Application Performance Mgmt in ambienti fortemente virtualizzati

Angela Molinari

Agenda

- Introduction to Tivoli Virtual Server Monitoring
- ITM for Virtual Servers Key Capabilities
- New Product: IBM Tivoli for Virtual Environments
- Scenarios

Monitoring and management requires visibility at different levels in the Cloud and the ability to cope with change

With cloud computing:

- A mix of physical and virtual servers, storage/network devices and applications
- Workloads appears, change and move at the click of a mouse
- Capacity planning for highly virtualized environments is more complex due to
  - High infrastructure sharing and advanced reservation capabilities
  - Resources dynamic change and cloudburst to public clouds
  - …

Cloud tenants want to understand and optimize performance of the services being delivered

Cloud administrator needs to support changes in demand and understand workload trends

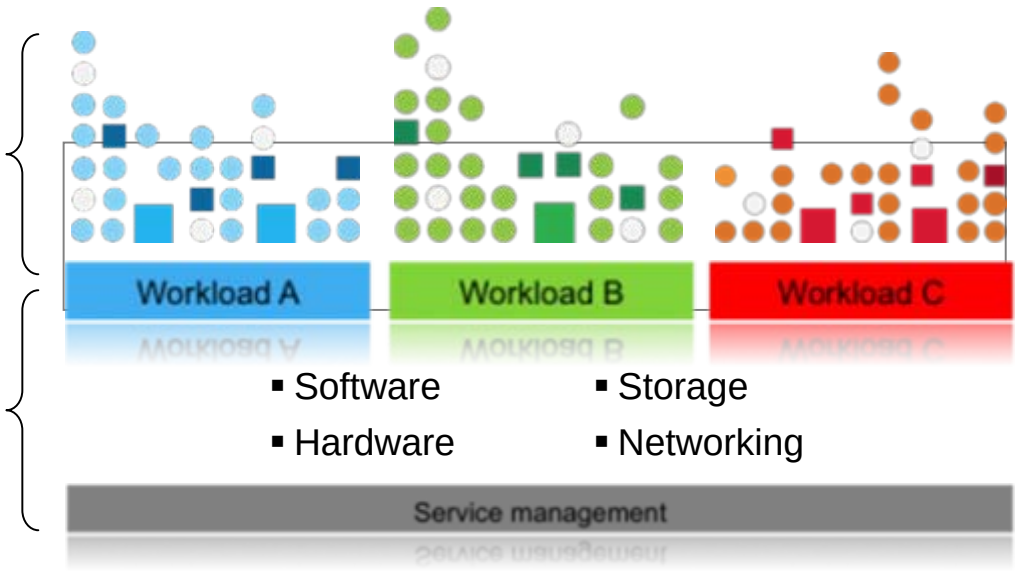

Tivoli Virtual Server Monitoring solutions available today…

- **ITM**
  - IBM PowerVM (CEC, VIOS, HMC, LPARs)
- **ITM for Virtual Servers**
  - Citrix (MetaFrame, XenServer)
  - Linux (KVM, RHEV-H)
  - NetApp (DataFabric Manager)
  - VMware (ESX, ESXi, vCenter)
- **ITCAM for Microsoft Applications**
  - HyperV, Citrix, Vmware
- **ITCAM for Applications bundle**
  - Includes ITM for Virtual Servers agents

Monitoring PowerVM

- **ITM AIX Agent – "Premium Agent"**
  - Monitors the availability, health and performance of key AIX system resources: LPAR configurations, CPU, memory, storage, network, printers, NIM and Workload Partitions
- **ITM VIOS Agent**
  - Monitors the availability and health of the VIOS resources:
    - CPU, memory, storage, and networks
    - View storage and network mappings between the VIOS Server and its clients
    - *pre-installed* on a VIOS system so only agent configuration is required.
- **ITM CEC Agent**
  - Monitor and view metrics on:
    - Number of LPARs per CEC
    - CPU, and memory allocations per LPAR
    - LPAR state, LPAR utilization, operating environment, CEC modes and CEC utilization.
- **ITM HMC Agent for System p**
  - Monitors the availability and health of the HMC resources:
    - CPU, memory, storage, and network
    - Runs on any available LPAR

Out-of-box alerts and expert advice

AIX LPAR

KPX_memrepage_Info
KPX_vmm_pginwait_Info
KPX_vmm_pgfault_Info
KPX_vmm_pgreclm_Info
KPX_vmm_unpin_low_Warn KPX_vmm_pgout_pend_Info
KPX_Pkts_Sent_Errors_Info KPX_Sent_Pkts_Dropped_Info
KPX_Pkts_Recv_Errors_Info KPX_Bad_Pkts_Recvd_Info
KPX_Recv_pkts_dropped_Info KPX_Qoverflow_Info
KPX_perip_InputErrs_Info KPX_perip_InputPkts_Drop_Info
KPX_perip_OutputErrs_Info KPX_TCP_ConnInit_Info
KPX_TCP_ConnEst_Info
KPX_totproc_cs_Info KPX_totproc_runq_avg_Info
KPX_totproc_load_avg_Info KPX_totnum_procs_Info
KPX_perproc_IO_pgf_Info KPX_perproc_nonIO_pgf_Info
KPX_perproc_memres_datasz_Info
KPX_perproc_memres_textsz_Info KPX_perproc_mem_textsz_Info
KPX_perproc_vol_cs_Info
KPX_Active_Disk_Pct_Info KPX_Avg_Read_Transfer_MS_Info
KPX_Read_Timeouts_Per_Sec_Info
KPX_Failed_Read_Per_Sec_Info
KPX_Avg_Write_Transfer_MS_Info
KPX_Write_Timeout_Per_Sec_Info
KPX_Failed_Writes_Per_Sec_Info
KPX_Avg_Req_In_WaitQ_MS_Info
KPX_ServiceQ_Full_Per_Sec_Info
KPX_perCPU_syscalls_Info KPX_perCPU_forks_Info
KPX_perCPU_execs_Info
KPX_perCPU_cs_Info
KPX_Tot_syscalls_Info
KPX_Tot_forks_Info
KPX_Tot_execs_Info
KPX_LPARBusy_pct_Warn KPX_LPARPhyBusy_pct_Warn
KPX_LPARvcs_Info
KPX_LPARfreepool_Warn KPX_LPARPhanIntrs_Info
KPX_LPARentused_Info KPX_LPARphyp_used_Info
KPX_user_acct_locked_Info KPX_user_login_retries_Info
KPX_user_idletime_Info

HMC

KPH_Busy_CPU_Info
KPH_Paging_Space_Full_Info
KPH_Disk_Full_Warn
KPH_Runaway_Process_InfoThe

VIOS

KVA_memrepage_Info
KVA_vmm_pginwait_Info
KVA_vmm_pgfault_Info
KVA_vmm_pgreclm_Info
KVA_vmm_unpin_low_Warn KVA_vmm_pgout_pend_Infov Networking
KVA_Pkts_Sent_Errors_Info KVA_Sent_Pkts_Dropped_Info
KVA_Pkts_Recv_Errors_Info KVA_Bad_Pkts_Recvd_Info
KVA_Recv_pkts_dropped_Info
KVA_Qoverflow_Info
KVA_Real_Pkts_Dropped_Info KVA_Virtual_Pkts_Dropped_Info
KVA_Output_Pkts_Dropped_Info KVA_Output_Pkts_Failures_Info
KVA_Mem_Alloc_Failures_Warn
KVA_ThreadQ_Overflow_Pkts_Info KVA_HA_State_Info
KVA_Times_Primary_Per_Sec_Info KVA_perip_InputErrs_Info
KVA_perip_InputPkts_Drop_Info KVA_perip_OutputErrs_Info
KVA_TCP_ConnInit_Info
KVA_TCP_ConnEst_Infov Process
KVA_totproc_cs_Info
KVA_totproc_runq_avg_Info KVA_totproc_load_avg_Info
KVA_totnum_procs_Info
KVA_perproc_IO_pgf_Info KVA_perproc_nonIO_pgf_Info
KVA_perproc_memres_datasz_Info
KVA_perproc_memres_textsz_Info KVA_perproc_mem_textsz_Info
KVA_perproc_vol_cs_Info
KVA_Firewall_Info
KVA_memrepage_Info
KVA_vmm_pginwait_Info
KVA_vmm_pgfault_Info
KVA_vmm_pgreclm_Info
KVA_vmm_unpin_low_Warn KVA_vmm_pgout_pend_Infov Networking
KVA_Pkts_Sent_Errors_Info KVA_Sent_Pkts_Dropped_Info
KVA_Pkts_Recv_Errors_Info KVA_Bad_Pkts_Recvd_Info
KVA_Recv_pkts_dropped_Info
KVA_Qoverflow_Info
KVA_Real_Pkts_Dropped_Info
KVA_Virtual_Pkts_Dropped_Info
KVA_Output_Pkts_Dropped_Info KVA_Output_Pkts_Failures_Info
KVA_Mem_Alloc_Failures_Warn
KVA_ThreadQ_Overflow_Pkts_Info KVA_HA_State_Info
KVA_Times_Primary_Per_Sec_Info KVA_perip_InputErrs_Info
KVA_perip_InputPkts_Drop_Info KVA_perip_OutputErrs_Info
KVA_TCP_ConnInit_Info
KVA_TCP_ConnEst_Infov Process
KVA_totproc_cs_Info
KVA_totproc_runq_avg_Info KVA_totproc_load_avg_Info
KVA_totnum_procs_Info
KVA_perproc_IO_pgf_Info KVA_perproc_nonIO_pgf_Info
KVA_perproc_memres_datasz_Info
KVA_perproc_memres_textsz_Info
KVA_perproc_mem_textsz_Info KVA_perproc_vol_cs_Info
KVA_Firewall_Info
KVA_Active_Disk_Pct_Info KVA_Avg_Read_Transfer_MS_Info
KVA_Read_Timeouts_Per_Sec_Info
KVA_Failed_Read_Per_Sec_Info
KVA_Avg_Write_Transfer_MS_Info
KVA_Write_Timeout_Per_Sec_Info
KVA_Failed_Writes_Per_Sec_Info
KVA_Avg_Req_In_WaitQ_MS_Info
KVA_ServiceQ_Full_Per_Sec_Info
KVA_perCPU_syscalls_Info
KVA_perCPU_forks_Info
KVA_perCPU_execs_Info
KVA_perCPU_cs_Info
KVA_Tot_syscalls_Info KVA_Tot_forks_Info
KVA_Tot_execs_Info
KVA_LPARBusy_pct_Warn KVA_LPARPhyBusy_pct_Warn
KVA_LPARvcs_Info
KVA_LPARfreepool_Warn
KVA_LPARPhanIntrs_Info
KVA_LPARentused_Info
KVA_LPARphyp_used_Info KVA_user_acct_locked_Info
KVA_user_login_retries_Info
KVA_user_idletime_Info

Major scenarios in Managing Cloud Environments

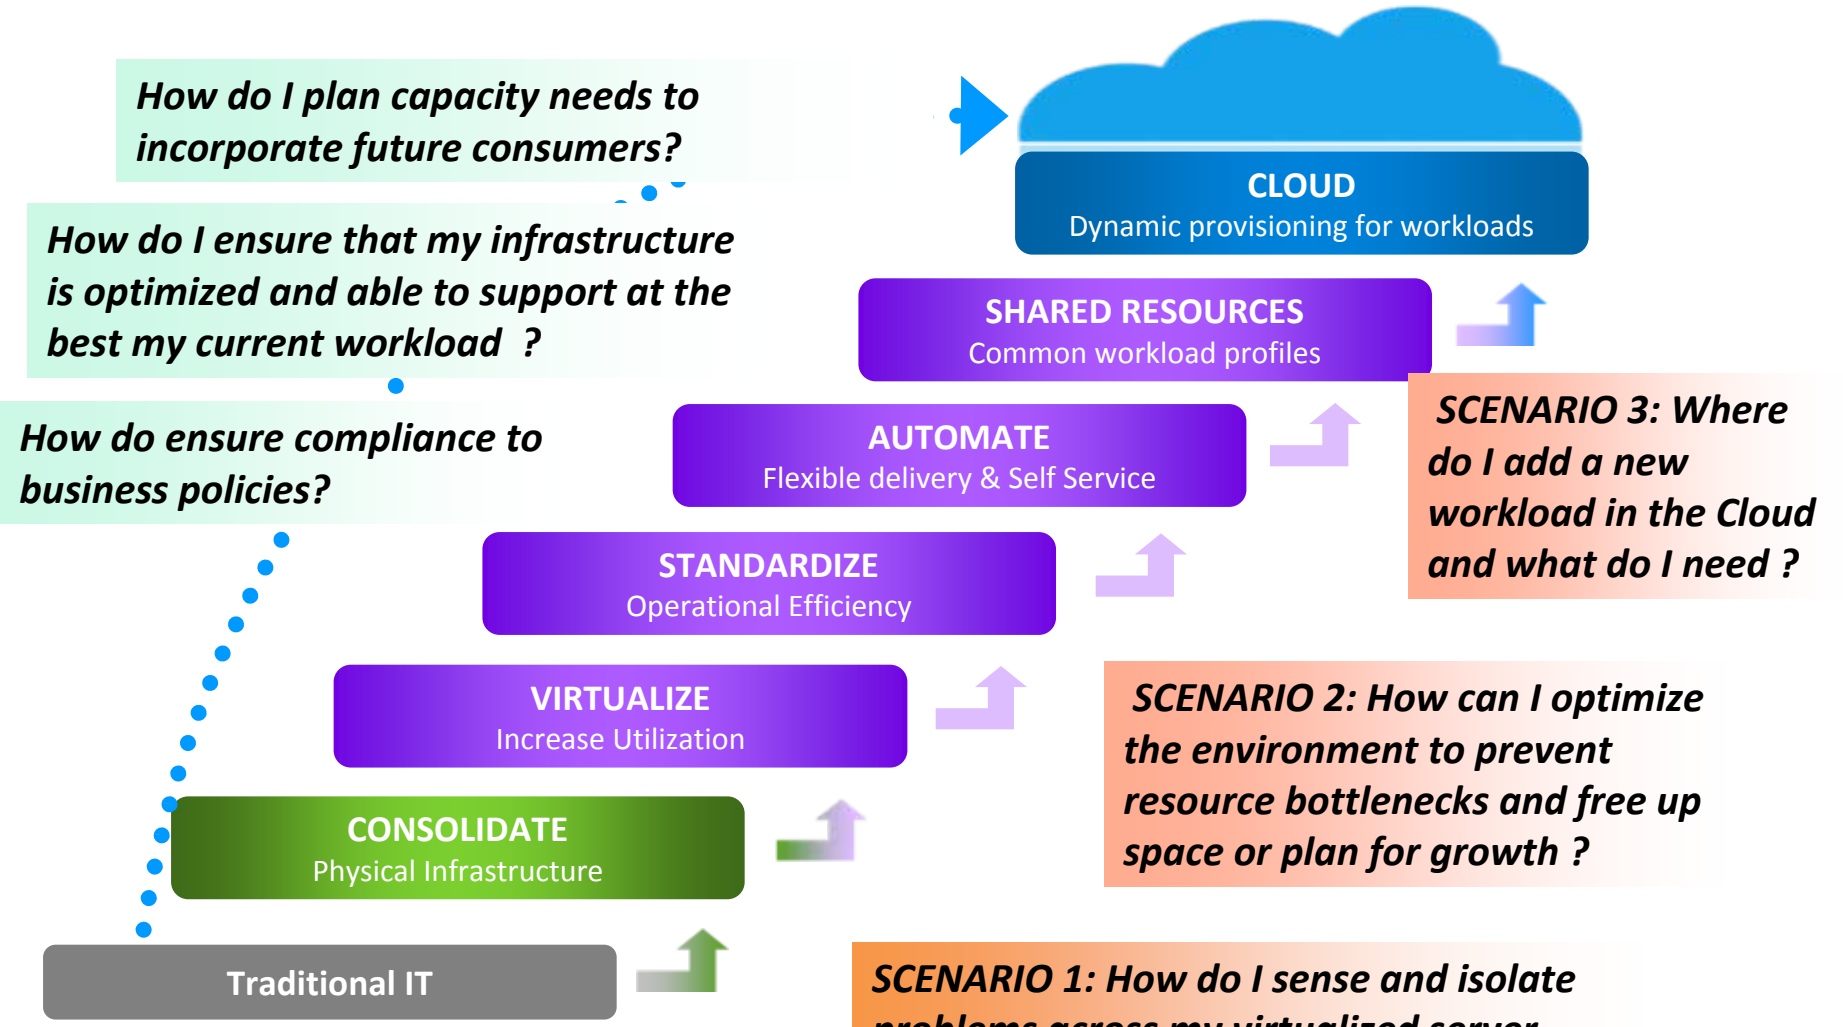

ITM for Virtual Environments

- TIP based dashboards with holistic view of health of whole environment
- Out of the box contextual views of health (availability, performance and capacity) in the complete context of the virtual environment to include physical and virtual server, storage and network resources.
- Integrates across our tool set to merge physical & virtual data – TPC, ITNM, ITM, TADDM & ITMfVS
- Views with performance and capacity reports for assessment of environment and long term trend analysis.
- Additional platform support including XenApp and XenDesktop

