

# Introducing SmartCloud Monitoring



*Antonio Sgro*  
*SAPM Best Practices Lead*  
*Email: [antonio.sgro@it.ibm.com](mailto:antonio.sgro@it.ibm.com)*

## Pulse2012

Optimizing the World's Infrastructure

**March 4-7** Las Vegas, Nevada



## Disclaimer

- *IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion. Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision. The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code or functionality. Information about potential future products may not be incorporated into any contract. The development, release, and timing of any future features or functionality described for our products remains at our sole discretion*



# Agenda

- Why are we having a Cloud conversation today?
- Challenges of managing a Cloud/Virtualized environment
- SmartCloud Monitoring Overview
- Key Features
- Integration into the SAPM Portfolio
- SmartCloud Monitoring Architecture
- Why is SmartCloud Monitoring Important
- VMware Health Dashboard
- Capacity Planning
- Demo
- Additional Resources

# Cloud computing is fundamental to changing the economics of business infrastructures and speeding the delivery of innovative products & services

## Improve security and compliance control posture

Embrace new business opportunities while maintaining control and mitigating risk.



## Improve speed and dexterity

Speed the delivery of new offerings and services by creating new models of self-service and deployment.

## Deliver IT without boundaries

Implement new systems and management processes that simplify access to information in order to deliver better business outcomes.



## Create new business value

Empower internal and external communities to define and create new offerings and services.

# What we are hearing from our customers

Infrastructure Teams



*How can I improve my resource utilization, simplify administration and reduce cost?*

Line of Business Teams



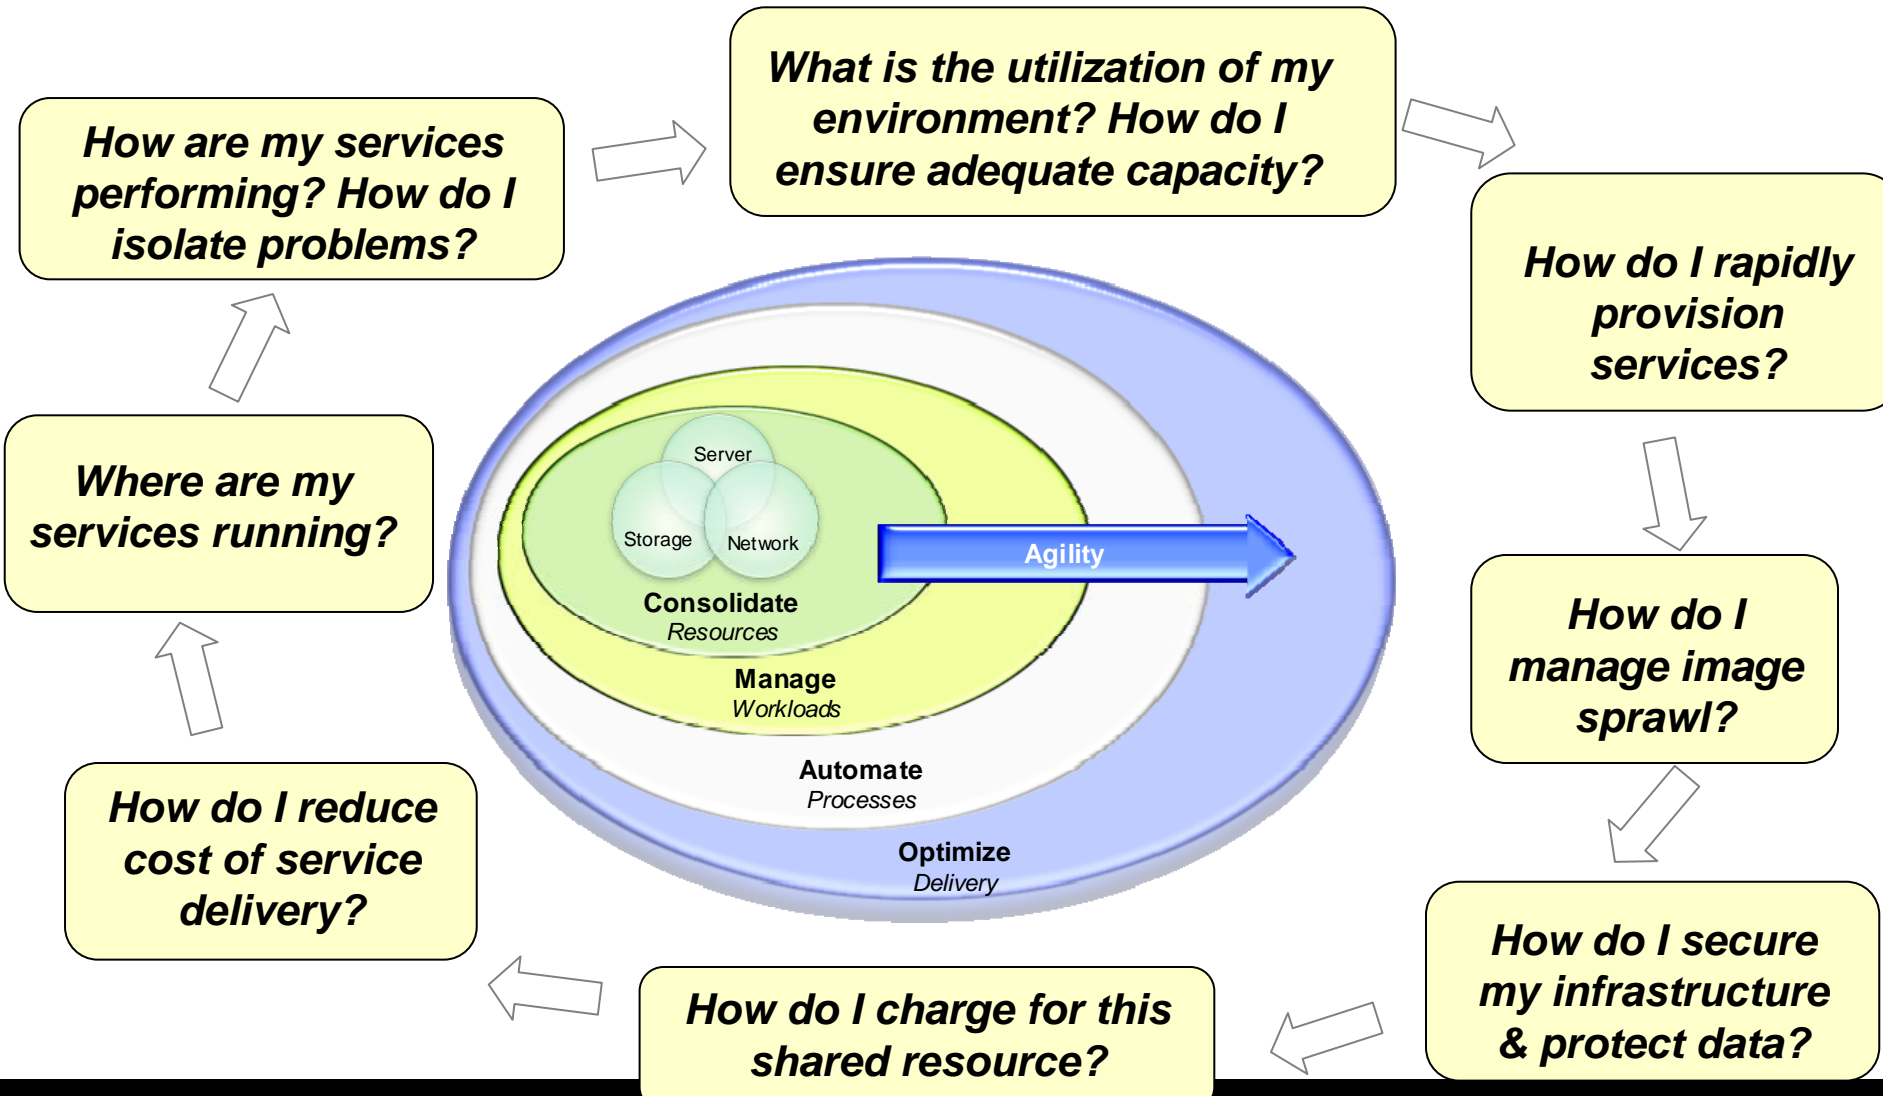
*How can I accelerate my application release cycle in an optimized fashion and in a quality fashion?*



