17664

Opening your eyes to how your Mainframe Tape environment is really performing.

Burt Loper
John Ticic

www.IntelliMagic.com

Agenda

Is Tape processing dead?

What data is available?

What can we observe in this data?

- Look at the z/OS and hardware view

What's important in our tape environment?

- Show examples of important aspects of tape processing, highlighting performance and problem investigation

Summary/Conclusions

Who is IntelliMagic

- A leader in Availability Intelligence
  - New visibility of threats to continuous availability by automatic interpretation of RMF/SMF/Config data using built-in expert knowledge
- Over 20 years developing storage performance solutions
- Privately held, financially independent
- Customer centric and highly responsive
- Products used daily at some of the largest sites in the world

IntelliMagic Vision

Protect availability with automated visibility inside your storage infrastructure

IntelliMagic Direction

Model how different storage hardware options would service your workloads

Presenter

- Burt Loper – Senior Technical Consultant
  - 35 years at IBM, latest experience architecting, installing and configuring TS7700 systems for customers
  - TS7700 Performance – authored the TS7700 Health Assessment
  - With IntelliMagic since January 2014

Is Tape Processing Dead?

Is Tape Processing Dead?

- Remains lowest cost per Terabyte
- Part of the Storage Hierarchy
- Legacy uses
  - Backup – possibly diminishing
  - Disaster Recovery – Last line of insurance
- Growing uses
  - Compliance – gov. or regulatory
  - Archive – older data being retained
  - Rapid growth in data – longer retentions

What data is available?

Tape Data Sources

- z/OS SMF and RMF general data
- IBM TS7700 BVIR history data
- Oracle VSM SMF data
- Tape catalog data

z/OS Tape Data Sources
SMF and TMS data

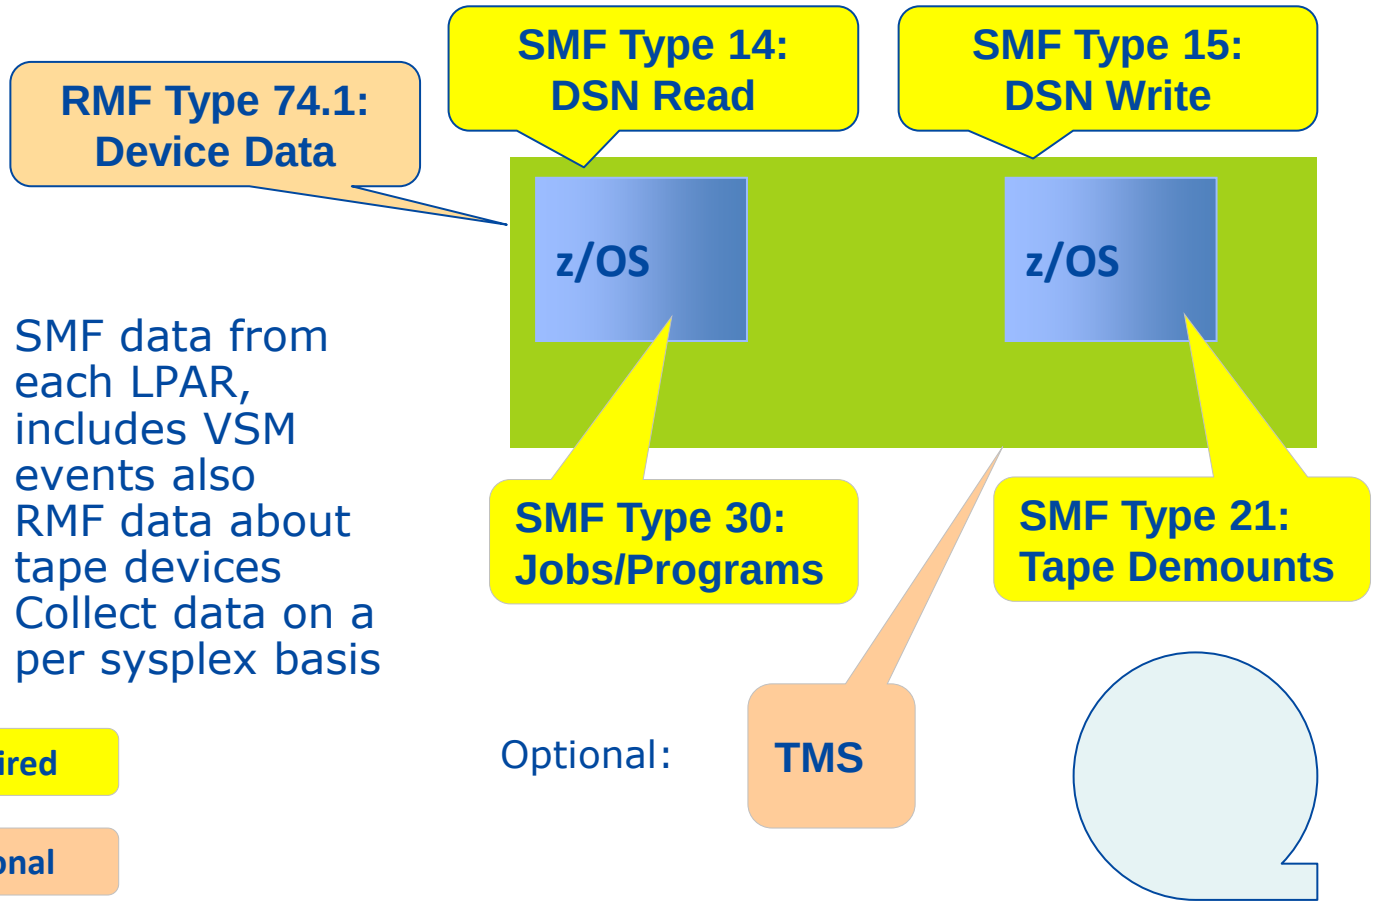

- SMF data from each LPAR, includes VSM events also
- RMF data about tape devices
- Collect data on a per sysplex basis

required

optional

Hardware Data Sources
TS7700 BVIR Data and VSM Records

- TS7700 BVIR collects data on a per Grid basis
- Consolidated by Grid/Library Cluster for reporting
- Oracle HSC writes special SMF records for VSM events (*see appendix for details*)

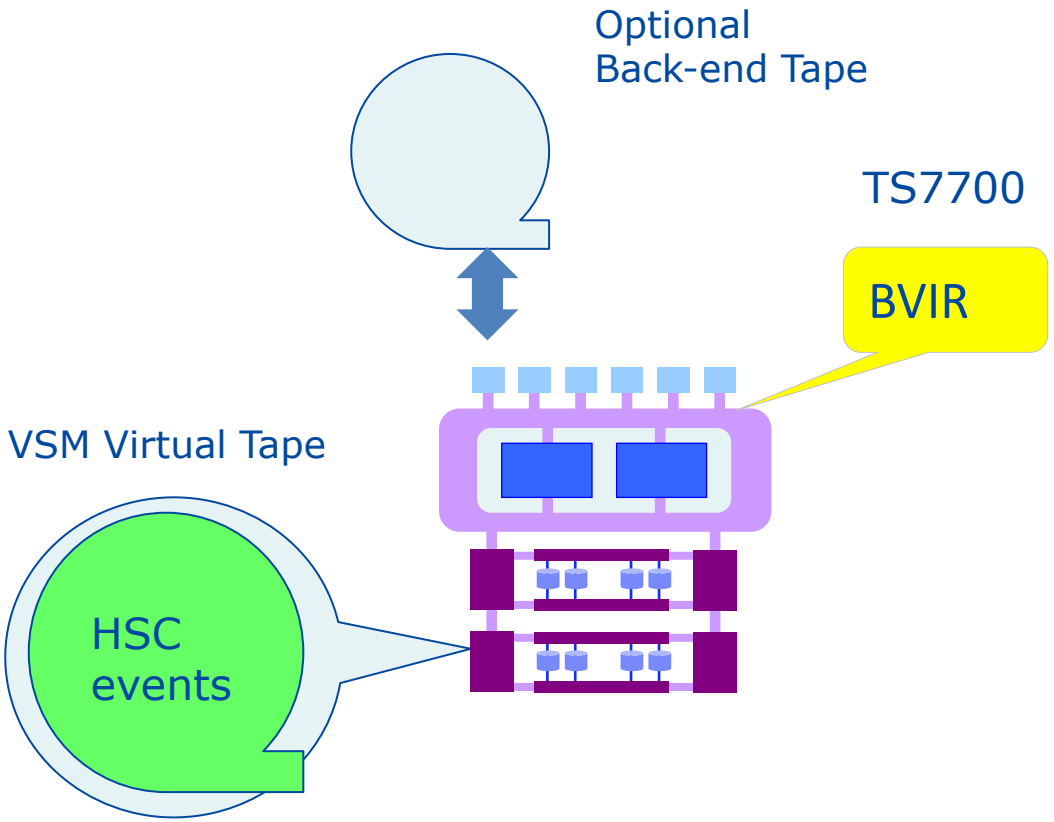

Tape Information is Everywhere

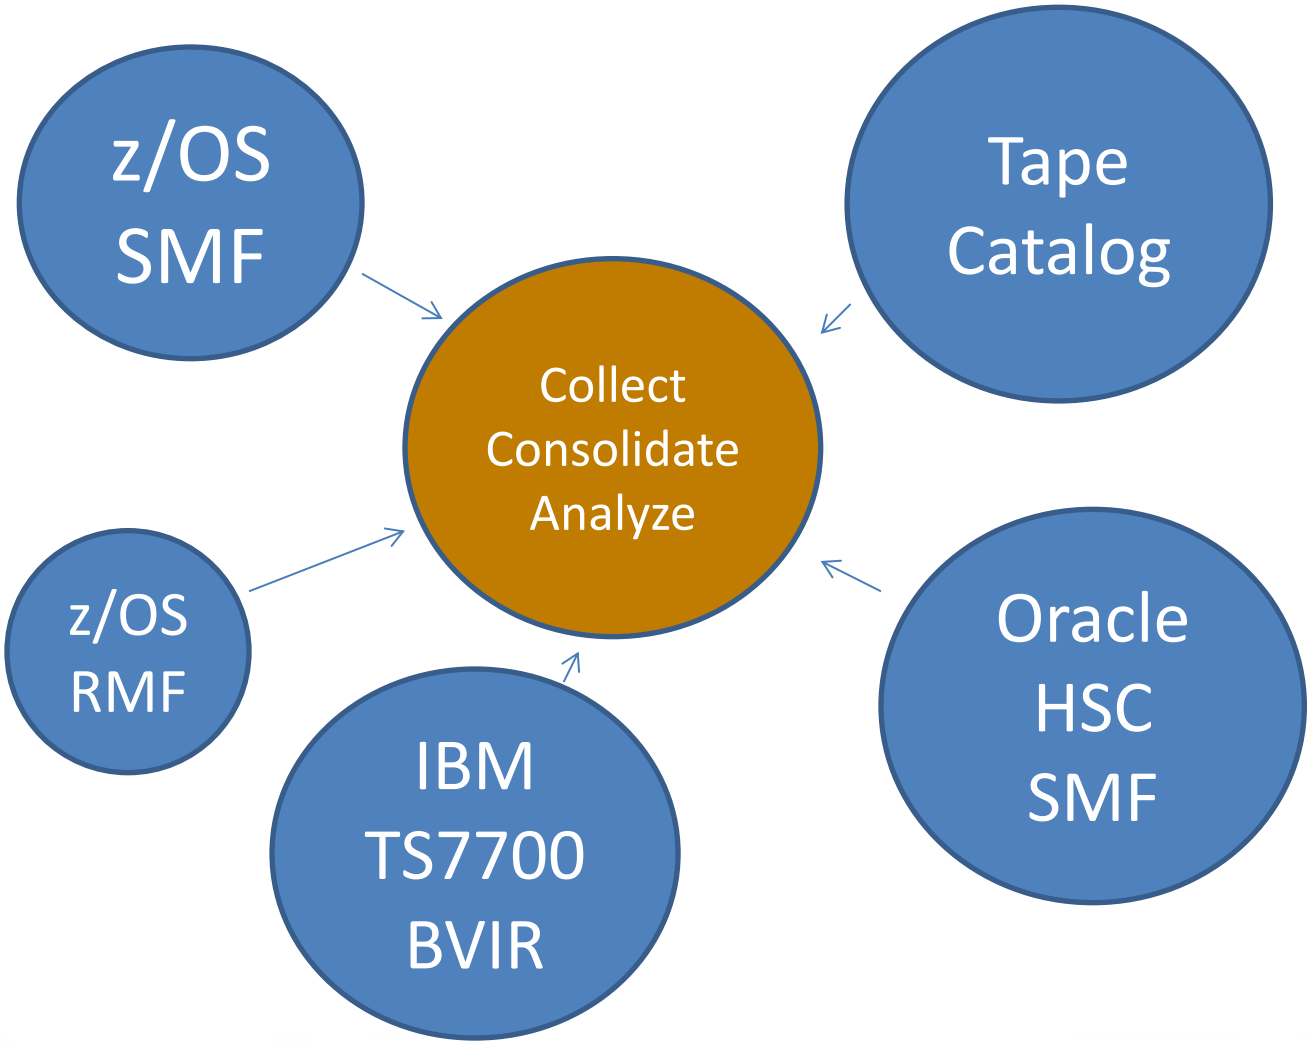

What's important in your Tape environment?

IBM TS7700 Performance

TS7700 Virtual Subsystem

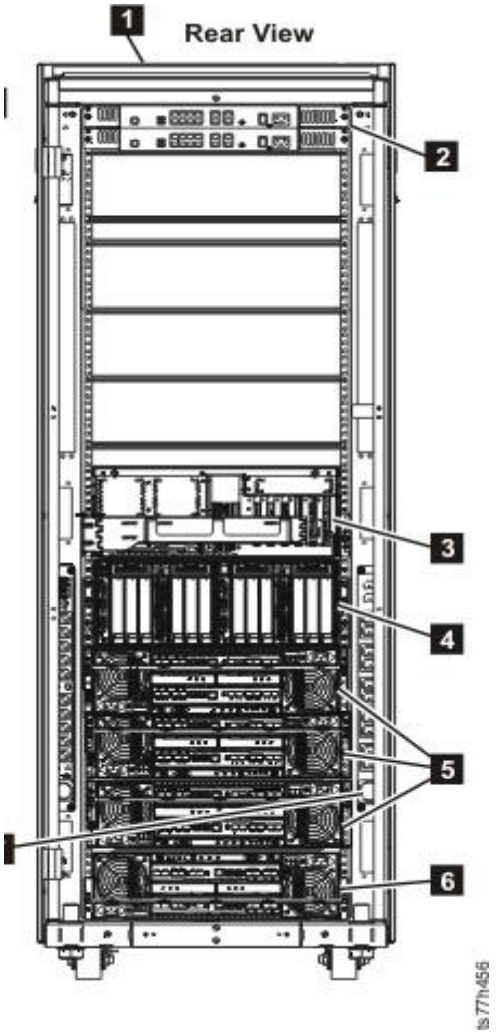

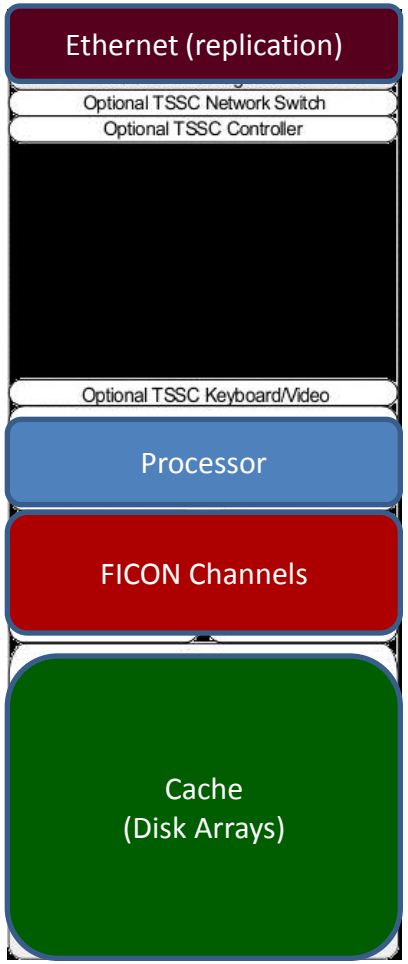

The TS7700 Dashboards Summarize the Analysis

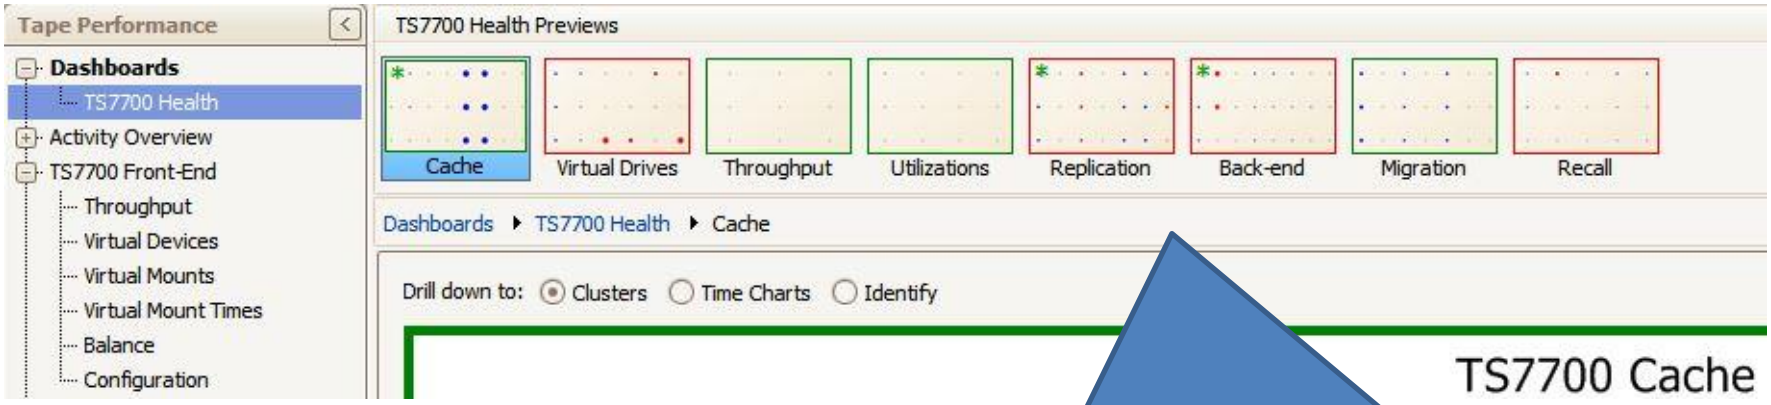

Each of these dashboards checks a particular aspect of the TS7700 performance and capacity

How Hard is my Hardware Running?