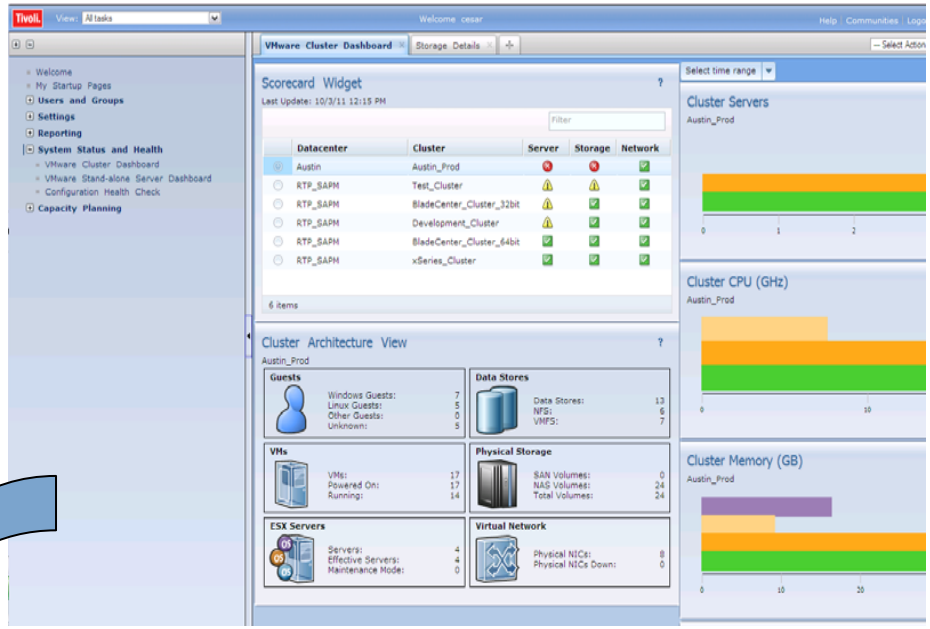

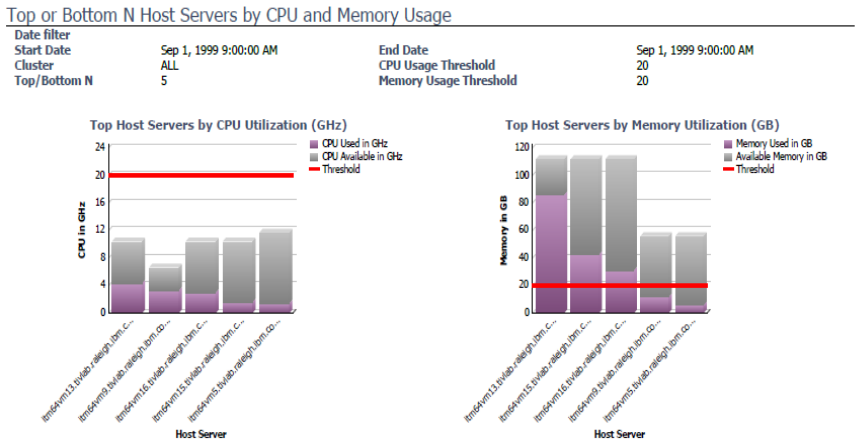

ITM for VE Health Summary Dashboard: VMware Sample

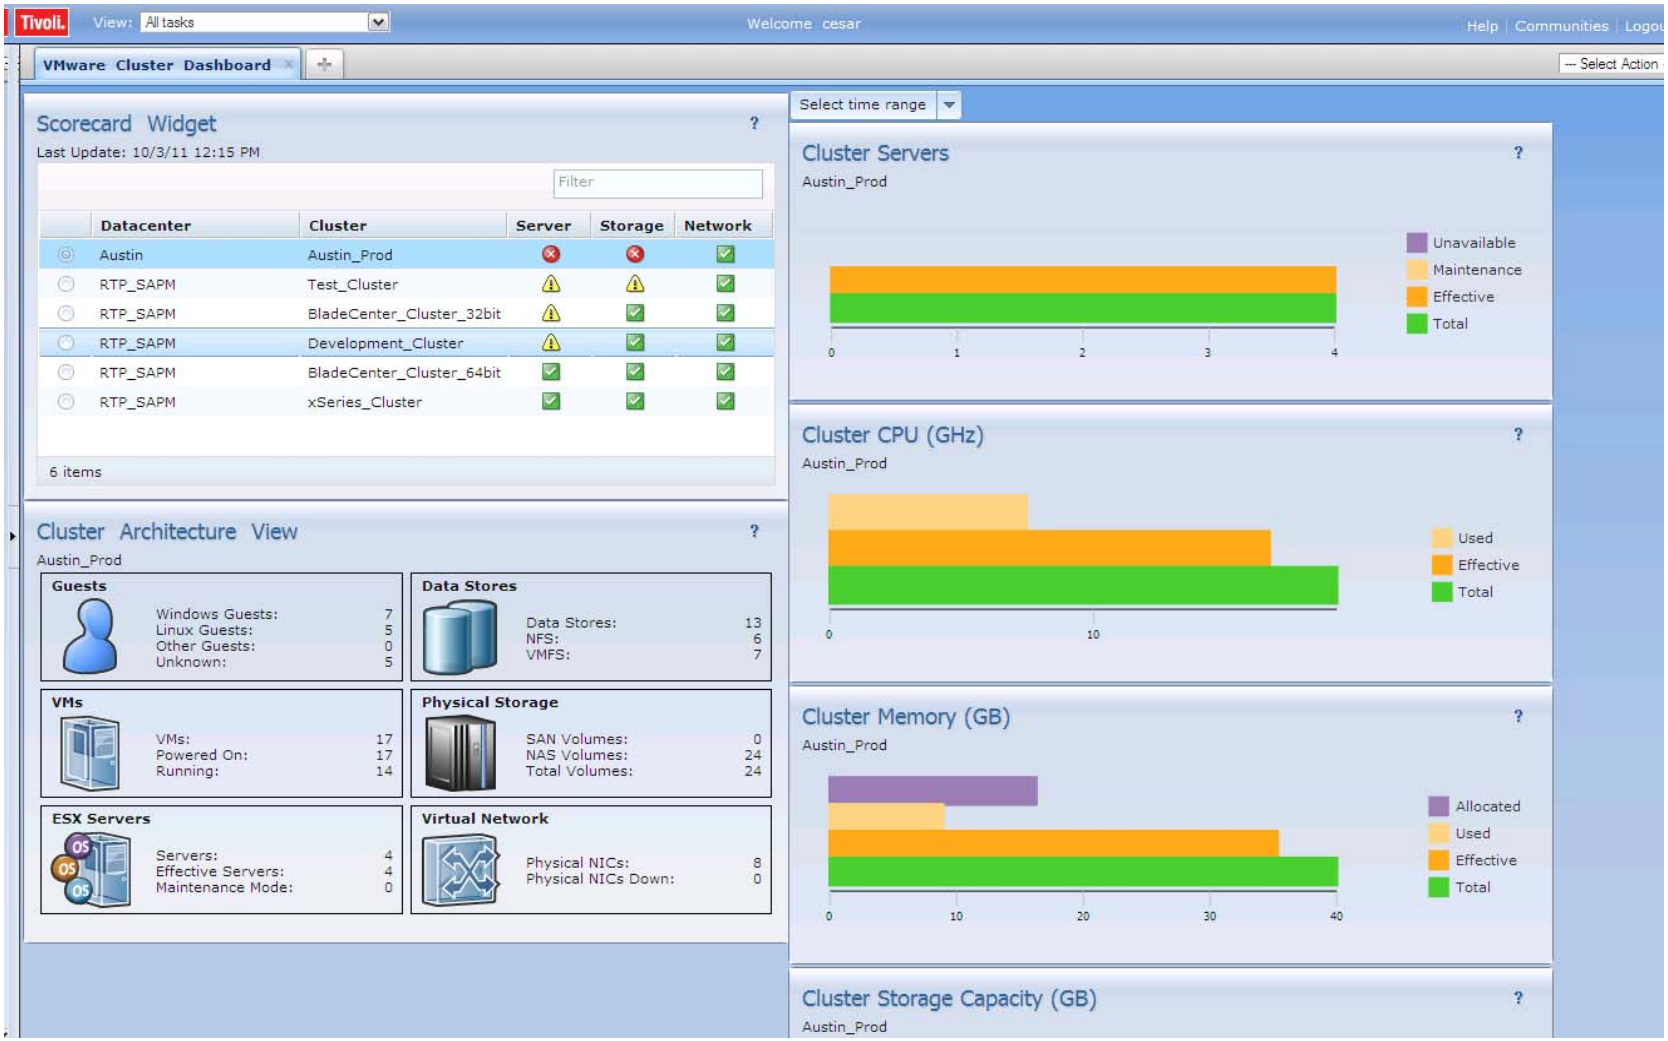

ITM for VE Problem Diagnose Page: VMware Storage Sample

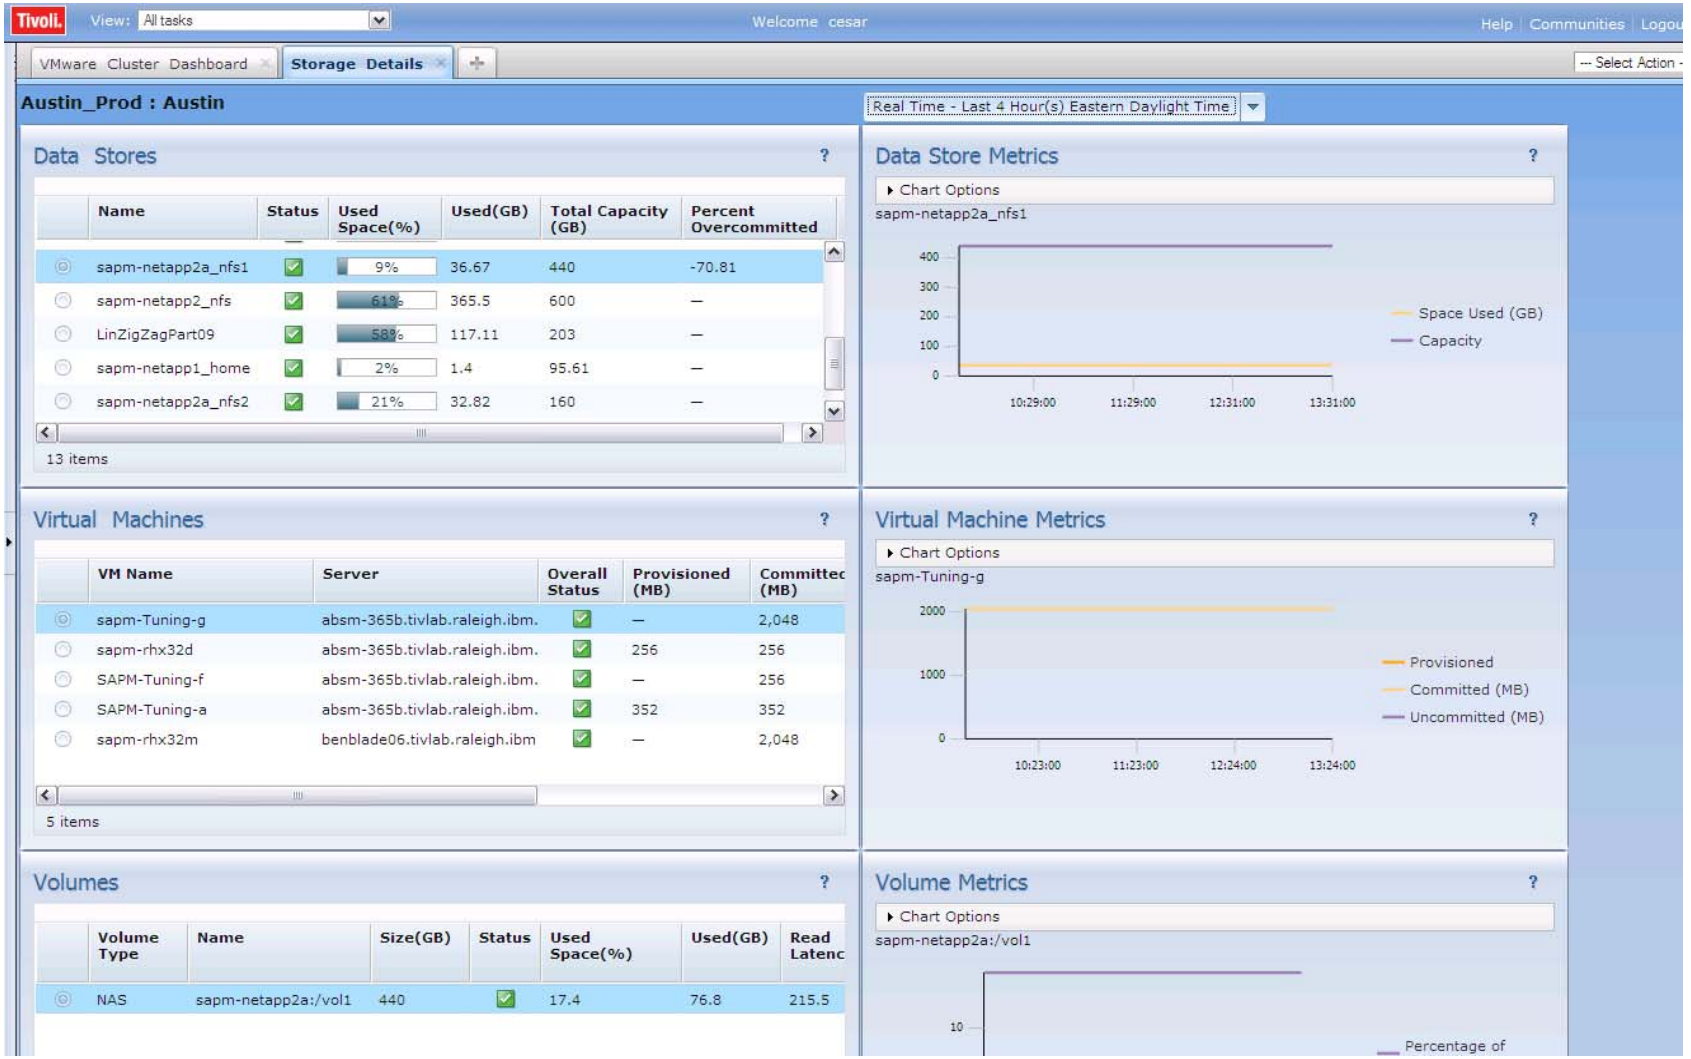

ITMfVE Launch in Context for Drill-down: Storage Sample

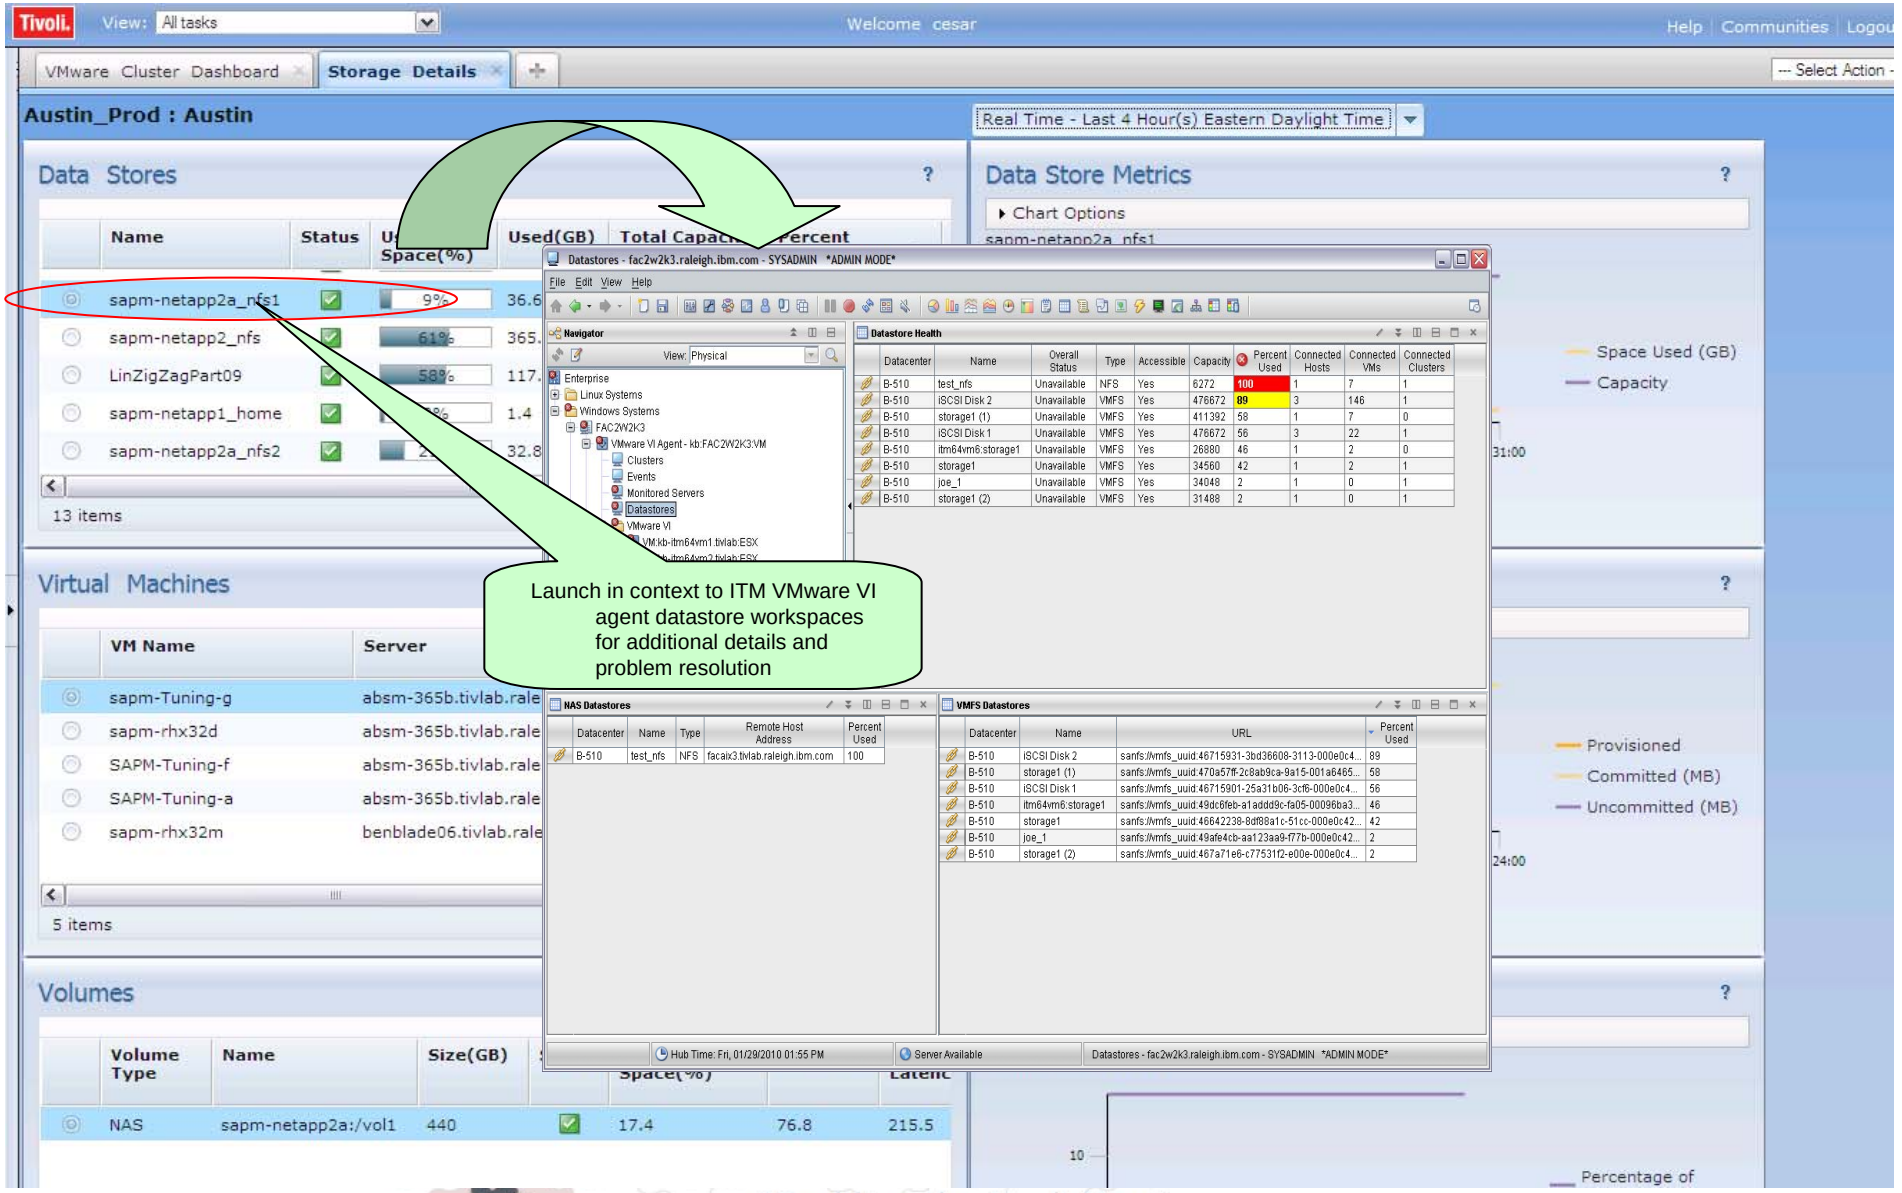