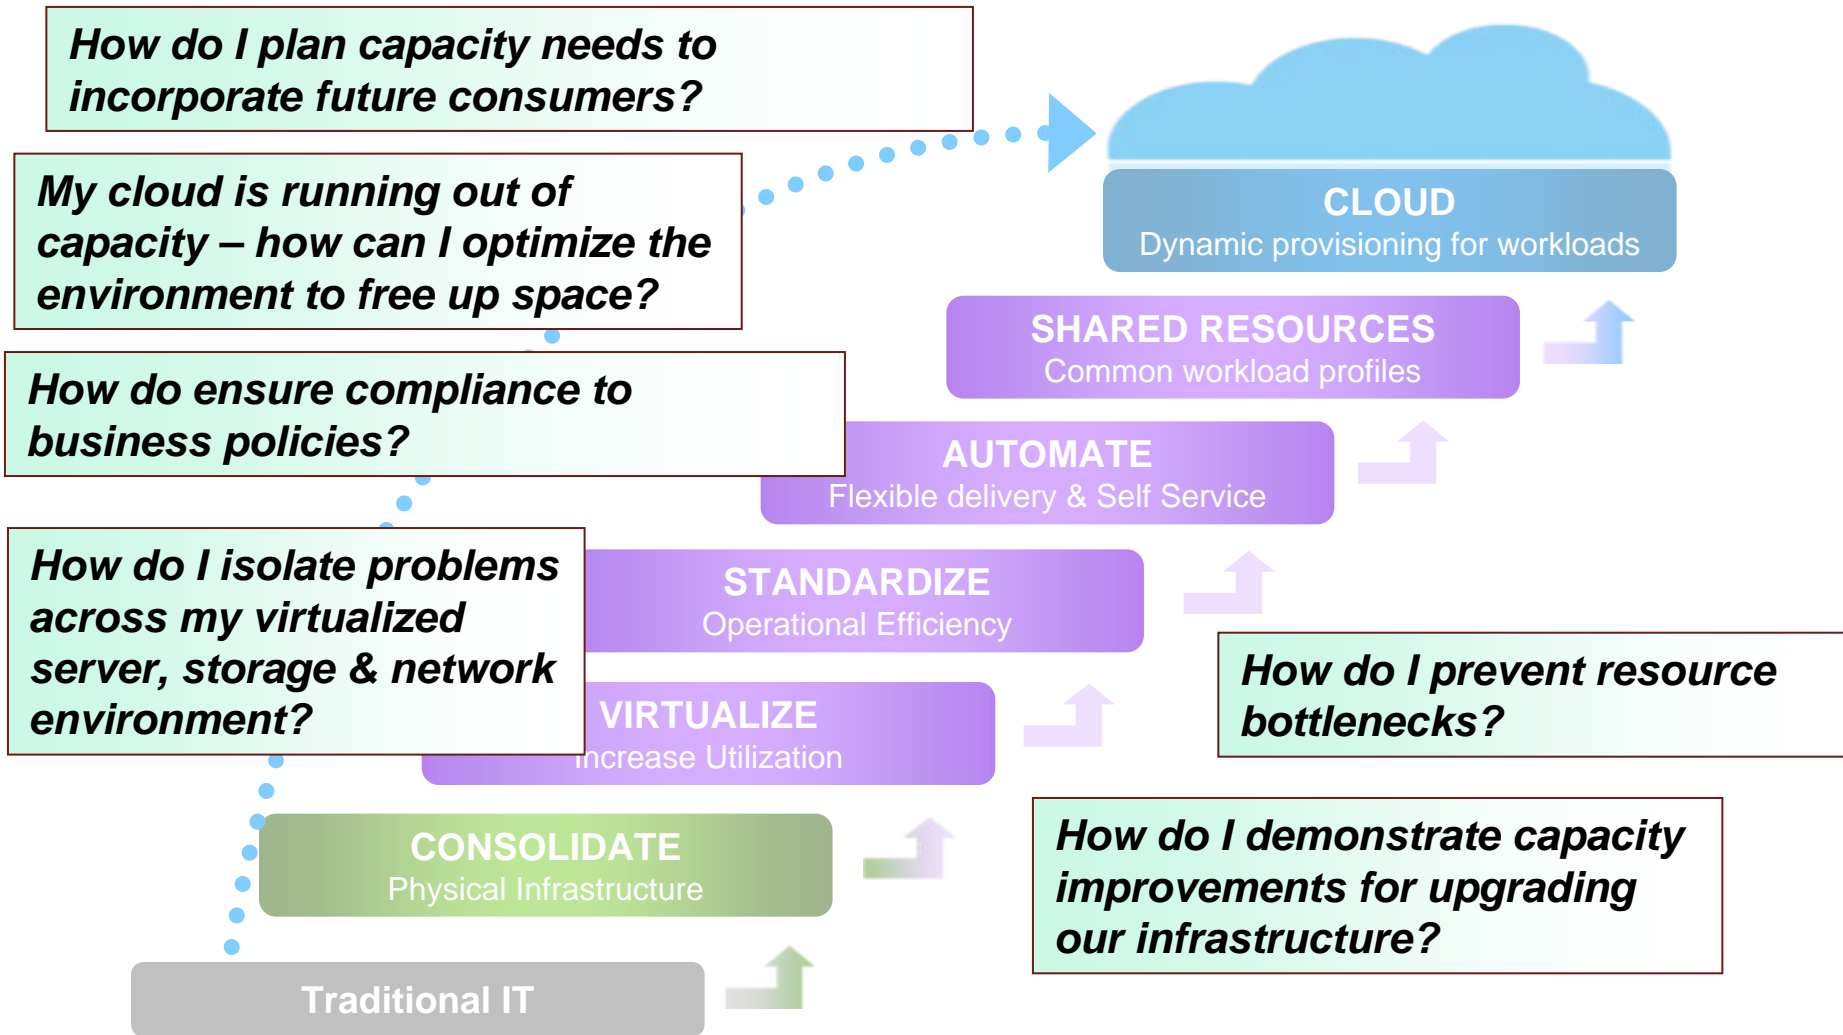
Operations Teams

*How can I improve responsiveness and drive productivity and efficiency while maintaining stringent qualities of service?*

# Addressing the challenges of managing a Cloud/Virtual environment



# Path from Traditional IT to Cloud






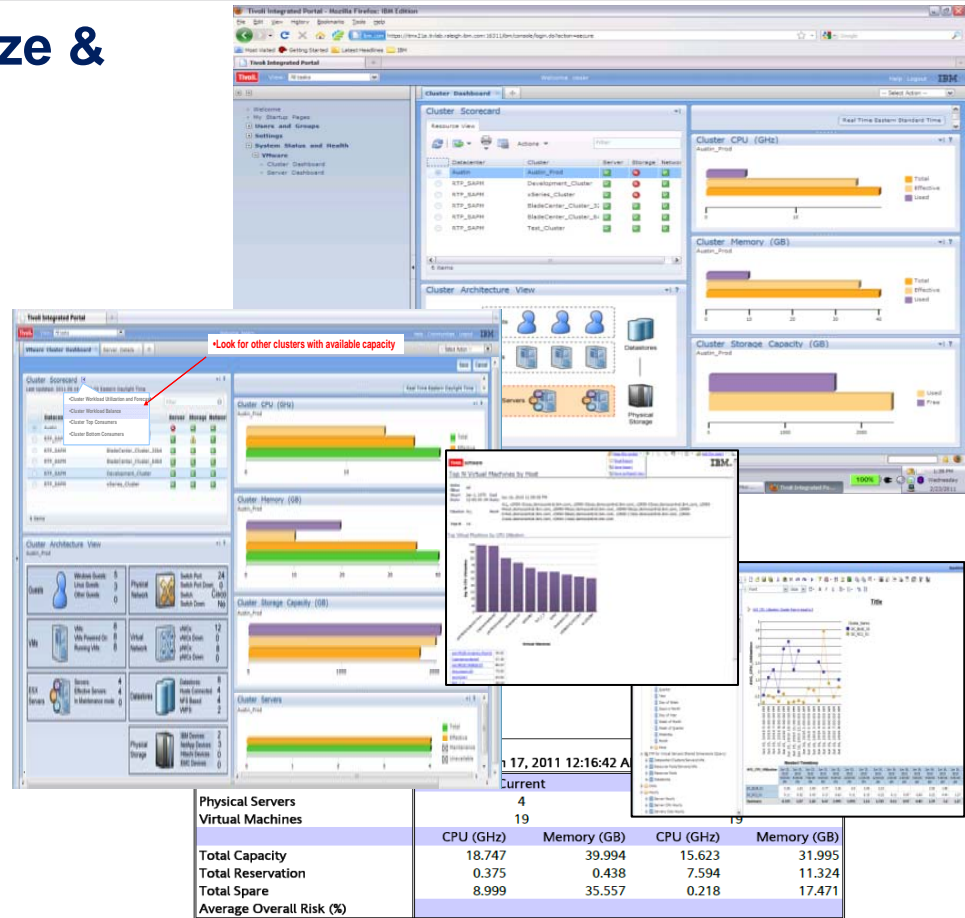
# SmartCloud Monitoring

- **What is it?** *SmartCloud Monitoring (SCM) is a bundling of IBM Tivoli Monitoring for Virtual Environments v7.1 and IBM Tivoli Monitoring OS agent monitoring v 6.2.3*
- **What does it do?** *SCM monitors the health and performance of a VMware infrastructure as well as the virtual machines running on it, tracking the availability of shared resources on the host at the same time it tracks the consumption of those resources by individual workloads.*
- **What business problem does this help to solve?** *In order to realize true cost savings from a virtualization or cloud investment, customers need to be able to run virtual machines densely enough to maximize consolidation, yet be assured that their workloads are still running as well as they were before being virtualized, with room for expansion.*
- **Are there any prerequisites or co-requisites that need to go with this offering?** *No. SCM contains a full ITM infrastructure and agents that can connect to VMware and other hypervisors out-of-the-box*
- **Product Integration:** *SmartCloud Monitoring is an IBM Tivoli Monitoring based solution. Therefore, it can be easily integrated with the Business Service Management part of the portfolio and can be integrated with SmartCloud Provisioning. The solution includes out of the box integration with NetApp storage, Tivoli Storage Productivity Center and Tivoli Application Dependency Discovery Manager (TADDM).*
- **Why do you need SmartCloud Monitoring?** *Because you're not getting true TCO out of your virtualization or cloud investment, or your key business services are at risk because of resource contention or performance problems in the cloud environment.*



# Key capabilities to optimize & maintain a private cloud

- **Health dashboards** to provide an instant, consolidated glimpse into cloud health
- **Topology views** of the key interrelated components of the cloud
- **Reports** on the health trends of cloud components and workloads, powered by Cognos
- **What-If** capacity planning scenarios
- **Policy-Based** optimization to put workloads where they'll perform best, not just where they'll fit
- **Performance Analytics** for right-sizing of virtual machines
- **Integration** with industry-leading Tivoli service management portfolio





# Monitoring Agents SmartCloud Monitoring



## ITM for Virtual Environments:


- VMware
- KVM
- NetApp
- Network Monitoring Agent
- Citrix XenServer \*
- Citrix XenApp \*
- Citrix XenDesktop \*
- Cisco UCS \*

## ITM 6.2.3 OS Monitoring

- Windows
- UNIX
- Linux (x, p, and z)
- iOS
- Agent-less OS Monitoring

## Pricing:

- SMC is priced per Virtual Machine (counting peak usage, since these are perpetual licenses)



*\* Agents are sold separately from SmartCloud Monitoring with RVU pricing, but will be included in a future version*



# Physical & Virtual Server, Storage & Network Problems

## Storage Problems

- Data store issues
  - May be caused by insufficient space - # of VMs, storage usage growth, changes in configuration
- Storage Latency/response time issues
  - May be caused by too many VMs associated to same LUN (Volume), HBA bottlenecks, backend storage issues (disk, etc.)

