtual Control Unit 1 (VCU1) CUADD=1		
FDID 100-17F FDID 180-1BF FDID 1C0-1FF	UNITADD=00 UNITADD=80 UNITADD=C0	ADDRESS=(300,128) ADDRESS=(280,64) ADDRESS=(9C0,64)
Virtual Control Unit 2 (VCU2) CUADD=2		
FDID 200-23F FDID 240-27F FDID 280-2BF FDID 2C0-2DF FDID 2E0-2FF	UNITADD=00 UNITADD=40 UNITADD=80 UNITADD=C0 UNITADD=E0	ADDRESS=(500,64) ADDRESS=(440,64) ADDRESS=(580,64) ADDRESS=(6E0,32) ADDRESS=(7E0,32)
Virtual Control Unit 3 (VCU3) CUADD=3		
FDID 300-33F FDID 340-35F FDID 360-37F FDID 380-39F FDID 3A0-3BF FDID 3C0-3FF	UNITADD=00 UNITADD=40 undefined UNITADD=80 undefined undefined	ADDRESS=(720,64) ADDRESS=(840,32) ADDRESS=(880,32)

Tip: Attempt non-sequential addressing schemes only if needed and thoroughly understood.

Alternative Restoration

B

DF/DSS cannot do a standalone restore to an SVA 3380K (6EA) devices. Because DF/DSS issues a Sense ID command, to which an SVA responds as a K device, DF/DSS never issues a Read Device Characteristics command to get the number of cylinders. Therefore, the host Define Extent parameters have an ending extent for a K device, which an SVA rejects.

If the customer uses DF/DSS and does not have an alternative system from which to do restores, the customer or service representative must work around this incompatibility to restore the data to a real K device. When a customer or service representative is defining (configuring) initially, he/she must do one of the following:

1. Use K devices and not use all of the cylinders
2. Use KE devices. (KE devices are K devices plus the number of cylinders defined.) If the customer must restore to 3380K (6EA) devices, then the customer or service representative must restore to a K device and then zap the DIRF bit and the VTOC, which forces a free space recalculation.

IOCP/HCD

C

Description

Hardware Configuration Definition (HCD) provides an interactive, system interface allowing definition of the I/O hardware configuration for both the channel subsystem (CSS) and the operating system. The configuration might consist of multiple processors, each containing multiple partitions.

HCD stores the entire configuration data in a central repository, the input/output definition file (IODF). The IODF, a VSAM linear data set, is a single source for all hardware and software definitions for a multi-processor or multi-partition system.

An input/output configuration data set (IOCDS) contains different configuration definitions for the selected processor. The IOCDS is built from the production IODF. The CSS uses the configuration data to control I/O requests.

Prior to HCD availability, IOCP (I/O configuration program) and MVSCP (MVS configuration program) were used. The hardware was defined to the channel subsystem with IOCP and to the MVS operating system with MVSCP.

The I/O hardware configuration must be defined to the host operating systems connected to the SVA subsystem.

- The RESOURCE, CHPID, CNTLUNIT, and IODEVICE statements are described in this section.
- Examples are included at the end of this section. Each example includes a scenario description, channel and ICE card configuration, and IOCP statements.
- A maximum of 4096 functional devices might be configured in a subsystem.
- I/O configuration rules, processor limits, and the operating system govern how configurations might be coded.

- New installations, additional processors, additional channels, and device type definitions might require changes to the configuration definition.

Physical Configuration and HCD

Up to 16 logical control units (VCUs) might be defined in the HCD for the SVA. Use the following criteria:

- To access all 4096 devices, all 16 VCUs must be defined.
- ESCON channel paths defined to a set of 4 LCUs must differ from those defined to another set of LCUs. This constraint is due to a maximum configuration of 1024 devices (4 LCUs) to an ESCON channel.
- A maximum of 8 channel paths to each VCU can be configured from each operating system. This is a host processor limit.
- Coordinate the physical channel configurations with the HCD definitions:
  - Verify that the physical configuration *matches* the HCD/IOCP definitions. For optimal performance, split the channels from each host between the clusters.
  - At least 2 channels from each operating system are recommended to avoid a single point of failure (unavailability). Configure the channels symmetrically between the clusters. At least 8 channels and up to 8 ICE3 cards are recommended for optimal performance and to take advantage of the architecture.

Clusters And Paths Selections

The CUNUMBR and the PATH parameters in the CNTLUNIT definitions for each VCU (CUADD) determine the path rotation to access a device:

- Control unit numbers of the control units (VCU) defined to a device
- Channel paths used by the control units

The CSS, channel subsystem, selects the first channel path specified in each control unit defined to the device, beginning with the lowest numbered control unit. This pattern continues with each control unit defined to the VCU (CUADD). Then CSS selects the second channel path in each control unit defined to the device, beginning with the lowest numbered control unit. This process repeats until all channel paths in all control units defined to the device have been tried.

To ensure distribution of the workload over the two logical components of the SVA, specify the HCD definitions to correlate with path rotation between cluster 0 and cluster 1 of the SVA. Verify optimal configuration of the channels to the ICE card interfaces.

The following examples show how two methods of definitions results in the same path rotation sequence:

20, 21, 22, 23, 60, 61, 62, 63

For optimal performance and availability, code the PATH parameter in the CNTLUNIT definition by using either of the following methods:

- One control unit for each VCU (CUADD) might be defined in the HCD. Code the channels in the PATH parameter so that they alternate from cluster 0 to cluster 1 of the SVA. The channel subsystem tries the channels in the sequence coded in the PATH parameter.

  Example: PATH=(20,21,22,23,60,61,62,63)

  Paths 20,22,60,62 must be cabled to cluster 0, and paths 21,23,61,63 must be cabled to cluster 1.

- If 2 control unit definitions for each different CUADD parameter are coded in the HCD, only 1 VCU is generated for use by the operating system and CSS. The paths in the first definition must be physically connected to cluster 0, and all the channels in the second definition must be connected to cluster 1 of the SVA.

  Example: 2 CNTLUNIT statements with PATH=(20,22,60,62) and PATH=(21,23,61,63).

  Paths 20,22,60,62 must be cabled to cluster 0, and paths 21,23,61,63 must be cabled to cluster 1.

Note: ICE 00, 01, 02, and 03 cards interface with cluster 0, and ICE 10, 11, 12, and 13 cards interface with cluster 1. The channel configuration as described previously allows for the proper rotation of channels between clusters.

Defining the HCD/IOCP

Tip: To prevent problems, keep the SVA configuration as simple as possible. Try to stay within the guidelines established in this document.

RESOURCE Statement

- Defines logical partition names and numbers for LPAR IOCDS mode.
- Not required for Basic IOCDS mode.
- Keywords:

  PARTITION=((name, partition number),...)

- Identifies the information utilized to generate a LPAR IOCDS. All logical partition names must match those specified in the RESOURCE statement.
- Specifies the name of the LPAR to be defined, 1-8 alphanumeric characters (non-numeric first character).
- Designates the number associated with the respective LPAR, a one-digit hexadecimal.

CHPID Statement

- Defines channel path identifiers (chpids).
- Keywords:

PATH=(chpid number,...)

Specifies a two-digit hexadecimal number for each CHPID number, maximum of eight CHPIDs.

TYPE=CNC

Designates the mode of I/O operation for the channel path (TYPE=CNC for ESCON).

Note: The ESCON channels can be configured to a maximum of 1024 devices (four LCUs)

Type=FCV

***Caution: Performance Impact* - Only FICON channels are allowed.**

Specifies the channel path is a FICON channel attached to an ESCON director. FCV channel paths operate the same as ESCON channel paths.

PARTITION=(name,....)

PARTITION=((access list),(candidate list),REC)

Specifies the name of the logical partitions, matching any names in the RESOURCE statement.

SHARED

- The Multiple Image Facility (MIF) allows logical partitions to share channel paths. If a processor complex has MIF capability and is running in LPAR mode, its logical partitions can share channel paths to reduce the number of physical connections.
- Indicates a channel path can be accessed by more than one logical partition (LP) at the same time (MIF).
- The channel paths to the same control unit must be all un-shared or all shared within a single IOCDS.

SWITCH=number

- Required for channel paths to have dynamic connections through a director.
- Specifies an arbitrary number for the Director to which all of the channel paths on the CHPID statement are assigned. The number can range from 00-FF for a maximum of 256 for each IOCDS. A channel path might be assigned to only one switch number.

CNTLUNIT Statement

- Describes the control unit characteristics.
- Identifies the attached channel paths (maximum of eight for each VCU).
- Specifies the unit addresses recognized by the control unit.
- Required for each group of 256 addresses, one CNTLUNIT per cluster.
- Keywords:

CUNUMBR=number

Specifies the unique hexadecimal number (0000-FFFE) assigned to the control unit. Assign the lower number cluster 0 and the higher number to cluster 1 for each group of 256 devices.

PATH=(chpid,...)

Identifies the channel paths attached to the control unit, two hexadecimal digits for each channel path identifier (CHPID). Define each set of paths in the same order for each group of 256 addresses. Code the channels so that they alternate between cluster 0 and cluster 1 of the SVA.

LINK=link address

-Must equal the link addresses associated with the physical ports on the director to which the control unit is attached.

- Designates the director link addresses attached to the control unit, corresponding to the same sequence as the channel paths in the PATH keyword. For example, the first link address must correspond to the first channel identified in the PATH keyword.
- The LINK keyword is optional if none of the channel paths in the PATH keyword attaches to an Director (SWITCH keyword).

Link addresses that correspond to channel paths not attaching to an Director can be specified as two asterisks (**) or any valid link address value. The valid link address range is 01-FE.

- UNITADD=((address,number))

UNITADD=((00,256))

- Defines the unit addresses of the I/O devices recognized by the control unit.
- Indicates the number of sequential unit addresses recognized by the control unit.
- Designates the full range of unit addresses that the control unit can address, whether or not all devices are defined.
- UNITADD=((00,256)) is required.

CUADD=address

- Specifies the unique logical address for the control unit, a one-digit hexadecimal number in the range 0-F. This address allows a CNC channel path on multiple CNTLUNIT statements for a given path (chpid.link or switch.link).
- CUADD=0 allows access to the 1st set of 256 devices (FDID 0-FF).
- CUADD=1 allows access to the 2nd set of 256 devices (FDID 100-1FF).
- CUADD=2 allows access to the 3rd set of 256 devices (FDID 200-2FF).
- CUADD=3 allows access to the 4th set of 256 devices (FDID 300-3FF).
- CUADD 0-F parameters can be defined according to the FDIDs which are defined. All logical control units are not required to be defined.

Note: To access all 4096 devices, all 16 VCUs must be defined.

***Performance Impact* - UNIT=type**

UNIT=3990

- Do not code 3990-3 because HCD does not allow the 256 parameter on the UNITADD keyword.
- Do NOT code 3390-6 as the host limits the logical paths to 128.
- UNIT=3990 is required.

IO DEVICE Statement

- Specifies I/O device addresses.

- Describes the device characteristic.
- Designates the assigned control units.
- Keywords:

ADDRESS=(address,number)

- Defines a unique (host) device address, 1-4 digit hexadecimal (0000-FFFF). Host 3-digit or 4-digit addresses can be used.
- Indicates the number of sequential device numbers, a decimal value in the range of 1-256. Specify either the full range of addresses designated in the CNTLUNIT statement or the sequential number of defined, functional devices.
- Defines the number of devices (unit control blocks) addressed by the host.
- The unit address for a device defaults to the last two digits in the address parameter of the ADDRESS keyword, unless the UNITADD keyword is specified.

UNITADD=address

- Designates the logical unit address transmitted on the ESCON channel path to select the I/O device, two hexadecimal digits in the range of 00-FF.
- MUST be designated if the last two digits of the ADDRESS parameter do not match the required ESCON address (00-FF).

CUNUMBR=number

- Specifies the control unit numbers assigned to the control units (specified in the CUNUMBR parameter of the CNTLUNIT statement) to which the device is connected.

STADET=Y

For ESCON devices, the default is STADET=Y. The Status Verification Facility is used to provide the system with a means of reporting presentation of device status with inappropriate combinations of bits, called device-status check.

Unit=device type

This defines the device type parameter as either 3380 or 3390.

PARTITION=(name,...)

- Optional keyword
- Identifies the name of logical partitions (LPAR) permitted to access the channel path to the device.

- Restricts unnamed partitions from accessing a device, even if the LPAR can access a channel path assigned to the device.

Device Parameters And Features

Supported device parameters and features are described in the following section.

From the HCD panel titled “Define Device Parameters / Features,” enter the appropriate values (YES or NO) for the site environment.

DYNAMIC – Specifies device eligibility for dynamic I/O configuration.

LOCANY – Specifies UCB capability to reside in 31 bit storage.

OFFLINE – Specifies device to be online or offline at IPL (optional).

SHARED – Device shared with other systems.

Note: SHARED is recommended for proper functioning of the RESERVE/ RELEASE mechanism if devices are shared between LPARs or processors. This parameter is applicable in a MVS environment only.

EXAMPLE

In the following example, device numbers 1000 to 10FF are defined with device type 3390 and logical unit addresses 00 to FF. Each device is connected to control unit number 1000 and can be shared between operating systems. Each volume is online at IPL with dynamic configuration support and capability for its UCB to reside beyond the 16 MB line.

```
IODEVICE ADDRESS=(1000,256),UNITADD=00,CUNUMBR=1000,STADET=Y,UNIT=3390,
FEATURE=SHARED,OFFLINE=YES,DYNAMIC=YES,LOCANY=NO
```

ESCON Director Definition

Define the ESCON Director (ESCD) as a control unit and a device to allow:

- channel commands sent from a host control program to an ESCD , and
- error information sent from an ESCD to a host.

For example, the IOCP statements to define an ESCD are the following:

```
CNTLUNIT CUNUMBR=1A0,
    UNIT=SWCH,
```

```
        UNITADD=((00,1)),
        PATH=(10,20),
        LINK=(FE,FE)
IODEVICE CUNUMBR=1A0,
        UNIT=SWCH,
        ADDRESS=A00,
        UNITADD=00
```

Specify FE, the control unit port, as the LINK parameter on the CNTLUNIT statement. The control unit port (FE) is the terminating link address for each channel path attached to this ESCON Director. This port supports the communication between the ESCD and the host, and, therefore, cannot be blocked, prohibited, or a dedicated connection.

Restriction -- Be aware of the following machine limit:

> The maximum number of devices that can be supported on an ESCON channel is 1024. The HCD and the processor set this limit. When a subsystem with 1024 addressing is configured to a channel, then an entry is unavailable for any other addresses.

The control unit port of an ESCON Director can be configured on any CHPID connected from a processor to that ESCD. However, the 1024 limitation cannot be exceeded.

