DSS

Data & Storage Services

CERN IT Department

TSM Monitoring @ CERN

Daniele Francesco Kruse
CERN IT/DSS

Presented by Giuseppe Lo Presti

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

20th HEPiX - Vancouver - October 2011

CERN

Outline

- TSM at CERN
- TSM Management Station
  - Overview
  - Main features
- TSMMSv2
  - Motivations
  - Design
  - New ideas

TSM at CERN (1/3)

- We back up:
  1. Network filesystems (60'000 AFS, 1'500 DFS volumes)
  2. Email (18'000 mailboxes)
  3. Web sites (12'000 websites)
  4. Databases (120 DB servers)
  5. Servers (1'000 Linux and Windows servers)
  6. Virtual Machines (120 hypervisors)

- We don't back up:
  1. Physics data (using CASTOR for this)
  2. User PCs (already backing up home AFS/DFS directories)

TSM at CERN (2/3)

- We currently have around 3.8 PB of backup data and 0.6 PB of archived data
  - … and growing superlinearly (last year 1 PB)
- Average daily traffic is 50 TB also growing steadily
- Around 1,200 nodes are backed up, for a total 1,500 million files

TSM at CERN (3/3)

17 TSM Servers in production on RHEL4/5

80 TB of disk storage

- 2 IBM TS3500 libraries
- 48 IBM drives
- 4'500 IBM 3952 cartridges

TSM Management Station

TSM monitoring tool developed in-house

- Gathers data from the TSM servers
- Generates graphs and reports with various statistics
- Sends e-mails to users and administrators to inform them about potential issues
- Very useful to manage the increasing number of TSM servers

TSM Management Station

- Home
- Nodes
  - Volume
  - Occupancy
  - Account
- Servers
  - Drives
  - Tapes
  - Storage Pools
  - Console
  - Schedule
  - Performance
- Operational
  - Daily Client
  - Daily Server
  - Weekly Server
- Trends
  - Volume
  - Server Volume
  - Node
  - Occupancy
  - Server Occupancy
  - Audit Occupancy
  - Server Database
  - Drives
  - Disk Space
  - Tapes
  - Storage Pools
  - Account Volume
  - Account Occupancy
- Reports
  - Top Client Volume
  - Top Client Volume Growth
  - Top Client Occupancy
  - Top Client Occupancy Growth
  - Top Media Cost by Department
  - License use by account
  - Slowest Client
  - Oldest Inactive Client
  - Account

TSM Management Station

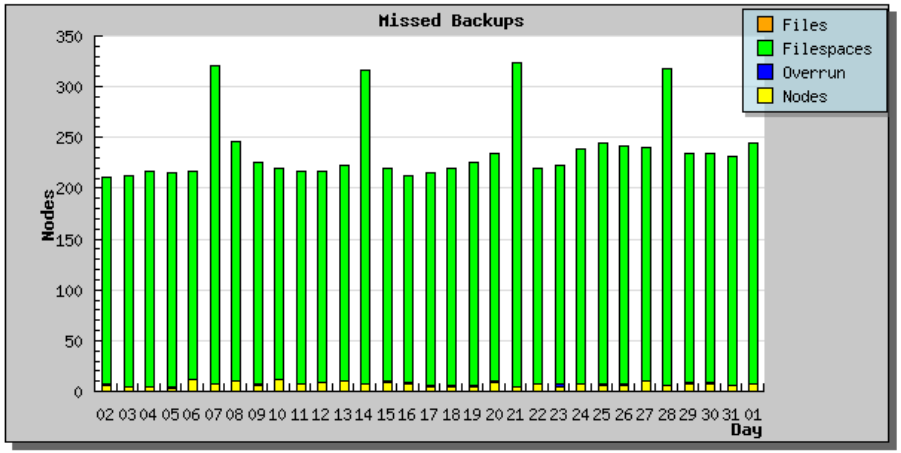

Backup Status	
Total Nodes	1239
Missed Nodes	7
Overrun Nodes	0
Missed Filespaces	237
Missed Files	