Managing Cloud Environments – Scenario 1

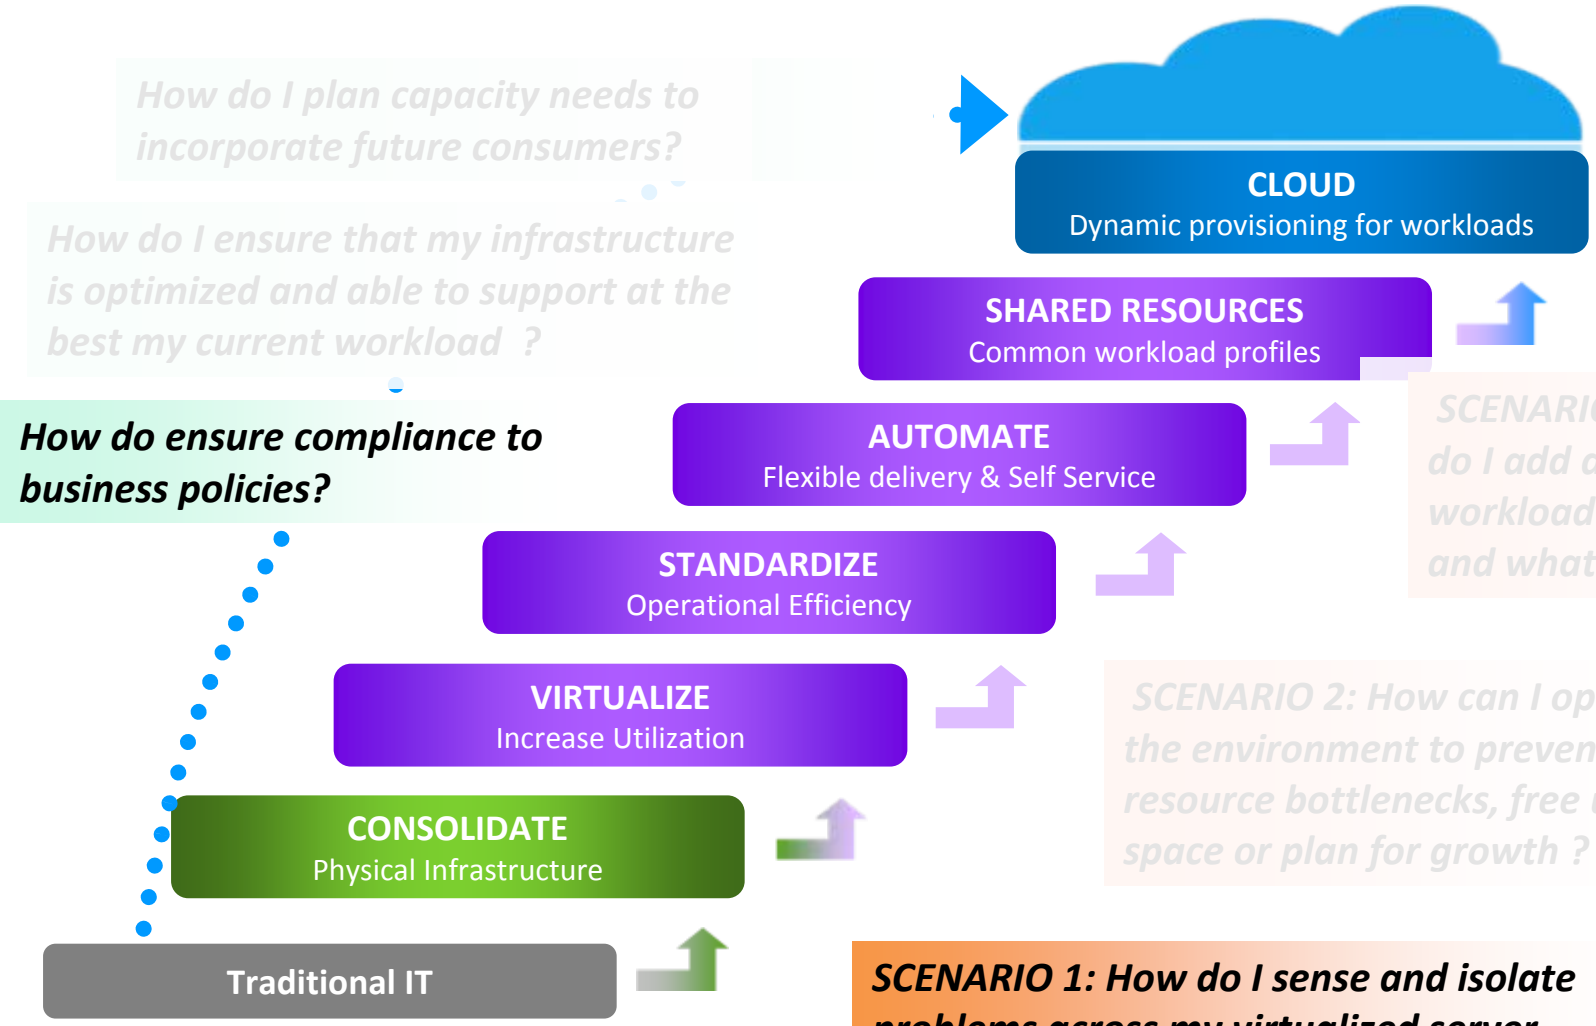

Sample Scenario

An IT Admin resolves an incident related to storage capacity

Cluster Health Scorecard showing Critical Storage Problem

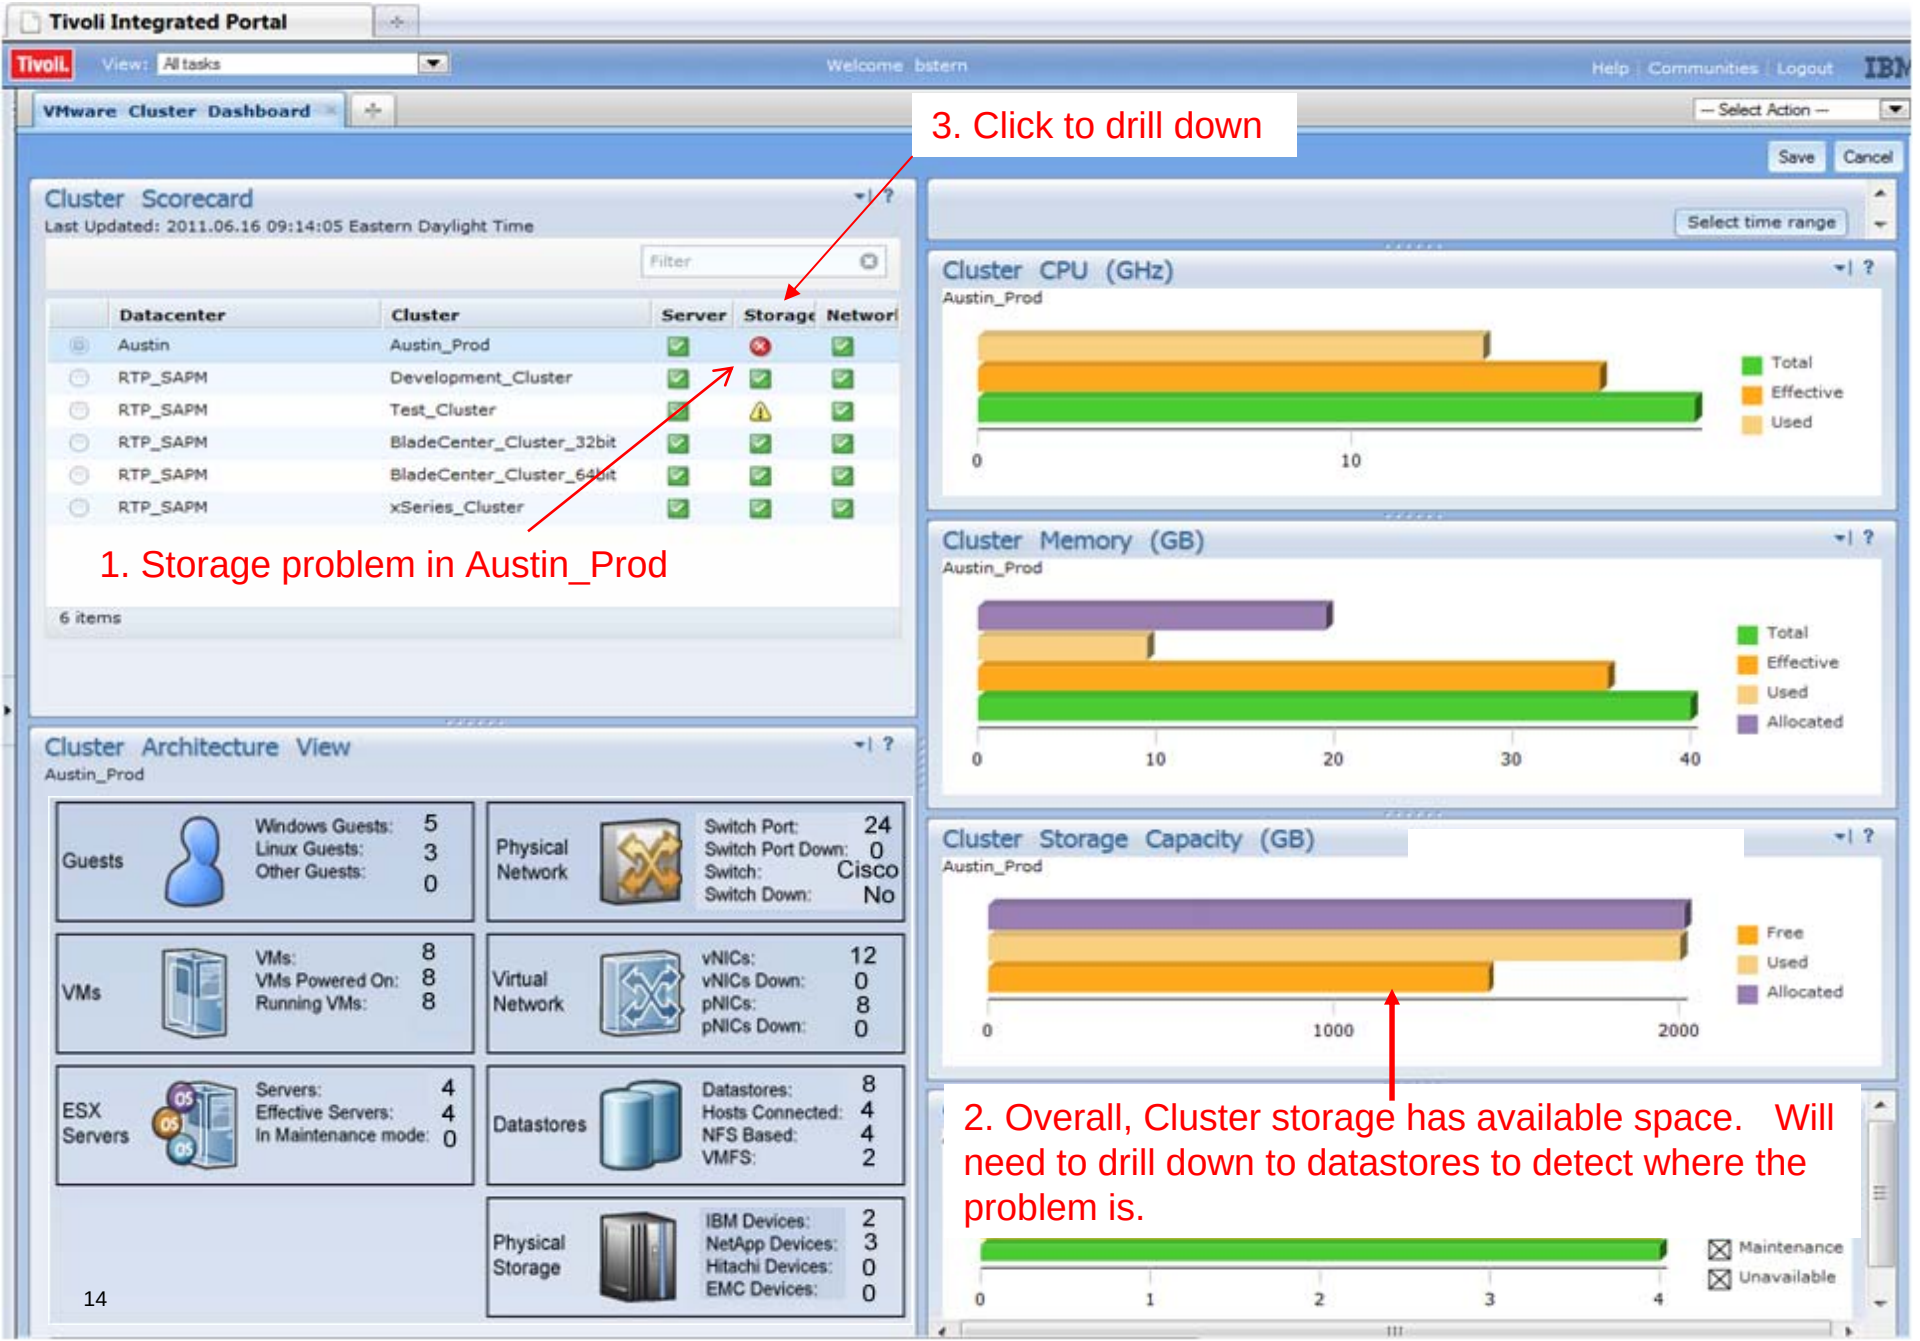

Diagnose Problem with Storage

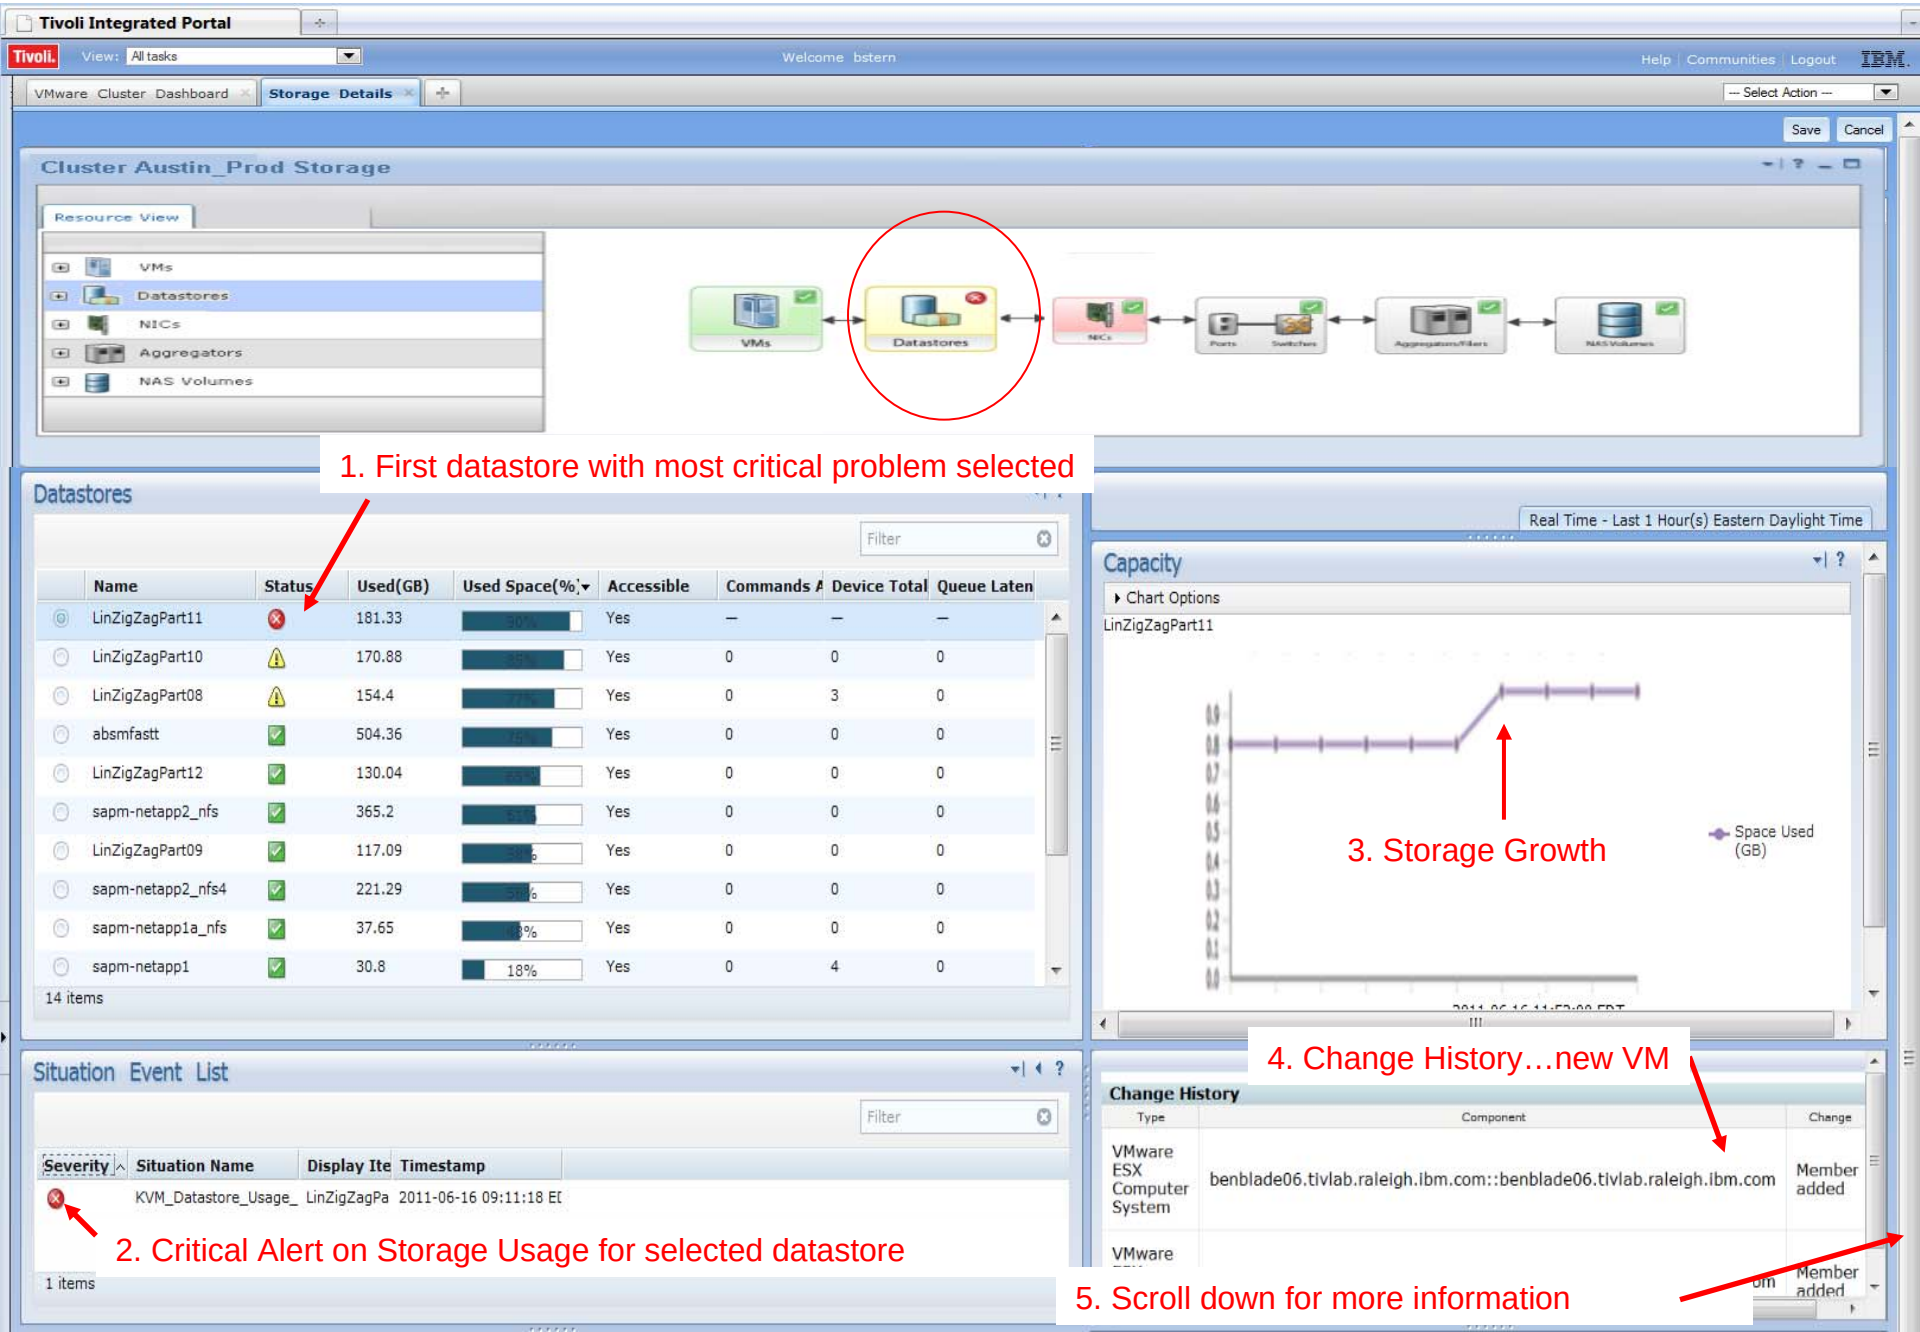

Find a Good Target Datastore for VMs

