May 2012

IBM® Virtualization Engine TS7700 Series Best Practices

TPF Host and TS7700 IBM Virtualization Engine V1.1

Gerard Kimbuende
gkimbue@us.ibm.com
TS7700 FVT Software Engineer

John Tarby
jtarby@us.ibm.com
TPF Tape support

Vanessa Earle
earlev@us.ibm.com
TS7700 FVT Software Engineer

1.1 Purpose

The TS7700 IBM Virtualization Engine can attach to a variety of host systems, including z/OS®, IBM z/VM, IBM z/VSE™, and IBM z/TPF. Each has its own configuration and requires different tasks to be performed to optimize the solution. The purpose of this document is to highlight some best practices with regards to the IBM Virtualization Engine TS7700 and the Transaction Processing Facility (TPF).

TPF processes real time transactions and is capable of handling large, continuous loads of transactions at very fast speeds. Oftentimes this high transaction load results in high data rates to tape for logging and recovery purposes. The ability of the TS7700 to emulate a large number of physical tape devices and accommodate high data rates makes it an attractive solution. However, there are a few things to consider as well. For TS7700 customers using TPF as their host, it may become necessary to pay special attention to characteristics that may negatively impact the performance of the machine.

When tape processing on the TS7700 slows down for any reason, host write requests begin to form a queue on TPF. The rate of the request in conjunction with the length of the delay determines the amount of resource that needs to be allocated to keep TPF operating without impact to the end users. In cases where the queue builds to the point where the amount of available blocks are reduced to critically low levels, host applications, may be delayed and/or the tape writes may be redirected to other tape devices. This may cause customers to experience longer execution times for some applications.

This document will explain what occurs during these periods as well as provide some recommendations on what can be done to help minimize timeouts in your environment.

1.2 Performance and the TS7700 Grid

There are some performance considerations that should be noted with regards to TS7700 grid environments. Certain operations submitted by TPF require all clusters be operational and able to communicate with each other. The majority of the time, communication between the host and the TS7700 occur without impact or delay to the host. One issue with TPF arises when a cluster is failing and the wait periods on the TS7700 exceed the timeout values on TPF. In situations where a cluster within a grid unexpectedly becomes inoperable, the TPF host may not be able to perform certain commands due to timing constraints including:

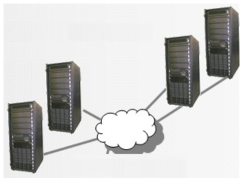

- Changing status of a volume
- Loading tapes
- Changing the category of a volume via ZTPLF command or user exit when a volume on a tape device is redirected

To protect against timeouts during a cluster failure, TPF should be configured to avoid issuing tape library commands to devices in a TS7700 grid along critical code paths within TPF, specifically during a tape device switch. It is recommended that customers set the user exits on TPF so that the category is changed when a volume is loaded by TPF then not changed again.

TPF implements a strict 16 second time-out during normal write/read operations. While the TS7700 has multiple redundant components designed for fail over and high availability of the system, there is a possibility that TPF could time out prior to the completion of the failover process, which would normally have prevented write time-outs. Historically, customers have configured their environment to accommodate the delays associated with physical tape (temp write errors, retries, etc). With the introduction of virtual tape, many of these I/O response delays have been eliminated. However, background tasks such as recovery of disk cache, DB health check out, and system paging are all contributing factors to a device slowdown. Occasional delays of 8 seconds or less can be considered normal but may affect TPF. During this window, it is common to see sharp performance fluctuations depending on resource availability.

A user exit has been developed which allows the TPF MIH value, also known as the TPF Lost Interrupt Value, to be dynamic and not restricted to the historical 16 second setting. The exit can alter this value based on device level failures on the TS7700 and other factors in the TPF system. In the event a device becomes unresponsive for longer than one second, the exit is called. Code added in the exit can determine if it is appropriate for TPF to continue to wait for this operation to complete or if the operation should be terminated. The determination to abort the operation should be made before the available TPF resources have been depleted to the point where it will impact overall system availability. Balancing the amount of resources that are allocated along with the time that TPF will wait for an operation to complete will result in continued, uninterrupted availability of TPF. This user exit, along with tuning of the TS7700, will either eliminate or at least minimize the effects of these write timeouts. Both these recommendations will be discussed later in this document.