	613									513								
Servers	*TSM41*	TSM609	TSM61	TSM62	TSM63	TSM64	TSM65	TSM66	*TSM66-OLD*	TSM51	*TSM511*	TSM52	TSM53	TSM54	TSM55	*TSM58*	TSM59	TSM91
DB BACKUP	0.3 days		0.4 days	0.4 days	0.4 days	0.3 days	0.3 days	0.3 days	0.3 days	0.4 days		0.4 days	0.3 days	0.4 days	0.3 days		0.3 days	0.4 days
DB RECLOG	0%		0%	0%	0%	0%	25%	0%	0%	0%		0%	0%	0%	0%		0%	0%
DB SPACE	9%		55%	94%	89%	88%	94%	68%	46%	76%		66%	74%	92%	82%		18%	79%
POOL ARCHIVEPOOL	0%		6%	33%	16%	0%	59%	45%	13%	37%	0%	52%	41%	0%	71%	0%	54%	39%
POOL DIRMCPOOL	0%	0%	20%	13%	21%	19%	71%	4%	2%	4%	0%	0%	34%	18%	38%	22%	0%	25%
POOL ARCHIVEFILE	0%					7%										0%		
POOL BACKUPFILE	30%					59%		0%										
POOL BACKUPPOOL	0%	48%	39%	40%	61%	36%	96%	68%	5%	39%	0%	39%	40%	37%	63%	70%	0%	42%
POOL BENCHPOOL	0%																	
POOL TESTPOOL												0%				100%		
POOL AFSARCHPOOL														31%				
POOL BACKUPDISK																2%		
POOL ORAPOOL							0%											
CART ARCHIVECART	0%		67%	53%	54%	48%	59%	0%	73%	35%	0%	72%	70%	0%	77%	0%	60%	58%
CART COPYCART		66%										26%		49%		15%		
CART BACKUPCART	1%	75%	12%	46%	14%	47%	18%	21%	62%	41%	0%	34%	45%	41%	57%	38%	0%	64%
CART BENCHCART	0%																	
CART SPACEMGCART	0%																	
CART AFSARCHCART														78%				
CART TESTCART																2%		
CART BACKUPCART_OLD								0%										
Library					TSM613									TSM513				
Drives	IBM0101	IBM0102	IBM0103	TSM65	TSM65	TSM52	IBM0107	TSM52	TSM609	IBM2101	TSM55	TSM55	TSM54	TSM55	TSM52	IBM2107	IBM2108	TSM55
	TSM609	IBM0403	IBM0404	IBM0405	IBM0406	TSM65	TSM65			IBM2110	TSM54	IBM2712	IBM2911	IBM2912				

TSM Management Station

TSMMS daily report example:

Msg	Target	Description
TSM101W	COPYCART@TSM54	Stage pool COPYCART free space 1365GB is below minimum 4444GB
TSM101W	COPYCART@TSM609	Stage pool COPYCART free space 270GB is below minimum 4444GB
TSM101W	COPYCART@TSM52	Stage pool COPYCART free space 1245GB is below minimum 4444GB
TSM115E	TSM613	Path IBM0405 is not online
TSM116W	TSM613	Drive IBM0405 is not online
TSM105W	TSM65	Database utilisation 94% exceeds 90%
TSM105W	TSM62	Database utilisation 94% exceeds 90%
TSM105W	TSM54	Database utilisation 92% exceeds 90%
TSM110W	DBSRVD227@TSM62	Node is not registered to backup but tried 288 times
TSM110W	LXFSRD0305@TSM63	Node is not registered to backup but tried 288 times
TSM110W	LXLIC06@TSM51	Node is not registered to backup but tried 288 times
TSM110W	TRTDAQ1@TSM64	Node is not registered to backup but tried 286 times
TSM110W	LXLIC07@TSM51	Node is not registered to backup but tried 288 times
TSM110W	DBSRVD101@TSM51	Node is not registered to backup but tried 288 times
TSM110W	DBSRVD228@TSM62	Node is not registered to backup but tried 286 times
TSM103W	AFSARCHCART@TSM54	Pool AFSARCHCART has insufficient empty volumes (1) and only 7 filling
TSM017W	CERNDATA58_MONTHLY@TSM61	Growth of backup files on tape over a month (was 114038 files, now 4510713 files, 3955%) [ask]
TSM016W	CAEVMSRV51@TSM66	Growth of backup data on tape over a month (was 2.1TB, now 6.6TB, 313%) [ask]
TSM016W	CAEVMSRV20@TSM65	Growth of backup data on tape over a month (was 1.2TB, now 3.9TB, 315%) [ask]

