

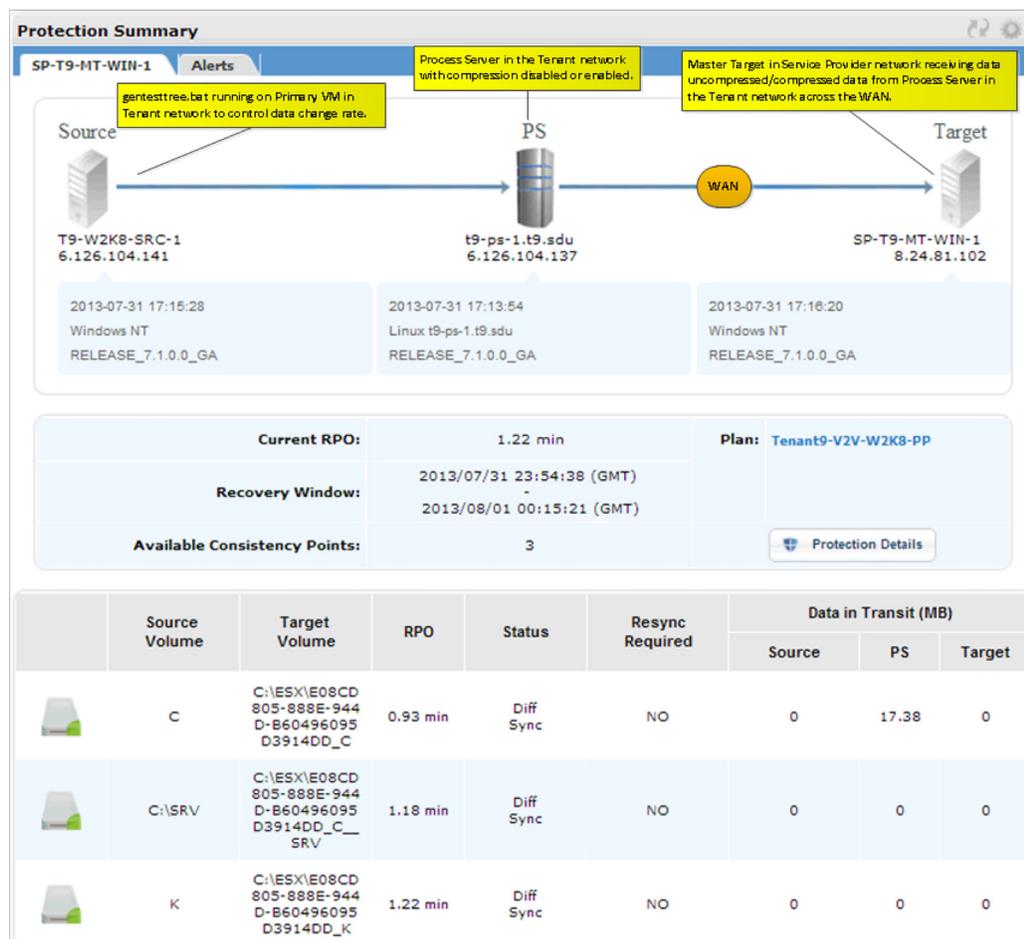


APPENDIX A

Characterization of Replication Process

As part of the testing effort for this DRaaS solution, the impact of the primary VM's data change rate on CPU utilization, network bandwidth, and disk I/O were investigated. The tests were executed on a single tenant with one Windows 2008 virtual server being protected by a MT running Windows 2008 in the SP's network. The process server, which was located in the tenant network, sent data to the MT across a simulated WAN link. [Figure A-1](#) shows this setup along with some details about the protection plan being used.

Figure A-1 Protection Summary for Characterization Test Setup



**Note**

The results presented in this section should be used to understand the general effects of data change rates and compression on the system and should be not used for planning purposes by themselves. InMage provides recommendations and analysis tools to assist with network and resource planning that should be used.

InMage has a script to generate data changes on the primary server to put a load on the disaster recovery system. The script calls an executable that writes data to a temporary folder, waits for a short duration, removes the written data, and then loops back to the start. The data writes are captured by the InMage agent on the primary server and sent to the local processing server. The processing server then compresses those changes, if compression is enabled, and the MT pulls those changes to the protection VMDK in the SP network.

