IBM® Storage

Using RDP with IBM FlashSystem to Debug Fibre Channel Optics Errors

IBM Storage Team

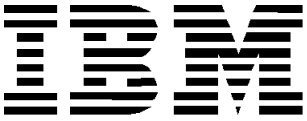

© Copyright International Business Machines Corporation 2022.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

© Copyright IBM Corp. 2022.

About this document

The focus of this IBM® blueprint is to showcase the Read Diagnostic Parameters (RDP) feature of the Fibre Channel protocol (FCP).

The data that is provided by RDP commands can simplify the process of managing and analyzing any issues on complex SAN fabrics. In this blueprint, we provide guidance to help users and administrators understand the meaning of RDP data and how to use it.

The intent of this blueprint is to help a user understand what RDP is, what data RDP represents, and how to use that data to identify potential issues within the SAN fabric that is hosted by that Fibre Channel (FC) switch.

© Copyright IBM Corp. 2022.

Executive summary

RDP is an FC primitive, which enables you to use the FC switch command-line interface (CLI) to monitor and debug the small form factor pluggable (SFP) optical modules that are installed on endpoint devices.

In an FC storage area network (SAN), bad optics and cabling tend to be one of the most common point of failures. This situation can result from SFP electronics wearing out over time, cables malfunctioning due to physical stress, or other issues. These component-facing issues might manifest as performance drops, loss of paths, congestion, and other issues.

In Fabric or FC SANs, it can be difficult to identify the exact item that causes such issues. An average compute node can have around 2 - 3 FC adapters, with up to four ports each. Modern SAN storage devices might have up to four quad-port FC adapters per node, which means that there are many physical ports in any SAN.

The IBM FlashSystem® family of products incorporate support to respond to these commands from version 8.5.2.0 and later.

Scope

The focus of this blueprint is to enable administrators and users to use certain RDP commands from the FC switches. It also describes what outputs to expect and the meanings of the data that is produced by the commands.

This blueprint provides guidance to help users or administrators foresee any upcoming failures based on the data that is returned by RDP commands.

Users must have working knowledge of the following items:

- Basic CLI usage
- Connecting to FC switches over a network

Introduction

RDP is an Extended Link Service (ELS) command that is implemented within the FC protocol standards.

When an endpoint receives an RDP command from another endpoint (typically a switch), it generates a response with all its internal SFP health and port error information populated into it. RDP has been part of the FC protocol for nearly a decade now, so it has been implemented by most popular FC switch vendors.

These parameters can be viewed on the switch and used to monitor or debug interfaces that are connected to it. A key aspect of RDP is that it is non-intrusive to the port itself. It is treated like a regular ELS by the firmware and driver stack, and therefore poses no risk to regular operation.

Figure 1 shows how the RDP command flow works.