## Network Problems

- Network response time problem
  - May be caused by too many VMs sharing NICs without necessary throughput capacity, changes in configuration
- Network connection problem
  - May be caused by unavailability of connection to physical network (physical switch / port)

## Server Problems

- VM, Host, Cluster CPU, Memory Utilization
- % ready
- Memory over commit (Active Memory/Physical Memory)
- Swapping/ballooning
- VM Swapping
- Absence of processes, growth of process resource consumption
- Host server failures



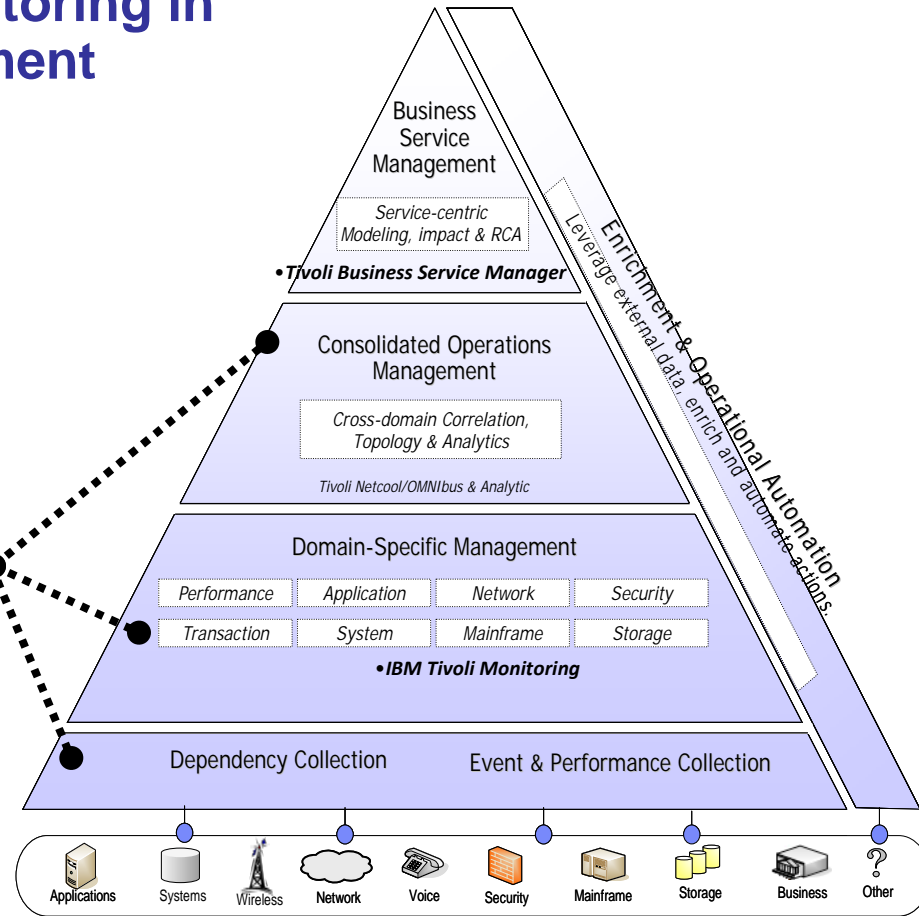
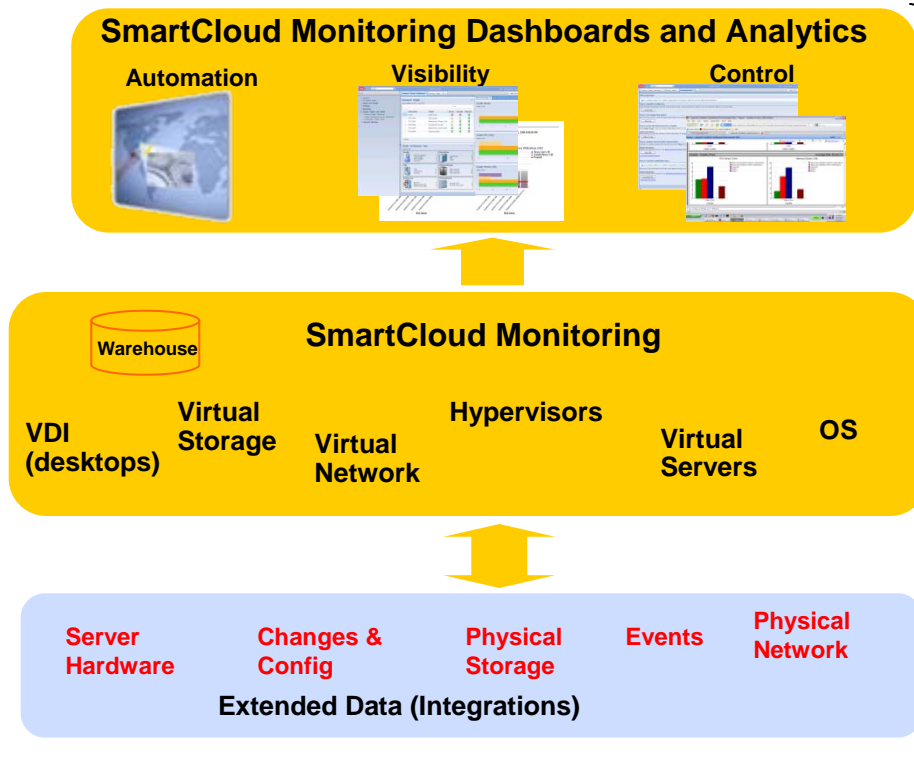
# Key Features for Virtual Environment Management

- **Dynamic thresholds:** Virtual environments by their very nature are dynamic. If a virtual machine is added to a host, a host added to a cluster or if virtual machines move from host to host, the baseline that was established is meaningless and needs to change with the new workload characteristics. In this example, ITM can reestablish thresholds to more accurately alert operators that there is a problem in the environment.
- **Predictive trending:** In addition to the dynamic nature of virtual environments, virtual environments can scale up or down at a moments notice. For this reason, it is critical to have a handle on when critical resources will near capacity in order to provision additional physical capacity with enough advance warning.
- **Capacity planning:** Customers need to right-size their virtual environment, balance workloads, and optimize or increase the density of their virtual environments.
- **Integrated physical and virtual resource monitoring:** We are finding from our customers that it is not enough to just manage the virtual server environment, but to have a handle on the impact of backing storage systems. It is for this reason we are providing out of the box integration of virtual and physical storage and network monitoring to isolate problems faster.
- **Multiple hypervisor support:** A large percentage (40-50%) of customers have deployed more than one hypervisor in their data center, and this percentage will likely grow as hypervisor technologies continue to mature. IBM Tivoli Monitoring provides broad hypervisor coverage to help reduce the cost of managing multiple hypervisors.
- **Guest and Application Monitoring:** Monitoring of the virtual environment in the context of the guest OS & application. Customers who want to place tier 1 and 2 applications onto virtual infrastructures require confidence of ensuring high availability for these applications and want to do so in an integrated fashion.
- **Storage and Network:** Storage and Network monitoring a critical element of Cloud and Virtualized Environments

# Design/Architecture

## Logical View

### Positioning of SmartCloud Monitoring in the Business Services Management



Note: All layers are inclusive of distributed and mainframe.