Utilization Dashboard: each Bubble Summarizes a Chart

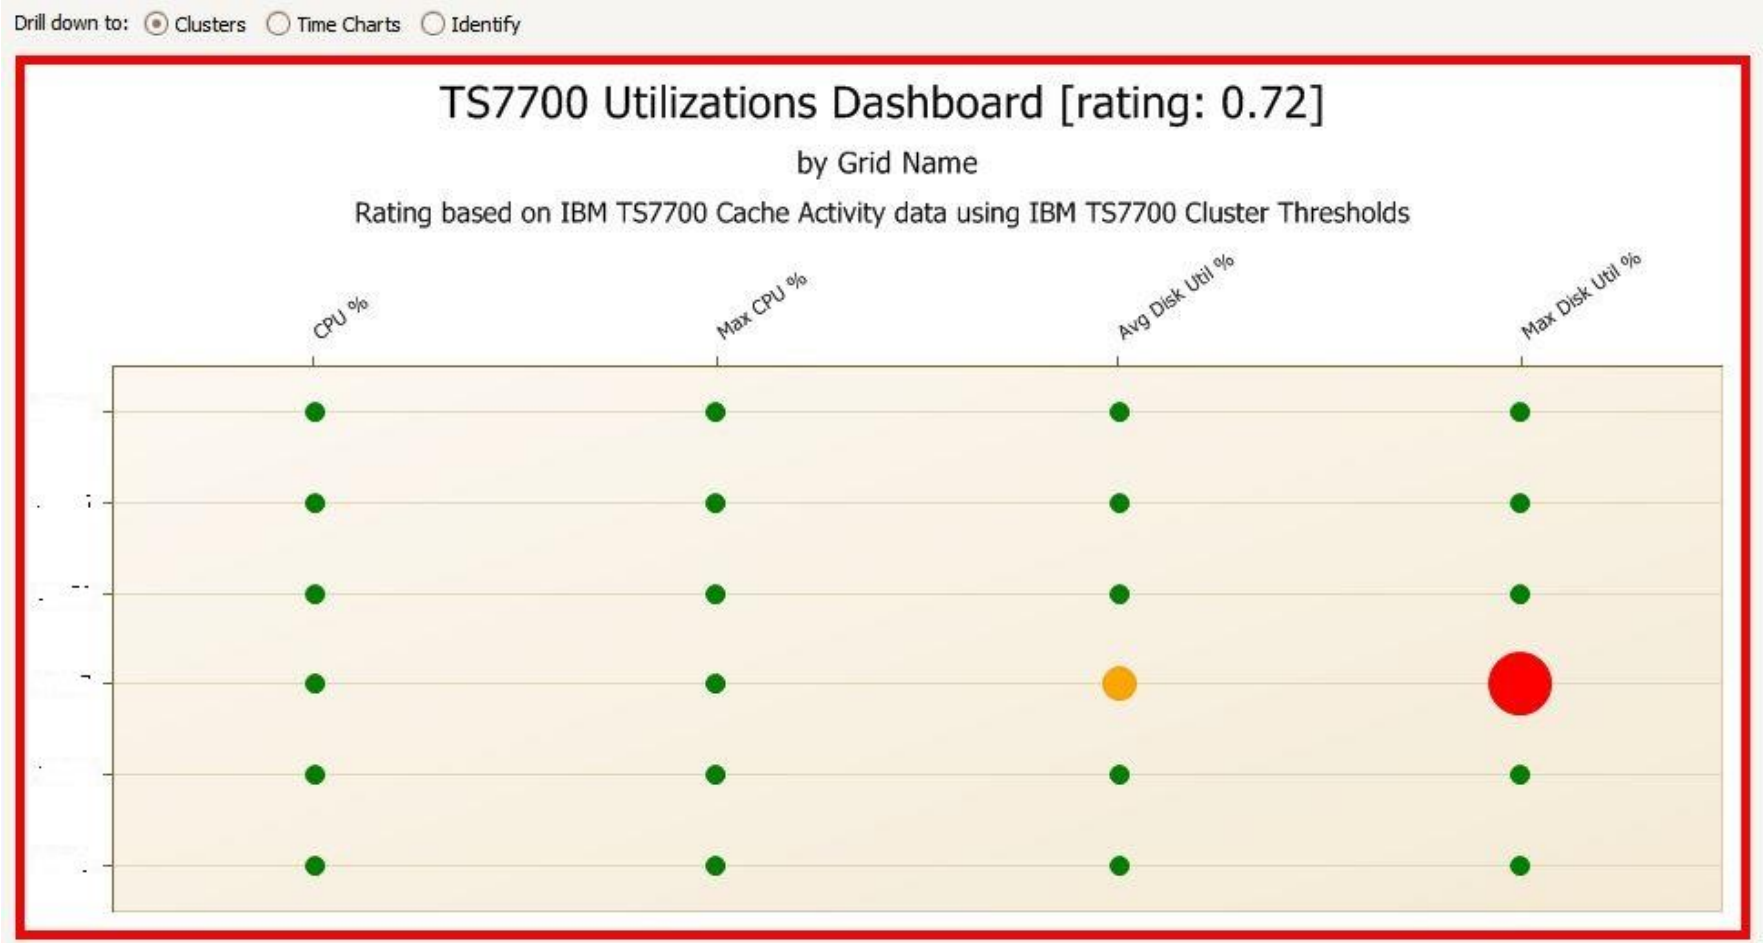

TS7700 Processor & Disk Utilizations

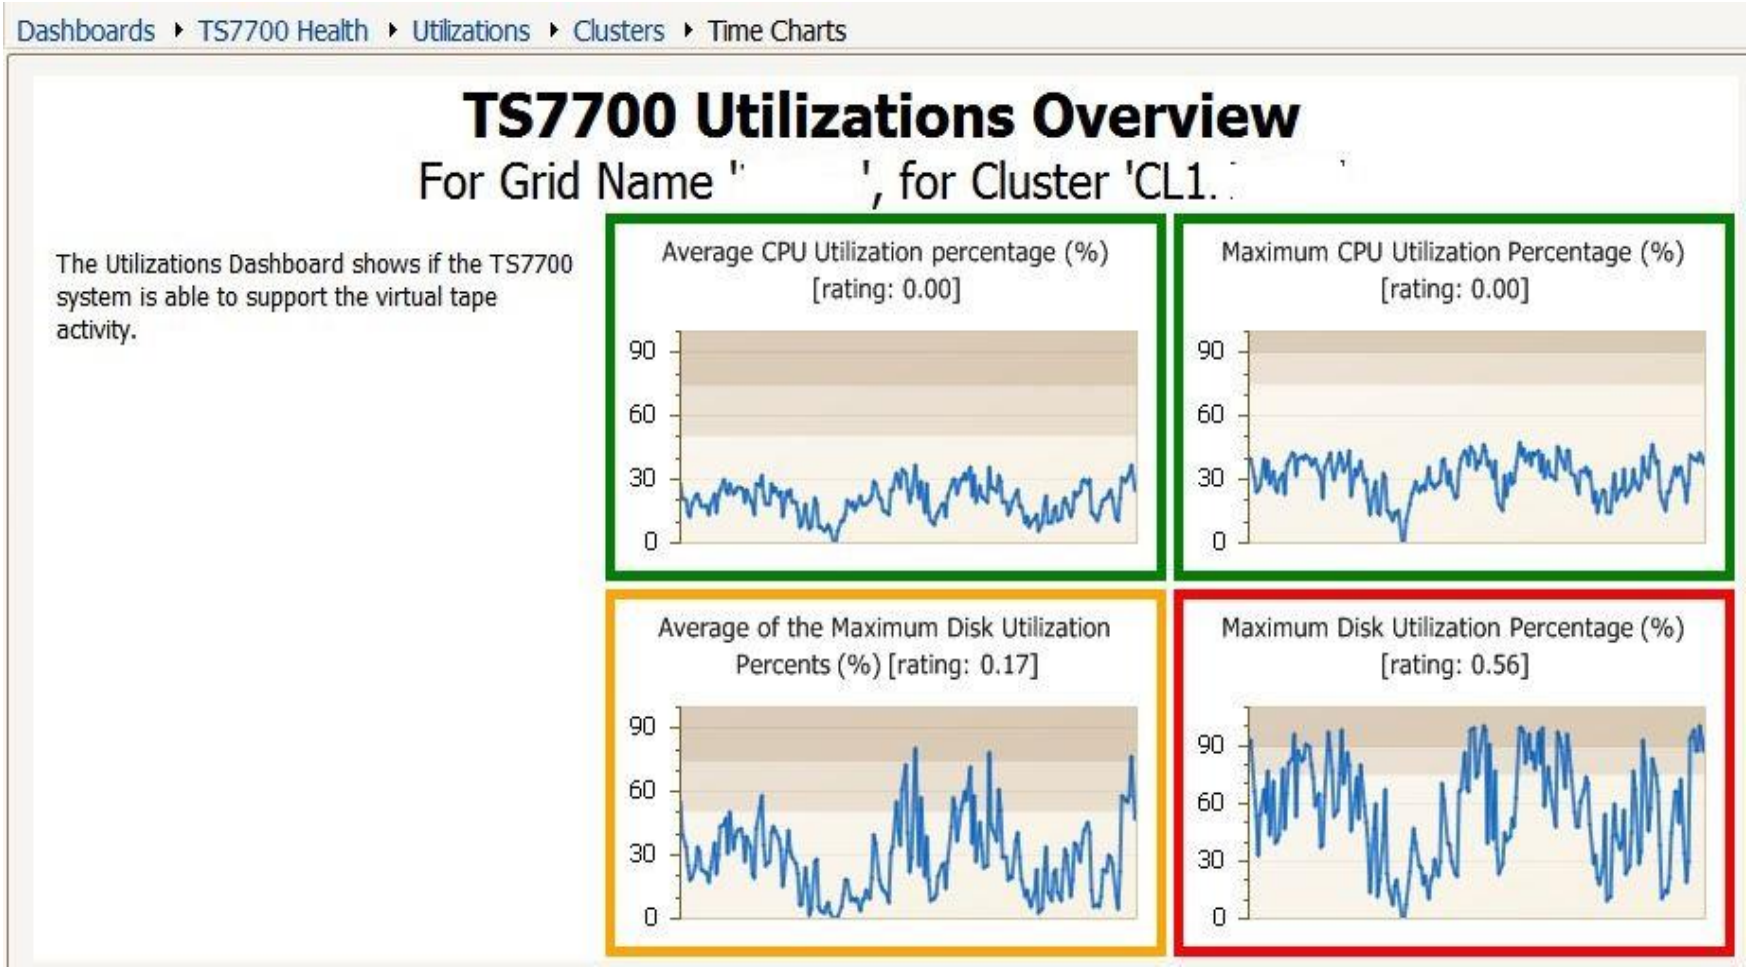

TS7700 Processor Utilizations

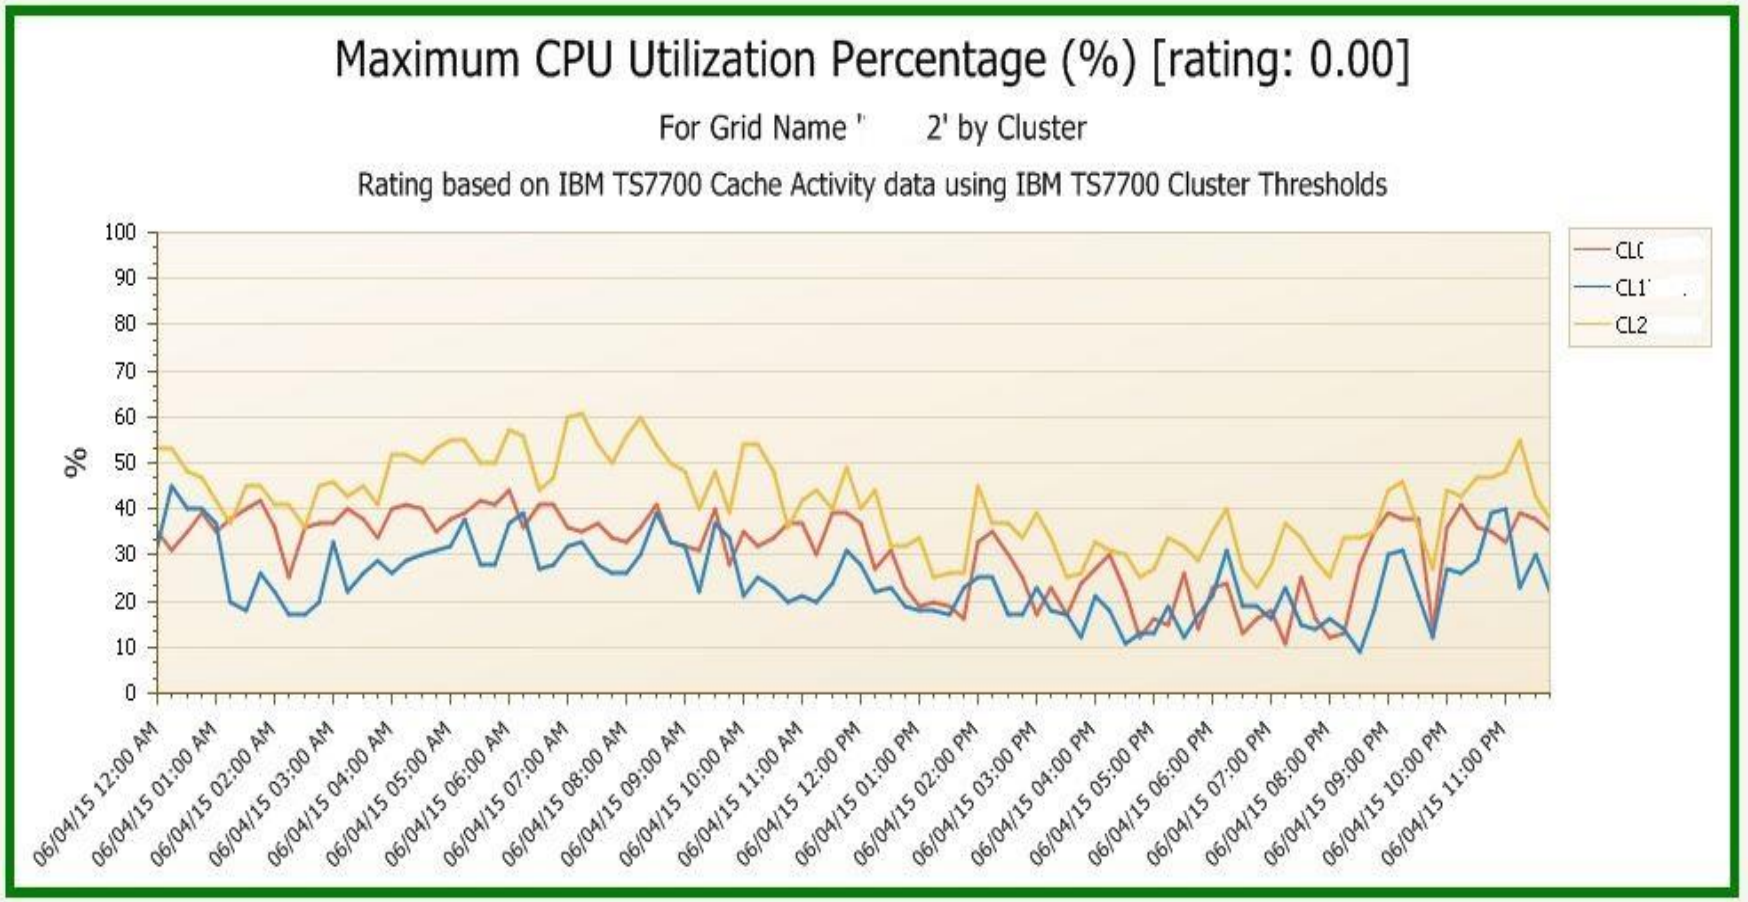

TS7700 Disk Utilization

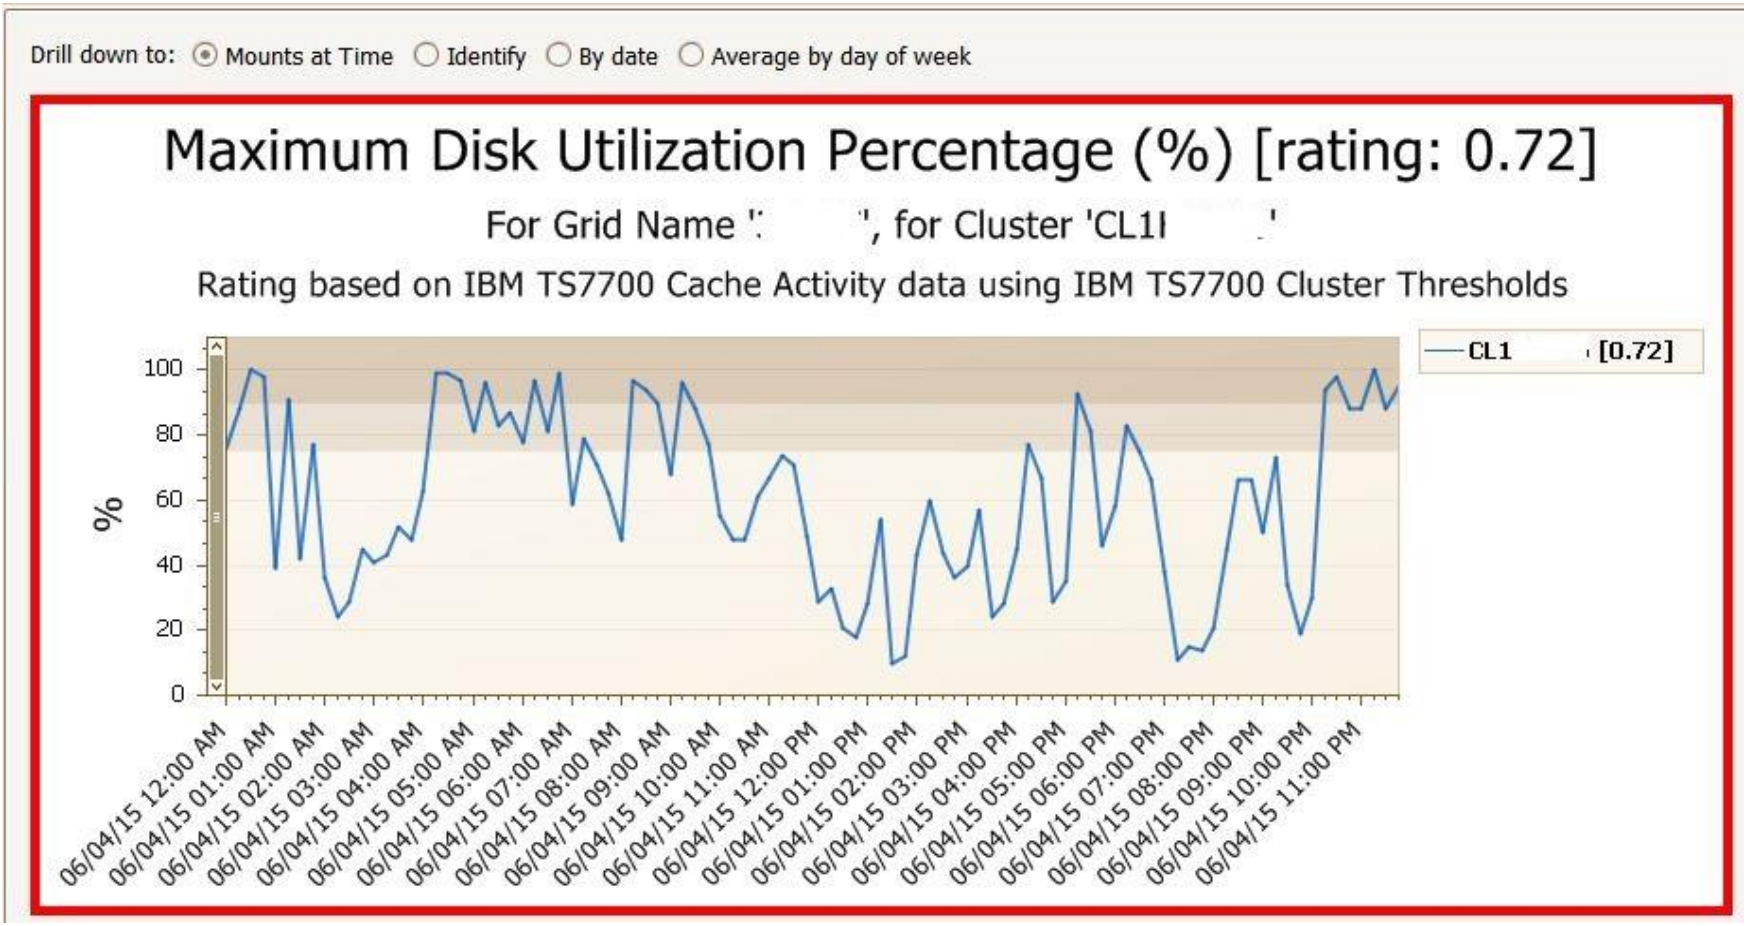

Channel Throughput (MB/s) for all TS7700 Grids

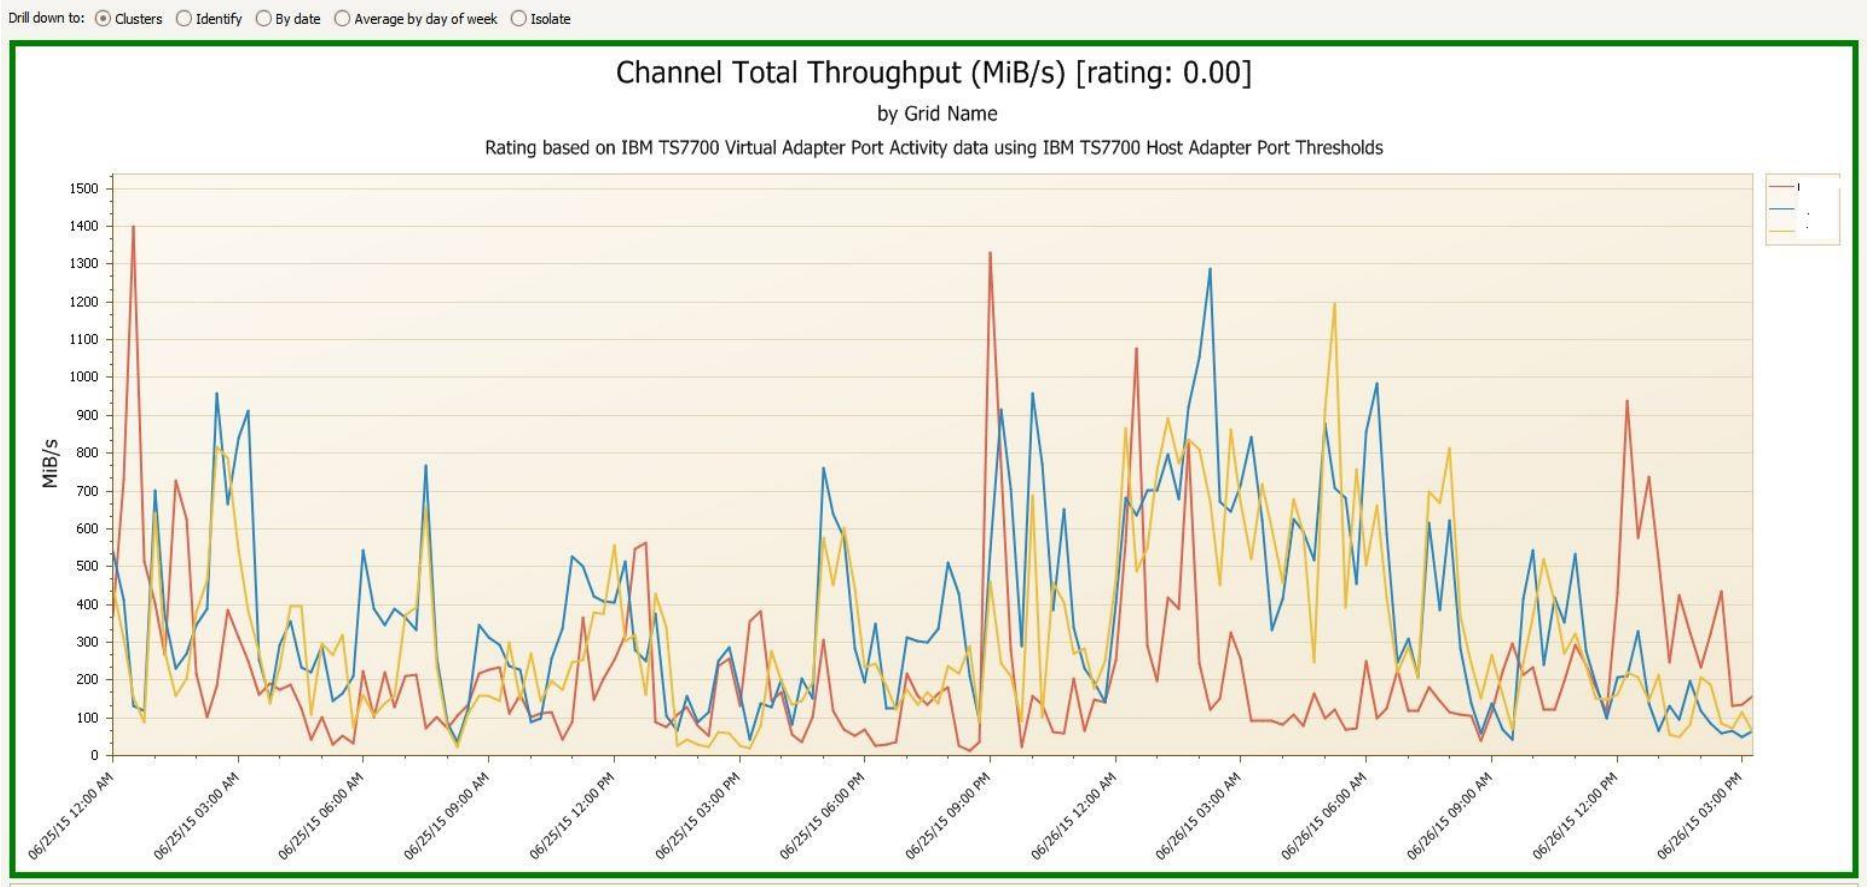

TS7700 Cache Flows

TS7700 System ▸ System ▸ Cache Rate ▸ Isolate ▸ Cache Flows ▸ Cache Flows

Drill down to: Identify

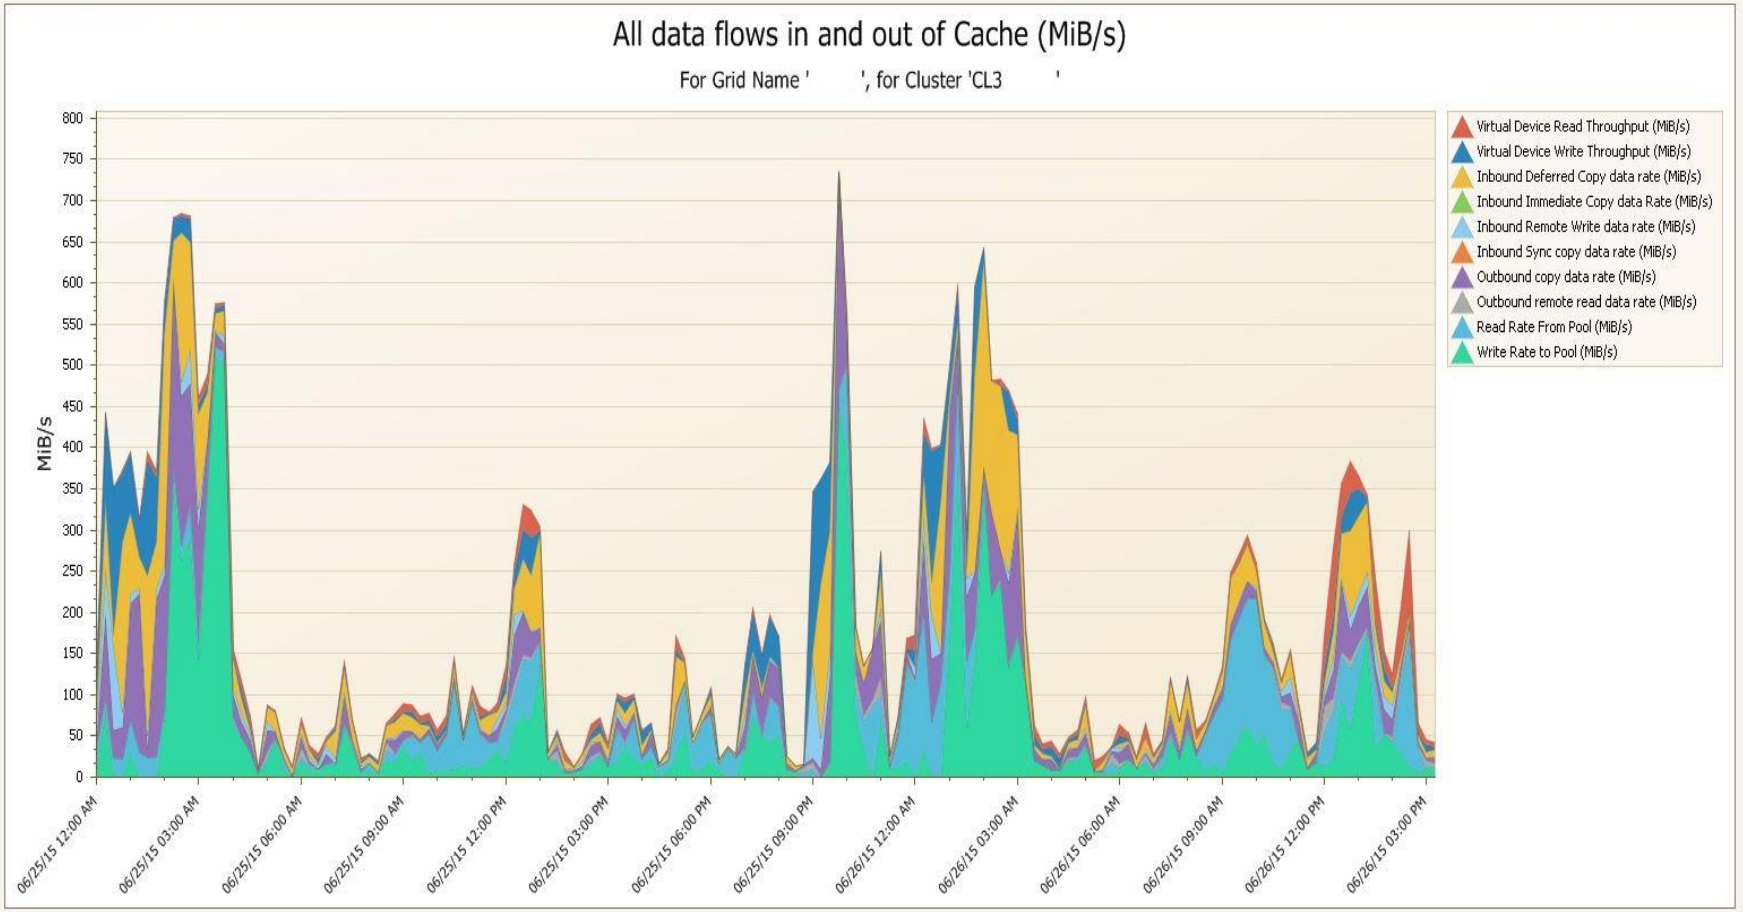

How Long is it taking for Data to be Replicated to my DR Site?

