Dell™ Compellent™ Storage Center 5.3 Data Progression with Exchange 2010

Dell Compellent Technical Solutions

November 2011

Contents

Figures

Introduction

The purpose of this document is to show results from a simulation of the data patterns of Exchange 2010 mailboxes hosted on a hosted on a Dell™ Compellent™ Storage Center clustered array enabled with Data Progression. The system included 3Gb SAS drives,two trays of 450GB 15k drives, and two trays of 1TB 7k drives. The simulation environment included 5,000 mailboxes using two 24 core servers. The storage profile used was the recommended, with daily replays. The front-end connectivity was 4Gb Fiber-Channel.

Though the mailboxes were configured to generate 2GB each, because of daily replays, the charts show a larger volume of data. Replays were used to simulate a real-world production environment, with recovery points available. The servers were configured with five databases each, which would support up to 10TB of email per server. The database and log files were moved to separate volumes.

Automated tiered storage from Dell Compellent, called Data Progression, is designed to optimize the use of storage in Exchange environments. Data Progression reduces the need for large numbers of high-performance, high-cost disks by moving frequently used data to higher performance tiers of storage while moving infrequently used data to lower cost, higher-density disks. The migration of data is done on the block level, so data within a volume can be moved based on performance characteristics.

The activity level for the mailboxes was set at 150 messages/day/user. In most cases this will be heavier than the standard customer load. The goal was to provide sufficient load to give a good picture of what Data Progression will accomplish, and to show how much data movement will occur.

Figure 1 below shows the logical configuration:

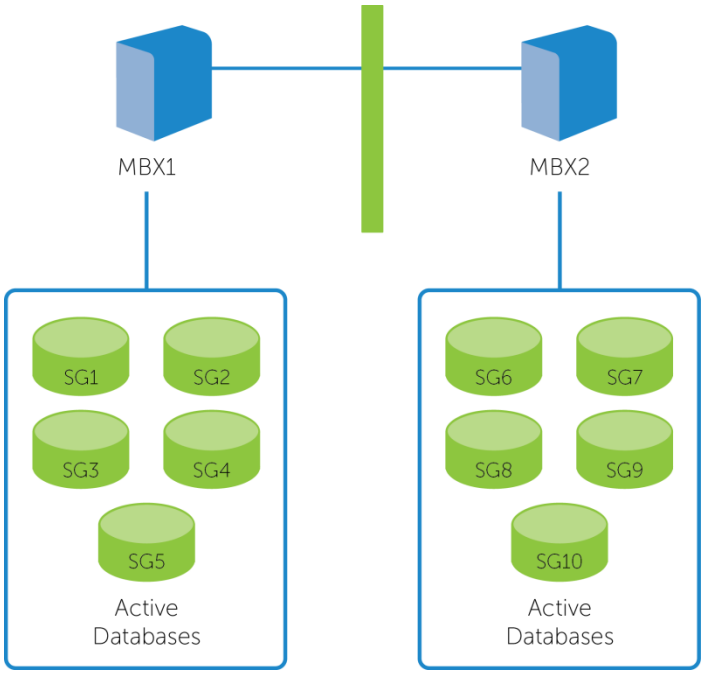

Figure 1

Background/overview

Industry studies show that as much of 80 percent of Exchange data is inactive. This means that a lot of fast, higher-cost storage is being unnecessarily utilized. This is where automated tiering technology like Data Progression can deliver immense value. The ability to automatically classify and migrate inactive portions of an Exchange Server database at the block-level onto lower-cost drives with D is causing many organizations to rethink their Exchange Server storage strategy.

Storage Center's patented Data Progression technology uses real-time intelligence to identify inactive blocks in an Exchange Server configuration and automatically migrates those blocks off an expensive storage tier to a lower-cost tier of storage based on frequency of access. This solution does not require time consuming data classification and the repetitive manual transfer of data between tiers. Each volume is configured by default with a recommended storage profile that manages the RAID configuration and provides optimal operation and performance for Exchange on the Dell Compellent Storage Center. With this configuration, all data written to each volume is written at RAID10 providing the best possible I/O performance for Exchange database and log operations.

Virtualization advantages

In Exchange Server environments using traditional SAN platforms, up to 60 percent of disk space ends up allocated but unused. A key feature in the performance of a Dell Compellent-backed Exchange Server solution is its advanced virtualization capabilities. Dell Compellent virtualizes at the disk level within the SAN, accelerating data access for Exchange Server by spreading read/write operations across all the disk drives so multiple requests can be processed in parallel. Full disk parallelization increases I/O performance and reduces hot spots in Exchange Server environments without requiring performance tuning. Each volume created spans multiple tiers as well as drives within the pool, with the option to utilize all shared drives or a specific subset of drives for each volume.

Exchange efficiencies

Based on the changes in behavior in Exchange 2010, best practices have also changed. A big part of the change involves the architecture and data structure of databases and the shift away from Single Instance Storage. More specifically, the database has been restructured and the schema flattened to produce a more efficient, sequential-like write pattern. This allows greater control over disk IO, as a mailbox is more self-contained within the database. Now when a mailbox is loaded, all indexes or views required to render a mailbox are contained within that mailbox, requiring fewer random reads, and allowing sequential reads to open a mailbox.

This change alone has helped to reduce the overall I/O of the database and reduced mailbox IO requirements by 70 percent. This reflects an IOPS per mailbox decrease of 90 percent since Exchange 2003. With all of the changes in the mailbox database structure, Exchange now has the capability to offer mailboxes that are 5GB-10GB each. The database page size has been increased to 32K (up from 4K in Exchange 2003 and 8K in Exchange 2007). This means that more data can be written to a page which in turn requires less overall reads.

To offset the storage increase, Exchange 2010 does fast page compression. This essentially offsets the increase in space required, but requires a lot of processing power to compress and decompress real-time. This is the reason for an increased number of servers. The user count per processor now matters as well.

Simulation findings

As data grows and usage patterns change, Data Progression can automatically move inactive blocks of data to a lower tier of storage (both disk class and RAID level) on-the-fly. With the recommended storage profile, active data is always written at RAID10, while any replays are initially stored at Tier 1 on RAID-5. This data eventually makes its way down the RAID levels and tiers. Figure 2 below shows the effect of Data Progression on Exchange 2010 volumes. Data will progress as it ages.

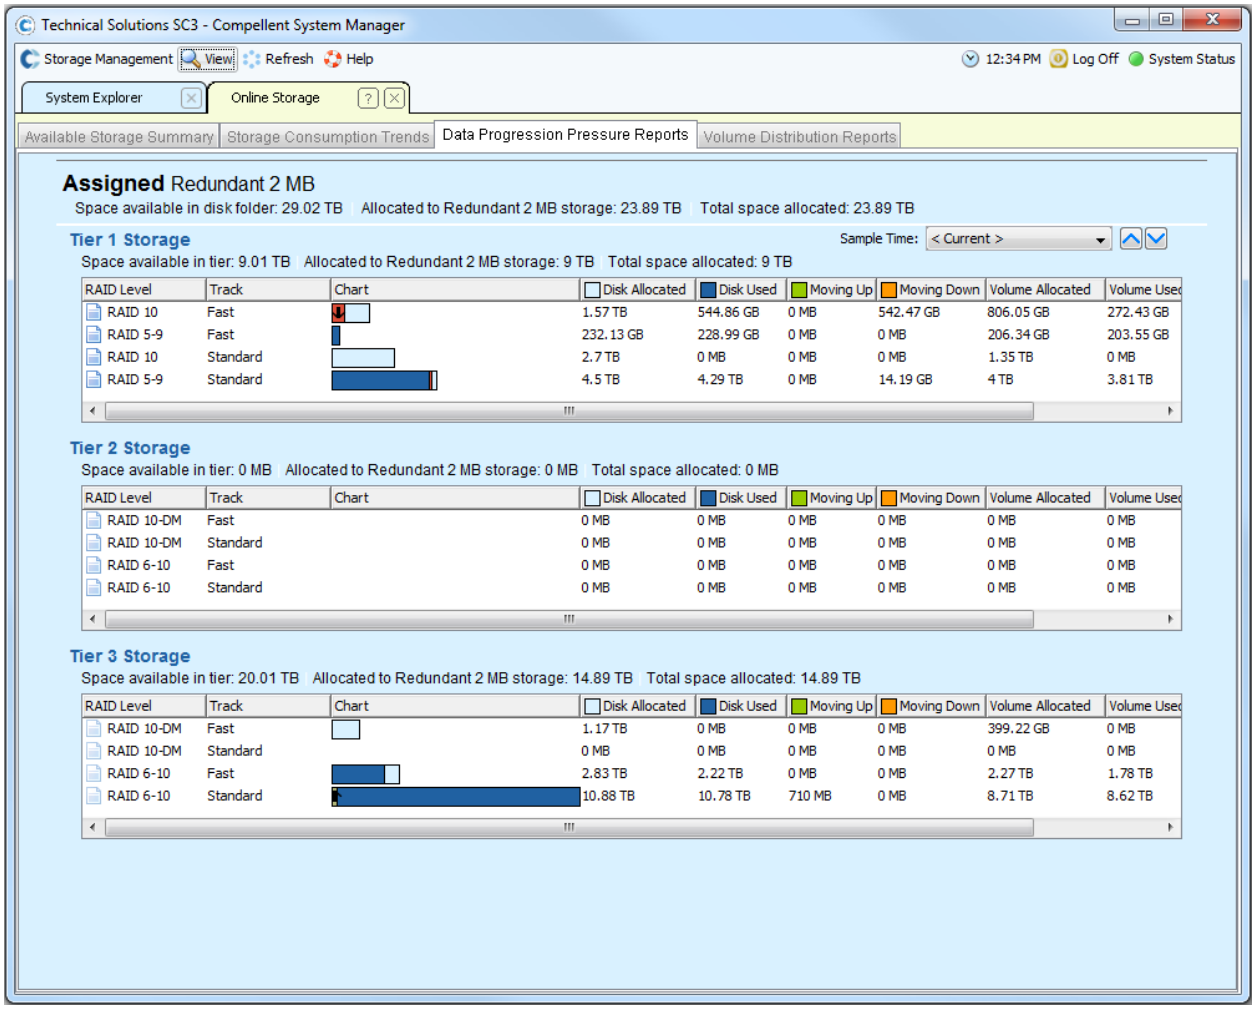

Figure 2

Leveraging a second or even third tier of storage for inactive data allows administrators to increase the number of recommended users without adding additional high-performance spindles. With Data Progression, increased user loads are a scalable benefit. As an Exchange Server configuration gets larger and user mailbox sizes increase, the amount of email that ages and is then automatically migrated to a lower tier increases. More users with more email can be managed with slower, less-expensive spindles on a second tier instead of constantly expanding by purchasing faster, more-expensive spindles.

