

Compaq SANworks™

Application Note Data Replication Manager over an ATM Link Supporting OSG Up To and Including Version 2.2.1

Visit Our Website for the Latest Information

At Compaq, we are continually making additions to our storage solution product line. Please check our website for more information on our Fibre Channel product line as well as the latest drivers, technical tips, and other documentation. Visit our website at:

<http://www.compaq.com/storageworks>

© 2000 Compaq Computer Corporation.

Compaq, the Compaq logo, and StorageWorks Registered in U. S. Patent and Trademark Office.

SANworks, Tru64 UNIX, and OpenVMS are trademarks of Compaq Information Technologies Group, L.P. in the United States and other countries.

Microsoft, MS-DOS, Windows, Windows NT are trademarks of Microsoft Corporation in the United States and other countries.

Intel, Pentium, Intel Inside, and Celeron are trademarks of Intel Corporation in the United States and other countries.

Motif, OSF/1, UNIX, the "X" device, IT DialTone, and The Open Group are trademarks of The Open Group in the United States and other countries.

All other product names mentioned herein may be trademarks of their respective companies.

Confidential computer software. Valid license from Compaq required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Compaq shall not be liable for technical or editorial errors or omissions contained herein. The information in this document is provided "as is" without warranty of any kind and is subject to change without notice. The warranties for Compaq products are set forth in the express limited warranty statements accompanying such products. Nothing herein should be construed as constituting an additional warranty.

Compaq service tool software, including associated documentation, is the property of and contains confidential technology of Compaq Computer Corporation. Service customer is hereby licensed to use the software only for activities directly relating to the delivery of, and only during the term of, the applicable services delivered by Compaq or its authorized service provider. Customer may not modify or reverse engineer, remove, or transfer the software or make the software or any resultant diagnosis or system management data available to other parties without Compaq's or its authorized service provider's consent. Upon termination of the services, customer will, at Compaq's or its service provider's option, destroy or return the software and associated documentation in its possession.

Printed in the U.S.A.

Data Replication Manager over an ATM Link

Application Note

Fourth Edition (November 2000)

Part Number: EK-DRMAL-AA. D01

About This Document

This Application Note provides an overview of the ATM (Asynchronous Transfer Mode) link and its effect on performance when used with Data Replication Manager (DRM). The DRM is a storage-based data replication and workload migration solution that copies data online and in real time to remote locations via an extended Storage Area Network (SAN). The ATM link has not been qualified for configurations other than Data Replication Manager. This document supports these configurations:

- Data Replication Manager configurations of up to two storage systems with up to four servers (four at the local site and four at the remote site).
- Open System Gateway revision level 2.2.1 and earlier.
- Fibre Channel switch version 2.1.7 or higher

For complete details on Data Replication Manager, refer to the *Compaq SANworks Data Replication Manager HSG80 ACS Version 8.5P Operations Guide* for the appropriate platform.

This Application Note contains the following sections:

- Introduction to DRM with ATM Transport
 - Fibre-Channel-to-ATM Gateway
 - Open Systems Gateway Interface
- ATM Link Topologies
 - Dual Links
 - Shared Virtual Channel Links
 - Fractional Virtual Channel Links
 - ATM Inverse Multiplexing
 - Encryption
- Case Study
 - Example Criteria
 - DRM Requirements
 - Obtaining Best Performance
- Installation and Configuration
 - Fibre Channel Switch Setup
 - Service Kit
 - Open Systems Gateway Setup
 - Site Readiness Checklist
 - Modifying VPI, VCI, QoS, and PCR Values
 - Modifying ActiveVciBits Values
- Websites
 - Compaq
 - ATM Forum
- Disclaimer
- Glossary

Introduction to Data Replication Manager with ATM Transport

During normal data processing, data is written to initiator (local) and target (remote) sites. While copies of data reside at both sites, host data access occurs through the initiator site, unless there is a failure or catastrophe that disables processing at that site. In the event of an initiator failure, another site can continue processing data in the interim.

Data Replication Manager provides rapid data access recovery and continued data processing after the loss of one or more components. Data Replication Manager uses the peer-to-peer remote copy function of the HSG80 controller to achieve data replication. HSG80 controller pairs at the initiator site are connected to their partner HSG80 controller pairs at the target site.

Data Replication Manager can replicate data up to 100 kilometers (approximately 62 miles) via an extended storage area network over direct Fibre Channel links or extend over greater distances with Fibre-Channel-to-ATM Gateways.