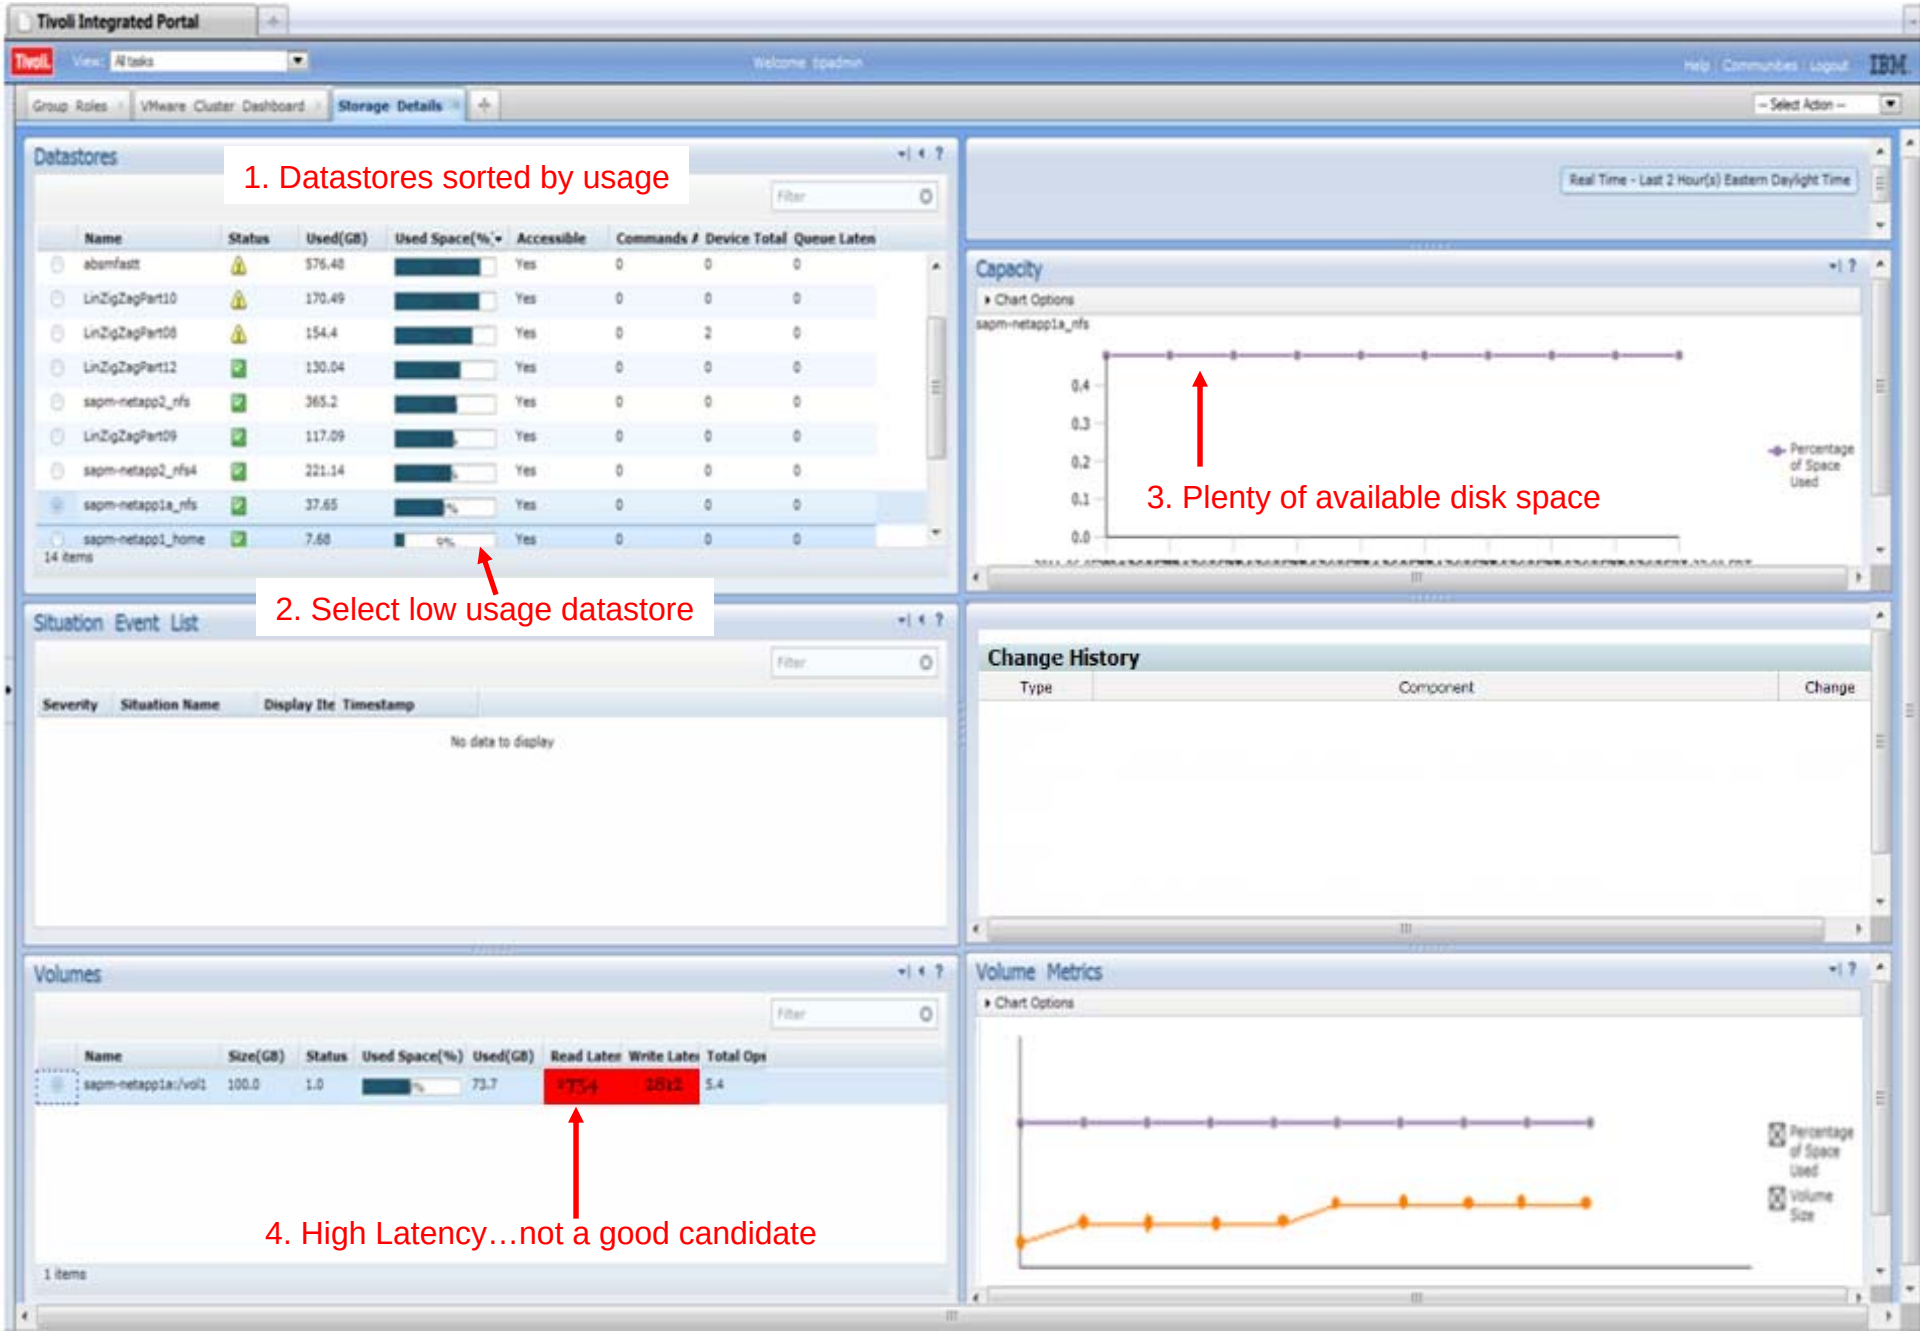

Found a Suitable Datastore

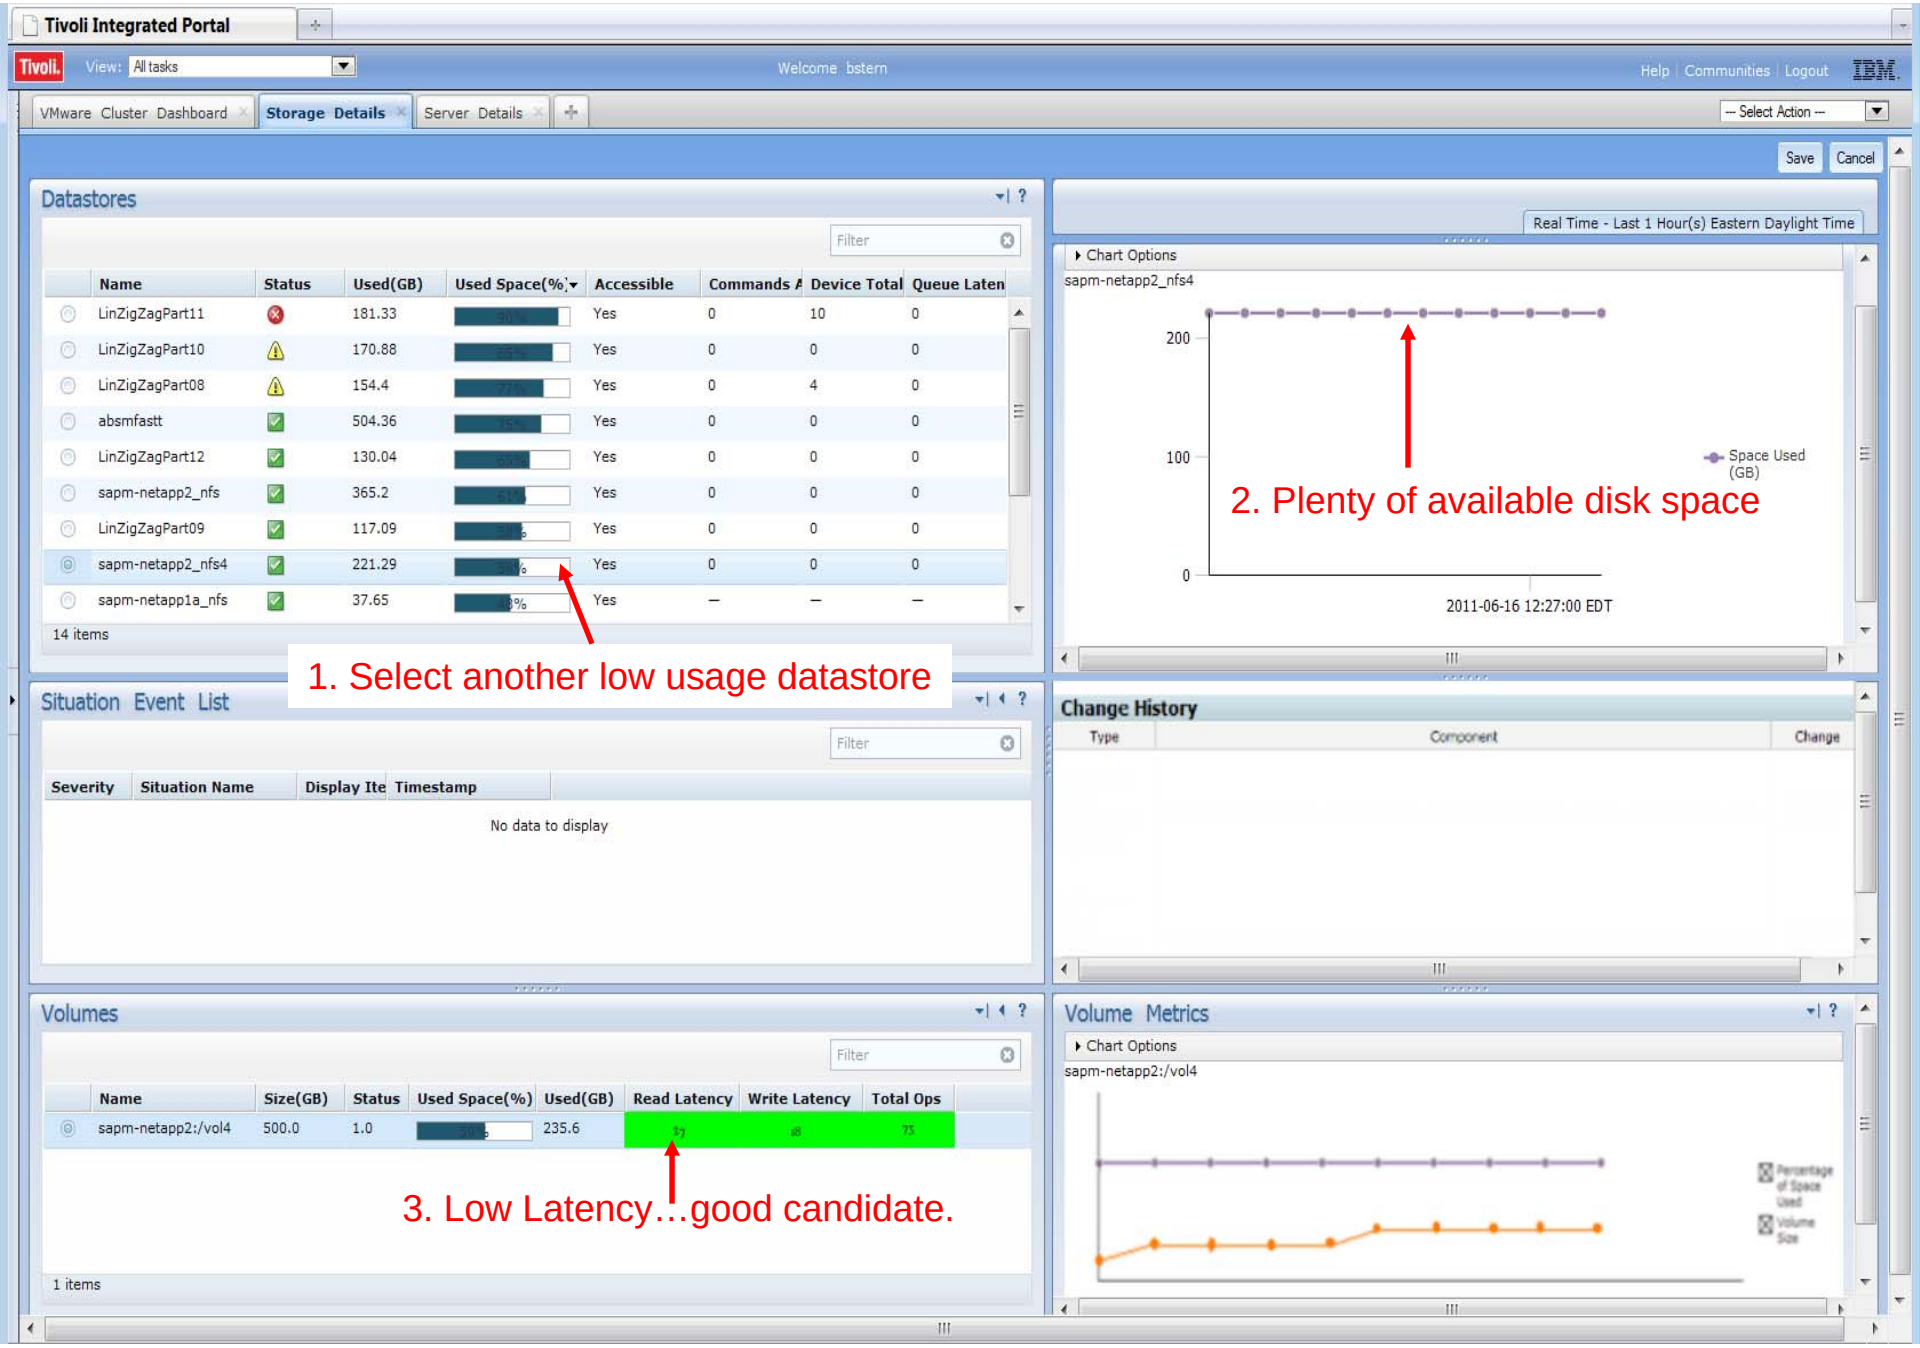

Storage Critical Issue Resolved

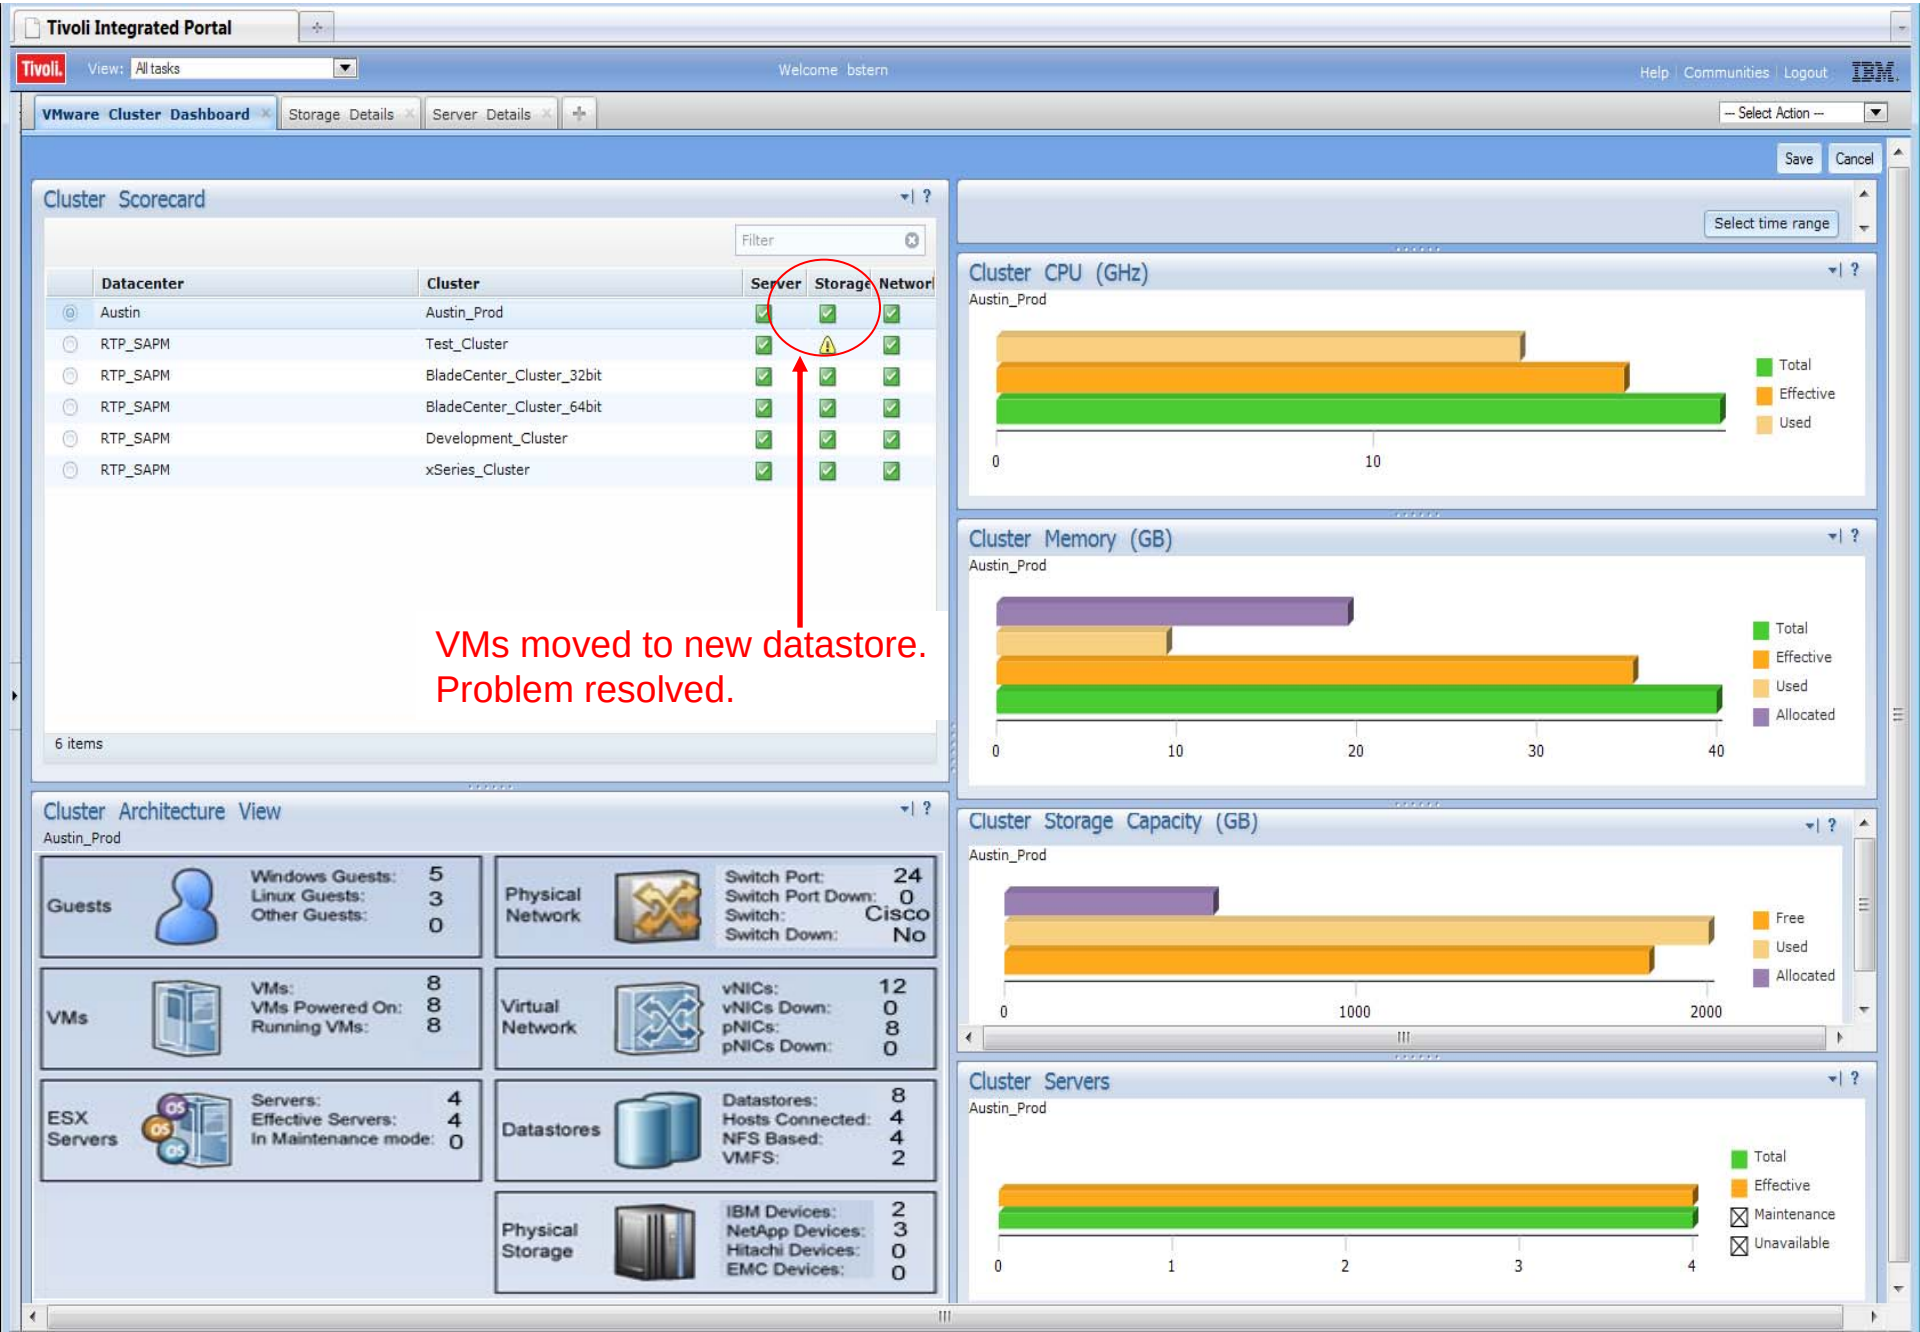

Managing Cloud Environments – Scenario 2

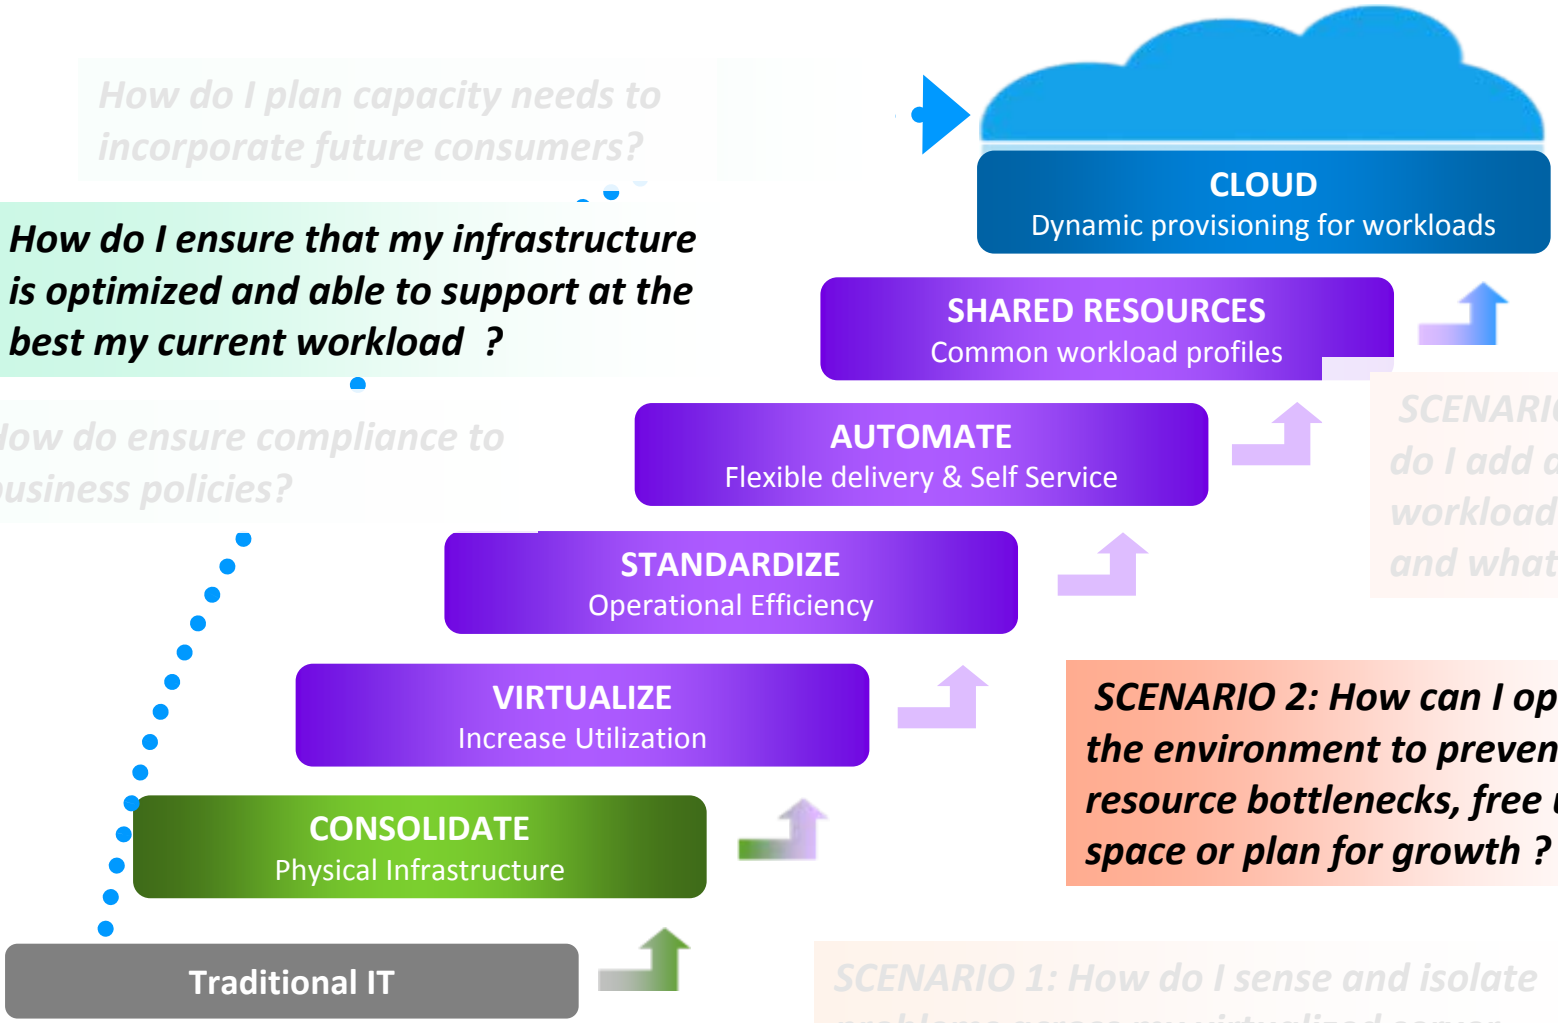

Sample Scenario

An IT Admin rebalances workload to avoid future performance and capacity bottlenecks