Replication – Receiving Cluster

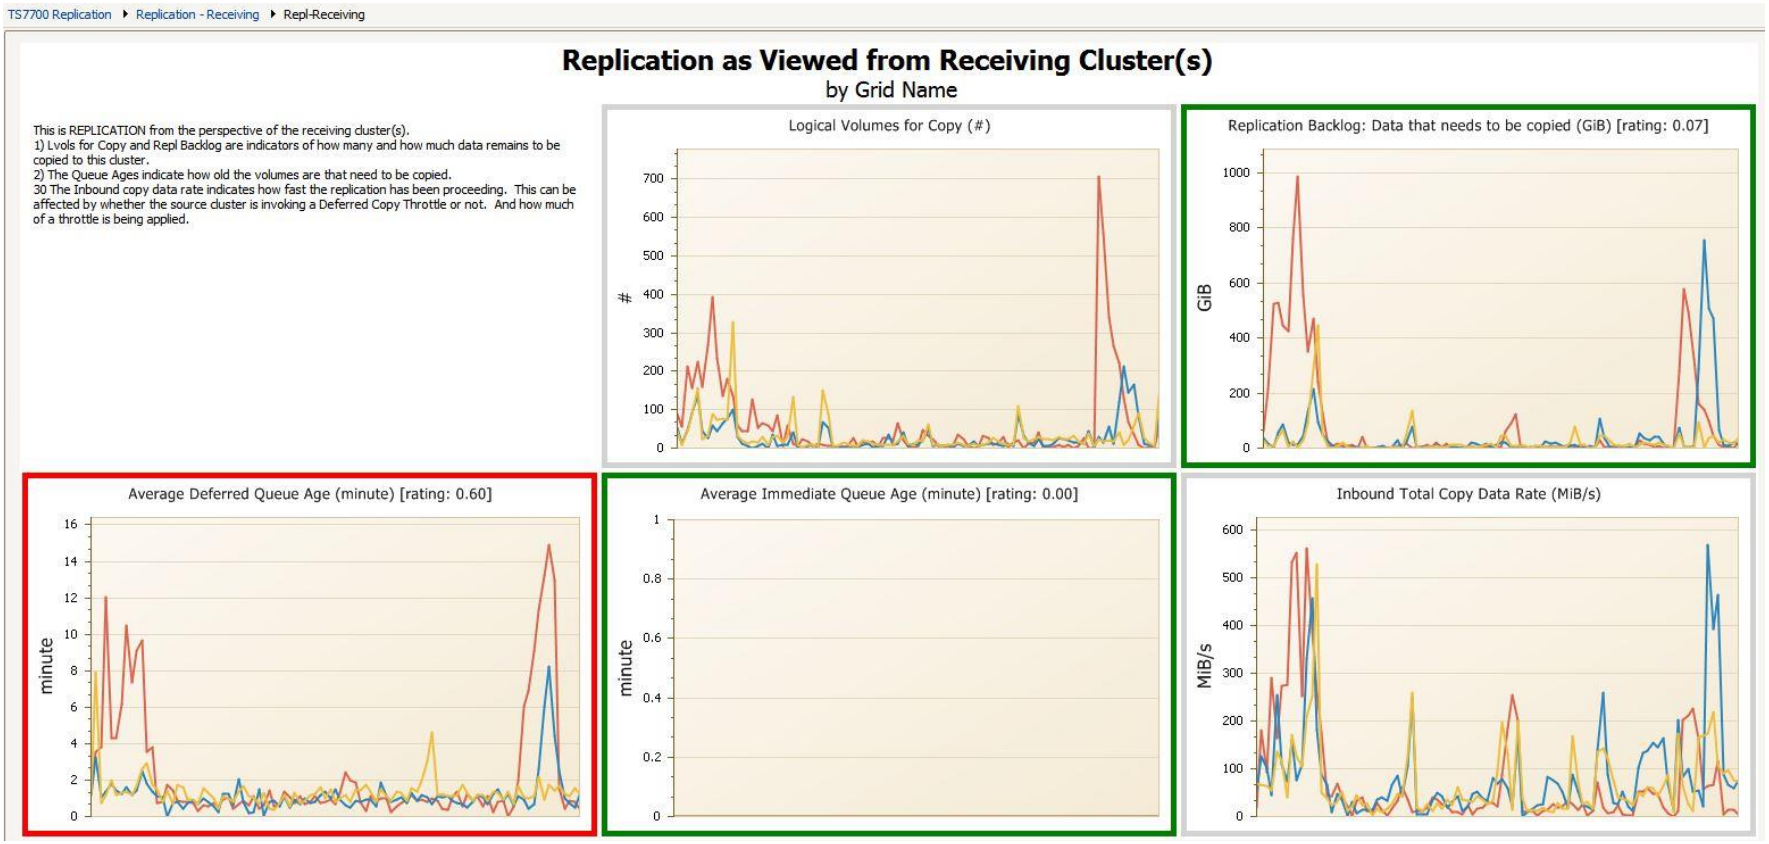

Replication – Sending Cluster

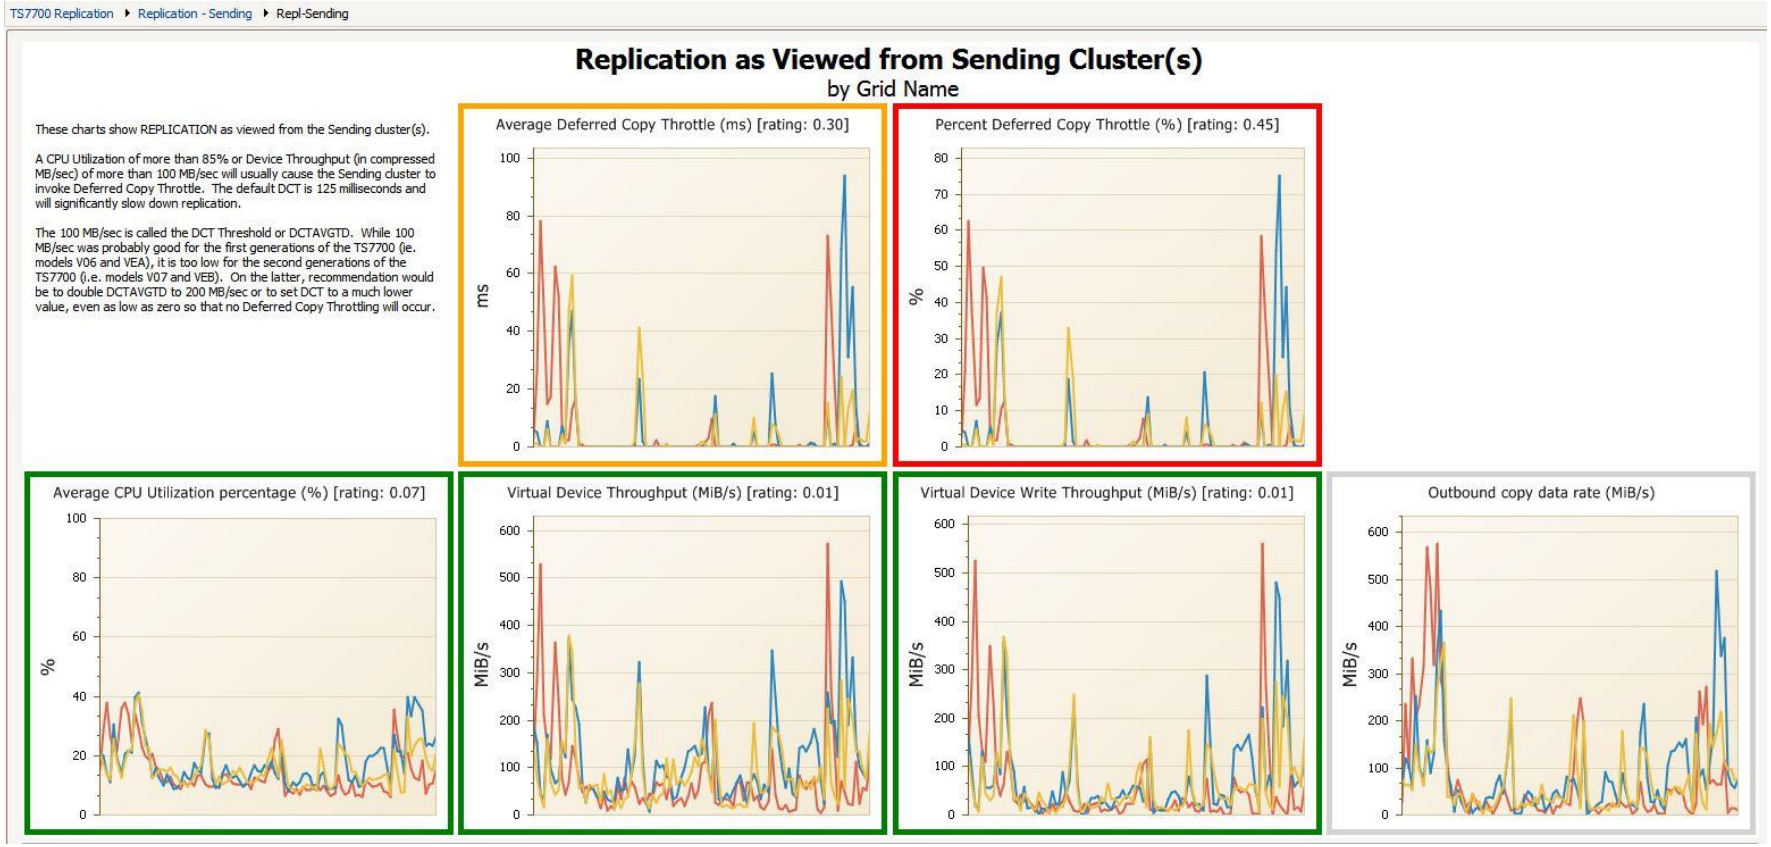

Is there enough Cache to adequately support your Tape Workloads?