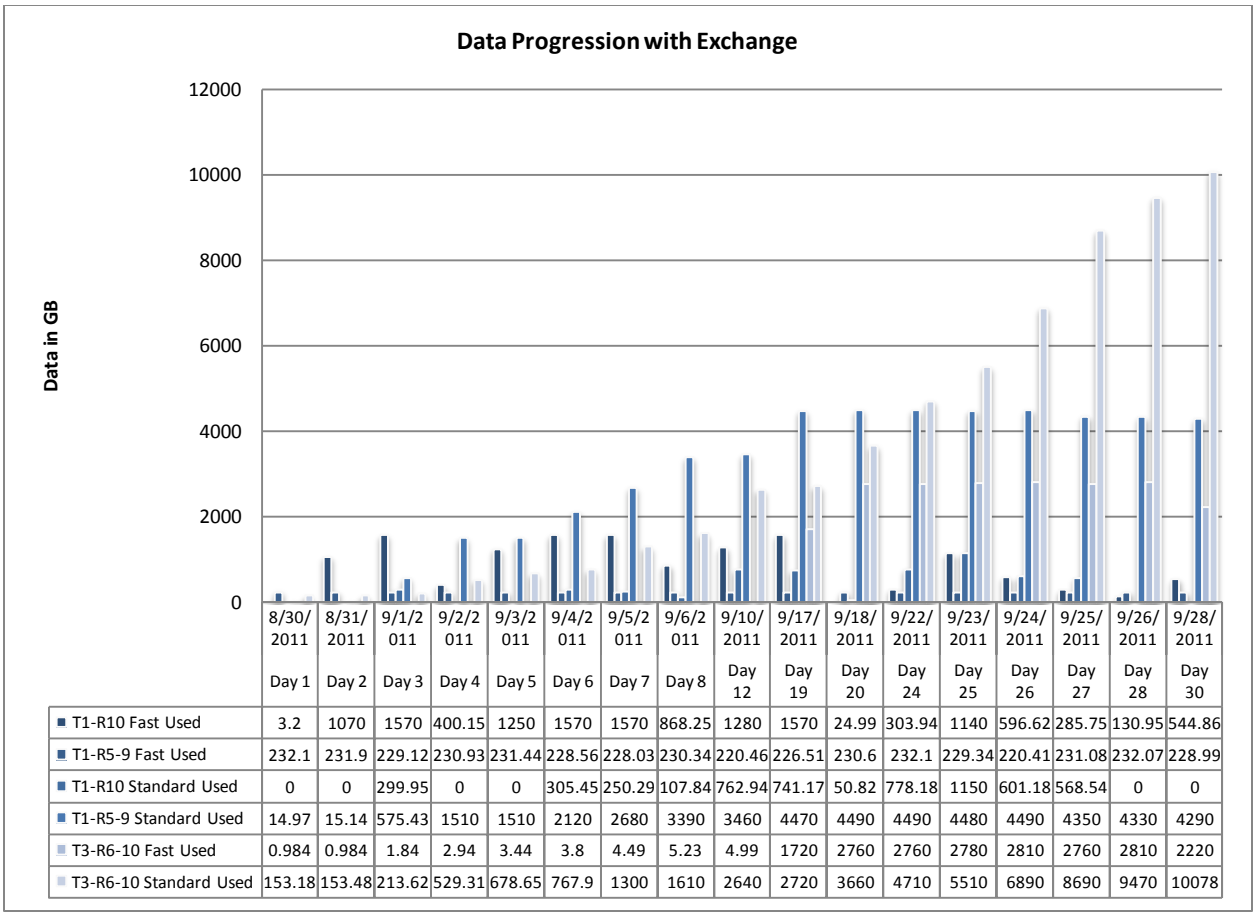

	8/30/2011 Day 1	8/31/2011 Day 2	9/1/2011 Day 3	9/2/2011 Day 4	9/3/2011 Day 5	9/4/2011 Day 6	9/5/2011 Day 7	9/6/2011 Day 8	9/10/2011 Day 12	9/17/2011 Day 19	9/18/2011 Day 20	9/22/2011 Day 24	9/23/2011 Day 25	9/24/2011 Day 26	9/25/2011 Day 27	9/26/2011 Day 28	9/28/2011 Day 30
T1-R10 Fast Used	3.2	1070	1570	400.15	1250	1570	1570	868.25	1280	1570	24.99	303.94	1140	596.62	285.75	130.95	544.86
T1-R5-9 Fast Used	232.1	231.9	229.12	230.93	231.44	228.56	228.03	230.34	220.46	226.51	230.6	232.1	229.34	220.41	231.08	232.07	228.99
T1-R10 Standard Used	0	0	299.95	0	0	305.45	250.29	107.84	762.94	741.17	50.82	778.18	1150	601.18	568.54	0	0
T1-R5-9 Standard Used	14.97	15.14	575.43	1510	1510	2120	2680	3390	3460	4470	4490	4490	4480	4490	4350	4330	4290
T3-R6-10 Fast Used	0.984	0.984	1.84	2.94	3.44	3.8	4.49	5.23	4.99	1720	2760	2760	2780	2810	2760	2810	2220
T3-R6-10 Standard Used	153.18	153.48	213.62	529.31	678.65	767.9	1300	1610	2640	2720	3660	4710	5510	6890	8690	9470	10078

Figure 3

Summary

For most companies, an Exchange Server database represents perhaps the most significant amount of stored and inactive data. Within a week after receiving an email, most users no longer access attachments and calendar items. In fact, for many, Exchange Server becomes a kind of online repository. Users don't have or take the time to delete inactive items.

For some companies, just migrating calendar data can result in significant savings. This unique approach to intelligently managing data is changing traditional views of storage management and Exchange Server storage platform implementation. Using Data Progression results in as much as a 74 percent drop in hardware costs alone. And by eliminating manual data classification and movement, Dell Compellent Data Progression cuts storage administration time in half, allowing companies to achieve the benefits of a tiered storage environment without the traditional costs or complexities.

Figure 3 shows the amount of space used and the progression of data within the tiers and RAID levels over time. Historical data moves down the tiers very similar to other types of data. The most active data stays on Tier 1, the historical data moves to Tier 3. The data will continue to progress over time.

At the end of the test runs, the majority of the data has moved down to Tier 3, (Figure 4) showing how Data Progression can save space on Tier 1 disk. By moving the historical data to Tier 3, the Tier 1 disk

will be available for other applications to use. Exchange can also continue to benefit from the Tier 1 disk available for new email.