# Design/Architecture

## Logical View

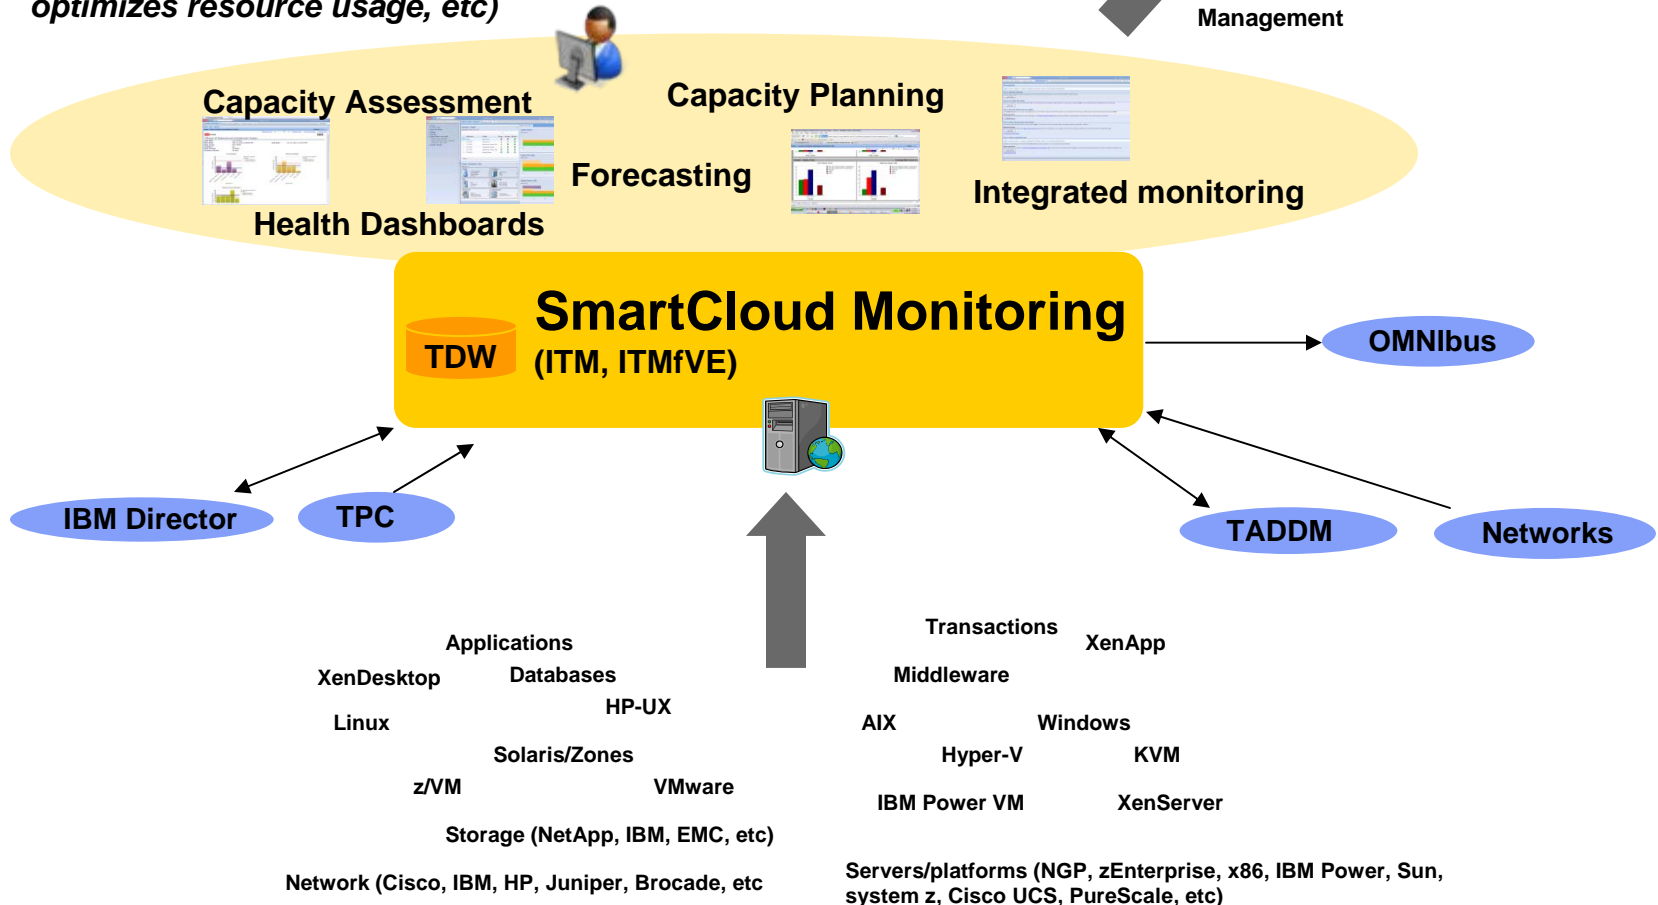
# SmartCloud Monitoring Overview

*Assures that health of cloud environment meets customer needs (reduces MTTR, lower operations cost, etc)*

*Helps to consolidate and reduce IT footprint (reduces TCO, optimizes resource usage, etc)*

Integrated Services Management **Tivoli**

Part of broader IBM Services Management

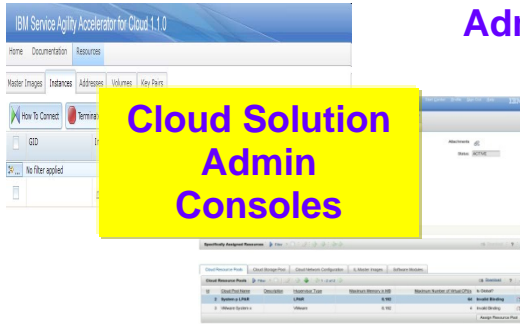


# Design/Architecture

## Operational View

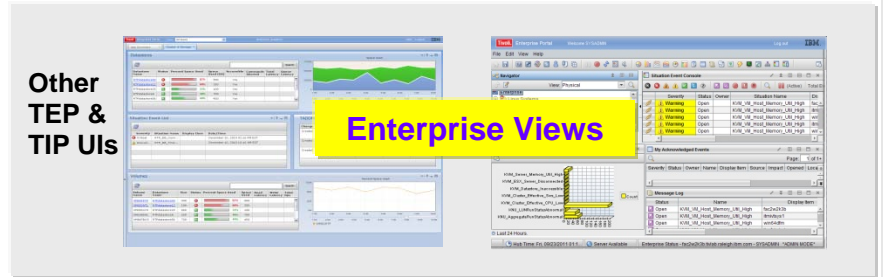
# Cloud Admin navigates amongst UIs for levels of Cloud health

Simple view of health of servers, services and components that make up the Cloud infrastructure



**Cloud Solution Admin Consoles**

Admin Views



Other TEP & TIP UIs

**Enterprise Views**

Launch to see overall cloud health

Drilldown to see details of physical and virtual resources in the cloud



Cloud Admin

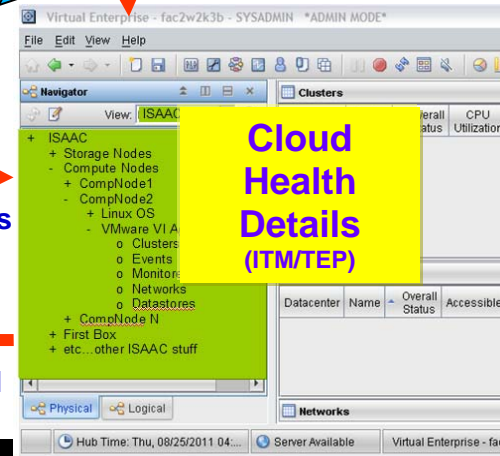
Overall health of cloud infrastructure and workloads, capacity info



**Cloud Health Dashboard (ITM/VE/TIP)**


Drilldown to see details of physical and virtual resources in the cloud

Launch to see overall cloud health



**Cloud Health Details (ITM/TEP)**

Detailed information about cloud components



Capacity Planner



IT Operator, App Owner

IBM Business Partner Summit 2012

Optimizing the World's Infrastructure  
March 4-7 Las Vegas, Nevada

# How do you get started?

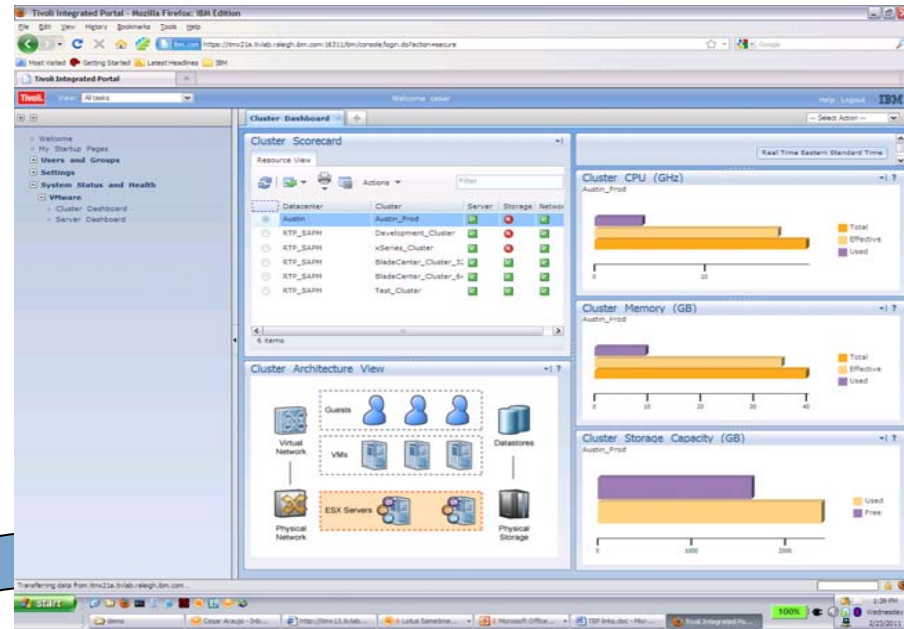
- SmartCloud Monitoring is an ITM based application
  - Deploy the included ITM infrastructure (management servers, data warehouse, agent depot, etc)  
*Customers with an existing ITM installation (v6.2.2 FP4 or later) can skip this step, and go right to deployment of the monitoring agents*
  - Deploy the ITMfVE 7.1 agents to monitor the cloud infrastructure and deploy the OS agents to the virtual machine workloads running there
  - Deploy TCR 2.1.1 server (unless one already exists)
  - Deploy the SmartCloud Monitoring Dashboard and Capacity Planning on the TCR server.
  - Optionally, install further integration with storage, Network Monitoring, OMNIbus, IBM Systems Director, TADDM, etc.




# Health Dashboard

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 – Storage, Network, ITM, TADDM & VMware
- Views with performance and capacity reports for assessment of environment and long term trend analysis



Top or Bottom N Host Servers by CPU and Memory Usage



# SmartCloud Monitoring Health Summary Dashboard: VMware Sample

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

VMware Cluster Dashboard + Select Action

### Scorecard Widget

Last Update: 10/3/11 12:15 PM







Filter

|                                  | Datacenter | Cluster                   | Server | Storage | Network |
|----------------------------------|------------|---------------------------|--------|---------|---------|
| <input checked="" type="radio"/> | Austin     | Austin_Prod               | ✘      | ✘       | ✔       |
| <input type="radio"/>            | RTP_SAPM   | Test_Cluster              | ⚠      | ⚠       | ✔       |
| <input type="radio"/>            | RTP_SAPM   | BladeCenter_Cluster_32bit | ⚠      | ✔       | ✔       |
| <input type="radio"/>            | RTP_SAPM   | Development_Cluster       | ⚠      | ✔       | ✔       |
| <input type="radio"/>            | RTP_SAPM   | BladeCenter_Cluster_64bit | ✔      | ✔       | ✔       |
| <input type="radio"/>            | RTP_SAPM   | xSeries_Cluster           | ✔      | ✔       | ✔       |