- Define the control unit port to a CHPID that does not have 1024 addresses assigned.
- If no other CHPIDs are attached to the ESCD from the processor being defined except for the 1024 address ones, only 1023 addresses can be configured to the SVA.

Esoteric

An eligible device table (EDT) is a named collection of the I/O devices eligible for allocation. The EDT includes the esoteric and generic names associated with these devices.

An esoteric names a grouping of I/O devices, usually from the same class and type. The data characteristics and allocation requirements at a site might determine esoteric group selections. For example, an esoteric might be assigned to a device group for allocation of temporary data sets. Using an esoteric, rather than a specific device number or a generic name, increases the prospect of satisfying an allocation request. Waiting for serialized access to a device delays allocation.

A generic name associates a group of devices with the same characteristics. For example, each 3390 device type is grouped with a generic name of 3390. A job specifying 3390 could allocate any device from that group.

Using z/OS HCD, the EDT and device group information is defined for the IODF. During IPL or dynamic configuration, the EDT used by the z/OS operating system is identified and built. After IPL, jobs can request device allocation from any of the device groups assigned to the selected EDT. Any of the esoteric names might be specified in the UNIT parameter of the job control language (JCL) DD statement. Based on this name, the job allocates a device assigned to the group.

HCD/IOCP Esoteric Parameters:

UNITNAME Statement

Specifies the esoteric

NAME = esoteric name

- Specifies the name of the esoteric, up to 8 alphanumeric characters.
- Do not use reserved esoteric names (SYSALLDA).
- UNIT=((devnum,n),devnum,n)..)
- Identifies a (host) device number, a hexadecimal address (0000-FFFF), defined in the configuration.
- Indicates the number of sequential device numbers.

TOKEN=token value

- Allows controlling the order of esoteric names in the EDT (optional-all or none).
- Avoids access problems for data sets using esoteric names in the catalogs.
- Permits flexibility from the alphabetical order, including all or none of the esoteric names.
- Specifies a number in the range 1 to 8999.

VIO=value

- States device eligibility for VIO with a yes or no value.
- Set to a value of YES requires the esoteric to include at least one DASD device type.
- VIO (virtual I/O) refers to temporary data set allocations in paging storage only.

The following example assigns 4 device numbers 0900 through 0903 to esoteric group TEMP01, which is not eligible for VIO.

```
UNITNAMENAME=TEMP01,
            UNIT=((0900,4)),
            VIO=NO
```

For further information, refer to the following IBM manuals:

OS/390 HCD User's Guide Document Number: SC28-1848

OS/390 HCD Planning Document Number: GC28-1750

Note: IBM OS/390 DFSMS and other host software products offer device allocation controls. Refer to the respective vendor manuals for further information.

SVA with FICON

Native FICON is not currently supported.

FICON channel in FCV mode allows the accessing of ESCON interface control units by the use of a 9032-5 ESCON Director FICON Bridge adapter. Processor support of the FICON channel is required.

Channel device addressing support

An ESCON channel operating in CNC mode supports 1024 device addresses. This device address support increases up to 16 K for a FICON channel operating in FCV mode. However, the path between the ESCD and the SVA ESCON interface supports 1024 devices.

FICON (FCV mode) configuration rule

- A logical CU and device cannot be accessed more than once from the same channel path (CHPID). Define a channel to be used once to a CU or logical control unit.
- The same channel might access different logical control units (CUADDs) within the same physical control unit subsystem.

I/O Data Transfers

- One FICON channel in FCV mode can be used for up to eight concurrent I/O transfers from the same physical control unit, where each transfer is from a device in a different logical control unit.
- One ESCON (CNC) channel path can be used for only one I/O transfer at a time.

The I/O performance to a device over a FICON (FCV mode) channel is similar to an ESCON channel.

Aggregation of ESCON Channel Paths

To convert the host paths on a set of logical control units from ESCON (CNC) to FICON (FCV), of one FICON to four ESCON channels do

NOT exceed a combined ESCON channel utilization of 200%. If each ESCON channel peak utilization is 20% to 25%, then the workload on 6 to 8 ESCON (CNC) to 1 FICON (FCV) might be a candidate. Connectivity of both CNC and FCV channels to the same logical control unit is allowed.

This subsystem configuration includes connectivity to the host system using FICON channels to the FICON Bridge. The following example has one SVA subsystem defined with 2048 devices. In this scenario, the FCV type channels have 2048 devices defined. Adhering to device support limitations and performance guidelines, the FCV channels might be defined to additional logical control units.

For HCD definitions, refer to the other examples in this appendix. The HCD definition is similar except for the channel type. The IOCP report from HCD is included in this example.

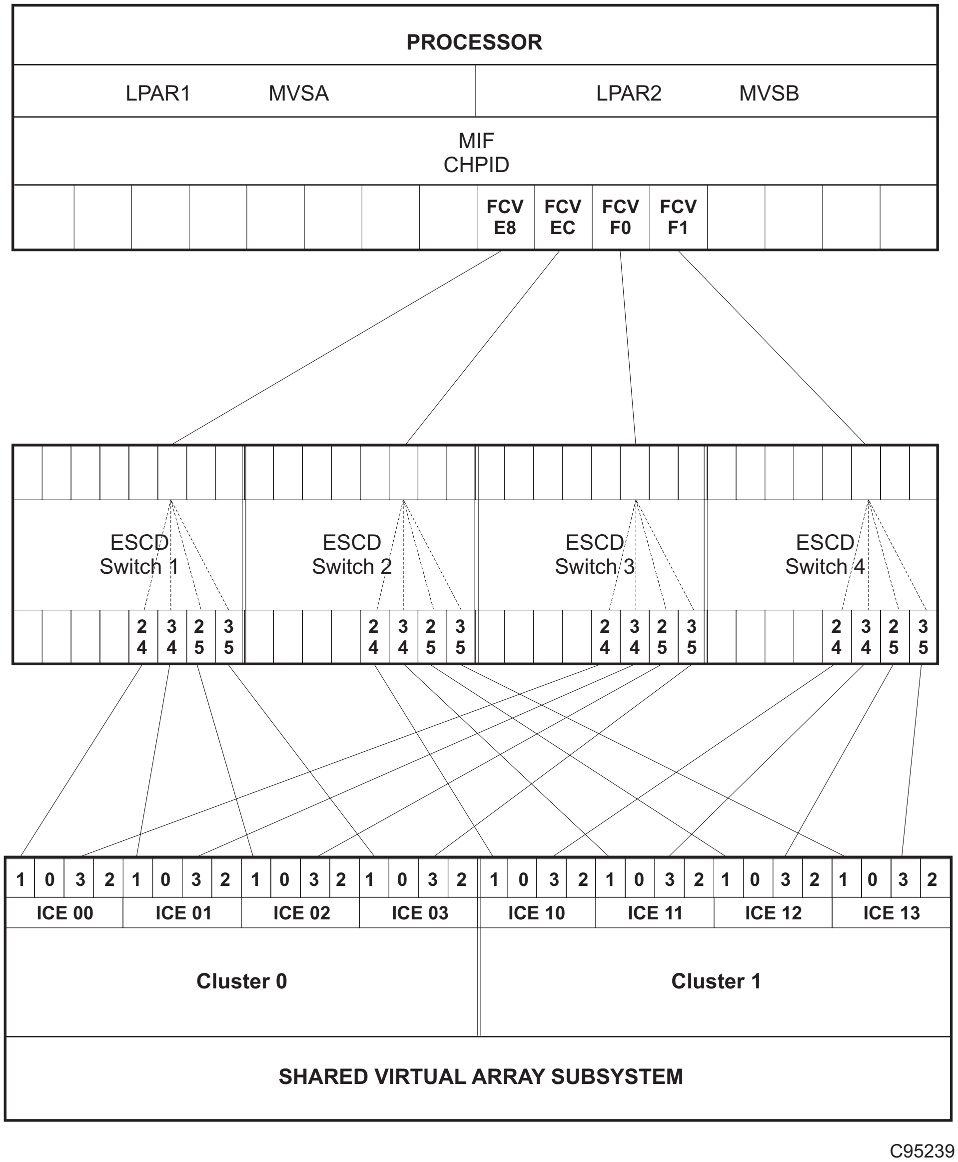

FICON Configuration Example

Example IOCP report:

```
ID SYSTEM = (9672,6)

RESOURCE PARTITION=((MVSA,1),MVSB,2))

CHPID PATH=(E8),SWITCH=01,TYPE=FCV,SHARED

CHPID PATH=(EC),SWITCH=02,TYPE=FCV,SHARED

CHPID PATH=(F0),SWITCH=03,TYPE=FCV,SHARED
```

```
CHPID PATH=(F1),SWITCH=04,TYPE=FCV,SHARED
CNTLUNIT CUNUMBR=1000,PATH=(E8,EC,F0,F1),UNITADD=((00,256)),  *
               LINK=(24,24,24,24),CUADD=0,UNIT=3990
IODEVICE ADDRESS=(1000,256), UNITADD=00,CUNUMBR=(1000),       *
               STADET=Y,UNIT=3390
CNTLUNIT CUNUMBR=1100,PATH=(E8,EC,F0,F1),UNITADD=((00,256)),  *
               LINK=(34,34,34,34),CUADD=1,UNIT=3990
IODEVICE ADDRESS=(1100,256),UNITADD=00,CUNUMBR=(1100),        *
               STADET=Y,UNIT=3390
CNTLUNIT CUNUMBR=1200,PATH=(E8,EC,F0,F1),UNITADD=((00,256)),  *
               LINK=(24,24,24,24),CUADD=2,UNIT=3990
IODEVICE ADDRESS=(1200,256),UNITADD=00,CUNUMBR=(1200),        *
               STADET=Y,UNIT=3390
CNTLUNIT CUNUMBR=1300,PATH=(E8,EC,F0,F1),UNITADD=((00,256)),  *
               LINK=(34,34,34,34),CUADD=3,UNIT=3990
IODEVICE ADDRESS=(1300,256),UNITADD=00,CUNUMBR=(1300),        *
               STADET=Y,UNIT=3390
CNTLUNIT CUNUMBR=1400,PATH=(E8,EC,F0,F1),UNITADD=((00,256)),  *
               LINK=(25,25,25,25),CUADD=4,UNIT=3990
IODEVICE ADDRESS=(1400,256),CUNUMBR=(1400),STADET=Y,          *
               STADET=Y,UNIT=3390
CNTLUNIT CUNUMBR=1500,PATH=(E8,EC,F0,F1),UNITADD=((00,256)),  *
               LINK=(35,35,35,35),CUADD=5,UNIT=3990
IODEVICE ADDRESS=(1500,256),UNITADD=00,CUNUMBR=(1500),        *
               STADET=Y,UNIT=3390
CNTLUNIT CUNUMBR=1600,PATH=(E8,EC,F0,F1),UNITADD=((00,256)),  *
               LINK=(25,25,25,25),CUADD=6,UNIT=3990
IODEVICE ADDRESS=(1600,256),UNITADD=00,CUNUMBR=(1600),        *
               STADET=Y,UNIT=3390
CNTLUNIT CUNUMBR=1700,PATH=(E8,EC,F0,F1),UNITADD=((00,256)),  *
               LINK=(35,35,35,35),CUADD=7,UNIT=3990
IODEVICE ADDRESS=(1700,256),UNITADD=00,CUNUMBR=(1700),        *
```

```
                  STADET=Y,UNIT=3390
9032-5 definition:
CNTLUNIT CUNUMBR=D001,PATH=(E8),UNITADD=((00,001)),LINK=(FE), *
                  UNIT=9032-5
CNTLUNIT CUNUMBR=D002,PATH=(EC),UNITADD=((00,001)),LINK=(FE), *
                  UNIT=9032-5
CNTLUNIT CUNUMBR=D003,PATH=(F0),UNITADD=((00,001)),LINK=(FE), *
                  UNIT=9032-5
CNTLUNIT CUNUMBR=D004,PATH=(F1),UNITADD=((00,001)),LINK=(FE), *
                  UNIT=9032-5
IODEVICE ADDRESS=D001,MODEL=5,UNITADD=00,CUNUMBR=(D001),      *
                  STADET=Y,UNIT=9032
IODEVICE ADDRESS=D002,MODEL=5,UNITADD=00,CUNUMBR=(D002),      *
                  STADET=Y,UNIT=9032
IODEVICE ADDRESS=D003,MODEL=5,UNITADD=00,CUNUMBR=(D003),      *
                  STADET=Y,UNIT=9032
IODEVICE ADDRESS=D004,MODEL=5,UNITADD=00,CUNUMBR=(D004),      *
                  STADET=Y,UNIT=9032
```

Examples:

Use the following examples to develop the SVA configuration for a specific site environment. All examples include a description, cable diagrams, IOCP, and HCD.

Example 1 - SVA With ESCON Directors And 8-Path VCU

This scenario illustrates the HCD panels and the respective IOCP statements to define a SVA. A processor with 4 LPARs sharing 8 channels (EMIF) and 2 ESCON directors are included in the configuration.

- 8 ESCON channels are connected to this subsystem via 2 ESCON directors (ESCD).
  - Channels 40, 42, 44, 46 are connected to Cluster 0.
  - Channels 41, 43, 45, 47 are connected to Cluster 1.
  - The channels are being shared by 4 partitions (EMIF) on 1 processor.
  - 8 ICE3 cards, 4 links per card.

- 16 logical paths per link are available, 512 logical paths total.
- 4 logical paths per link (one for each LPAR) might be established in this example.
- Device addresses (host) are 6000-63FF.
- Device types are defined as 3390.
- 1024 of the 4096 functional devices are defined.

• Define unique ESCON CHPIDs for each set of 1024 devices (host limit). If all devices are defined, assign a different collection of CHPIDs to each set of 4 LCUs:

CUADD = 0 1 2 3

CUADD = 4 5 6 7

CUADD = 8 9 A B

CUADD = C D E F

•For explanation purposes, the example displays the configuration for CUADD = 0, 1, 2, and 3. If the customer requirement is to define 1024 devices, then this configuration could be used.

• Figure 12 on page 137 illustrates the ESCON channel and ICE card configuration for 1024 devices only.

• To expand this example to 4096 devices, abide by the ESCON limitations of 1024 devices per channel and define up to sixteen LCUs.