Cluster Health Scorecard showing Server Problem

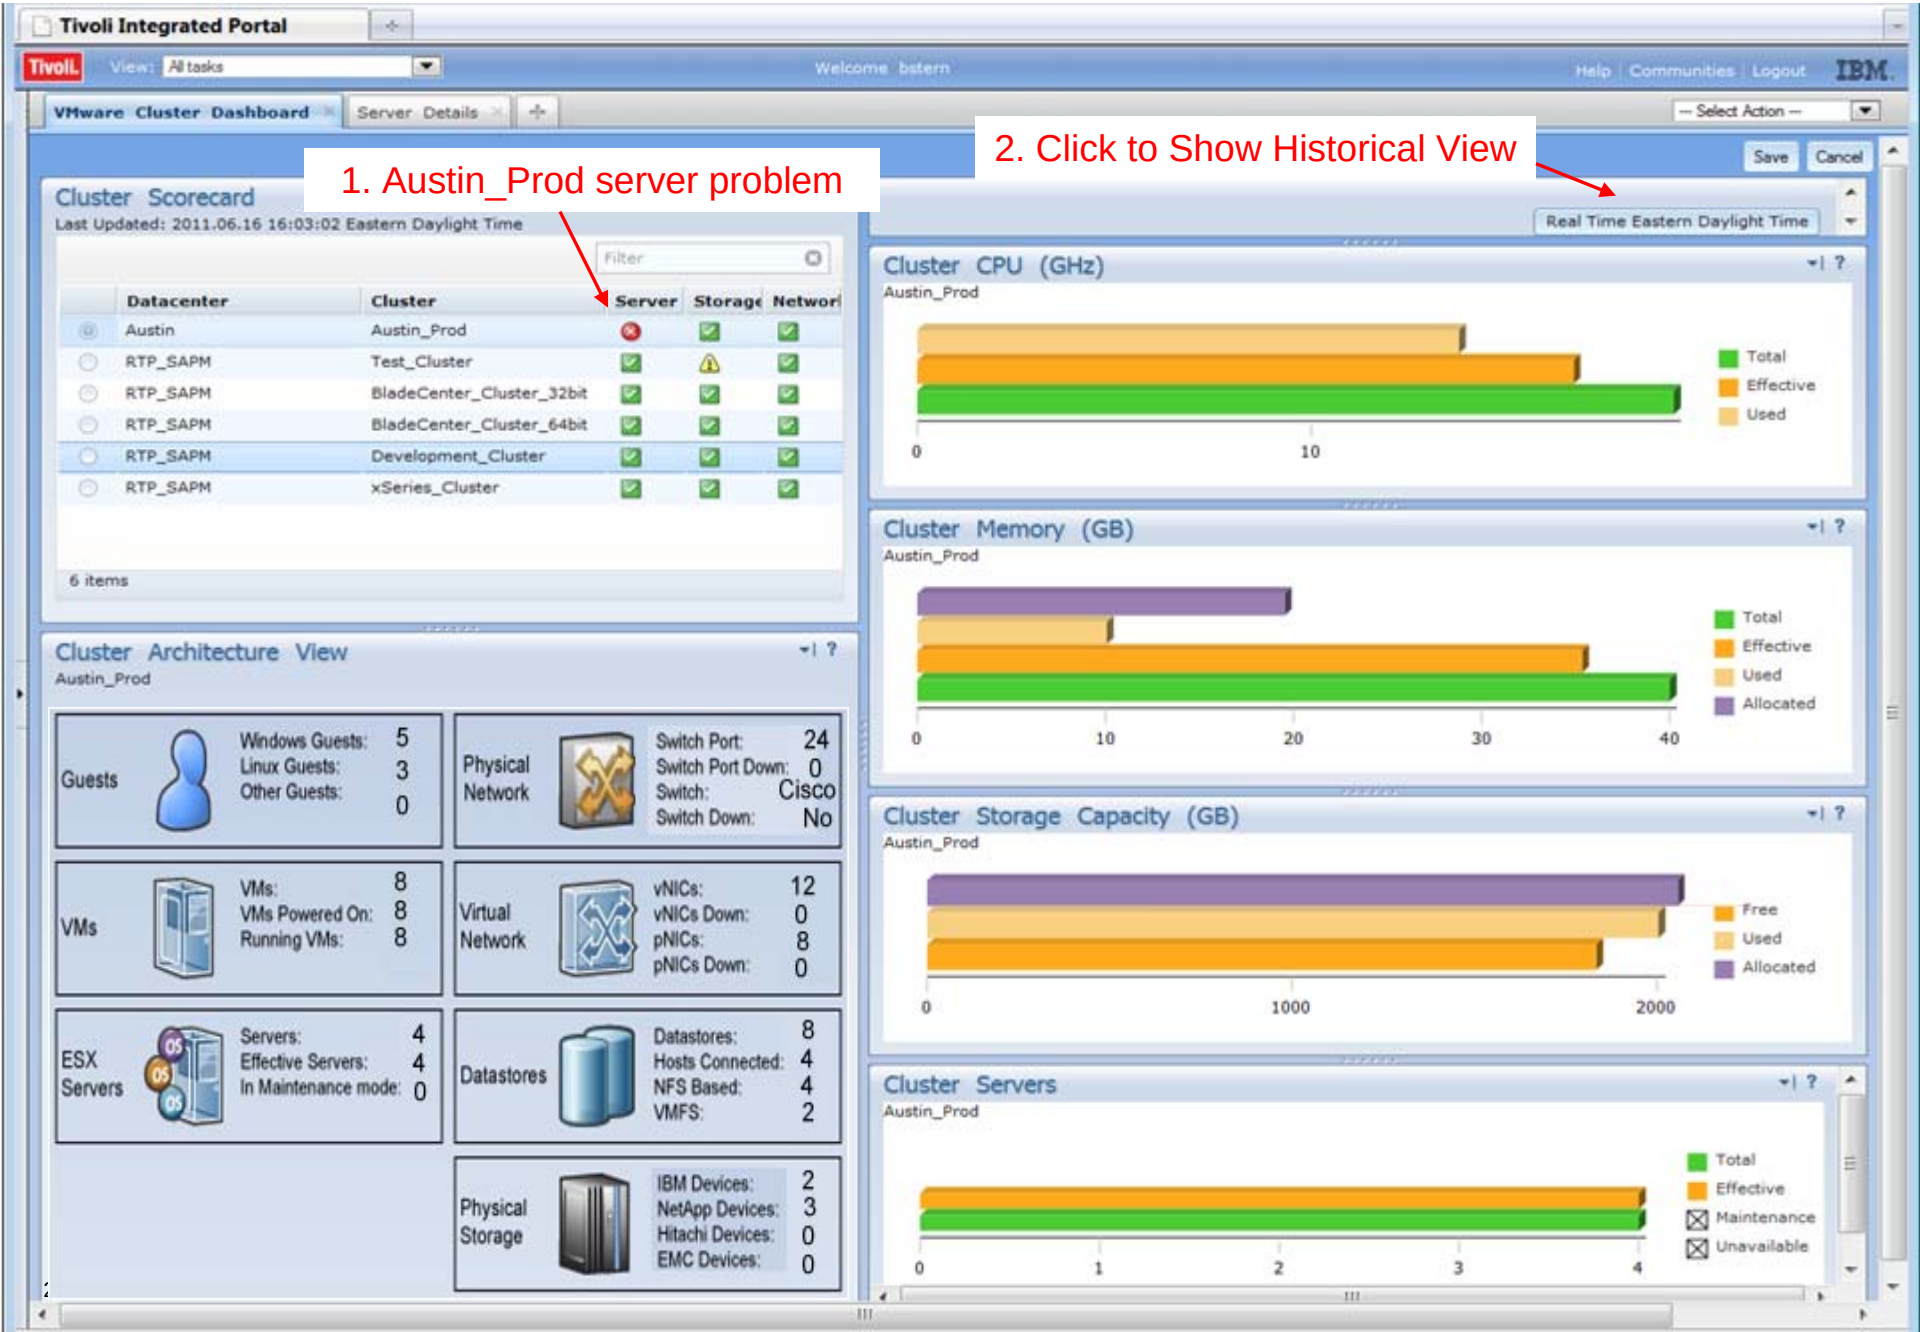

Predict CPU Utilization Critical for Cluster

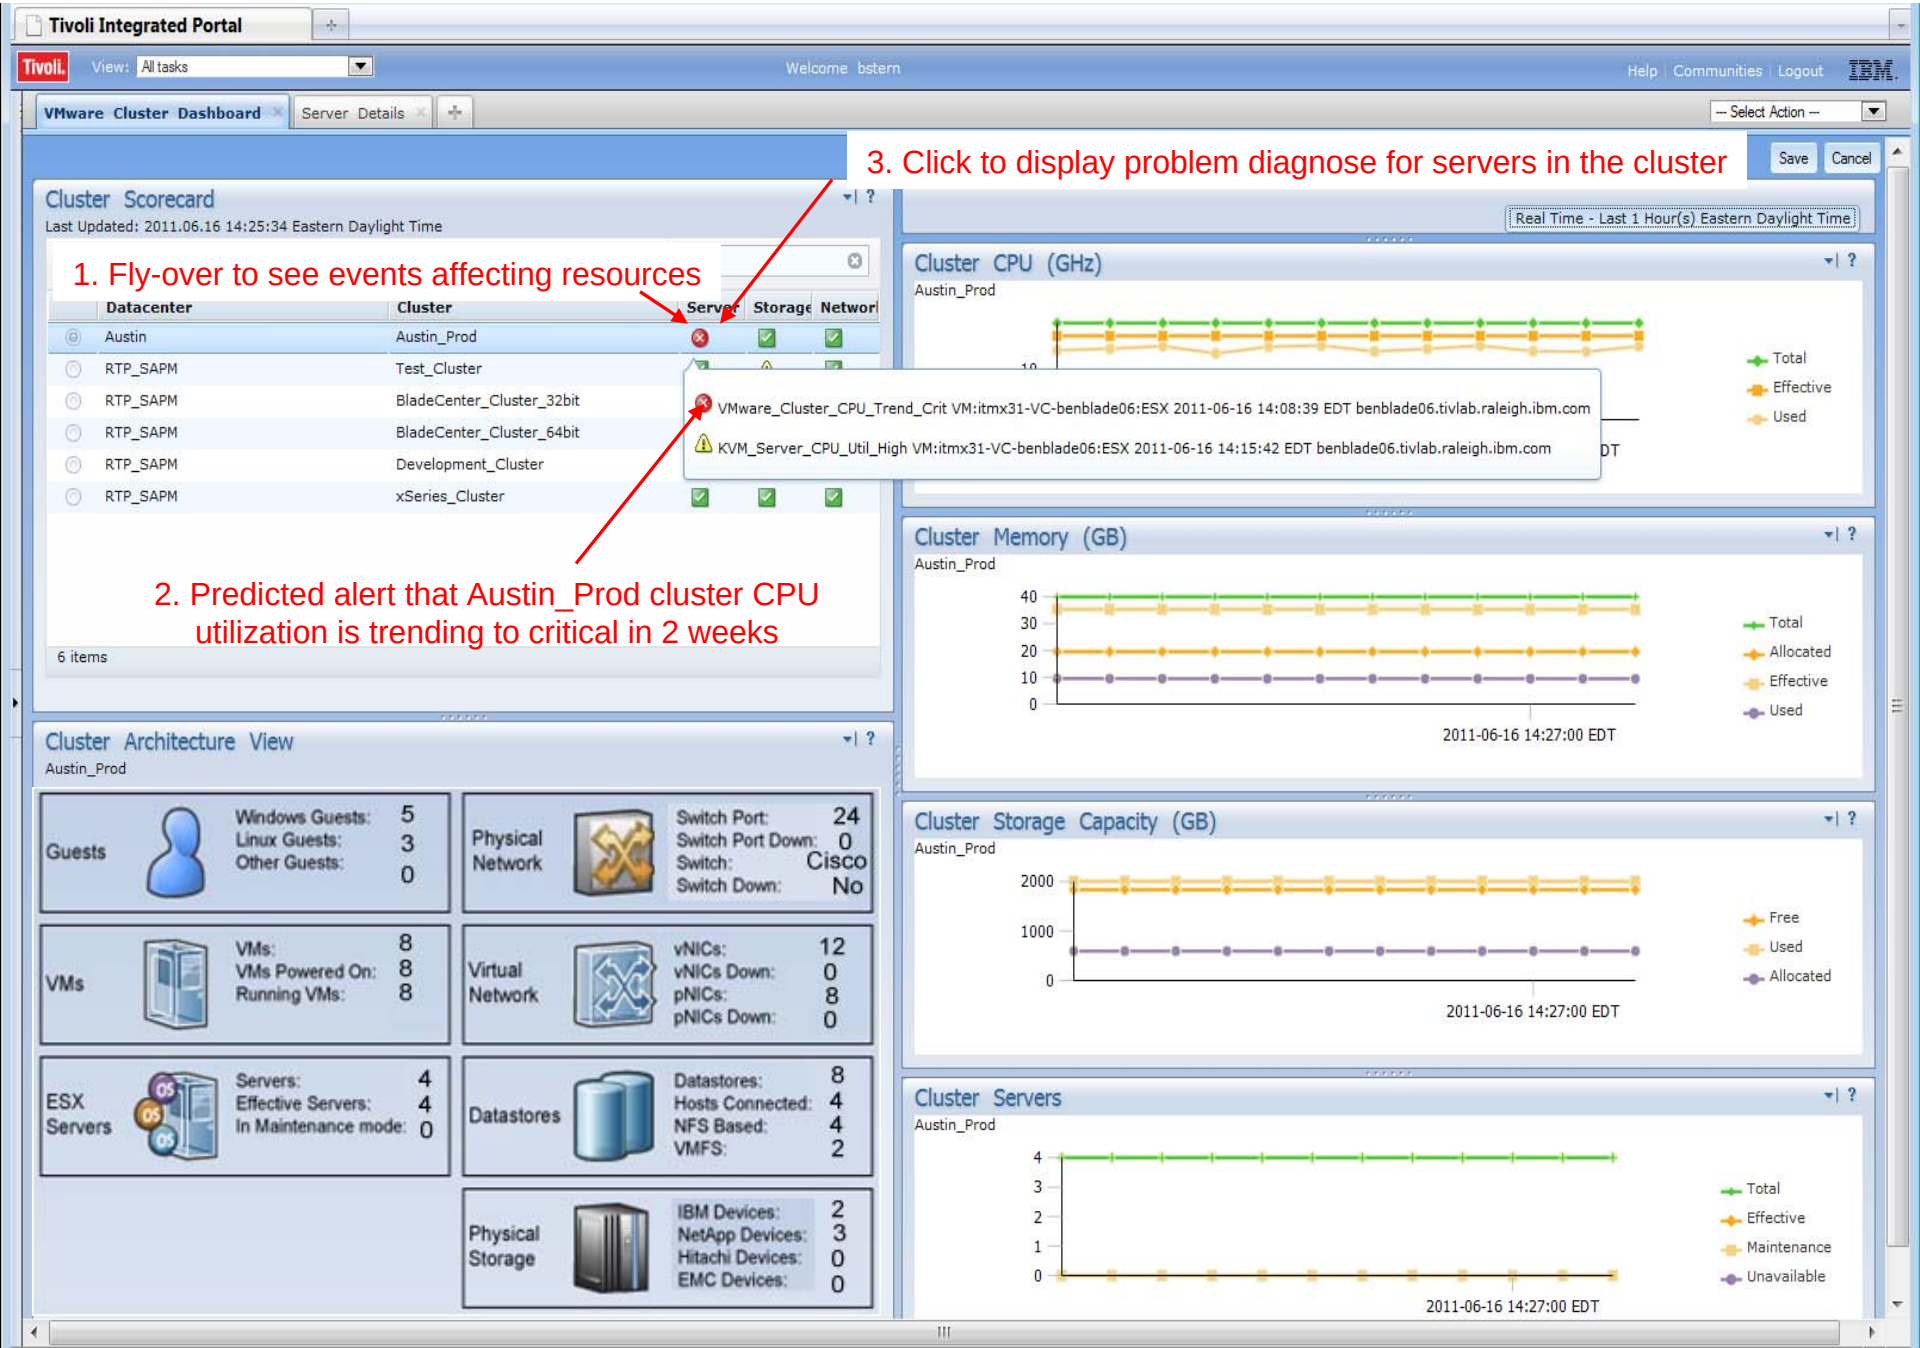

Diagnosing Server Problem for Austin_ Prod Cluster

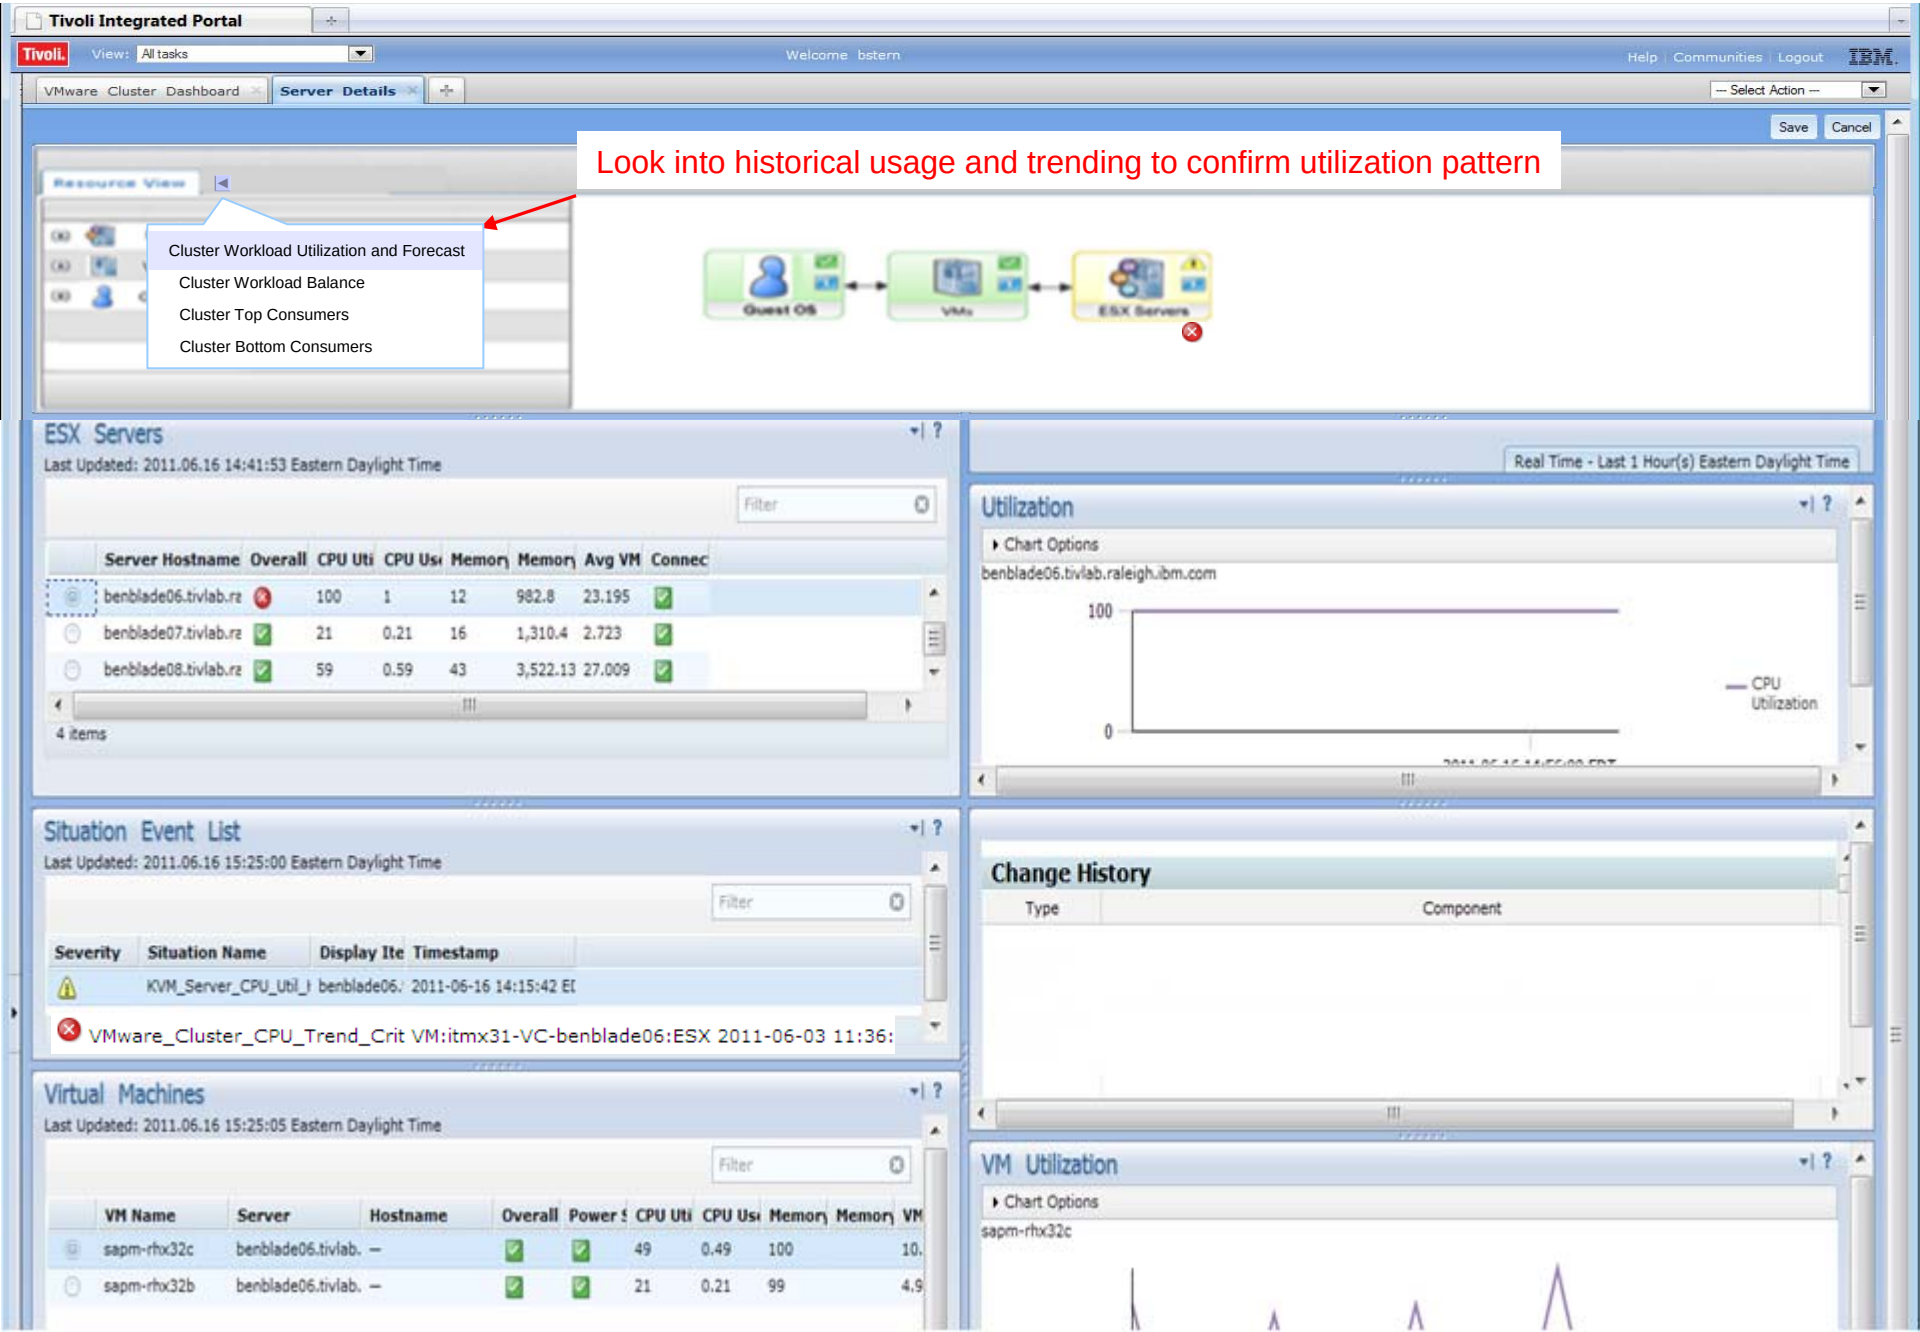

Austin_Prod Cluster Historical View

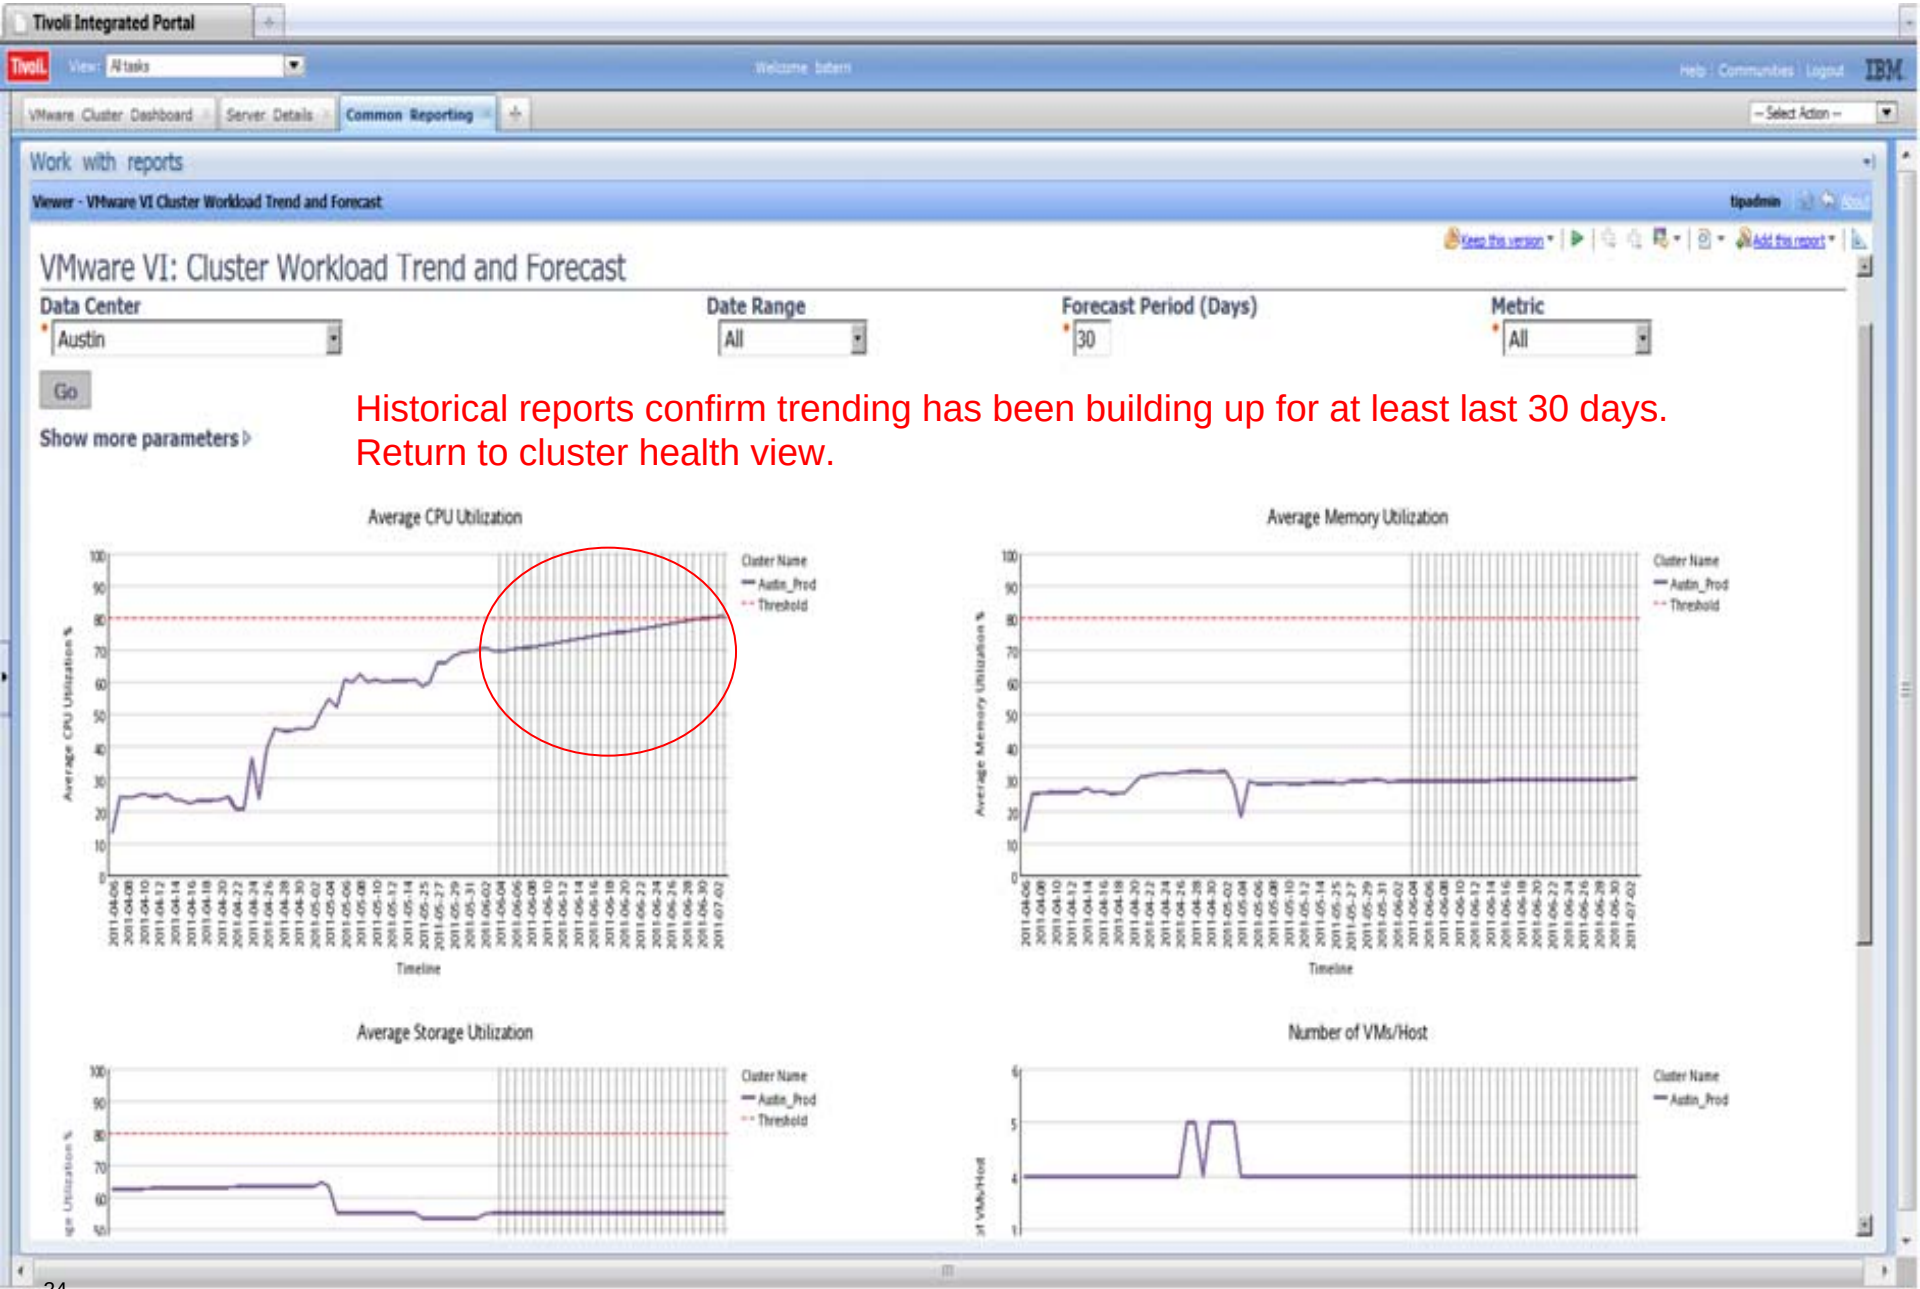

Need to Find Clusters to Spread VMs

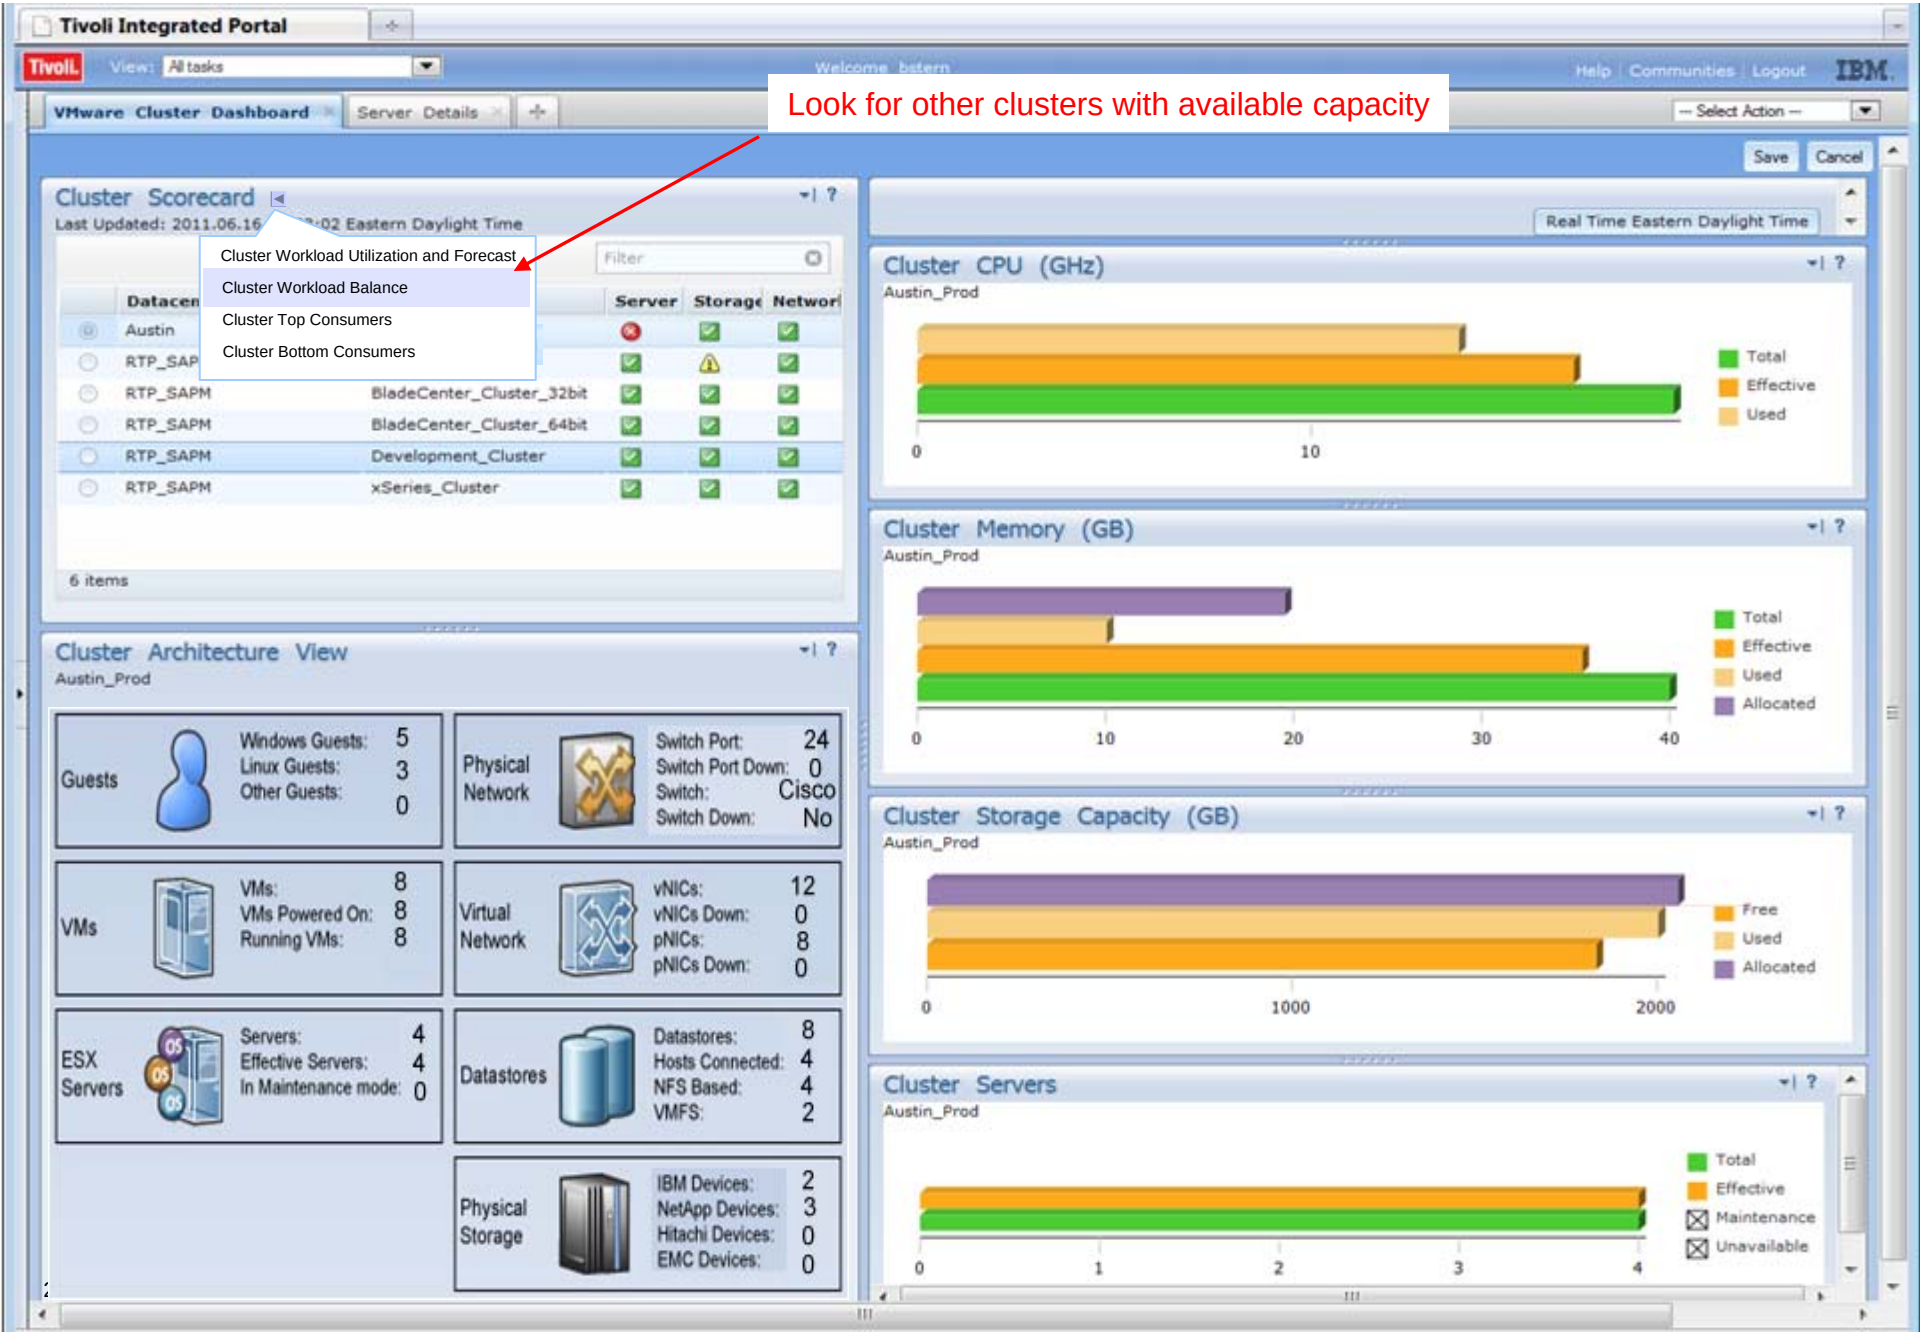

Found Clusters to Balance Load

Tivoli Integrated Portal

Tivoli. View: All tasks Welcome bstern Help | Communities | Logout

Common Reporting | VMware Cluster Dashboard | Server Details -- Select Action --

Work with reports