2.1 TPF Sizing and Planning

When determining the size of a TS7700 subsystem to be used by TPF, several factors should be taken into consideration. The first point to consider is that a TPF system must always have a set of system tapes to write to. This means that in any TPF configuration there must be at least two TS7700 clusters attached to TPF at all times. This allows for both planned and unplanned hardware outages to occur without impacting the TPF system. In the case of an outage, the remaining cluster should have enough device addresses and throughput capacity to allow the TPF system to continue to operate.

The second factor is whether the clusters should be configured in a grid or if they should be independent. The advantage to the grid configuration is that through the use of data replication the volumes created by TPF would always be available even if one of the clusters became unavailable. The drawback to a grid configuration is that for certain errors the two clusters in the grid will be unable to communicate with one another for several minutes. During this period, the TPF system may be unable to load new logical volumes into virtual drives when the previous volumes are unloaded. This condition can be mitigated by keeping an adequate supply of ALT tapes loaded on each cluster at all times. This ensures that the TPF system has a tape mounted if a tape switch is required and it is not possible to load a volume into the drive that just had a volume unloaded from it. The impact of these errors can also be reduced by using the category change user exit so that the tape category is not altered during tape switch processing. An ALT tape is a generic standby logical volume that can be used by different active volumes if needed. The use of ALT tapes can be further managed by using the ZTDEV, ZTLBL and ZTGRP operator commands to associate an ALT tape with a specific tape name or group of tape names.

If it is determined that a grid configuration is desired, then the copy mode needs to be determined: No Copy, Rewind Unload (RUN), Deferred, or Synchronous. With the deferred copy mode, volume replication occurs later in time after close processing allowing the device to be reused immediately for another volume access. This is the preferred copy mode if host performance is most important. In Rewind Unload mode, volume replication occurs as soon as the RUN command is received. The device itself is not available again until a copy of the volume has been made to the other clusters in the grid. Depending on the size of the volume, the speed of the link between the clusters in the grid, and the number of copies that are in progress at the same time, replication of volumes can take several minutes. Status of the RUN command is held during the replication process making it a poor candidate for TPF operations. Synchronous copy mode duplexes host writes simultaneously to two library locations and provides a zero recovery point objective (RPO) to the TS7700. Synchronous replication relies on remote cluster access and thus can impact TPF operations when network issues arise. Synchronous copy mode is only available for machines running microcode 8.21.0.xx or higher. If using Sync Mode Copy it would be best to enable the Synchronous Deferred on Write Failure option. This way, if the secondary sync copy cannot be made, the write does not fail.

2.2 Tuning TPF:

To efficiently configure the z/TPF system to use the TS7700, a user must identify the various types of devices that are mounted for output along with the method and rate in which records are written. Data that is written to tape can be divided into two general categories. Certain applications that request writes without waiting for the record to be written to the tape device, or into the tape blocking buffer, can cause large queues to grow. Alternatively, applications that wait for each record to be written before requesting the next write will never generate a large queue regardless of the rate at which they are writing to the tape device. It's the first situation that must be accommodated through tuning.

If possible, applications which request writes without waiting for them to complete, should throttle themselves or terminate processing if the module queue starts to build. The application should then resume writing when the queue has been processed. An example of a TPF system function that does this is PIU trace. If it detects that the queue is growing to an excessive depth then trace is quiesced until the queue has decreased.