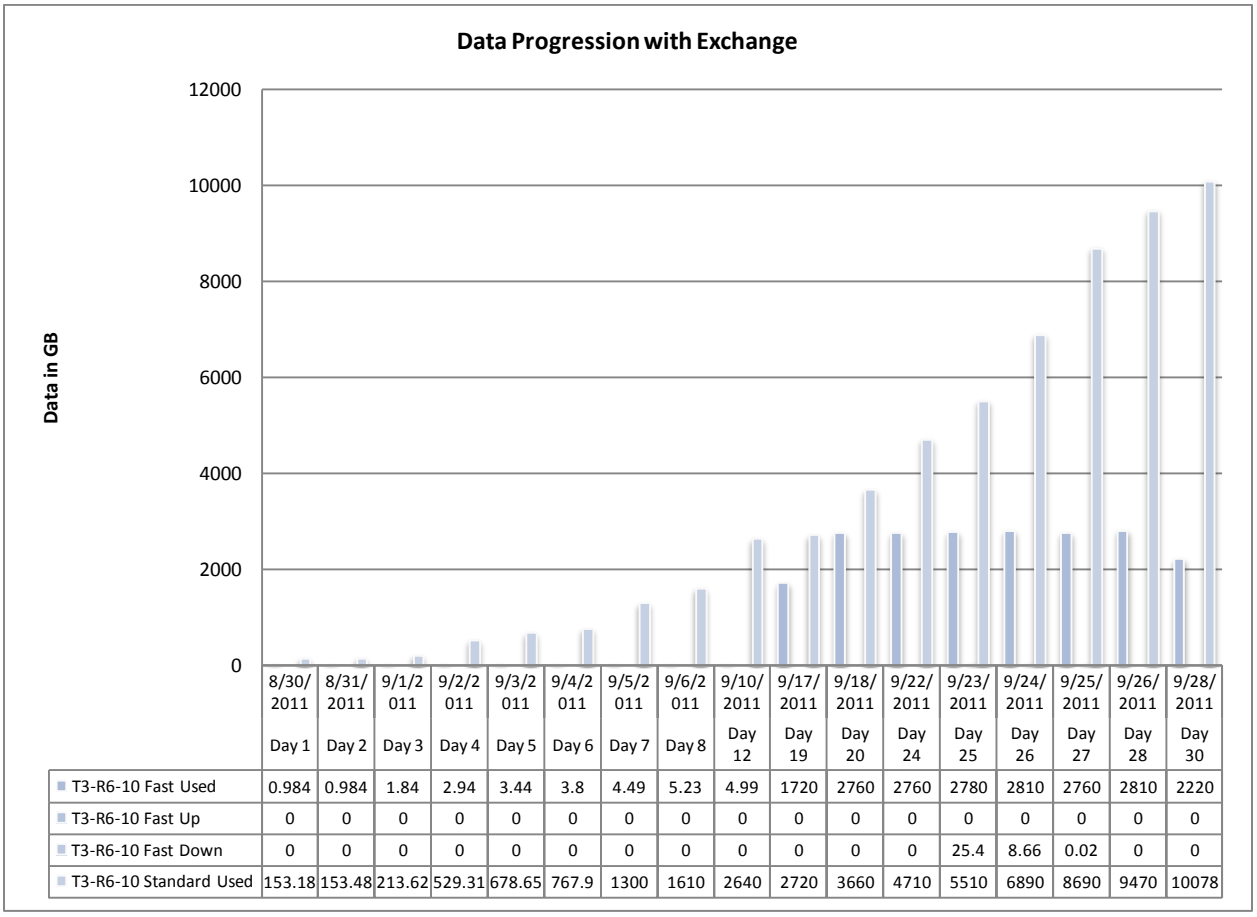

	8/30/2011 Day 1	8/31/2011 Day 2	9/1/2011 Day 3	9/2/2011 Day 4	9/3/2011 Day 5	9/4/2011 Day 6	9/5/2011 Day 7	9/6/2011 Day 8	9/10/2011 Day 12	9/17/2011 Day 19	9/18/2011 Day 20	9/22/2011 Day 24	9/23/2011 Day 25	9/24/2011 Day 26	9/25/2011 Day 27	9/26/2011 Day 28	9/28/2011 Day 30
T3-R6-10 Fast Used	0.984	0.984	1.84	2.94	3.44	3.8	4.49	5.23	4.99	1720	2760	2760	2780	2810	2760	2810	2220
T3-R6-10 Fast Up	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
T3-R6-10 Fast Down	0	0	0	0	0	0	0	0	0	0	0	0	25.4	8.66	0.02	0	0
T3-R6-10 Standard Used	153.18	153.48	213.62	529.31	678.65	767.9	1300	1610	2640	2720	3660	4710	5510	6890	8690	9470	10078

Figure 4

Appendix

Figures 5-13 in the attached appendix show the progression of the data over time as the mailboxes grow and data is generated in them. The time for progression was four weeks. These charts show Exchange in an isolated environment. Other applications running on the same Storage Center will affect data flow as well.

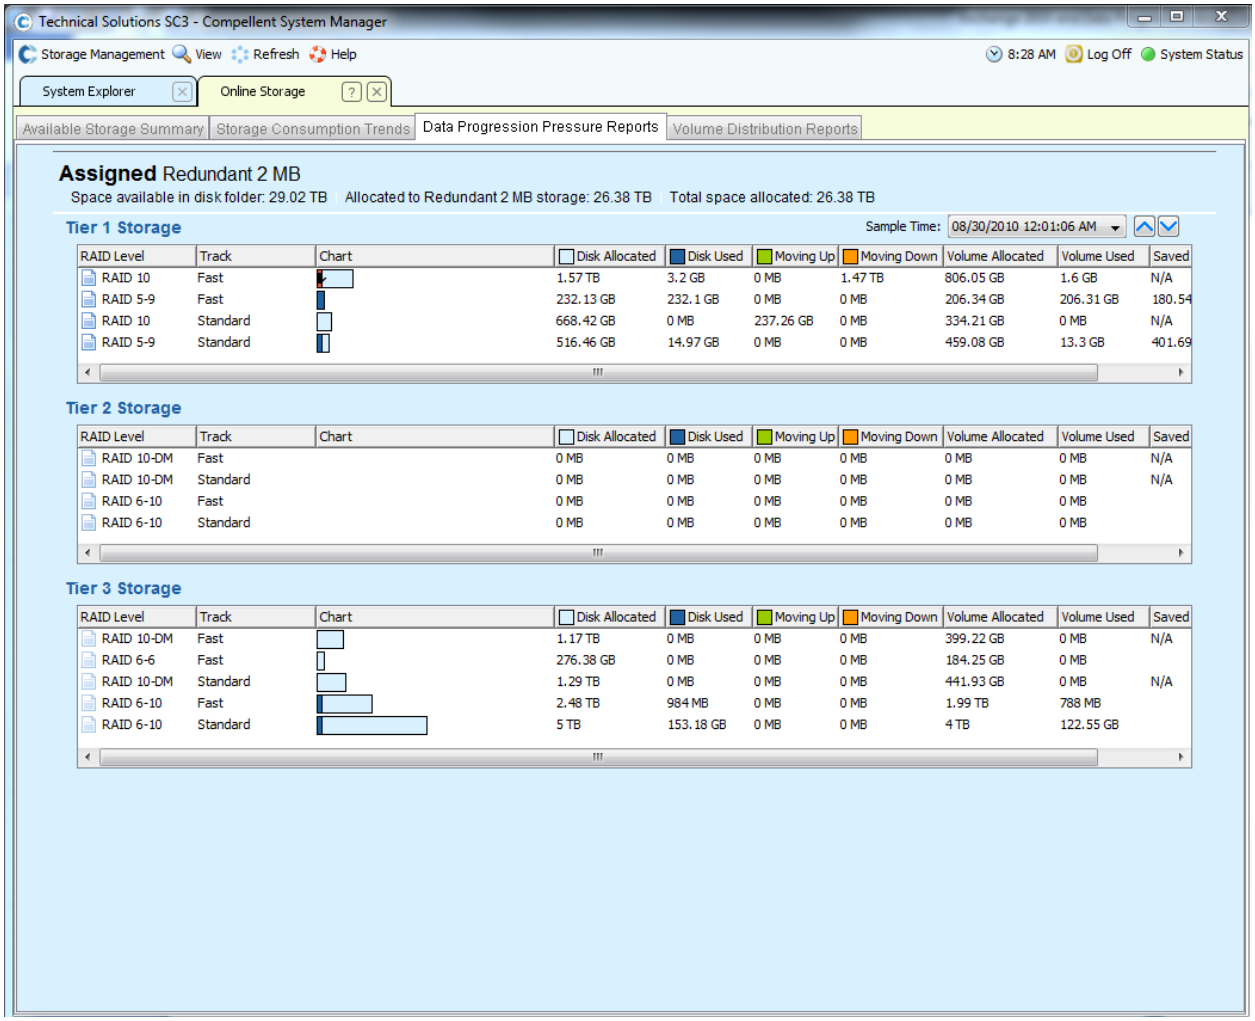
Technical Solutions SC3 - Compellent System Manager
Storage Management
View
Refresh
Help
8:28 AM
Log Off
System Status
System Explorer
Online Storage
Available Storage Summary
Storage Consumption Trends
Data Progression Pressure Reports
Volume Distribution Reports
Assigned Redundant 2 MB
Space available in disk folder: 29.02 TB
Allocated to Redundant 2 MB storage: 26.38 TB
Total space allocated: 26.38 TB
Tier 1 Storage
Sample Time: 08/30/2010 12:01:06 AM
RAID Level | Track | Chart | Disk Allocated | Disk Used | Moving Up | Moving Down | Volume Allocated | Volume Used | Saved
RAID 10 | Fast | 1.57 TB | 3.2 GB | 0 MB | 1.47 TB | 806.05 GB | 1.6 GB | N/A
RAID 5-9 | Fast | 232.13 GB | 232.1 GB | 0 MB | 0 MB | 206.34 GB | 206.31 GB | 180.54
RAID 10 | Standard | 668.42 GB | 0 MB | 237.26 GB | 0 MB | 334.21 GB | 0 MB | N/A
RAID 5-9 | Standard | 516.46 GB | 14.97 GB | 0 MB | 0 MB | 459.08 GB | 13.3 GB | 401.69
Tier 2 Storage
RAID Level | Track | Chart | Disk Allocated | Disk Used | Moving Up | Moving Down | Volume Allocated | Volume Used | Saved
RAID 10-DM | Fast | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | N/A
RAID 10-DM | Standard | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | N/A
RAID 6-10 | Fast | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB
RAID 6-10 | Standard | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB
Tier 3 Storage
RAID Level | Track | Chart | Disk Allocated | Disk Used | Moving Up | Moving Down | Volume Allocated | Volume Used | Saved
RAID 10-DM | Fast | 1.17 TB | 0 MB | 0 MB | 0 MB | 399.22 GB | 0 MB | N/A
RAID 6-6 | Fast | 276.38 GB | 0 MB | 0 MB | 0 MB | 184.25 GB | 0 MB
RAID 10-DM | Standard | 1.29 TB | 0 MB | 0 MB | 0 MB | 441.93 GB | 0 MB | N/A
RAID 6-10 | Fast | 2.48 TB | 984 MB | 0 MB | 0 MB | 1.99 TB | 788 MB
RAID 6-10 | Standard | 5 TB | 153.18 GB | 0 MB | 0 MB | 4 TB | 122.55 GB

