CHAPTER 5

Monitoring, Best Practices, Caveats, and Troubleshooting

This chapter includes the following major topics:

InMage Resource Monitoring

InMage components are deployed at both the SP and enterprise. Depending on where the components reside, resource utilization can be monitored with a combination of various tools. As the VMDC CLSA team has done extensive work in the area of service assurance in a SP VMDC cloud, interested readers should refer to the VMDC CLSA for additional details. On the Enterprise side, if an Enterprise already deploys a comprehensive and well-managed infrastructure and systems monitoring program to enable proactive and make better infrastructure planning decisions based on historical trending and detailed usage analysis, we recommend those Enterprises to simply incorporate InMage components into their existing monitoring framework. This section does not intend to repeat previous CLSA recommendations or provide guidance on Enterprise end-to-end monitoring; instead we are focusing on specific metrics that an Enterprise or SP can gather based on our lab implementation.

Metrics such as storage IOPS, IO size, WAN bandwidth utilization, CPU usage on the primary server, CPU usage on the processing server, and RPO are all important metrics to monitor for performance and capacity planning. Most of those statistics are available directly from InMage or can be accessed through the environment that InMage connects to:

- WAN Bandwidth: InMage CX-CS server, vCenter statistics (V2V), Netflow (P2V)
- LAN Bandwidth Per Virtual/Physical Machine: InMage CX-CS server, vCenter statistics (V2V), NetFlow (P2V)
- CPU (PS): InMage CX-CS server, vCenter statistics (V2V), SNMP/SSH
- CPU (Agent): vCenter (V2V), Windows PerfMon, SNMP/SSH
- RPO: InMage
- IO: perfmon, iostat and drvutil utility from InMage

This section will focus on statistics monitoring using InMage. vCenter monitoring is well documented; refer to vSphere Performance Monitoring for monitoring details and operations. NetFlow monitoring can be deployed at key observation points such as the server access layer, fabric path domains, and WAN to gain visibility into LAN / WAN bandwidth and application performance. The Cisco NetFlow Generation Appliance (NGA) 3240 introduces a highly scalable, cost-effective architecture for cross-device flow generation in today's high-performance data centers. Refer to the 3240 Datasheet for details.

Bandwidth Monitoring

Although in different formats, LAN/WAN bandwidth reporting can be generated directly from the CXCS server or from the RX server for a particular customer/tenant. From the CX-CS UI statistics are grouped based on roles:

Network traffic statistics for all ScoutOS-based devices, the PS and CX server, are available from the CX by accessing Monitor > Network Traffic Report. Statistics are stored in RRD databases maintained by the RRDtool, available in 24 hours, week, month and year intervals. As in most RRD implementations, statistics are more granular for the 24 hour interval and less granular for older statistics. Figure 5-1 is an example of network traffic rate for an PS server.

Figure 5-1 Sample Network Traffic Rate for a PS Server

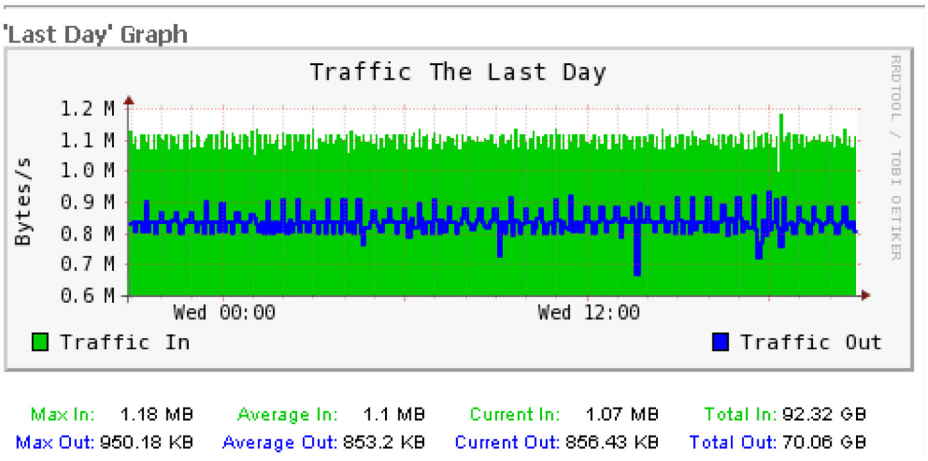

Network traffic statistics for Unified Agent-based devices, the source and MT server, are available from the CX by accessing Monitor > Bandwidth Report. Similar to ScoutOS statistics, data are stored in RRD databases maintained by the RRDtool. Daily aggregate statistics are also available. Figure 5-2 is an sample daily bandwidth report.

Figure 5-2 Sample Daily Bandwidth Report

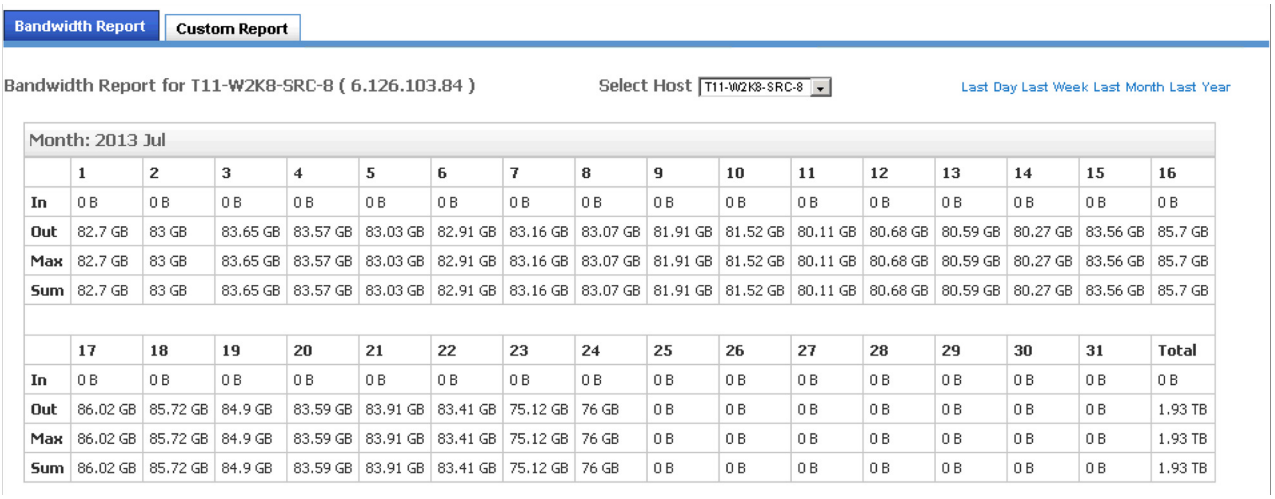

Bandwidth Report | Custom Report

Bandwidth Report for T11-W2K8-SRC-8 (6.126.103.84) Select Host T11-W2K8-SRC-8 Last Day Last Week Last Month Last Year

Month: 2013 Jul

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
In	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B
Out	82.7 GB	83 GB	83.65 GB	83.57 GB	83.03 GB	82.91 GB	83.16 GB	83.07 GB	81.91 GB	81.52 GB	80.11 GB	80.68 GB	80.59 GB	80.27 GB	83.56 GB	85.7 GB
Max	82.7 GB	83 GB	83.65 GB	83.57 GB	83.03 GB	82.91 GB	83.16 GB	83.07 GB	81.91 GB	81.52 GB	80.11 GB	80.68 GB	80.59 GB	80.27 GB	83.56 GB	85.7 GB
Sum	82.7 GB	83 GB	83.65 GB	83.57 GB	83.03 GB	82.91 GB	83.16 GB	83.07 GB	81.91 GB	81.52 GB	80.11 GB	80.68 GB	80.59 GB	80.27 GB	83.56 GB	85.7 GB