Viewer - VMware VI Balanced and Unbalanced Clusters tipadmin About

Keep this version | Add this report | Watch new versions | Add comments

Tivoli. software IBM.

VMware VI Balanced and Unbalanced Clusters

Date filter	Last 30 days		
Start Date	May 15, 2011 12:00:00 AM	**End Date**	Jun 14, 2011 11:59:59 PM
Data Center	RTP_SAPM		
Cluster(s)	ALL		
Shift Periods	All Shifts		
Vacation Periods	All days		

CPU Utilization

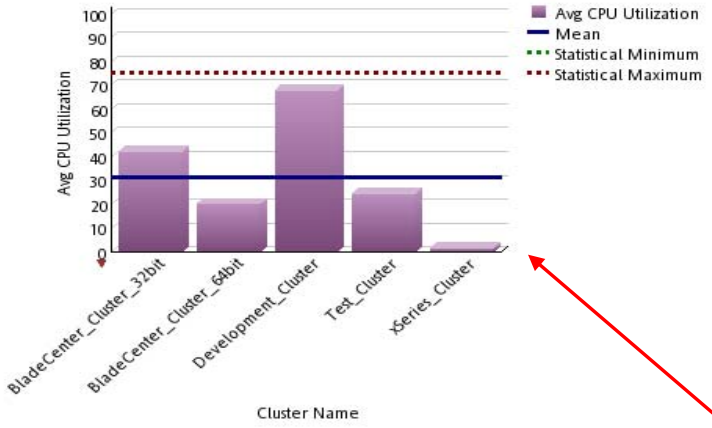

Memory Utilization

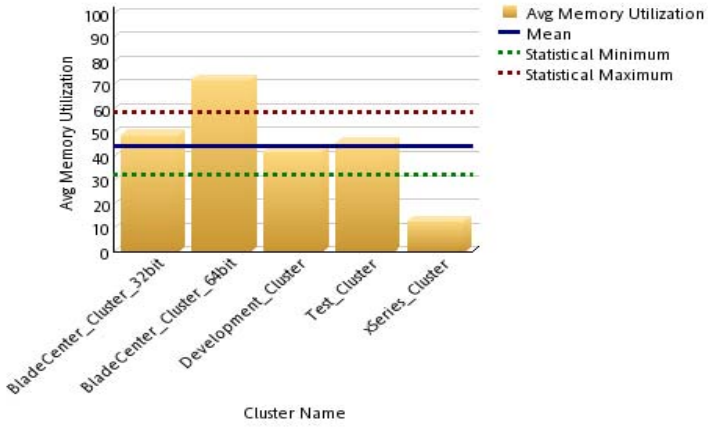

Datastore Space Utilization

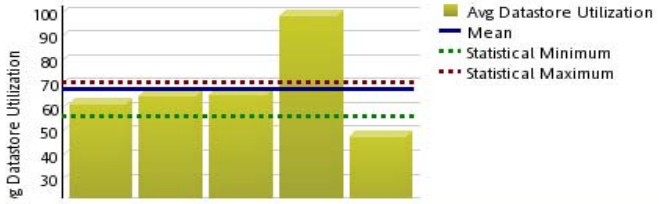

Couple of clusters are under-utilized in terms of CPU and within normal utilization for other resources.