Fibre-Channel-to-ATM Gateway

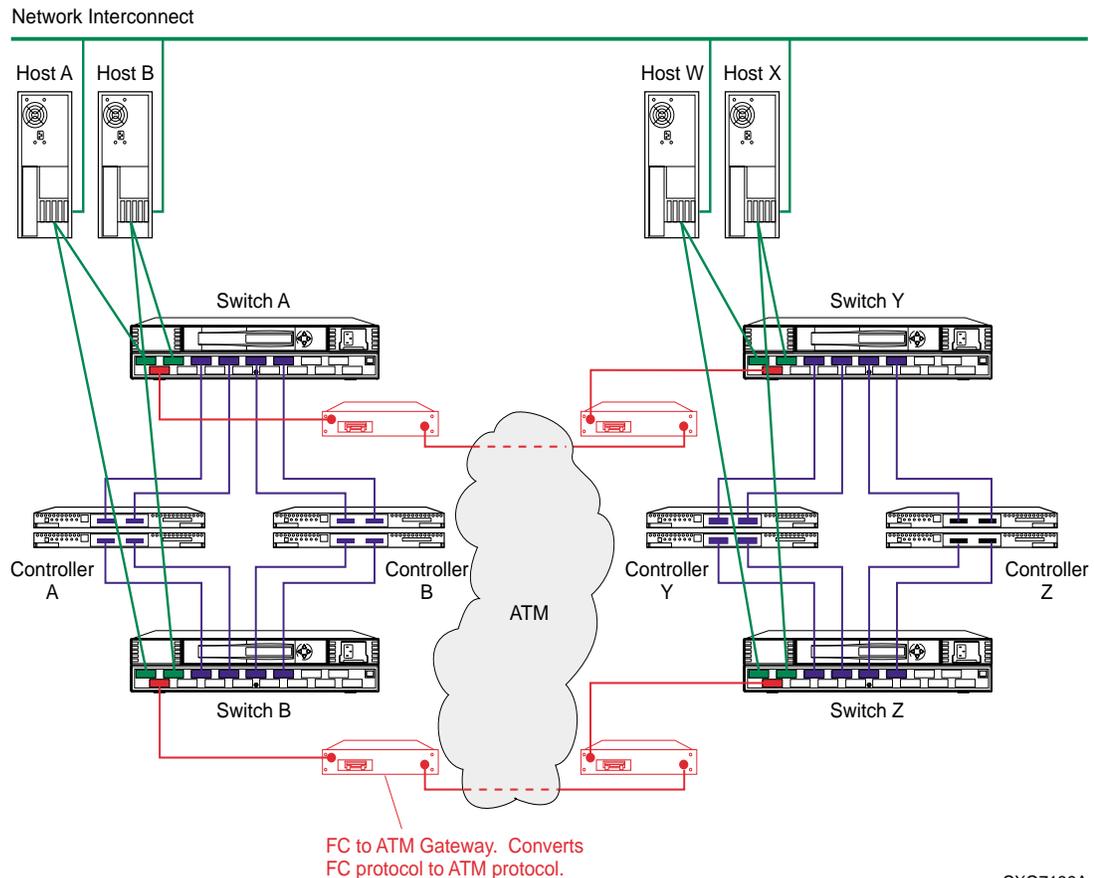
Fibre-Channel-to-ATM gateways, referred to as Open Systems Gateways (OSGs), provide ATM connectivity into telecommunication networks for the *Compaq SANworks* extended SAN/DRM solution. ATM connectivity is ideal, where contingency sites are required to be very long distances away, to protect against wide area disasters, or where dark fiber connectivity is unavailable.

ATM is a communications networking technology that carries information (voice, video, and data) in 53-byte cells. The fixed-length cell allows a network to carry any type of information within the cell and also to provide stringent service qualities that can differ by application.

Each virtual connection in an ATM network has a service category. The performance of the connection is measured by the established *Quality of Service* (QoS) parameters. *Constant Bit Rate* (CBR) is the service category used by the current DRM solution over ATM. The only requirement for CBR is *Peak Cell Rate* (PCR). This value is determined by performance requirements.

The DRM solution uses virtual connections in an ATM network that are predefined and permanently left in place. These connections are referred to as a *Permanent Virtual Circuit* (PVC). The 53-byte cell contains a 5-byte header, which contains all the information to transport the cell from the source to the destination. Every cell header contains an address identifier used by the ATM switches to track how cells move from one physical link to the next. Before any communication exchanges take place, the connectivity information stored in the switches must be loaded manually. The network administrator must select the end points that require connectivity to set up PVCs. These address identifiers are known as VPIs and VCIs (*Virtual Path Identifiers* and *Virtual Channel Identifiers*).

DRM requires two ATM links, each running at a Public Line Rate between 1.544 megabits per second (193 kilobytes per second) and 144 megabits per second (18.06 megabytes per second). See Figure 1.



CXO7139A

Figure 1 Data Replication Manager configured over an ATM link

Open Systems Gateway Interface

The Open Systems Gateway (OSG) interface supports:

- 50-micron multi-mode fiber, up to 500 meters (Fibre Channel and ATM sides)
- 62.5-micron multi-mode fiber, up to 200 meters (Fibre Channel and ATM sides)
- Performance target data rate from 1.554 megabits per second to 144 megabits per second over an ATM OC-3 link
- Point-to-point connectivity only (no alternate routing supported at present)
- Fibre Channel Class 3 traffic
- Frame data field size of 2048 bytes

The OSG interface does not support Class F frames.

ATM Link Topologies

Dual Links

The most disaster tolerant of the ATM Link configurations, dual links provide protection if one link becomes unavailable. No ATM switch is required. Each link has diverse routing to increase survivability. Figure 2 shows a dual-link configuration.