	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	Total
In	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B	0 B
Out	86.02 GB	85.72 GB	84.9 GB	83.59 GB	83.91 GB	83.41 GB	75.12 GB	76 GB	0 B	0 B	0 B	0 B	0 B	0 B	0 B	1.93 TB
Max	86.02 GB	85.72 GB	84.9 GB	83.59 GB	83.91 GB	83.41 GB	75.12 GB	76 GB	0 B	0 B	0 B	0 B	0 B	0 B	0 B	1.93 TB
Sum	86.02 GB	85.72 GB	84.9 GB	83.59 GB	83.91 GB	83.41 GB	75.12 GB	76 GB	0 B	0 B	0 B	0 B	0 B	0 B	0 B	1.93 TB

The RX UI reports only aggregate network traffic from Unified Agents based on a time interval. Traffic report for the PS and CX are not available from the RX.

Although the majority of the performance stats are not directly exportable from the GUI, fetching data directly from the RRD database is fairly simple and straightforward. The following is an example of fetching bandwidth data from Jul 24 2013 07:46:18 to 08:03:48:

```
[root@sp-t10-ps-1 052E4A5E-8195-1341-90A49C18364A0532]# pwd
/home/svsystems/052E4A5E-8195-1341-90A49C18364A0532
[root@sp-t10-ps-1 052E4A5E-8195-1341-90A49C18364A0532]# rrdtool fetch bandwidth.rrd
AVERAGE --start 1374651978 --end 1374653028
in  out
1374652200: 0.0000000000e+00 8.1258306408e+08 1374652500: 0.0000000000e+00
2.2297785464e+08 1374652800: 0.0000000000e+00 7.3103023488e+08 1374653100:
0.0000000000e+00 4.4805078912e+08
```

Customized graphs can be generated easily as well using the RRDtool:

```
[root@sp-t10-ps-1 052E4A5E-8195-1341-90A49C18364A0532]# rrdtool graph xgao.png --start
1374651978 --end 1374653028 DEF:myxgao=bandwidth.rrd:out:AVERAGE LINE2:myxgao#FF0000
```

Scout Server Health Monitoring

Scout Server statistics for CPU, memory, disk and free space are directly available from the CX-CS portal. Statistics are displayed at near real time at the CX-CS UI dashboard. Historical performance data are kept in round-robin databases (RRD) and maintained by the RRDtool similar to the bandwidth reports.

Table 5-1 Scout Server Health Statistics

Resource	Process Server	CX-CS Server
System Load	Yes	Yes
CPU Load	Yes	Yes
Memory Usage	Yes	Yes
Free Space	Yes	Yes

Table 5-1 Scout Server Health Statistics (continued)

Resource	Process Server	CX-CS Server
Disk Activity	Yes	Yes
PS Services	Yes	NA
Web Server	NA	Yes
Database Server	NA	Yes
CS Services	NA	Yes

In a single core system, your System Load/Load Average should always be below 1.0, meaning that when a process asks for CPU time it gets it without having to wait. It is important to keep in mind that PS/CS server are typically deployed with multiple cores. When monitoring such system, the rule of thumb is max load should not exceed number of cores. Use the following command to figure out the number of cores on your system:

```
[root@sp-t10-ps-1 052E4A5E-8195-1341-90A49C18364A0532]# cat /proc/cpuinfo|grep processor
    processor : 0
    processor : 1
    processor : 2
    processor : 3
```

CPU load or CPU percent is the amount of time in an interval that the system's processes were found to be active on the CPU. While it can be a good indicator of overall utilization when used in conjunction with system load, it is important to remember that CPU percent is only a snapshot of usage at the time of the measurement, this statistics alone is not a good indication of overall utilization.

Monitoring memory usage on a linux system can be tricky because RAM is used to not only store user application data, but also kernel data as well as cache/mirror data stored on the disk for fast access (Page Cache). Page Cache can consume large amount of memory in general, anytime a file is read, file data goes into memory in forms of page cache. Inode and Buffer cache are kernel data cached in memory. In a typical system it is perfectly normal to see memory usage increases linearly over time as in Figure 5-3:

Figure 5-3 Memory Usage

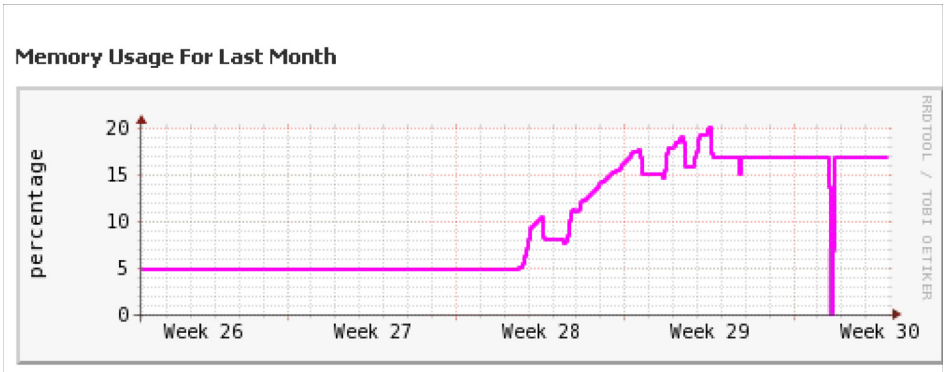

When memory usage reaches some watermark, the kernel starts to reclaim memory from the different cache described above. Swap statistics can be a good indication of overall memory health . vmstat can be used to monitor swap usage:

```
[root@sp-t10-ps-1 052E4A5E-8195-1341-90A49C18364A0532]# vmstat 5
procs memoryswapio---- --systemcpu
r b swpd free buff cache si sobibo in cs us sy id wa st
```

```
00      0 3726088  769220   717836    0     0     0     20     0     0   1   0  99   0   0
00      0 3725972  769220   717836    0     0     0     33   697   746   0   0 100   0   0
00      0 3726004  769220   717840    0     0     1     35   891   988   0   0  99   0   0
```

"si" and "so" are abbreviations for "swap in" and "swap out", the size of the swap can indicate the health of the system:

1. small "si" and "so" is normal, available memory is sufficient to deal with new allocation requests.
2. large "si" and small "so" is normal as well.
3. small "si" and large "so" indicates that when additional memory is needed, system is able to reclaim / allocate sufficient memory.
4. large "si" and "so" indicates system is "thrashing" and we are running out on memory Large "so" and "si" should be avoided. Small swap in and large swap out considered to be normal.

Dirty write is another metric to indicate if system is running low on memory. This information can be obtained from meminfo file:

```
[root@sp-t10-ps-1 052E4A5E-8195-1341-90A49C18364A0532]# more /proc/meminfo | grep
Dirty
Dirty:704 kB
[root@sp-t10-ps-1 052E4A5E-8195-1341-90A49C18364A0532]#
```

Disk activity should align with the underline storage configuration. Refer to "Implementation Best Practices" section on page 5-8.

RPO and Health Monitoring

Recovery Point Objective (RPO) is the maximum amount of data loss tolerated during disaster recovery. Depending on data change rate, WAN, and storage, RPO values can fluctuate. InMage, by default, will attempt to maintain close to zero RPO through CDP. In reality, achievable RPO is a function of available resources both at the Primary and secondary data center. Monitoring real time RPO is fairly simple using InMage CX-CS or RX. Achievable RPO is reported per volume for each virtual/physical server under protection. In addition to real time reporting, InMage also provides historical trending as part of the Health Report.