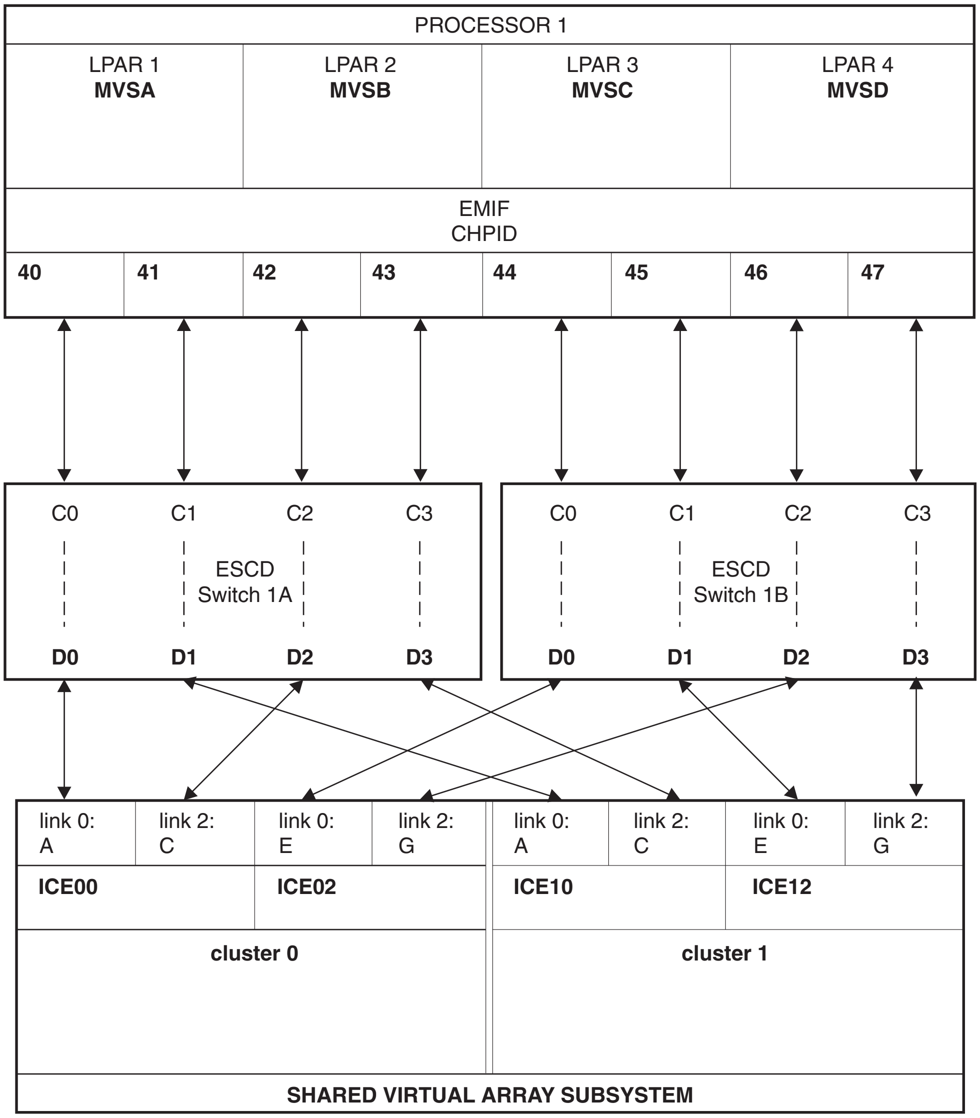

All ESCON Channels	
BASE	00
RANGE	4096
BFDID	00

C95233

Figure 12 Channel and ICE Card Configuration

IOCP Statements for Example 1 - 8 Paths To All Devices

Table 25 IOCP Statements

RESOURCE PARTITION=((MVSA,1),(MVSB,2),(MVSC,3),(MVSD,4))
CHPID
PATH=(40,41,42,43),TYPE=CNC,SWITCH=1A,PARTITION=(MVSA, MVSB),SHARED
PATH=(44.45.46.47),TYPE=CNC,SWITCH=1B,PARTITION=(MVSA, MVSB),SHARED
CNTLUNIT
CUNUMBR=6000,PATH=(40,41,42,43,44,45,46,47),LINK=(D0, D1, D2, D3, D0, D1, D2, D3), UNITADD=((00,256)),CUADD=0,UNIT=3990
CUNUMBR=6100,PATH=(40,41,42,43,44,45,46,47),LINK=(D0, D1, D2, D3, D0, D1, D2, D3), UNITADD=((00,256)),CUADD=1,UNIT=3990
CUNUMBR=6200,PATH=(40,41,42,43,44,45,46,47),LINK=(D0, D1, D2, D3, D0, D1, D2, D3), UNITADD=((00,256)),CUADD=2,UNIT=3990
CUNUMBR=6300,PATH=(40,41,42,43,44,45,46,47),LINK=(D0, D1, D2, D3, D0, D1, D2, D3), UNITADD=((00,256)),CUADD=3,UNIT=3990
IODEVICE
ADDRESS=(6000,256),UNITADD=00,CUNUMBR=6000,STADET=Y,UNIT=3390,FEATURE=SHARED
ADDRESS=(6100,256),UNITADD=00,CUNUMBR=6100,STADET=Y,UNIT=3390,FEATURE=SHARED
ADDRESS=(6200,256),UNITADD=00,CUNUMBR=6200,STADET=Y,UNIT=3390,FEATURE=SHARED
ADDRESS=(6300,256),UNITADD=00,CUNUMBR=6300,STADET=Y,UNIT=3390,FEATURE=SHARED

HCD Example

The definition method in the following example is driven from the Define, Modify, or View Configuration Data panel. Explanation of HCD panel navigation is beyond the scope of this document. Definitions for channel paths, control units, and devices of a SVA are shown.

Channel path definition for CHPIDs 40-43:

```
Add Channel Path

Specify or revise the following values.

Processor ID...: P1 Processor One

Configuration mode: LPAR

Channel path ID. . . . . . . . 40 +

Number of CHPIDs . . . . . . . 4

Channel path type. . . . . . . CNC +

Operation mode . . . . . . . . SHR +

Description. . . . . . . . . . Shared Virtual Array
with EMIF
```

```
Specify the following values only if connected to a
switch:

Dynamic switch ID  . . . . 1A + (00 - FF)

Entry switch ID  . . . . . 1A +

Entry port . . . . . . . . C0 +
```

Channel path definition for CHPIDs 44-47:

```
Add Channel Path

Specify or revise the following values.

Processor ID . . . : P1  Processor One

Configuration mode : LPAR

Channel path ID. . . . . . . . 44 +

Number of CHPIDs . . . . . . . 4

Channel path type. . . . . . . CNC +

Operation mode . . . . . . . . SHR +

Description  . . . . . . . . . Shared Virtual Array
with EMIF

Specify the following values only if connected to a
switch:

Dynamic switch ID  . . . . . . 1B + (00 - FF)

Entry switch ID  . . . . . . . 1B +

Entry port . . . . . . . . . . C0 +
```

Director entry port definitions:

```
Update Switch Connections

Specify or revise the following values.

Processor ID . . . : P1

        Dyn +      -- Entry +--

CHPID  Switch      Switch Port

40     1A          1A     C0

41     1A          1A     C1

42     1A          1A     C2

43     1A          1A     C3

44     1B          1B     C0

45     1B          1B     C1
```

```
46      1B          1B      C2

47      1B          1B      C3
```

Note: Specify the switch and entry port of each channel. Match the physical connections between the channels and the directors.

LPAR access list definition:

```
Define Access List

Select one or more partitions for inclusion in the
access list.

Channel path ID. . : 40     Channel path type . . :
CNC

Operation mode . . : SHR      Number of CHPIDs. . .
: 8

/ Partition Name      Number         Usage Descrip-
tion

/ MVSA                 1               OS    LPAR One

/ MVSB                2               OS    LPAR Two

/ MVSC               3               OS    LPAR Three

/ MVSD                4               OS    LPAR Four
```

Note: Select the partitions to be allowed access to the channel paths.

Logical control unit definition:

```
Add Control Unit

Specify or revise the following values.

Control unit number. . : 6000  +  Type . . : 3990  +

Processor ID . . . . . : P1  Processor One

Channel path IDs . . . . 40 41 42 43 44 45 46 47 _ _
_+

Link address . . . . . . D0 D1 D2 D3 D0 D1 D2 D3 _ _
_+

Unit address . . . . . . 00 __ __ __ __ __ __ __+

Number of units. . . . . 256_ __ __ __ __ __ __ __

Logical address. . . . . 0 + (same as CUADD)

Protocol . . . . . . . . . _ + (D, S, or S4)

I/O concurrency level. . _ + (1, 2, or 3)
```

Note: Define CUADD = 0, 1, 2, 3 with the same CHPIDs. Assign unique control unit numbers.

Device definition for host addresses 6000-6FFF:

```
Add Device

Specify or revise the following values.

Device number . . . . . . . . . 6000 (0000 - FFFF)

Number of devices . . . . . . . 256__

Device type . . . . . . . . . . 3390___________ +

Serial number . . . . . . . . . _______________

Description . . . . . . . . . . _______________

Connected to CUs. . . 6000 ___ ___ ___ ___ ___ ___ +
```

Note: Repeat for device numbers 6100, 6200, and 6300, which are connected to control unit numbers 6100, 6200, and 6300.

```
Device / Processor Definition

Select processor(s) to change device/processor def-
initions, then press Enter.

Device number . . : 6000      Number of devices . :
256

Device type . . . : 3390/Processor ID  UA +  Time-
Out  STADET  Preferred Explicit Device CHPID + Can-
didate List

/ P1            00    No        Yes     __    No
```

Notes:

1. UA is the unit address keyword.
2. Repeat for device numbers 6X00, where X is 1 - 3 for 1024 devices, or 1 - F for 4096 devices.

Example 2 - SVA With 2 Processors, ESCON Directors, And 4-Path VCU

This subsystem configuration includes connectivity to 2 physical processors, each with 2 LPARs. The channel configuration includes 16 ESCON channels (EMIF) from each processor via 2 ESCON directors. A collection of 8 paths accesses four LCUs with a total of 1024 devices.

***Caution: Performance Impact* - A maximum number of 8 channel paths to a device can be configured from each operating system. This is a host processor and HCD/IOCP limit.**

The following example illustrates the configuration for Processor 1 and 2.

The HCD and IOCP are shown for Processor 1. The definitions would be similar for Processor 2.

- 16 ESCON channels (EMIF) from processor 1 are connected to this subsystem via 2 ESCON directors (ESCD).
  - Channels 58, 59, 5C, 5D, 60, 61, 64, and 65 are connected to Cluster 0. Respective outbound ESCD links are SW2-C0, SW1-C0, SW2-C2, SW1-C2, SW2-C4, SW1-C4, SW2-C6, and SW1-C6.
  - Channels 5A, 5B, 5E, 5F, 62, 63, 66, and 67 are connected to Cluster 1. Respective outbound ESCD links are SW2-C1, SW1-C1, SW2-C3, SW1-C3, SW2-C5, SW1-C5, SW2-C7, and SW1-C7.
  - Channels 58, 5A, 5C, and 5E accesses 1024 functional devices 000 - 3FF.
  - Channels 59, 5B, 5D, and 5F accesses 1024 functional devices 400 - 7FF.
  - Channels 60, 62, 64, and 66 accesses 1024 functional devices 800 - BFF.
  - Channels 61, 63, 65, and 67 accesses 1024 functional devices C00 - FFF.
- 8 ICE3 cards (32 total links) The maximum availability of logical host paths is 16 per ICE card link, 512 total. This HCD configuration allows for 4 host paths per ICE card link.

Caution: *Performance Impact* - Do not exceed the total host logical paths, which might be established from all processors. Configure CHPIDs from the same LPAR to different links and only one of the top and one of the bottom 2 links on each ICE3 card.

- Device addresses (host) are 1000-1FFF.
  - Device types are defined as 3390-3.
  - Functional devices (FDIDs) 000-FFF, the maximum number of 4096 devices, are defined.