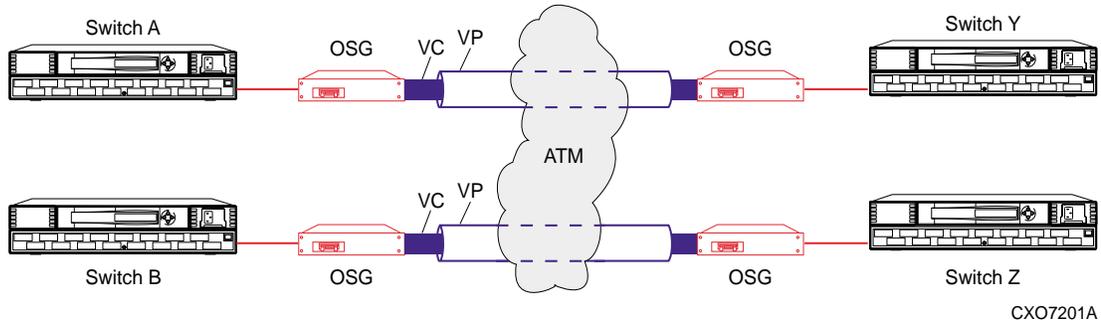


Figure 2 Dual-link configuration

CXO7201A

Shared Links

A shared link, like that shown in Figure 3, allows for one telecommunication link between sites, with the line bandwidth divided between the two Fibre Channel fabrics. The data rate is shared according to the Peak Cell Rate allocation. Because there is only one link, this configuration has less survivability than dual links.

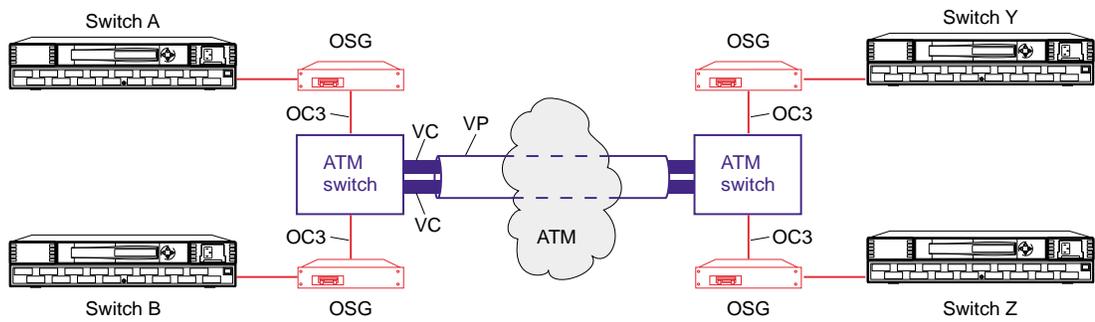
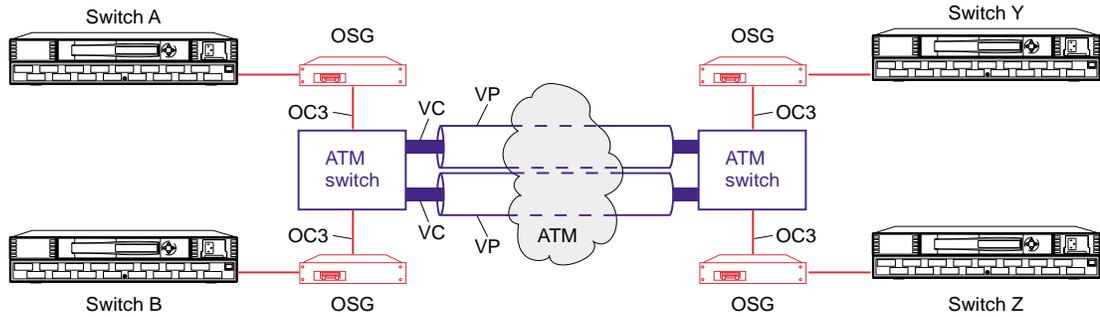


Figure 3 Shared virtual channel link configuration

CXO7202A

Fractional Links

Fractional links allow for leasing the exact bandwidth or access speeds needed to support a user's requirements (see Figure 4). This topology may also consist of shared links.

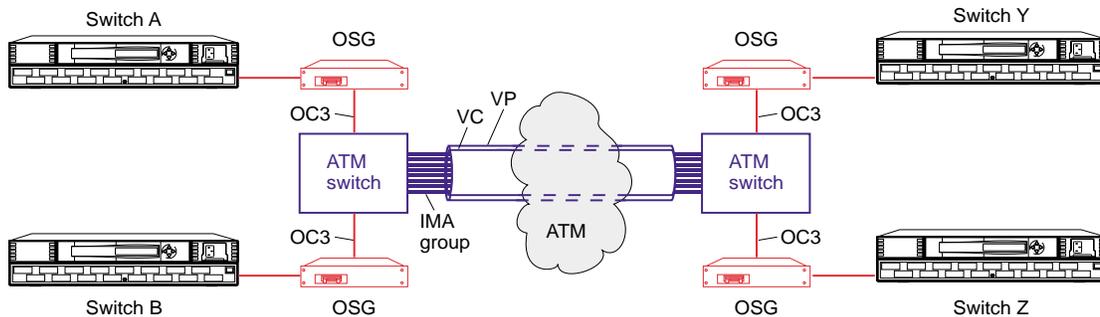


CXO7203A

Figure 4 Fractional link configuration

ATM Inverse Multiplexing

Multiple ATM T1/E1 speed circuits provide aggregate bandwidth when higher bandwidth is not available or when faster access rates are unnecessary (see Figure 5).



CXO7204A

Figure 5 ATM inverse multiplexing configuration

Encryption

Encryption is for customers who require strong confidentiality, authentication, and access control to ensure network and information security (see Figure 6).