6 items

### Cluster Architecture View

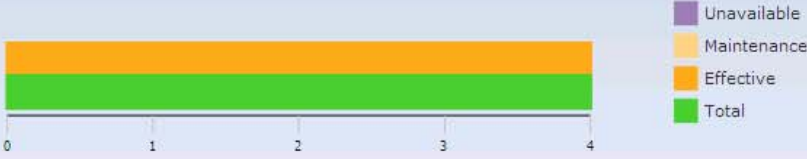
Austin\_Prod

|  |  |
|--|--|
| <h4>Guests</h4>  <ul style="list-style-type: none"> <li>Windows Guests: 7</li> <li>Linux Guests: 5</li> <li>Other Guests: 0</li> <li>Unknown: 5</li> </ul> | <h4>Data Stores</h4>  <ul style="list-style-type: none"> <li>Data Stores: 13</li> <li>NFS: 6</li> <li>VMFS: 7</li> </ul>                          |
| <h4>VMs</h4>  <ul style="list-style-type: none"> <li>VMs: 17</li> <li>Powered On: 17</li> <li>Running: 14</li> </ul>                                     | <h4>Physical Storage</h4>  <ul style="list-style-type: none"> <li>SAN Volumes: 0</li> <li>NAS Volumes: 24</li> <li>Total Volumes: 24</li> </ul> |
| <h4>ESX Servers</h4>  <ul style="list-style-type: none"> <li>Servers: 4</li> <li>Effective Servers: 4</li> <li>Maintenance Mode: 0</li> </ul>            | <h4>Virtual Network</h4>  <ul style="list-style-type: none"> <li>Physical NICs: 8</li> <li>Physical NICs Down: 0</li> </ul>                     |

### Cluster Servers

Austin\_Prod


Select time range



| Status      | Count |
|-------------|-------|
| Unavailable | 0     |
| Maintenance | 0     |
| Effective   | 4     |
| Total       | 4     |

### Cluster CPU (GHz)


Austin\_Prod



| Category  | Value |
|-----------|-------|
| Used      | ~15   |
| Effective | ~35   |
| Total     | ~40   |

### Cluster Memory (GB)

Austin\_Prod



| Category  | Value |
|-----------|-------|
| Allocated | ~15   |
| Used      | ~10   |
| Effective | ~35   |
| Total     | ~40   |

### Cluster Storage Capacity (GB)

Austin\_Prod

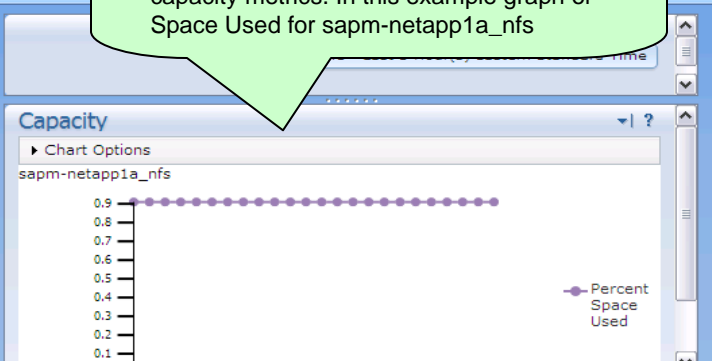
- Welcome
- My Startup Pages
- Users and Groups
- Settings
- System Status and Health

Cluster Dashboard Storage Details

### Datstores

Resource View

| Name              | Status | Used(GB) | Used Space(%) |
|-------------------|--------|----------|---------------|
| absmtast          | ✖      | 591.19   | 87%           |
| itmx33            | ✔      | 104.72   | 39%           |
| sapm-netapp2_nfs  | ✔      | 77.84    | 13%           |
| sapm-netapp1a_nfs | ✖      | 72.81    | 92%           |
| sapm-netapp1      | ✖      | 164.8    | 95%           |



Real-time or historical graph for selected sapm-netapp1a\_nfs datastore performance or capacity metrics. In this example graph of Space Used for sapm-netapp1a\_nfs

List of datastores of Austin\_Prod Cluster

### Situation Event List

Resource View

| Severity | Situation Name                   | Display Item | DateTime          |
|----------|----------------------------------|--------------|-------------------|
| CRITICAL | KVM_Datastore_Usage_ sapm-netapp |              | 02/23/11 09:39:50 |

List of storage problems (situations) associated with sapm-netapp1a\_nfs datastore.

### Change History


| Type            |                   |         |                         |
|-----------------|-------------------|---------|-------------------------|
| VMWareDataStore | sapm-netapp1a_nfs | Created | Febr 21, 2011 8:56 AM E |

List of changes associated with sapm-netapp1a\_nfs datastore

### Volumes

Resource View

| Name                | Size(GB) | Status | Used Space(%) | Used(GB) |
|---------------------|----------|--------|---------------|----------|
| sapm-netapp1a:/vol1 | 100.0    | 2.6    | 91%           | 72.8     |



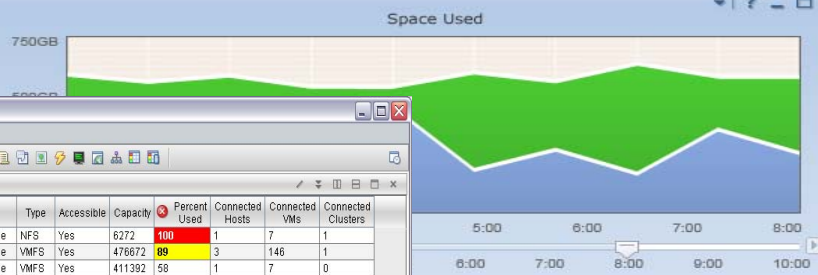
Real-time or historical graph for Netapp volume associated with sapm-netapp1a\_nfs datastore

Netapp NAS volume associated with sapm-netapp1a\_nfs datastore

App Summary Cluster A Storage

Datstores

| Datstore Name  | Status | Percent Space Used | Sp Us |
|----------------|--------|--------------------|-------|
| RTPdatastore00 | ✖      | 82%                |       |
| RTPdatastore12 | ✖      | 68%                |       |
| RTPdatastore25 | ✔      | 31%                |       |
| RTPdatastoreA  | ✔      | 40%                |       |
| RTPdatastore56 | ✔      | 49%                |       |



Datstores - fac2w2k3.raleigh.ibm.com - SYSADMIN \*ADMIN MODE\*

File Edit View Help

Navigator View: Physical

Datstore Health

| Datacenter | Name              | Overall Status | Type | Accessible | Capacity | Percent Used | Connected Hosts | Connected VMs | Connected Clusters |
|------------|-------------------|----------------|------|------------|----------|--------------|-----------------|---------------|--------------------|
| B-510      | test_nfs          | Unavailable    | NFS  | Yes        | 6272     | 100          | 1               | 7             | 1                  |
| B-510      | ISCSI Disk 2      | Unavailable    | VMFS | Yes        | 476672   | 89           | 3               | 146           | 1                  |
| B-510      | storage1 (1)      | Unavailable    | VMFS | Yes        | 411392   | 58           | 1               | 7             | 0                  |
| B-510      | ISCSI Disk 1      | Unavailable    | VMFS | Yes        | 476672   | 56           | 3               | 22            | 1                  |
| B-510      | itm64vm6 storage1 | Unavailable    | VMFS | Yes        | 26880    | 46           | 1               | 2             | 0                  |
| B-510      | storage1          | Unavailable    | VMFS | Yes        | 34560    | 42           | 1               | 2             | 1                  |
| B-510      | joe_1             | Unavailable    | VMFS | Yes        | 34048    | 2            | 1               | 0             | 1                  |
| B-510      | storage1 (2)      | Unavailable    | VMFS | Yes        | 31488    | 2            | 1               | 0             | 1                  |

Launch in context to ITM VMware VI agent datastore VMwares for additional details and problem resolution

Situation Event List

| Severity   | Situation Name | Display Item |
|------------|----------------|--------------|
| Critical   | RTP_DB_Conn... |              |
| Unavail... | RTP_DB_Time... |              |

NAS Datstores

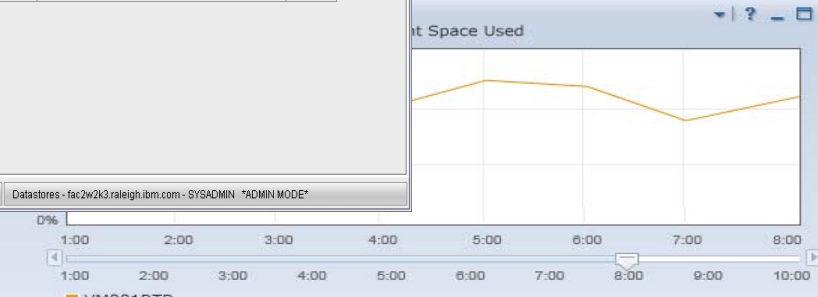
| Datacenter | Name     | Type | Remote Host Address      | Percent Used |
|------------|----------|------|--------------------------|--------------|
| B-510      | test_nfs | NFS  | fac2w2k3.raleigh.ibm.com | 100          |

VMFS Datstores


| Datacenter | Name              | URL   | Percent Used |
|------------|-------------------|---|--------------|
| B-510      | ISCSI Disk 2      | sanfs://mfs_uid46715931-3bd36809-3113-000e0c4...  | 89           |
| B-510      | storage1 (1)      | sanfs://mfs_uid4706576-2c9ab9ca-9a15-001a6465...  | 58           |
| B-510      | ISCSI Disk 1      | sanfs://mfs_uid46715901-25a31b06-3c8b-000e0c4...  | 56           |
| B-510      | itm64vm6 storage1 | sanfs://mfs_uid494r6feb-a1add98c-fa05-00096ba3... | 46           |
| B-510      | storage1          | sanfs://mfs_uid46642238-8d88a1c-51cc-000e0c42...  | 42           |
| B-510      | joe_1             | sanfs://mfs_uid494fe4cb-aal23aa9-f77b-000e0c42... | 2            |
| B-510      | storage1 (2)      | sanfs://mfs_uid467a71e5-c775312-e00e-000e0c4...   | 2            |

Volumes

| Volume Name | Datstore Name  | Size | Status | Perce |
|-------------|----------------|------|--------|-------|
| VM001RTP    | RTPdatastore00 | 400  | ✖      |       |
| VM002ATL    | RTPdatastore12 | 100  | ✖      |       |
| VM056CHI    | RTPdatastore25 | 800  | ✔      | 31%   |
| VM298DAL    | RTPdatastoreA  | 200  | ✔      | 40%   |
| VM887BOS    | RTPdatastore56 | 700  | ✔      | 49%   |