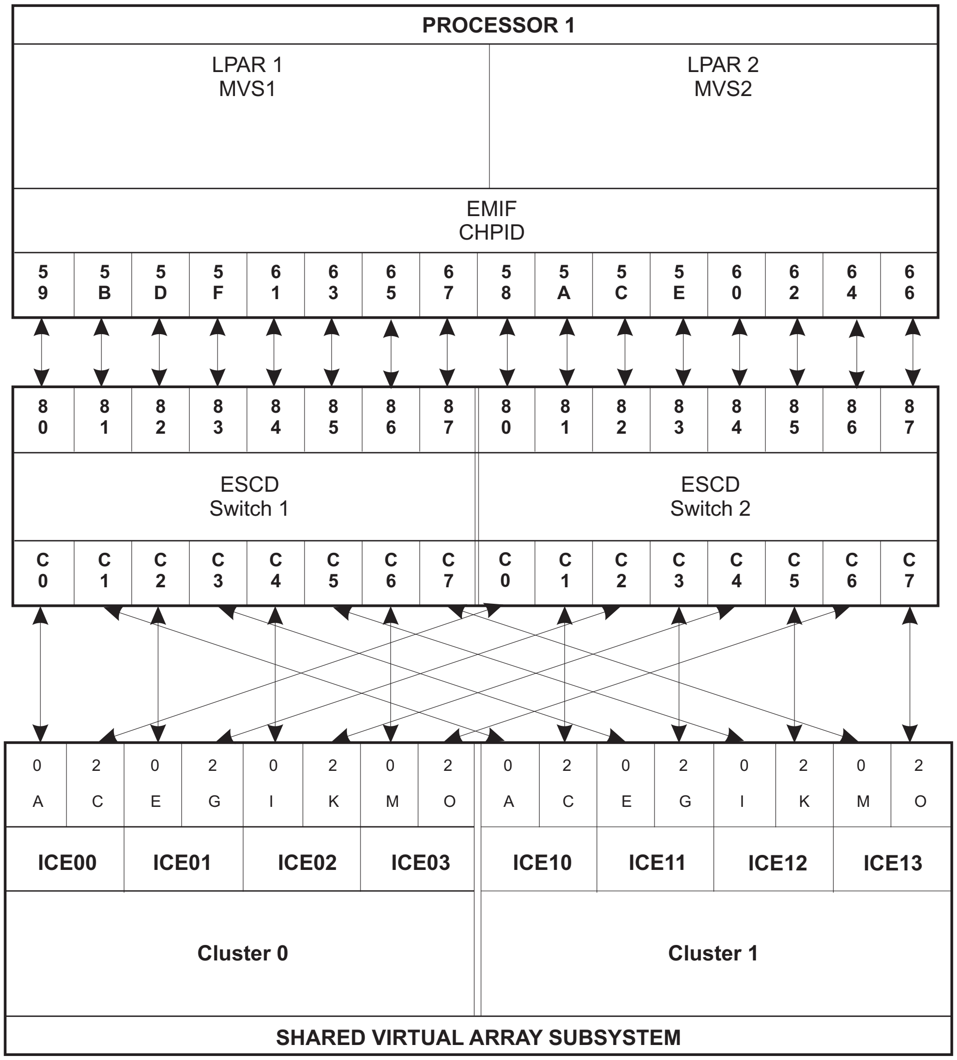

All ESCON Channels	
BASE	00
RANGE	4096
BFDID	00

C95237

Figure 13 Channel And ICE Card Configuration For Example 2

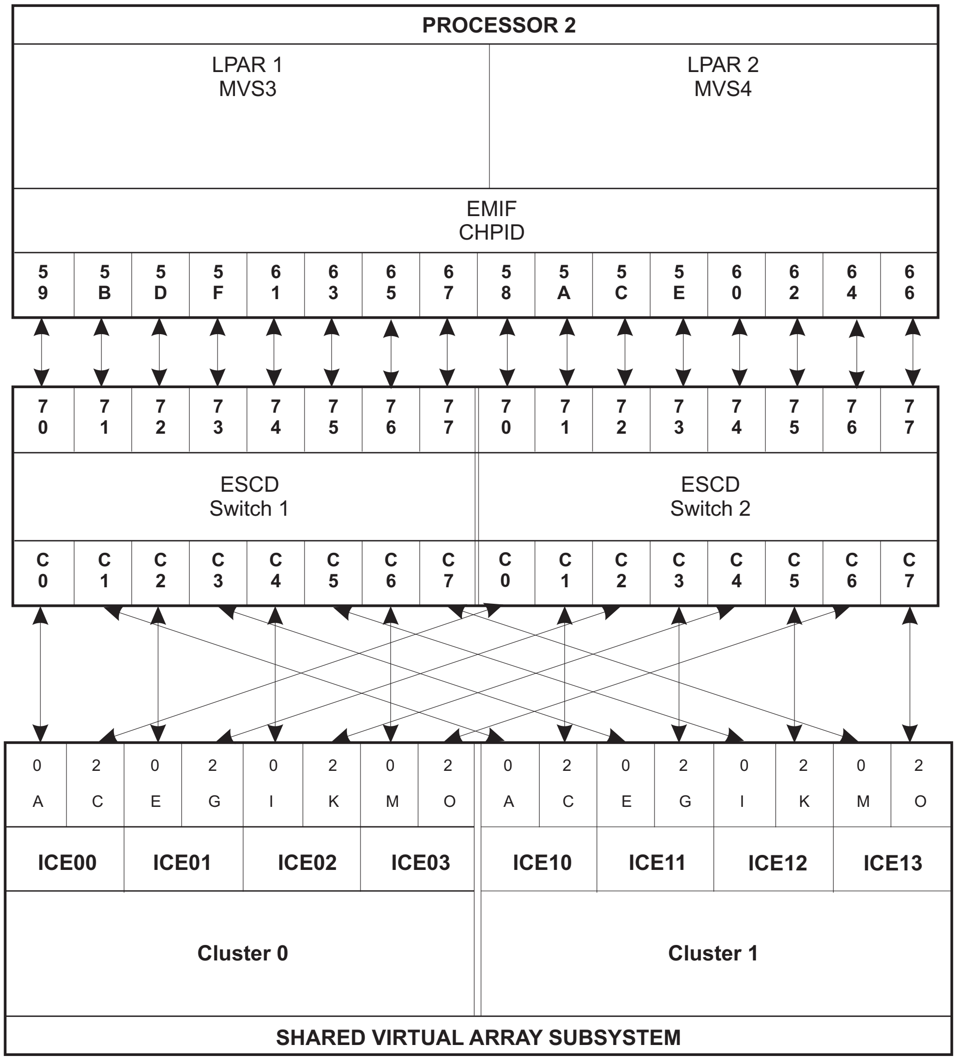

All ESCON Channels	
BASE	00
RANGE	4096
BFDID	00

C95236

Figure 14 Channel And ICE Card Configuration For Example 2

Note: Figure 13 on page 143 and Figure 14 display the same SVA subsystem with connectivity to two processors.

IOCP Statements

Table 26 IOCP Statements

RESOURCE PARTITION=((MVS1,1),(MVS2,2))
CHPID
PATH=(59, 5B, 5D, 5F, 61, 63, 65, 67),TYPE=CNC,SWITCH=01,PARTITION=(MVS1,MVS2),SHARED
PATH=(58, 5A, 5C, 5E, 60, 62, 64, 66),TYPE=CNC,SWITCH=02,PARTITION=(MVS1,MVS2),SHARED
CNTLUNIT
CUNUMBR=1000,PATH=(58,5A,5C,5E),LINK=(C0,C1,C2,C3),UNITADD=((00,256)),CUADD=0,UNIT=3990
CUNUMBR=1100,PATH=(58,5A,5C,5E),LINK=(C0,C1,C2,C3),UNITADD=((00,256)),CUADD=1,UNIT=3990
CUNUMBR=1200,PATH=(58,5A,5C,5E),LINK=(C0,C1,C2,C3),UNITADD=((00,256)),CUADD=2,UNIT=3990
CUNUMBR=1300,PATH=(58,5A,5C,5E),LINK=(C0,C1,C2,C3),UNITADD=((00,256)),CUADD=3,UNIT=3990
CUNUMBR=1400,PATH=(59,5B,5D,5F),LINK=(C0,C1,C2,C3),UNITADD=((00,256)),CUADD=4,UNIT=3990
CUNUMBR=1500,PATH=(59,5B,5D,5F),LINK=(C0,C1,C2,C3),UNITADD=((00,256)),CUADD=5,UNIT=3990
CUNUMBR=1600,PATH=(59,5B,5D,5F),LINK=(C0,C1,C2,C3),UNITADD=((00,256)),CUADD=6,UNIT=3990
CUNUMBR=1700,PATH=(59,5B,5D,5F),LINK=(C0,C1,C2,C3),UNITADD=((00,256)),CUADD=7,UNIT=3990
CUNUMBR=1800,PATH=(60,62,64,66),LINK=(C4,C5,C6,C7),UNITADD=((00,256)),CUADD=8,UNIT=3990
CUNUMBR=1900,PATH=(60,62,64,66),LINK=(C4,C5,C6,C7),UNITADD=((00,256)),CUADD=9,UNIT=3990
CUNUMBR=1A00,PATH=(60,62,64,66),LINK=(C4,C5,C6,C7),UNITADD=((00,256)),CUADD=A,UNIT=3990
CUNUMBR=1B00,PATH=(60,62,64,66),LINK=(C4,C5,C6,C7),UNITADD=((00,256)),CUADD=B,UNIT=3990
CUNUMBR=1C00,PATH=(61,63,65,67),LINK=(C4,C5,C6,C7),UNITADD=((00,256)),CUADD=C,UNIT=3990
CUNUMBR=1D00,PATH=(61,63,65,67),LINK=(C4,C5,C6,C7),UNITADD=((00,256)),CUADD=D,UNIT=3990
CUNUMBR=1E00,PATH=(61,63,65,67),LINK=(C4,C5,C6,C7),UNITADD=((00,256)),CUADD=E,UNIT=3990
CUNUMBR=1F00,PATH=(61,63,65,67),LINK=(C4,C5,C6,C7),UNITADD=((00,256)),CUADD=F,UNIT=3990
IODEVICE
ADDRESS=(1000,256),UNITADD=00,CUNUMBR=1000,STADET=Y,UNIT=3390,FEATURE=SHARED
ADDRESS=(1100,256),UNITADD=00,CUNUMBR=1100,STADET=Y,UNIT=3390,FEATURE=SHARED
ADDRESS=(1200,256),UNITADD=00,CUNUMBR=1200,STADET=Y,UNIT=3390,FEATURE=SHARED
ADDRESS=(1300,256),UNITADD=00,CUNUMBR=1300,STADET=Y,UNIT=3390,FEATURE=SHARED
ADDRESS=(1400,256),UNITADD=00,CUNUMBR=1400,STADET=Y,UNIT=3390,FEATURE=SHARED
ADDRESS=(1500,256),UNITADD=00,CUNUMBR=1500,STADET=Y,UNIT=3390,FEATURE=SHARED
ADDRESS=(1600,256),UNITADD=00,CUNUMBR=1600,STADET=Y,UNIT=3390,FEATURE=SHARED
ADDRESS=(1700,256),UNITADD=00,CUNUMBR=1700,STADET=Y,UNIT=3390,FEATURE=SHARED
ADDRESS=(1800,256),UNITADD=00,CUNUMBR=1800,STADET=Y,UNIT=3390,FEATURE=SHARED
ADDRESS=(1900,256),UNITADD=00,CUNUMBR=1900,STADET=Y,UNIT=3390,FEATURE=SHARED

Table 26 IOCP Statements (Continued)

ADDRESS=(1A00,256),UNITADD=00,CUNUMBR=1A00,STADET=Y,UNIT=3390,FEATURE=SHARED
ADDRESS=(1B00,256),UNITADD=00,CUNUMBR=1B00,STADET=Y,UNIT=3390,FEATURE=SHARED
ADDRESS=(1C00,256),UNITADD=00,CUNUMBR=1C00,STADET=Y,UNIT=3390,FEATURE=SHARED
ADDRESS=(1D00,256),UNITADD=00,CUNUMBR=1D00,STADET=Y,UNIT=3390,FEATURE=SHARED
ADDRESS=(1E00,256),UNITADD=00,CUNUMBR=1E00,STADET=Y,UNIT=3390,FEATURE=SHARED
ADDRESS=(1F00,256),UNITADD=00,CUNUMBR=1F00,STADET=Y,UNIT=3390,FEATURE=SHARED

HCD Example

The definition method in the following example is driven from the Define, Modify, or View Configuration Data panel. Explanation of HCD panel navigation is beyond the scope of this document. Definitions for channel paths, control units, and devices are shown in the following panels.

Channel path definition for CHPID 59:

```
Add Channel Path

Specify or revise the following values.

Processor ID . . . : P1 Processor One

Configuration mode: LPAR

Channel path ID. . . . . . . . 59 +

Number of CHPIDs . . . . . . . 1

Channel path type. . . . . . . CNC +

Operation mode . . . . . . . . SHR +

Description. . . . . . . . . . Shared Virtual Array with
EMIF

Specify the following values only if connected to a
switch:

Dynamic switch ID  . . . . 01 + (00 - FF)

Entry switch ID  . . . . . 01 +

Entry port . . . . . . . . 80 +
```

Note: Repeat for channels 5B 5D 5F 61 63 65 67 connected to switch id 01 (ESCON director). The associated entry ports are 81-87. The switch must be defined prior to connecting the channel path to the switch. A group of channel paths with consecutive identifiers and the same type and mode might be defined in one step by entering a value for Number of CHPIDs. The attributes for each CHPID might be modified in another panel.

Channel path definition for CHPID 58:

```
Add Channel Path

Specify or revise the following values.

Processor ID . . . : P1  Processor One

Configuration mode : LPAR

Channel path ID. . . . . . . . 58 +

Number of CHPIDs . . . . . . . 1

Channel path type. . . . . . . CNC +

Operation mode . . . . . . . . SHR +

Description  . . . . . . . . . Shared Virtual Array
with EMIF

Specify the following values only if connected to a
switch:

Dynamic switch ID  . . . . . . 02 + (00 - FF)

Entry switch ID  . . . . . . . 02 +

Entry port . . . . . . . . . . 80 +
```

Note: Repeat for channels 5A 5C 5E 60 62 64 66 connected to switch id 02 (ESCON Director). The associated entry ports are 81-87.

Director (switch) entry port definition:

```
Update Switch Connections

Specify or revise the following values.

Processor ID . . . : P1

Dyn         + --        Entry       +--

CHPID       Switch      Switch      Port

58          02          02          80
```

Note: This panel allows definition of the entry ports for all subsequent channel paths if a group of channel paths have been defined in one step. The switch and port values displayed initially are only valid for the first channel path.

LPAR access list definition:

```
Define Access List

Select one or more partitions for inclusion in the
access list.

Channel path ID. . : 58 Channel path type . . : CNC

Operation mode . . : SHR Number of CHPIDs. . . : 1
```

/ Partition Name	Number	Usage	Description
/ MVS1	1	OS	LPAR One
/ MVS2	2	OS	LPAR Two

Note: Each channel is shared, accessed by more than one logical partition at the same time (EMIF).

Channel path attribute listing:

```
Channel Path List Row 1 of 10 More: >

Command===>
```

```
     Select one or more channel paths, then press
Enter. To add use F11.

    Processor ID . . . : P1

    Configuration mode : LPAR


                Dyn + --Entry +--

    / CHPID Type + Mode + Switch Switch Port Con.
Description

    _ 58    CNC    SHR    02    02    80

    _ 59    CNC    SHR    01    01    80

    _ 5A    CNC    SHR    02    02    81

    _ 5B    CNC    SHR    01    01    81

    _ 5C    CNC    SHR    02    02    82

    _ 5D    CNC    SHR    01    01    82

    _ 5E    CNC    SHR    02    02    83

    _ 5F    CNC    SHR    01    01    83

    _ 60    CNC    SHR    02    02    84

    _ 61    CNC    SHR    01    01    84

******************************* Bottom of data
```

Note: Type over the appropriate values to modify the channel path attributes, except the channel path ID.

Channel Path List Row 1 of 10 More: <

```
Command ==>
```

```
Select one or more channel paths, then press Enter. To add, use
F11.

   1=MVS 1    2=MVS2     3=         4=         5=
   6=         7=         8=         9=         A=
   B=         C=         D=         E=         F=

                          --------------- Partitions -
  / CHPID Type + Mode + 1  2  3  4  5  6  7  8  9  A  B  C  D  E  F
  _ 58    CNC    SHR    a  a  _  _  _  _  _  _  _  _
  _ 58    CNC    SHR    a  a  _  _  _  _  _  _  _  _
  _ 59    CNC    SHR    a  a  _  _  _  _  _  _  _  _
  _ 5A    CNC    SHR    a  a  _  _  _  _  _  _  _  _
  _ 5B    CNC    SHR    a  a  _  _  _  _  _  _  _  _
  _ 5C    CNC    SHR    a  a  _  _  _  _  _  _  _  _
  _ 5D    CNC    SHR    a  a  _  _  _  _  _  _  _  _
  _ 5E    CNC    SHR    a  a  _  _  _  _  _  _  _  _
  _ 5F    CNC    SHR    a  a  _  _  _  _  _  _  _  _
  _ 60    CNC    SHR    a  a  _  _  _  _  _  _  _  _
  _ 61    CNC    SHR    a  a  _  _  _  _  _  _  _  _
  ******************************* Bottom of data
```

Note: The letter a or c indicates assignment of a partition to the access (a) or candidate (c) list for each channel path.