Figure 5

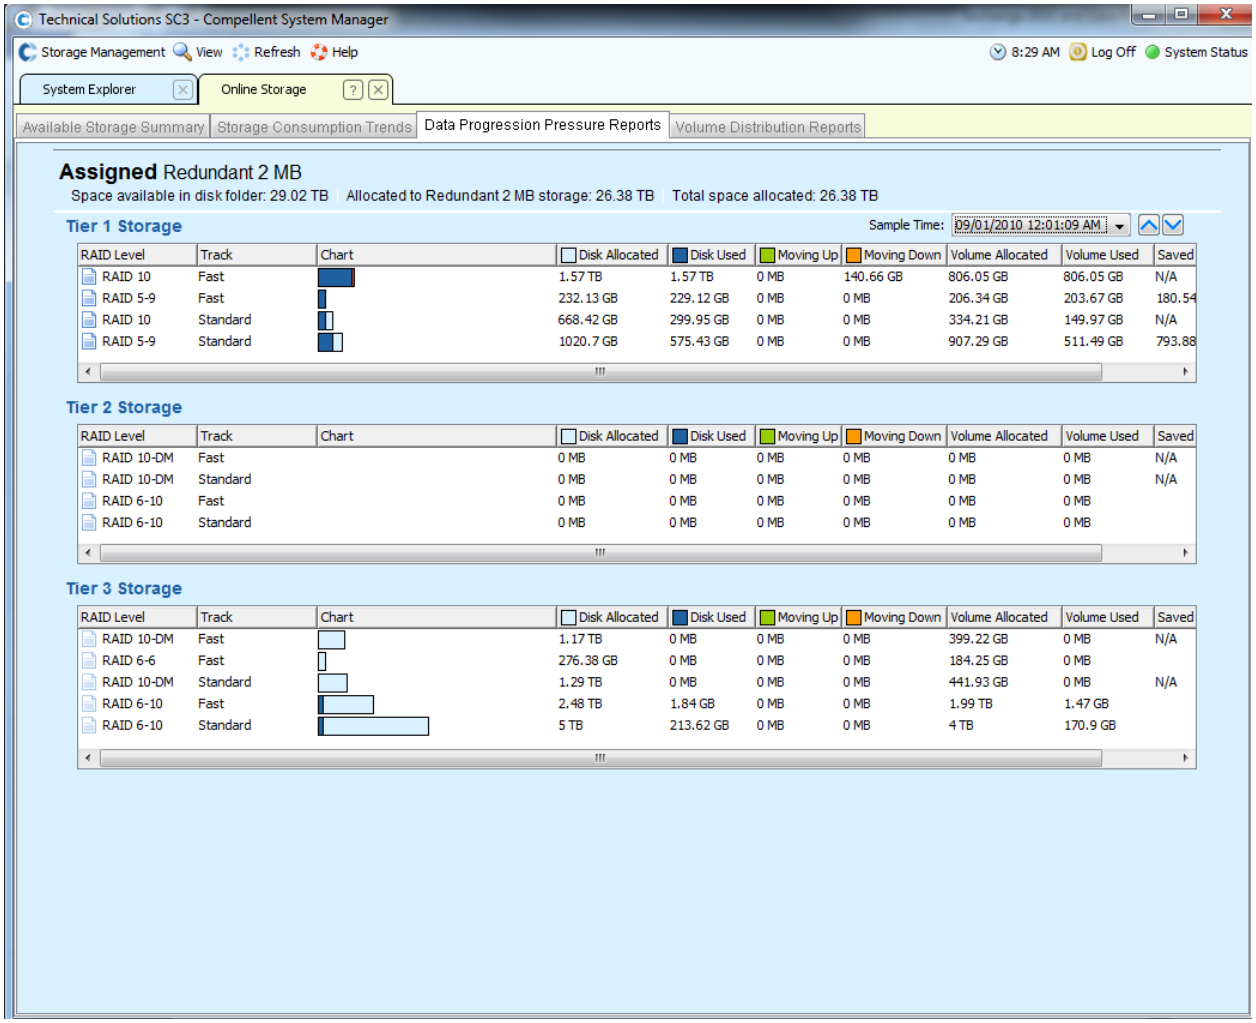
Technical Solutions SC3 - Compellent System Manager
Storage Management
View
Refresh
Help
8:29 AM
Log Off
System Status
System Explorer
Online Storage
Available Storage Summary
Storage Consumption Trends
Data Progression Pressure Reports
Volume Distribution Reports
Assigned Redundant 2 MB
Space available in disk folder: 29.02 TB
Allocated to Redundant 2 MB storage: 26.38 TB
Total space allocated: 26.38 TB
Tier 1 Storage
Sample Time: 09/01/2010 12:01:09 AM
RAID Level | Track | Chart | Disk Allocated | Disk Used | Moving Up | Moving Down | Volume Allocated | Volume Used | Saved
RAID 10 | Fast | 1.57 TB | 1.57 TB | 0 MB | 140.66 GB | 806.05 GB | 806.05 GB | N/A
RAID 5-9 | Fast | 232.13 GB | 229.12 GB | 0 MB | 0 MB | 206.34 GB | 203.67 GB | 180.54
RAID 10 | Standard | 668.42 GB | 299.95 GB | 0 MB | 0 MB | 334.21 GB | 149.97 GB | N/A
RAID 5-9 | Standard | 1020.7 GB | 575.43 GB | 0 MB | 0 MB | 907.29 GB | 511.49 GB | 793.88
Tier 2 Storage
RAID Level | Track | Chart | Disk Allocated | Disk Used | Moving Up | Moving Down | Volume Allocated | Volume Used | Saved
RAID 10-DM | Fast | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | N/A
RAID 10-DM | Standard | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | N/A
RAID 6-10 | Fast | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB
RAID 6-10 | Standard | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB
Tier 3 Storage
RAID Level | Track | Chart | Disk Allocated | Disk Used | Moving Up | Moving Down | Volume Allocated | Volume Used | Saved
RAID 10-DM | Fast | 1.17 TB | 0 MB | 0 MB | 0 MB | 399.22 GB | 0 MB | N/A
RAID 6-6 | Fast | 276.38 GB | 0 MB | 0 MB | 0 MB | 184.25 GB | 0 MB
RAID 10-DM | Standard | 1.29 TB | 0 MB | 0 MB | 0 MB | 441.93 GB | 0 MB | N/A
RAID 6-10 | Fast | 2.48 TB | 1.84 GB | 0 MB | 0 MB | 1.99 TB | 1.47 GB
RAID 6-10 | Standard | 5 TB | 213.62 GB | 0 MB | 0 MB | 4 TB | 170.9 GB

Figure 6

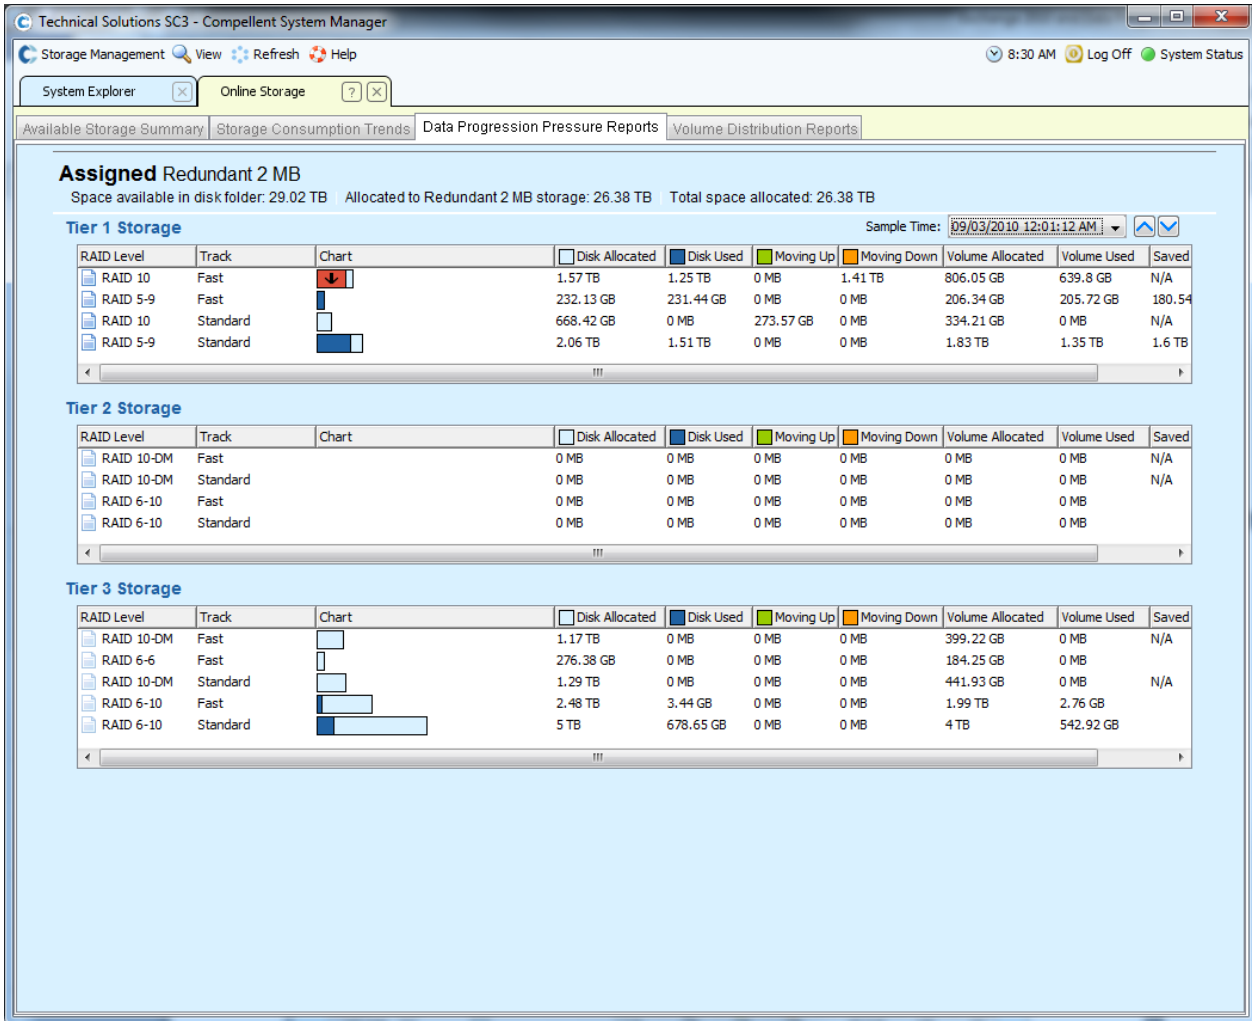
Technical Solutions SC3 - Compellent System Manager
Storage Management
View
Refresh
Help
8:30 AM
Log Off
System Status
System Explorer
Online Storage
Available Storage Summary
Storage Consumption Trends
Data Progression Pressure Reports
Volume Distribution Reports
Assigned Redundant 2 MB
Space available in disk folder: 29.02 TB
Allocated to Redundant 2 MB storage: 26.38 TB
Total space allocated: 26.38 TB
Tier 1 Storage
Sample Time: 09/03/2010 12:01:12 AM
RAID Level | Track | Chart | Disk Allocated | Disk Used | Moving Up | Moving Down | Volume Allocated | Volume Used | Saved
RAID 10 | Fast | 1.57 TB | 1.25 TB | 0 MB | 1.41 TB | 806.05 GB | 639.8 GB | N/A
RAID 5-9 | Fast | 232.13 GB | 231.44 GB | 0 MB | 0 MB | 206.34 GB | 205.72 GB | 180.54
RAID 10 | Standard | 668.42 GB | 0 MB | 273.57 GB | 0 MB | 334.21 GB | 0 MB | N/A
RAID 5-9 | Standard | 2.06 TB | 1.51 TB | 0 MB | 0 MB | 1.83 TB | 1.35 TB | 1.6 TB
Tier 2 Storage
RAID Level | Track | Chart | Disk Allocated | Disk Used | Moving Up | Moving Down | Volume Allocated | Volume Used | Saved
RAID 10-DM | Fast | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | N/A
RAID 10-DM | Standard | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | N/A
RAID 6-10 | Fast | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB
RAID 6-10 | Standard | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB
Tier 3 Storage
RAID Level | Track | Chart | Disk Allocated | Disk Used | Moving Up | Moving Down | Volume Allocated | Volume Used | Saved
RAID 10-DM | Fast | 1.17 TB | 0 MB | 0 MB | 0 MB | 399.22 GB | 0 MB | N/A
RAID 6-6 | Fast | 276.38 GB | 0 MB | 0 MB | 0 MB | 184.25 GB | 0 MB
RAID 10-DM | Standard | 1.29 TB | 0 MB | 0 MB | 0 MB | 441.93 GB | 0 MB | N/A
RAID 6-10 | Fast | 2.48 TB | 3.44 GB | 0 MB | 0 MB | 1.99 TB | 2.76 GB
RAID 6-10 | Standard | 5 TB | 678.65 GB | 0 MB | 0 MB | 4 TB | 542.92 GB