Real time RPO can be monitored directly from the CX-CS Dashboard under protection details on the **Plan Summary > Protection Details** page as shown in Figure 5-4.

Figure 5-4 RPO Protection Details

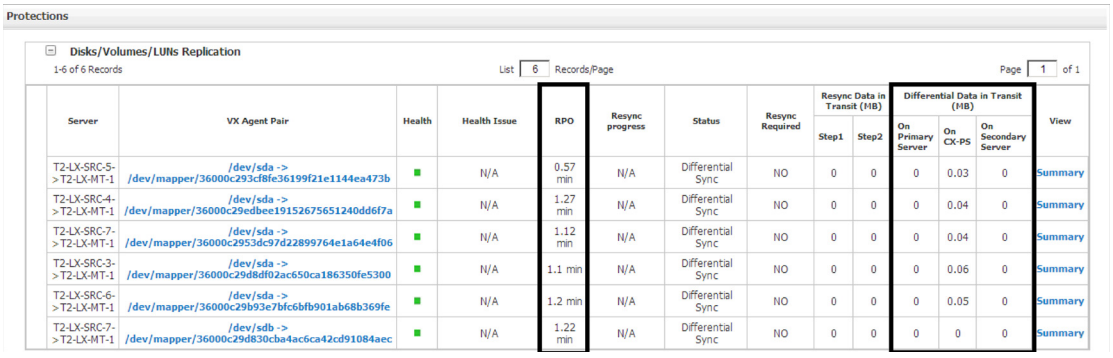

If the reported RPO is worse than the target RPO, it is important to find out where the bottleneck may reside. The "Differential Data in Transit" column in Figure 5-4 is designed to quickly identify potential bottlenecks:

- If the majority of the differential data resides on the Primary Server, the bottleneck is most likely within the primary site. Health of the Primary server, campus LAN connectivity, and Processing server health should be investigated.
- If the majority of the differential data resides on the CX-PS server, this could be an indication of WAN congestion or CX-PS server processing delay.
- If the majority of the differential data resides on the secondary server, then the health of the MT needs to be investigated. It is possible the MT couldn't process the incoming data fast enough or recovery site storage isn't sufficient to keep up with the change rate.
- If the majority of the differential data resides on both the CX-PS and secondary server, then the issue may be due to small I/O size, low compression ratio, or slow data draining at the target. Refer to "I/O Monitoring" section on page 5-7 for details.

Health Report provides a historical view of all volumes under protection. Data are sorted by date and by server. It displays health details such as data changes with and without compression, available retention window in days, RPO threshold and violation, available consistency points, and so on. Refer to Figure 5-5 for a complete list of options.

Figure 5-5 RPO Health Report

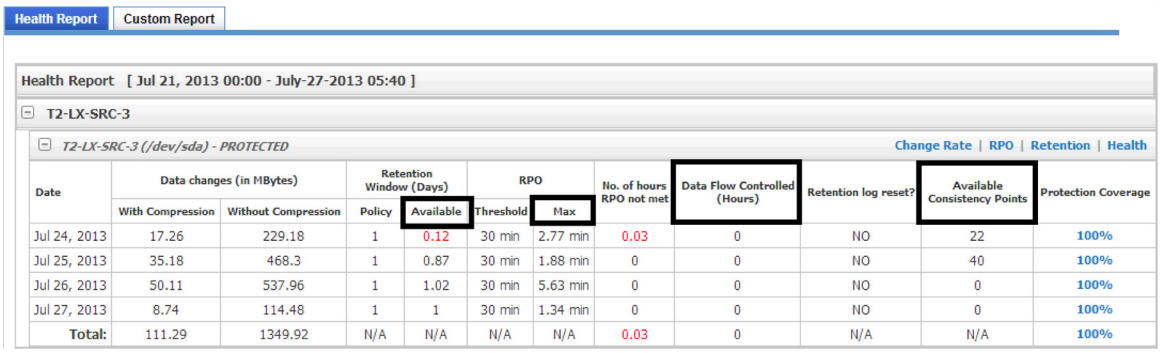

Date	Data changes (in MBytes)		Retention Window (Days)		RPO		No. of hours RPO not met	Data Flow Controlled (Hours)	Retention log reset?	Available Consistency Points	Protection Coverage
	With Compression	Without Compression	Policy	Available	Threshold	Max					
Jul 24, 2013	17.26	229.18	1	0.12	30 min	2.77 min	0.03	0	NO	22	100%
Jul 25, 2013	35.18	468.3	1	0.87	30 min	1.88 min	0	0	NO	40	100%
Jul 26, 2013	50.11	537.96	1	1.02	30 min	5.63 min	0	0	NO	0	100%
Jul 27, 2013	8.74	114.48	1	1	30 min	1.34 min	0	0	NO	0	100%
Total:	111.29	1349.92	N/A	N/A	N/A	N/A	0.03	0	N/A	N/A	100%

Key parameters to monitor are:

- **Retention Window Available:** If available value is less than the configured value, this may be an indication actual change rate exceeds the allocated retention journal space. Refer to Retention Volume Sizing, page 3-3 for additional details. For time-based retention policy, increase the size of the MT retention volume may provide a temporary relief.
- **Maximum RPO:** If Maximum RPO reached exceeds the RPO threshold, this could be an indication of WAN and storage delay / bottleneck. Refer to the previous discussion on real time RPO monitoring for details.
- **Data Flow Controlled:** This occurs when incoming change rate from the primary server exceeds the 8 GB buffer allocated by the PS server. When in data flow-controlled state, instead of sending all change blocks to the process server, the primary server will cache all changes locally in memory (up to 250 MB). Once the PS buffer utilization drops below 8 GB, the normal exchange of data between the primary and PS sever will resume. Flow controlled state is also known as thrashing state; it is intended to give the PS server sufficient time to process data it already has by backing off the primary server. This condition occurs if there's lack of sufficient WAN bandwidth or storage resources at CX-CS / MT. Refer to "Scout Server Storage and Compute Implementation" section on page 3-6 for details.

- Available Consistency Points: The available consistency point should align with retention window available. If the retention window has issues, there will be similar issues for available consistency points. Under rare circumstances where the retention window is normal, but available consistency points are low or does not exist, confirm if the proper version of VSS is installed with all required OS patches, remove third party backup software, if any, and finally confirm if sufficient disk space is available on the primary server. VSS requires at least 650MB of free disk. CX-CS log also keeps track of failure reasons, refer to "Configuration and Processing Server Logging" section on page 5-21 for additional details.

The Health Report is available from both the CX-CS UI and the RX.

I/O Monitoring

Applications will generally vary greatly in their I/O, and can change over time as the workload changes. Understanding the I/O is critical guarantee application performance. The first thing to understand is how much I/O the application under protection is doing, total I/Os per Second (IOPs), and the consistency of the I/O load.

- Is it constant, or bursty?
- If the load is bursty, what are the sizes of the largest spike, and duration / frequency of the spike?
- Does the spike occur predictably? or random?
- What is the growth trend of existing workloads under protection?
- Does IOPs growth correlate with capacity growth, or is performance growth out-pacing capacity growth?