Cache Dashboard – All Grids

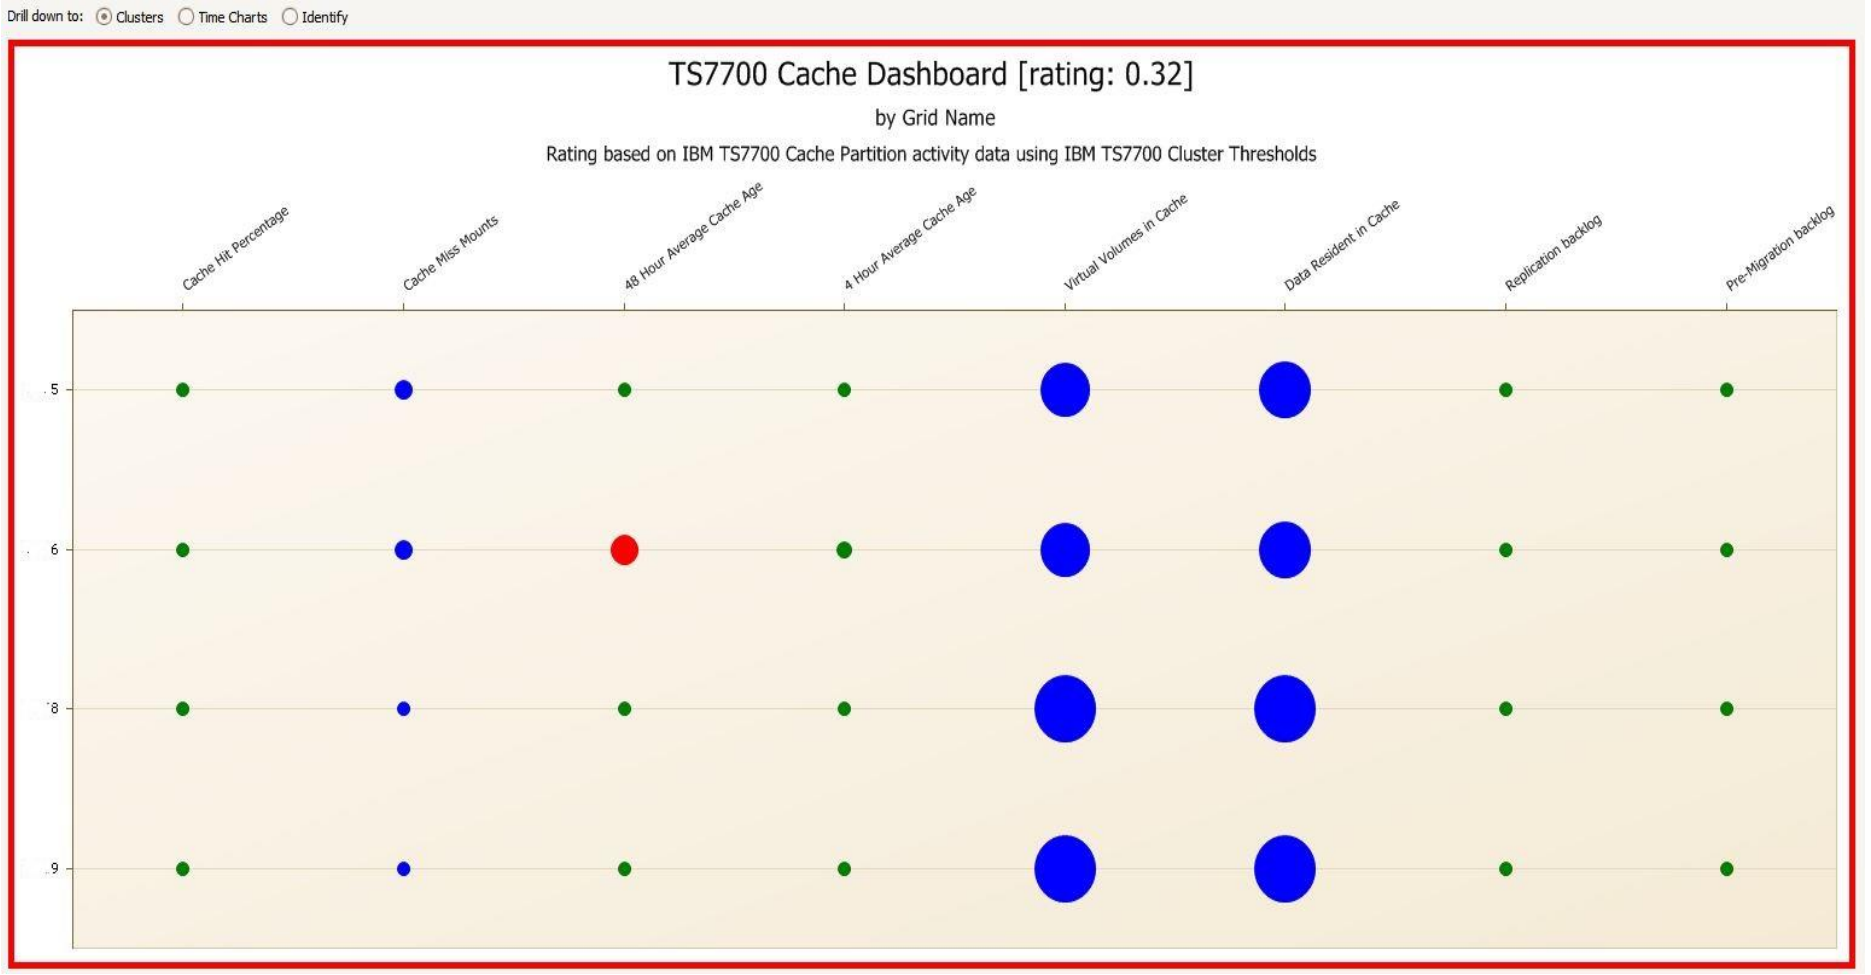

Cache Dashboard – Single Grid

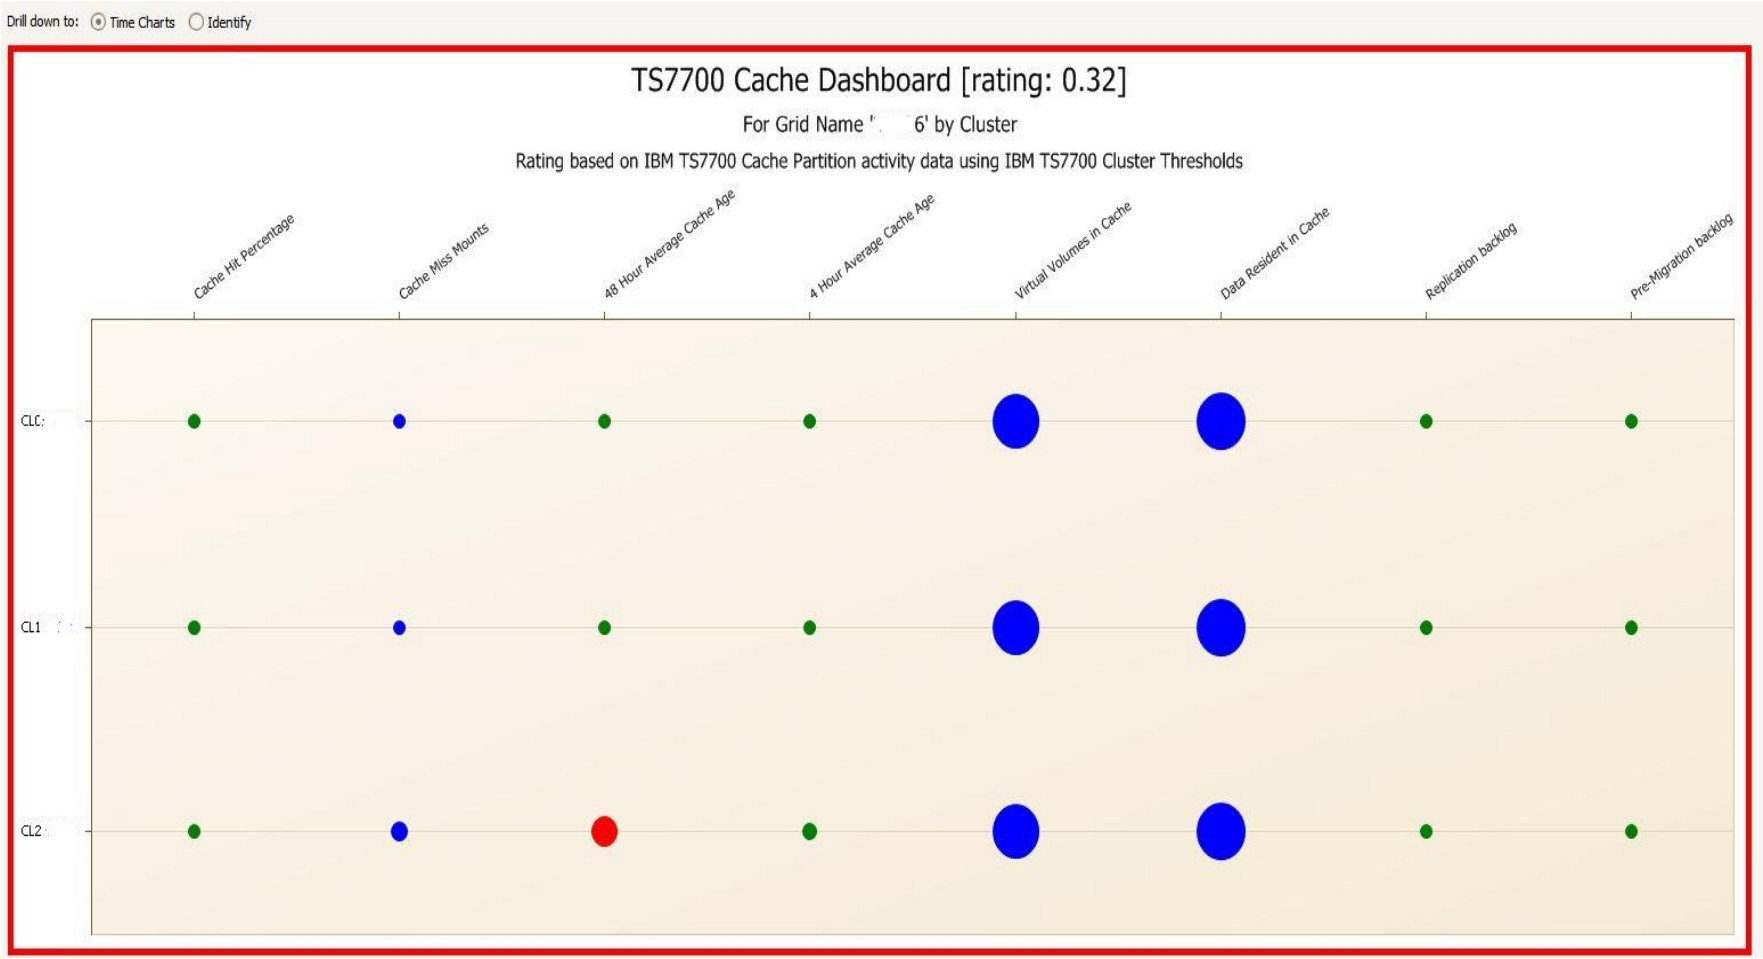

Cache Overview – Multi-chart

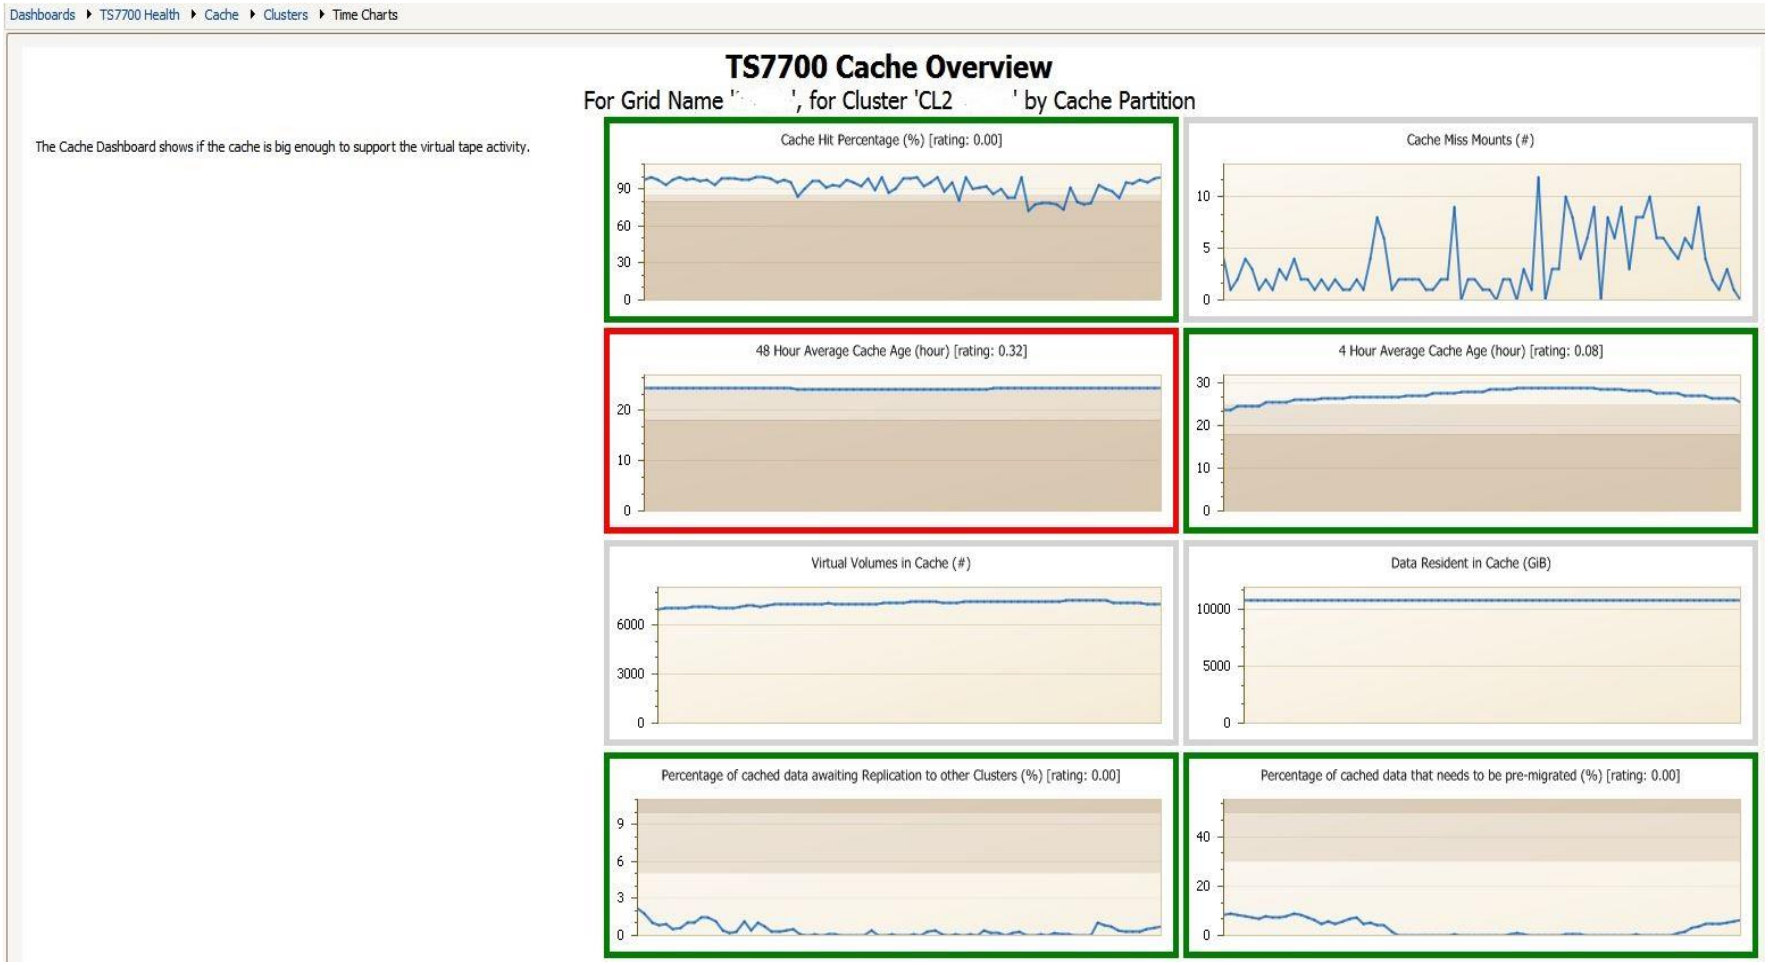

Avg. Cache Age – 18 hour interval

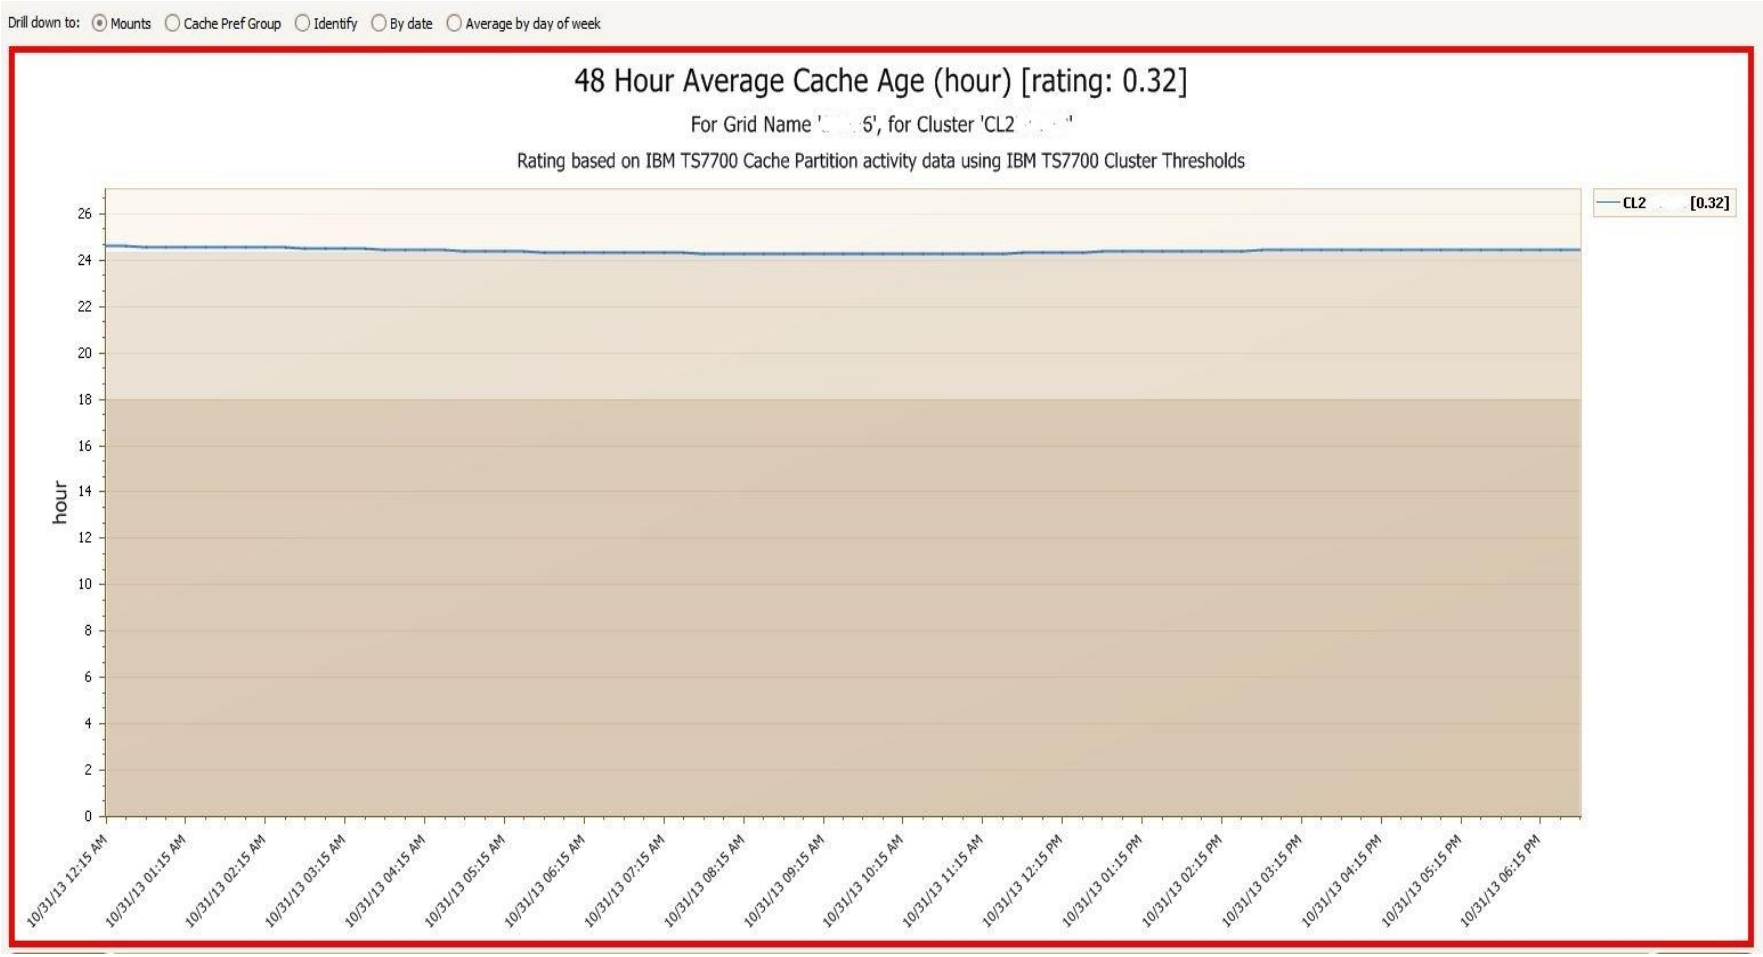

Avg. Cache Age – over 2 weeks

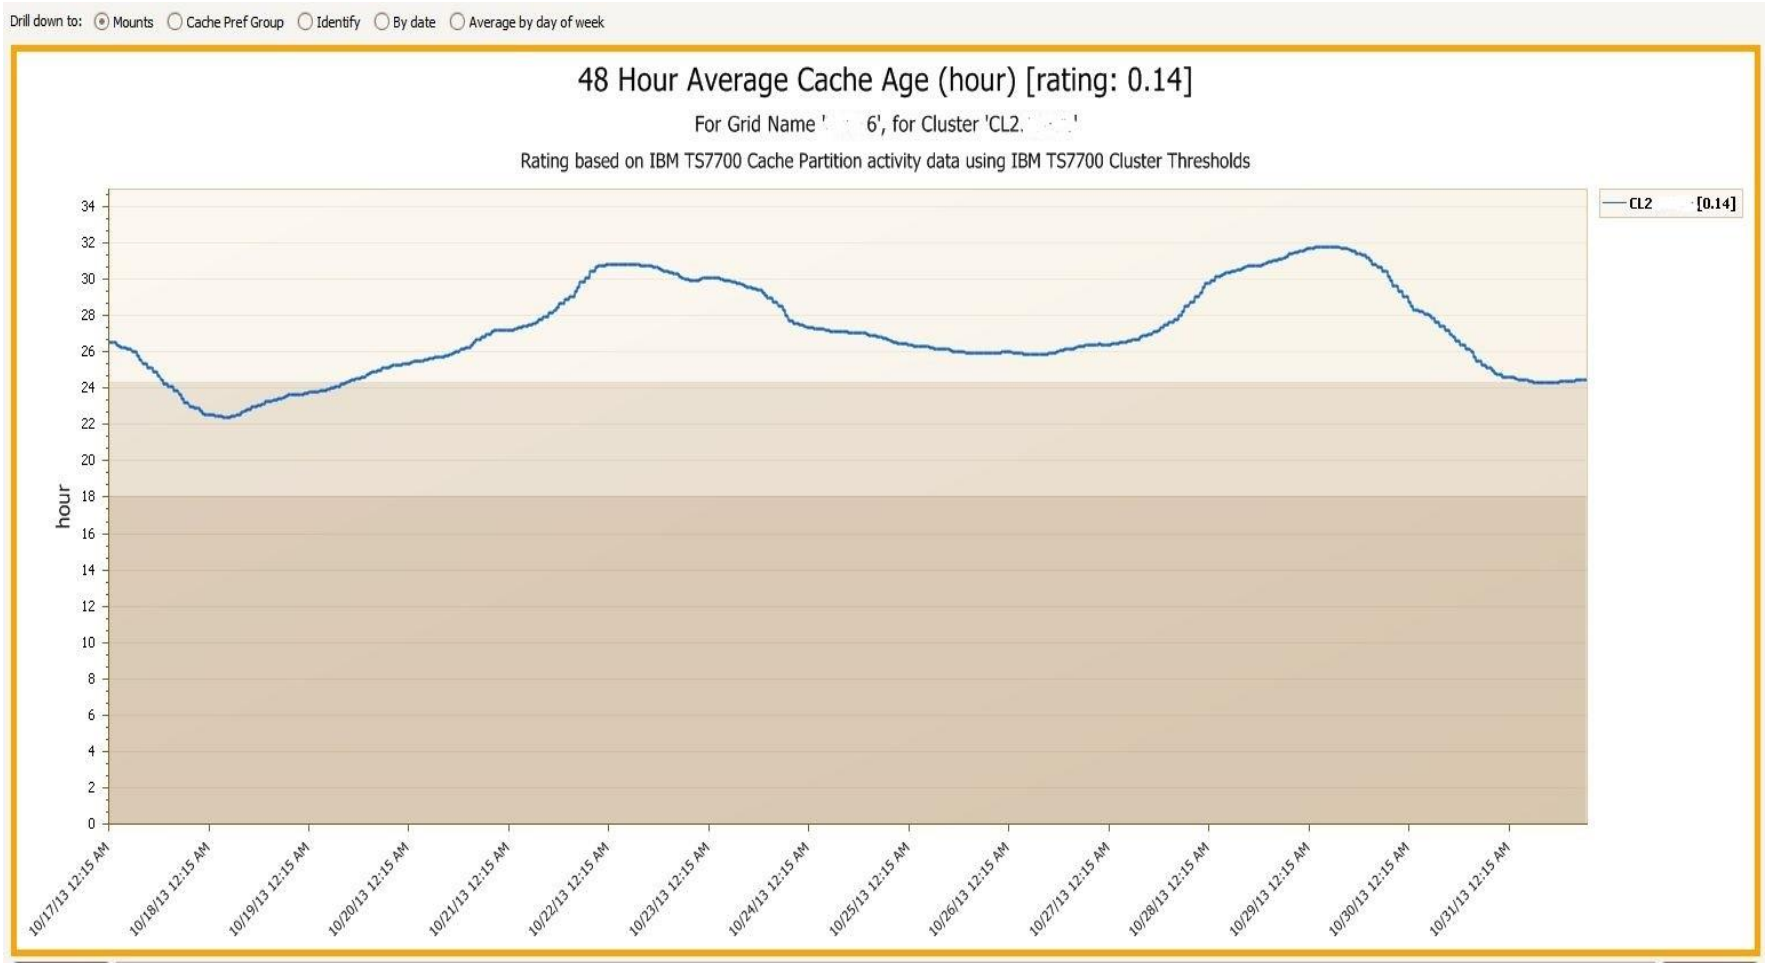

Oracle STK VSM

Presenter

John Ticic – Senior Technical Consultant

Started in Systems Programming in 1984.

Joined IntelliMagic in 2008 as a Senior Consultant

Specialties include:

- Disk/Tape performance
- z/OS Performance
- z/OS, zSeries implementation
- Presenting (I/O classes, SHARE, GSE,..)

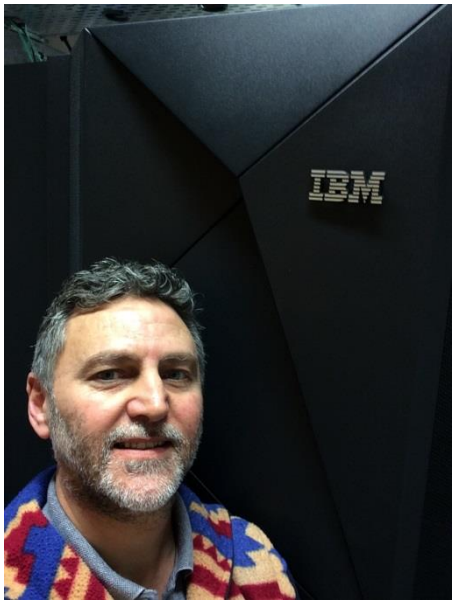

VSM Technology

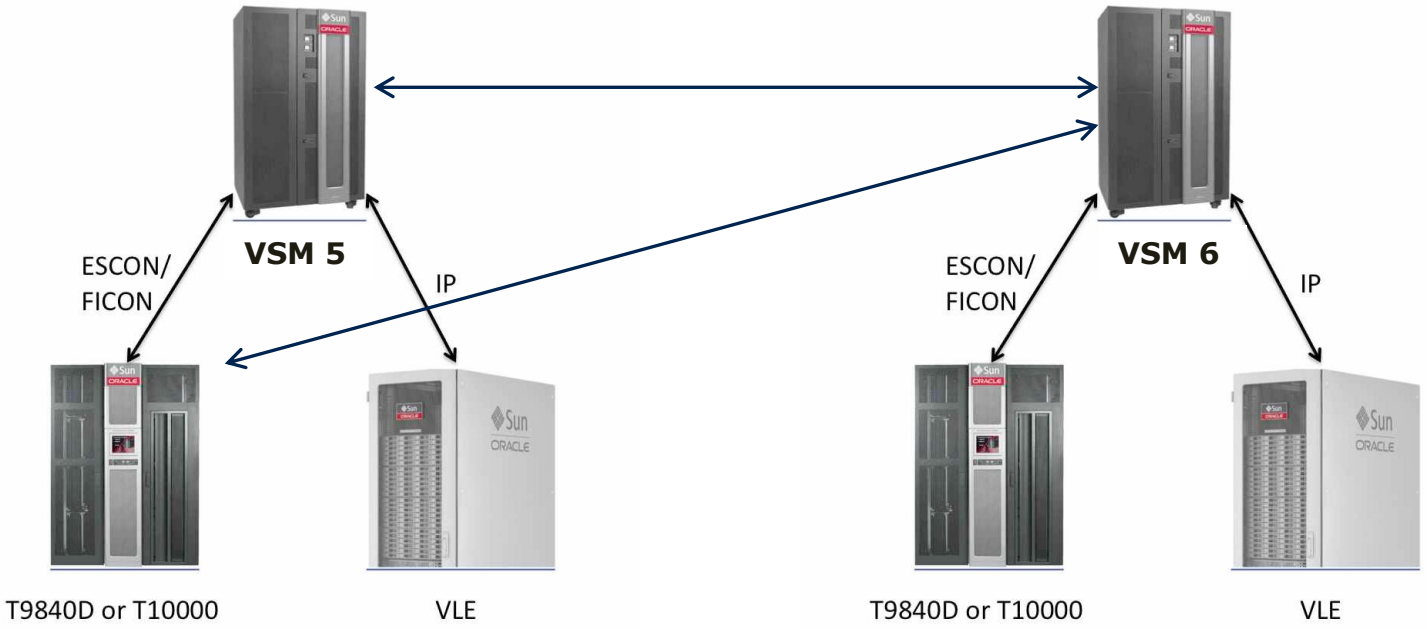

Different generations of hardware.

Different methods of replicating tapes.

Lot's of information in the STK user SMF records.

Why are some of my Jobs running slowly?

Why are some of my Jobs running slowly?

Yesterday, some batch Jobs took much longer! Why?

Well lots of possible reasons:

- Application changes
- Processing more data
- CPU (or storage) resource shortages
- Had to wait for devices
- Had to wait for volumes
- I/O contention

Let’s investigate tape mounts.

Mounts

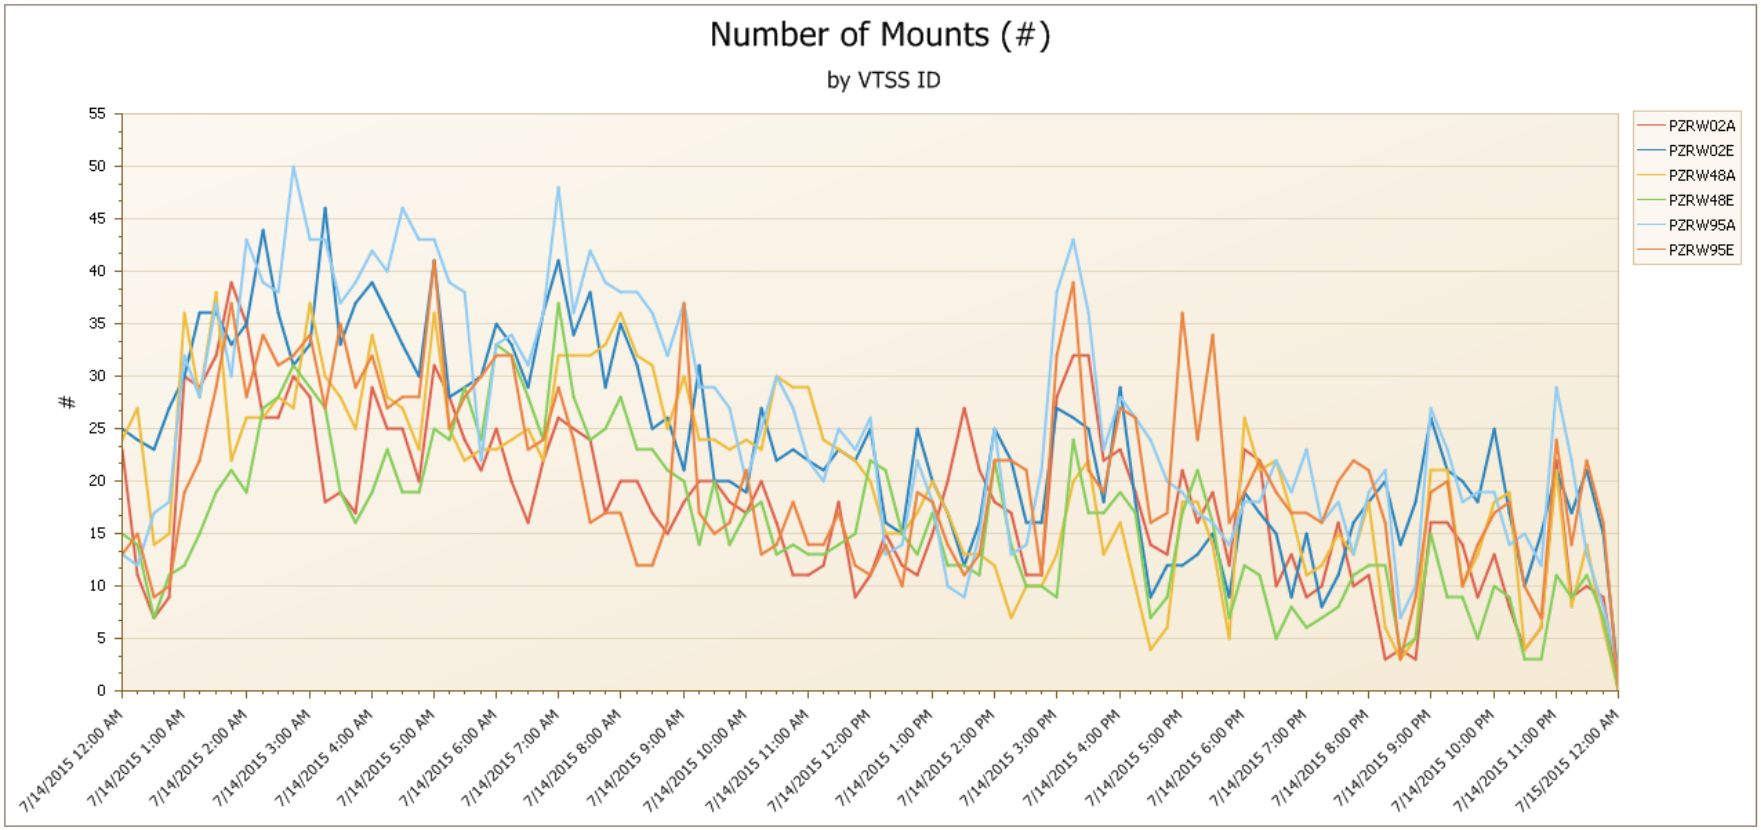

We can see our mount distribution (6 x VSM 6).

What are our Mount times like?

Average Mount Times

Drill down to: (•) Mount Types ( ) Mgt Class ( ) Job ( ) Systems ( ) Mount Events ( ) VTV at time ( ) By date ( ) Average by day of week ( ) Isolate

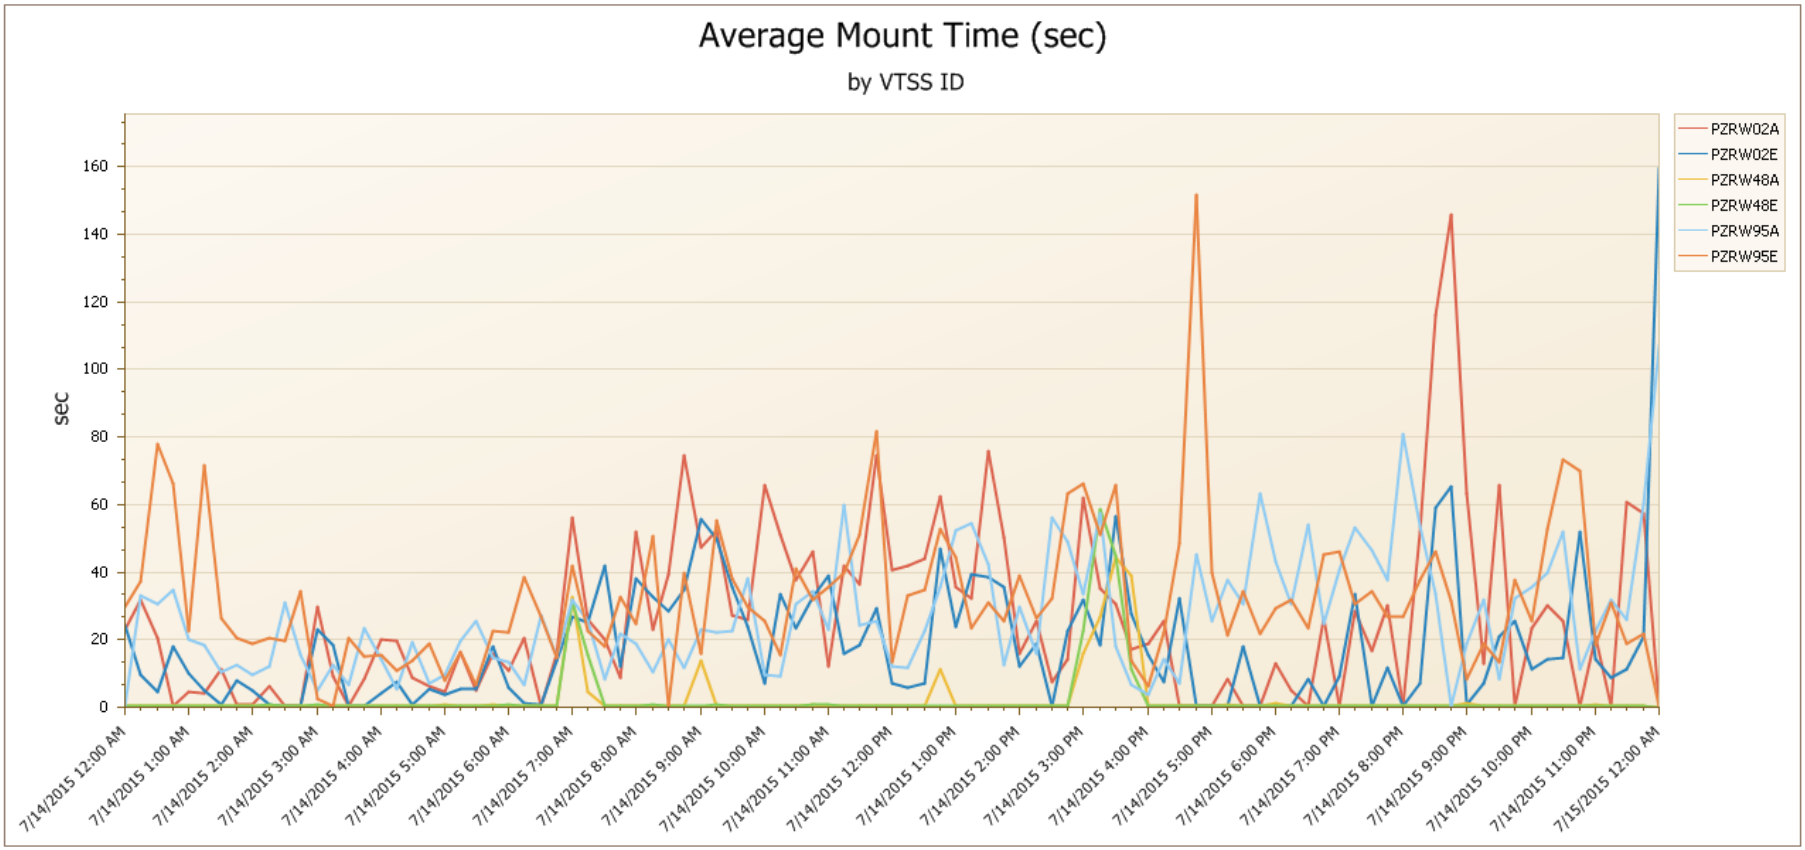

These are average times.

We can look at the maximums, but let's zoom into one VSM.

Average Mount Times

Drill down to: (•) Mount Events ( ) VTV at time

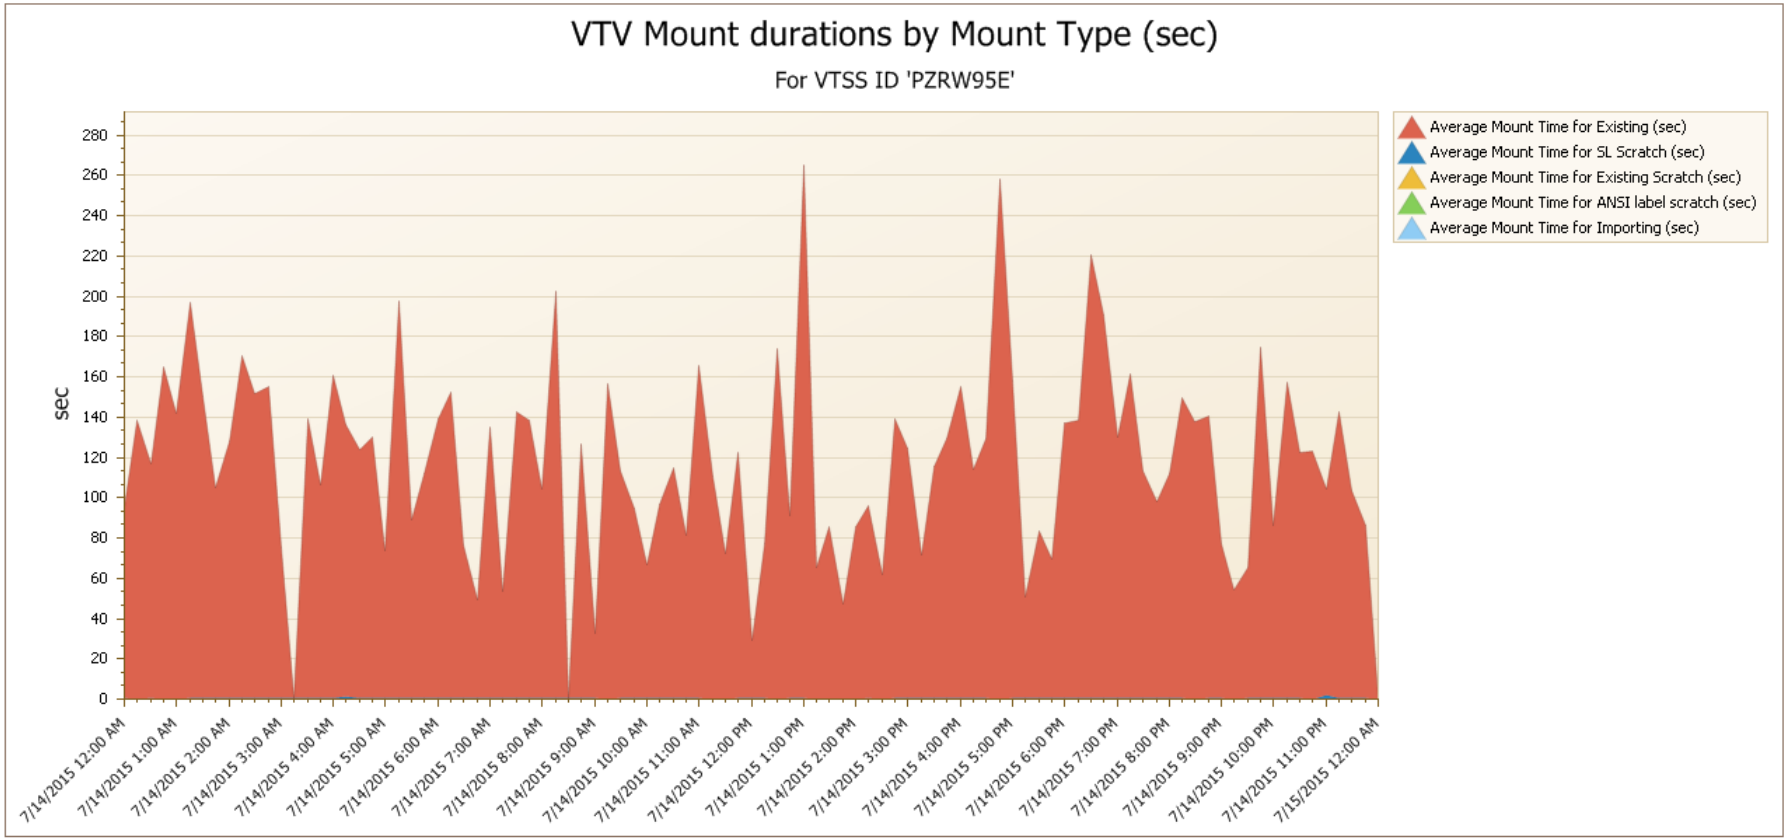

These are average times per Mount type.

Mounts for scratch tapes are almost not visible, but mounts for existing tapes need to be staged from real drives

VTV Mount Times

VTV Mounts by volser: For VTSS ID 'PZRW95E' by VTV volser ID

Drag a column header here to group by that column

VTV volser ID	VTV volser ID	Originat...	VTV Man...	VTD dev...	Virtual mount type	MVS jobname	MVS stepn...	MVS data set name	Mount Time...
0EPZWE	0EPZWE	D223	GBAA	F3AB	mount existing VTV	DFHSM	DFHSM	XXK06444.P301231.LF124.CFNWI.L5800073	830.89
0ERTTW	0ERTTW	D223	GBAA	F3D7	mount existing VTV	DFHSM	DFHSM	ZMS93254.P307076.LF673.FBKMO.L9834454	614.59
0ERXKT	0ERXKT	N223	GBAA	F3FC	mount existing VTV	OAM	OAM	MCK.HPZVCQJ.DDOC.JMCZGW63	376.64
0EZTII	0EZTII	N223	GBAA	F3BF	mount existing VTV	OAM	OAM	MCK.HPZVCQJ.DDOC.JMCZGW63	373.55
0EKRRT	0EKRRT	N223	GBAA	F3C2	mount existing VTV	OAM	OAM	ZCQ.UOSYCRP.CCXC.JMCZGW22	356.80
0ERTYG	0ERTYG	N223	GBAA	F39D	mount existing VTV	OAM	OAM	ZCQ.UOSYCRP.CCXC.JMCZGW22	301.60
0ADTKJ	0ADTKJ	D223	GBAA	F3C9	mount existing VTV	DFHSM	DFHSM	MHL26542.P162264.LF792.FEGHI.L7571364	277.77
0ELBEH	0ELBEH	D223	GBAA	F3CC	mount existing VTV	DFHSM	DFHSM	UVO46855.P146644.LF449.CCOTS.L1407053	261.88
0EICXK	0EICXK	N223	GBAA	F3E2	mount existing VTV	OAM	OAM	ZCQ.UOSYCRP.CCXC.JMCZGW22	168.08

We can look at specific volumes.

For example:

VTV 0EPZWE (DFHSM) is taking 830 seconds to mount.