Some applications such as exception recording or logging cannot be throttled because they are writing data that is a function of the workload on the TPF system itself. In these cases, adequate system resources need to be allocated to allow TPF to queue the pending write requests until the device has resumed normal processing. If sufficient resources are not allocated, the TPF system may enter into an 'input list shutdown' until the tape device resumes normal processing. As described earlier, for the TS7700 family of devices, under normal conditions there may be occasional periods of time where the device will not respond to a host requested write for up to 8 seconds. While this will not occur frequently, the TPF system must be configured to tolerate delays of this magnitude. In extreme cases, specifically during error scenarios, this delay could be even greater and exceed the default timeout values for TPF's lost interrupt processing which has been set to 16 seconds for a write command. When this happens the TPF system will abort the current operation and initiate tape device switch if applicable.

2.3 TPF Software Buffers:

Determining how many blocks need to be allocated to tolerate applications that request writes without waiting for the record to be written into the tape blocking buffer involves calculating the number of tape write requests that are made by the application on a per second basis. When the device is not responding, each of these requests will result in two blocks being used: one 4K frame to hold the data and one System Work Block (SWB) to hold the position of the write on the module queue. A 4K frame is used even if the size of the data block being written is smaller than 4K. Once the number of write requests per second is determined (for example 15,000 / second) the number of blocks can be calculated by multiplying that value by the default timeout value of 16 seconds.

$$\left(\frac{15{,}000}{\cancel{\text{seconds}}}\right)(16\ \cancel{\text{seconds}}) = \ 240{,}000\ \textit{blocks of buffer space needed}$$

At a minimum, enough 4K Frames and SWBs should be allocated so that 16 seconds worth of write requests could be added to the module queue without causing the TPF system to enter an 'input list shutdown' condition. If the active and standby tapes are typically mounted on the same TS7700 subsystem then it is possible for a cascading effect to occur where TPF attempts to use both drives on the same subsystem with each drive being unable to respond in that 16 second period. Now, instead of that tape being unavailable for 16 seconds it could potentially be unavailable for 32 seconds. Although this condition is rare, the TPF system should be configured to tolerate this case.

2.4 TPF User Exit:

It may be determined that high data rates of certain applications make it impractical to allocate enough SWBs and Frames to accommodate the size of the module queue that could build during the potential delays. If this is the case, the default lost interrupt value assigned by TPF can be overridden through the use of a user exit. The user exit was added with z/TPF APAR PJ39803 and it provides a mechanism to override the default lost interrupt values. The user exit is called any time it takes longer than one second for a tape operation to complete. By using this exit, a customer can decide to abort the current operation before the default time has elapsed and significantly reduce the number of blocks that would be required for TPF to tolerate these extreme cases. The decision to abort the operation can be based on any number of factors including the current queue depth as well as the resource levels for other blocks in the system. The consequences of aborting the operation is that TPF will determine that the status of the operation is unknown and needs to presume that it failed. Based on this presumption, TPF will stop writing to the first tape and begin writing to the standby or ALT tape. The first record written will be the one that timed out on the previous tape. There is a high likelihood that the record in question will exist as both the last record on the previous tape and the first record on the next tape.

By default the Tape Lost Interrupt user exit is set to not allow altering the values. Sample code is provided in copy member /base/cp/cusr.cpy that can be used to abort the operation if the queue depth grows above a preset value and the operation on the top of the module queue has been active for longer than one second. The sample code uses a tabular structure to hold a list of tape names along with the maximum queue depths that are allowed for each tape name. Additionally the exit will keep track of when the last operation timed out and what the queue depth was at that time. This allows the exit to provide a recovery period of a few seconds following a tape switch so that the performance of the new tape can be evaluated without immediately causing the operation to time out.