Understanding the I/O size is also critical: some applications do small block I/O (less than 64K), and while others do large streaming I/Os, sequential write of MBs of data to disk. There can be differences even within the same application. Take MS-SQL for example: online transaction processing (OLTP) workloads tend to select a small number of rows at a time. These transfers are fairly small in size. Data warehouse applications tend to access large portions of the data at a time. These operations result in larger I/O sizes than OLTP workloads do. I/O size is important because it impacts storage bandwidth and latency. 1K IOPS with 4K IO size results in 4 MB/s storage throughput, while 64K I/O size will drive 16x the bandwidth. As I/O size increases, amount of time to write data to disk (latency) also increases.

The information in Table 5-2 needs to be obtained before onboarding a new customer and should be constantly monitored for existing customers.

Table 5-2 I/O Monitoring Information

Counter	Description
Disk Reads/sec Disk Writes/sec	Measures the number of IOPs.
Average Disk sec/ Average Disk sec/Write	Measures disk latency. Numbers vary, but here are the optimal values for averages over time: • 1 - 5 milliseconds (ms) for Log (ideally 1 ms or less on average) • 5 - 20 ms for Database Files (OLTP) (Ideally 10 ms or less on average) • Less than or equal to 25-30 ms for Data (decision support or data warehouse)

Table 5-2 I/O Monitoring Information *(continued)*

Counter	Description
Average Disk Bytes/Read Average Disk Bytes/Write	Measures the size of I/Os being issued.
Current Disk Queue Length	Displays the number of outstanding I/Os waiting to be read or written from the disk.
Disk Read Bytes/sec Disk Write Bytes/sec	Measures total disk throughput.

In a Windows-based system, the counters above are available in Perfmon. For more information about Performance Monitoring on your specific version of Windows, refer to Microsoft support sites. Linux hosts can use iostat to gather similar performance statistics. Refer to Monitoring Storage with Iostat for additional details.

Implementation Best Practices

Master Target Total VMDK

A maximum of 60 VMDKs/RDMs can be attached to a MT, of which at least three are already in use (MT OS, cache, and retention VMDK). That leaves 57 slots free. The following factors should be considered when determining the maximum number of VMDKs a single MT should protect:

- DR Drill: When performing a drill, InMage attaches all the disk snapshots of a VM to the MT. That is, if one of the VMs being protected to the MT has three disks, then you will need at least three SCSI IDs open to be able to perform a drill for that VM.
- Future Growth: Application growth; move from a medium share point deployment to a large deployment. Total allocated SCSI slots should not exceed 25 - 30 (MT OS and retention included). Therefore, thirty more are open for DR drill, if the desire is to perform a DR drill for all VMs mapped to a MT all at once, as opposed to few at a time. As a best practice, InMage recommends to not exceed forty SCSI slots on a MT. The remaining slots are reserved for disks/ VM addition or DR drill.

Master Target Cache

As a design guideline, 500MB of disk space per VMDK under protection should be reserved on the cache volume. The total size of the cache volume is a function of total number of volumes under protection:

Size of Cache Volume = (Total number of volumes under protection) * (500MB per volume)

Recovery Plan

There is one recovery plan per protection plan. To achieve RTO SLAs for large environments, you can have pre-defined recovery plans created ahead of time with the "Recover later" option and trigger them all at the same time from within vContinuum or the Scout server UI.

Scout Server

To properly tune the system, monitor the data in transit per volume and hourly data change history. If a particular server is generating a large change rate or if there's a WAN/Storage bottleneck, it is recommended to proactively increase the cache disk on the processing server. The main intent behind increasing these thresholds is to cache data on the process server instead of the sources.

Retention Policy

At the time of protection you can provide space or time based policy or both for retention. You could change this policy whenever you need from CX-GUI. Best practice for consistency interval and retention length depends on the following factors:

- Disk space available for retention.
- Number of servers and change rate from those servers assigned to retention drive.
- Nearest point you can go back and recover data using book mark or tags for application servers.

Table 5-3 InMage Storage Configuration

Change Rate	Disk Size	Type of Disk	Number of Disk	RAID
300GB	390GB	10K /15K	8	RAID 1+0
700GB	790GB	10K /15K	12	RAID 1+0
1TB	790GB	10K /15K	24	RAID 1+0

Storage Array Sizing

Profiling on the I/O characteristics of all customer work loads to determine the type of storage configuration is required at the SP's cloud. Classify the workload into the following:

- Steady State
- Customer Onboard
- DR Recovery / Drill

For each use case, characterize the worst case number of VM, read / write ratio, and average IO size. As an example:

- Workload 1: Normal Operation
  - Maximum of 250 VMs
  - Each VM requires 96 IOPS on SP-provided storage average
  - Storage is characterized as 33% read / 33% random write / 34% sequential write
  - Average read/write size is 20KB
  - 75GB of space required per VM
- Workload 2: Onboarding
  - Maximum of 3 VMs at any given time
  - Each VM requires 80 IOPS average
  - 100% sequential write
  - Average read/write size is greater than 16KB
  - 75GB of space required per VM
- Workload 3: DR Operation / Recovery
  - Maximum of 250 VMs
  - Each VM requires 96 IOPS on SP-provided storage average
  - Storage is characterized as 67% read / 33% random write
  - Average block size is 20KB
  - 75GB space required per VM

- Workload 4: Mixed Normal Operation and Recovery
  - Maximum of 250 VMs
  - Each VM requires 96 IOPS on SP-provided storage average
  - Storage is characterized as 50% read / 25% random write / 25% random read
  - Average block size is 20KB
  - 75GB space required per VM

Based on each workload characteristic and storage vendor, determine if combinations of SSD / SAS / SATA could fit into your environment. Both FlexPod and vBlock offer auto tiering, but there are some major differences in terms of implementation.

- VMAX:
  - Performance Time Window: 24/7/365 for continuous analysis
  - Data Movement Window: 24/7/365 for continuous data movement
  - Workload Analysis Period: 24/7/365 for continuous analysis
  - Initial Analysis Period: Can be configured to be between 2 hours and 4 weeks, The default is 8 hours.
  - FAST-VP Relocation Rate: 1 to 10, 786KB chunks
  - Promotion to Cache: Immediate
- VNX:
  - Performance Time Window: 24/7/365 for continuous analysis
  - Data Movement Window: Once every 24 hours
  - Workload Analysis Period: Once an hour
  - Initial Analysis Period: Once an hour
  - FAST-VP Relocation Rate: Low/Medium/High, 1GB chunk
  - Promotion to Cache: Immediate
- NetApp:
  - Performance Time Window: 24/7/365 for continuous analysis
  - Data Movement Window: 24/7/365 for continuous data movement
  - Workload Analysis Period: 24/7/365 for continuous analysis
  - Promotion to Cache: Immediate
  - Random Write < 16K IO size - SSD
  - Random Write > 16K IO Size - SAS /SATA
  - Sequential Writes - SAS / SATA

Depending on storage platform and vendor, I/O size can influence if write cache can be optimized. Data movement window can influence whether a special on-boarding strategy needs to be implemented. Check with the storage vendor to determine the optimal storage configuration.

InMage Interaction with MS SQL

Replication and backup products truncate application logs through VSS. However, the VSS writer implementation for MS SQL does not expose the ability to truncate logs. This is different from the behavior of MS Exchange VSS Writer, for example, which exposes the API to truncate logs. Due to this, InMage does not have a way to truncate MS SQL logs. For DB type of apps, the recommendation is to continue native application backup on a regular basis to maintain log sizing.