In the following script, the GenerateTestTree.exe executable is called with three parameters; only the second parameter (e.g., size of the data to write each time) was changed for each iteration. This value was stepped from 1 up to 250.

gentesttree.bat Test Script

```
:loop1
@echo on
"c:\scripts\GenerateTestTree.exe" 0 10 8 C:\temp1
ping -n 5 localhost
"c:\Program Files (x86)\InMage Systems\rm.exe" -rf C:\temp1
goto loop1
```

The GenerateTestTree.exe executable has the following syntax:

```
GenerateTestTree.exe <mode=0(write)|1(verify)> <size MB> <random seed int> <dest dir>
```

At each interval of the test script, the system was allowed to settle for an hour or more and then a number of metrics was collected. Between some iterations, we observed data being cached at the process server so we forced a restart of the replication process to flush that data.

- Metrics captured from the Tenant vCenter:
 - **Primary VM Disk Write Rate (WR)**—Measured in MBps. This is the actual amount of data changes resulting from the gentesttree.bat test script.
 - **Process Server BW Input and Output**—Measured in MBps (bytes), but converted to Mbps (bits). The difference between these two values will be the amount of compression the process server can accomplish.
 - **Primary VM (Agent) CPU Utilization**—Measured in MHz. This is the amount of CPU required by the agent.
 - **Process VM (PS)CPU Utilization**—Measured in MHz. This is the amount of CPU required by the process server.
- Metrics captured from the SP vCenter:
 - **Master Target Disk Read Rate (RR)**—Measured in MBps. This is the read rate required by the MT to compare blocks.
 - **Master Target Disk Write Rate (WR)**—Measured in MBps. This is the write rate by the MT to write old block to log and new block to VMDK.
 - **Master Target CPU Utilization**—Measured in MHz, This is the amount of CPU required by the MT.

This section presents the following topics:

- [Replication with Compression Disabled, page A-3](#)
- [Replication with Compression Enabled, page A-4](#)
- [Comparison of Compression Enabled and Disabled, page A-5](#)

Replication with Compression Disabled

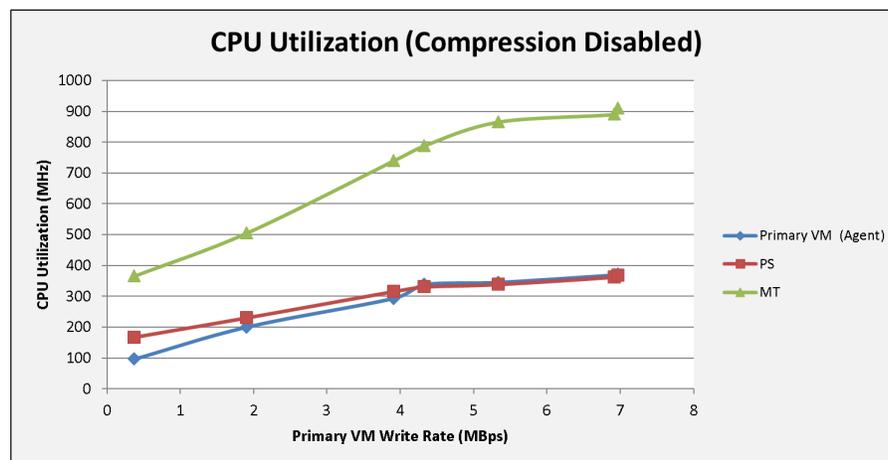
Figure A-2 shows the summary results for each of the seven (7) iterations when compression was disabled on the process server in the tenant network.

Figure A-2 Summary of Results with Compression Disabled

GenerateTestTree Params			Enterprise (Primary VM, PS)					Service Provider (MT)		
Mode	Size	Seed	Pri VM Disk WR	PS BW (In)	PS BW (Out)	Pri VM CPU	PS CPU	Disk RR	Disk WR	CPU
0	1	8	0.375	3.39968	3.883008	96	167	0.395	1.095	365
0	10	8	1.91	16.46592	16.1792	200	230	1.956	5.622	505
0	50	8	3.909	33.83296	32.768	293	315	4.027	11.784	739
0	100	8	4.33	37.35552	35.55328	338	330	6.65	12.728	788
0	150	8	5.336	41.598976	41.20576	345	338	5.05	14.96	865
0	200	8	6.922	46.465024	43.4176	369	362	5.34	15.745	890
0	250	8	6.968	46.61248	44.212224	372	367	5.399	15.991	910