Logical control unit definition:

```
Add Control Unit

Specify or revise the following values.

Control unit number. . : 1000   +    Type . . . :
3990  +

Processor ID . . . . . : P1  Processor One

Channel path IDs . . . . 58 5A 5C 5E __ __ __ __ +

Link address . . . . . . C0 C1 C2 C3 __ __ __ __ +

Unit address . . . . . . 00 __ __ __ __ __ __ __ +

Number of units. . . . . 256   __ __ __ __ __ __

Logical address. . . . . 0 + (same as CUADD)

Protocol . . . . . . . . _ + (D, S, or S4)
```

```
I/O concurrency level. . _ + (1, 2, or 3)
```

Note: Define CUADD = 0, 1, 2, and 3 with the same CHPIDs. The CHPIDs rotate between physical cluster 0 and cluster 1 connections to the SVA. Logical control unit CUADD=0 is defined in this panel. Assign unique control unit numbers. The protocol and I/O concurrency level parameters are ignored.

```
Add Control Unit

Specify or revise the following values.

Control unit number. . : 1400   +  Type . . . : 3990
+

Processor ID . . . . . : P1  Processor One

Channel path IDs . . . . 59 5B 5D 5F __ __ __ __ +

Link address . . . . . . C0 C1 C2 C3 __ __ __ __ +

Unit address . . . . . . 00 __ __ __ __ __ __ __ +

Number of units. . . . . 256 _ __ __ __ __ __ __

Logical address. . . . . 0 + (same as CUADD)

Protocol . . . . . . . .  + (D, S, or S4)

I/O concurrency level. .  + (1, 2, or 3)
```

Note: Define CUADD = 4, 5, 6, and 7 with the same CHPIDs. Logical control unit CUADD=4 is defined in this panel.

```
Add Control Unit

Specify or revise the following values.

Control unit number. . : 1800   +    Type . . . : 3990
+

Processor ID . . . . . : P1  Processor One

Channel path IDs . . . . 60 62 64 66 __ __ __ __ +

Link address . . . . . . C4 C5 C6 C7 __ __ __ __ +

Unit address . . . . . . 00 __ __ __ __ __ __ __ +

Number of units. . . . . 256   __ __ __ __ __ __

Logical address. . . . . 2 + (same as CUADD)

Protocol . . . . . . . . . _ + (D, S, or S4)

I/O concurrency level. . _ + (1, 2, or 3)
```

Note: Define CUADD = 8, 9, A, and B with the same CHPIDs. Logical control unit CUADD=8 is defined in this panel.

```
Add Control Unit
```

```
Specify or revise the following values.

Control unit number. . : 1C00   +    Type . . . : 3990
+

Processor ID . . . . . : P1  Processor One

Channel path IDs . . . . 61 63 65 67 __ __ __ __ +

Link address . . . . . . C4 C5 C6 C7 __ __ __ __ +

Unit address . . . . . . 00 __ __ __ __ __ __ __ +

Number of units. . . . . 256 _ __ __ __ __ __ __

Logical address. . . . . 2 + (same as CUADD)

Protocol . . . . . . . .  + (D, S, or S4)

I/O concurrency level. .  + (1, 2, or 3)
```

Note: Define CUAAD = C, D, E, and F with the same CHPIDs. Logical control unit CUADD=C is defined in this panel.

Device definition for host addresses 1000-10FF:

```
Add Device

Specify or revise the following values.

Device number . . . . . . . . . 1000 (0000 - FFFF)

Number of devices . . . . . . . 256__

Device type . . . . . . . . . . 3390___________ +

Serial number . . . . . . . . . _______________

Description . . . . . . . . . . _______________

Connected to CUs. . . 0A00 0A01 __ __ __ __ __ __ +
```

```
Device / Processor Definition

Select  processor(s)  to  change  device/processor  definitions,
then press Enter.

Device number . . : 1000 Number of devices . : 256

Device type . . . : 3390

/ Processor ID   UA +   Time-Out   STADET   Preferred  Explicit
Device

                                        CHPID + Candidate List

/ P1           00    No        Yes      __   No
```

Note: Repeat for device numbers 1x00 where x = 1 through F. UA is the unit address keyword, 00.

Example 3 – SVA With ESCON Directors and 4-Path VCU

This subsystem configuration includes connectivity to one physical processor, with 2 LPARs. The channel configuration includes 16 ESCON channels (EMIF) from the processor via 2 ESCON directors. A collection of 4 paths accesses each set of 1024 devices.

***Caution: Performance Impact -* A maximum number of 8 channel paths to a device can be configured from each operating system. This is a host processor and HCD/IOCP limit.**

The following example illustrates the configuration.

- 16 ESCON channels (EMIF) from the processor are connected to this subsystem via 2 ESCON directors (ESCD).
  - Channels 58, 59, 5C, 5D, 60, 61, 64, and 65 are connected to Cluster 0. Respective outbound ESCD links are SW2-C0, SW1-C0, SW2-C2, SW1-C2, SW2-C4, SW1-C4, SW2-C6, and SW1-C6.
  - Channels 5A, 5B, 5E, 5F, 62, 63, 66, and 67 are connected to Cluster 1. Respective outbound ESCD links are SW2-C1, SW1-C1, SW2-C3, SW1-C3, SW2-C5, SW1-C5, SW2-C7, and SW1-C7.
  - Channels 58, 5D, 5A, and 5F (CNTLUNIT CUADD=0, 1, 2, and 3) accesses functional devices 000 - 3FF.
  - Channels 59, 5C, 5B, and 5E (CNTLUNIT CUADD=4, 5, 6, and 7) accesses functional devices 400 - 7FF.
  - Channels 60, 65, 62, and 67 (CNTLUNIT CUADD=8, 9, A, and B) accesses functional devices 800 - BFF.
  - Channels 61, 64, 63, and 66 (CNTLUNIT CUADD=C, D, E, and F) accesses functional devices C00 - FFF.
- 8 ICE3 cards, quad link (32 total links)

  The maximum availability of host paths is 16 per ICE card link, 512 total. This configuration allows for 2 host paths per ICE card link, 32 total (2 LPARs x 16 links).

  Do not exceed the total host paths, which might be established from all processors.

  Device addresses (host) are 1000-1FFF.

  - Device types are defined as 3390-3.
  - Functional devices (FDIDs) 00-FFF, the maximum number of 4096, are defined.

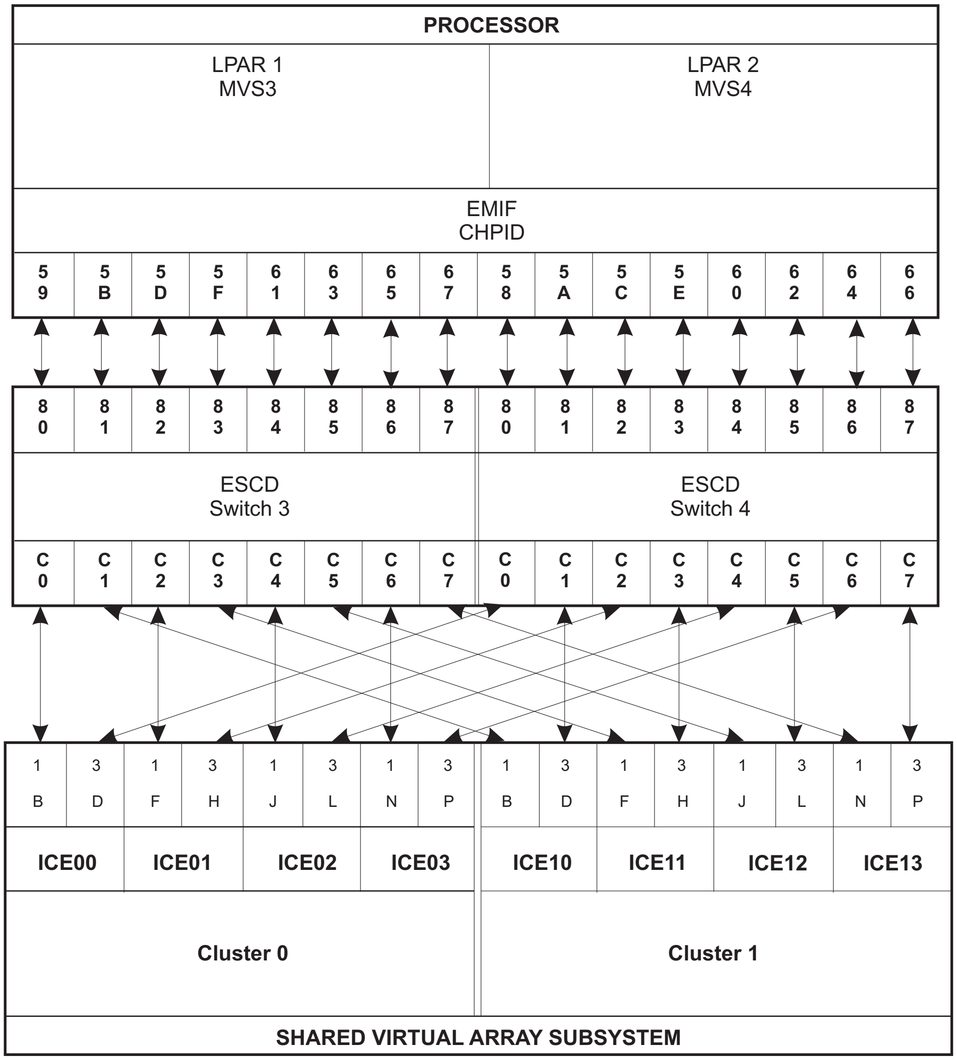

All ESCON Channels	
BASE	00
RANGE	4096
BFDID	00

C95235

Figure 15 Channel and ICE Card Configurations For Example 3

IOCP Statements

Table 27 IOCP Statements

RESOURCE PARTITION=((MVS1,1),(MVS2,2))
CHPID
PATH=(59,5B,5D,5F,61,63,65,67),TYPE=CNC,SWITCH=01,PARTITION=(MVS1,MVS2),SHARED
PATH=(58,5A,5C,5E,60,62,64,66),TYPE=CNC,SWITCH=02,PARTITION=(MVS1,MVS2),SHARED
CNTLUNIT
CUNUMBR=1000,PATH=(58,5A,5D,5F),LINK=(C0,C1,C2,C3),UNITADD=((00,256)),CUADD=0,UNIT=3990
CUNUMBR=1100,PATH=(58,5A,5D,5F),LINK=(C0,C1,C2,C3),UNITADD=((00,256)),CUADD=1,UNIT=3990
CUNUMBR=1200,PATH=(58,5A,5D,5F),LINK=(C0,C1,C2,C3),UNITADD=((00,256)),CUADD=2,UNIT=3990
CUNUMBR=1300,PATH=(58,5A,5D,5F),LINK=(C0,C1,C2,C3),UNITADD=((00,256)),CUADD=3,UNIT=3990
CUNUMBR=1400,PATH=(59,5B,5C,5E),LINK=(C0,C1,C2,C3),UNITADD=((00,256)),CUADD=4,UNIT=3990
CUNUMBR=1500,PATH=(59,5B,5C,5E),LINK=(C0,C1,C2,C3),UNITADD=((00,256)),CUADD=5,UNIT=3990
CUNUMBR=1600,PATH=(59,5B,5C,5E),LINK=(C0,C1,C2,C3),UNITADD=((00,256)),CUADD=6,UNIT=3990
CUNUMBR=1700,PATH=(59,5B,5C,5E),LINK=(C0,C1,C2,C3),UNITADD=((00,256)),CUADD=7,UNIT=3990
CUNUMBR=1800,PATH=(60,62,65,67),LINK=(C4,C5,C6,C7),UNITADD=((00,256)),CUADD=8,UNIT=3990
CUNUMBR=1900,PATH=(60,62,65,67),LINK=(C4,C5,C6,C7),UNITADD=((00,256)),CUADD=9,UNIT=3990
CUNUMBR=1A00,PATH=(60,62,65,67),LINK=(C4,C5,C6,C7),UNITADD=((00,256)),CUADD=A,UNIT=3990
CUNUMBR=1B00,PATH=(60,62,65,67),LINK=(C4,C5,C6,C7),UNITADD=((00,256)),CUADD=B,UNIT=3990
CUNUMBR=1C00,PATH=(61,63,64,65),LINK=(C4,C5,C6,C7),UNITADD=((00,256)),CUADD=C,UNIT=3990
CUNUMBR=1D00,PATH=(61,63,64,65),LINK=(C4,C5,C6,C7),UNITADD=((00,256)),CUADD=D,UNIT=3990
CUNUMBR=1E00,PATH=(61,63,64,65),LINK=(C4,C5,C6,C7),UNITADD=((00,256)),CUADD=E,UNIT=3990
CUNUMBR=1F00,PATH=(61,63,64,65),LINK=(C4,C5,C6,C7),UNITADD=((00,256)),CUADD=F,UNIT=3990
IODEVICE
ADDRESS=(1000,256),UNITADD=00,CUNUMBR=1000,STADET=Y,UNIT=3390,FEATURE=SHARED
ADDRESS=(1100,256),UNITADD=00,CUNUMBR=1100,STADET=Y,UNIT=3390,FEATURE=SHARED
ADDRESS=(1200,256),UNITADD=00,CUNUMBR=1200,STADET=Y,UNIT=3390,FEATURE=SHARED
ADDRESS=(1300,256),UNITADD=00,CUNUMBR=1300,STADET=Y,UNIT=3390,FEATURE=SHARED
ADDRESS=(1400,256),UNITADD=00,CUNUMBR=1400,STADET=Y,UNIT=3390,FEATURE=SHARED
ADDRESS=(1500,256),UNITADD=00,CUNUMBR=1500,STADET=Y,UNIT=3390,FEATURE=SHARED
ADDRESS=(1600,256),UNITADD=00,CUNUMBR=1600,STADET=Y,UNIT=3390,FEATURE=SHARED
ADDRESS=(1700,256),UNITADD=00,CUNUMBR=1700,STADET=Y,UNIT=3390,FEATURE=SHARED
ADDRESS=(1800,256),UNITADD=00,CUNUMBR=1800,STADET=Y,UNIT=3390,FEATURE=SHARED
ADDRESS=(1900,256),UNITADD=00,CUNUMBR=1900,STADET=Y,UNIT=3390,FEATURE=SHARED
ADDRESS=(1A00,256),UNITADD=00,CUNUMBR=1A00,STADET=Y,UNIT=3390,FEATURE=SHARED
ADDRESS=(1B00,256),UNITADD=00,CUNUMBR=1B00,STADET=Y,UNIT=3390,FEATURE=SHARED
ADDRESS=(1C00,256),UNITADD=00,CUNUMBR=1C00,STADET=Y,UNIT=3390,FEATURE=SHARED
ADDRESS=(1D00,256),UNITADD=00,CUNUMBR=1D00,STADET=Y,UNIT=3390,FEATURE=SHARED
ADDRESS=(1E00,256),UNITADD=00,CUNUMBR=1E00,STADET=Y,UNIT=3390,FEATURE=SHARED
ADDRESS=(1F00,256),UNITADD=00,CUNUMBR=1F00,STADET=Y,UNIT=3390,FEATURE=SHARED

HCD for Example 3 - 16 channels (4-path VCU)

The definition method in the following example is driven from the Define, Modify, or View Configuration Data panel. Explanation of HCD panel navigation is beyond the scope of this document. Definitions for channel paths, control units, and devices are shown in the following panels.