TSMMS also sends an email for each error in each TSM server

TSM Management Station

- Allows management of groups of nodes (by department and division) and generates graphs and stats for each group
- Sends alerts to nodes whenever an operation fails or whenever they miss their periodic backup
- Features options to suspend or stop the alerting system
- Gives information of each node about file spaces, backup history performance and stats, associated schedules, etc.
- … and many other stats and graphs

Motivations for a new TSMMS

- TSMMS provides 90% of all the information that is needed
- However:
  - not use-case oriented
  - not compatible with TSM v6.x
    (heavily depending on the TSM 5 database schema)
- The choice was then to start from scratch with a clean design and architecture
- ***Change in philosophy:*** **the focus is now on how to convey the relevant information for each use-case**

Splunk

- TSMMS takes care of the monitoring and the alerting system
- **TSMMSv2** will be only responsible for the monitoring while the alerting tasks will be moved to **Splunk**
- Splunk is a commercially available tool (with a free trial):
  - Log aggregator/mining
  - Search engine
  - New features: alerting and reporting
- TSMMSv2 and Splunk will work together to provide the TSM admin with proper information and alerts

Splunk

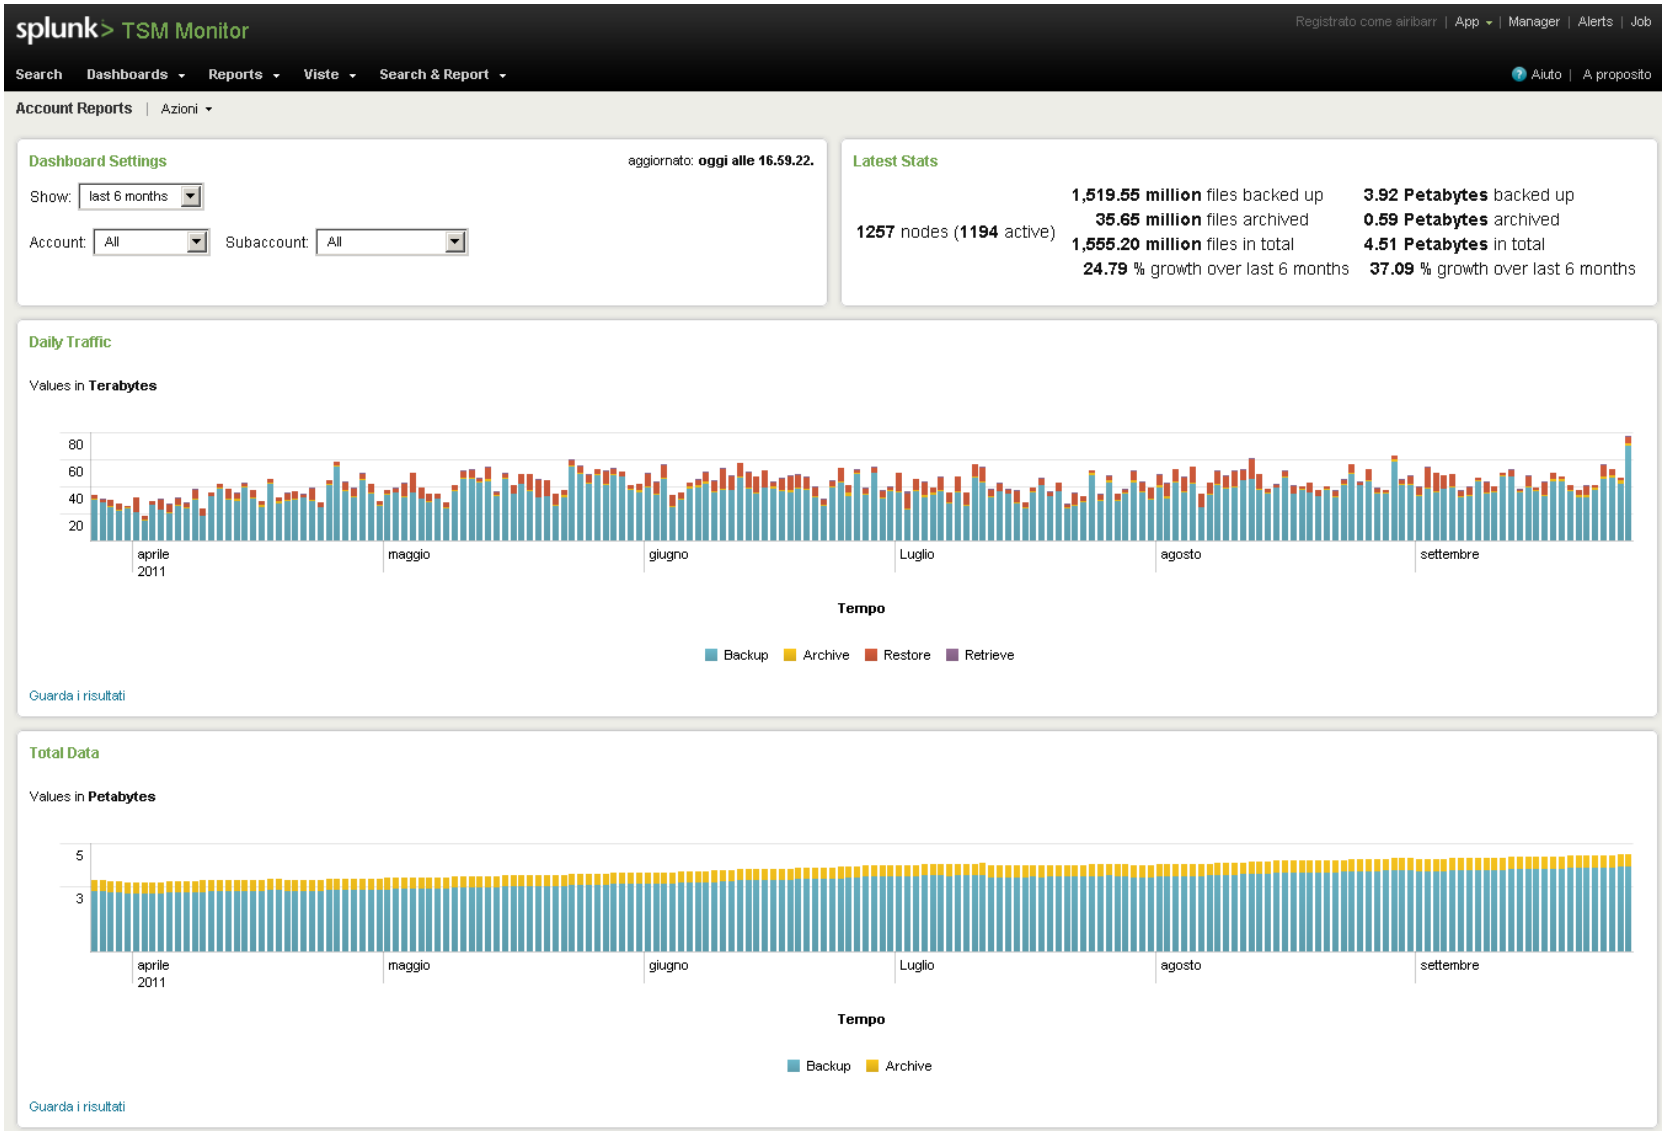
splunk> TSM Monitor
Registrato come airibarr | App | Manager | Alerts | Job
Search Dashboards Reports Viste Search & Report
Aiuto | A proposito
Account Reports | Azioni
Dashboard Settings
aggiornato: oggi alle 16.59.22.
Show: last 6 months
Account: All
Subaccount: All
Latest Stats
1257 nodes (1194 active)
1,519.55 million files backed up
35.65 million files archived
1,555.20 million files in total
24.79 % growth over last 6 months
3.92 Petabytes backed up
0.59 Petabytes archived
4.51 Petabytes in total
37.09 % growth over last 6 months
Daily Traffic
Values in Terabytes
80
60
40
20
aprile 2011
maggio
giugno
Luglio
agosto
settembre
Tempo
Backup Archive Restore Retrieve
Guarda i risultati
Total Data
Values in Petabytes
5
3
aprile 2011
maggio
giugno
Luglio
agosto
settembre
Tempo
Backup Archive
Guarda i risultati