Could also launch in context to Storage Productivity Center Console



# Monitoring the Cloud/Virtual Environments with IBM SmartCloud Monitoring - Why is it important?

- Provides availability, performance and capacity management of hypervisors, OS and applications in **ONE** tool
- Enables fast problem isolation with integrated monitoring and dashboards across physical & virtual environment bringing together and in context hypervisor, OS, server, network and storage health information
- Advanced monitoring increases resources availability with proactive and predictive monitoring as well as baseline and dynamic thresholds which enables management policies that adapt to business needs
- New capacity management tools provide assessment of capacity usage and planning for optimization, consolidation and workload placement, enabling better resource usage, reducing costs, footprint and licenses
- Reduces labor costs with lower MTTR and more efficiency in the operations with OOTB knowledge base best practices (situations, workflows, expert advices, etc)
- Provides broad coverage with heterogeneous hypervisor and platform support
- Well-integrated with Business Services Management



# Capacity Planning is **CRITICAL** to Cloud management!

- **Helps consolidate and reduce costs**

- Reduces HW and labor costs
- Reduces number of physical servers required to run workloads
- Reduce number of required licenses

- **Helps ensure application availability**

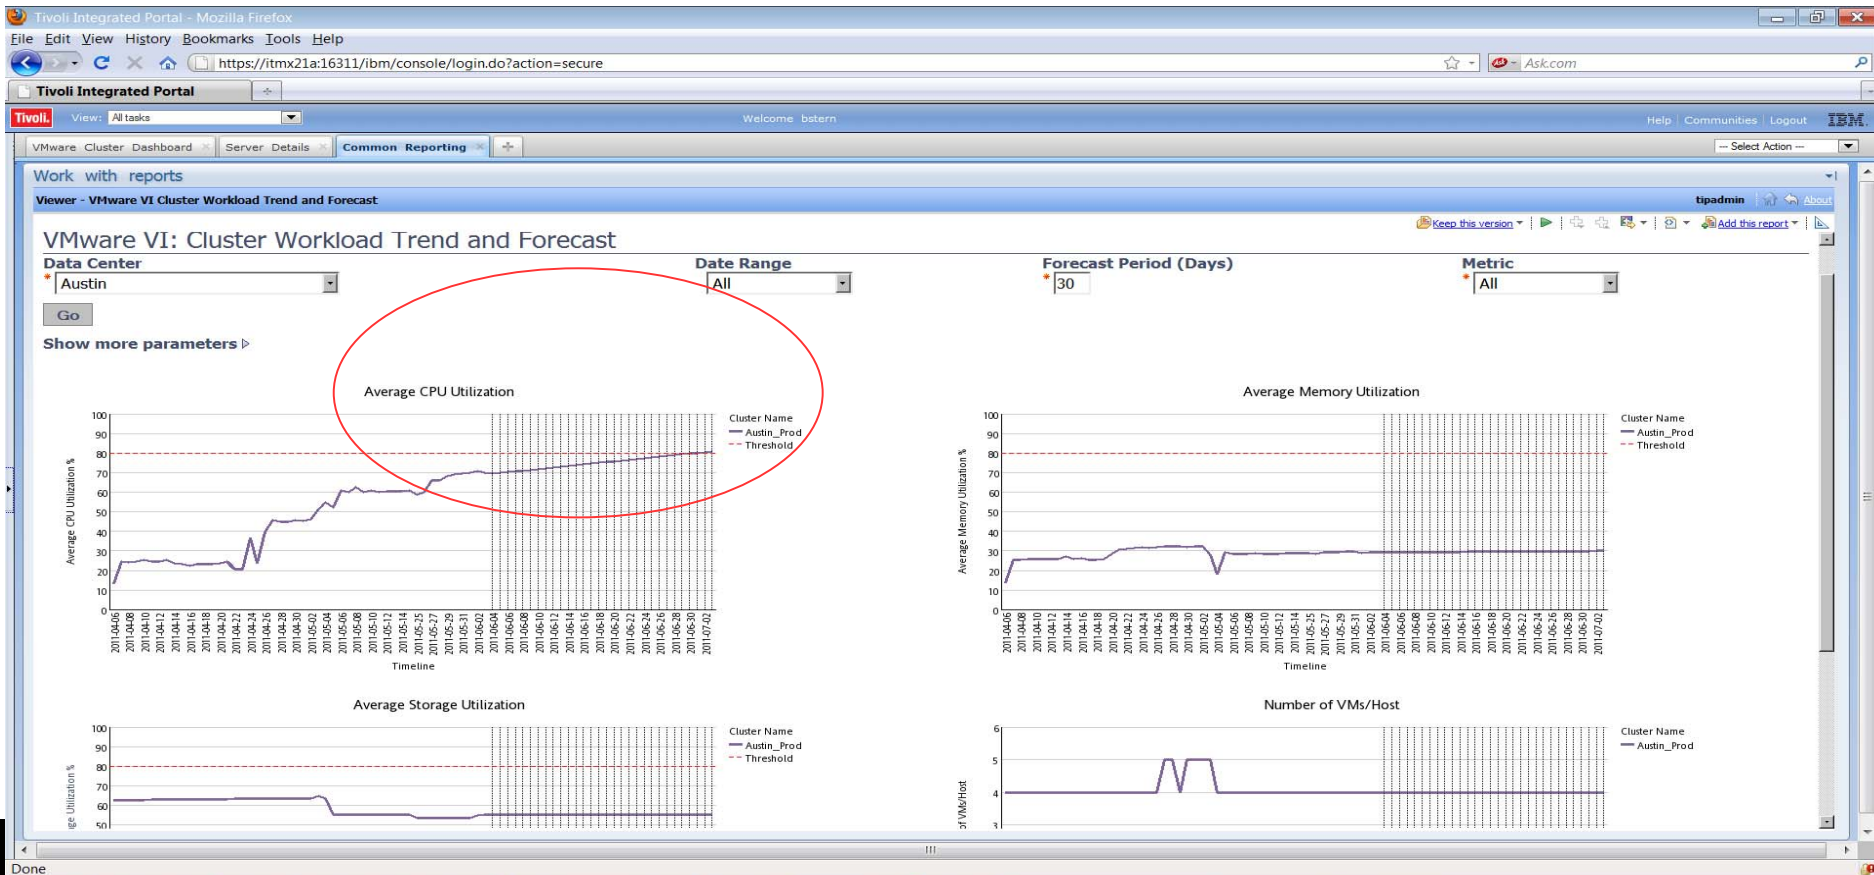
- Are any resources overloaded? When will physical resources reach their limits?
- Have there been any significant changes in my environment between two weeks?
- Ensure supply can meet demand
- Ensure business policies are met

- **Helps optimize resource utilization**

- Right size virtual machines
- Identify trends for workload balancing

# Scenario Summary

- An IT Admin may want to balance workloads to avoid a capacity bottleneck (Scenario 2).
- A Capacity Analyst may want to plan for an IT infrastructure that anticipates business growth.
- For this scenario assume --
  - The IT Dept has projected a 10% CPU growth for the applications in Austin\_Prod cluster over the next 1 month.
  - The goal is to prepare for a 20% growth in computing resources to be safe.



# Launch PlanningCenter

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

PlanningCenter [ + ] [ - Select Action ]

## PlanningCenter

**Step 1: Snapshot Config Data**

Load Config

Load the latest configuration data for physical servers and virtual machines for analysis. You can do some what-if analysis changes on this data.

**Step 2: Set Analysis time period**

Set Time

Set the time period for which the measurement data in the warehouse can be analyzed corresponding to the Virtual Machines loaded in **Step 1**. The measurement data is federated from the warehouse.

Beta note: Please select Set Time Range option in Timeselctor for setting time.

**Step 3: Scope the Infrastructure for Analysis**

Define Scope

Click the Scope button to go to an expert mode page where you can select the subset of physical servers that should be part of this analysis. Default scope includes all physical servers loaded in **Step 1**.

**Advanced options:**  
While scoping you can also edit the current configuration in [Current Environment > Edit](#) page to add new attributes or clean the data for physical servers and virtual machines, if required.

**Step 4: Size Virtual Machines**

Size VMs

[Current Topology](#)

Analyze the warehouse data within the time limit set in **Step 2**, to compute the VM level sizing estimates using default settings (average daily utilization).

**Advanced options:**  
Experts can go do custom sizing in [Current Environment > Edit](#) page. While on this edit page, several custom actions can be triggered to characterize virtual machines based on measurement data.

**Step 5: Generate Optimization Plan**

Generate Plan

[Recommendation Topology](#)

[Recommendation Report](#)

Generate a recommended environment using default settings.

**Advanced options:**  
Custom settings can be made in [Recommended Environment > Edit Settings](#) page where experts can select optimization strategies, such as applicable business and technical policies, optimization goal etc.

Any change in steps 1 to 3 needs a re-generation of the plan in step 4 to view the latest recommendation.

5-Step planning process supporting novice and expert users



# Select the overall data set we want to analyze

1. Load the latest configuration data for what-if

The screenshot shows the PlanningCenter interface with two main steps. Step 1, titled "Step 1: Snapshot Config Data", contains a "Load Config" button. Step 2, titled "Step 2: Set Analysis time period", contains a "Set Time" button. A red arrow points from the first step's instruction to the "Load Config" button. A red box highlights the "Set Time" button, with a red line extending from it to the "Time Selector" dialog box shown below.

**Step 1: Snapshot Config Data**

Load Config

Load the latest configuration data for physical servers and virtual machines for analysis. You can do some what-if analysis changes on this data.

**Step 2: Set Analysis time period**

Set Time

Set the time period for which the measurement data in the warehouse can be analyzed corresponding to the Virtual Machines loaded in **Step 1**. The measurement data is federated from the warehouse.

2. Set the time period to analyze measurement data in Warehouse

The "Time Selector" dialog box is open, showing the "Set Time Range" section. It includes fields for Start Date (3/6/2011), Start Time (9:56 PM), End Date (6/7/2011), End Time (9:56 PM), and Time zone (Local). Below these fields is a timeline slider with markers for 9:56PM on 10/6/2010, 11/6/2012, 12/6/2014, 1/6/2017, and 2/6/2019. The dialog has "OK" and "Cancel" buttons at the bottom right.