Let’s look at some more details for this volume.

Job Details

VTV Identify

For vtssid 'PZRW95E', for Volume Label '0EPZWE'

VTV Identify Mounts: For Volume Label '0EPZWE'

Time	Volume Label	Data Set Name	System	Jobname	z/OS Progra...	Alloc...	Blocksize ...	Total Tape Block...	Read...	Read ...	Write ...	Devic...	Devic...	Sp...	Non-specific ...
7/14/2015 5:01 PM	0EPZWE	HSMPROD.HMIGTAPE.DATASET	D223	DFHSM	BJBDPK	836	16384	649189	0.00	0.81	0.00	0.15	0.00	1	0

VTV Identify Mounts: For VTSS ID 'PZRW95E', for VTV volser ID '0EPZWE'

Time	VTSS ID	Ori...	VTSS ID	VT...	Virtual mount t...	VTV timestamp	Recall indicator	MVS ...	MVS ...	MVS data set name	Mount start	Mount end	VTV ...
7/14/2015 5:01 PM	PZRW95E	D223	PZRW95E	F3AB	mount existing VTV	7/14/2015 5:01 PM	mounted after a recall	DFHSM	DFHSM	SYS15195.T164720.RA000.DFHSM.R0697100	7/14/2015 4:47 PM	7/14/2015 5:01 PM	GBAA

VTV Identify Dismounts: For VTSS ID 'PZRW95E', for VTV volser ID '0EPZWE'

Time	VTSS ID	Ori...	VTSS ID	VTV state	MV...	Uncom...	Th...	The last time th...	Virtual mount ty...	MVS j...	MVS ...	M...	VTV...	Synchronous Replication Status	RUN Received
7/14/2015 5:01 PM	PZRW95E	D223	PZRW95E	VTV dismounted	F3AB	10593.28	3816	7/7/2015 1:10 AM	mount existing VTV	DFHSM	DFHSM		GBAA	Synchronous replication not requested	7/14/2015 5:01 PM

VTV Identify Recall: For VTSS ID 'PZRW95E', for VTV volser ID '0EPZWE'

Time	VTSS ID	HSC S...	VTSS ID	R...	MVC vol...	R...	Aver...	N...	The last time the VT...	Rec...	Recall start	Recall end	VTV M...	MVC St...	RTD ...	MVS a...
7/14/2015 5:01 PM	PZRW95E	O223	PZRW95E	8	192347	no	3816	4	7/7/2015 1:10 AM	auto	7/14/2015 4:57 PM	7/14/2015 5:01 PM	GBAA	LACFE	0003	0F32

Detailed information about the tape activity from both z/OS (SMF 14/15/21/30) and VSM.

Note: No replication information since no data was written.

Recall Details

VTSS ID	RTD ID	MVC volser ID	Average VTV media size (MiB)	Number of GiB recalled (GiB)
PZRW95E	8	192347	3816	4

Real Tape Drive #

Real Tape

Amount recalled

Amount needed

Recall start	Recall end	VTV Managemen...	MVC Storage Class	RTD Channel ...	MVS address of RTD
7/14/2015 4:57 PM	7/14/2015 5:01 PM	GBAA	LACFE	0003	0F32

3:41 Mins

Is this too long?

Average Recall Time

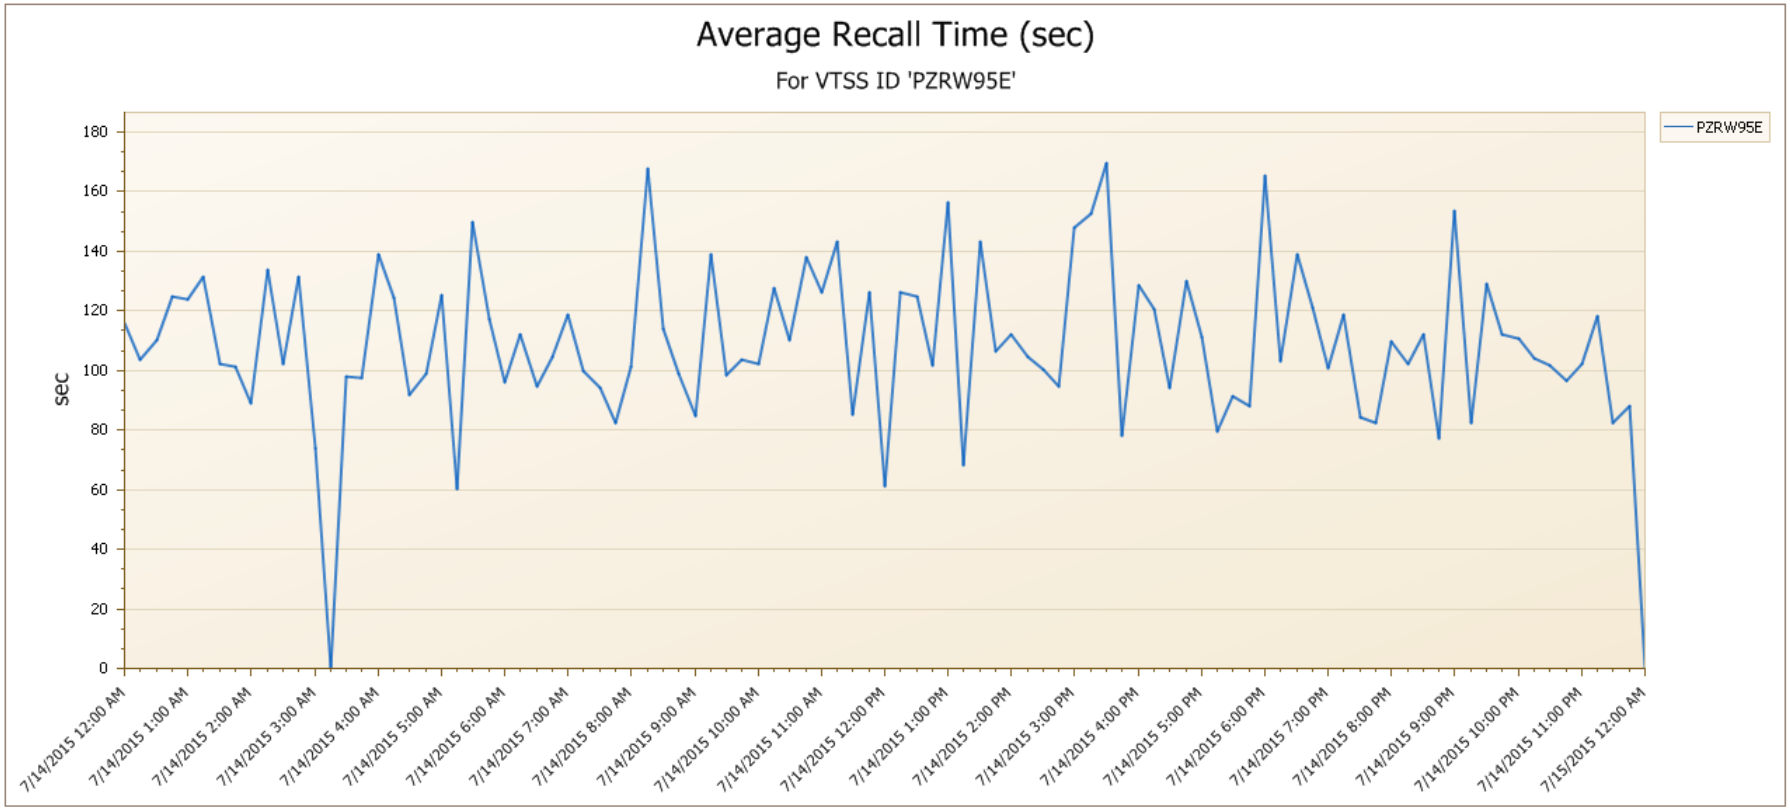

On average, the times look ok.

What about the peaks!

Maximum Recall Time

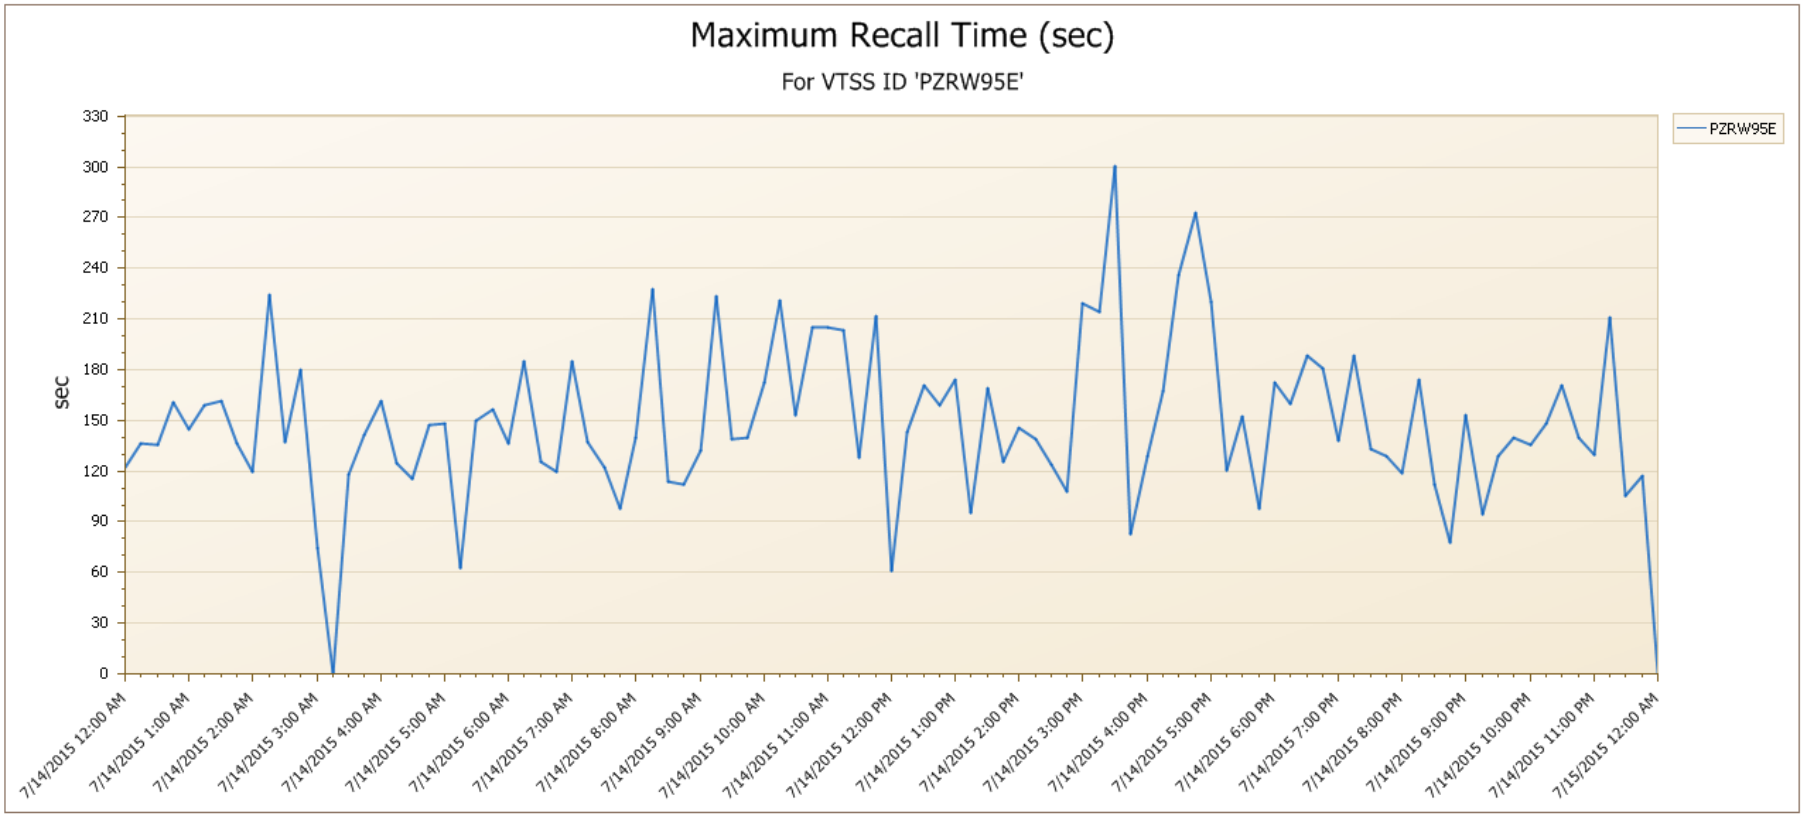

Yes, some peaks. We can look at the detailed records, but let's look at the mount distributions.

RTD Mount Time

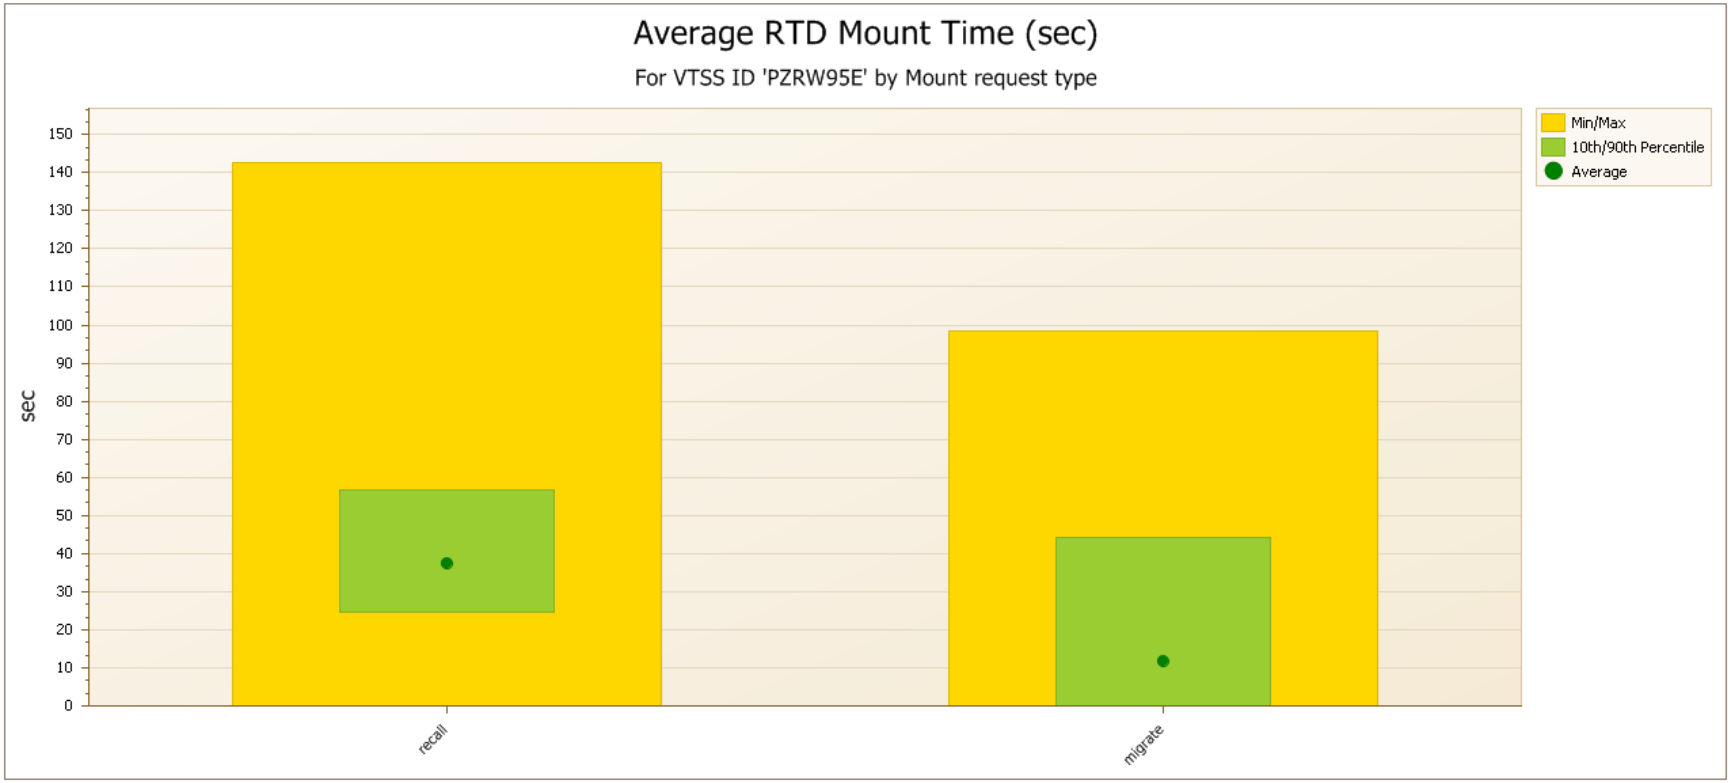

We see that this VSM is doing Recalls and Migrates. Let's look at RTD (Real Tape Device) #8 in detail.

Specific RTD Activity

RTD Mount events: For VTSS ID 'PZRW95E', for RTD ID '8'

Drag a column header here to group by that column

Time	VTSS ID	HSC S...	VTSS ID	MVC vol...	Actual v...	Read/write state	Mount r...	Mount start	RTD M...	MVC St...	RTD ...	MVS ...	RTD Loc...	Device ...	TapePlex Name
7/14/2015 6:09 AM	PZRW95E	O223	PZRW95E	788167	788167	read/write state	migrate	7/14/2015 6:08 AM	34.60	LACFE	0003	0F32	01:00	T1D34	YAHVHJE4
7/14/2015 6:23 AM	PZRW95E	O223	PZRW95E	804203	804203	read/write state	recall	7/14/2015 6:22 AM	55.57	LACFE	0003	0F32	01:00	T1D34	YAHVHJE4
7/14/2015 6:57 AM	PZRW95E	O223	PZRW95E	383950	383950	read/write state	recall	7/14/2015 6:56 AM	27.26	LACFE	0003	0F32	01:00	T1D34	YAHVHJE4
7/14/2015 7:20 AM	PZRW95E	O223	PZRW95E	300576	300576	read/write state	recall	7/14/2015 7:19 AM	26.21	LACFE	0003	0F32	01:00	T1D34	YAHVHJE4
7/14/2015 7:25 AM	PZRW95E	O223	PZRW95E	198818	198818	read/write state	recall	7/14/2015 7:24 AM	25.17	LACFE	0003	0F32	01:00	T1D34	YAHVHJE4
7/14/2015 7:30 AM	PZRW95E	O223	PZRW95E	090628	090628	read/write state	migrate	7/14/2015 7:29 AM	25.17		0003	0F32	01:00	T1D34	YAHVHJE4
7/14/2015 10:38 AM	PZRW95E	O223	PZRW95E	300576	300576	read/write state	recall	7/14/2015 10:37 AM	26.21	LACFE	0003	0F32	01:00	T1D34	YAHVHJE4
7/14/2015 10:52 AM	PZRW95E	O223	PZRW95E	722969	722969	read/write state	recall	7/14/2015 10:52 AM	23.07	LACFE	0003	0F32	01:00	T1D34	YAHVHJE4
7/14/2015 11:38 AM	PZRW95E	O223	PZRW95E	176923	176923	read/write state	recall	7/14/2015 11:38 AM	27.26	LACFE	0003	0F32	01:00	T1D34	YAHVHJE4
7/14/2015 12:53 PM	PZRW95E	O223	PZRW95E	457070	457070	read/write state	recall	7/14/2015 12:52 PM	36.70	LACFE	0003	0F32	01:00	T1D34	YAHVHJE4
7/14/2015 1:18 PM	PZRW95E	O223	PZRW95E	473050	473050	read/write state	recall	7/14/2015 1:18 PM	26.21	LACFE	0003	0F32	01:00	T1D34	YAHVHJE4
7/14/2015 1:28 PM	PZRW95E	O223	PZRW95E	364785	364785	read/write state	recall	7/14/2015 1:27 PM	30.41	LACFE	0003	0F32	01:00	T1D34	YAHVHJE4
7/14/2015 2:59 PM	PZRW95E	O223	PZRW95E	722969	722969	read/write state	recall	7/14/2015 2:59 PM	24.12	LACFE	0003	0F32	01:00	T1D34	YAHVHJE4
7/14/2015 3:05 PM	PZRW95E	O223	PZRW95E	801710	801710	read/write state	recall	7/14/2015 3:04 PM	44.04	LACFE	0003	0F32	01:00	T1D34	YAHVHJE4
7/14/2015 3:15 PM	PZRW95E	O223	PZRW95E	818861	818861	read/write state	recall	7/14/2015 3:14 PM	61.87	LACFE	0003	0F32	01:00	T1D34	YAHVHJE4
7/14/2015 3:26 PM	PZRW95E	O223	PZRW95E	815442	815442	read/write state	recall	7/14/2015 3:26 PM	51.38	LACFE	0003	0F32	01:00	T1D34	YAHVHJE4
7/14/2015 3:34 PM	PZRW95E	O223	PZRW95E	012938	012938	read/write state	recall	7/14/2015 3:33 PM	61.87	LACFE	0003	0F32	01:00	T1D34	YAHVHJE4
7/14/2015 4:42 PM	PZRW95E	O223	PZRW95E	192347	192347	read/write state	recall	7/14/2015 4:42 PM	36.70	LACFE	0003	0F32	01:00	T1D34	YAHVHJE4
7/14/2015 4:49 PM	PZRW95E	O223	PZRW95E	192347	192347	read/write state	recall	7/14/2015 4:48 PM	36.70	LACFE	0003	0F32	01:00	T1D34	YAHVHJE4
7/14/2015 4:58 PM	PZRW95E	O223	PZRW95E	192347	192347	read/write state	recall	7/14/2015 4:57 PM	88.08	LACFE	0003	0F32	01:00	T1D34	YAHVHJE4
7/14/2015 7:08 PM	PZRW95E	O223	PZRW95E	722969	722969	read/write state	recall	7/14/2015 7:07 PM	24.12	LACFE	0003	0F32	01:00	T1D34	YAHVHJE4

RTD ID 8 is mainly busy with Recalls. There are occasional Migrates.

Specific RTD Activity

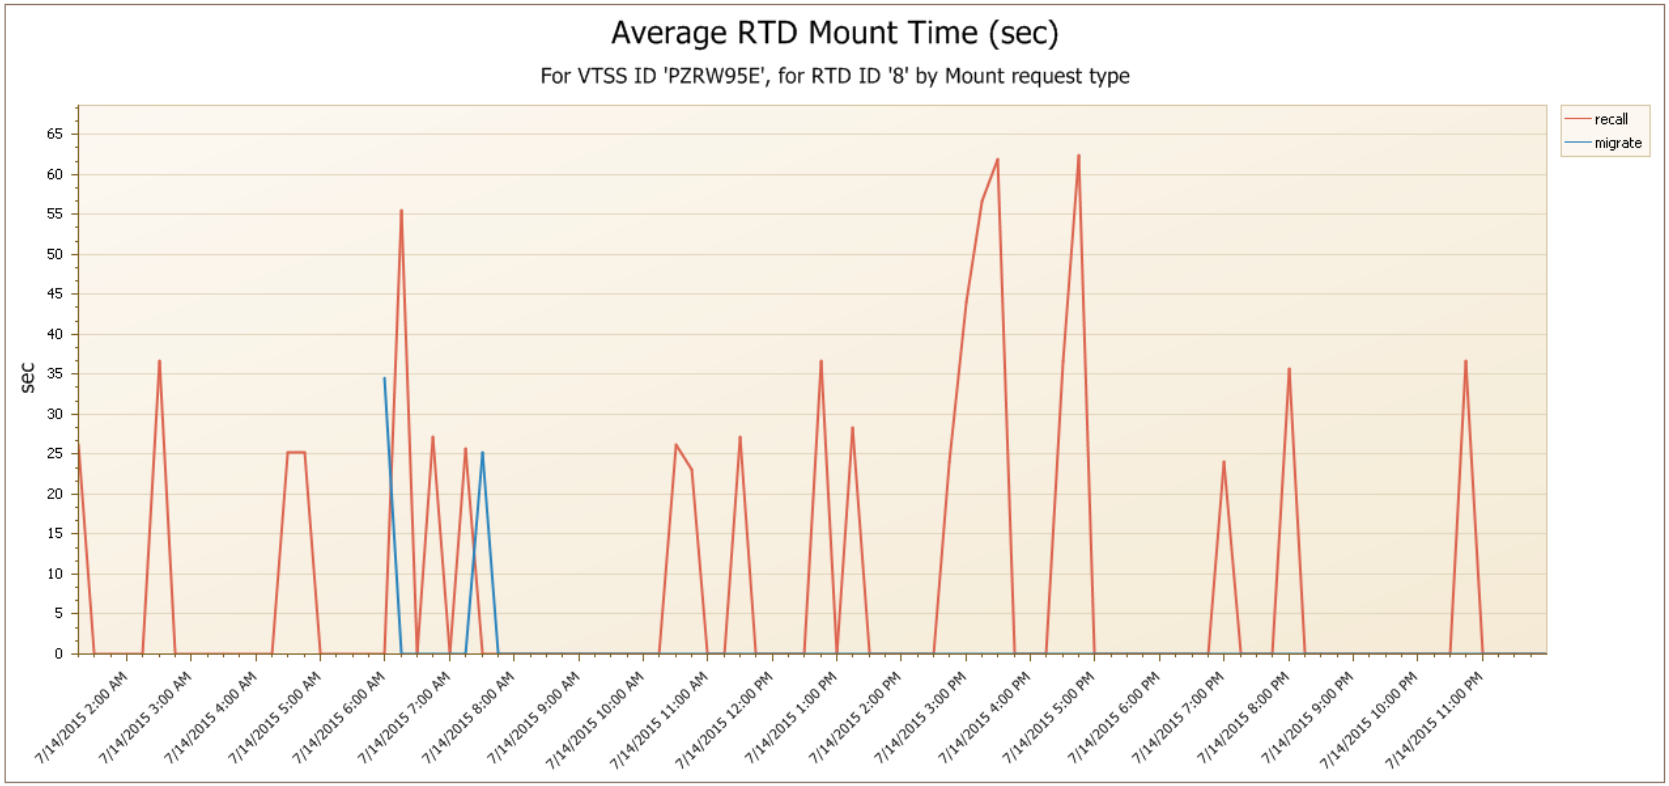

No apparent thrashing!

Summary

For very long mount times, there may be:

- Contention inside the VSM (large queues)
- Contention for RTDs (thrashing between Migrate, Recall, Reclaim)
- Robotic delays mounting the tape
- Delays positioning to the VTV on the MVC
- Media errors

Use of the SMF records highlights the possible cause.

How long is replication for my tapes taking?

Replication Challenges

- Minimize disruption to production Tape usage.
  - E.g. Should batch Jobs wait until Tape is fully replicated.
- Maximize Recovery Point Objective.
  - How much data loss can we accept.
- Minimize Recovery Time Objective.
  - How long until we are back up and running.

These decisions need to be made BEFORE a technology is selected and implemented.

Now the big question:

"How is my Tape replication running?"

How long is replication for my tapes taking?

So, you're replicating data synchronously.

How long is it taking?

Is it consistent during the day?