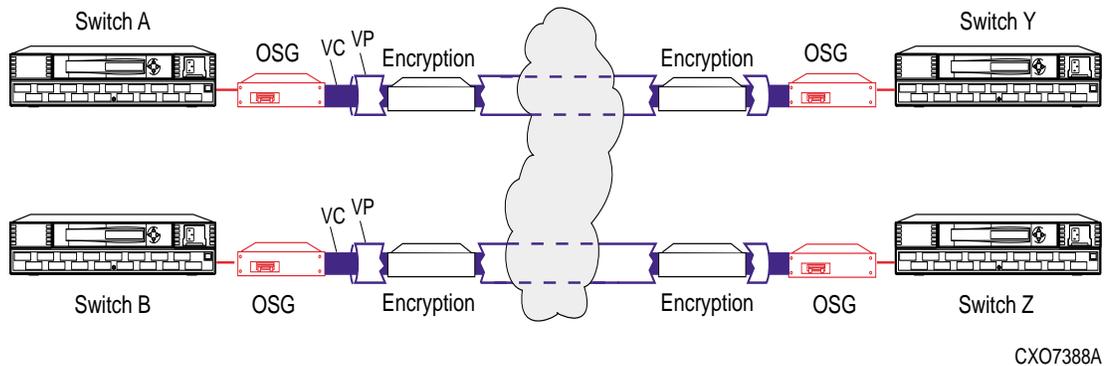


Figure 6 Encryption

Case Study

Example Criteria

- Distance between sites: 150 miles (241.35 kilometers)
- Write percentage: 30%
- I/O size: 8 kilobytes

DRM Requirements

DRM requires two separate PVCs. The telecommunications provider must supply VPIs and VCIs through their network for each PVC.

Obtaining Best Performance

When setting up a DRM configuration, changing the operational mode between synchronous and asynchronous can have a significant effect on performance as seen by the user. There are no hard and fast rules to determine which setting is best. The best setting depends on a combination of the I/O workload, the type of intersite communication, and the distance of the intersite link.

This section provides the background information necessary to obtain the best performance over a DRM connection with an ATM link.

Workload profile

In normal operation, read requests from the host are satisfied from the local storage system. Read performance is, therefore, similar to the performance in a non-DRM environment. Write operations, however, must be replicated to the remote site. This replication, which uses different communication transports over varying distances, introduces performance variations.

As the percentage of writes increases, the amount of data sent to the remote site also increases, potentially reducing performance. The first step in assessing the effect of DRM is, therefore, to analyze the existing I/O workload. The percentage of writes, the size of these writes, the write request rate, and the existing write response time should be determined. With this information it is possible to predict the resulting performance, using DRM under different configurations, with a reasonable degree of accuracy.

Distance considerations

When data is transferred from the local site to remote sites, the SCSI protocol, encapsulated into Fibre Channel frames, requires a minimum of four trips over the path. These are, conceptually:

1. Ask the remote site if it is ready to receive data.
2. Wait for the response from the remote site.
3. Send the data to the remote site.
4. Wait for acknowledgment from the remote site.

When sending data over fiber, the transmission time is approximately 5 microseconds per kilometer. Since a minimum of four trips is required for each SCSI data transfer, this translates to 20 microseconds per kilometer, or about 32.2 microseconds per mile. For example, if a remote site is located 150 miles away from the local site, the total time for every data transfer is 4830 microseconds. Since a typical I/O operation on a non-DRM configuration with write-back cache takes approximately 500 microseconds, long distances can have a significant effect on performance.

Communication type

The specific type of communication may also have a significant effect on performance. When using ATM over an OC-3 link, the data rate is limited to a maximum of 155 megabits per second. Straight Fibre Channel, on the other hand, has a native speed slightly over 1 billion bits (Gb) per second. If the communications link is ATM, the hardware that translates between Fibre Channel and ATM protocol introduces an additional latency of about 2.41 milliseconds combined for all four trips. The effects of these different transport mechanisms must also be taken into account when considering performance.

Taking the previous example of 150 miles (4.83 milliseconds), the addition of ATM increases the latency by 2.41 milliseconds to a total of 7.24 milliseconds. The maximum data rate is also lowered from the Fibre Channel maximum of 100 megabytes per second to approximately 13 megabytes per second (taking protocol overhead into account).

Synchronous operation

In the synchronous mode of operation, command completion is not returned to the host until the data has been stored on both the local and remote sites. If the host computer issues multiple write requests at the same time, all of these requests are pipelined (multiple requests on the connection at the same time). Even though multiple commands may be outstanding, completion for each command is returned only when all data for that command has been received and acknowledged by the remote site. Since the commands are pipelined, very little queuing occurs, so the response time of each request is determined by the total time for the four trips. The length and type of the communications path is the primary factor for the response time of each command.

The maximum number of commands that may be outstanding simultaneously is controlled by the `OUTSTANDING_IO` parameter on the controller. If the number of commands issued by the host exceeds this value, those commands are queued internally by the local controller for later transmission.

Asynchronous operation

To reduce response times, asynchronous operation returns I/O completion to the host as soon as the data is sent to the local controller but before it is received by the remote system. The result is that the host computer sees only the response time associated with the local system. The time delays associated with data replication are masked, since the replication is accomplished in the background.

Because completion of the transfer is returned to the host, the host may send additional requests, confident that the previous data is safely written. Because of this, the controller must ensure that all data that was previously acknowledged as being complete, but that is currently “in transit” to the remote system, is delivered in the correct order. This assurance is achieved by allowing only one data transfer at a time (for a particular association set) to be sent to the remote site. Any additional data transfer commands are queued internally within the controller and sent one at a time as the previous transfers complete.