The following is an example of this Exit routine being used to handle a hardware error while a high volume data write application was in progress:

Reboot 1 cluster within 2 cluster grid environment while TPF I/O running
Verify that when the TS7700 reboots with I/O and TPF running. When returning online, the system will be fully operational and will recover successfully. The cache subsystem will not reboot. This will simulate a check1 reboot.
1.1.1 Test Case Execution
While I/O and TPF is running force a reboot on the TS7700, simulating a check1.
1.1.2 Test Case Observation
This forced reboot will cut all connection to the host while I/O and TPF is running. Any write jobs that were running at the time of reboot will be moved to another TS7700 via the TPF user exit that is in place and the user will see an IFCC will. Once the machine is recovered and online, all I/O should resume successfully, and no permanent errors should be present.
• Verify the proper TPF tape swich occurs via TPF user exit.
• Verify proper IFCC's occur.
• Issued a forced shutdown on the TS7700.
• TPF user exit was called, and TPF queue grew over 50,000 threshold to trigger tape switch.
• No devices were avilable on the original cluster, so TPF I/O failed over to a secondary TS7700.
• I/O continued to run on secondary cluster until original cluster became available.

For the scenario above, TPF no longer has access to any device on the original cluster. The exit is called and TPF detects the queue growth is at critical levels. After only a few seconds it aborts the operation and initiates a tape switch. I/O then continues on the failover cluster until the original cluster is available again. The exit is coded to detect the maximum tolerable level of queuing which, in this case, represents approximately 3 seconds. Following the tape switch, the queue is emptied and the system continues to function normally. In the event the original cluster becomes available, I/O automatically switches back once the tape fills on the failover cluster.

3.1 The New TS7700 Servers

Part of the reason some customers may see an issue with their TPF environment is that many are still configured the same way that they were when they migrated from the Virtual Tape Server (VTS) to the TS7700. For customers that are at earlier releases of TS7700 hardware and microcode, an increase in workload can push the TS7700 to its capacity. The reality is these customers may need to allocate more resources to accommodate their growing environments. The chart below shows that the performance of the R2.0 TS7700 V07 with the newer disk cache CC8 is much faster today than any of its predecessors (refer to Figure 1 and 2). In fact, there can be more than a 2X performance boost depending on the current hardware! More specific performance measurements are available by going to 'IBM Virtualization Engine TS7720 and TS7740 Release 1.6 through R2.1 Performance White Paper'. It can be found at http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101938 .

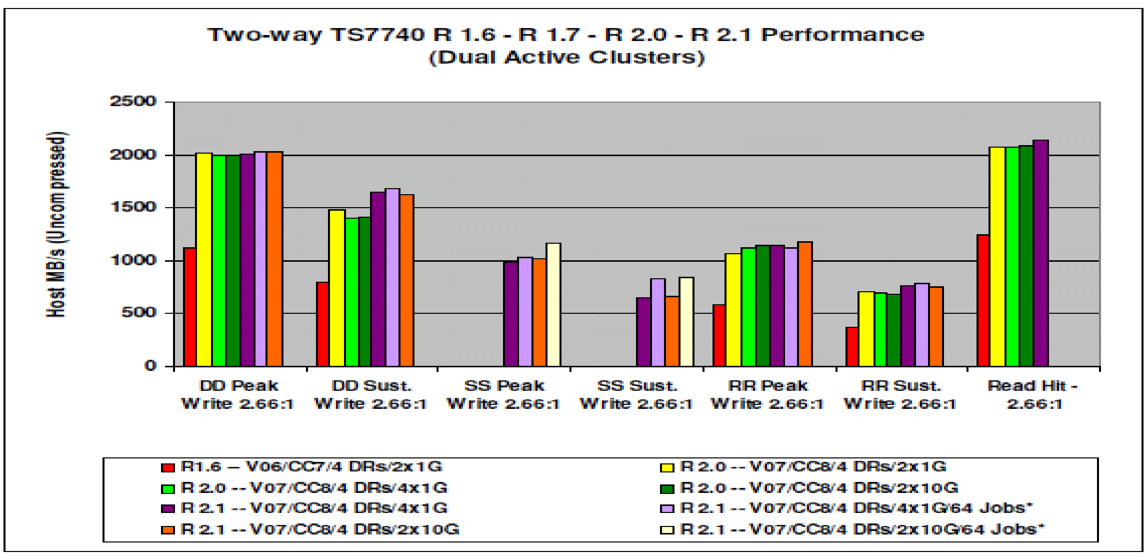

* 64 jobs – 128 jobs (64 jobs per active cluster)