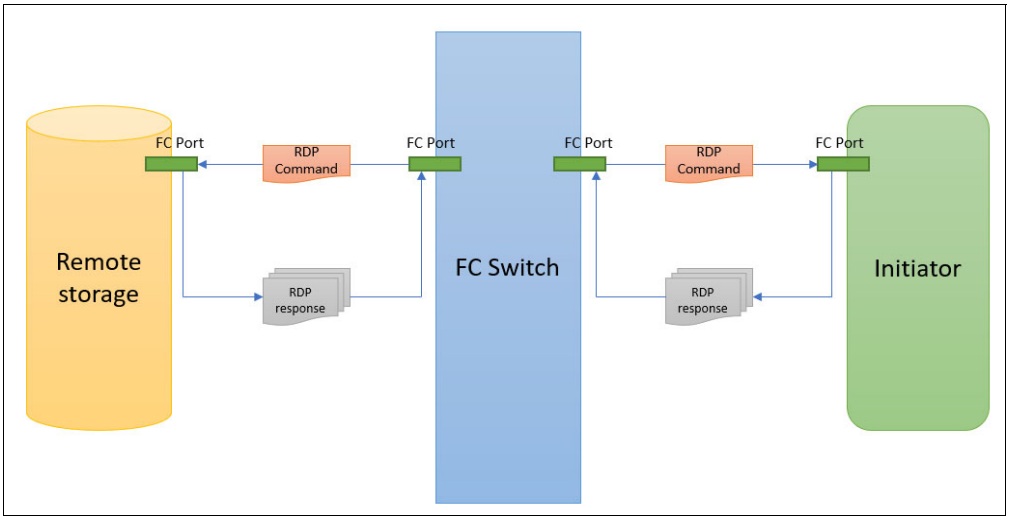

Figure 1 Basic RDP flow on a Fibre Channel SAN

IBM FlashSystem and IBM Spectrum® Virtualize products have enabled support for responding to RDP commands with version 8.5.2.0.

In this blueprint, we describe using RDP with two of the most popular FC switch vendors: Brocade and Cisco.

Prerequisites

This section describes the following prerequisites to use RDP from Cisco or Brocade FC switches with IBM FlashSystem, IBM Storwize®, or IBM SAN Volume Controller products:

- Ensure that you are running IBM Spectrum Virtualize 8.5.2.0 or later.

  Note: All the commands that are mentioned in this blueprint are non-destructive. They do not affect or modify the switch configuration in any way.

- Ensure that the FC switch version is Cisco NXOS 8.x or later or Brocade Fabric OS (FOS) 7.4 or later.

How to use RDP

This section describes the following topics:

- Brocade Fabric OS
- Cisco NXOS

Brocade Fabric OS

Brocade FOS delivers information that is collected through RDP by two commands, with the relevant fields under each command:

- The **sfpshow** command with the **-link** option. The syntax is as follows:

  `sfpshow <port_num> -link`

  This command internally sends an RDP command to the port in question, which generates the output that is shown in Figure 2.

```
IBM_8960_F64:FID128:admin> sfpshow 40 -link
Identifier:  3    SFP
Connector:   7    LC
Transceiver: 6804406000000000 8,16,32_Gbps M5 sw Inter,Short_dist
Encoding:    6    64B66B
Baud Rate:   280  (units 100 megabaud)
Length 9u:   0    (units km)
Length 9u:   0    (units 100 meters)
Length 50u (OM2):  2    (units 10 meters)
Length 50u (OM3):  7    (units 10 meters)
Length 62.5u:0    (units 10 meters)
Length 50u (OM4):  10   (units 10 meters)
Vendor Name: BROCADE
Vendor OUI:  00:05:1e
Vendor PN:   57-1000333-01
Vendor Rev:  A
Wavelength:  850  (units nm)
Options:     083a Loss_of_Sig,Tx_Fault,Tx_Disable
BR Max:      112
BR Min:      0
Serial No:   JAA718463077041
Date Code:   181117
DD Type:     0x68
Enh Options: 0xfa
Status/Ctrl: 0x0
Pwr On Time: 3.32 years (29084 hours)
E-Wrap Control: 0
O-Wrap Control: 0
Alarm flags[0,1] = 0x0, 0x0
Warn Flags[0,1] = 0x0, 0x0A
Temperature: 52       Centigrade
Current:     7.496    mAmps
Voltage:     3318.1   mVolts
RX Power:    -0.0     dBm (996.3uW)
TX Power:    -1.1     dBm (784.9 uW)

State transitions: 11
Port Speed Capabilities     4Gbps 8Gbps 16Gbps   32Gbps

PEER Port Gbic Info

Vendor Name: AVAGO
Serial num:  AD2109G0GKJ
Vendor PN:   AFBR-57G5MZ-ELX
Vendor Rev:
Date Code:   210319
        Laser Type:      Short Wave Laser
        SFP Type:        Optical Port Type
        Connecter Type: Other
Following SFP Parameters are Valid
        Temperature: 39       Centigrade [Range -128 - +128 C]
        Current:     7.496    mAmps     [Range 0 - 131 mAmps]
        Voltage:     3255.6   mVolts    [Range 0 - 3600 mVolts]
        Rx Power:    713.2    uW       [Range 0 - 6550 uW]
        Tx Power:    1171.9   uW       [Range 0 - 6550 uW]
Signal Loss (Upstream)  :  -11.4   dBm (71.7  uW)
Signal Loss (Downstream):  -7.6    dBm (175.6 uW)
Port Speed Capabilities       8Gbps 16Gbps   32Gbps
Last poll time: 09-30-2022 UTC Fri 13:05:30
```