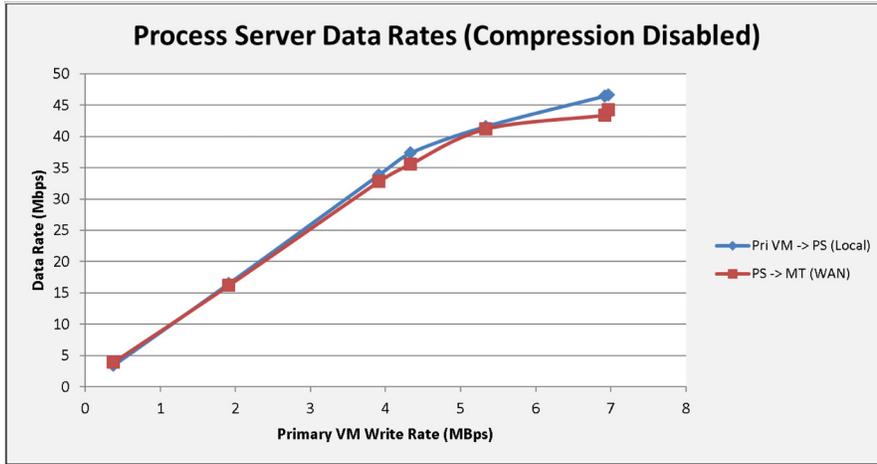
In Figure A-3, the CPU utilization for the primary VM, process server, and MT is plotted against the primary VM disk write rate. The utilization on the primary VM and process server are very close consuming 100-200 MHz for the lowest disk write load and approaches 400 MHz for the upper limit of the load script (7 MBps). On the SP side, the MT consumes more cycles for the same workload. The MT starts at a little under 400 MHz and climbs up to around 900 MHz at the highest loads.

Figure A-3 Chart of CPU Utilization vs Primary VM Disk Write Rate (Compression Disabled)



In Figure A-4, the input and output data rates for the network interface of the process server are plotted against the primary VM disk write rate. The data rate coming from the primary VM and going out to the MT are very close, which is expected since the process server is not performing any compression on the data.

Figure A-4 Chart of Bandwidth vs Primary VM Disk Write Rate (Compression Disabled)



Replication with Compression Enabled

Figure A-5 shows the summary results for each of the seven (7) iterations when compression was enabled on the process server in the tenant network.

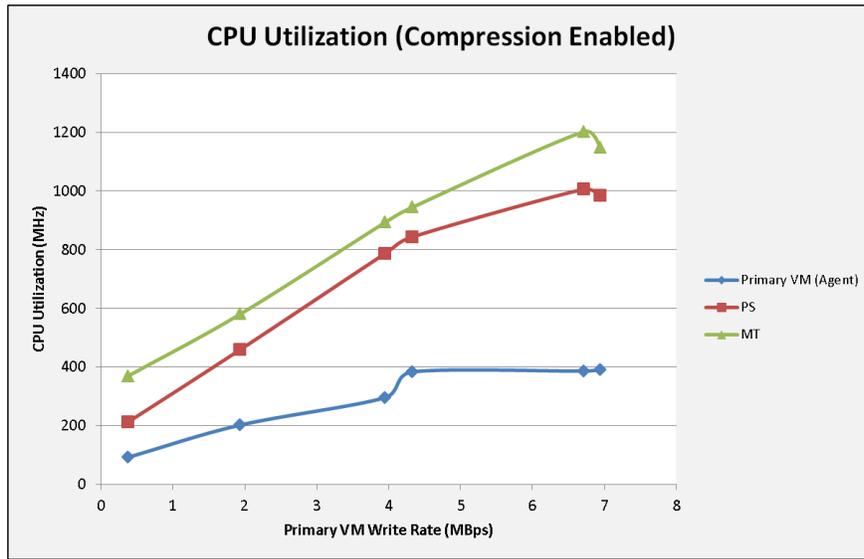
Figure A-5 Summary of Results with Compression Enabled

GenerateTestTree Params			Enterprise (Primary VM, PS)					Service Provider (MT)		
Mode	Size	Seed	Pri VM Disk WR	PS BW (In)	PS BW (Out)	Pri VM CPU	PS CPU	Disk RR	Disk WR	CPU
0	1	8	0.375	3.35872	2.58048	92	212	0.394	0.981	369
0	10	8	1.934	16.760832	12.509184	202	459	2.039	5.284	580
0	50	8	3.951	33.923072	27.189248	296	787	4.04	11.065	893
0	100	8	4.325	37.126144	28.844032	383	843	4.425	12.06	944
0	150	8	6.712	45.62944	38.985728	386	1006	5.683	16.891	1202
0	200	8	6.943	45.703168	38.928384	391	983	5.623	17.011	1148
0	250	8	7.796	59.523072	36.53632	423	1296	5.158	14.589	1116

In Figure A-6, the CPU utilization for the primary VM, process server, and MT is plotted against the primary VM disk write rate. The utilization on the primary VM consumes around 100-200 MHz for the lowest disk write load and approaches 400 MHz for the upper limit of the load script (7 MBps). As expected, this is similar to the no compression test results, since nothing has changed on the primary VM.