Because the host sees a very fast response time, it may issue I/O requests at a rapid rate. The replication, on the other hand, proceeds at the rate dictated by the intersite link. As a result, a large queue of requests waiting for replication may build within the controller.

The number requests in the queue is limited by the `OUTSTANDING_IO` parameter. The number of requests that have been acknowledged to the host by the controller as being completed, but have not yet been sent (or acknowledged) by the remote site, is compared to the value of `OUTSTANDING_IO`. If the number of requests is greater, then subsequent requests sent by the host are not acknowledged as completed until they are sent to and acknowledged by the remote site.

A simple example may help to clarify this:

Assume that `OUTSTANDING_IO` is set to five and that a host issues a request, waits for completion from the controller, then immediately issues another request. Assume also that the DRM configuration is on ATM over a very long distance, requiring a significant amount of time to perform replication. In asynchronous mode, each request issued by the host is acknowledged as completed by the controller very quickly. As a result, the host issues five requests before the remote site has completed the first request. If the host then issues another (sixth) request, it exceeds the value of `OUTSTANDING_IO`. The local controller does not return completion of this request to the host, but queues it behind the previous requests that are awaiting replication. Completion of the sixth request is not returned until all five preceding requests are sent to and acknowledged by the remote site and the sixth request itself has been sent and acknowledged. The result is that the host sees a response time for the first five requests equal to the time taken by the local controller, while the sixth request has a response time equal to the sum of the transit times from local to remote controller for all preceding requests.

Assume that the time for the local controller is 500 microseconds, while the total time for replication to the remote controller is 5 milliseconds. In this case, the host sees a response time of 500 microseconds for each of the first five requests, since completion is returned immediately. The sixth request has a response time slightly less than 30 milliseconds, since it must wait for all previous requests to be completed by the remote site. In other words, the sixth request has a response time sixty times greater than the first five.

Performance analysis

Taking the preceding information into account, the performance picture becomes quite complex, since there are so many factors that interact. To analyze the performance of a DRM configuration and select the optimum parameters, all of the contributing factors must be fully quantified. This analysis entails a multiple step procedure, and requires detailed knowledge of the I/O workload and proposed DRM configuration.

The steps required to configure a DRM solution for optimum performance consist of both preinstallation and postinstallation analyses. The preinstallation phase consists of measuring the current I/O workload on the units to be replicated, calculating the minimum transport bandwidth required to support the existing I/O load, and predicting the response times for the DRM system for both synchronous and asynchronous replication. Once the system is installed, additional steps may be taken to tune the configuration for maximum performance.

The preinstallation phase consists of the following steps:

1. Determine the I/O workload parameters that are relevant to DRM. These parameters should be measured only on those units that are to be replicated, since non-DRM units are irrelevant to this analysis.

The parameters to be measured for each unit include:

- Transfer size of writes
- Request rate of writes
- Either the average response time or the average queue depth of writes
- Either the average response time or the average queue depth of reads (for a later, more comprehensive analysis)
- Transfer size and request rate of reads (for a later, more comprehensive analysis)

The data may be measured with existing operating system utilities (such as the Windows NT Performance Monitor) or third party utilities (such as ViewPoint from Datametrics). It is important to gather this data during both typical and critical periods. Since the data gathered is used for later analysis, the results of the analysis depend on the quality of the collected information. If there are several distinct I/O phases (such as interactive and batch), data should be collected separately for each phase and the resulting data analyzed independently.

2. Determine the minimum transport bandwidth required. The bandwidth required is the product of the write request rates and the write transfer sizes for all replicated units. For example, if the busiest period shows a total write request rate of 150 requests per second with an average transfer size of 8 kilobytes, the required bandwidth would be 1200 kilobytes per second (1.2 megabytes per second), so neither ATM nor straight Fibre Channel would impose significant bottlenecks. Note that due to other considerations, this rate may not be maintainable. This step is merely to ensure the minimum bandwidth required to meet the existing data rate.
3. The latency for synchronous replication may be calculated once the distance between the local and remote site is known. For straight Fibre Channel, the latency in milliseconds is the distance in miles multiplied by 0.0322 (distance in kilometers must be multiplied by 0.02). If ATM is used, 2.41 milliseconds must be added to this figure to account for the ATM protocol conversions. For a 150-mile distance, for example, the fiber latency is 4.83 milliseconds (7.24 milliseconds for ATM as shown in the section "Communication type" on page 9).

Although there may be large latencies associated with synchronous replication, the ability to send multiple requests from the host simultaneously allows a very high request rate to be supported.

4. For asynchronous replication, the host sees very little latency since completion is returned when the local controller receives the data. Only one data transfer may be outstanding between the local and remote site in asynchronous replication mode however, so the maximum request rate is dependent on the amount of time taken for that replication. The maximum rate can be calculated by taking the reciprocal of the latency calculated in step 3. Using 4.83 milliseconds (0.00483 seconds) as an example, 1 divided by 0.00483 equals about 207 requests per second. If the connection is made via ATM, the latency from step 3 of 7.24 milliseconds results in a maximum request rate of about 138 requests per second.

