



MDI Site Planning Guide

(Version 2.5)

Legal Notices

© 2016 Luminex Software, Inc. All rights reserved. No part of this work may be reproduced or disclosed to third parties in any form or by any means, graphically, mechanically or electronically, including but not limited to photocopying, recording or taping without the prior written permission of Luminex Software, Inc.

Luminex Mainframe Data Integration (“MDI”) products are sold and licensed by Luminex Software, Inc. (“Luminex”) pursuant to the terms and conditions of Luminex’s standard Purchase and License Agreement or other applicable agreement (“Luminex Agreement”). Consult your Luminex Agreement for the specific terms and conditions governing the sale and license of your MDI. All software is covered by copyright, trade secret and/or patent protection and is owned by Luminex. No part may be copied, modified or transferred without the prior written permission of Luminex. You may not reverse engineer the hardware or reverse compile the software. Hardware and software described herein is confidential, proprietary and trade secret property of Luminex and may not be disclosed to third parties or used except as permitted in your Luminex Agreement.

Luminex and MDI are trademarks of Luminex Software, Inc. All other trademarks and trade names are the property of others.

DISCLAIMER: While every effort has been made to ensure the technical accuracy of this manual, Luminex makes no warranties or representations, either express or implied, with respect to the contents herein. Further, should the procedures set forth in this manual not be followed or be misapplied, Luminex disclaims any warranty for damage to the product or any other liability.

MDI Site Planning Guide

Software version 1.1

Document version 2.5

Luminex Software, Inc.

871 Marlborough Avenue

Riverside, CA 92507

Phone: +1 (951) 781-4100

Fax: +1 (951) 781-4105

info@luminex.com

www.luminex.com

October 2016

Table of Contents

1. Overview	1
2. Definitions	2
3. Site Preparation/Requirements	4
3.1 Communication-Network Connections	4
3.2 Equipment Rack Requirements.....	4
3.3 Storage-Network.....	4
3.3.1 Storage-Network: Switch.....	5
3.3.2 Storage-Network: Transceivers	5
3.3.3 Storage-Network: Cables	7
3.4 Mainframe Channels	7
3.4.1 ESCON Directors and FICON Switches	8
3.4.2 ESCON Cables.....	8
3.4.3 FICON Cables.....	9
3.5 Remote Access.....	9
3.5.1 VPN	9
3.5.2 WebEx (CISCO).....	9
4. MDI: Hardware Specifications	10
4.1 Control Unit Specifications.....	10
4.2 Power Requirements	10
5. Miscellaneous Components	11
6. Software Installation Prerequisites.....	12
7. Mainframe Configuration	13
7.1 Tape Device Characteristics	13
7.2 IOGEN Addresses.....	15
7.3 IOGEN Device Numbers	15
7.4 Logical Path Diagram	17
7.4.1 Simple Logical Path.....	17
7.4.2 Logical Paths through Switch	18
7.4.3 Logical Paths from Two Mainframes through Two Switches	21
7.4.4 Logical Paths through Cascaded Switches	23
7.5 MDI Configuration Specifications.....	26
7.6 IOCP Parameters.....	27
7.7 HCD IOGEN Settings.....	28
7.7.1 Channel Settings	30
7.7.2 Control Unit Settings	30
7.7.3 Device Settings	31
7.7.4 Device Settings/MVS Configurations.....	31
7.8 Missing Interrupt Handler (z/OS).....	32

7.9	Device Allocation Steering (z/OS)	33
7.10	Mainframe Operating System	34
7.11	Scratch Tapes	34
8.	Luminex Mainframe Software Installation.....	36
8.1	ADMIN+	36
8.2	LumAlert	36
Appendix A. Customer Questionnaire		
37		
A.1.	Equipment Infrastructure	37
A.2.	Networking Infrastructure	37
A.3.	Regarding Storage	37
A.4.	Regarding Device Allocation Steering	37
A.5.	Tape Management System	38
A.6.	Mainframe(s) Configuration	38
A.7.	Performance	38
A.8.	Applications	38
A.9.	Facilities	38
A.10.	Data Migration	38
A.11.	Installation Support	39
A.12.	Check-Off List	39
Appendix B.....Final Installation Check Off		
43		
B.1.	IP Addresses.....	43
B.2.	Remote Access.....	43
B.2.1.	Dial-up Method	43
B.2.2.	VPN Method	43
B.3.	Mainframe Configuration	43
B.4.	MDI Implementation	44
B.5.	Final Infrastructure.....	44

1. Overview

Luminex's Mainframe Data Integration (MDI) solution facilitates the exchange of data between the mainframe and open (distributed) systems platforms. The data to/from those platforms can be transformed off-mainframe by Luminex for platform compatibility. The Luminex mainframe devices are typically defined as 3590 tape drives on FICON. Through the use of a tape block size near 256K, the data transfer is virtually "MIPS free" compared to other data transfer methods, and all off-mainframe processing is completely MIPS free.

By leveraging the mainframe's I/O channels, MDI provides a secure data transport mechanism for transferring and transforming tape data between mainframes and open systems more efficiently.

2. Definitions

Below are commonly used terms in this document.

Driver	Refers to the Luminex FICON or ESCON driver that interfaces the Application to the respective interface board
ESCON	Enterprise System Connection (a type of IBM mainframe channel)
FC	Usually means Fibre Channel; however, in HCD and IOCP, "FC" means FICON, and "FCP" means Fibre Channel
FICON	Fibre Connection (a type of IBM mainframe channel capable of much greater throughput than ESCON)
HCD	Hardware Configuration Definition
IOCP	Input/Output Configuration Program
IOGEN	Process (or end result) of defining logical paths to devices and all relevant characteristics for use by the mainframe channel system and mainframe operating systems
IP	Internet Protocol
LPAR	Logical Partition
MIF	Multiple Image Facility is a feature of IBM mainframes that allows LPARs to share channels and devices
MIH	Missing-Interrupt Handler
MIM/MIA	Multi-Image Management/Multi-Image Allocation
MVS	Multiple Virtual Storage – a predecessor of the z/OS operating system, and is often used when referring to z/OS
NFS	Network File System
SCSI, iSCSI	Small Computer System Interface, Internet SCSI
UCB	Unit Control Block
UFS	UNIX file system (also called Berkeley fast file system, or FFS)

VOLSER	Volume Serial Number – a unique identifier of a tape or virtual tape
ZFS	Zettabyte file system (not the IBM z-series file system, zFS)
z/OS	Primary IBM mainframe operating system for batch and non-web online systems; includes TSO/ISPF/PDF for data processing users, and CICS, IMS, Websphere for online and web users

3. Site Preparation/Requirements

The MDI is often installed in paired sites—a production facility and a disaster recovery (DR) site. The documentation below discusses the site preparation for a production site. Installation at a DR location is similar. If you have a dedicated configuration at the DR site, it should mirror the production configuration. If you will have a shared configuration at the DR site, the Luminex installation team will work with you to tailor the site preparation and equipment requirements to meet your specific needs.

3.1 Communication-Network Connections

Most MDI installations will need network connectivity to the data-center TCP-IP network. The standard configuration consists of a minimum of 2 network connections per MDI. Ideally these connections would be available for remote access via a customer provided VPN. The first network port is used for general IP communication and for the browser-based GUI (MDI administration/tape management). The second network port is used as a service port utilizing the Service Processor (eg. Drac, iLO, ILOM) port of the MDI system. This port provides low-level service access to MDI.

The customer must provide IP addresses, gateway addresses, and net mask information for these ports.

3.2 Equipment Rack Requirements

MDIs fit into a standard 19” rack with a depth of at least 30 useable inches to accommodate the components and cables. Each MDI typically occupies 2U of rack space, but check with your Luminex installation team for exact rack requirements.

3.3 Storage-Network

MDI is usually linked to the storage device via direct cable or through a switched network that is isolated from the communication-network. Depending on the storage device, this will either be block-level access (SAN) or file-level access (NAS). SAN usually requires clustering software in MDI, and Fiber Channel or SCSI protocol. NAS requires a file system implemented at the storage device and Ethernet IP protocol.

The physical connection usually needs to be high speed to support the mainframe throughput requirements. The physical layer consists of adapters (ports) and cables. It may also include switches and transceivers.

Cloud-based storage has its own requirements and is covered in documentation specific to that cloud storage vendor.

Unless otherwise specified in the sales agreement, all storage-network cables, switches, and switch-transceivers must be provided by the customer or storage device vendor. Unless otherwise specified, the storage device vendor will provide the storage-device transceivers. Unless otherwise specified, Luminex will provide MDI transceivers. The type of transceivers must be agreed upon by all parties.

For optical transceivers, single/multi mode, and light frequency (nm) must match for the transceivers at both ends of each cable, and the maximum speed should match.

3.3.1 Storage-Network: Switch

Usually each port on each storage device is connected to more than one MDI. This requires at least one switch to which multiple cables are attached. The switch needs to provide sufficient point-to-point bandwidth to meet the (mainframe) throughput requirements. The type of switch depends on the protocol, cable type, and transceiver requirements.

This storage-network switch may be located in the equipment rack for the MDI systems or in a location closer to the storage devices.

Best performance can be realized if the MDI data port, network storage devices, and the switch are configured for an MTU size of 9000 (Jumbo Frames).


3.3.2 Storage-Network: Transceivers

The most common type of removable transceiver is an SFP module (small form-factor pluggable transceiver). Some switch vendors require their own brand of SFP in the switch. The labeling on an SFP is usually not sufficient to describe all of the characteristics of the transceiver.


Usually it is necessary to look up the model number in the SFP vendor's documentation to verify that it will meet all of the requirements of the switch or adapter in which it will be inserted.