Application Consistency

In situations where Windows 2003 (Base, SP1) source machines have applications that require application quiesce, it is strongly suggested to upgrade to Windows 2003 SP2 to overcome the VSSrelated errors.

vSphere Version

To provide failover from enterprise to SP, the secondary vSphere (SP) version should be either the same or higher than the source (enterprise) vSphere server. To perform a failback from SP to

Enterprise, enterprise vSphere version should be either the same or higher than the SP vSphere. vSphere server may need to be upgraded if failback is required.

OS Recommendations

For new installations, InMage recommends:

- Secondary ESXi Platform: ESXi 5.1
- MT Platform for Windows: Windows 2008 R2 Enterprise Edition
- MT Platform for Linux: CentOS 6.2 64-bit
- CX Scout OS: CentOS 6.2 64-bit
- vContinuum: Windows 2008 R2

Protect Windows 2012 VM with ReFS Filesystem. This requires matching the 2012 MT.

SSH Client

Linux bulk agent install requires SSH access to originate from the primary server. Ensure SSH client is installed on all Linux primary servers.

Tenant Self Service

One way to achieve complete self service capability is to allow tenants to have access to the tenant-specific vContinuum Server. For security reasons you may prefer to create a non-administrator role on the vCenter for each tenant vContinuum user. The following privileges should be selected:

  - Datacenter
  - Datastore
  - Folder
  - Host
  - Network
  - Resource
  - Storage views
  - Virtual machine
  - vSphere Distributed Switch

Using vSphere RBAC, assign tenant-specific vSphere resources to the newly created tenant user/role.

Upgrade Sequence

General upgrade sequence follows:

- Upgrade the CX-CS.
- Upgrade the processing server.
- Upgrade the MT.
- Upgrade the agent on source physical or virtual server. Use the CX-CS UI instead of vContinuum to upgrade the agents.

Always refer to the release notes for the exact sequence.

Caveats

Refer to the following InMage documents:

- InMage RX Release Notes and Compatibility Matrix Documents

  http://support.inmage.net/partner/poc_blds/14_May_2013/Docs/RX/
- InMage CX Release Notes and Compatibility Matrix Documents
  http://support.inmage.net/partner/poc_blds/14_May_2013/Docs/Scout/
- InMage vContinuum Release Notes and Compatibility Matrix Documents
  http://support.inmage.net/partner/poc_blds/14_May_2013/Docs/vContinuum/

Disk Removal

vContinuum does not support disk removal. To remove a disk, first remove the VM from the protection plan and then add VM without the removed disk back into the protection plan. This operation requires a complete resync between enterprise and service provider.

Storage Over Allocation

vContinuum does not check for available capacity on the MT retention drive at the first protection. It is possible to reserve capacity beyond what is available. No alarms are generated for this condition; use the following procedure to confirm if you are running into this condition:

1. Create a new protection plan from the vContinuum.
2. Select any random Primary VM.
3. Select secondary ESX host.
4. When selecting data stores, confirm if the retention drive is available as an option.

If the retention drive is available, then the available capacity has not been over-subscribed. If the OS volume is the only available option, it is strongly recommended to manually reduce the retention size on protected VMs from the CX UI.

Protection Plan

A protection plan can only map to a single MT. No mechanisms exist to migrate a protection plan between MTs.

Continuous Disk Error

Linux MT continuously generates the following error:

```
Feb 26 11:08:32 mtarget-lx-22 multipathd: 36000c29f5decda4811ca2c34f29a9fdc: sdf -
directio checker reports path is down
Feb 26 11:08:32 mtarget-lx-22 kernel: sd 4:0:0:0: [sdf] Unhandled error code
Feb 26 11:08:32 mtarget-lx-22 kernel: sd 4:0:0:0: [sdf] Result:
hostbyte=DID_NO_CONNECT driverbyte=DRIVER_OK
Feb 26 11:08:32 mtarget-lx-22 kernel: sd 4:0:0:0: [sdf] CDB: Read(10): 28 00 00 00 00
00 00 00 08 00
```

- **Root Cause:** Each MT can have up to four SCSI controllers; each controller can have up to 15 disks. When adding a new SCSI controller, VMware requires a disk to be associated with the controller. Since disks are only added during protection time, the InMage workaround for this is to associate a dummy disk to the SCSI controller and delete the dummy disk once the controller is added. When OS attempts to acquire locks to disks that no longer exist, this causes the continuous error log. The error above is cosmetic in nature and can safely be ignored.

CX Integration with RX

If a CX-CS server is configured with dual NICS (one to communicate with the primary server and the second to communicate with the RX), use the push method instead of pull when adding the CX-CS to the RX. This is a known limitation.

VMware Tools

For virtual-to-virtual protection, updated VMware tools are required for InMage. vContinuum will not add a server to a protection plan if VM tools are not started or out of date.

Statistics Export

The network traffic rates cannot be exported from the CX-CS UI; it is only available from the dashboard as a daily, monthly and yearly graph. However, it is possible to access the raw data (RRD format) from the CS/PS Server.

Agent Status not reflected in CX Alerts and Notifications Window

CX-CS server does not raise any alarms in the CX Alerts and Notifications window when a source machines fails into bitmap mode. The machine's bit map mode cannot be recovered in the SP VPC.

Recovery Plan

1. Tenant-created recovery plans are not visible from vContinuum until execution. Once executed, recovery status can be monitored from vContinuum.
2. SP-created recovery plans (recover later) from vContinuum are intended to be SP managed and not visible to the tenant via multi-tenant portal.
3. Protect > Manage protected Files/Folders is a single source of truth. Both tenant-created and SP-created recovery plan can be viewed.

For recover later plans, outside of the Multi-Tenant Portal, there is no mechanism to view VMs included in a recovery plan. The assumption is the recovery plan should map exactly to a protection plan. If protection plan changes, a new recovery plan should be created.

Protection Plan Span Across Multiple Master Targets

All protected VMs disk have to reside in a single MT and cannot be spanned over multiple MTs.

Vmotion

Vmotion of MT from one ESX host to another is supported. Vmotion of MT storage will not work due to UUID changes.

Multi-tenant Portal

Multi-tenant portal does not support protection plan modifications. The following are InMage roadmap items:

- Add disk to a protected VM (Windows / Linux).
- Add VM (Windows / Linux) to an existing plan.
- Add physical servers (Windows / Linux) to an existing protection plan.

Portal Update

The default sync interval between the CX-CS and RX is 10 minutes. The tenant portal can lag behind the CX-CS UI by up to 10 minutes.

vContinuum Agent Install

vContinuum agent install wizard has the following limitations:

- Indicates if a server already has agents installed.
- Ability to select which network interface to use to push agent.
- NAT configuration for servers with multiple IPs.

vContinuum agent install may fail on servers with multiple NICs. Workaround is to manually install agent.

VM Name

During Windows failback, VM name configured in the primary vCenter may be modified to match the VM name in the secondary site. This issue is not observed on Linux machines.

Recovery Plan

If a disaster recovery is launched from the CX portal, depending on sequence of execution, successful recovery job maybe reported as failed. This is an UI reporting issue. To ensure accurate recovery job reporting, use the vContinuum to initiate the recover job.

Failback Readiness Check

When performing bulk failback recovery from secondary site to primary site using tag consistency, it is possible that readiness check may pass even when portion of the VM(s) have not reached tag consistent status. Workaround is to visually inspect the vContinuum log, ensure there are no errors before proceeding with live failback.

Recovery Readiness Check from the RX Portal

RX does not validate vCenter access as part of recovery readiness check.