Are all volumes being replicated synchronously?

Interesting questions. Let's have a look.

Average Replication Time

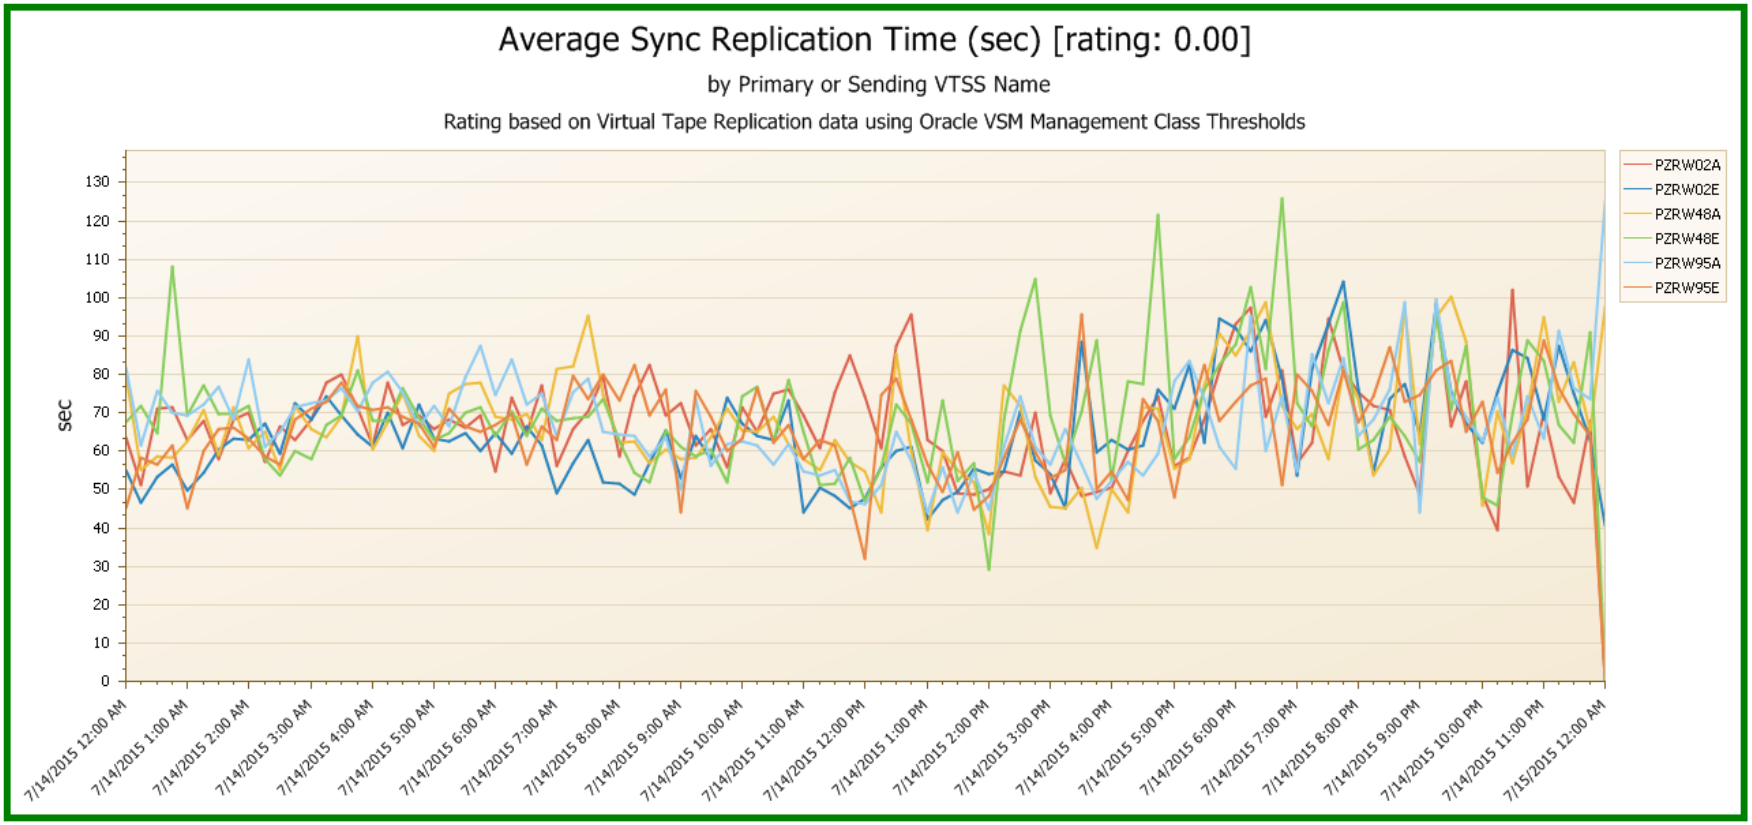

6 x VSM 6 systems, replicating synchronously.
Average time is around 70 seconds, a little more during the batch window.
This is the time per volume (VTV).

Average Replication Time (Normalized)

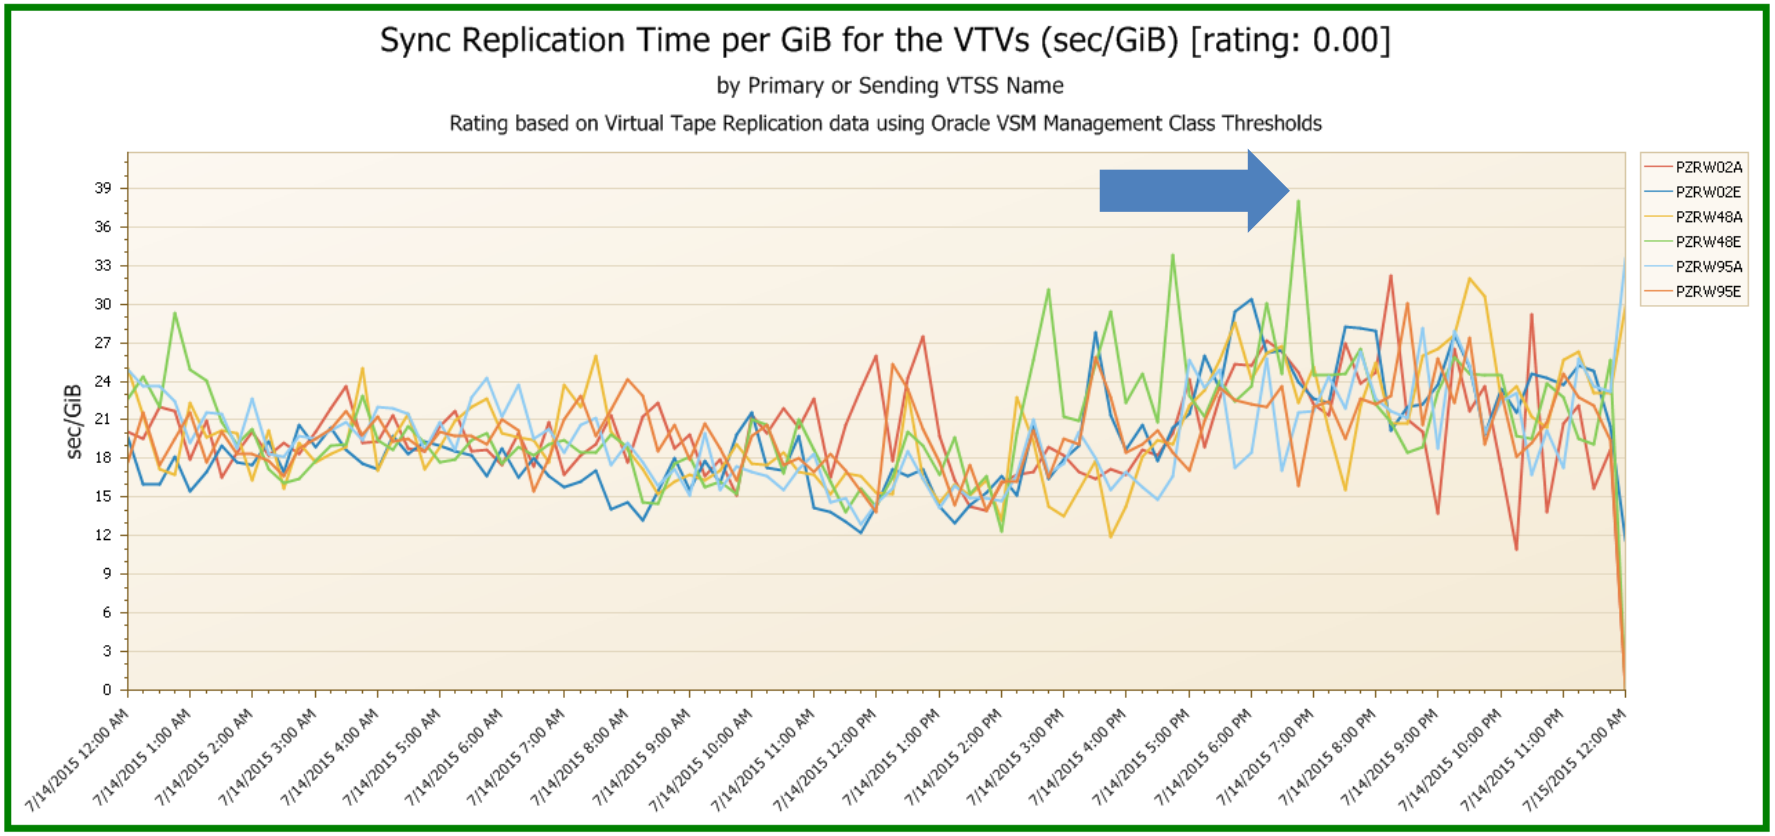

An average of 20 seconds per GiB.
It's taking longer during the batch window ☹
There are a few peaks ☹

Concentrate on a Specific VSM

Let’s concentrate on one VSM (the batch peaks).

VSM PZRW48E is taking longer to replicate at times.

How is this VSM doing?

Average Replication Time – PZRW48E

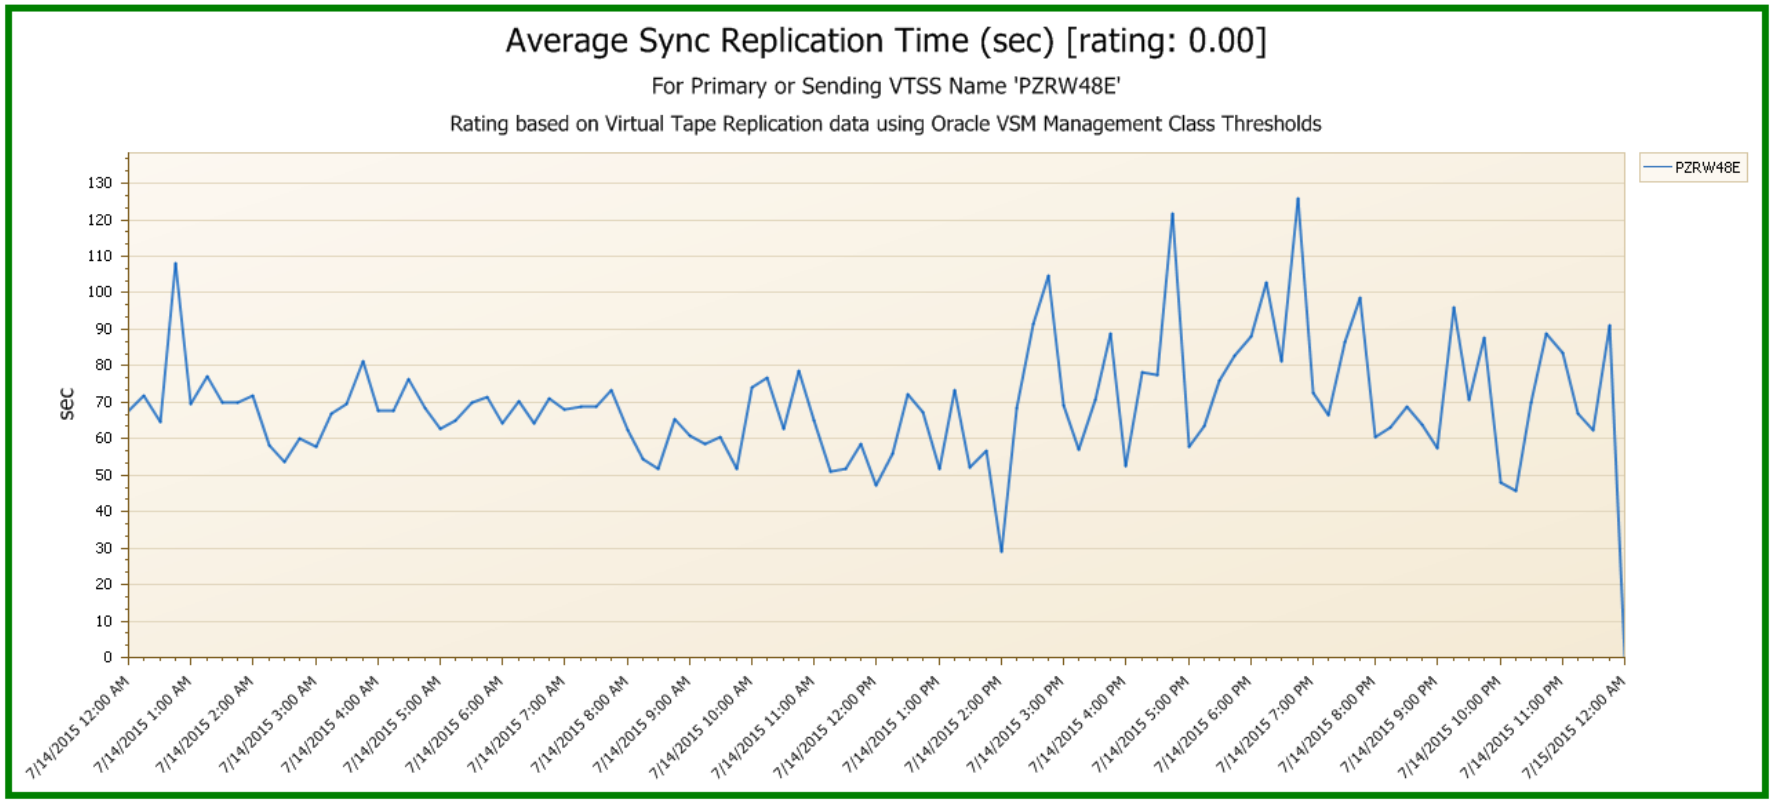

It certainly looks different later in they day.

Replication during the Day

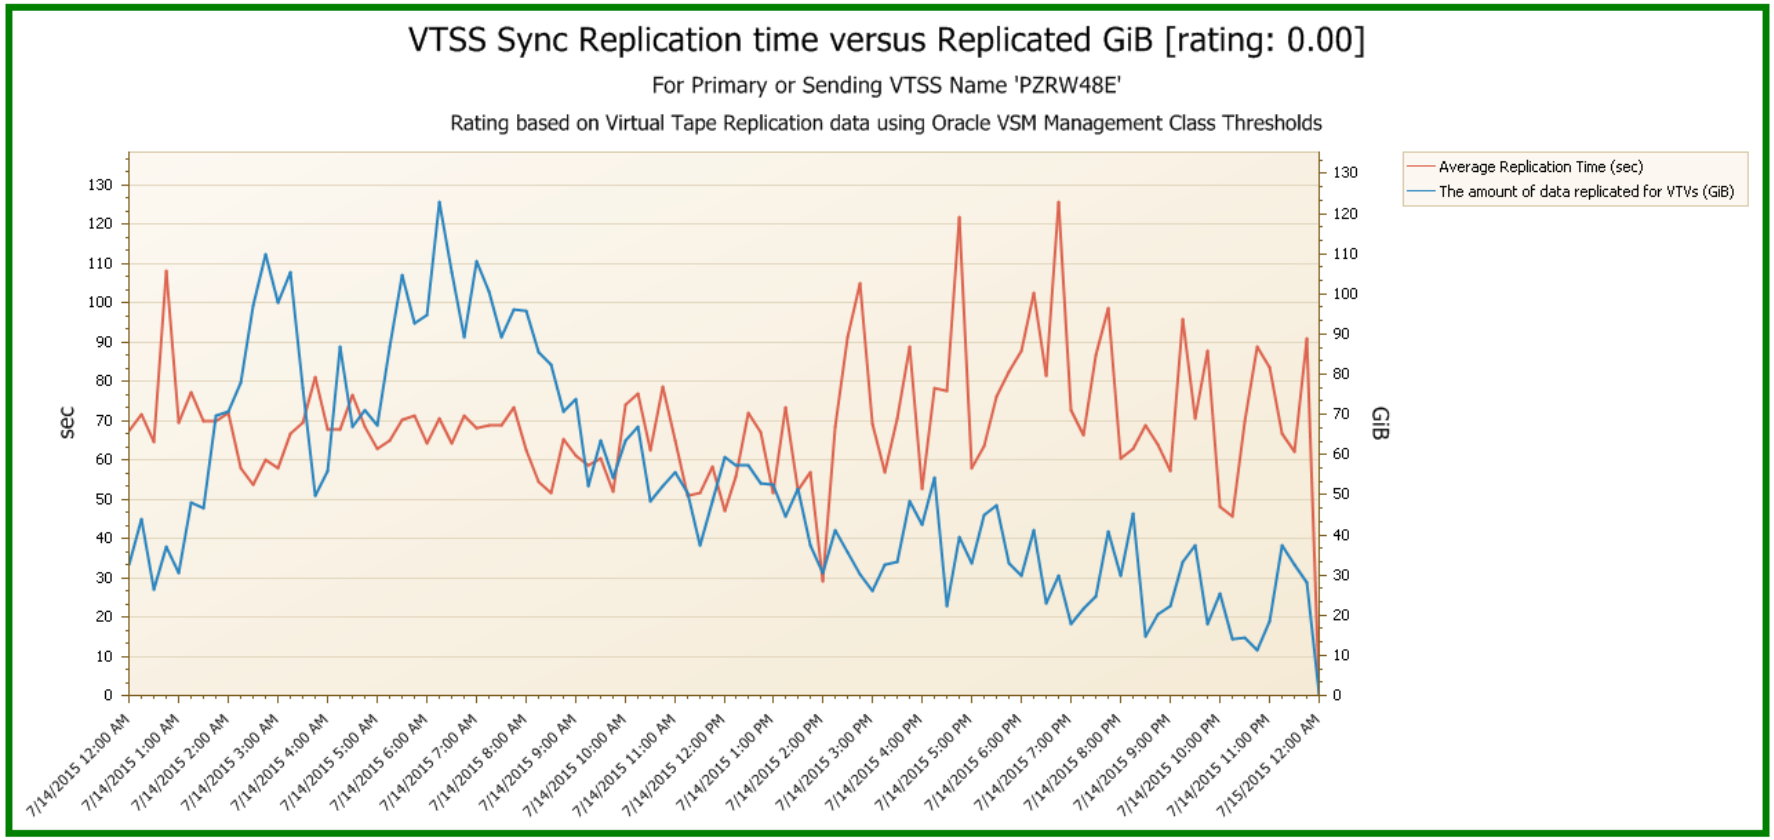

Replication seems to take longer when there is less data!

VTVs Replicated

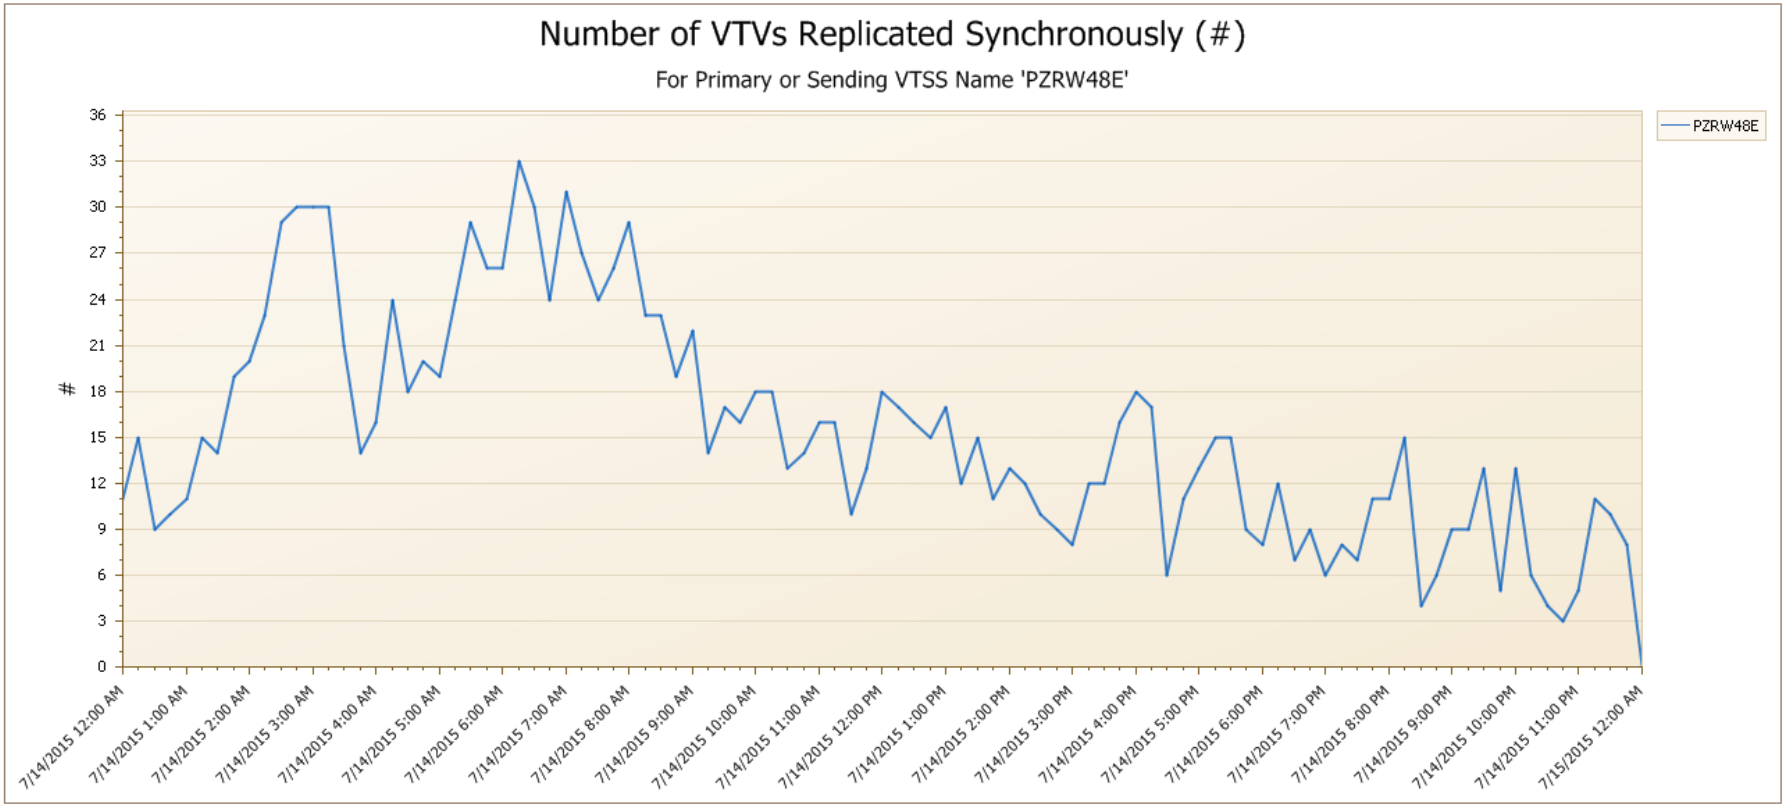

Many more VTVs replicated during the morning.

What else is this VSM up to?

Other Activity

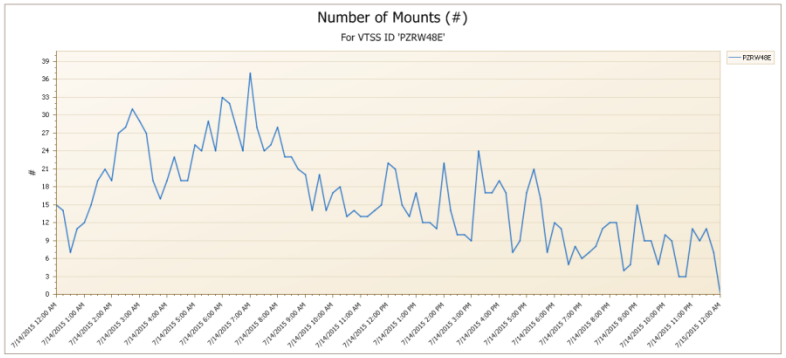

Front-end Mounts

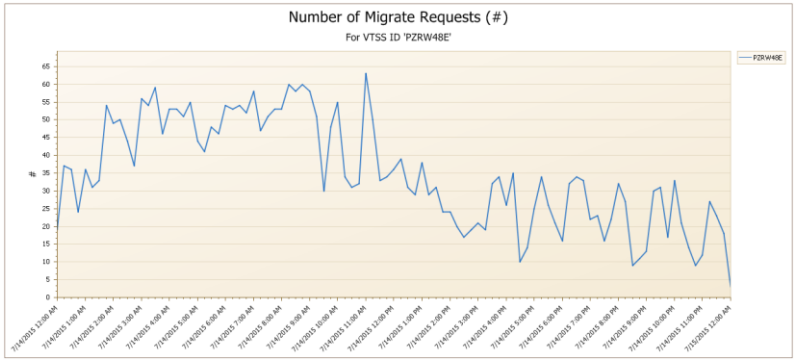

Migration

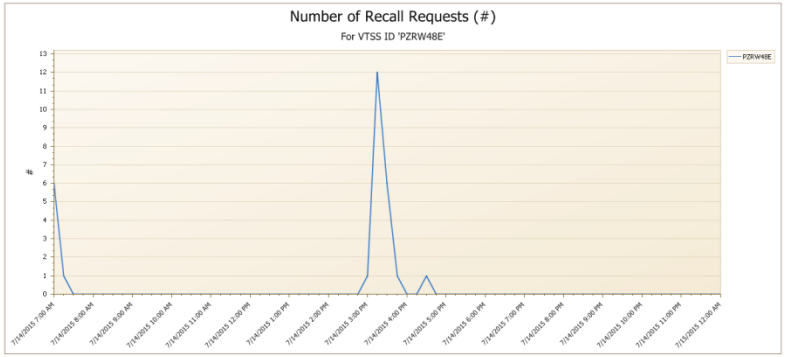

Recalls

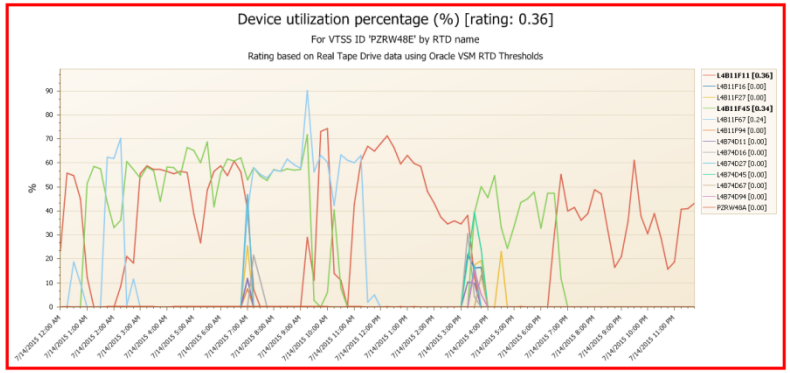

RTD activity/Utilization

We can also investigate the times for VTV Mounts, Migrates, Recalls and RTD mounts.