Transceivers that support 2/4/8 Gigabits are usually not compatible with those that support 1/2/4 Gigabits; they negotiate to the lowest common speed if at all.

Luminex typically uses optical SFPs, however there are some similar sized SFP+ transceivers that support copper cables. SFP+ may have additional requirements and restrictions that are not covered here. Some SFP+ transceivers are built into the ends of the copper cable.



Short-wave (top left and bottom left), long-wave (top right), 10 GbE (bottom right)



CISCO SFP+ with attached copper cable

3.3.3 Storage-Network: Cables

The cables used for the storage-network are either optical or copper.

The appropriate type of cables must be in place prior to installing MDI. Any cables that must be routed under floors, through ceilings, etc., must be obtained by the customer and be in place prior to installing MDI.

With the appropriate adapter cards, MDI can support 10-gigabit copper or optical fiber Ethernet (10GbE). However, 10GbE copper can only be used if the storage equipment is within the distance limit for the selected type of cable/port (physical layer).

The distance limits for optical fiber transceivers and cables are listed below.

Optical Transceiver Wave Length	Optical Mode	Optical Core Diameter (in micrometers) -- Cable Type	Most Restrictive Cable Length Limit
850nm	multi-mode	62.5 -- FDDI (obsolete)	26 meters
		62.5 -- OM1 (orange cable)	33 meters
		50 -- OM2 (orange cable)	82 meters
		50 -- OM3 (aqua cable)	300 meters
		50 -- OM4 (aqua cable)	400 meters
1310nm	single-mode	8.3 -- (yellow cable)	10 kilometers

If optical couplers are used to directly join cables (max length is reduced, and the coupler is a likely source of trouble), or if the cable-run distances approach the corresponding limit, then further research will be required to confirm that the cable/transceiver/adaptor is sufficient for the required distance.

3.4 Mainframe Channels

MDI can support FICON or ESCON mainframe channels in a point-to-point or networked configuration. The mainframe defines what control units and devices it expects attached to each channel by virtue of control unit address and device address. This is known as an IOGEN.

Typically the mainframe acts as though it is communicating directly to each device, however it is the control unit that interprets the mainframe requests and gives direction to the device. The physical MDI can act as one or more logical control units and emulates 16 tape devices per logical control unit.

There will be one or more channel interface ports on the MDI system. Each port can be connected directly to a mainframe channel or to a channel-protocol switch. If the switch is configured to logically tie more than one mainframe channel to MDI, or to logically tie more than one MDI port to a mainframe channel, then the switch number and switch port (on the MDI side of the switch) must be defined in the mainframe IOGEN in addition to the control unit address and device address.

ESCON is an older type of optical mainframe channel that is expected to be discontinued in future mainframes. It supports a bandwidth of 17 Megabytes per second, although most of the original ESCON devices cannot achieve this rate.

FICON is the most common type of optical mainframe channel. The FICON bandwidth depends on the capabilities of the FICON channel card installed in the mainframe.

3.4.1 ESCON Directors and FICON Switches

Both of these are switches for the corresponding mainframe channel protocol. It is up to the customer to provide direct connection from the mainframe channel to the corresponding port on the MDI system, or to provide this connection through a switch. When high redundancy is required, there is usually a separate switch for each port on the MDI system, and these switches are connected to multiple channels on the mainframe.

The specific design of connections between mainframe channels, switches, and MDIs must be defined by the customer and implemented in their mainframe IOGEN.

3.4.2 ESCON Cables

The MDI channel interface for ESCON requires cables with a small MTRJ plug. Cables with the large "D"-shaped plug (Retractable Shroud Duplex , RSD) will require a conversion cable to the small MTRJ plug. These should be obtained and in place prior to the installation of MDI. The customer will have to determine which type of plug is needed at the mainframe (or ESCON director) end of the cable.



ESCON small MTRJ plug (left, standard for MDI) and large RSD plug (right, conversion cable required)

3.4.3 FICON Cables

MDI supports both the long-wave (typically a yellow cable) and short-wave (typically an orange or aqua cable) interfaces. The customer must provide the type of interface on the mainframe, or director/switch, port that is connected to the MDI system.



The table below indicates the various cable specifications that are supported, and their distance limitations. These are the same cables as Fibre Channel but are usually more restrictive on distance than Fibre Channel. Although mode-conversion adapters exist (rare), the normal usage is multi-mode (cheap LASER or LED) on short-wave fibers for short distances, and single-mode (LASER) on long-wave fibers for long distances.

Specification	Wave-length	Fiber Diameter (microns)	Color	Maximum Distance (meters)
OM1	short-wave (850nm)	62.5	Orange	21
OM2	short-wave (850nm)	50	Orange	50
OM3	short-wave (850nm)	50	Aqua	150
OM4	short-wave (850nm)	50	Aqua	190
-	long-wave (1310nm)	9	Yellow	20000

3.5 Remote Access

Remote access is required by Luminex to provide the highest level of support quality. Remote access is used to diagnose hardware or software issues, and as a means for performing file transfers for logs and upgrades.

3.5.1 VPN

The preferred means of remote access is through the use of a customer provided VPN. Luminex support has experience with many different VPNs.

3.5.2 WebEx (CISCO)

This method usually requires a PC at the customer site that can be dedicated to WEBEX for the duration of the support session. Use of a user's workstation will impact that their ability to get other work done during the session.

4. MDI: Hardware Specifications

Below are the components that comprise a MDI system.

4.1 Control Unit Specifications

MDI is a rack mountable system. The system is configured with dual internal hard drives required for mirroring, and dual power supplies.

4.2 Power Requirements

Each MDI system has 2 redundant power supplies that can accommodate 100-240V with a typical current draw of 7-9A at 100V. The equipment requires a power cable with a IEC60320 C13 (female) connector.

The system is normally shipped with power cables containing the IEC60320 C13 (female) connector and a NEMA 5-15P (male) plug. Other power cables can be provided upon request.



IEC60320 C13 female connector (MDI) and NEMA 5-15P male plug (rack power)



Some equipment racks (Power Distribution Units) require this IEC60320 C14 male plug (rack power) instead of the NEMA 5-15P male.

5. Miscellaneous Components

Exclusive of the actual storage device, below is a list of potential components that may be required in the installation:

- 1 GbE or 10 GbE Network Switch
- Fibre Channel Switch
- ESCON or FICON cables, and transceivers for customer supplied switches

6. **Software Installation Prerequisites**

The installation of MDI software requires information prior to its configuration regarding the mainframe IOGEN. Using the IOGEN to determine the configuration is described in the *Mainframe Configuration* section.

Additionally the storage must be pre-configured prior to the completion of the MDI installation. This is described in the *Configuring Storage* section.

7. Mainframe Configuration

Each physical MDI appears to the mainframe as a collection of tape control units (each control unit has 16 tape devices). Each of these tape control units will be configured by Luminex to set the device characteristics and storage device characteristics appropriate to the intended purpose as defined by the customer.

The process of configuring the mainframe (IOGEN, tape management, allocation steering) involves the following steps.

- Identify the different sets of device/storage characteristics.
- Assign control units with these characteristics to each MDI.
- Design the logical paths connections from the mainframes to the MDI systems.
- Define the design in the mainframe IOGEN.
- Make the corresponding changes to tape management and allocation steering.
- Install and connect the equipment.

7.1 Tape Device Characteristics

Typically, all of the tape devices under a given control unit have the same tape model, storage device(s), and replication characteristics. For every needed variation of these, there should be at least one logical control unit (and its associated 16 tape devices). Common variations are:

- the need for 3590 and 3490 tape models,
- drives associated with VOLSERS that need to be replicated (versus those that do not), or
- tape drives associated with storage reserved for a special purpose.

Usually, one thinks in terms of VOLSER ranges and the assignment of a VOLSER range to a particular tape storage system. More than one tape storage system can be created, effectively, inside MDI by associating a group of tape drives with each set of characteristics (variation).

The 3490 drives use the lower 22 bits of the block ID to note the absolute block number within a tape cartridge. The upper 10 bits are used for rough positioning and other status, but these 10-bits can usually be ignored. The 3590 drives use all 32-bits of the block ID to note the absolute block number. The difference in purpose of the upper 10 bits causes incompatibilities between 3490 and 3590 devices. For this reason, VOLSERS for each must stay within their family of tape drives (34xx vs. 35xx).

After their initial introduction, the 3590 drives came out with a feature to emulate 3490 drives. This permitted much more data to be placed on a cartridge yet continue to use the lower 22-bit method of block IDs. For the purposes of this planning guide, we will use "3E90" to signify this type of virtual emulation. It is not common except for IBM DFHSM customers who needed more cartridge capacity but were unable to introduce native 3590 drives into their environment.

Here is a checklist of characteristics to consider. Use it to identify all of the variations that will be needed.

Characteristic	Choice
3490, 3E90, or 3590	
Owning TMC	
Connected LPARs/MFs	
Peak number of Allocated Tape Drives (per LPAR)	
VOLSER Range	
Maximum Gigabytes of Data per VOLSER	
Storage Device(s)	
Writable Storage Units	
Replication Destination	
Encrypt?	
Maximum Data Created per Day	
Department/Client	
Read-Only Storage Units	

The "Peak number of Allocated Tape Drives" will dictate the minimum number of tape devices that must be defined for this group of characteristics. More can be defined to reduce the need for device switching across LPARs/Mainframes and to provide for growth or unanticipated peaks.