Channel path definition for CHPID 59:

```
Add Channel Path

Specify or revise the following values.

Processor ID . . . : P1  Processor One

Configuration mode : LPAR

Channel path ID. . . . . . . . 59 +

Number of CHPIDs . . . . . . . 1

Channel path type. . . . . . . CNC +

Operation mode . . . . . . . . SHR +

Description. . . . . . . . . . Shared Virtual Array
with EMIF

Specify the following values only if connected to a
switch:

Dynamic switch ID  . . . . 01 + (00 - FF)

Entry switch ID  . . . . . 01 +

Entry port . . . . . . . . 80 +
```

Note: Repeat for channels 5B 5D 5F 61 63 65 67 connected to switch id 01 (ESCON director). The associated entry ports are 81-87. The switch must be defined prior to connecting the channel path to the switch. A group of channel paths with consecutive identifiers and the same type and mode might be defined in one step by entering a value for Number of CHPIDs. The attributes for each CHPID might be modified in another panel.

Channel path definition for CHPID 58:

```
Add Channel Path

Specify or revise the following values.

Processor ID . . . : P1  Processor One

Configuration mode : LPAR

Channel path ID. . . . . . . . 58 +

Number of CHPIDs . . . . . . . 1
```

```
Channel path type. . . . . . . . CNC +

Operation mode . . . . . . . . . SHR +

Description  . . . . . . . . . . Shared Virtual Array
with EMIF

Specify the following values only if connected to a
switch:

Dynamic switch ID  . . . . . . . 02 + (00 - FF)

Entry switch ID  . . . . . . . . 02 +

Entry port . . . . . . . . . . . 80 +
```

Director (switch) definition:

```
Update Switch Connections

Specify or revise the following values.

Processor ID . . . : P1

Dyn          + --         Entry        +--

CHPID        Switch       Switch       Port

58           02           02           80
```

Note: This panel allows definition of the entry ports for all subsequent channel paths if a group of channel paths have been defined in one step. The switch and port values displayed initially are only valid for the first channel path.

LPAR access list definition:

```
Define Access List

Select one or more partitions for inclusion in the
access list.

Channel path ID. . : 58 Channel path type . . : CNC

Operation mode . . : SHR Number of CHPIDs. . : 1

/ Partition Name  Number  Usage Description

/ MVS1        1           OS          LPAR One

/ MVS2        2           OS          LPAR Two
```

Note: Each channel is shared, accessed by more than one logical partition at the same time (EMIF).

Channel path attribute listing:

```
Channel Path List      Row 1 of 10 More: >

Command                                          ===>
__________________________________________
```

Select one or more channel paths, then press Enter. To add use F11.

Processor ID . . . : P1

Configuration mode : LPAR

				Dyn +	--Entry +--	
/	CHPID	Type +	Mode +	Switch	Switch Port	Con. Description
_	58	CNC	SHR	02	02	80
_	59	CNC	SHR	01	01	80
_	5A	CNC	SHR	02	02	81
_	5B	CNC	SHR	01	01	81
_	5C	CNC	SHR	02	02	82
_	5D	CNC	SHR	01	01	82
_	5E	CNC	SHR	02	02	83
_	5F	CNC	SHR	01	01	83
_	60	CNC	SHR	02	02	84
_	61	CNC	SHR	01	01	84

******************************* Bottom of data

Note: Type over the appropriate values to modify the channel path attributes, except the channel path ID.

Channel Path List Row 1 of 10 More: <

Command ===> ____

Select one or more channel paths, then press Enter. To add, use F11.

1=MVS 1	2=MVS2	3=	4=	5=
6=	7=	8=	9=	A=
B=	C=	D=	E=	F=

---------------- Partitions -

/	CHPID	Type +	Mode +	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
_	58	CNC	SHR	a	a	_	_	_	_	_	_	_	_					
_	58	CNC	SHR	a	a	_	_	_	_	_	_	_	_					
_	59	CNC	SHR	a	a	_	_	_	_	_	_	_	_					
_	5A	CNC	SHR	a	a	_	_	_	_	_	_	_	_					

```
_ 5B    CNC    SHR    a  a  _  _  _  _  _  _  _  _
_ 5C    CNC    SHR    a  a  _  _  _  _  _  _  _  _
_ 5D    CNC    SHR    a  a  _  _  _  _  _  _  _  _
_ 5E    CNC    SHR    a  a  _  _  _  _  _  _  _  _
_ 5F    CNC    SHR    a  a  _  _  _  _  _  _  _  _
_ 60    CNC    SHR    a  a  _  _  _  _  _  _  _  _
_ 61    CNC    SHR    a  a  _  _  _  _  _  _  _  _
 ****************************** Bottom of data
```

Note: The letter a or c indicates assignment of a partition to the access (a) or candidate (c) list for each channel path.

Logical control unit definition:

```
Add Control Unit

Specify or revise the following values.

Control unit number. . : 1000 + Type . . . : 3990 +

Processor ID . . . . . : P1  Processor One

Channel path IDs . . . . 58 5A 5D 5F __ __ +

Link address . . . . . . C0 C1 C2 C3 __ __ +

Unit address . . . . . . 00 __ __ __ __ __ __ __ +

Number of units. . . . . 256   __ __ __ __ __ __

Logical address. . . . . 0 + (same as CUADD)

Protocol . . . . . . . . _ + (D, S, or S4)

I/O concurrency level. . _ + (1, 2, or 3)
```

Note: Define CUADD = 0, 1, 2, and 3 with the same CHPIDs. The CHPIDs alternate between the physical cluster 0 and cluster 1 connections to the SVA. Logical control unit CUADD=0 is defined in this panel. Assign unique control unit numbers. The protocol and I/O concurrency level parameters are ignored.

```
Add Control Unit

Specify or revise the following values.

Control unit number. . : 1400  +  Type . . . : 3990
+

Processor ID . . . . . : P1  Processor One

Channel path IDs . . . . 59 5B 5C 5E __ __ +

Link address . . . . . . C0 C1 C2 C3 __ __ +
```

```
Unit address . . . . . . 00 __ __ __ __ __ __ __ +

Number of units. . . . . 256   __ __ __ __ __ __

Logical address. . . . . 4 + (same as CUADD)

Protocol . . . . . . . . _ + (D, S, or S4)

I/O concurrency level. . _ + (1, 2, or 3)
```

Note: Define CUADD = 4, 5, 6, and 7 with the same CHPIDs. CUADD=4 is defined in this panel. Assign unique control unit numbers.

```
Add Control Unit

Specify or revise the following values.

Control unit number. . : 1800  +  Type . . . : 3990
+

Processor ID . . . . . : P1  Processor One

Channel path IDs . . . . 60 62 65 67 __ __ +

Link address . . . . . . C4 C5 C6 C7 __ __ +

Unit address . . . . . . 00 __ __ __ __ __ __ __ +

Number of units. . . . . 256   __ __ __ __ __ __

Logical address. . . . . 8 + (same as CUADD)

Protocol . . . . . . . . _ + (D, S, or S4)

I/O concurrency level. . _ + (1, 2, or 3)
```

Note: Define CUADD = 8, 9, A, and B with the same CHPIDs. Logical control unit CUADD=8 is defined in this panel.

```
Add Control Unit

Specify or revise the following values.

Control unit number. . : 1C00   +    Type . . . :
3990  +

Processor ID . . . . . : P1  Processor One

Channel path IDs . . . . 61 63 64 66 +

Link address . . . . . . C4 C5 C6 C7 +

Unit address . . . . . . 00 __ __ __ __ __ __ __ +

Number of units. . . . . 256   __ __ __ __ __ __

Logical address. . . . . C + (same as CUADD)

Protocol . . . . . . . . _ + (D, S, or S4)

I/O concurrency level. . _ + (1, 2, or 3)
```

Note: Define CUADD = C, D, E, F with the same CHPIDs. Logical control unit CUADD=C is defined in this panel.

Device definition for host addresses 1000-1FFF:

```
Add Device

Specify or revise the following values.

Device number . . . . . . . . . 1000 (0000 - FFFF)

Number of devices . . . . . . . 256__

Device type . . . . . . . . . . 3390___________ +

Serial number . . . . . . . . . _______________

Description . . . . . . . . . . _______________

Connected to CUs. . . 1000 ___ ___ ___ ___ ___ ___ +

Device / Processor Definition

Select processor(s) to change device/processor def-
initions, then press Enter.

Device number . . : 1000 Number of devices . : 256

Device type . . . : 3390

/ Processor ID  UA +  Time-Out  STADET  Preferred
Explicit Device

                                            CHPID +
Candidate List

/ P1            00    No        Yes     __   No
```

Note: Repeat for device numbers 1x00, where x = 1-F. Specify assigned CU number. UA is the unit address keyword 00.

Example 4 – SVA With Direct Host Connectivity

This subsystem configuration includes connectivity to 1 physical processor with 3 LPARs. The channel configuration includes 8 ESCON channels (EMIF) from the processor.

A maximum number of 8 channel paths to a device can be configured from each operating system. This is a host processor and HCD/IOCP limit.

The customer requirement is to define all 4096 devices.

- Device Host addresses are 2000 - 2FFF.
- Device types are defined as 3390-3.
- Functional device (FDIDs) are defined as 000 - FFF.

The following example illustrates the configuration for Processor 1.

- Eight ESCON channels (EMIF) from processor 1 are connected to this subsystem.
  - Channels 20, 22, 24, and 26 are connected to Cluster 0.
  - Channels 21, 23, 25, and 27 are connected to Cluster 1.
  - Eight ICE3 cards, quad link (32 total links)
- The maximum availability of host paths is 16 per ICE card link, 512 total. This configuration allows for 3 logical host paths per ICE card link, 24 total (3 LPARs x 8 links). Each link is not required to be connected to a channel.
- Define unique CHPIDs for each set of 1024 devices (host limit). If all devices are defined, assign a different collection of CHPIDs to each set of 4 LCUs:

  CUADD = 0 1 2 3

  CUADD = 4 5 6 7

  CUADD = 8 9 A B

  CUADD = C D E F
- For explanation purposes, the example displays the configuration for CUADD = 0, 1, 2, and 3. If the customer requirement is to define 1024 devices, then this configuration could be used.
- To expand this example up to 4096 devices, abide by the ESCON limitation of 1024 devices per channel and define up to 16 LCUs.

***Caution: Performance Impact* - Do not exceed the total logical paths, which might be established from all processors.**

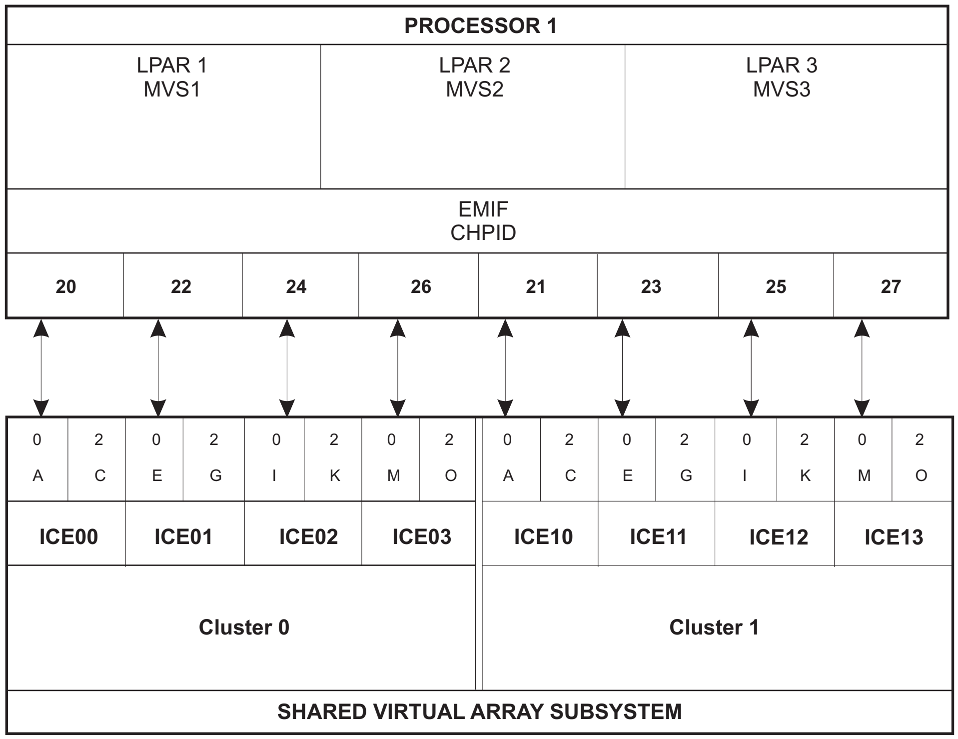

All ESCON Channels	
BASE	00
RANGE	4096
BFDID	00

C95234

Figure 16 Channel and ICE Card Configurations For Example 4

IOCP Statements

Table 28 IOCP Statements

RESOURCE PARTITION=((MVS1,1),(MVS2,2),(MVS3,3))
CHPID
PATH=(20, 21, 22, 23, 24, 25, 26, 27),TYPE=CNC,PARTITION=(MVS1, MVS2, MVS3),SHARED
CNTLUNIT
CUNUMBR=2000,PATH=(20,21,22,23,24,25,26,27),UNITADD=((00,256)),CUADD=0,UNIT=3990
CUNUMBR=2100,PATH=(20,21,22,23,24,25,26,27),UNITADD=((00,256)),CUADD=1,UNIT=3990
CUNUMBR=2200,PATH=(20,21,22,23,24,25,26,27),UNITADD=((00,256)),CUADD=2,UNIT=3990
CUNUMBR=2300,PATH=(20,21,22,23,24,25,26,27),UNITADD=((00,256)),CUADD=3,UNIT=3990
IODEVICE
ADDRESS=(2000,256),UNITADD=00,CUNUMBR=2000,STADET=Y,UNIT=3390,FEATURE=SHARED
ADDRESS=(2100,256),UNITADD=00,CUNUMBR=2100,STADET=Y,UNIT=3390,FEATURE=SHARED
ADDRESS=(2200,256),UNITADD=00,CUNUMBR=2200,STADET=Y,UNIT=3390,FEATURE=SHARED
ADDRESS=(2300,256),UNITADD=00,CUNUMBR=2300,STADET=Y,UNIT=3390,FEATURE=SHARED

HCD for Example 4

The definition method in the following example is driven from the Define, Modify, or View Configuration Data panel. Explanation of HCD panel navigation is beyond the scope of this document. Definitions for channel paths, control units, and devices are shown in the following panels.