Figure 1 . Two-way TS7740 Grid Maximum Bandwidth for Dual Active Clusters. Unless otherwise stated, all runs were made with 256 concurrent jobs (128 jobs per active cluster). Each job writing or reading 800 MiB (300 MiB volumes @ 2.66:1 compression) using 32 KiB block size, QSAM BUFFNO = 20, using eight 4Gb FICON channels from a z10 LPAR (four 4Gb channels per active cluster). Clusters are located at zero or near zero distance to each other in laboratory setup. DCT=125.

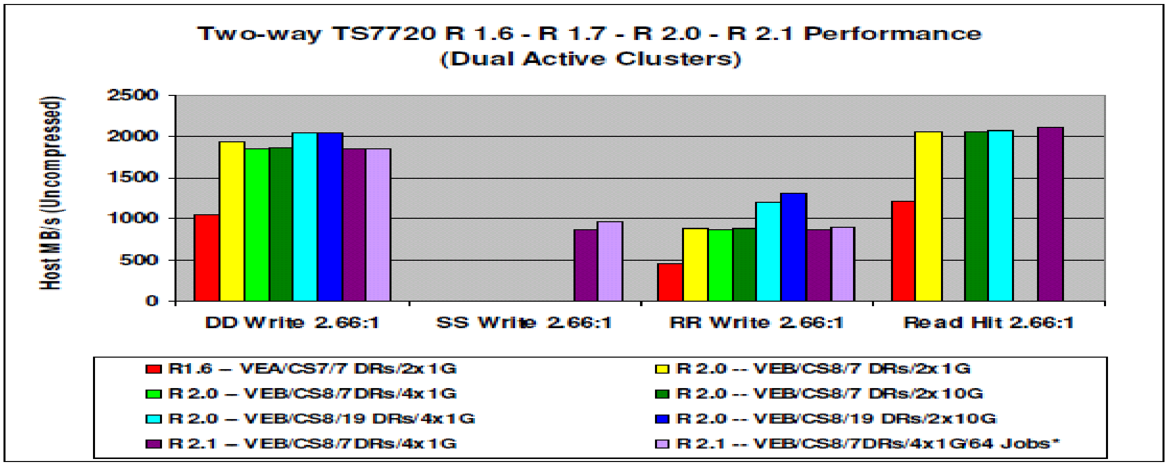

* 64 jobs – 128 jobs (64 jobs per active cluster)

Figure 2 . Two-way TS7720 Grid Maximum Bandwidth for Dual Active Clusters. Unless otherwise stated, all runs were made with 256 concurrent jobs (128 jobs per active cluster). Each job writing or reading 800 MiB (300 MiB volumes @ 2.66:1 compression) using 32 KiB block size, QSAM BUFFNO = 20, using eight 4Gb FICON channels from a z10 LPAR (four 4Gb channels per active cluster). Clusters are located at zero or near zero distance to each other in laboratory setup. DCT=125.

For 2-way TS7720 dual active clusters SS run, reducing the current jobs per active cluster from 128 to 64 give about 10% performance improvement.

3.2 Sizing and Planning for the TS7700:

Users must consider their current TS7700 environment when deciding workload and throughput requirements. This is especially true for TPF hosts due to the timeout functionality. Table 1 below highlights the different generations of TS7700 hardware and their approximated cache I/O rates. The 3957-V07 and 3957-VEB (highlighted in green) are the preferred hardware for TPF hosts. However, the older 3957-V06 and 3957-VEA can be used as long as I/O rates are determined to work within the bounds of the hardware's abilities. The table below represents theoretical, non-compressed cache IO data. It is **NOT** host performance nor is it workload data but rather internal gpfs calculations.