An analysis of peak throughput is typically conducted to determine the required number of mainframe channels and MDIs (and ports). This is usually done based on RMF, SMF, and tape catalog information. Once the number of MDIs has been determined, the layout of logical control units on each MDI can begin. Typically, it is beneficial to configure each MDI identically. This makes it possible to swap MDIs or to use the same configurations settings on each MDI. So, if two MDIs are required, variation 1 consists of 32 x 3490 drives, and variation 2 consists of 128 x 3590 drives, the following configuration could be set up on each MDI.

Control Unit x'00', Devices x'00' to x'0F', Type=3490
Control Unit x'01', Devices x'00' to x'0F', Type=3590
Control Unit x'02', Devices x'00' to x'0F', Type=3590

Control Unit x'03', Devices x'00' to x'0F', Type=3590

Control Unit x'04', Devices x'00' to x'0F', Type=3590

On each MDI there would be 16 x 3490 and 64 x 3590 tape devices. The combination of both results in the desired number.

For redundancy purposes, it is helpful to plan extra tape devices on each MDI (or extra MDIs) so that the entire work load can be handled by the remaining MDIs if one should fail or be taken down for maintenance.

The next section discusses unit addressing and device numbering, which is followed by diagrams of logical paths in which these items are presented.

7.2 IOGEN Addresses

For the purpose of routing IO to the correct device on the mainframe channel, each control unit must be assigned a logical control unit address. If there are to be 16 control units, then these would typically be numbered x'00' to x'0F'. The tape devices should be assigned unit addresses of x'00' to x'0F'. Taken together the control unit address (CU) and unit address (UA) uniquely identify a particular tape device within the Luminex MDI. CU/UA x'0407' is not the same tape device as CU/UA x'0C07'. These CU/UA numbers will appear in the device tag field of the mainframe DevServ QueryTape RCD command.

If there is an ESCON director or FICON switch between the mainframe and MDI, it does NOT have to be identified in the IOGEN if it is configured to have a static (non-switchable) route between two switch ports. Since the IO traffic from the channel can only go to the attached MDI, it is as if the MDI was directly connected to the mainframe channel.

Most of the time, the ESCON director or FICON switch has dynamic switching. This means that the switch inspects the destination address in the mainframe IO frame and passes it to the designated port on the MDI side of the switch. When this frame arrives at the MDI system, the identity of the port on the mainframe side of the switch is provided, and any responses from MDI to that logical path are sent back through the switch to the cable attached to the mainframe-side port of the switch. Each switch is assigned a unique number for clarity in the IOGEN.

When the logical control unit is defined in the IOGEN, the connecting information of which channel path ID (CHPID), and any switch number and the switch port number (on the MDI-side of the switch) is specified. This is sufficient information for the mainframe to send the IO frames to the correct tape device. Since each mainframe has its own channel adapters and a different logical path to the MDI logical control units, there must be an IOGEN for each mainframe.

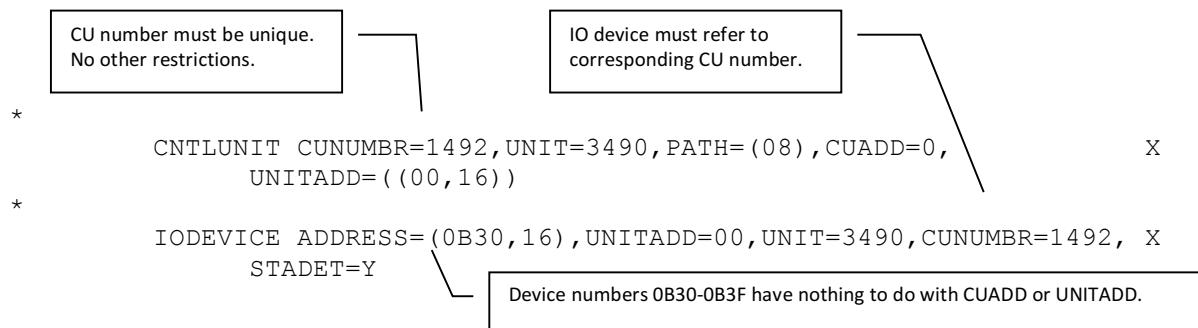
7.3 IOGEN Device Numbers

The IOGEN does not require that UCB device numbers have any similarity to the addresses that are used within the logical path to get IO frames to the correct equipment. It is a typical practice

to group all of the device numbers for a particular vendor's tape equipment into consecutive device numbers, no matter how many switches, channels, or logical control units have been defined. Generally all devices for a particular group of characteristics are given consecutive device numbers. However, some customers choose to give each MDI its own set of consecutive device numbers.

Therefore, it is extremely helpful to put the device numbers on the channel path map next to the corresponding logical control unit and its devices.

The IOGEN also requires an arbitrary number for each logical control unit. This number is only used by the DEVICE definitions to point to the corresponding control unit.



IMPORTANT: Luminex requires that each set of 16 tape devices point to only one logical control unit (only one value in the IODEVICE CUNUMBR parameter). This is because no two MDI logical control units can talk to the same device. The devices in each MDI logical control unit are completely independent of any other logical control unit (even if within the same MDI). The old-style tape drives with dual control units cannot be defined for MDI tape devices.

NOTE: HCD uses the terminology convention "addresses" which are used on the mainframe channel, IOGEN "numbers" are not. IOCP statements usually follow this convention. The exception is that "IODEVICE ADDRESS=" actually specifies "device numbers" which are not used on the mainframe channel.

7.4 Logical Path Diagram

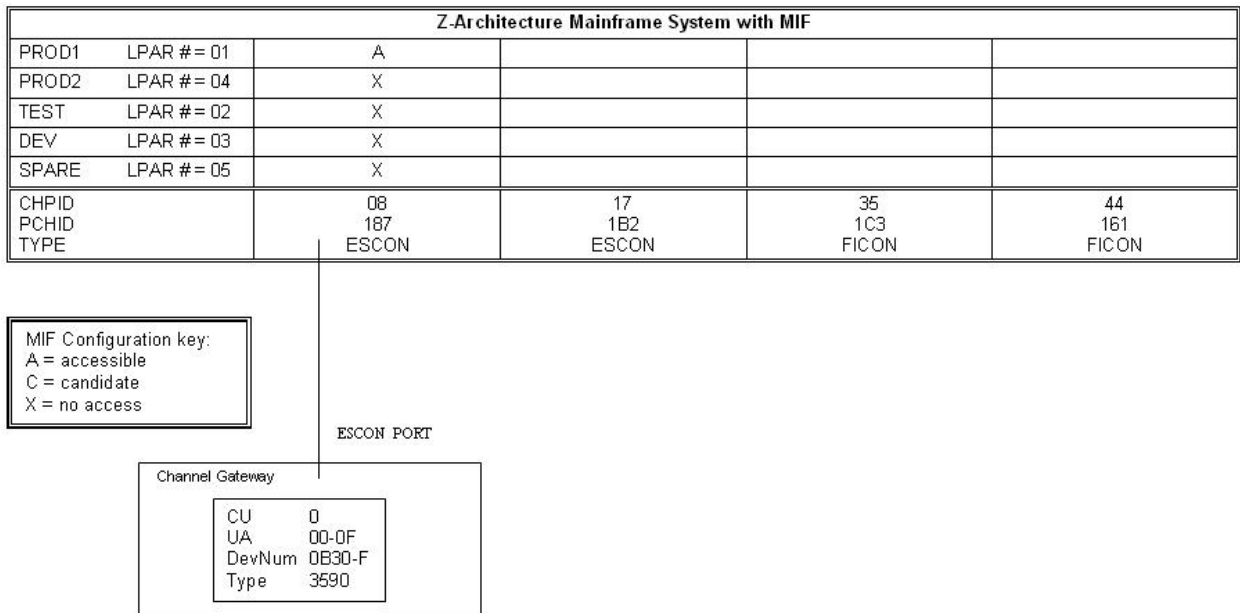
A logical path is one complete series of connections from a mainframe LPAR through a channel adapter port, through any intervening channel-protocol switches, through the channel-interface port on a MDI system to a specific logical control unit. When there are multiple logical control units within a MDI system, we normally just draw one path and enumerate the control units at the MDI system. Likewise, when there are multiple LPARs on the mainframe we normally just draw one path from the mainframe channel adapter port.

IOCP examples are provided to illustrate the essence of the relationship between the logical path diagrams and the IOGEN structures. An explanation of IOCP statement syntax can be found in IBM manual SB10-7037, "System z, Input/Output Configuration Program User's Guide for ICP IOCP." However, it is rare to actually code IOCP statements. Most z/OS IOGENs are accomplished via mainframe HCD or HCM using a panel-driven application that has little similarity to these IOCP statements. HCD can generate IOCP statements which can be used for comparison to these examples.

7.4.1 Simple Logical Path

Here is a simple example of an ESCON logical path from one LPAR to one MDI logical control unit (CU), directly connected (no switch).

Figure 7-1. One LPAR to one logical control unit



If this had been a FICON logical path, the diagram would be exactly the same except it would have said FICON instead of ESCON. The adapter cards, cables, connectors on the ends of the cables, and any intervening switches are completely different between ESCON and FICON.

The only difference in the IOGEN on the mainframe is that the mainframe channel must be defined as CNC for ESCON, whereas the FICON channel must be defined as FC (not FCP).