Time Selector

Set Time Range

| Start    |         | End      |         | Time zone |
|----------|---------|----------|---------|-----------|
| Date     | Time    | Date     | Time    |           |
| 3/6/2011 | 9:56 PM | 6/7/2011 | 9:56 PM | Local     |

9:56PM 10/6/2010    9:56PM 11/6/2012    9:56PM 12/6/2014    9:56PM 1/6/2017    9:56PM 2/6/2019

OK Cancel

### Step 3: Scope the Infrastructure for Analysis

Define Scope



Click the Scope button to go to an expert mode page where you can select the subset of physical servers that should be part of this analysis. Default scope includes all physical servers loaded in **Step 1**.

#### Advanced options:

While scoping you can also edit the current configuration in [Current Environment > Edit](#) page to add new attributes or clean the data for physical servers and virtual machines, if required.

1. Click opens a new tab

3. Optionally add new attributes to be used in analysis, e.g. for defining policies

2. Choose servers in Austin\_Prod to analyze

Create custom column views. Sort, filter by any attribute. Edit/clean data.

#### Edit Current Environment

View: Inventory >> Physical Servers

| Select Server                       | Status | Data Center | Cluster                   | Host Name*                        | Line Of Business | Architecture*   | Number Of CPU Cores* | CPU Speed [MHz]* | Memory |
|-------------------------------------|--------|-------------|---------------------------|-----------------------------------|------------------|---|----------------------|------------------|--------|
| <input checked="" type="checkbox"/> |        | Austin      | Austin_Prod               | absm-365b.tivlab.raleigh.ibm.com  | Billing Service  | Intel(R) Xeon(TM) MP CPU 3.00GHz                        | 4                    | 3000             |        |
| <input checked="" type="checkbox"/> |        | Austin      | Austin_Prod               | benblade06.tivlab.raleigh.ibm.com | Billing Service  | Intel(R) Xeon(TM) CPU 2.00GHz                           | 1                    | 1999             |        |
| <input checked="" type="checkbox"/> |        | Austin      | Austin_Prod               | benblade07.tivlab.raleigh.ibm.com | Billing Service  | Intel(R) Xeon(TM) CPU 2.00GHz                           | 1                    | 1999             |        |
| <input checked="" type="checkbox"/> |        | Austin      | Austin_Prod               | benblade08.tivlab.raleigh.ibm.com | Billing Service  | Intel(R) Xeon(TM) CPU 3.20GHz                           | 1                    | 3199             |        |
| <input type="checkbox"/>            |        | RTP_SAPM    | xSeries_Cluster           | absm-255a.tivlab.raleigh.ibm.com  |                  | Intel(R) Xeon(TM) MP CPU 2.50GHz                        | 8                    | 2485             |        |
| <input type="checkbox"/>            |        | RTP_SAPM    | xSeries_Cluster           | absm-365a.tivlab.raleigh.ibm.com  |                  | Intel(R) Xeon(TM) MP CPU 3.00GHz                        | 4                    | 3000             |        |
| <input type="checkbox"/>            |        | RTP_SAPM    | Test_Cluster              | absm-366a.tivlab.raleigh.ibm.com  |                  | Intel(R) Xeon(TM) MP CPU 3.60GHz                        | 8                    | 3669             |        |
| <input type="checkbox"/>            |        | RTP_SAPM    | Development_Cluster       | benblade02.tivlab.raleigh.ibm.com |                  | Intel(R) Xeon(TM) CPU 2.00GHz                           | 1                    | 1999             |        |
| <input type="checkbox"/>            |        | RTP_SAPM    | BladeCenter_Cluster_32bit | benblade03.tivlab.raleigh.ibm.com |                  | Intel(R) Xeon(TM) CPU 2.00GHz                           | 1                    | 1999             |        |
| <input type="checkbox"/>            |        | RTP_SAPM    | BladeCenter_Cluster_32bit | benblade04.tivlab.raleigh.ibm.com |                  | Intel(R) Xeon(TM) CPU 2.00GHz                           | 1                    | 1999             |        |
| <input type="checkbox"/>            |        | RTP_SAPM    | BladeCenter_Cluster_32bit | benblade05.tivlab.raleigh.ibm.com |                  | Intel(R) Xeon(TM) CPU 2.00GHz                           | 1                    | 1999             |        |
| <input type="checkbox"/>            |        | RTP_SAPM    | Development_Cluster       | benblade09.tivlab.raleigh.ibm.com |                  | Intel(R) Xeon(TM) CPU 3.20GHz                           | 1                    | 3199             |        |
| <input type="checkbox"/>            |        | RTP_SAPM    | BladeCenter_Cluster_64bit | benblade11.tivlab.raleigh.ibm.com |                  | Advanced Micro Devices Thunder K8QW (S4881) Opteron 890 | 2                    | 2591             |        |
| <input type="checkbox"/>            |        | RTP_SAPM    | BladeCenter_Cluster_64bit | benblade12.tivlab.raleigh.ibm.com |                  | Advanced Micro Devices Tyan YR190B8228 AMD Opteron 4122 | 8                    | 2392             |        |

1 - 15 of 15 items

5 | 10 | 25 | 50 | 100 | All

1

| VM CPU [MHz]* | Minimum Memory [MB]* | VMID                                  | Operating System   | Operating System Version |
|---------------|----------------------|---------------------------------------|--|--------------------------|
| 0             | 0                    | 564DB7A1-B41B-C28D-0BD1-A85D86F4AF A1 | Microsoft Windows Server 2003, Enterprise Edition (32-bit) |                          |
| 0             | 0                    | 423A2C35-1CA2-CD7D-3C78-79634143C3D6  | Microsoft Windows Server 2003, Enterprise Edition (32-bit) |                          |

Detailed view can also launch context sensitive reports from the menu in a new tab

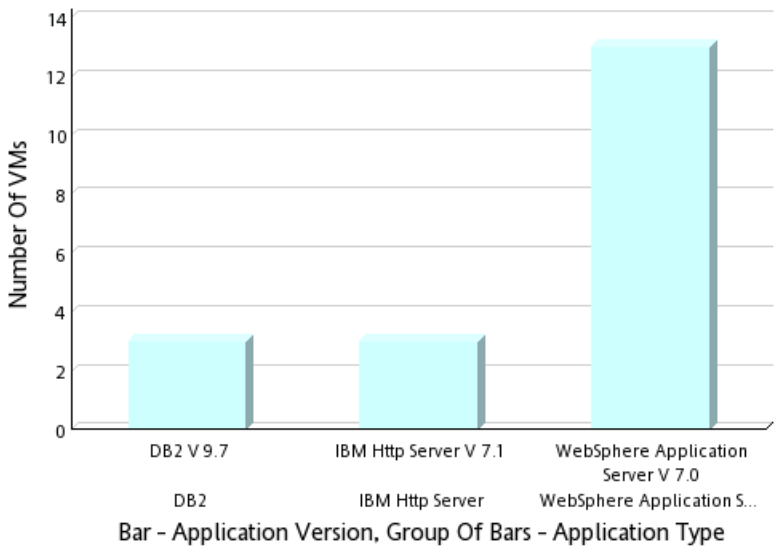
### Virtual Machine Inventory

Report As Of : Jun 16, 2011 12:53:03 PM

### Virtual Machine Summary

|                                 |    |
|---------------------------------|----|
| Number of Virtual Machines      | 19 |
| Number of Virtual CPUs          | 21 |
| Current CPU Reservation (GHz)   | 0  |
| Current Memory Reservation (GB) | 0  |

### VM Distribution: By Application Type And Version




Types of middleware on the 19 VMs on 4 hosts in Austin\_Prod

## Step 4: Size Virtual Machines

Size VMs

[Current Topology](#)



Click opens  
"interactive" view  
in new tab





Analyze the warehouse data within the time limit set in **Step 2**, to compute the VM level sizing estimates using default settings (average daily utilization).

### Advanced options:



Experts can go do custom sizing in [Current Environment > Edit](#) page. While on this edit page, several custom actions can be triggered to characterize virtual machines based on measurement data.

## CURRENT VIRTUAL ENVIRONMENT

### Data Centers



### Clusters of - Austin





### Physical Servers of - Austin\_Prod




Summary view of selection



Host at risk as it has no free capacity based on usage



### Virtual Machines of - absm-365b.tivlab.raleigh.ibm.com



VM not at risk as typical CPU usage < current reservation



## Step 5: Generate Optimization Plan

Generate Plan

[Recommendation Topology](#)

[Recommendation Report](#)



Generate a recommended environment using default settings.

### Advanced options:

Custom settings can be made in [Recommended Environment > Edit Settings](#) page where experts can select optimization strategies, such as applicable business and technical policies, optimization goal etc.



Any change in steps 1 to 3 needs a re-generation of the plan in step 4 to view the latest recommendation.

### Recommendation Generation

Step 1: Select rules to apply in optimization

#### Colocation/Anti-colocation

Active

#### Rule Instances

▶ Do not colocate VMs with DB2 and WAS

▶ Seperate High and Low Critical VMs



#### Boundary

Active

#### Rule Instances

▶ Create a Boundary for Critical VMs

▶ Create a Boundary for Win2003 32-bit VMs



#### Utilization

Active

#### Rule Instances

▶ Provide 50% more CPU for Critical VMs

▶ Use 20% for growth on MQ servers



Step 2: Select optimization goal

Minimize systems

Balance server utilization

Step 3: Select options

Keep existing recommendations

Ignore existing recommendations

Choose rules  
for what-if  
scenario



Click opens settings page  
in new tab – enables edits  
on optimization strategies.  
Will also link to a Rule Editor.



Step 5: Generate Optimization Plan

Generate Plan

[Recommendation Topology](#)

[Recommendation Report](#)

Click opens plan in new tab



Generate a recommended environment using default settings.

**Advanced options:**

Custom settings can be made in [Recommended Environment > Edit Settings](#) page where experts can select optimization strategies, such as applicable business and technical policies, optimization goal etc.