Documented Procedure for Extending a Linux Volume/Disk

The procedure for extending a Linux volume/disk is missing in the online documentation. The procedure is documented in Appendix B Extending a Linux Volume.

vContinuum V2P Recover Does Not Complete

The vContinuum wizard V2P recovery plan fails to reboot the physical server after the recovery is complete. At this point in the recovery plan, the recovery is complete, yet the recovery script never completes. Reboot the physical server from the console to complete the V2P recover plan.

Linux P2V and V2P Recover Fails to Bring Up Ethernet Interface

The P2V and V2P recover plan for a Linux physical server fails to restore the network configuration. The interface network configuration can be recovered by entering the service network restart command. However, this is work-around is not persistent and requires reentering the service network restart command after each reboot.

V2P Failback Plan Requires Resync

The V2P failback for a Linux physical server requires a Resync immediately after volume replication is complete. Once the volume replication achieves differential sync status, use the CX UI to manually Resync the volume replications.

Recovery Plan Readiness Check

The vContinuum wizard does not fail the recovery plan when any primary VM in the recovery plan fails the readiness check. Instead, the vContiuum wizard removes any primary VMs from the recovery plan that fail the readiness check and allows the user to proceed.

In situations where a large number of primary VMs are being recovered, it is possible that the user may not be aware that some of primary VMs in the recovery plan failed the readiness check. This can happen when the primary VM that fails the readiness check is not visible within the current vContinuum window and the user must scroll down to see the error.

Troubleshooting

Troubleshooting of DR events, both successful and failed, is critical to maintain business continuity. In an environment where business service level agreements (SLAs) are tied to speed to recovery, effective troubleshooting is a integral part of accomplishing those goals. Successful troubleshooting starts with the ability to trace an event from beginning to end, by starting from the highest level and drilling down towards each component. The ability to correlate multiple events and presentation of those events is a fundamental requirement.

Configuring and setting up protection plan is a multi-step process, it starts from tenant-side discovery and reaches steady state with Differential Sync. Table 5-4 provides a brief workflow description, software components involved, and relevant log files to review when failure occurs.

Table 5-4 Troubleshooting DR Events

Sequence	Workflow Description	Software Components	Relevant Logs
1	Primary Site Discovery	vContinuum, vCenter/ESXi/ Physical Server - Primary Site	1. vContinuum a. C:\Program Files (x86)\InMage Systems\vContinuum\logs \vContinuum.log b. C:\Program Files (x86)\InMage Systems\vContinuum\Latest \MasterConfigFile.xml 2. vCenter Status (tenant) a. VMware Tools Status b. vCenter Security
2	Secondary Site Discovery	vContinuum, vCenter/ESXi - Service Provider Secondary Site	1. vContinuum a. C:\Program Files (x86)\InMage Systems\vContinuum\logs \vContinuum.log b. C:\Program vFiles (x86)\InMage Systems\vContinuum\Latest \MasterConfigFile.xml 2. vCenter Status (tenant) a. VMware Tools Status b. vCenter Security
3	Datastore Selection	vContinuum, vCenter/ESXi - Service Provider Secondary Site	1. vContinuum a. C:\Program Files (x86)\InMage Systems\vContinuum\logs \vContinuum.log b. C:\Program Files (x86)\InMage Systems\vContinuum\Latest \MasterConfigFile.xml 2. Retention drive size. Refer to "Storage Over Allocation" in Caveats. 3. vCenter Security
4	Target VM and Network Configuration	vContinuum, vCenter/ESXi - Service Provider Secondary Site	1. vContinuum a. C:\Program Files (x86)\InMage Systems\vContinuum\logs \vContinuum.log b. C:\Program Files (x86)\InMage Systems\vContinuum\Latest \MasterConfigFile.xml 2. vCenter Security
5	Readiness Check vContinuum, vCenter/ESXi/ Physical Server (Primary and Secondary Site)	vContinuum	1. vContinuum a. C:\Program Files (x86)\InMage Systems\vContinuum\logs \vContinuum.log

Table 5-4 Troubleshooting DR Events (continued)

Sequence	Workflow Description	Software Components	Relevant Logs
6	Activate Plan • Export and Import VMX • Create VMDKs • Attach VMDKs to MT • Create protection pairs	vContinuum, vCenter, ESXi/ Physical Server (Primary and Secondary Site), CS, MT	1. vContinuum a. C:\Program Files (x86)\InMage Systems\vContinuum\logs 2. ESXi/vCenter (SP) a. vCenter Log 3. CX-CS a. Monitor - CX logs - job_log_xxxx
7	Unified Agent Pulling config	CX-CS, Unified Agent	1. Unified Agent a. configuratorapitestBed.exe. Refer to InMage Tools 2. CX-CS a. Monitor - CX logs - configurator_register_host_static_info
8	Resync Step 1	UA (source & MT), CS, PS	1. UA host logs (source & MT) a. Monitor - Host Logs 2. CS/PS logs a. /home/svsystems/target host id/target volume/resync. Look for the oldest files that are not being processed If completed_hcd files are not processed, look for transfer errors at the source (Monitor - Host Logs) If completed_sync files are not processed, look for transfer errors at the destination b. Check for WAN / bpm policies (Setting - Manage Bandwidth Usage) c. /home/svsystems/transport/log/cxps.err.log d. mysql e. tmanager(Monitor -> CX logs -> volsync)

Table 5-4 Troubleshooting DR Events (continued)

Sequence	Workflow Description	Software Components	Relevant Logs

9	Resync Step 2	UA (source & MT), CS, PS	1. UA host logs (source & MT)
			a. Monitor - Host Logs
			2. CS/PS logs
			a. /home/svsystems/<target host id>/<target volume>/diffs
			`[root@ent1-scout-1 diffs]# ls`
			`completed_diff_P130197922886501968_13708331`
			`1_E130197922887125969_13708350_WE1.dat.gz`
			`completed_diff_P130197922887125969_13708350`
			`1_E130197922887593970_13708428_WE1.dat.gz`
			`completed_diff_P130197922887593970_13708428`
			`1_E130197922888061971_13708513_WE1.dat.gz monitor.txt`
			`pre_completed_diff_P130197922888061971_1370`
			`1_E130197922888997972_13708670_WE1.dat`
			b. Check for WAN / bpm policies (Setting - Manage Bandwidth Usage
			c. /home/svsystems/transport/log/cxps.err.log
			d. mysql
			e. tmanager (Monitor - CX logs - volsync)

10	Differential Sync	UA (source & MT), CS, PS	1. UA host logs (source & MT) a. Monitor -> Host Logs 2. CS/PS logs a. /home/svsystems/<target host id>/<target volume>/diffs b. Check for WAN / bpm policies (Setting -> Manage Bandwidth Usage) c. /home/svsystems/transport/log/cxps.err.log d. mysql e. tmanager (Monitor - CX logs - volsync)

Table 5-5 Troubleshooting Recovery Plan

Sequence	Workflow Step	Software Components Involved	Logs
1	VM, Recovery point (Tag) selection and readiness check	vContinuum, vCenter / ESXi (Secondary Site), MT (Secondary Site)	1. vContinuum a. C:\Program Files (x86)\InMage Systems\vContinuum\logs b. C:\Program Files (x86)\InMage Systems\vContinuum\Latest \Recovery.xml 2. MT a. cdpcli for retention information. Refer to InMage Tools
2	Target VM network configuration and recovery sequencing	vContinuum, vCenter/ESXi (SP)	1. vContinuum a. C:\Program Files (x86)\InMage Systems\vContinuum\logs
3	Activation (Rollback disks, Apply network configuration, Detach disks from MT, Power on VMs)	vContinuum, vCenter/ESXi (SP), MT, CS	1. vContinuum a. C:\Program Files (x86)\InMage Systems\vContinuum\logs 2. ESXi/vCenter (SP) a. vCenter Log 3. MT a. Monitor -> Hosts 4. CX-CS a. Monitor -> CX logs -> job_log_xxxx

This section will introduce important log files and dependencies for InMage components:

- VMware Tools, page 5-19
- vContinuum Logging, page 5-20
- Configuration and Processing Server Logging, page 5-21
- InMage Tools, page 5-22

VMware Tools

For V2V protection, VMware Tools are required for source-side discovery. To confirm if VMware Tools is running on Windows, log onto the source server and under Services confirm that the VMtool is started, as shown in Figure 5-6.