Also, ESCON is limited to control unit addresses x'00' to x'0F', but FICON can go up to x'FE'. IOGEN will be covered in more detail in another section. Here are the IOCP statements that correspond to the logical path this diagram.

```

RESOURCE PART=( (CSS (0) , (PROD1,1) , (PROD2,4) , (TEST,2) , (DEV,3) ,    X
                (SPARE,5) ) )
CHPID  PATH=(CSS (0) , 08) , PCHID=187 , TYPE=CNC,                        X
        PARTITION=(CSS (0) , (PROD1) , (=) )
*
CNTLUNIT CUNUMBR=1492 , UNIT=3490 , PATH=(08) , CUADD=0,                  X
        UNITADD=( (00,16) )
*
IODEVICE ADDRESS=(0B30,16) , UNITADD=00 , UNIT=3490 , CUNUMBR=1492 , X
        STADET=Y

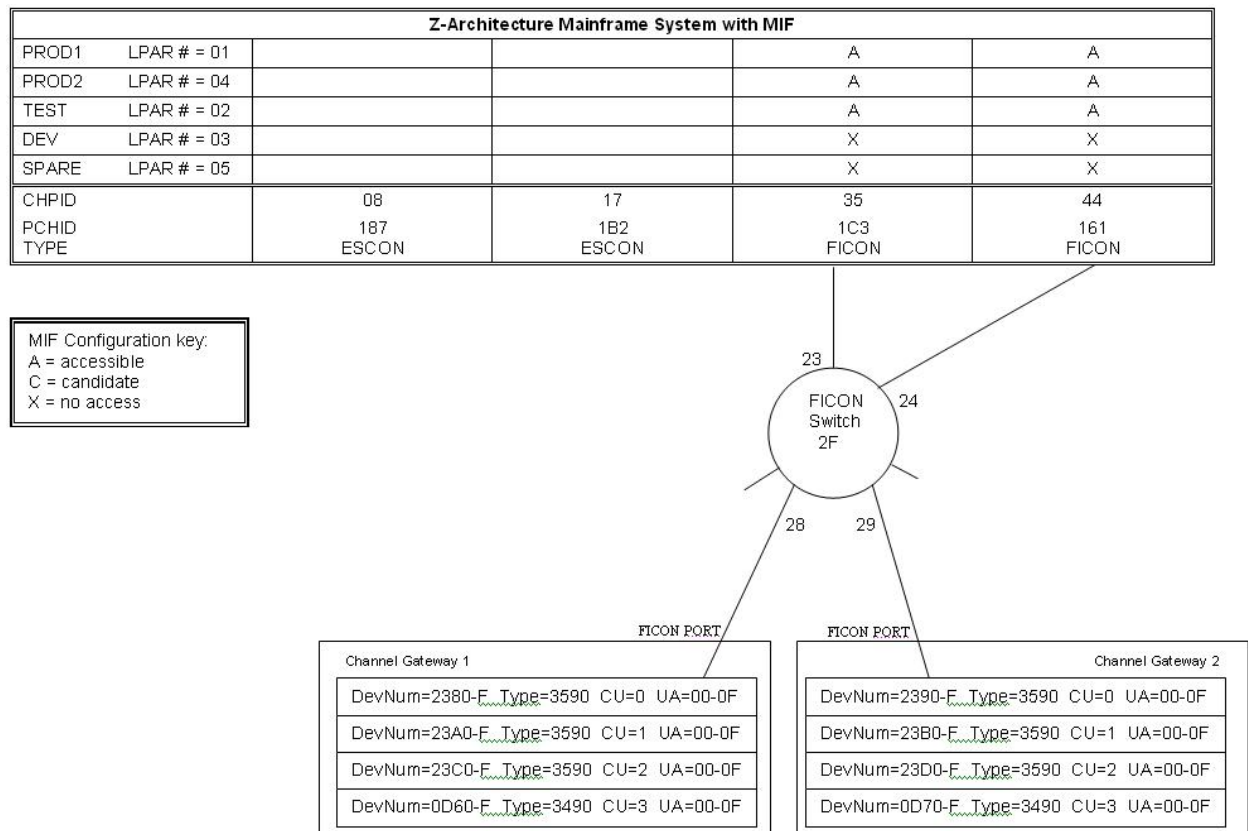
```

The CUNUMBR= is completely arbitrary. The only requirement is that the same value be used in the CNTLUNIT and IODEVICE definitions.

7.4.2 Logical Paths through Switch

Here is an example of FICON logical paths from three LPARs through two mainframe channels to four logical control units (CU) in each single-port MDI through a FICON switch. The first three control units are device type 3590, the last is 3490.

Figure 7-2. Three LPARs, one switch, 2 to one logical control unit



From the diagram, it should be apparent that the number of LPARs and the number of logical control units does not affect the link lines connecting the mainframe, switch, and MDIs. However, it is very important to label the following items on the path diagram:

- channel path IDs (CHPID),
- switch number,
- switch ports,
- control unit addresses,
- device unit addresses, and
- device numbers.

PCHID is only needed if the diagram will be used to verify that the FICON cable is plugged into the correct channel port on the mainframe. Older mainframes did not have this flexible mapping of CHPIDs to physical (PCHID) ports.

The following IOCP statements correspond to the logical paths in this diagram.

```

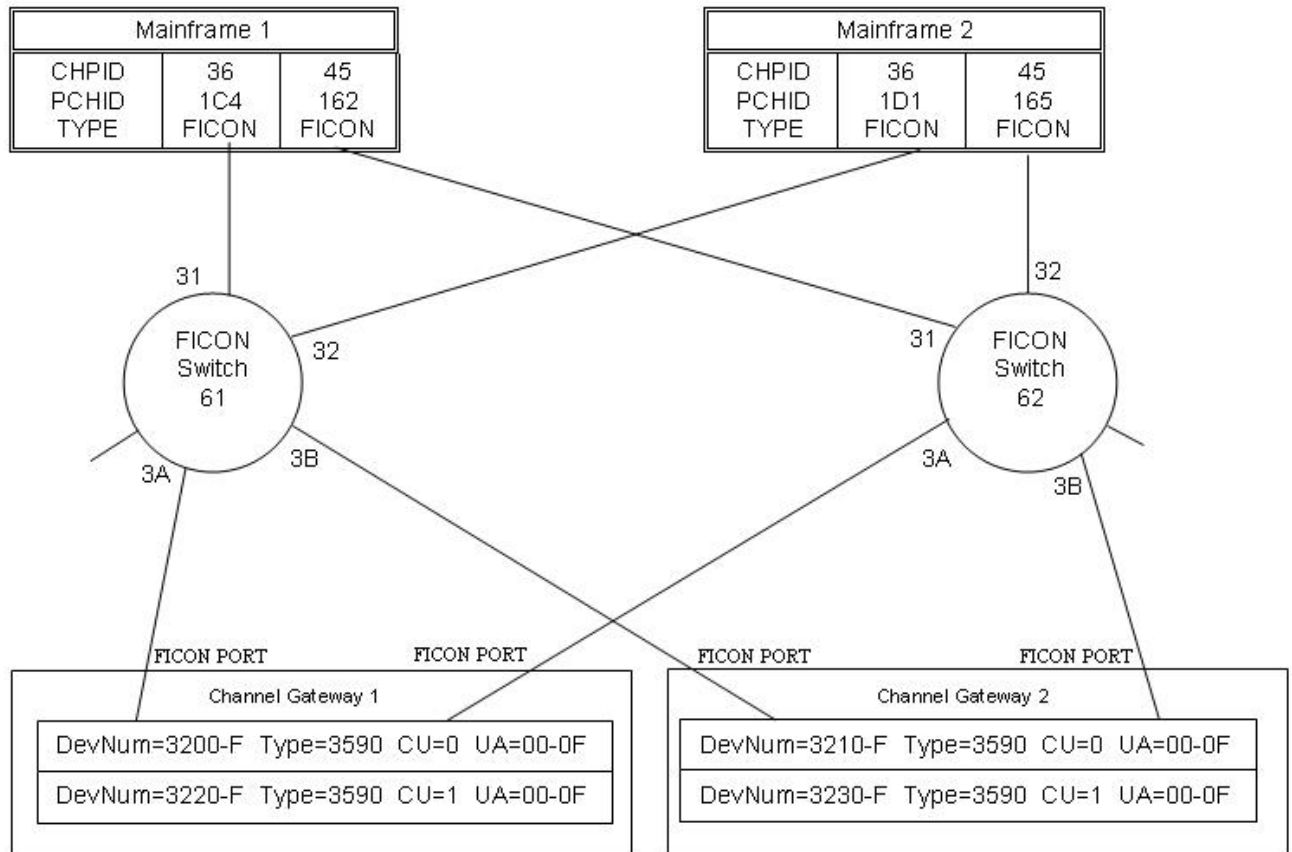
RESOURCE PART=((CSS(0),(PROD1,1),(PROD2,4),(TEST,2),(DEV,3),
              (SPARE,5)))
CHIPID PATH=(CSS(0),35),PCHID=1C3,TYPE=FC,SWITCH=2F,
PARTITION=(CSS(0),(PROD1,PROD2,TEST),(=))
CHIPID PATH=(CSS(0),44),PCHID=161,TYPE=FC,SWITCH=2F,
PARTITION=(CSS(0),(PROD1,PROD2,TEST),(=))
*
CNTLUNIT CUNUMBR=2380,UNIT=3590,CUADD=0,UNITADD=((00,16)),
PATH=(35,44),LINK=(28,28),
UNITADD=((00,16))
CNTLUNIT CUNUMBR=23A0,UNIT=3590,CUADD=1,UNITADD=((00,16)),
PATH=(35,44),LINK=(28,28),
UNITADD=((00,16))
CNTLUNIT CUNUMBR=23C0,UNIT=3590,CUADD=2,UNITADD=((00,16)),
PATH=(35,44),LINK=(28,28),
UNITADD=((00,16))
CNTLUNIT CUNUMBR=0D60,UNIT=3490,CUADD=3,UNITADD=((00,16)),
PATH=(35,44),LINK=(28,28),
UNITADD=((00,16))
*
CNTLUNIT CUNUMBR=2390,UNIT=3590,CUADD=0,UNITADD=((00,16)),
PATH=(35,44),LINK=(29,29),
UNITADD=((00,16))
CNTLUNIT CUNUMBR=23B0,UNIT=3590,CUADD=1,UNITADD=((00,16)),
PATH=(35,44),LINK=(29,29),
UNITADD=((00,16))
CNTLUNIT CUNUMBR=23D0,UNIT=3590,CUADD=2,UNITADD=((00,16)),
PATH=(35,44),LINK=(29,29),
UNITADD=((00,16))
CNTLUNIT CUNUMBR=0D70,UNIT=3490,CUADD=3,UNITADD=((00,16)),
PATH=(35,44),LINK=(29,29),
UNITADD=((00,16))
*
IODEVICE ADDRESS=(2380,16),UNITADD=00,UNIT=3590,CUNUMBR=2380, X
STADET=Y
IODEVICE ADDRESS=(23A0,16),UNITADD=00,UNIT=3590,CUNUMBR=23A0, X
STADET=Y
IODEVICE ADDRESS=(23C0,16),UNITADD=00,UNIT=3590,CUNUMBR=23C0, X
STADET=Y
IODEVICE ADDRESS=(0D60,16),UNITADD=00,UNIT=3490,CUNUMBR=0D60, X
STADET=Y
IODEVICE ADDRESS=(2390,16),UNITADD=00,UNIT=3590,CUNUMBR=2390, X
STADET=Y
IODEVICE ADDRESS=(23B0,16),UNITADD=00,UNIT=3590,CUNUMBR=23B0, X
STADET=Y
IODEVICE ADDRESS=(23D0,16),UNITADD=00,UNIT=3590,CUNUMBR=23D0, X
STADET=Y
IODEVICE ADDRESS=(0D70,16),UNITADD=00,UNIT=3490,CUNUMBR=0D70, X
STADET=Y