Channel path definition for CHPID 20-27:

```
Add Channel Path

Specify or revise the following values.

Processor ID . . . : P1  Processor One

Configuration mode : LPAR

Channel path ID. . . . . . . . 20 +

Number of CHPIDs . . . . . . . 8

Channel path type. . . . . . . CNC +

Operation mode . . . . . . . . SHR +

Description. . . . . . . . . Shared Virtual Array
with EMIF
```

```
Specify the following values only if connected to a
switch:

Dynamic switch ID  . . . . __ + (00 - FF)

Entry switch ID  . . . . . __ +

Entry port . . . . . . . . __ +
```

Note: A group of channel paths with consecutive identifiers and the same type and mode might be defined in one step by entering a value for Number of CHPIDs. The attributes for each CHPID might be modified in another panel. If a director is included in the configuration, the switch must be defined prior to connecting the channel path to the switch.

LPAR access list definition:

```
Define Access List

Select one or more partitions for inclusion in the
access list.

Channel path ID. . : 20  Channel path type . . : CNC

Operation mode . . : SHR Number of CHPIDs. . . : 8

/ Partition Name Number   Usage     Description

/ MVS1         1          OS        LPAR One

/ MVS2         2          OS        LPAR Two

/ MVS3         3          OS        LPAR Three
```

Note: Each channel is shared, accessed by more than one logical partition at the same time (EMIF).

Channel path attribute listing:

```
                                  Channel Path List
Row 1 of 10 More: >

                                Command        ===>
__________________________________________

     Select one or more channel paths, then press
Enter. To add use F11.

    Processor ID . . . : P1

    Configuration mode : LPAR
```

```
                              Dyn +  --Entry +--
       / CHPID Type + Mode + Switch Switch Port Con.
   Description
      _ 20     CNC     SHR     __      __      __
      _ 21     CNC     SHR     __      __      __
      _ 22     CNC     SHR     __      __      __
      _ 23     CNC     SHR     __      __      __
      _ 24     CNC     SHR     __      __      __
      _ 25     CNC     SHR     __      __      __
      _ 26     CNC     SHR     __      __      __
      _ 27     CNC     SHR     __      __      __
       ****************************** Bottom of data
```

Note: Type over the appropriate values to modify the channel path attributes, except the channel path ID.

Channel Path List Row 1 of 10 More: <

```
   Command ===> ___________________________________________

   Select one or more channel paths, then press Enter. To add, use
F11.

   1=MVS 1     2=MVS2      3= MVS3     4=          5=
   6=          7=          8=          9=          A=
   B=          C=          D=          E=          F=

                     ---------------- Partitions -
   / CHPID Type + Mode + 1  2  3  4  5  6  7  8  9  A  B  C  D  E  F
    _ 20     CNC     SHR     a   a   a   _   _   _   _   _   _   _
    _ 21     CNC     SHR     a   a   a   _   _   _   _   _   _   _
    _ 22     CNC     SHR     a   a   a   _   _   _   _   _   _   _
    _ 23     CNC     SHR     a   a   a   _   _   _   _   _   _   _
    _ 24     CNC     SHR     a   a   a   _   _   _   _   _   _   _
    _ 25     CNC     SHR     a   a   a   _   _   _   _   _   _   _
    _ 26     CNC     SHR     a   a   a   _   _   _   _   _   _   _
    _ 27     CNC     SHR     a   a   a   _   _   _   _   _   _   _
     ****************************** Bottom of data
```

Note: The letter a or c indicates assignment of a partition to the access (a) or candidate (c) list for each channel path.

Logical control unit definition

```
Add Control Unit

Specify or revise the following values.

Control unit number. . : 2000  +  Type ...: 3990 +

Processor ID . . . . . : P1  Processor One

Channel path IDs . . . . 20 21 22 23 24 25 26 27 +

Link address . . . . . . __ __ __ __ __ __ __ __ +

Unit address . . . . . . 00 __ __ __ __ __ __ __ +

Number of units. . . . . 256   __ __ __ __ __ __

Logical address. . . . . 0 + (same as CUADD)

Protocol . . . . . . . . _ + (D, S, or S4)

I/O concurrency level. . _ + (1, 2, or 3)
```

Note: Define CUADD = 0, 1, 2, and 3 with the same CHPIDs. The CHPIDs alternate between the physical cluster 0 and cluster 1 of the SVA. Logical control unit CUADD = 0 is defined in this panel. Assign unique control unit numbers. The protocol and I/O concurrency level parameters are ignored by ESCON.

Device definition for host addresses

```
Add Device

Specify or revise the following values.

Device number . . . . . . . . . 2000 (0000 - FFFF)

Number of devices . . . . . . . 256__

Device type . . . . . . . . . . 3390___________ +

Serial number . . . . . . . . . _______________

Description . . . . . . . . . . _______________

Connected to CUs. . . 2000 ___ ___ ___ ___ ___ ___ +
```

Note: Device numbers 2000-20FF connected to control unit numbers 2000 are defined in this panel.

```
Device / Processor Definition

Select processor(s) to change device/processor def-
initions, then

press Enter.

Device number . . : 2000 Number of devices . : 256
```

```
Device type . . . : 3390

/ Processor ID  UA +  Time-Out  STADET  Preferred
Explicit Device

                                          CHPID +
Candidate List

/ P1            00    No        Yes     __   No
```

Note: Repeat for device numbers 2x00, where x = 1-F. Specify assigned CU numbers. UA is the unit address keyword, 00.

Device Characteristics

D

3390 Devices

Table 29 3390 Device Geometry and Characteristics

Description	Functional 3390-1 device	Functional 3390-2 device	Functional 3390-3 device	Functional 3390-9 device
Usable track capacity (bytes)	57326	57326	57326	57326
Maximum record one (R1) length (bytes/track)	56664	56664	56664	56664
Data cylinders per actuator[a]	1113	2226	3339	10017
Tracks per cylinder	15	15	15	15
Capacity per actuator (megabytes)	946	1891	2838	8514
Sectors per track	224	224	224	224
Range of valid sector numbers	0-223	0-223	0-223	0-223
Rotational position sensing	Yes	Yes	Yes	Yes
Defective/alternate track	Yes	Yes	Yes	Yes
Diagnostic track	Yes	Yes	Yes	Yes
Device support track	Yes	Yes	Yes	Yes
Primary Track[b] Address Range Low High	 00000000 0458000E	 00000000 08B1000E	 00000000 0D0A000E	 00000000 2720000E
Alternate Track[b.] Address Range Low High	 04590000 0459000E	 08B20000 08B2000E	 0D0B0000 0D0B000E	 27210000 2723000E

Table 29 3390 Device Geometry and Characteristics (Continued)

Description	Functional 3390-1 device	Functional 3390-2 device	Functional 3390-3 device	Functional 3390-9 device
Diagnostic Track[b.] Address Range Low High	 045B0000 045B000E	 08B40000 08B4000E	 0D0D0000 0D0D000E	 27270000 2732000E
Device Support Tracks[b.] Range Low High	 04810000 0482000E	 08D90000 08DA000E	 0D190000 0D1A000E	 274B0000 2750000E
Status Track	0482000E	08DA000E	08DA000E	2750000E

a. For FlexVolume cylinder ranges, see Table 10, “FlexVolume Cylinder Options per Device Type,” on page 50.
b. Indicates the device’s specific track address ranges in hexadecimal, as the four-byte cylinder and head address (CCHH).

3380 Devices

Table 30 3380 Device Geometry and Characteristics

Description	3380-6EA K [2]	3380J	3380K
Usable track capacity (bytes), maximum record zero (R0)	47988	47988	47988
Maximum record one (R1) length (bytes per track)	47476	47476	47476
Data cylinders per actuator	1770	885	2655
Tracks per cylinder[a]	15	15	15
Capacity per actuator (megabytes)	1260	630	1890
Sectors per track	222	222	222
Range of valid sector numbers	0-221	0-221	0-221
Rotational position sensing	Yes	Yes	Yes
Defective/alternate track	Yes	Yes	Yes
Diagnostic track	Yes	Yes	Yes
Device support track	Yes	Yes	Yes

Table 30 3380 Device Geometry and Characteristics (Continued)

Description	3380-6EA K [2]	3380J	3380K
Primary Track[b] Address Range Low High	 00000000 06E9000E	 00000000 0374000E	 00000000 0A5E000E
Alternate Track[b.] Address Range Low High	 06EA0000 06EA000E	 03750000 0375000E	 0A5F0000 0A5F000E
Diagnostic Track[b.] Address Range Low High	 06EB0000 06EB000E	 03760000 0376000E	 0A620000 0A62000E
Device Support Tracks[b.] Range Low High	 06F40000 06F5000E	 FFFD0000 FFFD000E	 0A6B0000 0A6D000E
Status Track	06F5000E	FFFD000E	0A6D000E
Note: Indicates the device's specific track address ranges in hexadecimal, as CCHH (4 bytes).			

a. For FlexVolume cylinder ranges, see Table 10, "FlexVolume Cylinder Options per Device Type," on page 50.

b. SVA provides a virtual device image for a 3380E device and refers to it as a 6EA K device (where 6EA is the number of primary cylinders for a 3380E). Note that the storage control and device facilities, Device Model, MDR record ID, and OBR record ID of this device type are reported as those of a 3380K device type, whereas the remaining characteristics are reported as those of a 3380E device type.

Combination Examples

Table 31 3390 Device Combinations

MOD3	MOD9	DEVICES
4096	0	4096
4000	32	4032
3904	64	3968
3808	96	3904
3712	128	3840
3616	160	3776
3520	192	3712
3424	224	3648
3328	256	3584
3232	288	3520
3136	320	3456

Table 31 3390 Device Combinations (Continued)

MOD3	MOD9	DEVICES
3040	352	3392
2944	384	3328
2848	416	3264
2752	448	3200
2656	480	3136
2560	512	3072
2464	544	3008
2368	576	2944
2272	608	2880
2176	640	2816
2080	672	2752
1984	704	2688
1888	736	2624
1792	768	2560
1696	800	2496
1600	832	2432
1504	864	2368
1408	896	2304
1312	928	2240
1216	960	2176
1120	992	2112
1024	1024	2048
928	1056	1984
832	1088	1920
736	1120	1856
640	1152	1792
544	1184	1728
448	1216	1664
352	1248	1600
256	1280	1536
160	1312	1472
64	1344	1408
0	1365	1365

Note: Device combinations other than those in this device matrix are possible.

Access Authorization Address Range

E

Table 32 3380 Access Authorization Address Range

Description	3380-6E A K	3380J	3380K
Normal Access			
Low	00000000	00000000	00000000
High	06EA000E	0375000E	0A5F000E
Device Support			
Low	00000000	00000000	00000000
High	06F5000E	0376000E	0A6D000E
Diagnostic			
Low	06EB0000	03760000	0A620000
High	06F5000E	0376000E	0A62000E

Note:

1. Shown are the device's specific track address ranges in hexadecimal, as CCHH (4 bytes).
2. Invalid addresses within a range are not valid as extent or seek addresses, but might be included within the extent range.
3. Tracks that can be accessed under Normal Access Authorization might also be accessed under Diagnostic Access Authorization if the Seek precedes a Set File Mask command.
4. SVA provides a virtual device image for a 3380E device and refers to it as a 6EA K device (where 6EA is the number of primary cylinders for a 3380E). Note that the storage control and device facilities, Device Model, MDR record ID, and OBR record ID of this device type are reported as those of a 3380K device type, whereas the remaining characteristics are reported as those of a 3380E device type.

Table 33 3390 Access Authorization Address Range

Description	3390 Model 1	3390 Model 2	3390 Model 3	3390 Model 9
Normal Access				
Low	00000000	00000000	00000000	00000000
High	0458000E	08B1000E	0D0A000E	2720000E
Device Support				
Low	00000000	00000000	00000000	00000000
High	0482000E	08DA000E	0D1A000E	2750000E
Diagnostic				
Low	045C0000	08B40000	0D0D0000	27200000
High	045C000E	08B4000E	0D0D000E	2732000E

ICE/ICF Card Combinations

F

Table 34 ICE2, ICE3 and ICF Normal Locations

Feature Numbers	# of ICE3 Cards	# of ICE2 Cards	# of ICF Cards	Cluster 0				Cluster 1			
				Slot 00	Slot 01	Slot 02	Slot 03	Slot 10	Slot 11	Slot 12	Slot 13
ICE3 Cards Only											
	2	0	0	ICE3				ICE3			
	4	0	0	ICE3	ICE3			ICE3	ICE3		
	6	0	0	ICE3	ICE3	ICE3		ICE3	ICE3	ICE3	
	8	0	0	ICE3	ICE3	ICE3	ICE3	ICE3	ICE3	ICE3	ICE3
ICE2 and ICE3 Cards											
	2	2	0	ICE3	ICE2			ICE3	ICE2		
	4	2	0	ICE3	ICE2	ICE3		ICE3	ICE2	ICE3	
	6	2	0	ICE3	ICE2	ICE3	ICE3	ICE3	ICE2	ICE3	ICE3
	4	4	0	ICE3	ICE2	ICE3	ICE2	ICE3	ICE2	ICE3	ICE2
ICF Cards Only											
	0	0	2	ICF				ICF			
	0	0	4	ICF	ICF			ICF	ICF		
	0	0	6	ICF	ICF	ICF		ICF	ICF	ICF	
	0	0	8	ICF	ICF	ICF	ICF	ICF	ICF	ICF	ICF

Table 34 ICE2, ICE3 and ICF Normal Locations (Continued)