Figure 5-6 VMware Tools

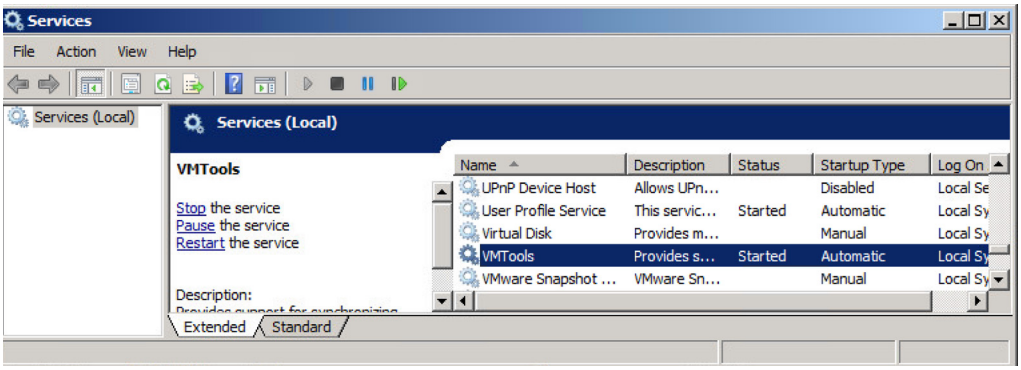

On Linux, issue:

```
ps -ef | grep vmtoolsd
```

vContinuum Logging

vContinuum is a stateless application that communicates with the CX-CS server to create and manage various protection and recovery plans. Based on user inputs, it pulls/pushes configuration changes to/from the CX-CS server. Detailed exchange between vContinuum and CX-CS server log are all logged on the vContinuum server under:

```
C:\Program Files (x86)\InMage Systems\vContinuum\logs
C:\Program Files (x86)\InMage Systems\vContinuum\Latest:
```

Important vContinuum log files are:

1. vContinuum.log - Use this log file to follow detailed events for jobs triggered from the vContinuum. The following is an example of Recovery Plan:

```
7/30/2013 10:58:53 AM Parameter grp id Task1 Initializing Recovery Plan: 7/30/2013
10:58:53 AM Name
This will initializes the Recovery Plan.It starts the EsxUtil.exe
binary for Recovery: 7/30/2013 10:58:53 AM Description
Completed: 7/30/2013 10:58:53 AM TaskStatus
7/30/2013 10:58:53 AM Logpath /home/svsystems/vcon/Demo_recovery_35790/
EsxUtil.log
7/30/2013 10:58:53 AM Parameter grp id Task2
Downloading Configuration Files: 7/30/2013 10:58:53 AM Name
The files which are going to download from CX are1.
Recovery.xml: 7/30/2013 10:58:53 AM Description
Completed: 7/30/2013 10:58:53 AM TaskStatus
7/30/2013 10:58:53 AM Logpath /home/svsystems/vcon/Demo_recovery_35790/
EsxUtil.log
7/30/2013 10:58:53 AM Parameter grp id Task3
Starting Recovery For Selected VM(s): 7/30/2013 10:58:53 AM Name
The following operations going to perform in this task:1.
Remove pairs for all the selected VMs2. Completes network
related changes for all VMs3. Deploys the source disk layout
on respective target disk(in case of windows): 7/30/2013 10:58:53 AM Description
Completed: 7/30/2013 10:58:53 AM TaskStatus
7/30/2013 10:58:53 AM Logpath /home/svsystems/vcon/Demo_recovery_35790/ EsxUtil.log
7/30/2013 10:58:53 AM Parameter grp id Task4
Powering on the recovered VM(s): 7/30/2013 10:58:53 AM Name
This will power-on all the recovered VMs1. It will detach
```

```
all the recoverd disks from MT2. Power-on the recovered VMs:
7/30/2013 10:58:53 AM Description
InProgress: 7/30/2013 10:58:53 AM TaskStatus
7/30/2013 10:58:53 AM Logpath /home/svsystems/vcon/Demo_recovery_35790/ EsxUtil.log
7/30/2013 10:59:18 AM paragroup <FunctionRequest Name="MonitorESXProtectionStatus"
Id="" include="No"><Parameter Name="HostIdentification"
Value="0451173A-C182-8D46-9C18A7E2E844E42D"/ ><Parameter Name="StepName"
Value="Recovery"/><Parameter Name="PlanId" Value="9"/></FunctionRequest>
7/30/2013 10:59:18 AM Count of parametergroup 2
```

Files in the "C:\Program Files (x86)\InMage Systems\vContinuum\Latest" are XML files that includes inventory information regarding:

- Primary Server under protection.
- Secondary Service Provider ESXi Environment.
- Detailed Storage information in Primary and Secondary site.
- Detailed network profile information in Primary and Secondary site. Figure 5-7 is an example of MasterConfigFile (detailed protection plan).

Figure 5-7 MasterConfigFile Example

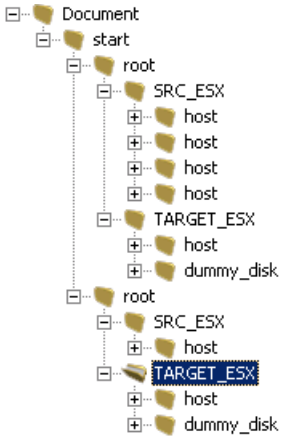

Configuration and Processing Server Logging

Logs required to troubleshoot the processing server and CS-CX can be accessed directly from the CS-CX UI. Summary of available logs can be found by navigating to Monitor > CX Logs. Table 5-6 shows the logs that are most relevant for troubleshooting:

Table 5-6 Configuration and Processing Server Logs

Log	Description
volsync	tmanager logs
Jobs_log	File replication logs. Logs associated with protection and recovery plan.
gentrends	Logs responsible for RRD trending graph.
perf_fr_job_cfg	Performance data gathering and management for gentrend.

Table 5-6 Configuration and Processing Server Logs

Log	Description
bmptrace	Bandwidth shaping module.
TrapLog	SNMP messages.