```


7.4.3 Logical Paths from Two Mainframes through Two Switches

Here is an example of two mainframes connected through two switches to three MDIs. Each mainframe is connected to each switch for redundancy. Each switch is connected to each MDI for redundancy. To reduce the complexity of the sample IOCP statements, we will assume that all LPARs on each mainframe have access to all control units on each MDI.

Figure 7-3. Three LPARs, one switch, 2 to one logical control unit



Although it is not necessary to use similar switch port numbers for the paths through each switch, many systems programmers follow this practice.

The following IOCP statements are for mainframe 1. Since all of the same port numbers and CHPIDs are for both mainframes, the IOCP statements would be the same except for the PCHID values on mainframe 2.

```

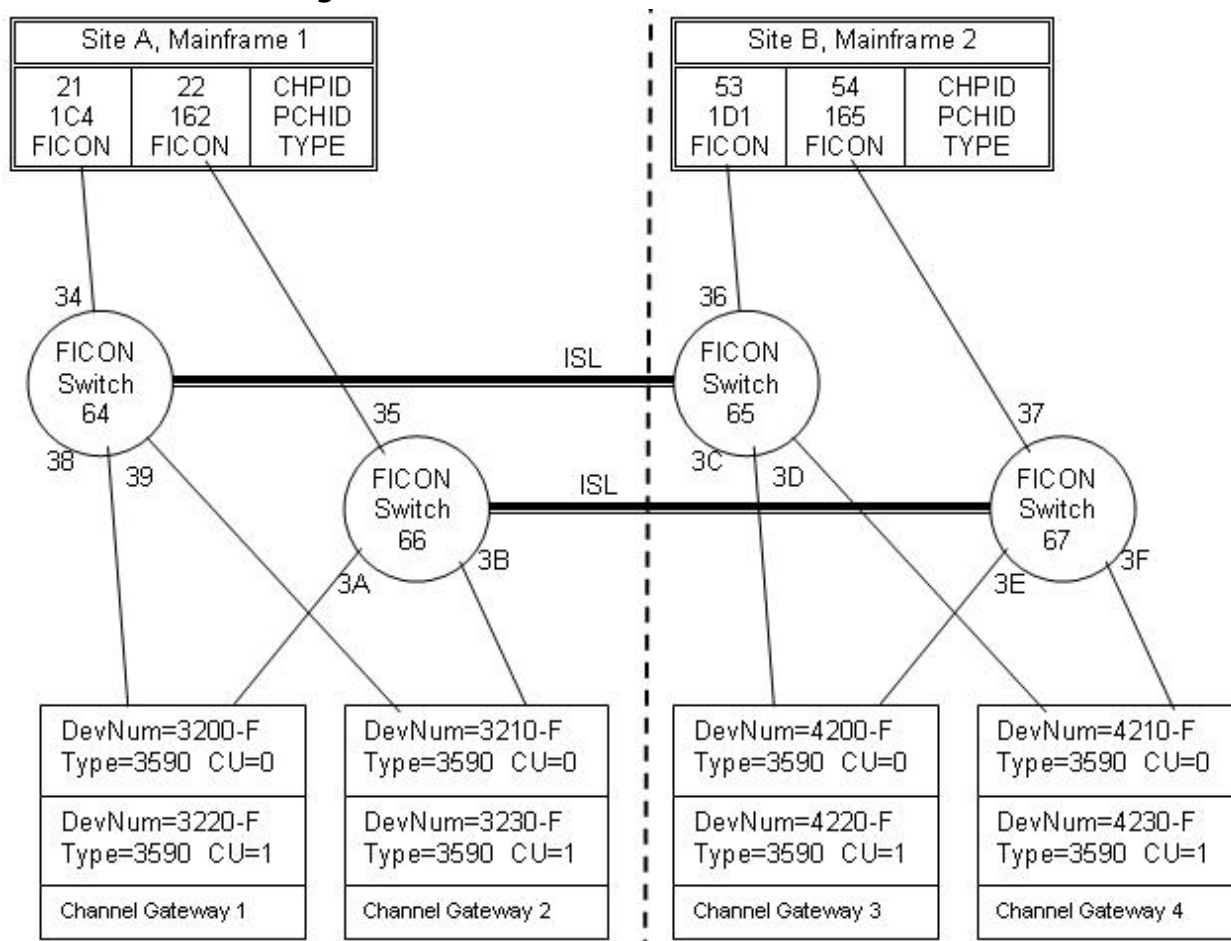
CHPID  PATH=(CSS(0),36),PCHID=1C4,TYPE=FC,SWITCH=61,      X
        PARTITION=(CSS(0),(*),(=))
CHPID  PATH=(CSS(0),45),PCHID=162,TYPE=FC,SWITCH=62,      X
        PARTITION=(CSS(0),(*),(=))
*
CNTLUNIT CUNUMBR=3200,UNIT=3590,CUADD=0,UNITADD=((00,16)),  X
        PATH=(36,45),LINK=(3A,3A),                        X
        UNITADD=((00,16))
CNTLUNIT CUNUMBR=3220,UNIT=3590,CUADD=1,UNITADD=((00,16)),  X
        PATH=(36,45),LINK=(3A,3A),                        X
        UNITADD=((00,16))
*
CNTLUNIT CUNUMBR=3210,UNIT=3590,CUADD=0,UNITADD=((00,16)),  X
        PATH=(36,45),LINK=(3B,3B),                        X
        UNITADD=((00,16))
CNTLUNIT CUNUMBR=3230,UNIT=3590,CUADD=1,UNITADD=((00,16)),  X
        PATH=(36,45),LINK=(3B,3B),                        X
        UNITADD=((00,16))
*
IODEVICE ADDRESS=(3200,16),UNITADD=00,UNIT=3590,CUNUMBR=3200, X
        STADET=Y
IODEVICE ADDRESS=(3210,16),UNITADD=00,UNIT=3590,CUNUMBR=3210, X
        STADET=Y
IODEVICE ADDRESS=(3220,16),UNITADD=00,UNIT=3590,CUNUMBR=3220, X
        STADET=Y
IODEVICE ADDRESS=(3230,16),UNITADD=00,UNIT=3590,CUNUMBR=3230, X
        STADET=Y

```

7.4.4 Logical Paths through Cascaded Switches

Here is an example of cascaded FICON switches. This is only possible if the FICON switches are made by the same vendor and can be cascaded. Check the switch vendor documentation (or IBM Redbook SG24-6266) for other cascaded FICON switch requirements (Insistent Domain ID, High Integrity).

Figure 7-4. Two sites, Cascaded FICON Switches



The use of cascaded FICON switches across sites is not necessary if site B is only used in disaster recovery situations, and there is continuous replication of tape data from site A to site B (not shown on the diagram).

Cascading is useful if the cost of the ISL links is not prohibitive, and either (a) each site has its own pool of VOLSERS, or (b) the cross-site VOLSER locking feature has been installed in the Luminex MDIs.