Table 1:

TS7700 Type	Cache Type	MC Level	Internal Memory	Cache IO Rates		
				Read	Writes	Expansion Drawers
3957-V06	CC6/CX6	8.2x.x.xxx and Higher	16 Gig	386 MB/s	194 MB/s	0 Expansion
3957-V06	CC7/CX7	8.2x.x.xxx and Higher	16 Gig	663 MB/s	418 MB/s	1 Expansion
3957-V06	CC8/CX8	8.2x.x.xxx and Higher	16 Gig	759 MB/s	587 MB/s	1 Expansion
3957-V07	CC8 and Higher	8.2x.x.xxx and Higher	16 Gig	981 MB/s	575 MB/s	3 Expansion Drawers
3957-VEA	CS7	8.2x.x.xxx and Higher	16 Gig	551 MB/s	243 MB/s	6 Expansion
3957-VEA	CS8	8.2x.x.xxx and Higher	16 Gig	736 MB/s	543 MB/s	1 Expansion
3957-VEA	CS8	8.2x.x.xxx and Higher	16 Gig	758 MB/s	413 MB/s	6 Expansion
3957-VEB	CS8	8.2x.x.xxx and Higher	16 Gig	384 MB/s	224 MB/s	0 Expansion
3957-VEB	CS8	8.2x.x.xxx and Higher	16 Gig	531 MB/s	390 MB/s	1 Expansion
3957-VEB	CS8 and Higher	8.2x.x.xxx and Higher	16 Gig	911 MB/s	430 MB/s	6 Expansion Drawers

- The cache IO rates displayed are approximations derived within a test environment. Customers may experience different rates depending on their specific configurations

From Table 1, the V06 server with the CC6/CX6 cache types, produce approximately 385 MB/sec of compressed read throughput and 194 MB/s of write throughput. In comparison, the V07 server with the CC8 cache controller and three expansion drawers yield approximately 981 MB/s of read throughput and 575 MB/sec of reads. Reads tend to be faster than writes because they require less system overhead. With regards to the tapeless solutions (VEA and VEB), the older CS7 disk produces a write average of 243 MB/s with 6 expansion drawers. The VEB server with latest CS8 disks produces an average of 430 MB/s, also with 6 expansion drawers and a much higher 911 MB/s for reads. Overall, the VEB and V07 show similar IO rates, but one must also remember that with the VEB solution, the TS7720 will not be consumed by functions such as pre-migrations and recalls, thus reducing overall CPU utilization.

3.3 Tuning the TS7700:

Although the TS7700 is capable of handling many different tasks concurrently, there are certain activities that utilize more CPU than others, thus causing a bottleneck in the TS7700. As was mentioned before, single I/O responses of 8 seconds or more could occur within the TS7700. However, depending on the importance of certain tasks to others, customers have the ability to tune their environment. For customers concerned with host performance, there are some simple ways in which you can tune the TS7700 to reduce the amount of overhead experienced (refer to Table 2).

Table 2:

More in depth information available by reading IBM Virtualization Engine TS7700 Series Best Practices – Understanding, Monitoring, and Tuning the TS7700 Performance

Tasks that consume CPU / Reduces Host Performance	Tunable Parameters
AIX operating system	Increasing Memory can help
Host Read/Write from/to cache	Add # of performance increments FC 5268 and Ficon adapters installed
Cache Full (uncopied resident data)	Adding expansion drawers will reduce need for recalls from physical tape
Replicating data from one cluster to another	Immediate/Sync copies utilize more resouces then other copy modes
Remote Mounts	Accessing volumes from local cluster reduces network traffic
Copy Export	Avoid Copy Export during heavy workload periods
Pre-Migration	Limit # of premigration drives per pool
Recall	Keeping data in cache for longer periods lessens the need for recalls from tape
Reclamation	Inhibit reclaim should be used during peak periods

The tasks above are easily tunable and can significantly impact performance of the TS7700. For example, performance degradation can be a factor of the operating system. By maximizing the physical memory of the TS7700, you can reduce operating system paging activity which can increase the speed at which the machine can process data; thus decreasing the likelihood of waits and timeouts. The VO6/VEA can be upgraded from 8 GB to 16 GB of memory while the V07/VEB solutions come pre-installed with the maximum 16 GB of memory. Next, when looking at host performance, it is important to maximize the number of performance increment feature codes (FC 5268) in the TS7700. The V06 supports a maximum of 600 MB/sec while the new V07 servers can go up to 1000 MB/sec. Adding the maximum number of features results in unrestricted peak data throughput capabilities, meaning you may be able to achieve performance exceeding 1000 MB/s.