The calculated rate is the maximum rate that can be supported while retaining full asynchronous response times. If the rate rises above the calculated value, then more requests are arriving at the controller every second than can be sent sequentially over the communications link. This results in queues of requests building at the local site and, when the number of requests exceeds the value of `OUTSTANDING_IO`, the request is delayed for a significant period of time (as shown in the section “Asynchronous operation” on page 10), resulting in a large increase in response time.

It is important to note that most I/O is quite bursty. Although the average arrival rate may be below the thresholds calculated in this step, bursts of I/O requests may cause an instantaneous peak rate and lengthy response times.

A rule of thumb is to aim for an average rate of no more than 80% of the rates calculated in this step. From the examples given, an average write rate of more than 165 requests per second (fiber) or 110 write requests per second (ATM) may result in long response times. Once the maximum write request rates have been calculated, the minimum transport bandwidth should (once again) be determined. As shown in step 2, this determination is accomplished by multiplying the average write transfer size by the write request rate. For the examples given, a maximum rate of 165 write requests per second yields a data rate of 1.32 megabytes per second. For a rate of 110 write requests per second on ATM, the minimum required bandwidth is 880 kilobytes per second.

5. If the data gathered in step 1 indicates that the write request rate for an association set exceeds the maximum request rate for asynchronous requests calculated in step 4, then better performance can be realized by setting the operation mode to synchronous. On the other hand, if the request rate from step 1 is less than the maximum rate for asynchronous mode, then setting the operation mode to asynchronous provides the best response time.
6. The optimum value of `OUTSTANDING_IO` depends on both the type of unit and the spatial locality of requests sent to individual units.

For JBOD units, if there is a possibility that two or more requests will be sent rapidly to the same location on the unit, the value of `OUTSTANDING_IO` should be set fairly low. A value of 10 or less is recommended.

For RAID-5 sets, if there is a possibility that two or more requests will be sent rapidly to the same strip, the value of `OUTSTANDING_IO` should also be set fairly low. To calculate the size of a strip, multiply the chunk size of the RAID-5 set by one less than the number of members in the set. For a six member RAID-5 set with a chunk size of 256, for example, the size of a strip would be 256 times 5, or 1,280.

In any event, performance of the unit should be monitored with tools such as `VTDPY` and `DSTAT`, and the value of `OUTSTANDING_IO` adjusted as required to attain the maximum throughput.

7. Once the preceding calculations are complete, it may be of interest to calculate the overall maximum request rate. Since all calculations are based on the write request rate only, the combined request rate is determined by dividing the write request rate by the fraction of writes. The result is the maximum request rate that this configuration supports.

For example, if the write fraction is 0.3 (30%), then a write rate of 165 per second for fiber results in a maximum request rate of 165 divided by 0.3, or 550 requests per second. For a write rate of 110 on ATM, the maximum request rate is 367 requests per second.

Summary

When choosing between synchronous and asynchronous replication, the best performance depends on a variety of workload factors. Since the setting of this parameter is dynamic, it is relatively easy to switch between the two operational modes and view the resulting performance change. Guidelines for setting the mode are not intended to be hard and fast rules, but should be considered a starting point.

Note that as workload conditions change over time, changes may be needed in the operation mode to maintain optimum performance.

Installation and Configuration

Fibre Channel Switch Setup

Set up the Fibre Channel switch as required in the *Compaq SANworks Data Replication Manager HSG80 ACS Version 8.5P Operations Guide* for your platform.

The Fibre Channel switch configuration must be modified to correctly communicate with the OSGs. Verify or change the following as appropriate:

1. Open a telnet session and log into the Fibre Channel switch.
2. Type `version` and then press **Enter**. You will see a display similar to the example below.

```
Kernel:      5.3.1
Fabric OS:   v2.1.8_rc2
Made on:     Fri Jun 23 13:40:45 PDT 2000
Flash:      Fri Jun 23 13:41:17 PDT 2000
BootProm:   Tue Dec 29 17:32:38 PST 1998
```

Verify that the switch (Fabric OS) is running version 2.1.7 or greater.

3. Type `licenseshow` and then press **Enter**. You will see a display similar to the example below.

```
Web license
  Zoning license
  SES license
RdSebdezQ9SeTfcQ:
  Web license
  Zoning license
  SES license
  QuickLoop license
  Fabric license
  Remote Switch license
```

Verify that the remote switch license is loaded.

If the switch license is not present, load the license by typing `licenseadd "license-key"` and then press **Enter**. To obtain the license key, refer to the Remote Switch Option Kit provided with your Fibre Channel switches. Repeat step 3 to verify that the license has been added.

4. Configure the switch for E-port connectivity with OSGs:
 - a. Type `switchdisable` and then press **Enter**.

- b. Type **configure** and then press **Enter**. You will see a display similar to the example below.
Configure...
Fabric parameters (yes, y, no, n): [no]
Type yes and then press **Enter**.
- c. Press **Enter** several times until you see a display similar to the following:
E_D_TOV: (1000..5000) [5000]
Type 5000 for this value and then press **Enter**.
- d. You will see a display similar to the following:
Data field size: (256..2112) [2048]
Type 2048 for this value and then press **Enter**.
- e. Press **Enter** several times until you see a display similar to the following:
Suppress Class F Traffic: (0..1) [1]
Type 1 for this value and then press **Enter**.
- f. Press **Enter** several times until you see a display similar to the following and you are back to the telnet prompt. The configuration process is complete:
Virtual Channel parameters (yes, y, no, n): [no]
Arbitrated Loop parameters (yes, y, no, n): [no]
System services (yes, y, no, n): [no]
No changes.
- g. Type **switchenable** and then press **Enter**. The switch setup is now complete.