Usually a systems programmer will use the same switch port numbers at both sites, but this example is intended to show the relationship of the logical path diagram to the corresponding IOCP statements below. Since there is a cascaded FICON switch on each channel, the LINK= parameters must specify the switch and port closest to the corresponding MDI.

```

* SITE A, MAINFRAME 1
  CHPID PATH=(CSS(0),21),PCHID=1C4,TYPE=FC,SWITCH=64, X
    PARTITION=(CSS(0),(*),(=))
  CHPID PATH=(CSS(0),22),PCHID=162,TYPE=FC,SWITCH=66, X
    PARTITION=(CSS(0),(*),(=))
*
  CNTLUNIT CUNUMBR=3200,UNIT=3590,CUADD=0,UNITADD=((00,16)), X
    PATH=(21,22),LINK=(6438,663A), X
    UNITADD=((00,16))
  CNTLUNIT CUNUMBR=3220,UNIT=3590,CUADD=1,UNITADD=((00,16)), X
    PATH=(21,22),LINK=(6438,663A), X
    UNITADD=((00,16))
*
  CNTLUNIT CUNUMBR=3210,UNIT=3590,CUADD=0,UNITADD=((00,16)), X
    PATH=(21,22),LINK=(6439,663B), X
    UNITADD=((00,16))
  CNTLUNIT CUNUMBR=3230,UNIT=3590,CUADD=1,UNITADD=((00,16)), X
    PATH=(21,22),LINK=(6439,663B), X
    UNITADD=((00,16))
*
  CNTLUNIT CUNUMBR=4200,UNIT=3590,CUADD=0,UNITADD=((00,16)), X
    PATH=(21,22),LINK=(653C,673E), X
    UNITADD=((00,16))
  CNTLUNIT CUNUMBR=4220,UNIT=3590,CUADD=1,UNITADD=((00,16)), X
    PATH=(21,22),LINK=(653C,673E), X
    UNITADD=((00,16))
*
  CNTLUNIT CUNUMBR=4210,UNIT=3590,CUADD=0,UNITADD=((00,16)), X
    PATH=(21,22),LINK=(653D,673F), X
    UNITADD=((00,16))
  CNTLUNIT CUNUMBR=4230,UNIT=3590,CUADD=1,UNITADD=((00,16)), X
    PATH=(21,22),LINK=(653D,673F), X
    UNITADD=((00,16))
*
  IODEVICE ADDRESS=(3200,16),UNITADD=00,UNIT=3590,CUNUMBR=3200, X
    STADET=Y
  IODEVICE ADDRESS=(3210,16),UNITADD=00,UNIT=3590,CUNUMBR=3210, X
    STADET=Y
  IODEVICE ADDRESS=(3220,16),UNITADD=00,UNIT=3590,CUNUMBR=3220, X
    STADET=Y
  IODEVICE ADDRESS=(3230,16),UNITADD=00,UNIT=3590,CUNUMBR=3230, X
    STADET=Y
*
  IODEVICE ADDRESS=(4200,16),UNITADD=00,UNIT=3590,CUNUMBR=4200, X
    STADET=Y
  IODEVICE ADDRESS=(4210,16),UNITADD=00,UNIT=3590,CUNUMBR=4210, X
    STADET=Y
  IODEVICE ADDRESS=(4220,16),UNITADD=00,UNIT=3590,CUNUMBR=4220, X
    STADET=Y
  IODEVICE ADDRESS=(4230,16),UNITADD=00,UNIT=3590,CUNUMBR=4230, X
    STADET=Y

```

Here are the IOCP statements for the other mainframe.

```
* SITE B, MAINFRAME 2
  CHPID PATH=(CSS(0),53),PCHID=1D1,TYPE=FC,SWITCH=65, X
    PARTITION=(CSS(0),(*),(=))
  CHPID PATH=(CSS(0),54),PCHID=165,TYPE=FC,SWITCH=67, X
    PARTITION=(CSS(0),(*),(=))
*
  CNTLUNIT CUNUMBR=3200,UNIT=3590,CUADD=0,UNITADD=((00,16)), X
    PATH=(53,54),LINK=(6438,663A), X
    UNITADD=((00,16))
  CNTLUNIT CUNUMBR=3220,UNIT=3590,CUADD=1,UNITADD=((00,16)), X
    PATH=(53,54),LINK=(6438,663A), X
    UNITADD=((00,16))
*
  CNTLUNIT CUNUMBR=3210,UNIT=3590,CUADD=0,UNITADD=((00,16)), X
    PATH=(53,54),LINK=(6439,663B), X
    UNITADD=((00,16))
  CNTLUNIT CUNUMBR=3230,UNIT=3590,CUADD=1,UNITADD=((00,16)), X
    PATH=(53,54),LINK=(6439,663B), X
    UNITADD=((00,16))
*
  CNTLUNIT CUNUMBR=4200,UNIT=3590,CUADD=0,UNITADD=((00,16)), X
    PATH=(53,54),LINK=(653C,673E), X
    UNITADD=((00,16))
  CNTLUNIT CUNUMBR=4220,UNIT=3590,CUADD=1,UNITADD=((00,16)), X
    PATH=(53,54),LINK=(653C,673E), X
    UNITADD=((00,16))
*
  CNTLUNIT CUNUMBR=4210,UNIT=3590,CUADD=0,UNITADD=((00,16)), X
    PATH=(53,54),LINK=(653D,673F), X
    UNITADD=((00,16))
  CNTLUNIT CUNUMBR=4230,UNIT=3590,CUADD=1,UNITADD=((00,16)), X
    PATH=(53,54),LINK=(653D,673F), X
    UNITADD=((00,16))
*
  IODEVICE ADDRESS=(3200,16),UNITADD=00,UNIT=3590,CUNUMBR=3200, X
    STADET=Y
  IODEVICE ADDRESS=(3210,16),UNITADD=00,UNIT=3590,CUNUMBR=3210, X
    STADET=Y
  IODEVICE ADDRESS=(3220,16),UNITADD=00,UNIT=3590,CUNUMBR=3220, X
    STADET=Y
  IODEVICE ADDRESS=(3230,16),UNITADD=00,UNIT=3590,CUNUMBR=3230, X
    STADET=Y
*
  IODEVICE ADDRESS=(4200,16),UNITADD=00,UNIT=3590,CUNUMBR=4200, X
    STADET=Y
  IODEVICE ADDRESS=(4210,16),UNITADD=00,UNIT=3590,CUNUMBR=4210, X
    STADET=Y
  IODEVICE ADDRESS=(4220,16),UNITADD=00,UNIT=3590,CUNUMBR=4220, X
    STADET=Y
  IODEVICE ADDRESS=(4230,16),UNITADD=00,UNIT=3590,CUNUMBR=4230, X
    STADET=Y
```

7.5 MDI Configuration Specifications

The configuration of a MDI is very flexible in the number of devices supported and the configuration of these devices. Below are the maximum configurations and the typical values of each.

Name	Maximum	Typical
Number of Logical Paths per Interface Port*	32,768	varies
Number of LPARs per Interface Port	256	1-8
Number of Control Unit Images per Interface Port	FICON: 256, ESCON: 16	1-2
Maximum Devices per Control Unit Image**	16	16

* The total number of MDI Logical Paths is the number of LPARs times the number of Control Units times the number of CHPIDs. For example, if there are 4 LPARs each with 5 Control Units defined to 2 CHPIDs attached to the MDI system, the total number of Logical Paths is $4 \times 5 \times 2 = 40$. This amount cannot exceed the maximum.

** For tape control units, the mainframe HCD will not allow more than 16 devices to be defined per 3490 or 3590 control unit image. The control unit type would have to be defined as NOCHECK to exceed 16 devices for a given control unit image. MDI supports up to 256 devices per control unit, but most customers adhere to the IBM convention of 16 tape devices per CU.

Increasing the number of Control Unit Images or Devices does not increase the performance or data accessibility of the system. It is unusual for a site to require more than 128 total devices in a single MDI.

7.6 IOCP Parameters

Although seldom used, IOCP statements may be used by mainframe systems programmers as shorthand for conveying IOGEN requirements for devices and their control units. After the sample worksheet are some sample IOCP statements that show different ways of connecting MDIs to the mainframe.

A diagram of the configuration represented by the sample IOCP statements is also shown below.

Configuration Rules:

- Each ESCON/FICON interface can support up to 32,768 logical paths.
- Each CHPID/logical partition (LPAR)/control unit combination requires a separate logical path.
- Each control unit image can support up to 256 devices, however HCD prefers 16 devices.
- Each path to a given control unit address must list the same number of devices in the last column.
- A Dynamic ESCON/FICON switch (director) is required if more than one mainframe or more than one CHPID is connected to a single ESCON/FICON interface. An alternative is to install one ESCON/FICON interface for each mainframe/CHPID combination that will be connected to this MDI.
- Define extra (stand-by) devices to be prepared for a channel interface failure.
- If the underlying storage is not partitioned, all like devices can get to any corresponding media.