Figure 7

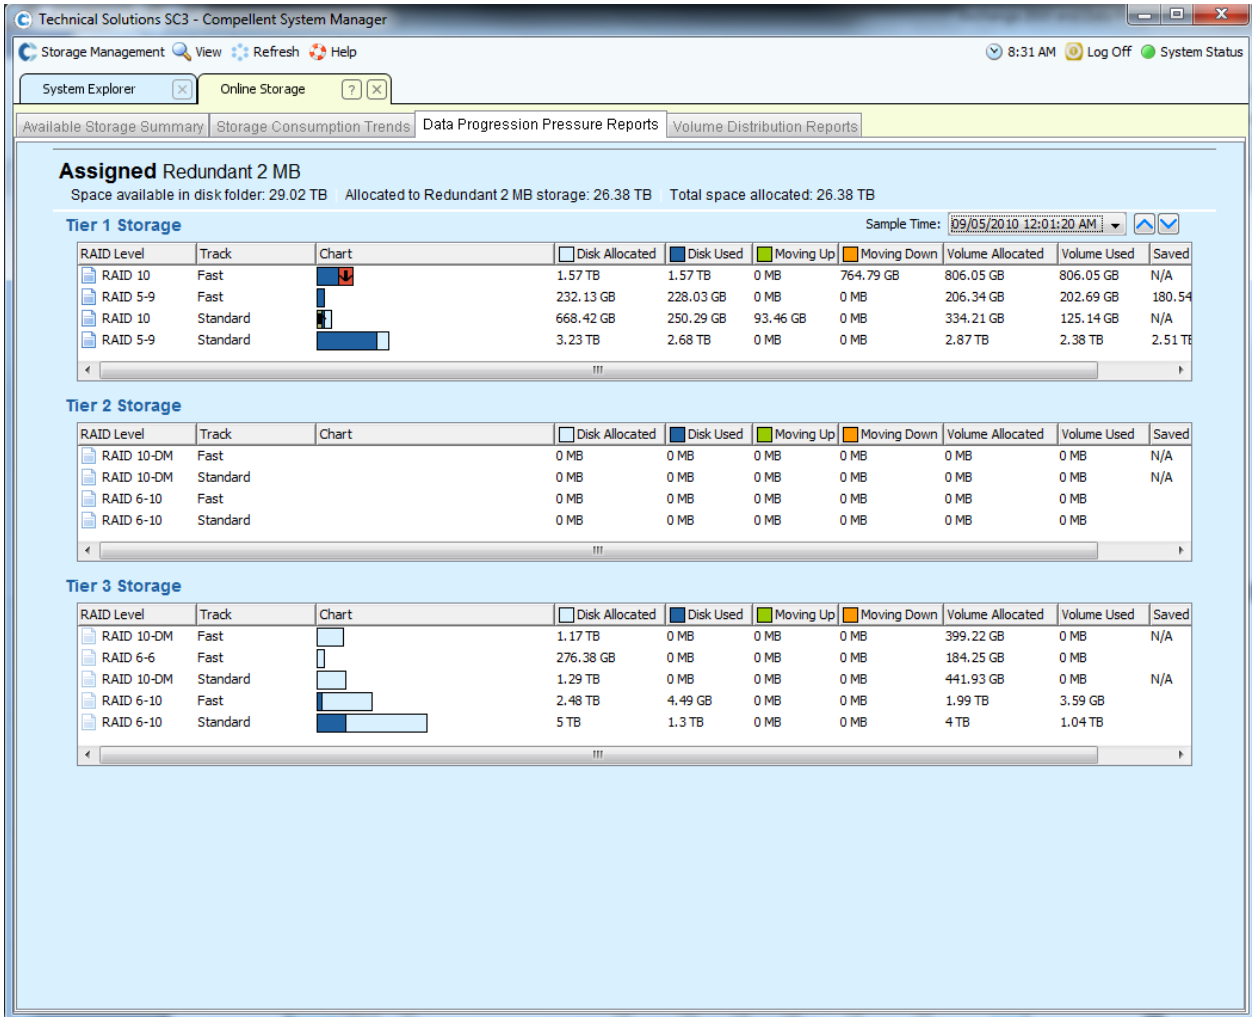
Technical Solutions SC3 - Compellent System Manager
Storage Management
View
Refresh
Help
8:31 AM
Log Off
System Status
System Explorer
Online Storage
Available Storage Summary
Storage Consumption Trends
Data Progression Pressure Reports
Volume Distribution Reports
Assigned Redundant 2 MB
Space available in disk folder: 29.02 TB
Allocated to Redundant 2 MB storage: 26.38 TB
Total space allocated: 26.38 TB
Tier 1 Storage
Sample Time: 09/05/2010 12:01:20 AM
RAID Level | Track | Chart | Disk Allocated | Disk Used | Moving Up | Moving Down | Volume Allocated | Volume Used | Saved
RAID 10 | Fast | 1.57 TB | 1.57 TB | 0 MB | 764.79 GB | 806.05 GB | 806.05 GB | N/A
RAID 5-9 | Fast | 232.13 GB | 228.03 GB | 0 MB | 0 MB | 206.34 GB | 202.69 GB | 180.54
RAID 10 | Standard | 668.42 GB | 250.29 GB | 93.46 GB | 0 MB | 334.21 GB | 125.14 GB | N/A
RAID 5-9 | Standard | 3.23 TB | 2.68 TB | 0 MB | 0 MB | 2.87 TB | 2.38 TB | 2.51 TB
Tier 2 Storage
RAID Level | Track | Chart | Disk Allocated | Disk Used | Moving Up | Moving Down | Volume Allocated | Volume Used | Saved
RAID 10-DM | Fast | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | N/A
RAID 10-DM | Standard | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | N/A
RAID 6-10 | Fast | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB
RAID 6-10 | Standard | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB
Tier 3 Storage
RAID Level | Track | Chart | Disk Allocated | Disk Used | Moving Up | Moving Down | Volume Allocated | Volume Used | Saved
RAID 10-DM | Fast | 1.17 TB | 0 MB | 0 MB | 0 MB | 399.22 GB | 0 MB | N/A
RAID 6-6 | Fast | 276.38 GB | 0 MB | 0 MB | 0 MB | 184.25 GB | 0 MB
RAID 10-DM | Standard | 1.29 TB | 0 MB | 0 MB | 0 MB | 441.93 GB | 0 MB | N/A
RAID 6-10 | Fast | 2.48 TB | 4.49 GB | 0 MB | 0 MB | 1.99 TB | 3.59 GB
RAID 6-10 | Standard | 5 TB | 1.3 TB | 0 MB | 0 MB | 4 TB | 1.04 TB