Service Kit

The Service Kit is required to remotely service and configure the Open Systems Gateway (OSG) through a modem. The kit consists of a Western Telematics (WTI) switch and necessary RJ-45 cables.

The WTI switch is a 16-port, asynchronous switch that mounts in a standard 19-inch rack along with the (OSG). A customer-supplied modem must be installed to enable remote servicing of the OSG. You must obtain a local phone number from your telecommunications provider for this modem.

The kit does not supply the cables required to connect locally to a server or terminal device. See your CNT hardware *Installation User Guide* for remote and local connectivity options.

Open Systems Gateway Setup

Use the Fibre-Channel-to-ATM Configuration Wizard to set up both gateways. Run *fcadm* from the Wizard menu. Answer “no” to all other questions except for the following required information:

- Node number (must use two digits)
- Ethernet IP address
- Ethernet subnetwork mask
- Internal clocking (set to *no* when connected to ATM switch)
- Default gateway IP address

NOTE: Prior to running *fcadm* wizard, verify that console message “active received” has been displayed. After *fcadm* wizard has terminated, wait for another console message “active received.” The “active received” message may take several minutes in both cases. Once the last “active received” message appears, it is safe to abort the OSG. After the OSG has rebooted, all changes are in effect.

Site Readiness Checklist

Verify the follow information before starting the configuration process:

- Copy of ATM QoS contract supplied to Compaq available
- VPI: _____ VCI: _____ PCR: _____ CDVT: _____
(minimum 3000 microseconds)
- ATM switch manufacturer: _____
- Model number of the switch: _____

Modifying VPI, VCI, QoS, and PCR Values

Edit the *atmlink* file to modify the VPI, VCI, QoS, and Peak Cell Rate (PCR) values. Refer to the Site Readiness Checklist in this Application Note for VPI and VCI values.

1. Type the following commands at the root level to access the *atmlink* file:

```
fs cd /cnt/configs/active/slot01/cfg
ed atmlink
```

2. The *atmlink* file contains the following:

```
AtmPvcVPI = value
AtmPvcVCI = value
AtmPvcQOS = value (change to CBR)
```

3. The *atmlink* file also contains the following:

```
AtmPvcUcrPcr = value
```

This value is the PCR of the gateway and is expressed in kilobits per second. It is also defined as the Premise Line Rate. Enter a value in the acceptable range of 1544 to 126000 kilobits per second. The value required for this field is the result of a performance analysis similar to the one performed in the previous Case Study.

Efficiency values that should be considered when specifying Line Rate with the telecommunications provider are as follows:

- SONET efficiency = 150.336 megabits per second / 155.52 megabits per second = 0.967
- ATM efficiency = 48 cells / 53 cells = 0.9
- Fibre Channel 8b/10b encoding = 8 bits / 10 bits = 0.80
- Fibre Channel efficiency = 2048 bytes / (2048 + 48 bytes) = 0.977

These values as well as any additional telecommunication-specific values must be used when calculating the bandwidth that the telecommunications provider must furnish to ensure the required DRM bandwidth.

As in our Case Study, if the required data rate is 880 kilobytes per second (7040 kilobits per second), then the following values must be used for correct premise and line rate setup:

- Premise Line Rate = $8 * 880 / (0.8 * 0.977) = 9007$ kilobits per second (Peak Cell Rate), which is calculated by adding ATM and SONET efficiencies.
- Line Rate = $8 * 880 / (0.967 * 0.9 * 0.8 * 0.977) = 10350$ kilobits per second (the telecommunications provider must supply 2 PVCs at this rate), which is calculated by adding Fibre Channel, 8b/10b, ATM, and SONET efficiencies.

Two PVCs with a QoS of CBR, and a PCR of 10350 kilobits per second, would be required to support this case study.

The DRM Solution manages data using remote copy sets (RCSs). Each RCS requires a unit to be associated with it. The combination of all RCSs (units) must not create a data load greater than the link can support. In this case, the link can support 10350 kilobits per second.

NOTE: Extreme care must be used when creating and managing RCSs on an ATM link, especially at the low speeds near the 1.544 megabits per second region. If these links are overstressed, unexpected behavior may result, such as loss of remote copy sets and link failures.

Table 1 Example - Premise and Line Rate Setting vs. Data Rate

Max Data Rate (Based on No Delay) (Kbps)	Premise Line Rate (AtmPvcUcrPcr) (Kbps)	Public Line Rate (cells per sec)	Public Line Rate (Mbps)
1050	1344	3499	1.544 (ATM T1)
1393	1782	4642	2.048 (ATM E1)
3401	4352	11332	5
6802	8703	22664	10
13605	17406	45328	20
20407	26109	67992	30
23378	29910	77892	34.368 (ATM E3)
27209	34812	90656	40
30431	38934	101390	44.736 (ATM T3)
34011	43515	113320	50
40814	52218	135984	60
47616	60921	158648	70
54418	69624	181313	80
61220	78327	203977	90
68023	87030	226641	100
81627	104436	271969	120
95232	121842	317297	140
98289	125754	327484	144.495 (ATM OC-3)

Modifying ActiveVciBits Values

Edit the *atmoc3* file to modify the *ActiveVciBits* values as indicated in Table 2. Refer to the Site Readiness Checklist in this Application Note for VPI and VCI values.