Use **PlanningCenter** to rebalance and optimize workload across the clusters

Major scenarios in Managing Cloud Environments

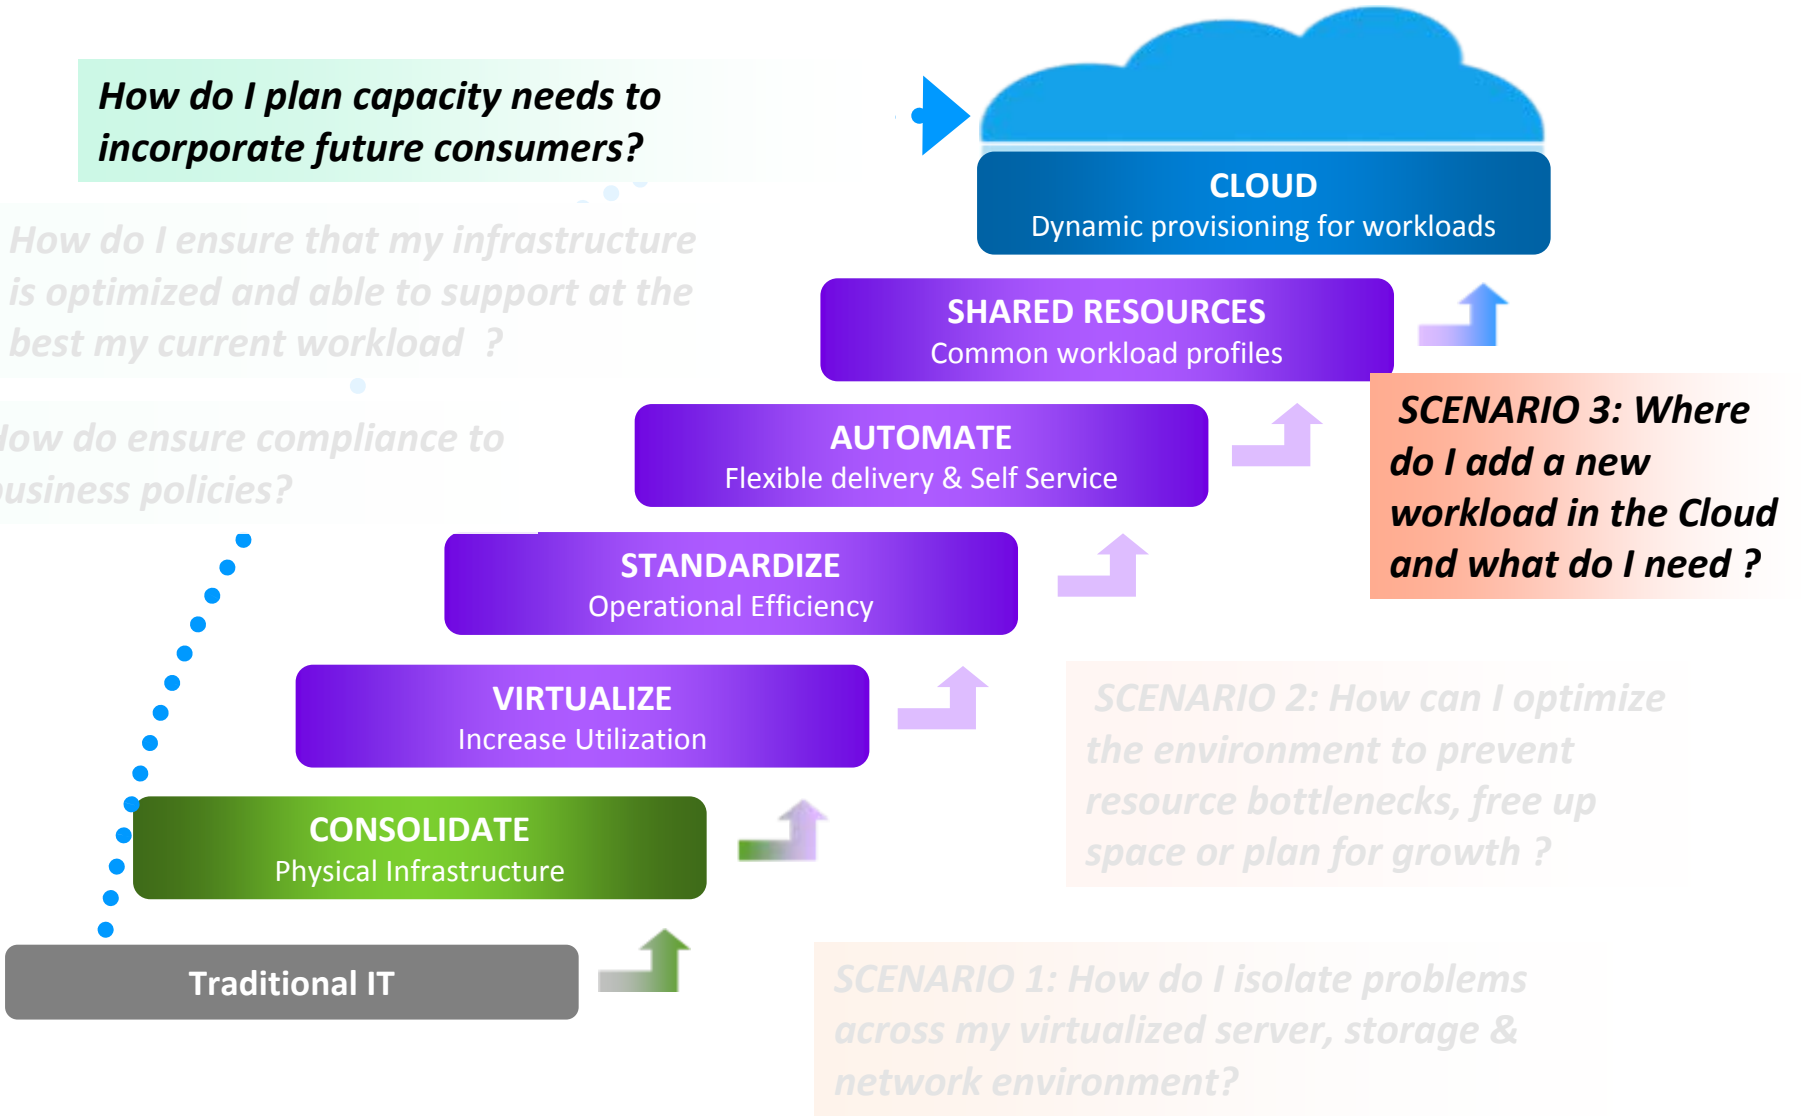

Sample Scenario

An Administrator performs a "what-if" analysis to understand if it possible to accommodate capacity growth

How many new VMs can be hosted on my current infrastructure

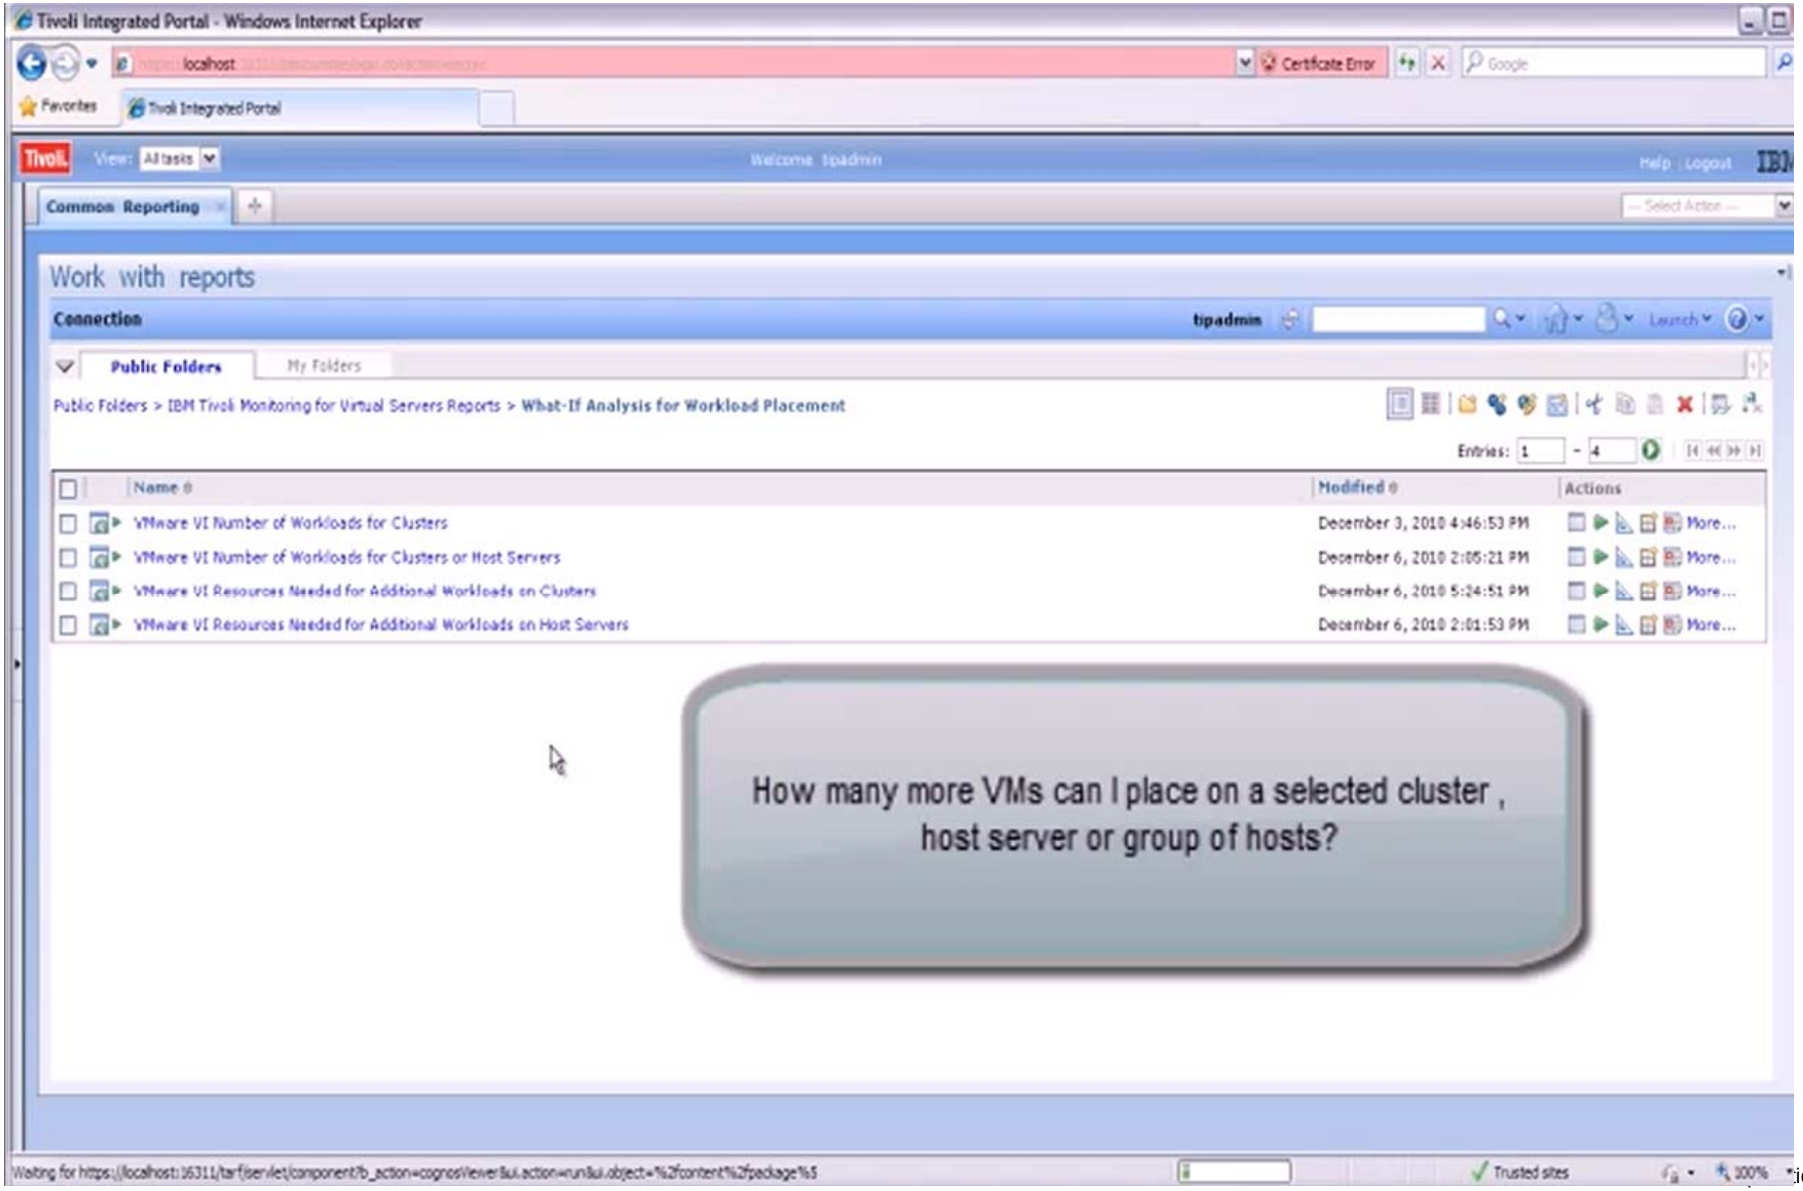

Specify the infrastructure that should host the new VMs

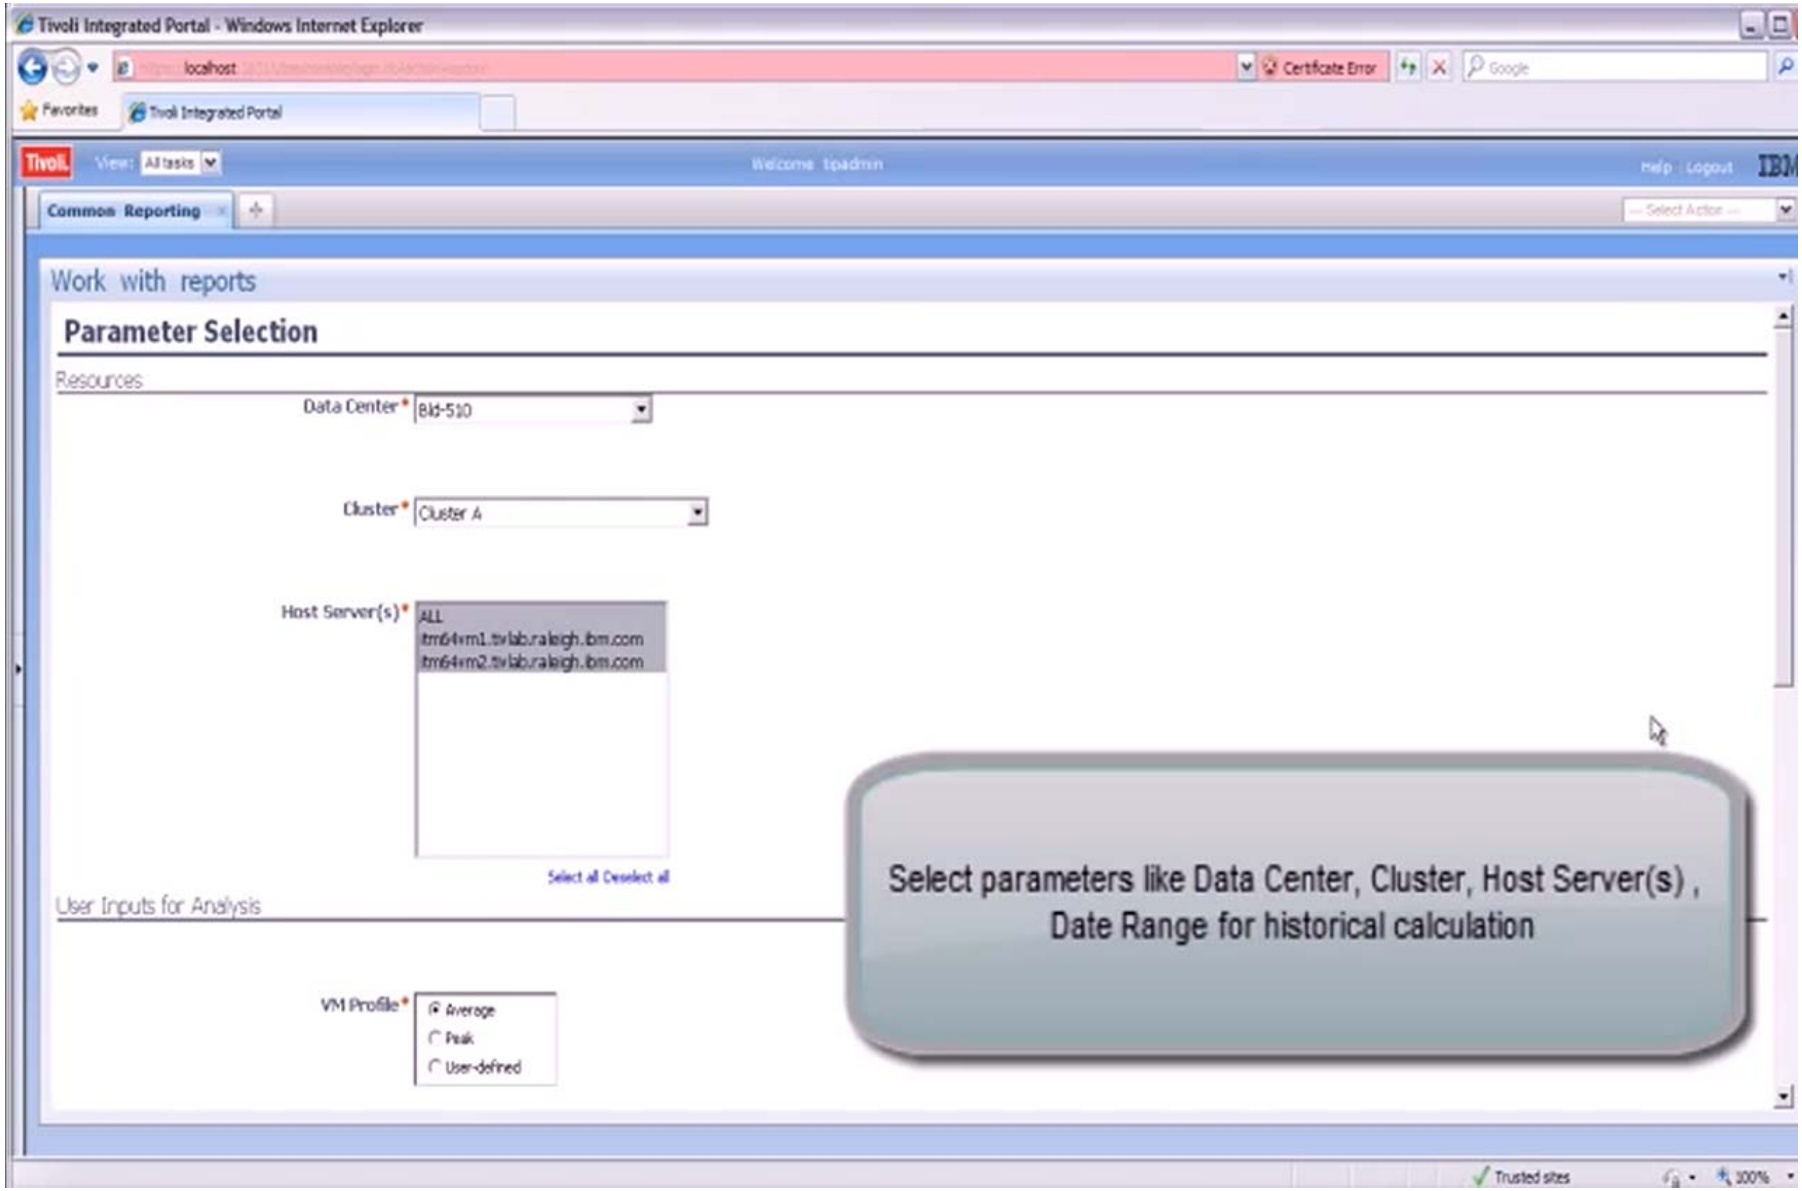

Specify typical VM size or use default templates

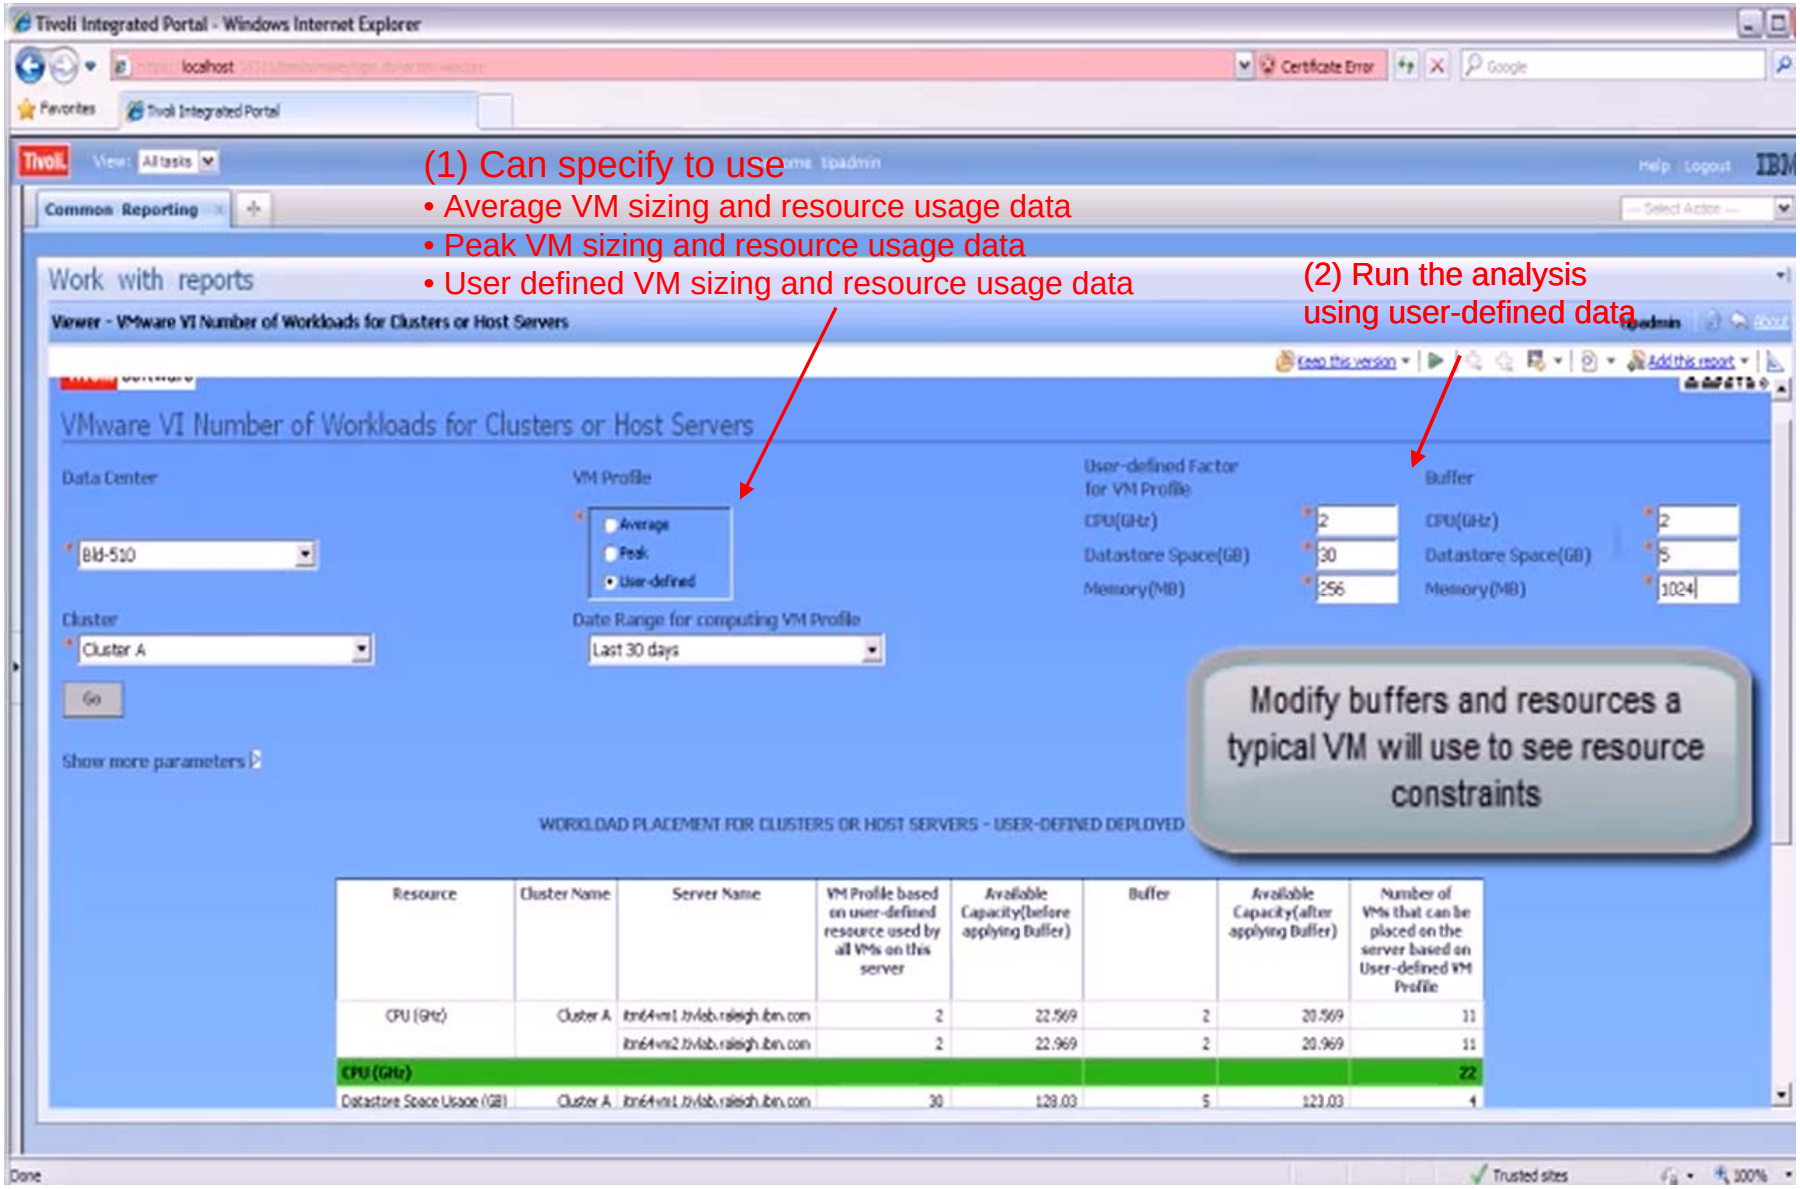

See how many new VMs can be hosted based on the resource constraints of the current infrastructure