VTV Compression Factor

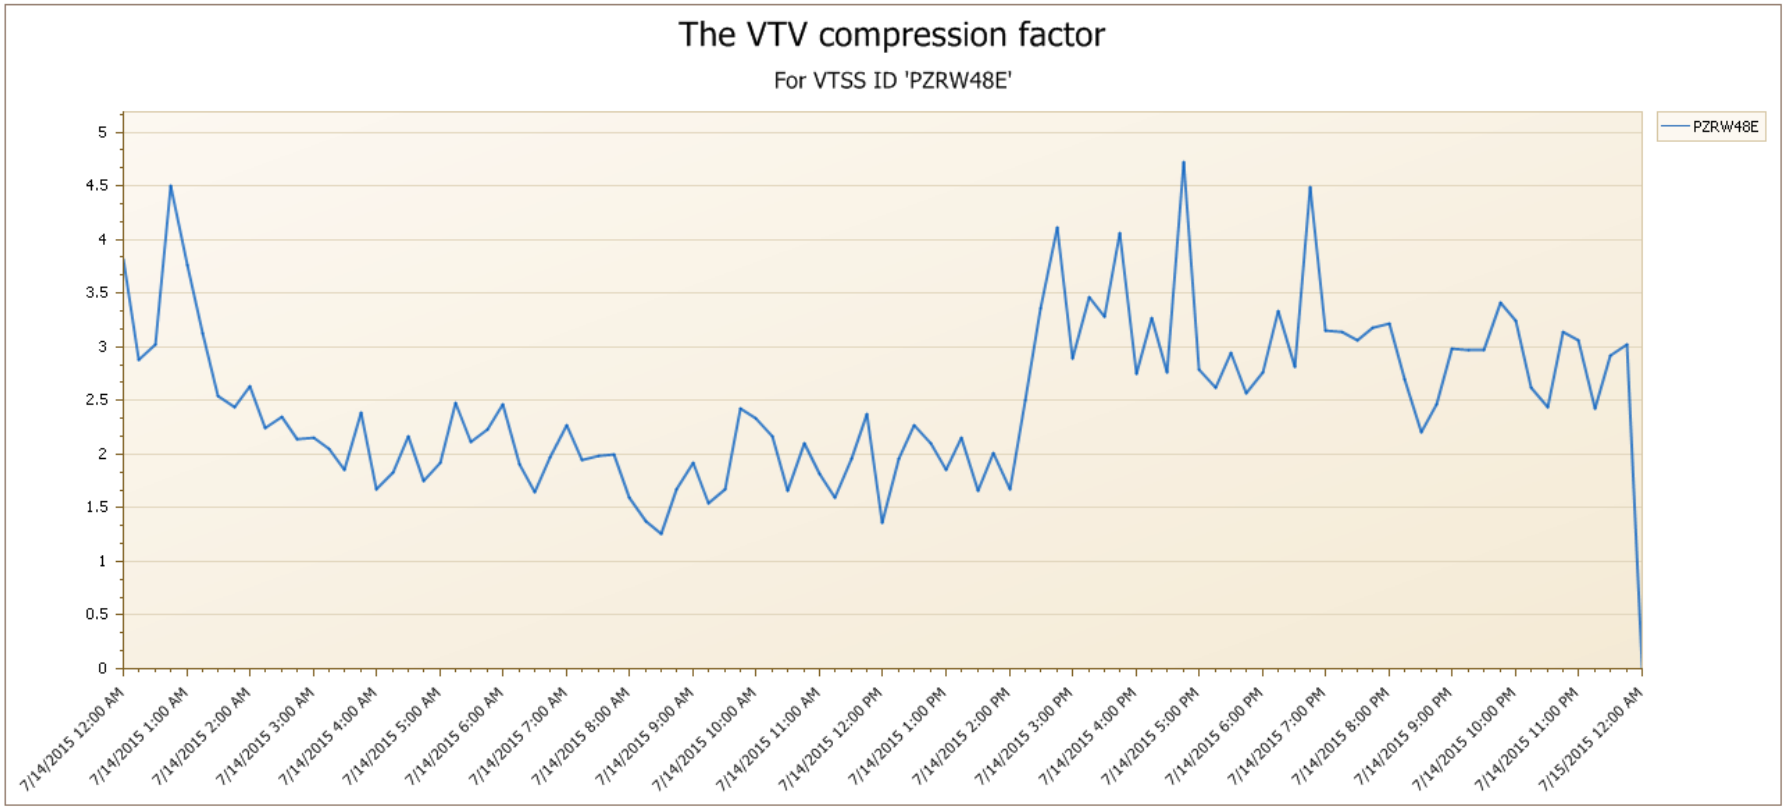

Some VTVs can compress more favorably.
(Also available from z/OS SMF records)

What's happening?

There doesn't seem to be a clear problem explanation.

Late in the day:

Fewer GiB to replicate
Fewer VTVs to replicate

More time per GiB
Higher compression factor

Maximum Replication Delay

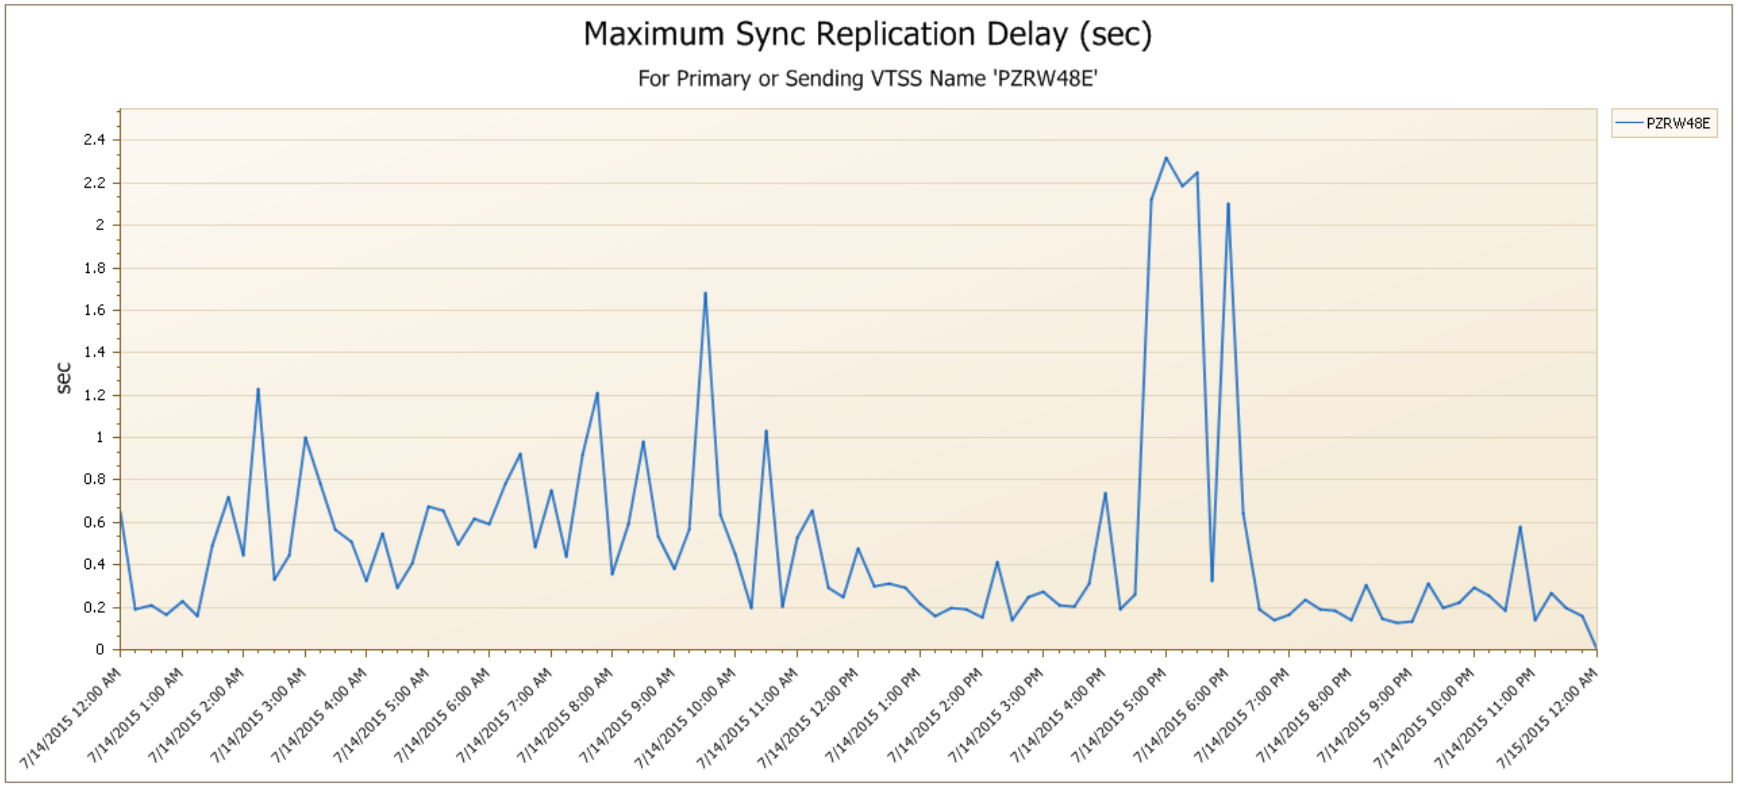

We'd been looking at the average replication time.
The maximum delay time certainly stands out.

Replication Delay Time

Replication delay time is the time between closing the Volume (Tape rewind indicator received from z/OS) and the start of the replication process.

This should normally be very low (less than 1 second.)

We have peaks approaching 2.5 seconds.

This VSM is probably quite busy internally at this time and this is probably resulting in a large queue for the replication tasks.

VSM Summary

These were two examples of what is important for VSM Tape operations, and how they can be investigated.

Performance data is available, but needs to be properly mined and presented.

Connecting the z/OS view to the virtual hardware view is critical to understand and manage a VSM environment.

z/OS Data

What can z/OS Supply?

Tape activity is recorded in:

SMF 14	Input
SMF 15	Output
SMF 21	Error Statistics by Volume (Mount)
SMF 30	Common Address Space
RMF	Various records for channels, devices, …

Bandwidth

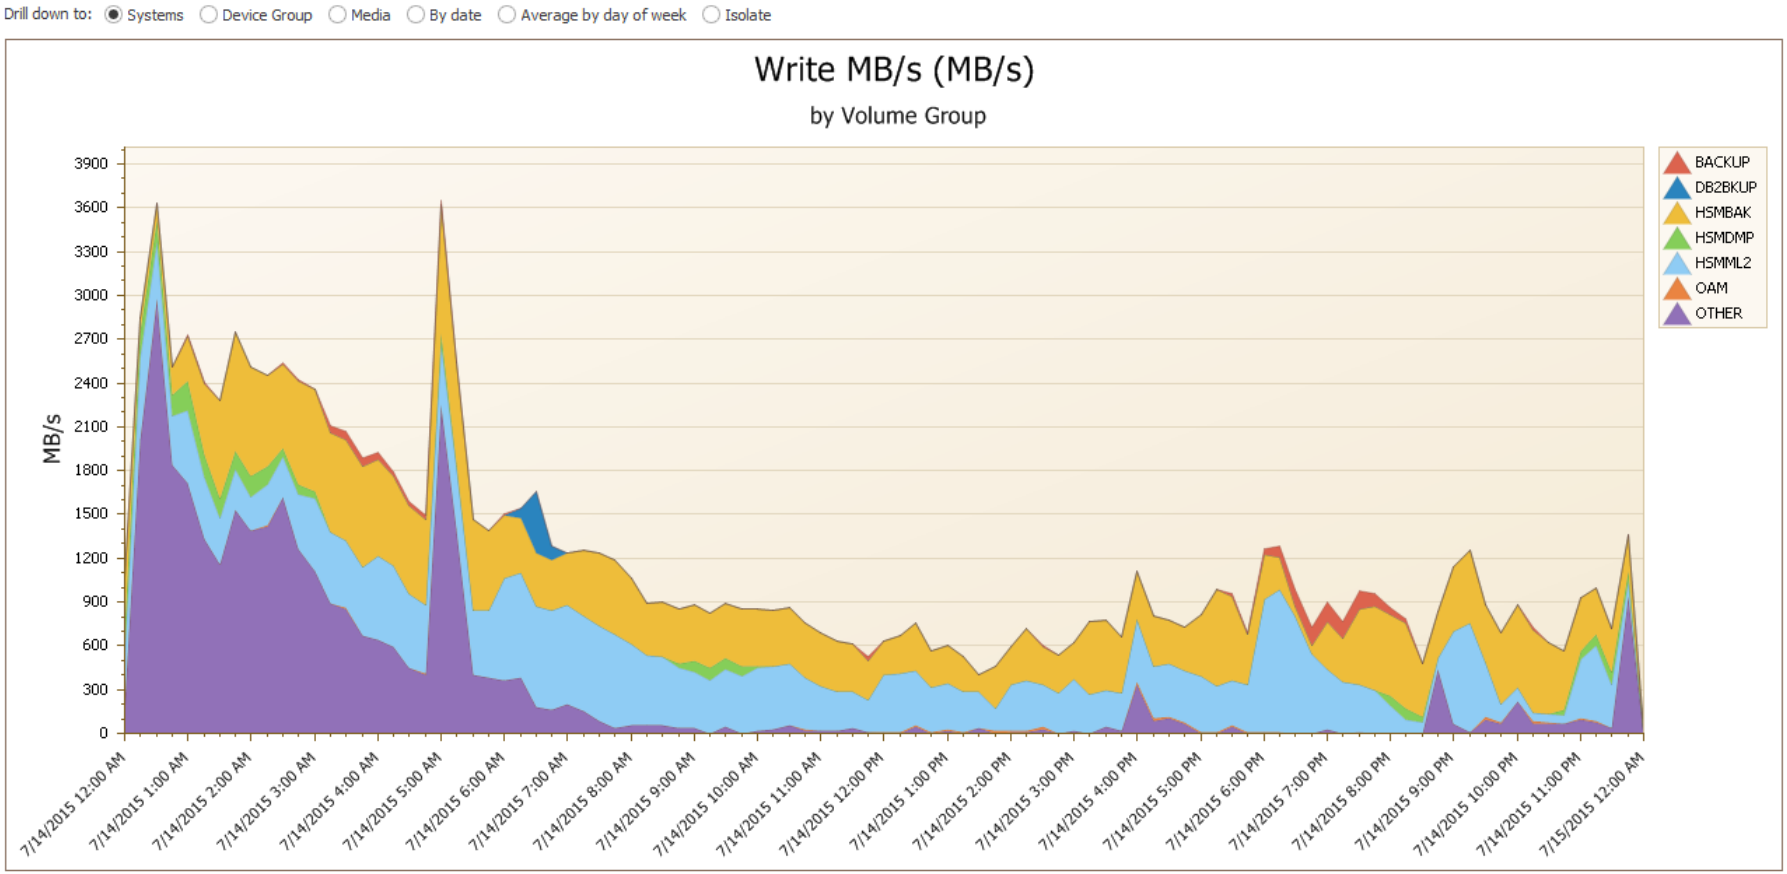

Based on z/OS data only, the write MB/s can be obtained (optionally, by group)

Device Throughput (Compressed)

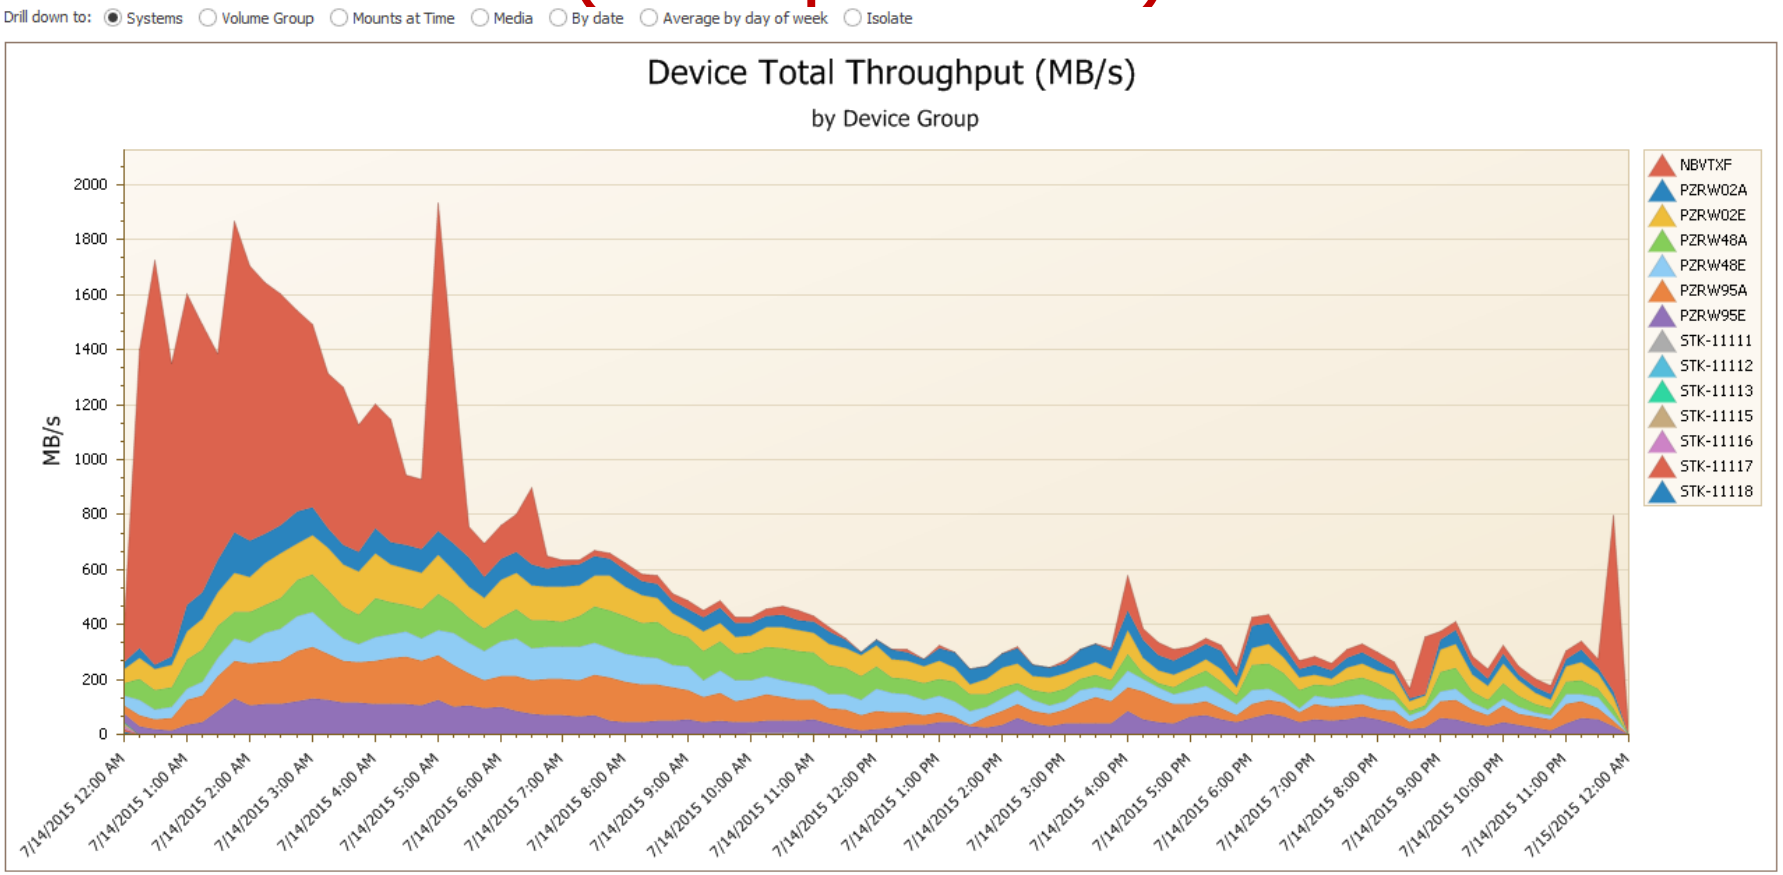

A breakdown of throughput by devices.

Very useful during technology migration.

Device Concurrency

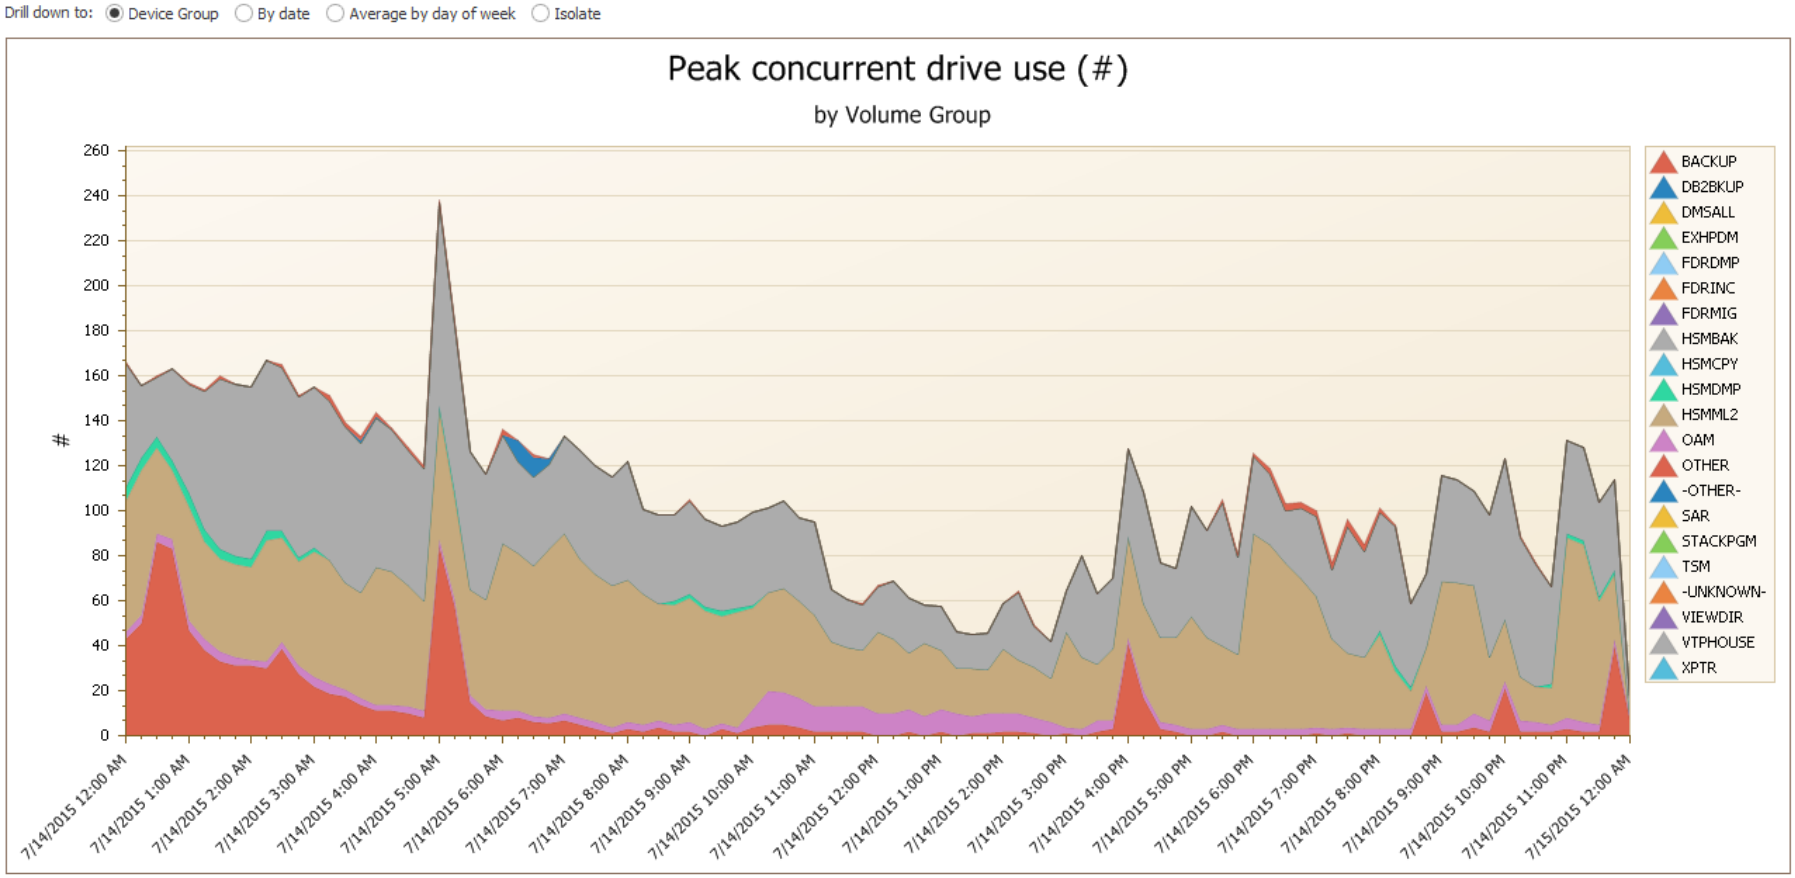

So how many tape devices do you need, and when.

Very relevant for Job scheduling.