1. Type the following commands at the root level to access the *atmoc3* file:

```
fs cd /cnt/configs/active/slot01/cfg
ed atmoc3
```

2. The *atmoc3* file will contain the following:

```
AtmVciBits = value
```

Table 2 ActiveVciBits Values

ActiveVciBits	VPI Range	VCI Range
12	0	0- 4095
11	0-1	0- 2048
10	0-3	0-1023
4	0-255	0-15

NOTE: For information about ed commands, refer to the *UltraNet Open Systems Gateway Command Reference Guide*. For detailed installation information consult your *UltraNet Open Systems Gateway Hardware Installation Users Guide*.

Websites

Compaq

Check the Compaq website for more information on the complete line of Fibre Channel storage products, product certification, technical information, updates, and documentation. This information can be accessed through our website at:

<http://www.compaq.com/products/storageworks/>

ATM Forum

Over 600 organizations belong to the ATM Forum, an international standards organization that promotes the use of ATM products and services. ATM Forum can be accessed at the following website:

<http://www.atmforum.com/>

Disclaimer

Compaq does not warrant third party products, which customers may purchase in the making of a DRM solution. These products are used on an “as is” basis unless otherwise specified in the Third Party’s Warranty Attributes. Third party products may be warranted by the third party as specified in the documentation provided with the third party products.

Glossary

ATM (<i>Asynchronous Transfer Mode</i>)	A communications networking technology for both LANs and WANs that carries information in fixed-size cells. ATM uses logical connections to provide quality of service guarantees that enable disparate traffic such as data, voice, and video to be carried over the same local or wide area network.
CBR (<i>Constant Bit Rate</i>)	A category of ATM service. CBR supports a constant or guaranteed data rate. CBR is designed to support applications that need a highly predictable transmission rate.
CDVT (<i>Cell Delay Variation Tolerance</i>)	The amount of delay, expressed in microseconds, of a cell from its actual delivery time.
Data Rate	The bandwidth required without any protocol overhead (8-bit bytes).
Default Gateway	The default path that a computer or router uses to forward and route data between two or more networks that have different protocols.
Fibre Channel	Technology for very high-speed, switching-based serial transmissions.
IP address	The numeric address that identifies a particular computer on the internet.
Latency	The amount of time required for a transmission to reach its destination.
OC-3	An optical carrier that provides high-speed bandwidth at 155 megabits per second.
PCR (<i>Peak Cell Rate</i>)	The maximum transmission speed of a virtual connection. PCR is a required parameter for the CBR service category.
Premise Line Rate	The speed of the link between the premise edge device and user equipment.
Premise Network	All transmission connections between user equipment and edge devices to the public network.
Public Line Rate	The speed of the link between the premise and public networks.
PVC (<i>Permanent Virtual Circuit</i>)	A logical connection manually defined by the network administrator. The PVC is created by specifying the VPI and VCI.
Public Network	All transmission connections supplied by the telecommunications provider terminating at a customer edge device.
QoS (<i>Quality of Service</i>)	Each virtual connection in an ATM network has a service category. The performance of the connection is measured by the established QoS parameters, which are outlined by the ATM Forum. Performance issues include data rate, cell loss rate, cell delay, and delay variation (jitter). Categories of ATM service are Constant Bit Rate (CBR), Variable Bit Rate-Real Time (VBR-RT), Variable Bit Rate- Non-Real Time (VBR-NRT), Available Bit Rate (ABR), and Unspecified Bit Rate (UBR).

SCSI (Small computer system interface)	A processor-independent, standard protocol for system-level interfacing between a computer and intelligent devices, including hard drives, floppy disks, CD-ROMs, printers, scanners, and others.
SDH	European version of SONET.
SONET	A standard optical interface for transmission systems. SONET is a synchronous protocol and is the physical layer for ATM LANs operating over fiber at 155 Mbps or higher. Note: OSGs default to SONET. If SDH is required, edit the ATMOC3 configuration file.
Speed of light through fiber	Approximately 200,000 kilometers per second or 5 microseconds per kilometer.
Subnet mask (also known as address mask)	An IP network that can be reached through a single IP address. All the members of the subnet share the mask value. Members of the subnet can then be referenced more easily. A subnetwork is a network that is part of another network, connected through a gateway, bridge, or router.
UBR (Unspecified Bit Rate)	A category of ATM service. UBR supports connections that have no performance requirements.
UltraNet Wizard Fibre Channel-to-ATM Configuration Wizard	An <i>UltraNet</i> application that lets the user designate the default configuration settings for Fibre-Channel-to-ATM on the Open Systems Gateway.
VCC (Virtual Channel Connection)	Defines a logical networking path between two endpoints on the network. ATM cells travel over this connection.
VCI (Virtual Channel Identifier)	The field of the cell header that stores the virtual channel address.
VPC (Virtual Path Connection)	A series of virtual paths joined together to form the logical groups of circuits defined for each link of the network. VPC moves traffic from one link to another.
VPI (Virtual Path Identifier)	The field of the cell header that stores the virtual path address.
WTI Switch	A 16-port serial switch that can be used to configure or service the OSG unit locally or remotely. The Western Telematics (WTI) switch must be installed to set up and service the ATM gateway.