Figure 2 Brocade FOS sfpshow command output

- The **portshow** command with the **-link** option. The syntax is as follows:

  `portshow <port number> -link`

  Figure 3 shows the command output.

```
IBM_8960_F64:FID128:admin> portshow 40 -link
portIndex:  40
portName: port40
portHealth: HEALTHY

Authentication: None
portDisableReason: None
portCFlags: 0x1
portFlags: 0x24b03       PRESENT ACTIVE F_PORT G_PORT U_PORT NPIV LOGICAL_ONLINE LOGIN NOELP LED ACCEPT FLOGI
LocalSwcFlags: 0x0
portType:  26.0
POD Port: Port is licensed
portState: 1    Online
Protocol: FC
portPhys:  6    In_Sync         portScn:   32   F_Port
port generation number:    51948
state transition count:    2085

portId:    012800
portIfId:    43020016
portWwn:   20:28:88:94:71:8c:7c:30
portWwn of device(s) connected:
        50:05:07:68:10:43:9b:0f
        50:05:07:68:10:83:9b:0f
        50:05:07:68:10:03:9b:0f
16b Area list:
Distance:  normal
portSpeed: N32Gbps

FEC: Active
Credit Recovery: Active
Aoq: Inactive
FAA: Inactive
F_Trunk: Inactive
LE domain: 0
Peer beacon: Off
FC Fastwrite: OFF
Interrupts:        9474      Link_failure: 107       Frjt:         0
Unknown:           2104      Loss_of_sync: 9         Fbsy:         0
Lli:               9474      Loss_of_sig:  957
Proc_rqrd:         130823    Protocol_err: 0
Timed_out:         0         Invalid_word: 3
Tx_unavail:        0         Invalid_crc:  0
Delim_err:         0         Address_err:  0
Lr_in:             1178      Ols_in:       107
Lr_out:            107       Ols_out:      201
nodeWwn: 10:00:88:94:71:8c:7c:30
Pn Phy Port Type: FC-FS-3 PN_Port or PF_Port
FEC Uncorrected Blocks: 12
Advertised Buffer Credit       :       20
Advertised Peer Buffer Credit  :       40

PEER PORT SHOW

portWwn: 50:05:07:68:10:43:9b:0f
nodeWwn: 50:05:07:68:10:00:9b:0f
Pn Phy Port Type: FC-FS-3 PN_Port or PF_Port
Link_failure: 4   Loss of Sync 6   Loss of Sig 0
Protocol Error: 0   Invalid Word 34   Invalid CRC 0
FEC Uncorrected Blocks: 0
Advertised Buffer Credit       :       40
Advertised Peer Buffer Credit  :       20
```

Figure 3 Brocade FOS portshow command output

Cisco NXOS

RDP can be triggered on Cisco FC switches by running the following command:

`show rdp fcid <fcid(nport id) vsan <vsan>`

Figure 4 shows the output that is generated by the command.

```
show rdp fcid 0x660022 vsan 1
---------------------------------------------------------------
                RDP frame details
---------------------------------------------------------------
Link Service Request Info:
------------------------------

Port Speed Descriptor Info:
------------------------------
Port speed capabilities : 32G 16G 8G
Port Oper speed         : 32000 Mbps

Link Error Status:
------------------------------
VN PHY port type               : FC
Link failure count             : 0
Loss of sync count             : 0
Loss of signal count           : 0
Primitive sequence proto error : 0
Invalid Transmission word      : 0
Invalid CRC count              : 0

Port Name Descriptor:
------------------------------
Node WWN            : 50:05:07:68:0b:00:ac:4a
Port WWN            : 50:05:07:68:0b:26:ac:4a
Attached Node WWN   : 20:01:00:de:fb:da:4c:21
Attached Port WWN   : 20:03:00:de:fb:da:4c:20

SFP Diag params:
------------------------------
SFP flags       : Optical
SFP Tx Type     : Short Wave

FEC Status:
------------------------------
Corrected blocks   : 0
Uncorrected blocks : 0

Buffer Credit Descriptor:
------------------------------
Rx B2B credit   : 186
Tx B2B credit   : 32
Port RTT        : 0 ns

Optical Product Data:
------------------------------
Vendor Name    : FINISAR CORP.
Model No.      : FTLF8532P4BNV-QL
Serial No.     : P2GA3CN
Revision       : A
Date           : Q

  ------------------------------------------------------------------------------
                Current              Alarms                 Warnings
                Measurement     High        Low         High          Low
  ------------------------------------------------------------------------------
  Temperature   47.25 C        90.00 C     -5.00 C     85.00 C        0.00 C
  Voltage        3.36 V         3.63 V      2.97 V      3.46 V        3.13 V
  Current        8.01 mA       12.00 mA     1.00 mA    11.50 mA       2.00 mA
  Tx Power      -1.46 dBm       5.00 dBm  -12.20 dBm    2.00 dBm     -8.20 dBm
  Rx Power      -1.14 dBm       5.00 dBm  -15.20 dBm    2.00 dBm    -11.20 dBm
  ------------------------------------------------------------------------------
  Note: ++  high-alarm; +  high-warning; --  low-alarm; -  low-warning
```

Figure 4 Cisco NXOS show rdp command output

Debug guide

All the parameters that are displayed in the RDP output have two associated thresholds: warning and alarm. Under each threshold, there are two more levels: high level and low level.

For each of these combinations, first a warning event, and then an alarm event can be generated at the respective endpoint (for example, temperature-low-warning, current-high-alarm, and so on).

With the RDP data, users can correlate the current measurement against the respective threshold values to better their understanding about the status of optics at the remote endpoint, and foresee any potential issues.

Temperature

The current temperature of the SFP. The limits of an SFP are decided by its category. The category is written on the SFP label as “COM”, “EXT”, or “IND”. The safe operating temperature ranges for them are as follows:

- COM: Commercial => 0° - 70 °C
- EXT: Extended => -20° - 85 °C
- IND: Industrial => -40° - 85 °C

SFP temperatures are typically caused by manufacturing or material quality, ambient temperature, and cooling efficiency.

High temperature

Causes a spike in optical power. A higher temperature can lead to signal errors or burn out the optical module, which results in a failed SFP and a loss of the link.

Low temperature

A low temperature can result in a module becoming unstable, which might lead to signal loss.

A temperature that is read from an RDP command can be used to monitor the state of the SFP and act as an early warning if the temperature is too high or low. You might need to adjust the ambient temperature.

Current

The amount of current that the SFP is consuming. Faulty SFPs have a value of “0”.

A higher current might be associated with higher transmit (Tx) power, which might lead to an optical element failure.

A lower current might lead to signal loss and an unreliable link.

Voltage

The amount of voltage that the SFP is consuming. A faulty SFP has a value of “0”.

Higher voltage levels might be associated with higher Tx power, which might lead to an optical element failure and a link loss.

Lower voltage levels might cause signal loss and an unreliable link.

Transmit power

Tx is the amount of power that is used in the transmit side of the SFP. It can be shown as absolute power (microwatts (uW)) or a ratio (decibels (dBm)).

Constant high Tx power might result in bit errors on the receiving side.

Low Tx power might indicate an unreliable or failing or failed SFP.

Receive power

Receive power (Rx) is the amount of power that is received on the transmit side of the SFP. It can be shown as absolute power (microWatts (uW)) or a ratio (dBm).