Tape Usage

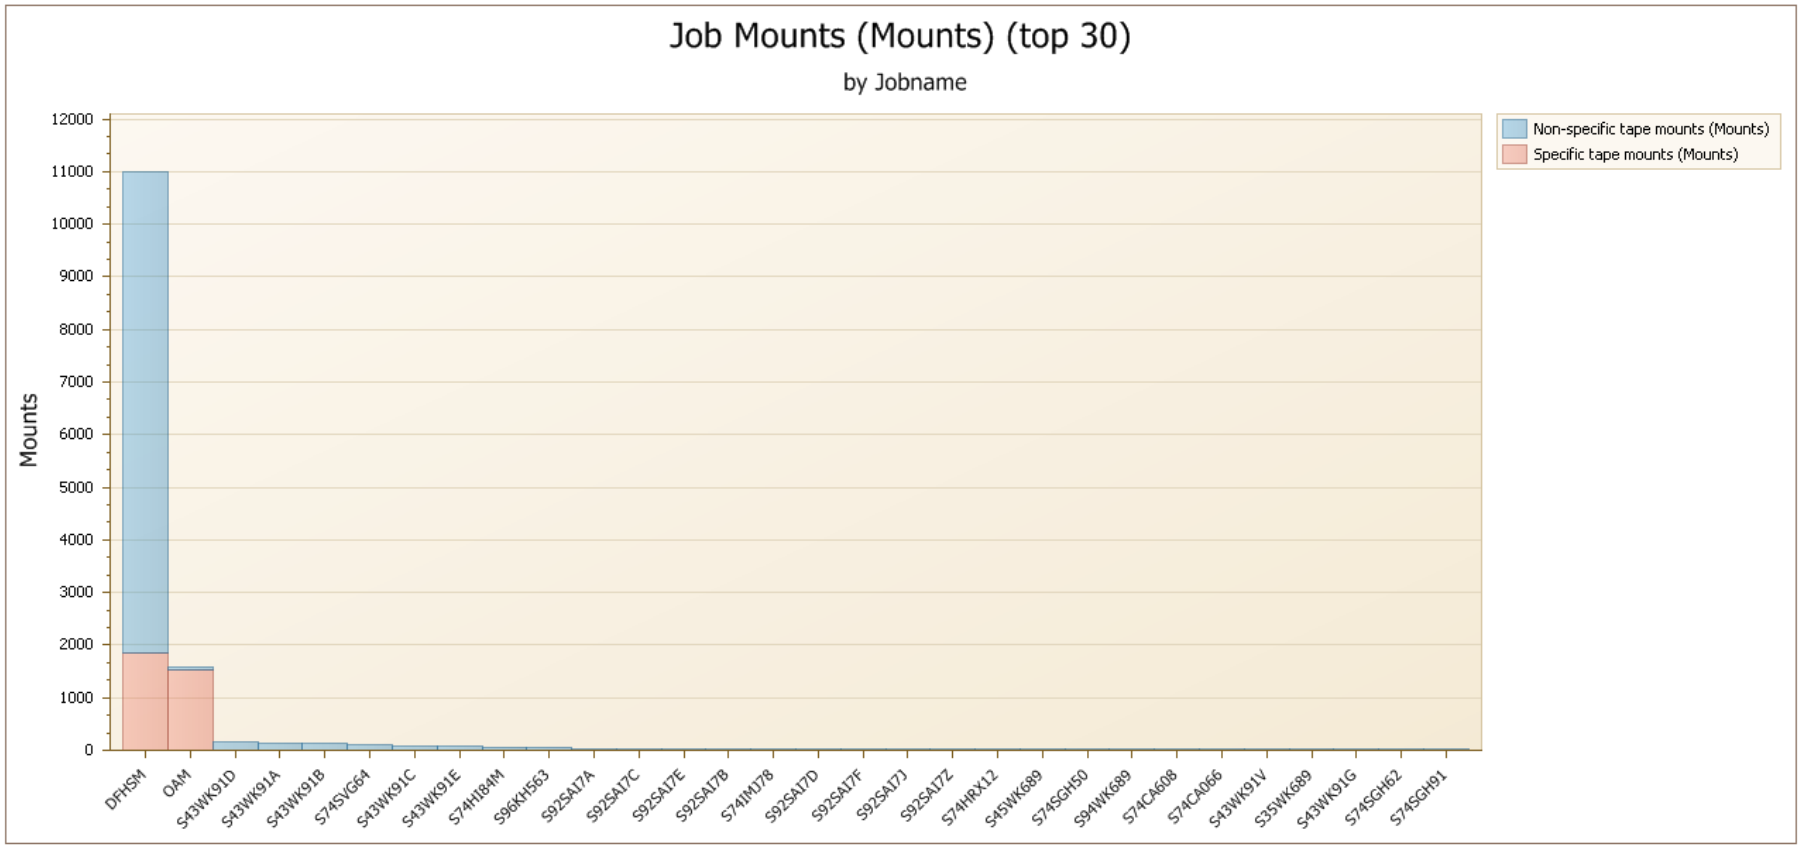

No surprises here!

z/OS Summary

Just looking at the z/OS specific SMF data can reveal very interesting information about tape processing.

- How much bandwidth do I need for replication
- How many tape devices per LPAR do I need
- Who are my major tape users, and when
- Investigate problems when they occur

Summary/Conclusions

Summary/Conclusion

Tape is not dead.

It’s just a little hard to see what is happening under the covers.

But, as we’ve seen today, there is information available. It is hard to manually process and interpret so you should implement reporting/performance tools

Whitepapers available at www.intellimagic.com

Thank You

www.intellimagic.com

Appendix

TS7700

Are my tape mount times reasonable?

Virtual Mounts vs. Mount Times

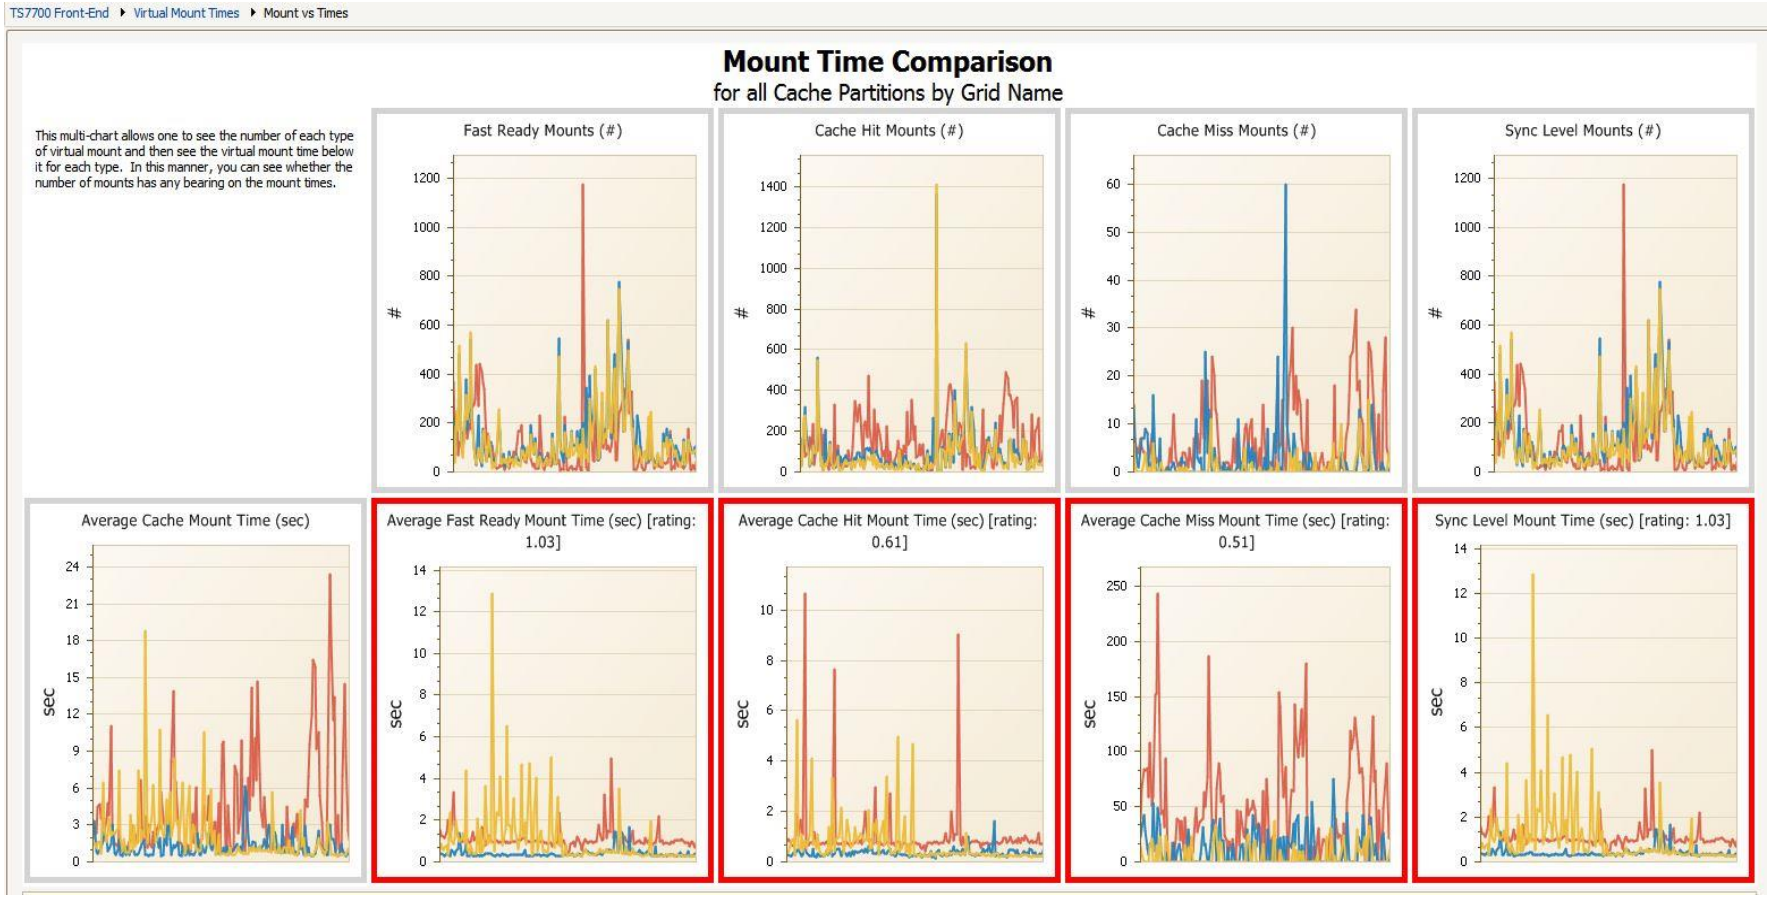

Are there enough back-end drives to support the migration, recall, and reclaim workloads?

Back-end Overview multi-chart

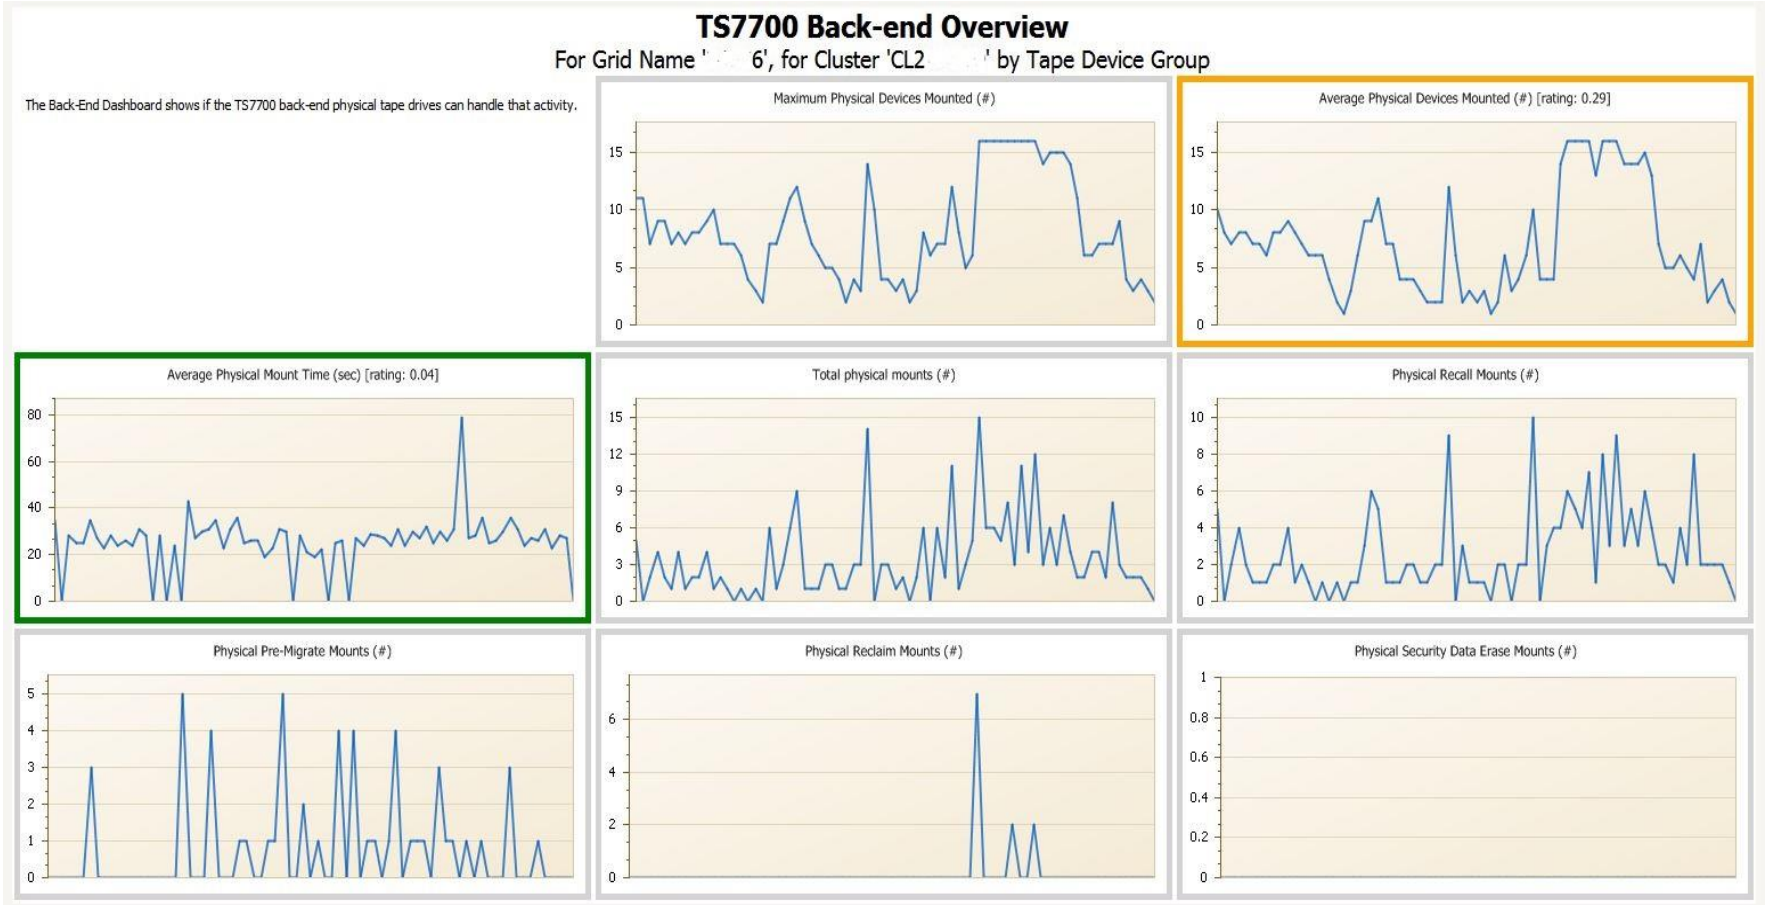

Migration Perspectives

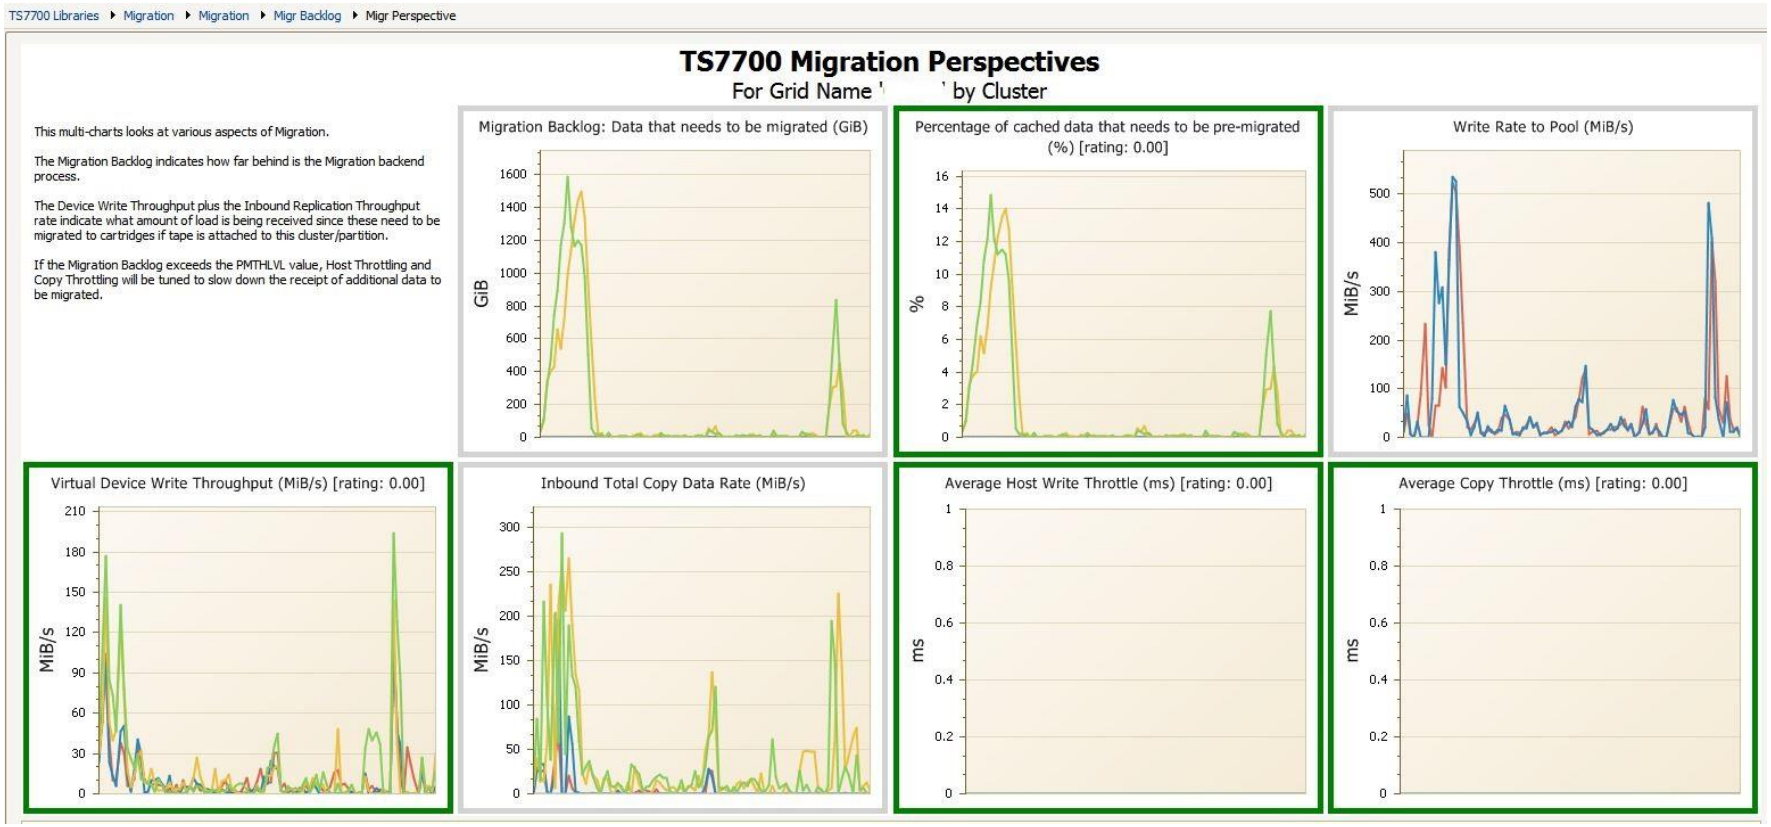

VSM

Oracle SMF Data

Subtype	Description
1	BLOS (LSM) Operation Statistics
2	Vary Station
3	Modify LSM Command
4	LMU Read Statistics
5	Cartridge Eject
6	Cartridge Enter
7	Move Detail
8	View Statistics
9	VTCS Configuration Change
10	VTSS subsystem performance
11	VTSS channel interface performance
13	VTV mount request
14	VTV dismount request
15	Delete VTV request
16	RTD mount request
17	RTD dismount request
18	Migrate VTV request

Subtype	Description
19	Recall VTV request
20	RTD performance request
21	Vary RTD
25	MVC status
26	VTV movement
27	VTV scratch status
28	VTV replication
29	VTV and MVC unlink event
30	Vary Clink event
31	Dynamically added/deleted transports
32	Internal use

Useful VSM Acronyms

Automated Cartridge System (ACS) The library subsystem consisting of one or two LMUs, and from 1 to 16 attached LSMs.

Cartridge Access Port (CAP) An assembly which allows an operator to enter and eject cartridges during automated operations. The CAP is located on the access door of an LSM.

Host Software Component (HSC) Software running on the Library Control System processor that controls the functions of the ACS.

library An installation of one or more ACSs, attached cartridge drives, volumes placed into the ACSs, host software that controls and manages the ACSs and associated volumes, and the library control data set that describes the state of the ACSs. (*See* TapePlex)

Useful VSM Acronyms

Library Control Unit (LCU) The portion of an LSM that controls the movements of the robot.

Library Management Unit (LMU) A hardware and software product that coordinates the activities of one or more LSMs/LCUs.

Library Storage Module (LSM) The standard LSM (4410) a twelve-sided structure with storage space for up to around 6000 cartridges. It also contains a free-standing, vision-assisted robot that moves the cartridges between their storage cells and attached transports.

Real Tape Drive (RTD) The physical transport attached to the LSM. The transport has a data path to a VTSS and may optionally have a data path to MVS or to another VTSS.

Useful VSM Acronyms

Storage Management Component (SMC) Software interface between IBM's z/OS operating system and Oracle StorageTek real and virtual tape hardware. SMC performs the allocation processing, message handling, and SMS processing for the ELS solution.

TapePlex (formerly "library"), a single Oracle StorageTek hardware configuration, normally represented by a single HSC Control Data Set (CDS). A TapePlex may contain multiple Automated Cartridge Systems (ACSs) and Virtual Tape Storage Subsystems (VTSSs).

Virtual Storage Manager (VSM)— A storage solution that virtualizes volumes and transports in a VTSS buffer in order to improve media and transport use.

Virtual Tape Control System (VTCS)— The primary host code for the Virtual Storage Manager (VSM) solution. This code operates in a separate address space, but communicates closely with HSC.

Useful VSM Acronyms

Virtual Tape Drive (VTD) An emulation of a physical transport in the VTSS that looks like a physical tape transport to MVS. The data written to a VTD is really being written to DASD. The VTSS has 64 VTDs that do virtual mounts of VTVs.

Virtual Tape Storage Subsystem (VTSS)— The DASD buffer containing virtual volumes (VTVs) and virtual drives (VTDs). The VTSS is a StorageTek RAID 6 hardware device with microcode that enables transport emulation. The RAID device can read and write "tape" data from/to disk, and can read and write the data from/to a real tape drive (RTD).

Virtual Tape Volume (VTV) A portion of the DASD buffer that appears to the operating system as a real tape volume. Data is written to and read from the VTV, and the VTV can be migrated to and recalled from real tape.

Connect z/OS and VSM SMF records

VSM Front-End ▸ VTV ▸ Mount Time ▸ Isolate ▸ VTV at time ▸ Identify

VTV Identify

For vtssid 'PZRW48E', for Volume Label '0ADRBE'

VTV Identify Mounts: For Volume Label '0ADRBE'

Time	Volume ...	Data Set Name	System	Jobname	z/OS Progr...	Alloc...	Blocksize...	Total Tape ...	Read+...	Read ...	Write MB (M...	Device...	Device Writ...	Sp...	Non-spec...
7/14/2015 4:44 PM	0ADRBE	HSMPROD.HMIGTAPE.DATASET	D223	DFHSM	BJBDPK	591	16384	593103	16.44	0.00	9717.40	0.00	3847.19	0	1

Volume 0ADRBE (scratch tape) is mounted by DFHSM for Migration .

593103 Blocks (16K) written at 16.44 MB/s

Total written by DFHSM	9717 MB
Total written onto the tape unit	3847 MB

Connect z/OS and VSM SMF records

VTV Identify Mounts: For VTSS ID 'PZRW48E', for VTV volser ID '0ADRBE'

Time	VTSS ID	Ori...	VTSS ID	VT...	Virtual mount type	VTV timestamp	Recall indicator	MVS j...	MVS ...	MVS data set name	Mount start	Mount end	VTV ...
7/14/2015 4:35 PM	PZRW48E	D223	PZRW48E	F1C0	mount sl scratch VTV	7/14/2015 4:35 PM	mounted without a recall	DFHSM	DFHSM	HSMPROD.HMIGTAPE.DATASET	7/14/2015 4:35 PM	7/14/2015 4:35 PM	GBAA

DFHSM mount is issued on VSM PZRW48E at 4:35:26 PM.

Mount is completed at 4:35:26 PM.

No recall is required since this is a scratch mount.

We also see the VTV unit (MVS device) and initiating z/OS host.

Connect z/OS and VSM SMF records

VTV Identify Dismounts: For VTSS ID 'PZRW48E', for VTV volser ID '0ADRBE'

Time	VTSS ID	Ori...	VTSS ID	VTV state	MV...	Uncom...	Th...	The last time the ...	Virtual mount type	MVS j...	MVS ...	M...	VTV ...	Synchronous Replication Status	RUN Received
7/14/2015 4:45 PM	PZRW48E	D223	PZRW48E	VTV dismounted	F1C0	9267.23	3816	7/14/2015 4:35 PM	mount sl scratch VTV	DFHSM	DFHSM		GBAA	Synchronous replication complete	7/14/2015 4:44 PM

DFHSM completes writing to the logical tape at 4:44:14 PM (RUN received)

We see the original and compressed data size (9267 MiB and 3816 MiB)

We also see that synchronous replication was requested (and issued) for this tape.

Connect z/OS and VSM SMF records

VTV Identify Migrate: For VTSS ID 'PZRW48E', for VTV volser ID '0ADRBE'

Time	VTSS ID	HSC S...	VTSS ID	R...	MVC vol...	VTV Siz...	The ...	The last time the VTV ...	Migrate re...	Migrate start	Migrate end	VTV M...	MVC St...	RTD Ch...
7/14/2015 4:45 PM	PZRW48E	O223	PZRW48E	8	155040	3816.4	3816	7/14/2015 4:35 PM	immediate	7/14/2015 4:45 PM	7/14/2015 4:45 PM	GBAA	LACFE	0003

This VTV is now migrated to a physical tape (MVC) labeled 155040. This was an “immediate” migration.

The HSC software on LPAR O223 is managing the migrate.

Physical drive (RTD) # 8 is being used.

Migration start (4:45:25) and end (4:45:51) are shown.

Connect z/OS and VSM SMF records

VTV Identify Replication: For VTV volser ID '0ADRBE' by Primary or Sending VTSS Name

Time	Primary ...	HSC ...	Primary ...	Seconda...	Cluster...	C...	C...	Flag	The am...	VTV last updated t...	Replicate Start Ti...	Replicate End Time	VTV ...	Rewind Received
7/14/2015 4:45 PM	PZRW48E	O223	PZRW48E	PZRW48A	VZRO18	0	0	SYNCHRONOUS REPLICATE	3816.39	7/14/2015 4:35 PM	7/14/2015 4:44 PM	7/14/2015 4:45 PM	GBAA	7/14/2015 4:44 PM

This VTV is replicated synchronously to VSM PZRW48A successfully.

Replication start (4:44:16 PM) and end (4:45:23 PM) are shown as is the amount of data (3816 MiB.)

Note: VTV Recall and VTV deleted are not applicable to this mount.