Any change in steps 1 to 3 needs a re-generation of the plan in step 4 to view the latest recommendation.

Tivoli software



Recommended Optimized Environment

Report As Of : Jun 17, 2011 12:16:42 AM


|                          | Current   |             | Recommendation |             |
|--------------------------|-----------|-------------|----------------|-------------|
|                          | CPU (GHz) | Memory (GB) | CPU (GHz)      | Memory (GB) |
| Physical Servers         | 4         |             | 3              |             |
| Virtual Machines         | 19        |             | 19             |             |
| Total Capacity           | 18.747    | 39.994      | 15.623         | 31.995      |
| Total Reservation        | 0.375     | 0.438       | 7.594          | 11.324      |
| Total Spare              | 8.999     | 35.557      | 0.218          | 17.471      |
| Average Overall Risk (%) |           |             |                |             |

Accommodated growth and consolidated further

Detailed Placement

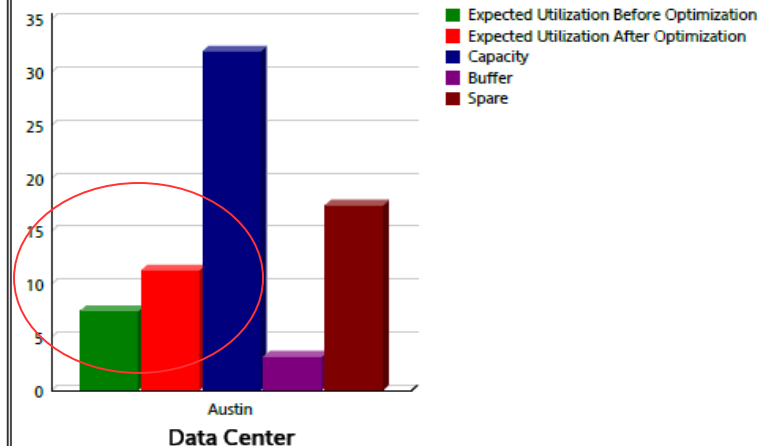
Data Center : Austin

CPU Details (GHz)



Growth in Austin\_Prod

Memory Details (GB)



Generate Plan

[Recommendation Topology](#)

[Recommendation Report](#)

Click opens topology view in new tab



Generate a recommended environment using default settings.

**Advanced options:**



Custom settings can be made in [Recommended Environment > Edit Settings](#) page where experts can select optimization strategies, such as applicable business and technical policies, optimization goal etc.





Any change in steps 1 to 3 needs a re-generation of the plan in step 4 to view the latest recommendation.

**RECOMMENDED VIRTUAL ENVIRONMENT**

**Data Centers**





**Clusters of - Austin**





**Physical Servers of - Austin\_Prod**



All hosts and VMs risk-free after optimization



**Virtual Machines of - absm-365b.tivlab.raleigh.ibm.com**



Consolidated 14 servers to 6

Tivoli software

Recommended Optimized Environment

Environment Summary

|                      | Current   |             | Recommendation       |             |
|----------------------|-----------|-------------|----------------------|-------------|
| Physical Servers     | 14        |             | 6                    |             |
| Virtual Machines     | 52        |             | 45                   |             |
|                      | CPU (GHz) | Memory (GB) | CPU (GHz)            | Memory (GB) |
| Total Capacity       | 47,000    | 139,795     | 33,805               | 47,988      |
| Total Reservation    | 0,000     | 0,000       | 21,804               | 46,247      |
| Total Spare          | 47,000    | 139,795     | 11,900               | 1,742       |
| Average Overall Risk |           |             |                      |             |
|                      |           |             | Average Overall Risk |             |
|                      |           |             | Average Risk per VM  |             |

Virtual Machines Not Allocated To Physical Server

| VM Name          | Current Reservation CPU (GHz) | Current Reservation Memory (GB) | Recommended Reservation CPU (GHz) | Recommended Reservation Memory (GB) | Expected Storage Usage (Kbps) | Expected Network Usage (Kbps) | Operating System                       | Application |
|------------------|-------------------------------|---------------------------------|-----------------------------------|-------------------------------------|-------------------------------|-------------------------------|--|-------------|
| Temporary        | 0.000                         | 0.000                           |                                   |                                     |                               |                               | Unavailable                            |             |
| Test_Scenario_VM | 0.000                         | 0.000                           |                                   |                                     |                               |                               | Unavailable                            |             |
| W2K32            | 0.000                         | 0.000                           |                                   |                                     |                               |                               | Unavailable                            |             |
| W2K33            | 0.000                         | 0.000                           |                                   |                                     |                               |                               | Unavailable                            |             |
| W2K8-84-01       | 0.000                         | 0.000                           |                                   |                                     |                               |                               | Microsoft Windows Server 2008 (84-bit) |             |
| W2K80            | 0.000                         | 0.000                           |                                   |                                     |                               |                               | Unavailable                            |             |
| W2K82            | 0.000                         | 0.000                           |                                   |                                     |                               |                               | Unavailable                            |             |

Detailed Placement

| Physical Server Summary | VM Name | # vCPUs | Recommended CPU Reservation (GHz) | Recommended Memory Reservation (GB) | Expected Storage Usage (Kbps) | Expected Network Usage (Kbps) | Operating System Name | Application |
|-------------------------|---------|---------|-----------------------------------|-------------------------------------|-------------------------------|-------------------------------|-----------------------|-------------|
|                         |         |         |                                   |                                     |                               |                               |                       |             |

Top Page up Page down Bottom

Done



# Why is Capacity Management Important?

- **Helps ensure application availability**
  - Are any resources overloaded? When will physical resources reach their limits?
  - Have there been any significant changes in my environment between two weeks?
- **Helps optimize resource utilization**
  - Right size virtual machines
  - Identify trends for workload balancing
- **Helps achieve the promise of resource availability**
  - Ensure supply can meet demand – pipeline management





# Your capacity management questions:

- System/Workload Characteristics, Performance and Trending
  - How is my environment performing overall?
  - Which are my most used servers for a given resource type?
  - Are there any bottlenecks in my current environment and where?
  - Am I reaching capacity on resources and which resource? When will I exhaust capacity?
  - Which is my top VM resource consumers for a given resource type?
  - Which are my least used servers for a given resource type?
  - Which are my bottom VM resource consumers for a given resource type?
  - Do I have any outstanding abnormal behavior this week compared to last week (other periods can be used)?
  - Are my systems/workloads balanced or unbalanced?
- System/Workload Estimation and Optimization (optimize and keep optimized – what if)
  - How many more VMs can I add to a cluster/server based on usage history?
  - How much more resources do I need to add additional VMs to environment?
  - How, where do I add capacity if existing systems are not enough for future growth for optimized capacity usage?
  - Where do I place new workloads? Do I really need to add more resources?
  - How can I optimize the VM placement to maximize usage and minimize costs?
  - How can I optimize the app placement to maximize usage and minimize costs?

# When will physical resources reach the breaking point?

**Tivoli software** **IBM**

## VMware VI: Host Server Forecast Alerts

Data Center: Bid-510 | Date Range: Last 30 days | Forecast Period (Days): 30 | Metric: All

Go

Show more parameters ▾

### Highest CPU Utilization Trend Alert for Host Servers in Forecasted Period

| Data Center | Cluster Name | Server Hostname                | Average CPU Utilization |
|-------------|--------------|--------------------------------|-------------------------|
| Bid-510     | Cluster A    | tn64vm1.tivoli.raleigh.ibm.com | 3.39                    |
| Bid-510     | Cluster A    | tn64vm2.tivoli.raleigh.ibm.com | 1.64                    |

### Highest Memory Utilization Trend Alert for Host Servers in Forecasted Period

| Data Center | Cluster Name | Server Hostname                | Average Memory Utilization |
|-------------|--------------|--------------------------------|----------------------------|
| Bid-510     | Cluster A    | tn64vm1.tivoli.raleigh.ibm.com | 29.67                      |
| Bid-510     | Cluster A    | tn64vm2.tivoli.raleigh.ibm.com | 44.61                      |

### Highest Storage Utilization Trend Alert for Host Servers in Forecasted Period

| Data Center | Cluster Name | Server Hostname                | Average Datastore Utilization |
|-------------|--------------|--------------------------------|-------------------------------|
| Bid-510     | Cluster A    | tn64vm1.tivoli.raleigh.ibm.com | 80.93                         |
| Bid-510     | Cluster A    | tn64vm2.tivoli.raleigh.ibm.com | 85.98                         |

This report alerts the user on when a server or group of selected servers will reach capacity limitations. It calculates a linear trend for the next 30 days and determines if any of the servers will exceed the user defined threshold for the server's CPU, Memory and Storage.

Jan 6, 2011 | 1 | 5:15:06 PM

Top | Page up | Page down | Bottom

Done | Local intranet | 100%

# How many more workloads can I place in this cluster or host? What-if I need to add X VMs of X size?

Data Center

Bld-510

Clusters

Cluster A

Go

Show more parameters

### RESOURCES NEEDED FOR ADDITIONAL WORKLOADS - AVERAGE DEPLOYED VM PROFILE

| Resource                   | VM Profile based on average resource used by all VMs on this cluster | Resources Needed by Additional VMs | Available Cluster Capacity (before applying Buffer) | Buffer | Available Cluster Capacity (after applying Buffer) | Capacity Needed |
|----------------------------|--|------------------------------------|---|--------|--|-----------------|
| CPU (GHz)                  | 0.235  | 1.177                              | 44.96   | 2      | 42.96  | 0               |
| Datastore Space Usage (GB) | 0  | 0                                  | 211.759   | 5      | 206.759  | 0               |
| Memory Usage (MB)          | 1,221,087  | 6,105,433                          | 9,020,455   | 256    | 8,764,455  | 0               |

12 : 00 AM

End Date

Earliest date

Dec 21, 2010

11 : 59 PM

Latest date

9



# Demo Dashboard and Capacity Planning



# Questions?

**Pulse2012**  
Optimizing the World's Infrastructure