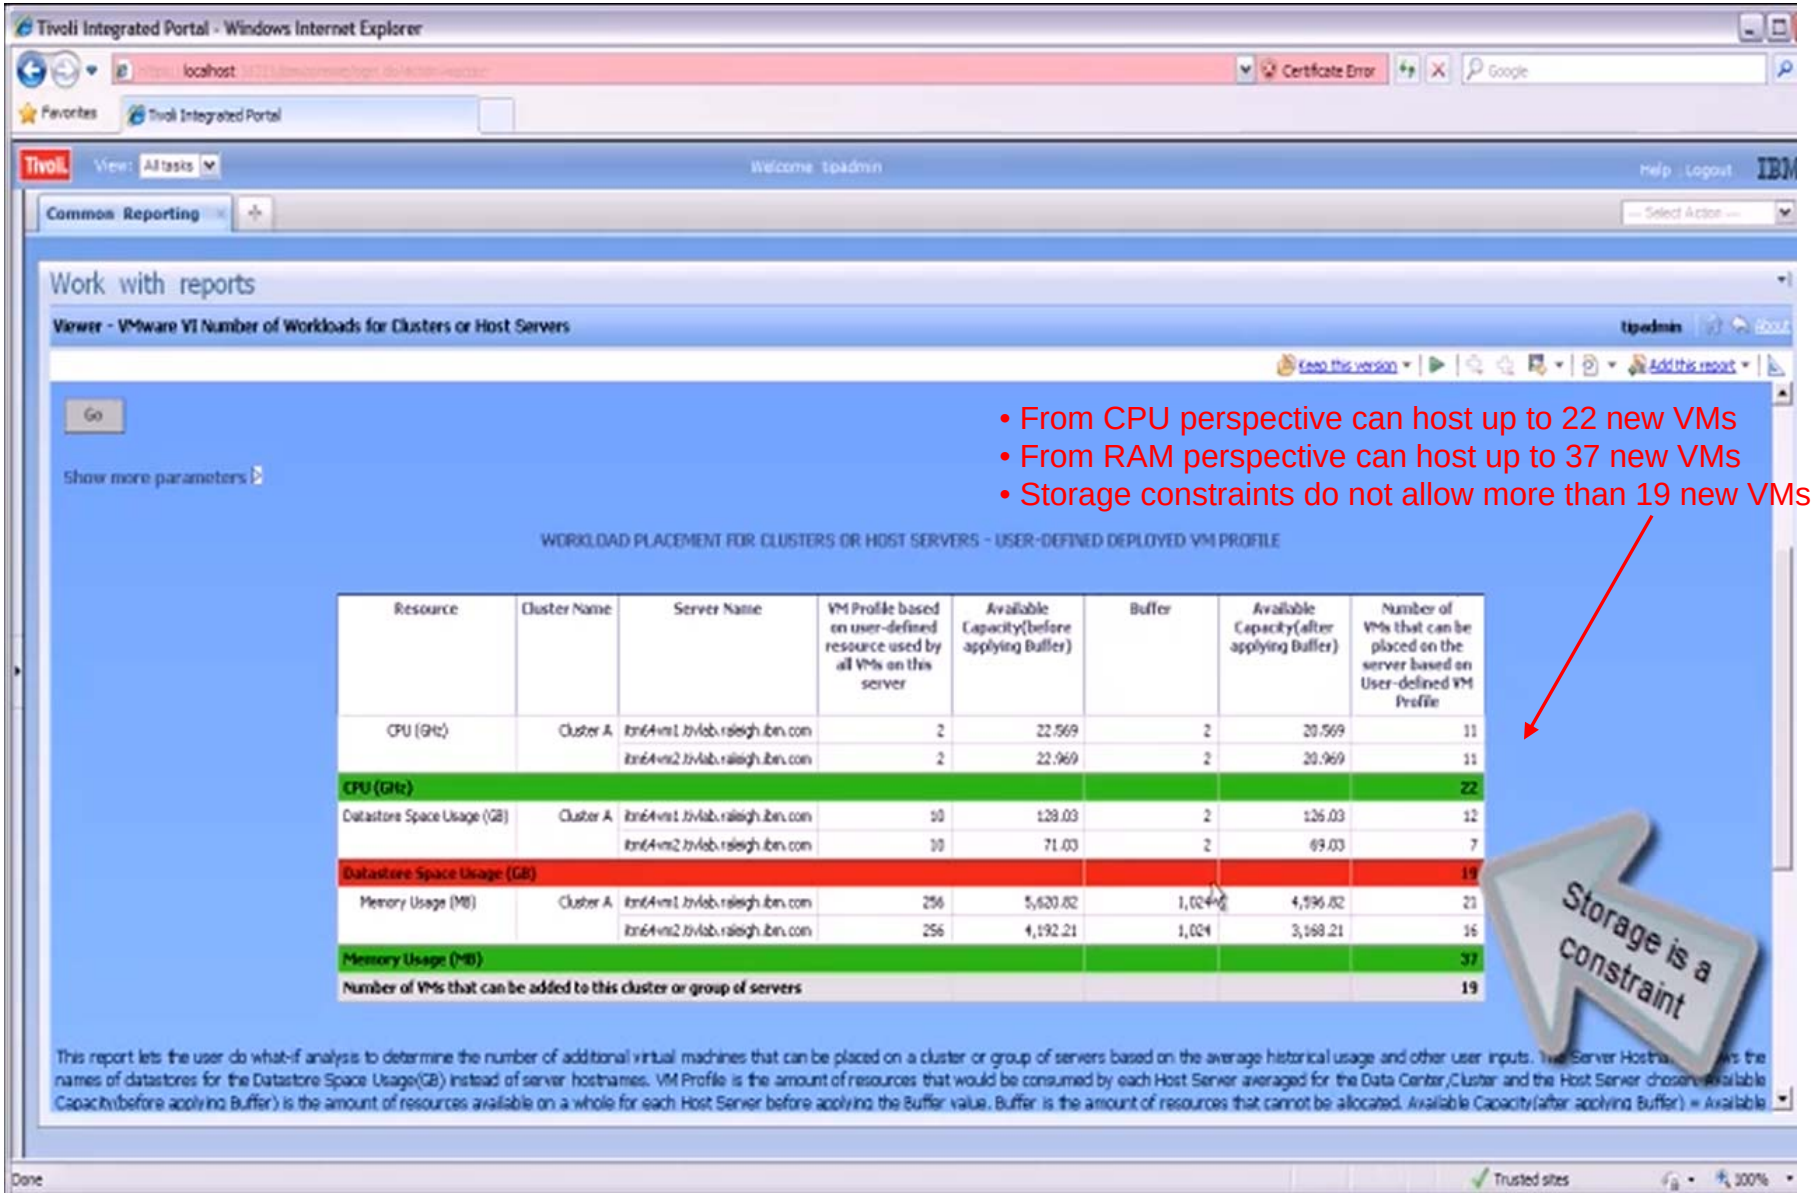

Need to add X new VMs, is the current infrastructure able to support the new Workload ? How much additional resources are needed ?

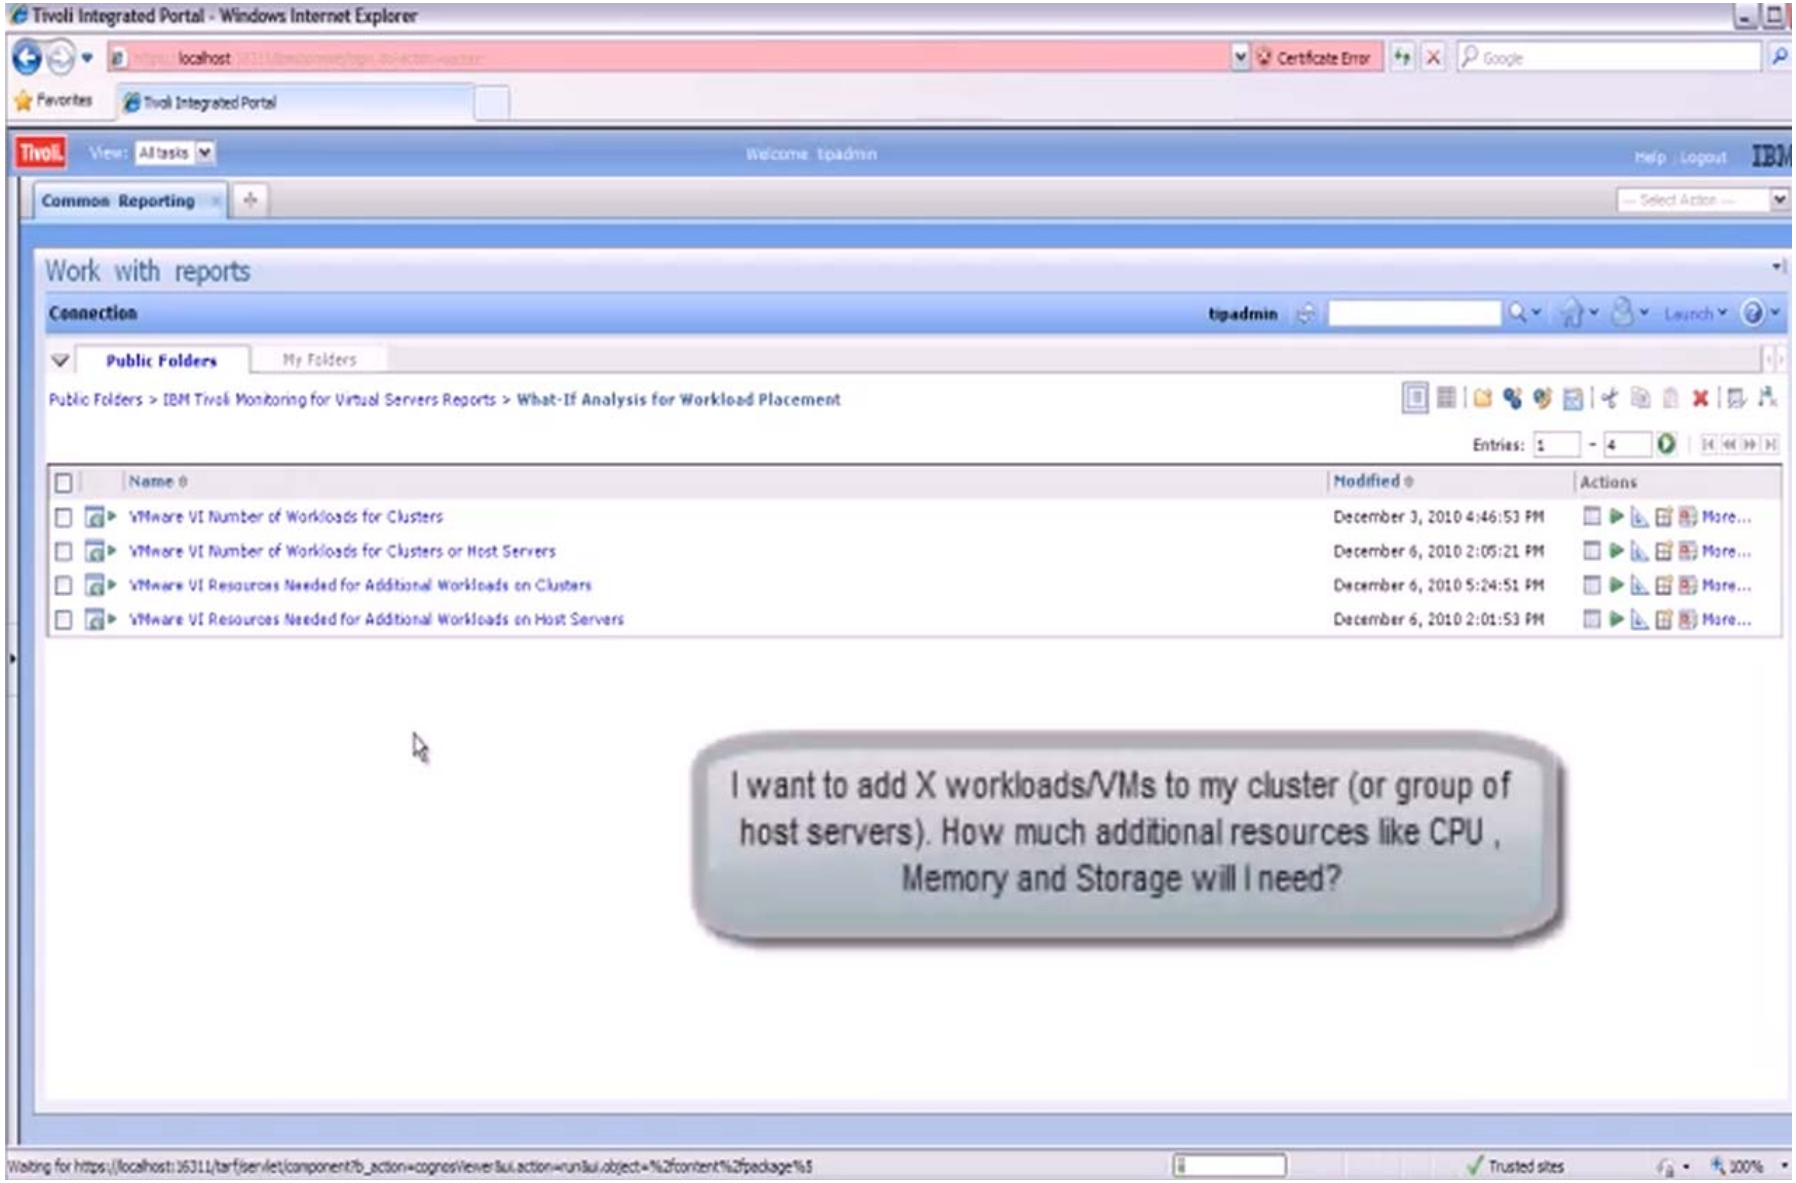

Current infrastructure is able to host 5 new VMs

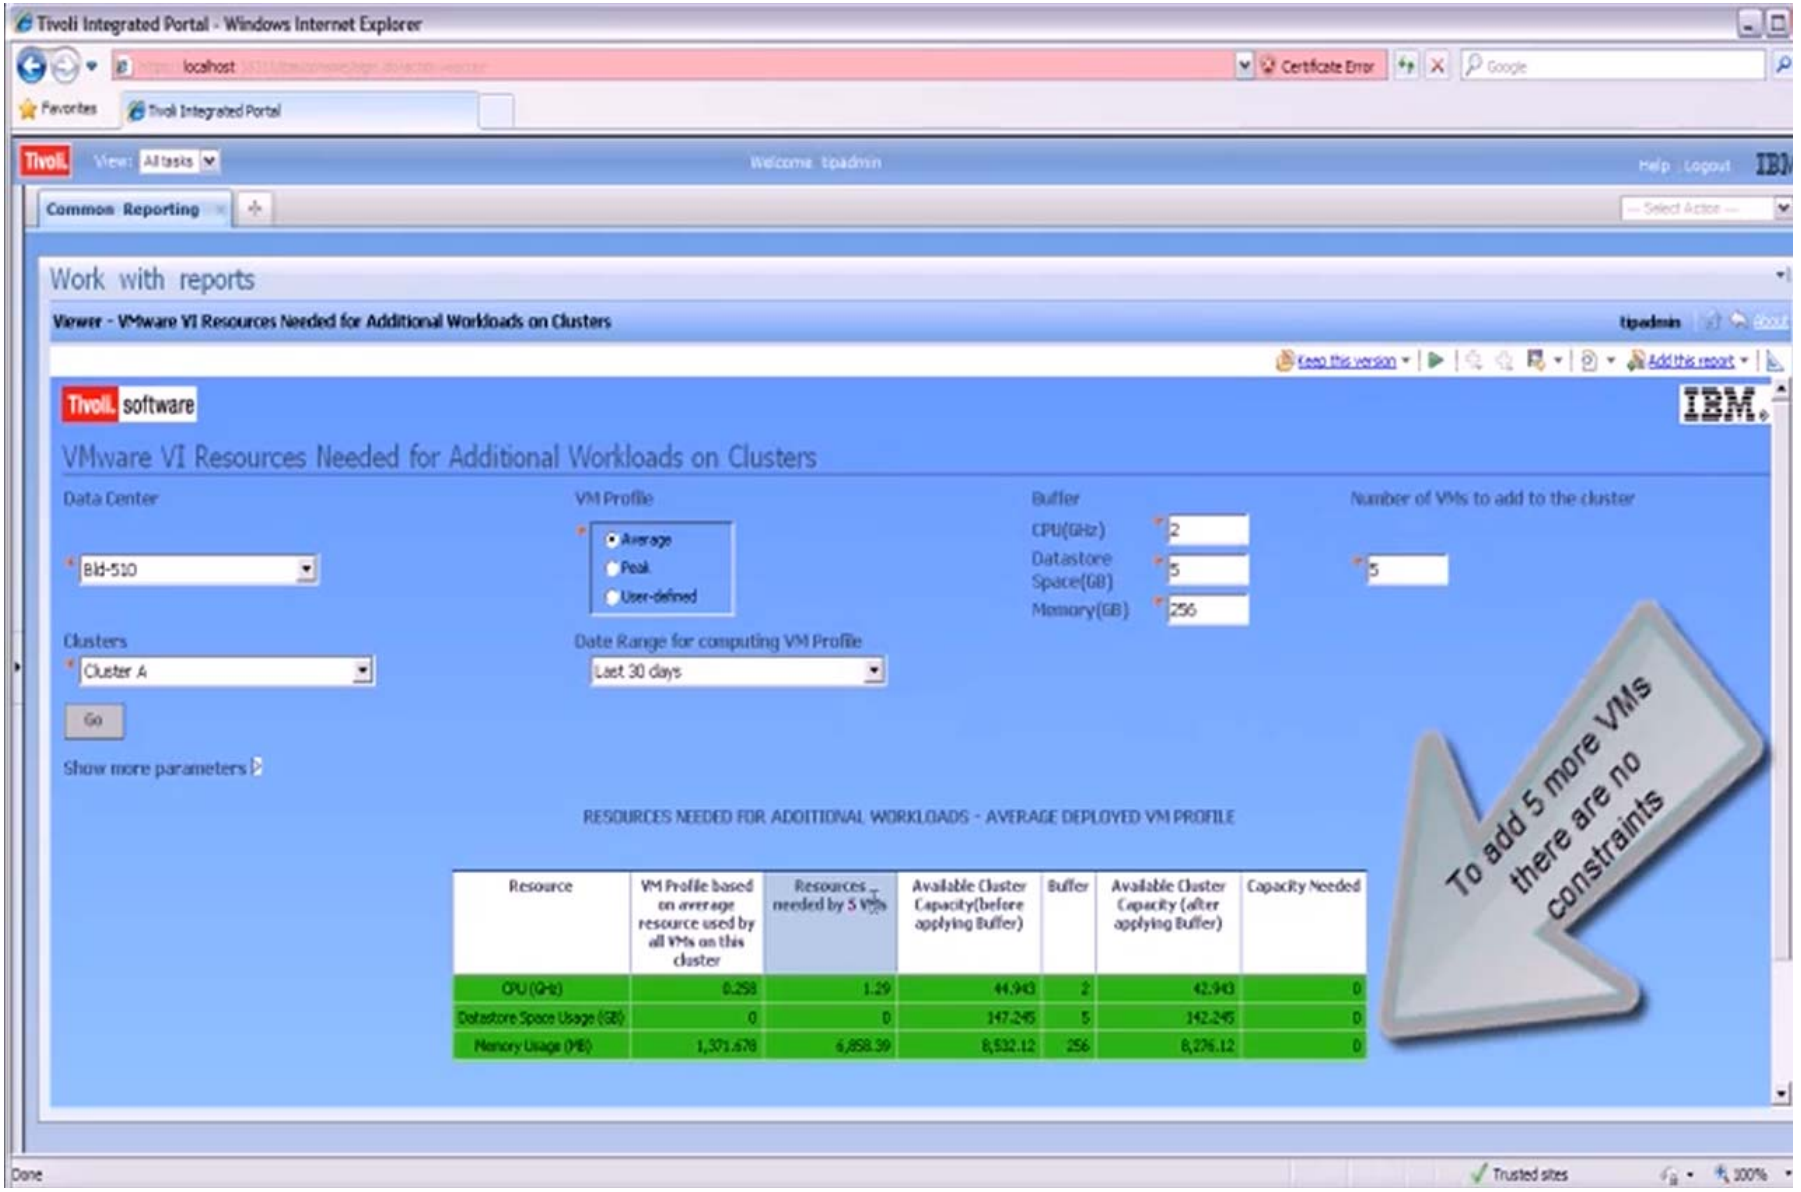

To add 50 VMs you need to add x Memory