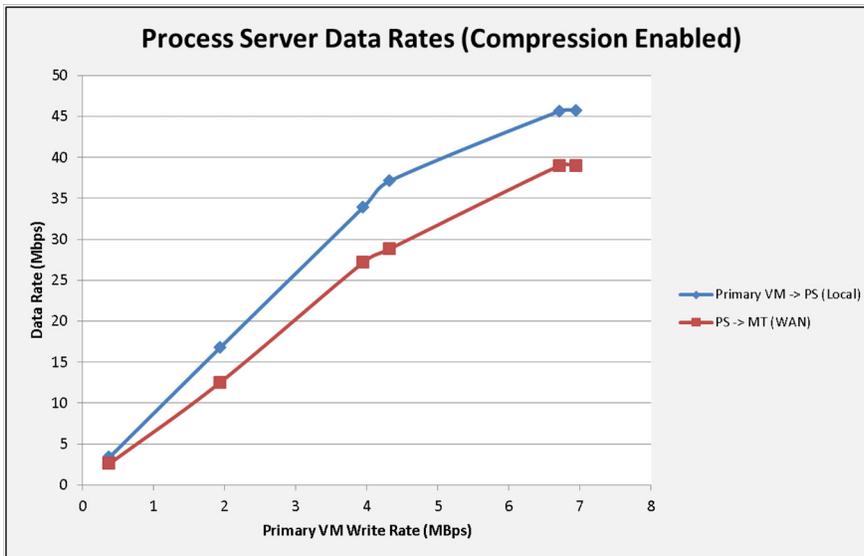
Looking at the CPU utilization for the process server, we see that the starting consumption is close to the no compression results (e.g., around 200 MHz), but almost triples in consumption near the upper limit of the load script. The highest consumption is around 1000 MHz, while the no compression results yielded around 350 MHz. The CPU utilization of the MT is about the same as the compression disabled case for the smallest loads, but consumes significantly more CPU resources at the higher loads.

Figure A-6 Chart of CPU Utilization vs Primary VM Disk Write Rate (Compression Enabled)



In [Figure A-7](#), the input and output data rates for the network interface of the process server is plotted against the primary VM disk write rate. The data difference between the two lines is the bandwidth savings on the WAN link due to the compression being applied by the process server.

Figure A-7 Chart of Bandwidth vs Primary VM Disk Write Rate (Compression Enabled)



Comparison of Compression Enabled and Disabled

In this section, results from the compression disabled and compression enabled test cases will be directly compared. In [Table A-1](#), we see that at the lower data change rates, there is a 26.95% CPU resource cost on the process server when compression is enabled. The cost approaches 200% at the higher data change rates.

Table A-1 Comparison of Process Server CPU Costs with Compression Enabled

PS CPU MHz (Compression Disabled)	PS CPU MHz (Compression Enabled)	% CPU Cost to Enable Compression on PS
167	212	26.95
230	459	99.57
315	787	149.84
330	843	155.45
338	1006	197.63
362	983	171.55

In [Table A-2](#), the CPU utilization of the MT on the SP side for both the compression disabled and compression enabled test cases is compared. The right-most column shows the % CPU resource cost when compression is used by the process server for data it sends to the MT. The MT consumes about the same amount of cycles at the lower data change rates (e.g., only about 1% CPU resource cost) when compression is enabled, but consumes 30-40% more CPU resources at the higher data rates. The maximum consumption was around 1200 MHz when compression was enabled and only around 890 MHz when compression was disabled.

Table A-2 Comparison of Master Target CPU Costs with Compression Enabled

MT CPU MHz (Compression Disabled)	MT CPU MHz (Compression Enabled)	% CPU Cost to Enable Compression on PS
365	369	1.10
505	580	14.85
739	893	20.84
788	944	19.80
865	1202	38.96
890	1148	28.99