Constant high Rx power might result in bit errors codes due to the photocurrent saturation phenomenon, which is where a photodiode sensor receives more energy than it can accurately convert into an electrical signal. This situation might lead to bit errors or signal loss.

Low Rx power might be caused by a faulty SFP on the sending side, which might manifest as bit errors. You might need to replace the remote-side SFP.

Signal loss (Brocade only)

A measure of the attenuation loss on the FC link. This situation might be caused by any of the above issues, or attenuation due to FC loss. Brocade represents it in dBm.

Higher numbers indicate a bad SFP or cable, which results in loss-of-signal errors.

Link error status

Both outputs (Figure 3 on page 6 and Figure 4 on page 7) also show the link error statistics for the specified port.

Link failure

The number of times that the FC link has gone down. This error might be a part of routine maintenance, such as when the peer port is replaced or serviced, but it also might indicate a failing or failed SFP or cable.

Loss of sync

Indicates a synchronization loss on the FC link. Similar to a link failure, it can be intended or erroneous. The SFP or cable to be evaluated.

Loss of signal

Indicates that a signal was lost on the FC link. This error can be a problem if it is not noticed immediately after any planned maintenance activity.

Primitive sequence protocol error

The number of times an error in the primitive sequence is detected. Primitive sequences are transmitted repeatedly on all active links. This error might indicate an issue with the receiver port.

Invalid transmission word

This count relates to the number of times that a bit-level error was detected in a frame. Higher counts can indicate a problematic link.

Cyclic Redundancy Check error

This count indicates the number of frames that is received that have a Cyclic Redundancy Check (CRC) error in them. This situation can occur in cases where a frame was re-transmitted. This error does not directly imply a bad link, but the link should be monitored for activity.

Summary

This blueprint explains the usage of FC switch-based RDP commands to monitor and analyze port SFP metrics, and how to use that data to predict certain known errors or failures.

Authors

This blueprint guide was produced by developers from the FC driver team working at IBM Systems Labs, Pune, India.

Ashish Jagdale is a software engineer at IBM Systems Labs, Pune, India. He has more than 7 years of experience working for the FC driver team on IBM SVC and IBM FlashSystem products. He holds a bachelors degree in electronics engineering from Shivaji University.

Mohit M Chitlange is a software engineer at IBM Systems Labs, Pune, India. He has more than 7 years of experience working for the FC driver team on IBM SVC and IBM FlashSystem products. He holds a masters degree in engineering from Pune University.

The authors thank the following contributor for their design and architectural support:

Rahul Fiske
Fibre Channel architect, ISDL Pune, IBM India

Resources

The following websites provide useful references to supplement the information that is contained in this blueprint:

- IBM FlashSystem 9200

  https://www.ibm.com/mysupport/s/topic/0TO0z000000ZTOiGAG/flashsystem-9200
- IBM FlashSystem 9500

  https://www.ibm.com/in-en/products/flashsystem-9500
- *Cisco MDS 9000 Series Command Reference, Release 8.x*

  https://www.cisco.com/c/en/us/td/docs/switches/datacenter/mds9000/sw/8_x/command/cisco_mds9000_command_ref_8x.html
- *Brocade Fabric OS Command Reference Manual, 8.2.x*

  https://docs.broadcom.com/doc/FOS-82x-Command-RM

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, MD-NC119, Armonk, NY 10504-1785, US

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

© Copyright IBM Corp. 2022.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

IBM®	IBM Spectrum®	Storwize®
IBM FlashSystem®	Redbooks (logo) ®	

The following terms are trademarks of other companies:

Other company, product, or service names may be trademarks or service marks of others.

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Privacy policy considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at http://www.ibm.com/privacy and IBM's Online Privacy Statement at http://www.ibm.com/privacy/details in the section entitled "Cookies, Web Beacons and Other Technologies," and the "IBM Software Products and Software-as-a-Service Privacy Statement" at http://www.ibm.com/software/info/product-privacy.

© Copyright IBM Corporation

December 2022

US Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Please recycle

ISBN 0738460982

REDP-5699-00