Figure 8

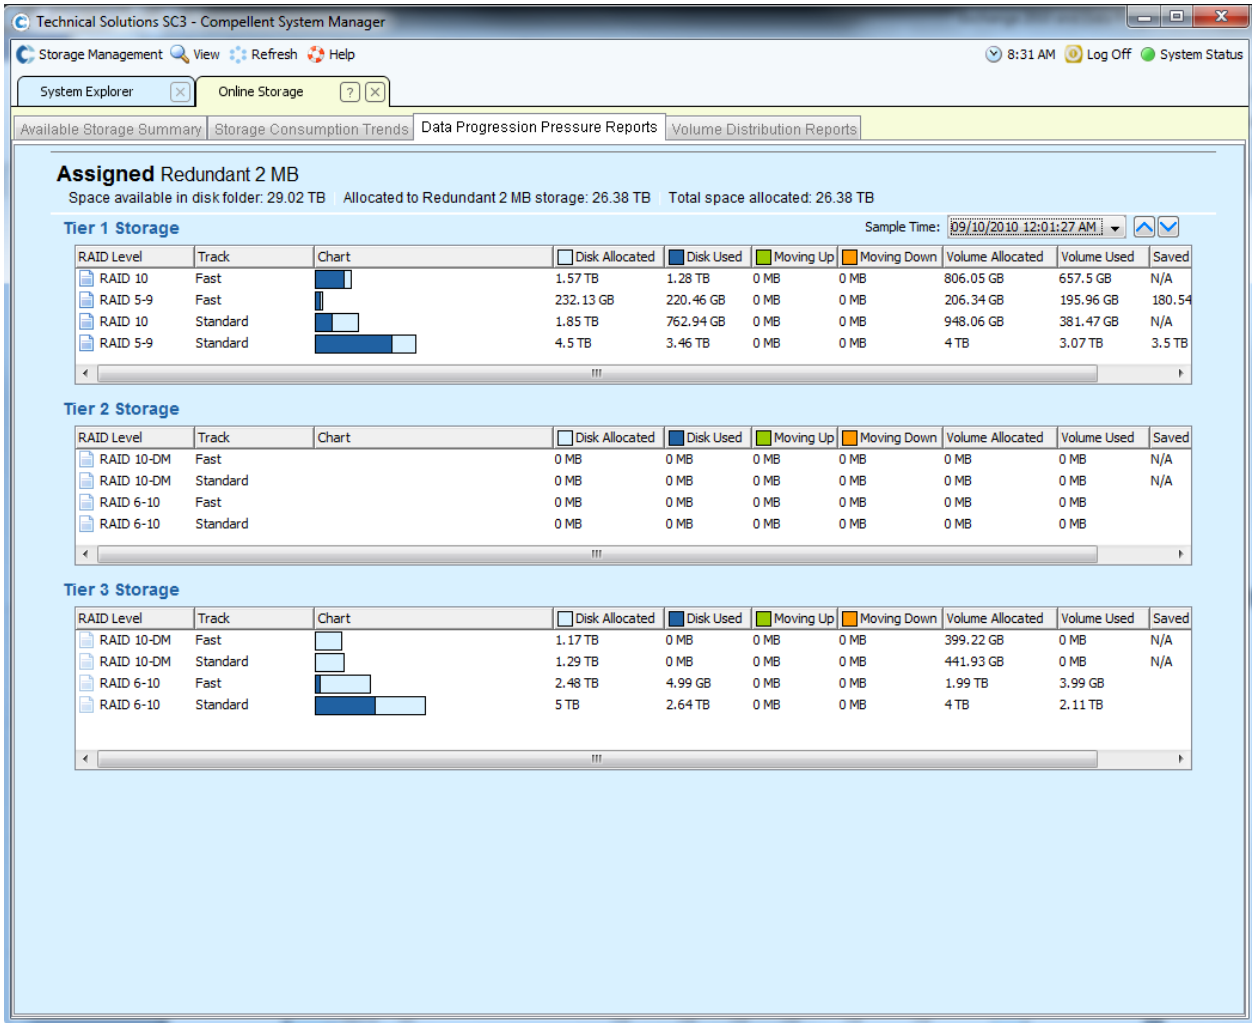
Technical Solutions SC3 - Compellent System Manager
Storage Management
View
Refresh
Help
8:31 AM
Log Off
System Status
System Explorer
Online Storage
Available Storage Summary
Storage Consumption Trends
Data Progression Pressure Reports
Volume Distribution Reports
Assigned Redundant 2 MB
Space available in disk folder: 29.02 TB
Allocated to Redundant 2 MB storage: 26.38 TB
Total space allocated: 26.38 TB
Tier 1 Storage
Sample Time: 09/10/2010 12:01:27 AM
RAID Level | Track | Chart | Disk Allocated | Disk Used | Moving Up | Moving Down | Volume Allocated | Volume Used | Saved
RAID 10 | Fast | 1.57 TB | 1.28 TB | 0 MB | 0 MB | 806.05 GB | 657.5 GB | N/A
RAID 5-9 | Fast | 232.13 GB | 220.46 GB | 0 MB | 0 MB | 206.34 GB | 195.96 GB | 180.54
RAID 10 | Standard | 1.85 TB | 762.94 GB | 0 MB | 0 MB | 948.06 GB | 381.47 GB | N/A
RAID 5-9 | Standard | 4.5 TB | 3.46 TB | 0 MB | 0 MB | 4 TB | 3.07 TB | 3.5 TB
Tier 2 Storage
RAID Level | Track | Chart | Disk Allocated | Disk Used | Moving Up | Moving Down | Volume Allocated | Volume Used | Saved
RAID 10-DM | Fast | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | N/A
RAID 10-DM | Standard | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | N/A
RAID 6-10 | Fast | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB
RAID 6-10 | Standard | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB
Tier 3 Storage
RAID Level | Track | Chart | Disk Allocated | Disk Used | Moving Up | Moving Down | Volume Allocated | Volume Used | Saved
RAID 10-DM | Fast | 1.17 TB | 0 MB | 0 MB | 0 MB | 399.22 GB | 0 MB | N/A
RAID 10-DM | Standard | 1.29 TB | 0 MB | 0 MB | 0 MB | 441.93 GB | 0 MB | N/A
RAID 6-10 | Fast | 2.48 TB | 4.99 GB | 0 MB | 0 MB | 1.99 TB | 3.99 GB
RAID 6-10 | Standard | 5 TB | 2.64 TB | 0 MB | 0 MB | 4 TB | 2.11 TB

Figure 9

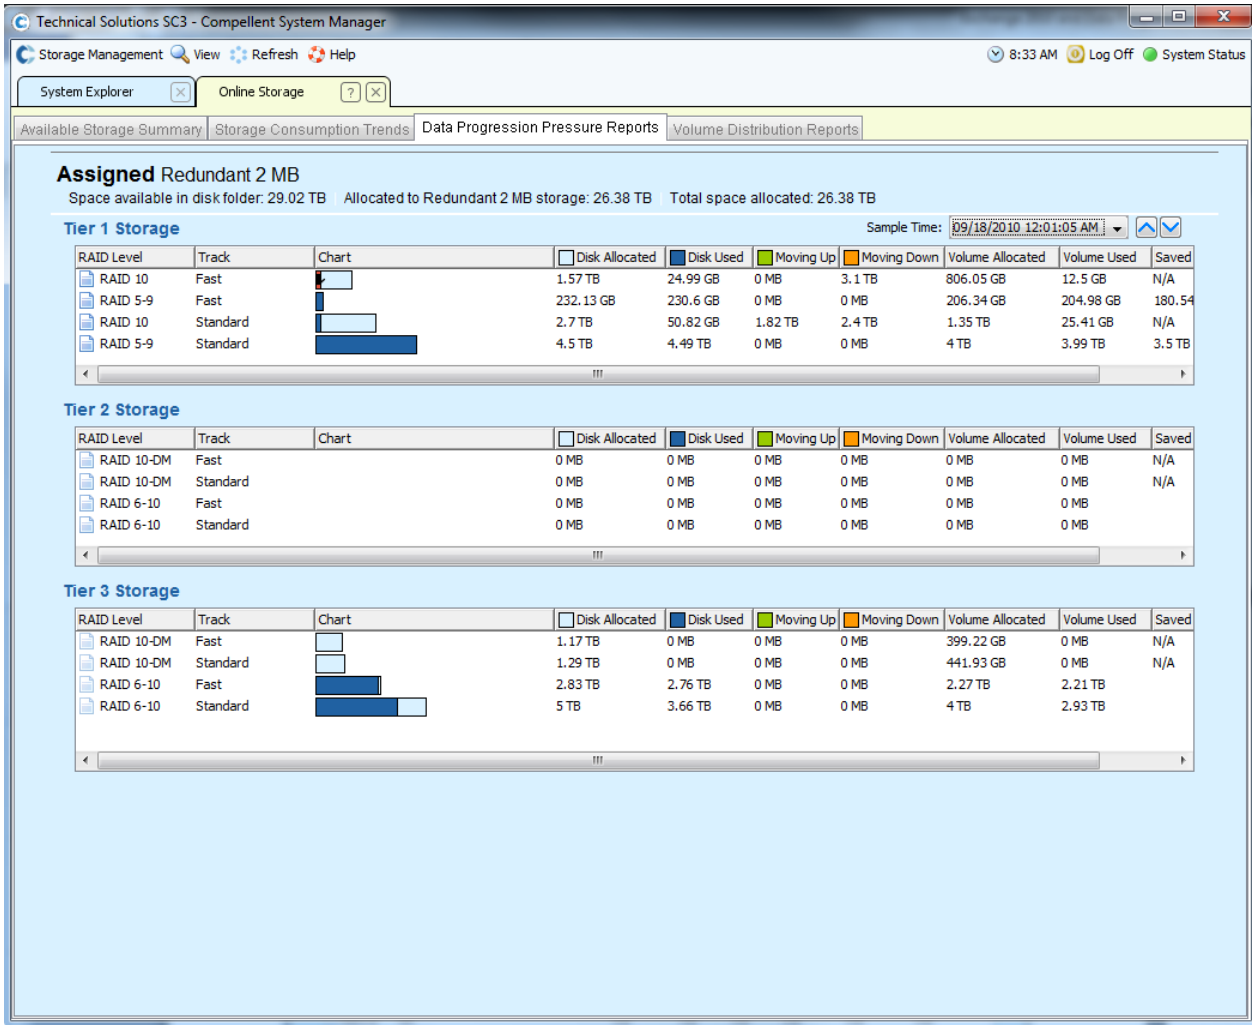
Technical Solutions SC3 - Compellent System Manager
Storage Management
View
Refresh
Help
8:33 AM
Log Off
System Status
System Explorer
Online Storage
Available Storage Summary
Storage Consumption Trends
Data Progression Pressure Reports
Volume Distribution Reports
Assigned Redundant 2 MB
Space available in disk folder: 29.02 TB
Allocated to Redundant 2 MB storage: 26.38 TB
Total space allocated: 26.38 TB
Tier 1 Storage
Sample Time: 09/18/2010 12:01:05 AM
RAID Level | Track | Chart | Disk Allocated | Disk Used | Moving Up | Moving Down | Volume Allocated | Volume Used | Saved
RAID 10 | Fast | 1.57 TB | 24.99 GB | 0 MB | 3.1 TB | 806.05 GB | 12.5 GB | N/A
RAID 5-9 | Fast | 232.13 GB | 230.6 GB | 0 MB | 0 MB | 206.34 GB | 204.98 GB | 180.54
RAID 10 | Standard | 2.7 TB | 50.82 GB | 1.82 TB | 2.4 TB | 1.35 TB | 25.41 GB | N/A
RAID 5-9 | Standard | 4.5 TB | 4.49 TB | 0 MB | 0 MB | 4 TB | 3.99 TB | 3.5 TB
Tier 2 Storage
RAID Level | Track | Chart | Disk Allocated | Disk Used | Moving Up | Moving Down | Volume Allocated | Volume Used | Saved
RAID 10-DM | Fast | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | N/A
RAID 10-DM | Standard | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | N/A
RAID 6-10 | Fast | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB
RAID 6-10 | Standard | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB
Tier 3 Storage
RAID Level | Track | Chart | Disk Allocated | Disk Used | Moving Up | Moving Down | Volume Allocated | Volume Used | Saved
RAID 10-DM | Fast | 1.17 TB | 0 MB | 0 MB | 0 MB | 399.22 GB | 0 MB | N/A
RAID 10-DM | Standard | 1.29 TB | 0 MB | 0 MB | 0 MB | 441.93 GB | 0 MB | N/A
RAID 6-10 | Fast | 2.83 TB | 2.76 TB | 0 MB | 0 MB | 2.27 TB | 2.21 TB
RAID 6-10 | Standard | 5 TB | 3.66 TB | 0 MB | 0 MB | 4 TB | 2.93 TB