With regards to the disk cache, increasing the amount of cache and/or upgrading to new cache systems can improve the overall performance of the box. As was presented in Table 1, the newer CC8 cache drawers with expansion frames attached have much higher compressed data I/O rates than the CC6 generation. Faster disk throughput can allow tasks to complete sooner.

For TS7700s that are within a grid environment, the different copy modes can play a big role in the TPF host performance as well, as was articulated in section 2.1. Many customers prefer that volumes be copied immediately to their peer clusters; however this can create performance degradation. Immediate copies (Rewind Unload copy mode) can cause host write throttling to occur since the occupied device is made unavailable until a copy of the volume exists on the other clusters of the grid. Limiting the amount of immediate copies will allow those same resources to be utilized elsewhere. The immediate copies also hold the status of the RUN command during the replication processes, which is not ideal for TPF workloads.

Whenever possible, remote mounts between clusters in a grid should also be avoided. Remote mounts occur when the local cluster attempts to access a volume from another cluster's cache. Keeping data within the local cache reduces the need for the TS7700 to utilize its grid network to obtain data from other clusters. To reduce remote mounts, click the "Prefer local cache for fast ready mount requests" option on the Management Interface GUI. Clicking this option will cause scratch mounts to be mounted on the local cluster, thus avoiding any grid network delays during writes. Insufficient bandwidth, latency within the grid, and overall network infrastructure can further exacerbate host performance. The network infrastructure also contains internal timeout elements far exceeding the 16 second MIH timer for TPF, making it at times problematic for TPF workloads. Synchronous copies are also highly sensitive to network infrastructure and thus may be susceptible to time-outs issues.

Lastly, installing the proper number of back-end drives is also very important. The TS7740 requires an adequate number of physical drives to handle normal back-end tape processes such as copy export, pre-migrations, recalls, and reclamation. Equally important, customers should be aware of when not to use these features. For example, it is recommended that copy export and pre-migrations be limited during peak workload periods and reclamation be disabled. This will allow the TS7700 to focus on host writes and allocate more CPU to these tasks. More in depth information on best practices is available by reading the IBM Virtualization Engine TS7700 series Best Practices – Understanding, Monitoring, and Tuning the TS7700 Performance white paper, available on Techdocs at http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101465.

4.1 Conclusion

The purpose of this document was to give TPF customers a better understanding of their TPF/TS7700 environment, especially with regards to why host write time-outs and IO waits can occur. However, by utilizing the best practices mentioned within this document, the effects of these issues can either be eliminated or significantly reduced. Unfortunately, not all write timeouts are created the same. There can be situations where more information is needed to diagnose specific timeouts. In such cases where delays over 8 seconds occur and/or the TPF exits do not switch their host IO to available tape devices, it is important to contact your next level of IBM support. For delays that are 8 seconds or more a hardware PMR should be opened. TPF traces as well as TS7700 logs should be extracted at the exact time of failure.

References:

IBM® Virtualization Engine TS7700 Series Best Practices - Understanding, Monitoring and Tuning the TS7700 Performance:
http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101465

IBM Virtualization Engine™ TS7720 and TS7740 Release 1.6 through R2.1 Performance White Paper:
http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101938
IBM TPF Product Information Center:
http://publib.boulder.ibm.com/infocenter/tpfhelp/current/index.jsp

IBM Virtualization Engine TS7700 with R 2.0:
http://www.redbooks.ibm.com/redbooks/pdfs/sg247975.pdf

Disclaimers:

End of Document