Sample IOCP Statements generated by HCD:

```
RESOURCE PART=((PROD1,1),(PROD2,4),(TEST,2),(DEV,3),(SPARE,5))
CHPID PATH=08,TYPE=CNC, X
PARTITION=((DEV),(PROD1,PROD2,TEST),REC)
CHPID PATH=17,TYPE=CNC, SWITCH=3E, X
PARTITION=((PROD1,PROD2),(TEST,DEV,SPARE))
CHPID PATH=35,TYPE=FC, SWITCH=2F, X
PARTITION=((PROD1,PROD2),(TEST,DEV,SPARE))
CHPID PATH=44,TYPE=FC, X
PARTITION=((PROD1,PROD2),(TEST,DEV,SPARE))
*
CNTLUNIT CUNUMBR=1492,UNIT=3490,PATH=(08),CUADD=0, X
UNITADD=((00,16))
CNTLUNIT CUNUMBR=1776,UNIT=3490,PATH=(17),CUADD=0,LINK=(E1), X
UNITADD=((00,16))
CNTLUNIT CUNUMBR=1783,UNIT=3490,PATH=(17),CUADD=1,LINK=(E1), X
UNITADD=((00,16))
CNTLUNIT CUNUMBR=1812,UNIT=3490,PATH=(35),CUADD=0,LINK=(28), X
UNITADD=((00,16))
CNTLUNIT CUNUMBR=2001,UNIT=3490,PATH=(35),CUADD=0,LINK=(29), X
UNITADD=((00,16))
CNTLUNIT CUNUMBR=2002,UNIT=3490,PATH=(44),CUADD=0, X
UNITADD=((10,16))
```

*

```
IODEVICE ADDRESS=(0B30,16),UNITADD=00,UNIT=3490,CUNUMBR=1492, X
  STADET=Y,TIMEOUT=Y
IODEVICE ADDRESS=(0C40,16),UNITADD=00,UNIT=3490,CUNUMBR=1776, X
  STADET=Y,TIMEOUT=Y
IODEVICE ADDRESS=(0C50,16),UNITADD=00,UNIT=3490,CUNUMBR=1783, X
  STADET=Y,TIMEOUT=Y
IODEVICE ADDRESS=(0D60,16),UNITADD=00,UNIT=3490,CUNUMBR=1812, X
  STADET=Y,TIMEOUT=Y
IODEVICE ADDRESS=(2370,16),UNITADD=00,UNIT=3490,CUNUMBR=2001, X
  STADET=Y,TIMEOUT=Y
IODEVICE ADDRESS=(2380,16),UNITADD=10,UNIT=3490,CUNUMBR=2002, X
  STADET=Y,TIMEOUT=Y
```

Comments:

- IODEVICE STADET and TIMEOUT can be set to Y or N. Typically, these are set to Y.
- CHPID can be SHARED, REConfigurable, or DEDICATED. Usually it is SHARED.
- CNC (ESCON) and FC (FICON) channel control units (CNTLUNIT) do not use the following parameters since they only apply to parallel channels (bus and tag): PROTOCOL, SHARED

MDI restrictions:

- The IODEVICE CUNUMBER= must specify only one CNTLUNIT CUNUMBER.

MDI recommendations:

- Always specify the IODEVICE UNITADD= parm. If not specified, then the unit address defaults to the last two digits of the IODEVICE ADDRESS= parm, which may not be what you want.
- Use IODEVICE UNITADD=00 unless there is some need to have a different starting unit address. MDI supports unit addresses in the range 00 to FF. Starting at 00 follows the IBM conventions for ESCON and FICON devices.

7.7 HCD IOGEN Settings

Terminology used below is based on the (IBM) HCD Planning manual for your operating system(s). You may be able to access IBM manuals on the Internet via the following Web link or some equivalent:

<http://www-03.ibm.com/servers/eserver/zseries/zos/bkserv/>

These settings presume that your mainframe supports MIF (or some equivalent), and that you have multiple LPARs. If this is not the case, then just skip the settings that are not available to you.

Any settings not listed below can be set according to your needs.

7.7.1 Channel Settings

When defining each ESCON/FICON channel for MDI, the following settings are relevant.

Channel Path Type	CNC or FC
SPAN	Supported
Operation Mode	DED, REC, or SHR
Managed	NO

7.7.2 Control Unit Settings

When defining the tape control unit, the following settings are relevant.

Control Unit Type	3490, or 3590 (preferred).
Define more than eight ports	NO
Unit Address; Number of Units	<p>Must match the configuration of MDI. Luminex prefers 16 units starting at address 00 for ALL defined control units.</p> <p>MDI has no restriction on the number of units, however HCD may require 16 units for some tape control unit types. You can bypass this restriction by changing the control unit type to NOCHECK; however some of the device settings for tape devices will not be available to you.</p>
Logical Address (CUADD)	<p>The unit settings can be:</p> <p>Implicit: Assigned in order as CUADD 0, CUADD 1, CUADD 2, etc.</p> <p>Explicit: Assigned, as desired, such CUADD 2, CUADD 3, CUADD 0, etc. In this example sequence, it is not required to assign CUADD 1.</p> <p>Note: The supported values of CUADD are 0-F for ESCON, and 00-FE for FICON.</p>
Protocol; I/O Concurrency Level	Both settings are IGNORED for CNC ESCON and FICON channels

7.7.3 Device Settings

When defining the tape devices, the following settings are relevant.

Device Type	3490 or 3590
Connected to CUs	Should be the control unit number you defined for MDI
TIMEOUT	Ignored
STADET	YES
Explicit dev candidate list	YES
Define device candidate list	Select only one partition name. This prevents other LPARs from using this device. All devices on the control unit must select the same partition name. However, you can select multiple partitions if you have a method for only allowing one partition at a time to use the tape device. (online/offline, MIM)
Define device-to-operating system configuration	Select the operating system configurations for which this device should appear in that IOGEN. You can define this device to all of your operating system configurations, even though you have only put it in the candidate list for one partition (see define device candidate list in the row above).

7.7.4 Device Settings/MVS Configurations

These settings are specific to each device, but only for z/OS or Multiple Virtual Storage (MVS) LPAR. They appear when you select an MVS operating system configuration. The panel title is “Define Device Parameters/Features”. Settings not listed below can be set according to your needs.

Settings	Y/N	Notes/Information
DYNAMIC	YES	This allows you to dynamically reconfigure this device without re-IPLing your mainframe if you have implemented HCD dynamic reconfiguration correctly. See the HCD manual for instructions on how to accomplish this if you have not yet done this. This is a very useful feature to implement.
LOCANY	YES	Use YES unless you have UCB 24-bit addressing restrictions in your operating system, with tape management software, or with application programs that do I/O directly to the device.
LIBRARY	NO	Even if MTL=YES is specified (below), this value must be NO. LIBRARY=YES is only supported by IBM equipment.
AUTOSWITCH	YES/NO	The system is compatible with "Autoswitch=yes" which allows automatic switching of tape devices across LPARs within a z/OS Sysplex. Please consult the z/OS Sysplex guides for Sysplex and GRS tape sharing considerations.
LIBRARY-ID		(do not specify a value unless MTL=YES)
LIBPORT-ID		(do not specify a value unless MTL=YES)
MTL	YES/NO	If you want to treat these devices as a manual tape library, you may change the appropriate settings accordingly. However, you are responsible for testing and documenting how to do this, since there are different considerations depending on which tape management software you are using.
SHARABLE	NO	
COMPACT	NO	

7.8 Missing Interrupt Handler (z/OS)

By default, the missing interrupt handler (MIH) is generally set to 3 minutes on the mainframe. This is the time in which a command must complete. With higher capacity tapes, this time may not be long enough to perform time-consuming operations such as rewind, unload, or space commands. Some NFS storage may also have significant delays. It is recommended that the MIH be set to 20 minutes, to give ample time for these commands and any error recovery that may be required. To determine the current MIH, at the mainframe system console enter:

```
D IOS,MIH,DEV=xxxx
```

where xxxx is the device address. An example of a result from z/OS is:

```
RESPONSE=ZOS1
IOS086I 16.57.02 MIH DEVICE TIMES 199
xxxx=03:00
```

In this example, the MIH for this device address is 3 minutes. The MIH can also be configured on the MDI system on a per control unit basis.

7.9 Device Allocation Steering (z/OS)

Steering tape mounts to the correct devices may require the installation of additional software if adequate tools are not already available within your environment. Although an "esoteric" device group name can be used to steer allocations for DISP=NEW data sets, this method usually does not work properly when reading from cataloged tape data sets. The reason for this is that the MVS catalog records the device type of the tape drive and not the esoteric name. When data set allocation attempts to find the data set, it requests the tape VOLSER on any drive of the model type specified in the MVS catalog. If you have more than one tape system of a particular model (3490 as an example) and the tape media cannot be mounted on any drive of that type, then retrievals for existing tape data sets can cause mount requests to the wrong tape system.

Therefore, an enhanced method of steering allocations is needed whenever there are independent tape systems whose tape data sets are (MVS) cataloged as identical tape models.

Some of the methods that can be used are:

- Luminex TMACS software;
- DFSMS-managed tape (groups tape drives as an MTL in the IOGEN);
- Oracle's NCS/SMC 6.2 or ELS/SMC software (TAPEREQ);
- Computer Associates MIM/MIA software product (device reserves);
- explicit UNIT= for every reference to a tape device, including cataloged tape data sets;
- or be the only tape devices of that type not controlled by any of the above.

The method you choose depends on what already exists in the environment. Determining a viable method should be discussed with the systems programmers and Luminex prior to the installation.

7.10 Mainframe Operating System

MDI supports all versions of z/OS mainframe operating systems.

7.11 Scratch Tapes

For tape management systems that use the non-specific method of requesting scratch tape mounts, MDI needs a list of scratch VOLSERs from which it can select a VOLSER and mount that tape in response to MSCRTCH, MPRIVAT, or other types of scratch pool mount load display messages. MDI removes the mounted scratch VOLSER from the scratch list if the tape is written upon.

The scratch list should be repopulated as VOLSERs expire and return to scratch status. A "scratch update" VOLSER is defined in MDI for receiving updates to the scratch list. The mainframe just has to write a simple file (LRECL=6, one scratch VOLSER per record) to this special tape VOLSER.

The best way to do this is to extract the list of scratch VOLSERs from the appropriate tape management system scratch tape report. Luminex can provide a program (LSCRUP) that extracts this list from the scratch reports from most z/OS tape management systems. Once this list is extracted, a subsequent step can write the list to the special "scratch update" VOLSER.

When MDI updates its scratch list, it can either (A) merge the new list with the current scratch list, sort, and remove duplicates, or (B) completely replace the current scratch list with the new list, and sort it. The A or B method must be specified in the MDI configuration.

The data on a scratch tape is not released (purged) until the tape is reused (rewritten). Since the scratch list is always resorted, the lower VOLSERs are reused first.