Figure 10

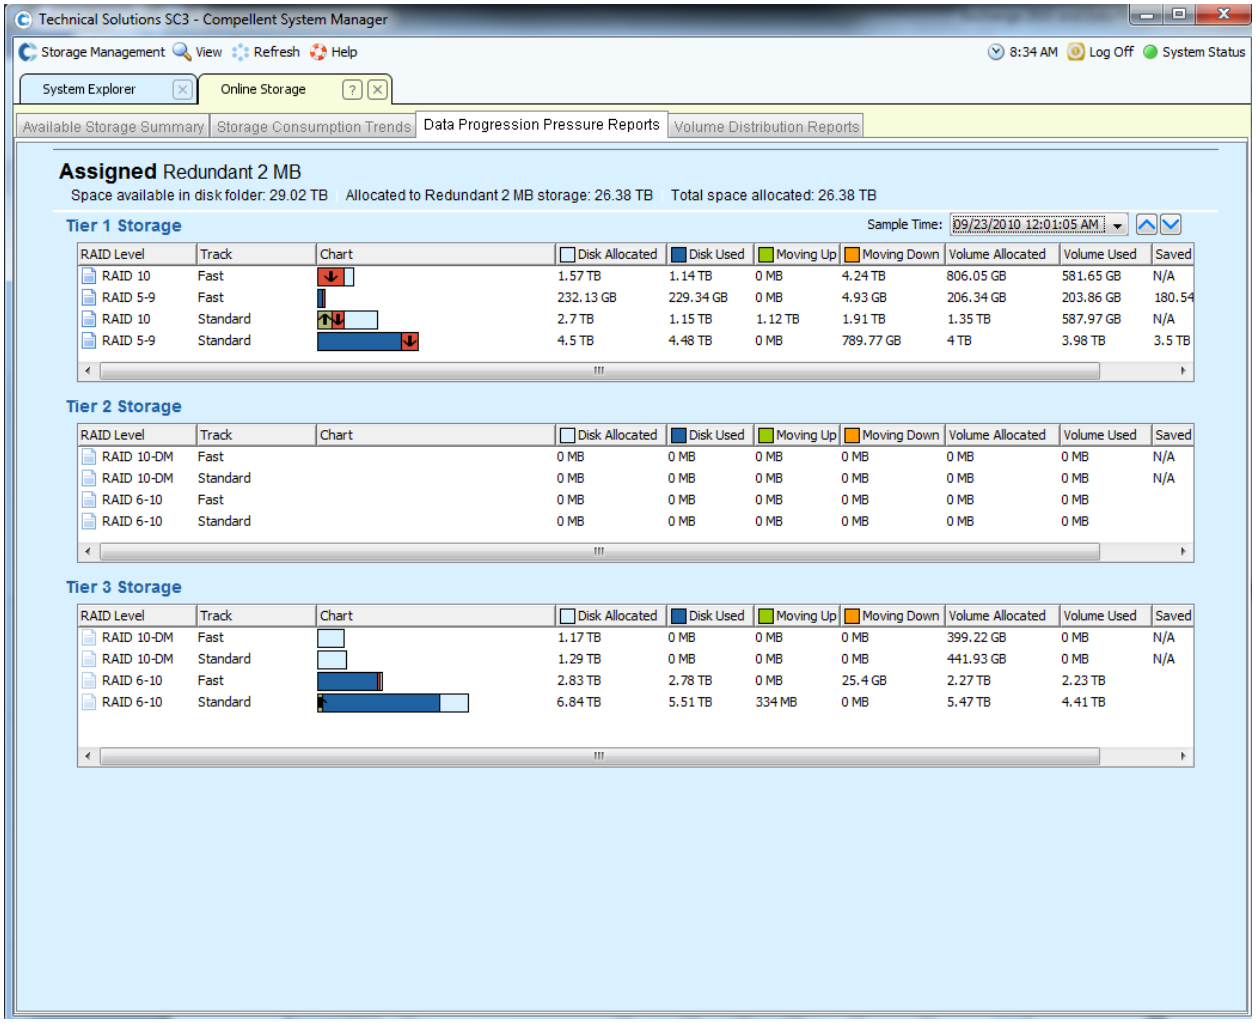
Technical Solutions SC3 - Compellent System Manager
Storage Management
View
Refresh
Help
8:34 AM
Log Off
System Status
System Explorer
Online Storage
Available Storage Summary
Storage Consumption Trends
Data Progression Pressure Reports
Volume Distribution Reports
Assigned Redundant 2 MB
Space available in disk folder: 29.02 TB
Allocated to Redundant 2 MB storage: 26.38 TB
Total space allocated: 26.38 TB
Tier 1 Storage
Sample Time: 09/23/2010 12:01:05 AM
RAID Level | Track | Chart | Disk Allocated | Disk Used | Moving Up | Moving Down | Volume Allocated | Volume Used | Saved
RAID 10 | Fast | 1.57 TB | 1.14 TB | 0 MB | 4.24 TB | 806.05 GB | 581.65 GB | N/A
RAID 5-9 | Fast | 232.13 GB | 229.34 GB | 0 MB | 4.93 GB | 206.34 GB | 203.86 GB | 180.54
RAID 10 | Standard | 2.7 TB | 1.15 TB | 1.12 TB | 1.91 TB | 1.35 TB | 587.97 GB | N/A
RAID 5-9 | Standard | 4.5 TB | 4.48 TB | 0 MB | 789.77 GB | 4 TB | 3.98 TB | 3.5 TB
Tier 2 Storage
RAID Level | Track | Chart | Disk Allocated | Disk Used | Moving Up | Moving Down | Volume Allocated | Volume Used | Saved
RAID 10-DM | Fast | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | N/A
RAID 10-DM | Standard | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | N/A
RAID 6-10 | Fast | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB
RAID 6-10 | Standard | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB
Tier 3 Storage
RAID Level | Track | Chart | Disk Allocated | Disk Used | Moving Up | Moving Down | Volume Allocated | Volume Used | Saved
RAID 10-DM | Fast | 1.17 TB | 0 MB | 0 MB | 0 MB | 399.22 GB | 0 MB | N/A
RAID 10-DM | Standard | 1.29 TB | 0 MB | 0 MB | 0 MB | 441.93 GB | 0 MB | N/A
RAID 6-10 | Fast | 2.83 TB | 2.78 TB | 0 MB | 25.4 GB | 2.27 TB | 2.23 TB
RAID 6-10 | Standard | 6.84 TB | 5.51 TB | 334 MB | 0 MB | 5.47 TB | 4.41 TB