TSMMSv2 modeled on a typical TSM admin day

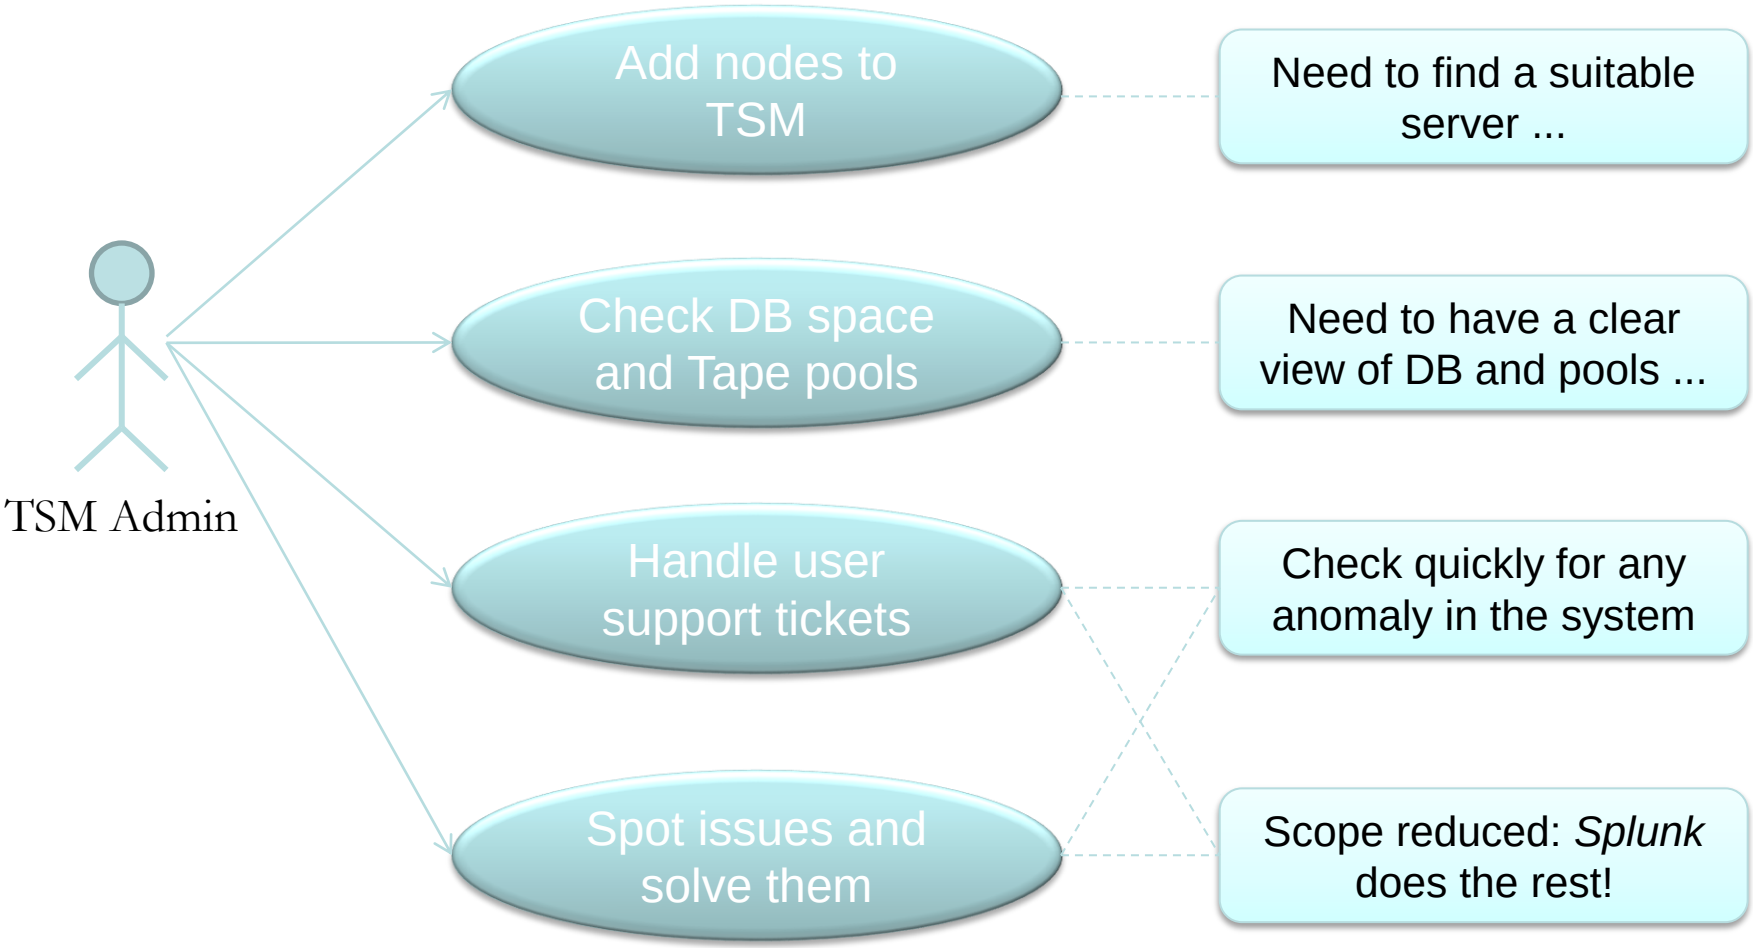

Structure of TSMMSv2

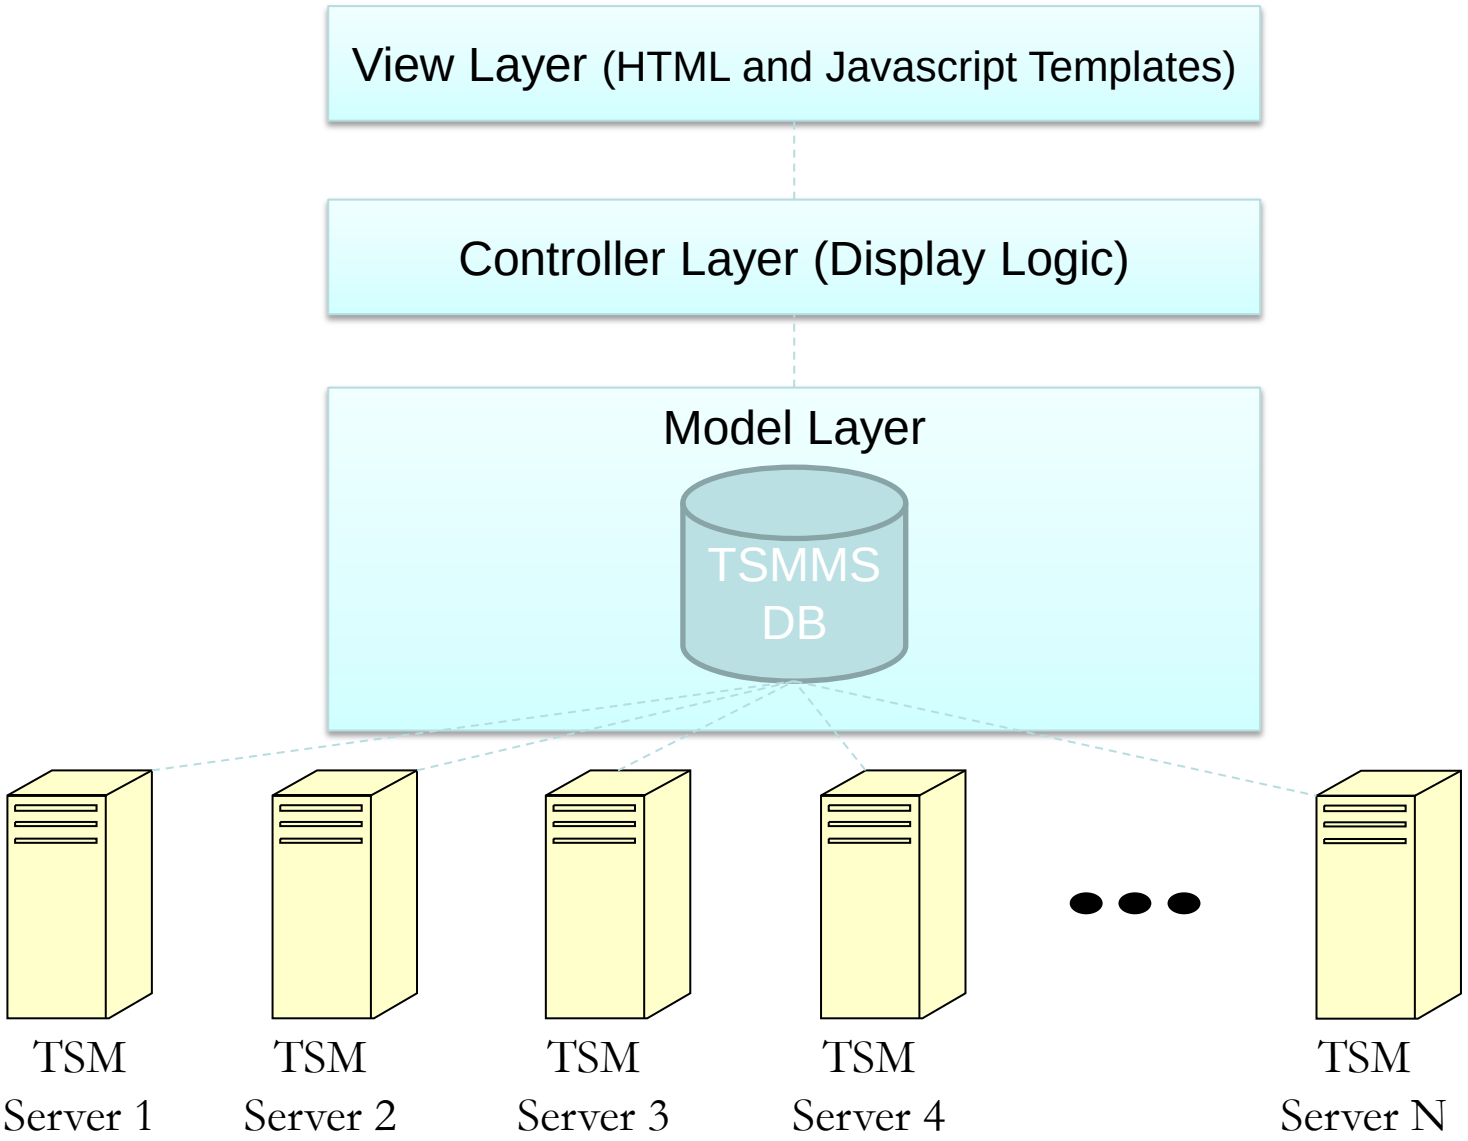

TSMMSv2 New Ideas

- TSMMSv2 will focus on helping TSM admins with daily tasks
- Display only relevant information (not everything else) for the most important issues that may arise
- Not only monitoring → also GUI for selected common administrative tasks
  - Add new nodes to approriate server
- Automation of certain tasks, such as:
  - Add new storage space where needed (ex. DB)
  - Automatically deal with faulty tapes or drives

Thank you, Questions ?