Troubleshooting directly from Monitor > CX logs can be overwhelming since it is a aggregation point for all logs. The best way to retrieve and view logs is to have some context around the failure. Having the right log file corresponding to the failure event simplifies the root cause isolation and avoids going down the wrong path due to conflicting error messages across different log files. To simplify the debugging process:

1. Always start from Monitor > File Replications.
2. Filter replication logs based on Status. Status > Failed > Search.
3. Find the corresponding failed job based on application name and then click on view details.
4. From the details windows, clock on log to open the corresponding log file relating to the failure event. Figure 5-8 is a screen capture:

Figure 5-8 View Log

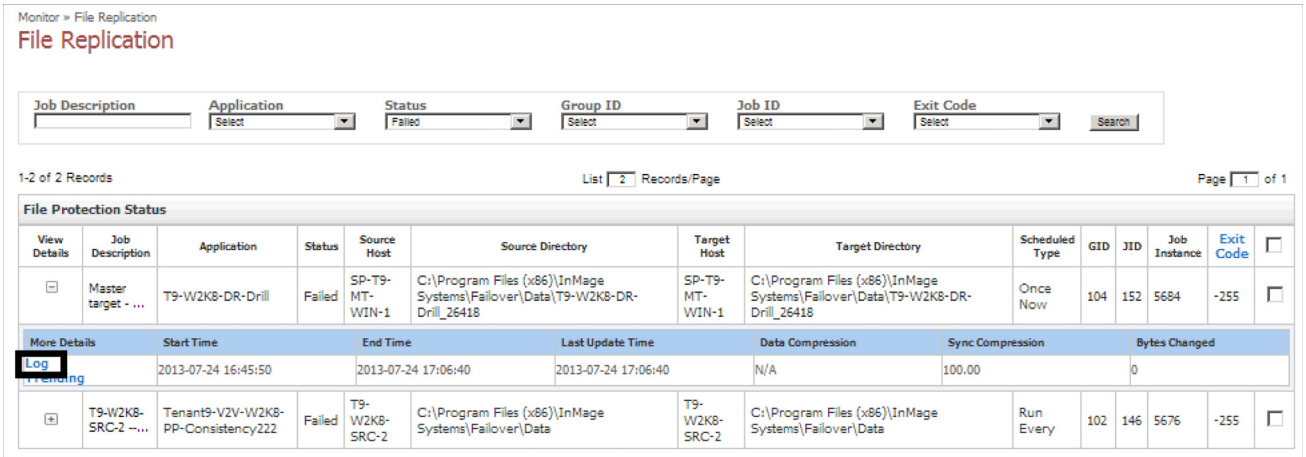

InMage Tools

The following InMage tools can be used to for troubleshooting:

1. **DrvUtil - Primary Server.** DrvUtil utility can be used to inspect and modify InMage DataTap driver settings on the primary server. For monitoring and troubleshooting purposes, Drvutil --ps and --pas options can be useful to monitor IO Size, number of dirty blocks in queue, size of the change block and so on. Other features of drvutil should not be used unless instructed by InMage support team.
2. **ConfiguratorAPITestBed - Primary Server.** This tool is most useful when troubleshooting initial data protection issues (Resync). First confirm if "dataprotection" process is running; if it's not, use configuratorapitestbed.exe to check if svagent is receiving configuration settings from the CX-CS correctly.

```
C:\Program Files (x86)\InMage Systems>ConfiguratorAPITestBed.exe --default
C:\Program Files (x86)\InMage Systems>ConfiguratorAPITestBed.exe --custom --ip
```

```
8.24.81.101 --port 80 --hostid 3F8E5834-AA0C-F246-B915D07CFB5D49CC
```

This is a Windows-based utility; there isn't a equivalent Linux version. To find out configuration settings for a Linux primary servers, use any available Windows DataTap agent and run the ConfiguratorAPITestBed command using the --hostid option where the hostid is the id of the Linux host.

3. **cdpcli - Master Target.** Use this utility on the MT to gather information regarding replication statistics, IO pattern, and protected volume.

```
c:\Program Files (x86)\InMage Systems>cdpcli.exe --listtargetvolumes
C:\ESX\3F8E5834-AA0C-F246-B915D07CFB5D49CC_C
C:\ESX\4FC26A25-683F-DC48-86F0C31180D4A5C0_C
C:\ESX\85F7759A-2EC9-9348-B8C0AD1E9A3F0331_C
C:\ESX\B6688D89-7471-EA43-AC76FBA7EDB948C6_C
C:\ESX\E08CD805-888E-944D-B60496095D3914DD_C
C:\ESX\3F8E5834-AA0C-F246-B915D07CFB5D49CC_C__SRV
C:\ESX\4FC26A25-683F-DC48-86F0C31180D4A5C0_C__SRV
C:\ESX\85F7759A-2EC9-9348-B8C0AD1E9A3F0331_C__SRV
C:\ESX\B6688D89-7471-EA43-AC76FBA7EDB948C6_C__SRV
C:\ESX\E08CD805-888E-944D-B60496095D3914DD_C__SRV
C:\ESX\3F8E5834-AA0C-F246-B915D07CFB5D49CC_K
C:\ESX\B6688D89-7471-EA43-AC76FBA7EDB948C6_K
C:\ESX\E08CD805-888E-944D-B60496095D3914DD_K
c:\Program Files (x86)\InMage Systems>cdpcli.exe --displaystatistics -- vol="C:\E
SX\3F8E5834-AA0C-F246-B915D07CFB5D49CC_C"
C:\ESX\3F8E5834-AA0C-F246-B915D07CFB5D49CC_C\ is a symbolic link to C:\ESX \3F8E
5834-AA0C-F246-B915D07CFB5D49CC_C


   ############## REPLICATION
     Target Volume Name:       STATISTICS #############

                                 C:\ESX\3F8E5834-AA0C-F246-B915D07CFB5D49CC_C
     Diffs pending in CX:        30391423
     Diffs pending in Target:    0
     Current RPO (secs):         96
     Apply rate (Bytes/sec):     12191886
     Apply time (secs):          0

c:\Program Files (x86)\InMage Systems>cdpcli.exe --showsummary --vol="C:\ESX \3F8
E5834-AA0C-F246-B915D07CFB5D49CC_C"
Database:E:\Retention_Logs\catalogue
\2460F4D5-7C71-5745-9804B2F
FB039366A\C\ESX\3F8E5834-AA0C-F246-B915D07CFB5D49CC_C\ef118abbc9\cdpv3.db
Version:3
Revision:2
Log Type:Roll-Backward
Disk Space (app):235020800 bytes
Total Data Files:10
Recovery Time Range(GMT): 2013/7/31 13:37:8:730:417:7 to
2013/7/31 13:49:13:5:533:1
c:\Program Files (x86)\InMage Systems>
c:\Program Files (x86)\InMage Systems>cdpcli.exe --iopattern --vol="C:\ESX \3F8E5
834-AA0C-F246-B915D07CFB5D49CC_C"
Io Profile:


size       %Access %Read %Random Delay Burst  Alignment Reply
     512B            0100           1000           1sector none
     512B              00             00           1sector none
     512B            0100             00           1sector none
     512B              20           1000           1sector none
      4KB           12100           1000           1sector none
      4KB             120             00           1sector none
      4KB           12100             00           1sector none
      4KB             140           1000           1sector none
```

```
   8KB          2100          1000          1sector none
   8KB            20            00          1sector none
   8KB          2100            00          1sector none
   8KB            50          1000          1sector none
  16KB          2100          1000          1sector none
  16KB            20            00          1sector none
  16KB          2100            00          1sector none
  16KB            40          1000          1sector none
  64KB          3100          1000          1sector none
  64KB            30            00          1sector none
  64KB          3100            00          1sector none
  64KB            60          1000          1sector none
 256KB          2100          1000          1sector none
 256KB            20            00          1sector none
 256KB          2100            00          1sector none
 256KB            40          1000          1sector none
   1MB          0100          1000          1sector none
   1MB            00            00          1sector none
   1MB          0100            00          1sector none
   1MB            10          1000          1sector none
   4MB          0100          1000          1sector none
   4MB            00            00          1sector none
   4MB          0100            00          1sector none
   4MB            10          1000          1sector none

c:\Program Files (x86)\InMage Systems>
```