Each group of tape devices can be assigned a separate scratch list (and corresponding special scratch update VOLSER). This provides support for multiple tape management (scratch) pools. If each scratch pool has a corresponding non-specific load display message (such as MMEDIA3 or MMEDIA4), MDI can be configured to recognize these and mount a scratch tape for the corresponding scratch list.

Typically, there is a scratch list for each independent VOLSER range. Each device type family (3490 versus 3590) usually has its own VOLSER range (and scratch list). Due to the difference in block ID format between 3480/3490 versus 3590 tape drives, VOLSER ranges are not shared across these device type families.

If MTL=YES has been specified, this means that the IBM Tape Catalog Data Base (TCDB) is also required to satisfy OAM lookups. The TCDB is also known as the VOLCAT. It is a special catalog that only contains IDCAMS VOLUMEENTRIES and other tape library definition values. When the tape management system changes a VOLSER to scratch status, it also has to change the VOLSER in the TCDB to scratch status.

Problems with VOLSER information in the TCDB often result in error messages similar to the following.

```
IEC518I SOFTWARE ERRSTAT: MLNOTVR  D02F,,SL,PRODJOB1,SYSUT2  
IEC502E R D02F,555001,SL,SYSUT2-code,PRODJOB1,STEP3
```

IEC502E code 9 means the scratch VOLSER is not defined in the TCDB.

IEC502E code 10 means the VOLSER is defined in the TCDB but not in this library.

IEC502E code 11 means the VOLSER is defined in the TCDB but is not in scratch status.

8. Luminex Mainframe Software Installation

Luminex provides optional software packages that can be installed on the mainframe to gain additional functionality from MDI.

8.1 ADMIN+

The mainframe feature, ADMIN+ for MVS, provides a means to perform functions on MDI and to gather MDI system information for use and display on the mainframe. Some features are: determining storage capacity, viewing system health, scratch pool manipulation, and more. Contact Luminex for more information on the use and installation of the ADMIN+ feature.

8.2 LumAlert

MDI's Alert Notification service sends alerts messages via email and/or snmp trap when triggered by a software or hardware event. An alert message will be sent when a condition is first detected and again daily until the condition is resolved. The conditions monitored are MDI hardware and software failures.

There may need to be authorization from the SMTP server to permit MDI to send email, particularly, external email to Luminex Software.

In order to configure the LumAlert package, the following items need to be provided:

- SMTP server DNS name or IP address
- Local domain name
- Recipient email address(es)

Appendix A. Customer Questionnaire

The following questions, and answers, will give the Luminex Sales Engineers an idea on how the customer intends to use MDI and the infrastructure involved in the installation. This is not intended to be an exhaustive list of questions but the questions are to be used to provoke more discussion. Use the table at the end of the section for the answers.

A.1. Equipment Infrastructure

Q: Is the customer providing rack space for the MDI systems?

The customer is usually responsible for providing rack space and electrical power connections. Each MDI is a 2U, 19" rack mounted system.

Q: What type of power outlets are in the rack?

Does the (rack) power outlet require a NEMA 5-15P plug (typical 3-prong outlet), or the IEC60320 C14 plug (shrouded 3-prong).

A.2. Networking Infrastructure

Q: Will the system(s) be available on the customer network?

On-site configuration is generally done over the network. Remote access usually requires that the MDI's main interface be on the customer network. The customer should be prepared to provide IP addresses, gateway IP address, domain name, and email server information.

Q: What is the means for remote access for Luminex Support:

The most common methods are:

- VPN
- WebEx
- Modem

A.3. Regarding Storage

Q: Is the storage to be shared with open systems?

Q: Is all data to be replicated?

If some data is replicated and some not, the mainframe configuration must accommodate different devices with different policies.

A.4. Regarding Device Allocation Steering

Q: What method will the customer be using:

- MTL configuration
- If customer currently is using ATLs (IBM VTS for example), then MTL is the best choice
- Esoteric
- STK HSC/SMC or ELS allocation software
- Other

A.5. Tape Management System

Q: What is the tape management system?

The most common are RMM, CA-1, CA-TLMS.

Q: How many tape management catalogs are to be supported?

MDI will have to be configured specifically to handle multiple tape management systems.

A.6. Mainframe(s) Configuration

Q: What is the OS? (z/OS is required)

Q: How many mainframes?

Q: How many LPARs will access the system?

Q: Are the FICON channels routed through switches/directors?

If so, what is the manufacturer?

Q: How many total devices are expected to be needed?

A.7. Performance

Q: What is the expected requirement for read/write performance?

SMF data may be required to determine peak read/write rates.

A.8. Applications

Q: What applications will be used?

Q: Is FDR Upstream being used?

Support for FDR Upstream with some deduplication storage may require additional configuration to optimize deduplication performance.

A.9. Facilities

Q: Is the Data Center managed by the customer or by third-party management?

Q: How often are systems programmers able to IPL or activate new IOGENs?

Q: Are there any change window restrictions?

For both the installation and HCD change.

A.10. Data Migration

Q: What is the intended method to migrate data?

- Allow old data to expire?

- Self-migrate?
- Use Luminex Migration service?
- If so, are there encrypted tapes that need to be migrated?
- Write only new data to MDI and continue use of old devices?

A.11. Installation Support

Q: Will a systems programmer be available?

Q: Is there an installation time window?

Q: Are there test jobs available?

Most tests will require that the tape management system be configured with the scratch tapes that will be used.

Q: Are their special security clearances required?

Q: Can PC laptops and cell phones be brought into the data center?

A.12. Check-Off List

Please provide as much information as possible to the Luminex installation team.

EQUIPMENT INFRASTRUCTURE Is the user providing rack space? What power plug is required?	YES NO 5-15P / C14 / Other: _____
NETWORKING INFRASTRUCTURE Will the equipment be on the customer network? What is the method for remote access?	YES NO VPN / Webex / Other
STORAGE Is the storage to be shared with other systems? Is the data to be replicated?	YES NO YES NO
DEVICE ALLOCATION STEERING What method will be used?	MTL / Esoteric / Other
TAPE MANAGEMENT SYSTEM What is the Tape Management System? How many Tape Management Catalogs?	RMM / CA-1 /CA-TLMS / BIM-EPIC / Other _____
MAINFRAME CONFIGURATION What is the OS? (z/OS is required) How many mainframes: How many LPARs Are there any directors or switches? How many total devices expected?	z/OS / Other _____ _____ YES NO _____
PERFORMANCE What are the expected read/write rate requirements?	_____

<p>APPLICATIONS</p> <p>List the applications to be used:</p> <p>_____</p> <p>_____</p> <p>_____</p> <p>_____</p> <p>_____</p> <p>Is FDR Upstream being used?</p>	<p>_____</p> <p>_____</p> <p>_____</p> <p>_____</p> <p>_____</p> <p>YES NO</p>
<p>FACILITIES</p> <p>Is the data center managed by a third-party?</p> <p>How often are the systems programmers able to do IPLs or activate new IOGENs?</p> <p>Are there any change window restrictions?</p>	<p>YES NO</p> <p>_____</p> <p>_____</p>
<p>DATA MIGRATION</p> <p>Do you intend to migrate data?</p> <p>If doing migration, how do you intend to do it?</p> <p>Self Migrate</p> <p>Allow old data to expire</p> <p>Use Luminex Tape Migration Services</p>	<p>YES NO</p> <p>YES NO</p> <p>YES NO</p> <p>YES NO</p>
<p>INSTALLATION SUPPORT</p> <p>Will a systems programmer be available?</p> <p>Is there an installation time window?</p> <p>Are test jobs available?</p> <p>Are their security protocols for the installer?</p> <p>Can PC laptops and cell phones be used in the computer room?</p>	<p>YES NO</p> <p>YES NO</p> <p>YES NO</p> <p>YES NO</p> <p>YES NO</p>

OTHER INFORMATION

Note all other pertinent information here:

Appendix B. Final Installation Check Off

The following items are required prior to installation. Having this information before the installation will ensure a timely and successful install.

B.1. IP Addresses

If any part of the solution will be on a customer maintained network, the following information will be required.

Port Description	IP	Gateway	Netmask
Management Port – This port will be used for both remote access and access to the Luminex browser interface.			
Service Port – This port provides low-level access to MDI.			
Network Storage Port – This port will connect to network storage, if any.			

B.2. Remote Access

If remote access to MDI is to be through phone dial-up, a phone number will be required. If the remote access is to be through VPN, specific information regarding the software and access method will be required.

B.2.1. Dial-up Method

Phone Number _____

B.2.2. VPN Method

Item	
If required, has client software been provided by customer?	YES NO
Security Clearance for Remote Access provided by customer?	YES NO
VPN Access instructions provided by customer?	YES NO

B.3. Mainframe Configuration

The IOGEN information should be verified by Luminex prior to the installation.

Item	
Has the IOGEN been provided by the customer?	YES NO
Has the IOGEN been confirmed by Luminex?	YES NO

B.4. MDI Implementation

The successful implementation of MDI requires that the user understand some important concepts and be prepared to test these soon after the installation

Item	
Does the customer understand the implementation of the scratching process?	YES NO
Does the customer understand their methods of device allocation steering?	YES NO

B.5. Final Infrastructure

The following items are needed to complete the installation in a timely manner and confirm the success of the installation.

Item	
Are all necessary network cables in place?	YES NO
Are all ESCON/FICON cables in place?	YES NO
Is the customer ready to run test jobs?	YES NO
Is the tape management system configured for the new devices?	YES NO



871 Marlborough Avenue
Riverside, CA 92507

1.888.LUMINEX
1.951.781.4100

www.luminex.com