Feature Numbers	# of ICE3 Cards	# of ICE2 Cards	# of ICF Cards	Cluster 0				Cluster 1			
				Slot 00	Slot 01	Slot 02	Slot 03	Slot 10	Slot 11	Slot 12	Slot 13
ICE3 and ICF Cards											
	2	0	2	ICE3	ICF			ICE3	ICF		
	2	0	3	ICE3	ICF	ICF		ICE3	ICF		
	3	0	2	ICE3	ICF	ICE3		ICE3	ICF		
	3	0	3	ICE3	ICF	ICE3	ICF	ICE3	ICF		
	2	0	4	ICE3	ICF	ICF		ICE3	ICF	ICF	
	4	0	2	ICE3	ICF	ICE3		ICE3	ICF	ICE3	
	3	0	4	ICE3	ICF	ICE3	ICF	ICE3	ICF	ICF	
	4	0	3	ICE3	ICF	ICE3	ICF	ICE3	ICF	ICE3	
	5	0	2	ICE3	ICE3	ICE3	ICF	ICE3	ICE3	ICF	
	2	0	5	ICE3	ICF	ICF	ICF	ICE3	ICF	ICF	
	4	0	4	ICE3	ICF	ICE3	ICF	ICE3	ICF	ICE3	ICF
	5	0	3	ICE3	ICE3	ICE3	ICF	ICE3	ICF	ICE3	ICF
	3	0	5	ICE3	ICF	ICE3	ICF	ICE3	ICF	ICF	ICF
	6	0	2	ICE3	ICE3	ICE3	ICF	ICE3	ICE3	ICE3	ICF
	2	0	6	ICE3	ICF	ICF	ICF	ICE3	ICF	ICF	ICF
ICE2 and ICF Cards											
	0	2	2	ICE2	ICF			ICE2	ICF		
	0	2	3	ICE2	ICF	ICF		ICE2	ICF		
	0	2	4	ICE2	ICF	ICF		ICE2	ICF	ICF	
	0	2	5	ICE2	ICF	ICF	ICF	ICE2	ICF	ICF	
	0	2	6	ICE2	ICF	ICF	ICF	ICE2	ICF	ICF	ICF
	0	4	4	ICE2	ICF	ICE2	ICF	ICE2	ICF	ICE2	ICF

Table 34 ICE2, ICE3 and ICF Normal Locations (Continued)

Feature Numbers	# of ICE3 Cards	# of ICE2 Cards	# of ICF Cards	Cluster 0				Cluster 1			
				Slot 00	Slot 01	Slot 02	Slot 03	Slot 10	Slot 11	Slot 12	Slot 13
ICE3, ICE2, and ICF Cards											
	2	2	2	ICE3	ICF	ICE2		ICE3	ICF	ICE2	
	2	2	3	ICE3	ICF	ICE2	ICF	ICE3	ICF	ICE2	
	3	2	2	ICE3	ICE2	ICE3	ICF	ICE3	ICE2	ICF	
	2	2	3	ICE3	ICF	ICE2	ICF	ICE3	ICF	ICE2	
	3	2	2	ICE3	ICE2	ICE3	ICF	ICE3	ICF	ICE2	
	2	2	4	ICE3	ICF	ICE2	ICF	ICE3	ICF	ICE2	ICF
	3	2	3	ICE3	ICE2	ICE3	ICF	ICE3	ICF	ICE2	ICF
	4	2	2	ICE3	ICE2	ICE3	ICF	ICE3	ICE2	ICE3	ICF
	2	4	2	ICE3	ICE2	ICE2	ICF	ICE3	ICE2	ICE2	ICF

Note: ICE2 cards are required for PPRC connections. ICE3 cards are required for standard ICE2CON connections.

System Block Diagram

G

A system block diagram is shown on the following page. Not all cards are shown. Only cards relating to data movement operations are shown. This system block diagram also indicates there is a 1 GB/sec. back end architecture.

Later versions of the SVA have a faster 2 GB/sec. back end and it is done with a four loop architecture. This architecture is shown in Figure 18 on page 181. The numbers show in that figure in the upper part of the IFF card boxes represent the "zones." With that in mind, the drives are connected in the 2 GB/sex. in the following fashion:

Loop 0 Zone 0 – even disk drives

Loop 0 Zone 1 – odd disk drives

Loop 1 Zone 0 – even disk drives

Loop 1 Zone 1 – odd disk drives

Note: Figure 18 on page 181 implies 2 cables between each IFF and IFES, and 2 cables from IFES to IFES, but in reality it is just one cable between each set of cards. The figure is more of a logical representation than a physical representation.

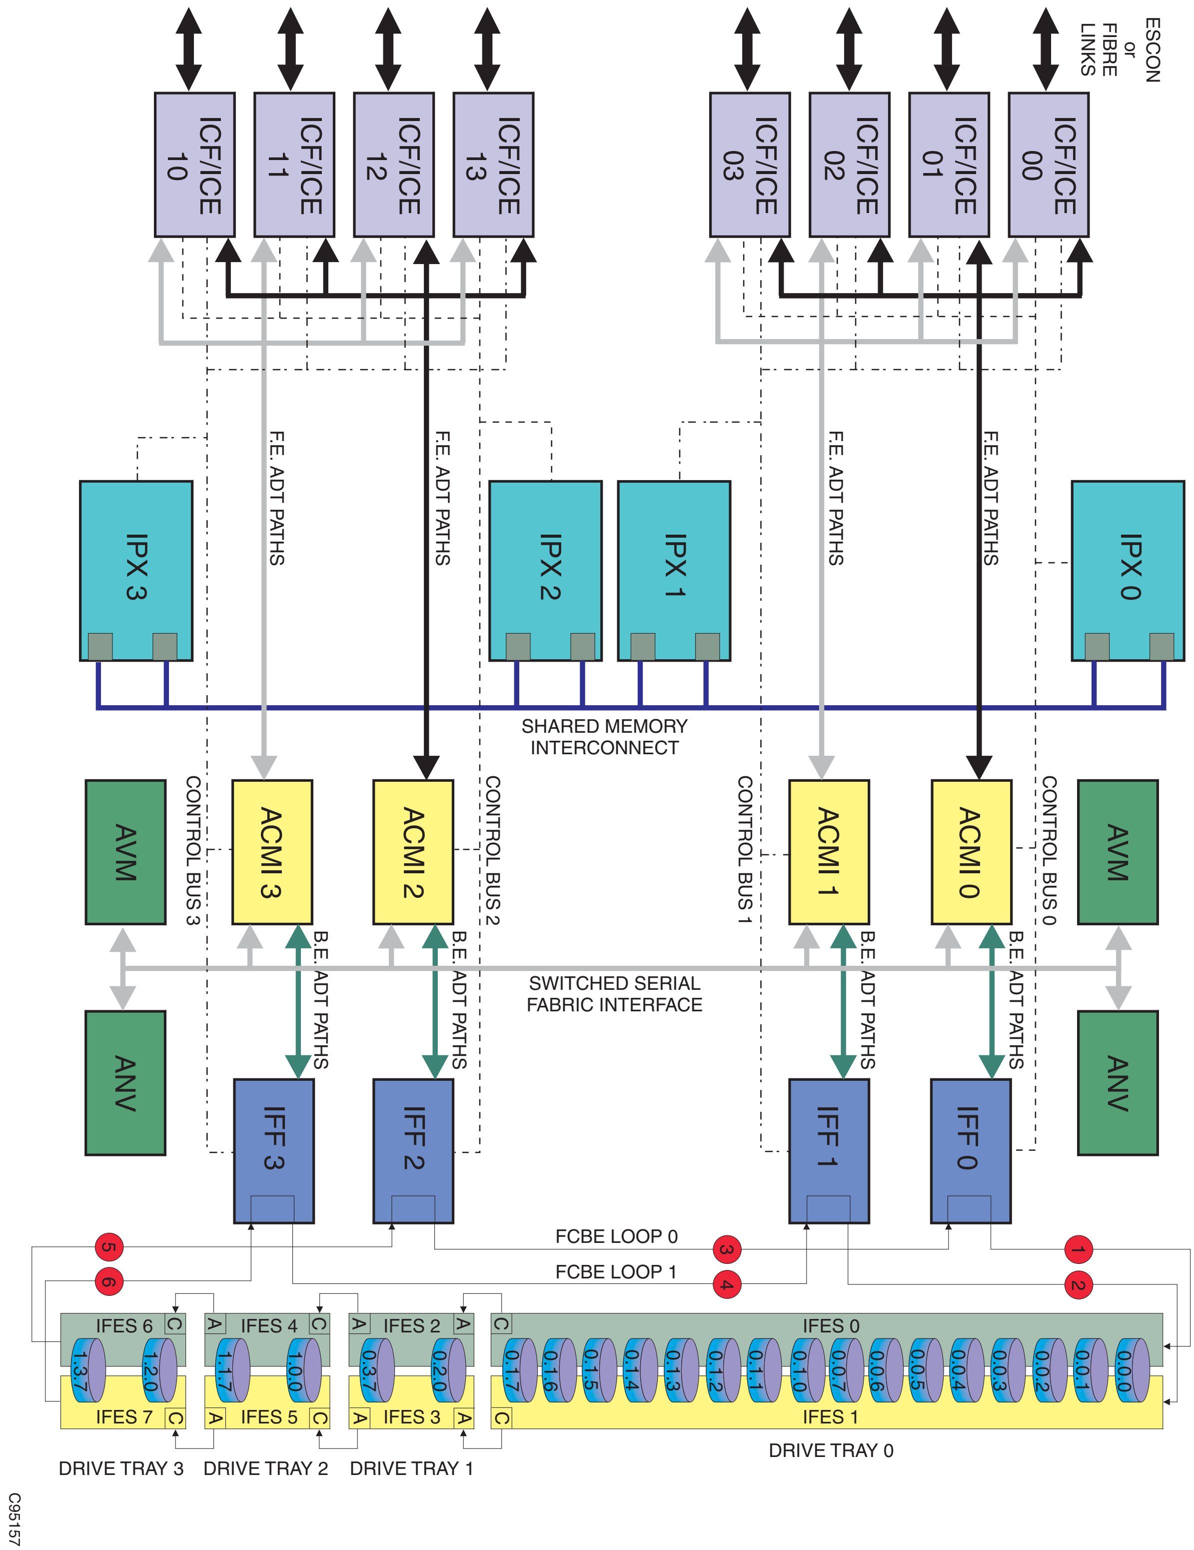

Figure 17 System Block Diagram

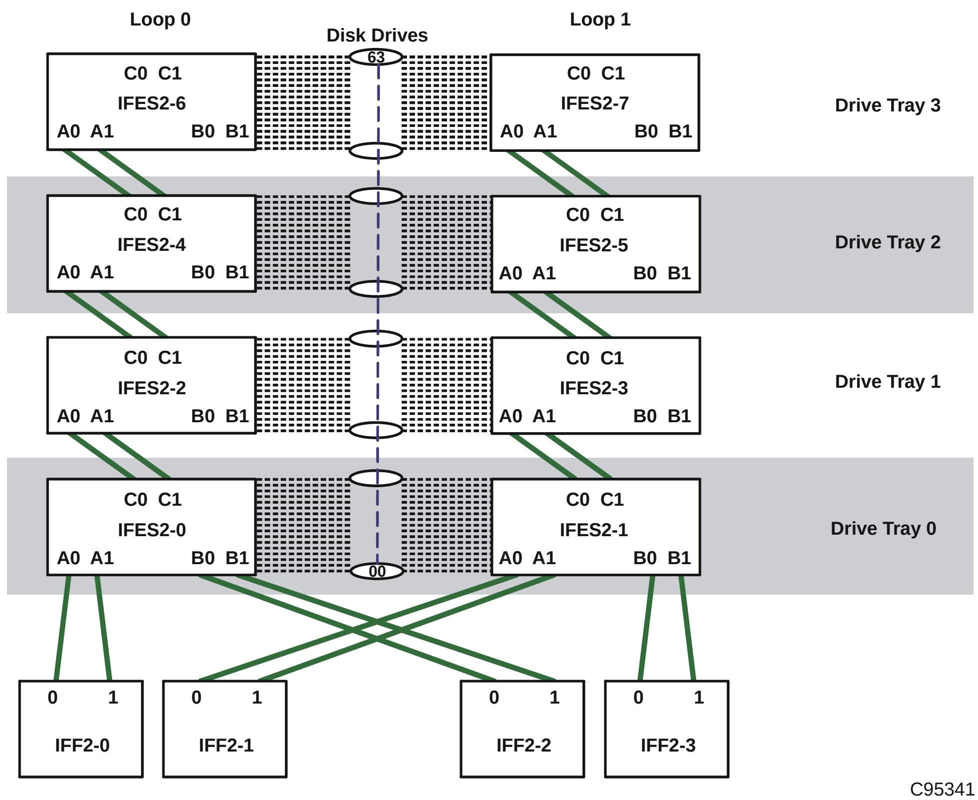

Figure 18 2 GB/sec. SVA Back End

Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 USA **Phone** 1-650-960-1300 or 1-800-555-9SUN **Web** sun.com

Argentina: 5411-4317-5636 • **Australia:** 1-800-550-786 • **Austria:** 43-1-601-26-0 • **Balkans:** 301-6188-111 • **Belgium:** 32 2-704 89 83 • **Brazil:** 55-11-51872100 • **Brunei:** 65-216-8333 • **Canada:** 1-800-422-8020 **(General);** 416-964-2001 **(Learning Management System Sales,Toronto)** • **Chile:** 562-372-4500 • **Colombia:** 571-629-2323 **Czech Republic:** 420 2 33009311 • **Denmark:** 45 4556 5040 • **Egypt:** 00 202 570 9442 • **Finland:** 358-9-525-561 • **France:** 33-1-41-33-17-17 • **Germany:** 49-89-460-08-2788 • **Greece:** 30-01-6188101 • **Hong Kong:** 852-2877-7077 • **Hungary:** 361-202-4415 • **India:** 91-80-229-8989 • **Indonesia:** 65-216-8333 • **Ireland:** 353-1-668-4377 **Israel:** 972-9-9710500 • **Italy:** 39-02-9259511 • **Japan:** 81-3-5779-1820 • **Korea:** 82-2-3453-6602 • **Malaysia:** 603-2116-1887 • **Middle East:** 00 9714 3366333 • **Mexico:** 525-261-0344 • **Netherlands:** 31-33-4515200 • **New Zealand:** 0800-786-338 • **North West Africa:** 00 9714 3366333 • **Norway: From Norway:** 47-22023950**, To Norway:** 47-23369650 • **Pakistan:** 00-9714-3366333 • **People's Republic of China:** 8610-6803-5588 • **Philippines:** 632-885-7867 • **Poland:** 48-22-8747848 • **Portugal:** 351-21-413-4000 • **Russia:** 7-095-935-8411 • **Saudi Arabia:** 00 9714 3366333 • **Singapore:** 65-216-8300 • **South Africa:** 27-11-256-6300 • **Spain:** 34-902-210-412 • **Sri Lanka:** 65-2168333 • **Sweden:** 46-8-631 22 00 • **Switzerland:** 41-1-908-90-50 **(German)** 41-22-999-0444 **(French)** • **Taiwan:** 886-2-25185735 • **Thailand:** 662-344-6855 • **Turkey:** 90 212 335 22 00 • **United Kingdom:** 44-1276-416-520 • **United States:** 1-800-422-8020 • **Venezuela:** 582-905-3800 • **Vietnam:** 65-216-8333 • **Worldwide Headquarters:** 1-650-960-1300