Figure 11

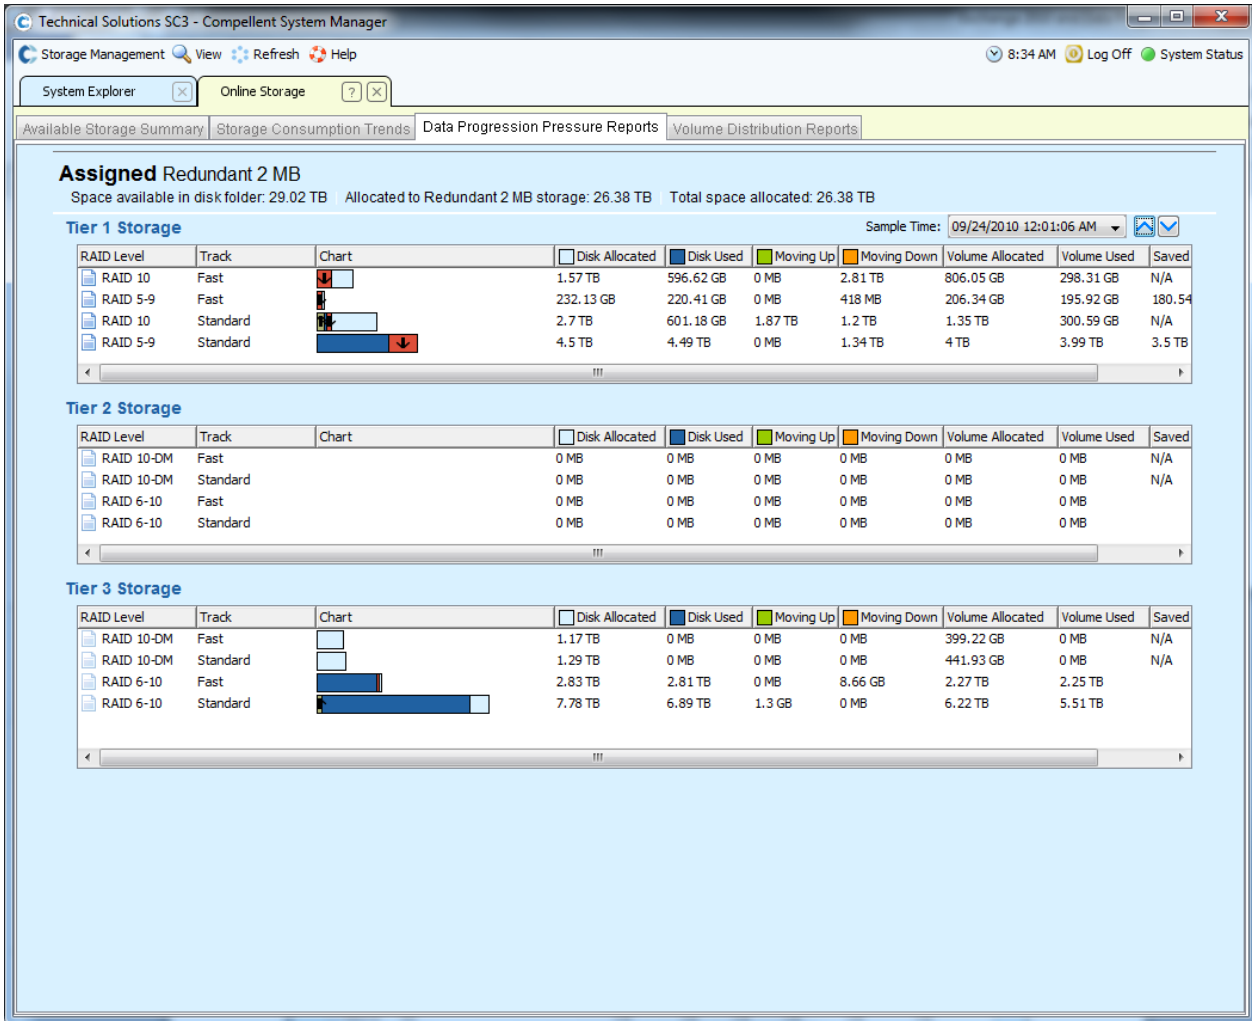
Technical Solutions SC3 - Compellent System Manager
Storage Management
View
Refresh
Help
8:34 AM
Log Off
System Status
System Explorer
Online Storage
Available Storage Summary
Storage Consumption Trends
Data Progression Pressure Reports
Volume Distribution Reports
Assigned Redundant 2 MB
Space available in disk folder: 29.02 TB
Allocated to Redundant 2 MB storage: 26.38 TB
Total space allocated: 26.38 TB
Tier 1 Storage
Sample Time: 09/24/2010 12:01:06 AM
RAID Level | Track | Chart | Disk Allocated | Disk Used | Moving Up | Moving Down | Volume Allocated | Volume Used | Saved
RAID 10 | Fast | 1.57 TB | 596.62 GB | 0 MB | 2.81 TB | 806.05 GB | 298.31 GB | N/A
RAID 5-9 | Fast | 232.13 GB | 220.41 GB | 0 MB | 418 MB | 206.34 GB | 195.92 GB | 180.54
RAID 10 | Standard | 2.7 TB | 601.18 GB | 1.87 TB | 1.2 TB | 1.35 TB | 300.59 GB | N/A
RAID 5-9 | Standard | 4.5 TB | 4.49 TB | 0 MB | 1.34 TB | 4 TB | 3.99 TB | 3.5 TB
Tier 2 Storage
RAID Level | Track | Chart | Disk Allocated | Disk Used | Moving Up | Moving Down | Volume Allocated | Volume Used | Saved
RAID 10-DM | Fast | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | N/A
RAID 10-DM | Standard | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | N/A
RAID 6-10 | Fast | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB
RAID 6-10 | Standard | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB
Tier 3 Storage
RAID Level | Track | Chart | Disk Allocated | Disk Used | Moving Up | Moving Down | Volume Allocated | Volume Used | Saved
RAID 10-DM | Fast | 1.17 TB | 0 MB | 0 MB | 0 MB | 399.22 GB | 0 MB | N/A
RAID 10-DM | Standard | 1.29 TB | 0 MB | 0 MB | 0 MB | 441.93 GB | 0 MB | N/A
RAID 6-10 | Fast | 2.83 TB | 2.81 TB | 0 MB | 8.66 GB | 2.27 TB | 2.25 TB
RAID 6-10 | Standard | 7.78 TB | 6.89 TB | 1.3 GB | 0 MB | 6.22 TB | 5.51 TB

Figure 12

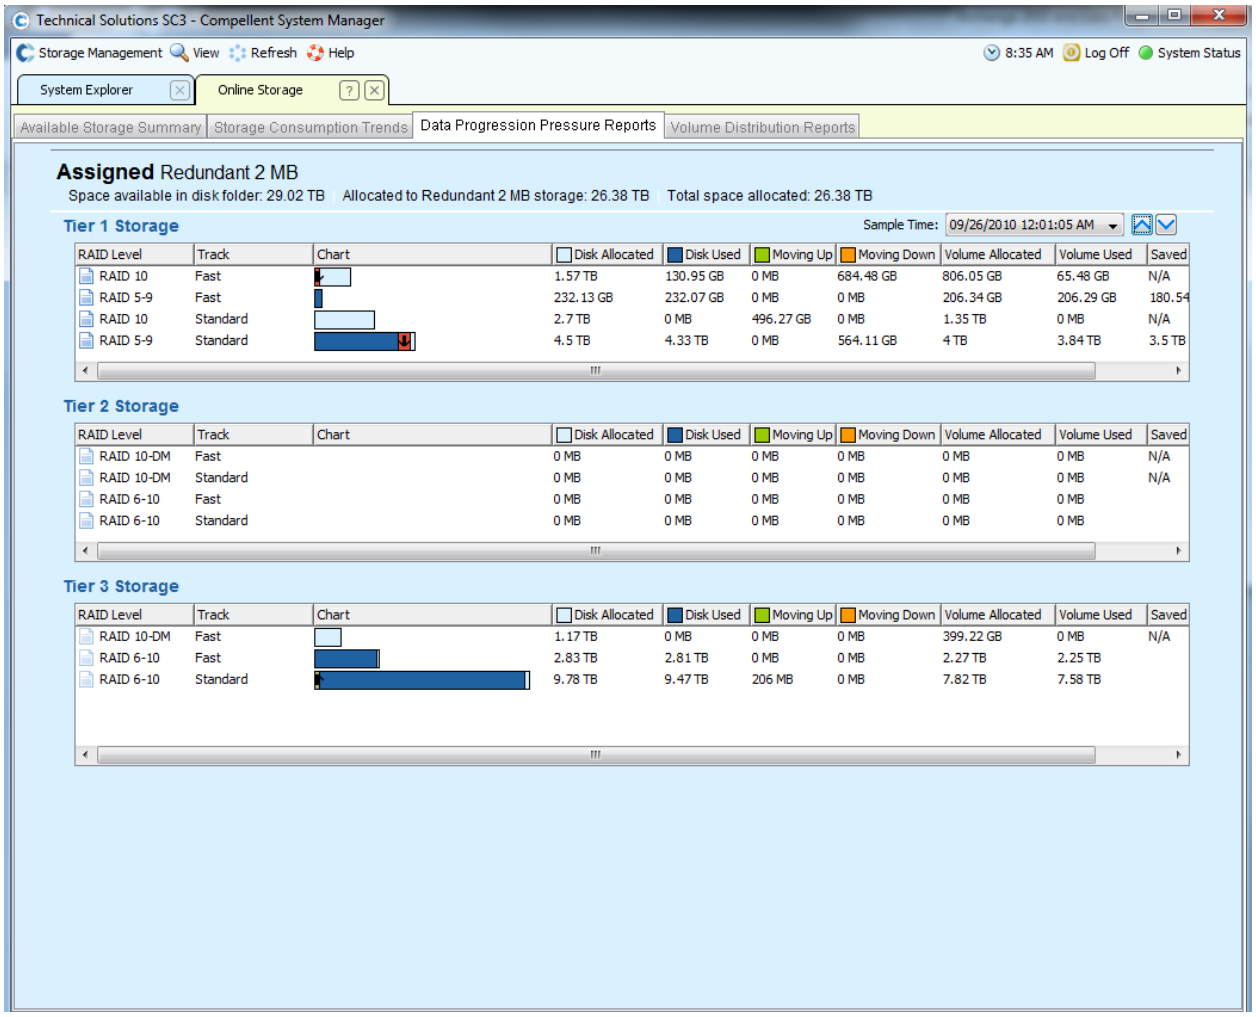
Technical Solutions SC3 - Compellent System Manager
Storage Management
View
Refresh
Help
8:35 AM
Log Off
System Status
System Explorer
Online Storage
Available Storage Summary
Storage Consumption Trends
Data Progression Pressure Reports
Volume Distribution Reports
Assigned Redundant 2 MB
Space available in disk folder: 29.02 TB
Allocated to Redundant 2 MB storage: 26.38 TB
Total space allocated: 26.38 TB
Tier 1 Storage
Sample Time: 09/26/2010 12:01:05 AM
RAID Level | Track | Chart | Disk Allocated | Disk Used | Moving Up | Moving Down | Volume Allocated | Volume Used | Saved
RAID 10 | Fast | 1.57 TB | 130.95 GB | 0 MB | 684.48 GB | 806.05 GB | 65.48 GB | N/A
RAID 5-9 | Fast | 232.13 GB | 232.07 GB | 0 MB | 0 MB | 206.34 GB | 206.29 GB | 180.54
RAID 10 | Standard | 2.7 TB | 0 MB | 496.27 GB | 0 MB | 1.35 TB | 0 MB | N/A
RAID 5-9 | Standard | 4.5 TB | 4.33 TB | 0 MB | 564.11 GB | 4 TB | 3.84 TB | 3.5 TB
Tier 2 Storage
RAID Level | Track | Chart | Disk Allocated | Disk Used | Moving Up | Moving Down | Volume Allocated | Volume Used | Saved
RAID 10-DM | Fast | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | N/A
RAID 10-DM | Standard | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | N/A
RAID 6-10 | Fast | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB
RAID 6-10 | Standard | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB | 0 MB
Tier 3 Storage
RAID Level | Track | Chart | Disk Allocated | Disk Used | Moving Up | Moving Down | Volume Allocated | Volume Used | Saved
RAID 10-DM | Fast | 1.17 TB | 0 MB | 0 MB | 0 MB | 399.22 GB | 0 MB | N/A
RAID 6-10 | Fast | 2.83 TB | 2.81 TB | 0 MB | 0 MB | 2.27 TB | 2.25 TB
RAID 6-10 | Standard | 9.78 TB | 9.47 TB | 206 MB | 0 MB | 7.82 TB | 7.58 TB

Figure 13