

Managing Traditional Workloads Together with Cloud Computing Workloads

Teraproc

Table of Contents

- Introduction 3
- Cloud Management Challenges 3
- Re-thinking of Cloud Management Solution 4
- Teraproc Cloud Manager 4
 - Teraproc Cloud Manager unique capabilities 5
- Teraproc Cloud Manager Capabilities 6
 - Resource Management 6
 - Resource Configuration 6
 - Resource Monitoring 7
 - Resource Reports 8
- Projects 8
 - Project Based Operations 8
 - Project Reports 9
- Instance & Service Management 9
 - Instances 9
 - User Request Management 10
 - Service Templates 10
 - Service Catalog 10
- Application Template Management 11
- User & System Management 11
 - Users 11
 - User Roles 11
 - System Parameters 12
- Alarms & Events 12
- Self-Service Portal 12
- Typical Scenarios for Teraproc Cloud Manager 12
 - Scenario 1: Infrastructure with applications running on physical servers 12
 - Scenario 2: Software development & testing environment 13
 - Scenario 3: Internal application hosting 13
- Summary 13

Introduction

In today's business with competition on the market, IT infrastructure needs to be more agile to meet the fast changing business needs, more reliable to deliver near zero down time, and more cost effective. Cloud computing provides low cost flexible alternatives for enterprises to run their IT infrastructure. Although public cloud is now getting pervasive and mature, it is still cheaper to have infrastructure inside a firewall to:

- Lower the OpEx for workloads running all year long
- Ensure data security
- Avoid transferring large amount of data between the public cloud and the on-premises systems

The private cloud solutions and their practice mimic the user experiences of public cloud such as AWS (Amazon Web Services), automate common IT tasks, and leverage modern technologies such as virtualization and industry standard hardware for cost effectiveness.

Private cloud computing technology and its practice are critical in certain business areas including, but not limited to:

- Shared IT infrastructure supporting multiple different applications, some running in a virtualized environment and others running on bare metal.
- Big data analytics: The infrastructure is dynamic and needs to grow when new data comes in on a daily basis. Analysis job run by multiple user groups often requires security to fence the data access.
- Application development: The development and test environments change on a daily or even an hourly basis.
- High performance computing: An infrastructure shared by many research or engineering projects that requires running applications on bare metal for performance reason rather on virtual machines.

Cloud Management Challenges

Organizations that want to build and implement a cloud solution are still facing some challenges.

The first hurdle an organization needs to overcome is to convert the existing workload from traditional managed environment into a cloud managed environment. Majority of cloud infrastructure management solutions assume users have already virtualized their environment, i.e. all workloads are running on virtual machines. If existing workloads are running on traditional UNIX servers, or on x86 servers but with dependencies on legacy network equipment or storage systems, these workloads are difficult to be migrated onto x86 virtual machines. Some workloads require access to the bare metal devices (e.g. InfiniBand) for performance reason. Examples are: high performance computing applications, real time transactional applications, and high performance analytics. It would be difficult for these workloads to be smoothly migrated to a cloud computing environment that relies on virtualization technology.

The second challenge for an organization to deploy cloud computing is that IT management process and practice change. Such change can not happen overnight. Such change can not be done by just deploying an open source based cloud management solution that is designed only for virtualized workloads.

Thirdly, cloud computing built on top of virtualization technology requires an organization to acquire technology expertise in multiple disciplines as well as integrations of such multiple technology layers. It includes virtualization technology, dynamic machine provisioning, cloud management, etc. Acquiring the expertise and integrating technologies could cost an organization money and time, which could diminish the return on investment (ROI) for deploying a cloud computing solution.

These challenges prevent organizations moving forward to deploying a cloud computing solution for their critical business applications. With the ever pressing needs from end users demanding highly nimble and flexible IT infrastructure, as well as the pressure from the company management on IT budget, CIOs need to rethink of a cloud computing strategy for company's complex application environments.

Rethinking of Cloud Management Solution

To accelerate return on investment for a cloud computing solution, it is highly demanded that the new IT management solution can manage both traditional workloads on non-virtualized environment, as well as the emerging workloads that fully leverage the elasticity of the modern cloud computing infrastructure. Specifically, the solution needs to be able to manage resource pools that are virtualized, resource pools without using virtualization technology, as well as resource pools that consists of traditional UNIX servers.

With such a solution, different applications can be migrated from a non-cloud environment to a cloud environment in its own pace, while the entire environment is centralized managed with a single management console.

The management solution needs to be able to manage multiple resource pools across different geographical locations. It should be compatible with commonly used open source cloud infrastructure management tools like OpenStack, which benefits from active development by multiple vendors.

Teraproc Cloud Manager

Teraproc Cloud Manager is a complete solution for managing heterogeneous cloud computing environments. Whether you have traditional workloads, high-performance computing workloads, or cloud workloads, Teraproc Cloud Manager helps you manage all your workloads in one cloud environment. It leverages commonly used cloud infrastructure management solutions like OpenStack, CloudStack, and Elaster as the foundation, and extends the capability to manage resources, services, applications, tenants, and projects across multiple data centers.

Figure 1 illustrates functional blocks of Teraproc Cloud Manager.

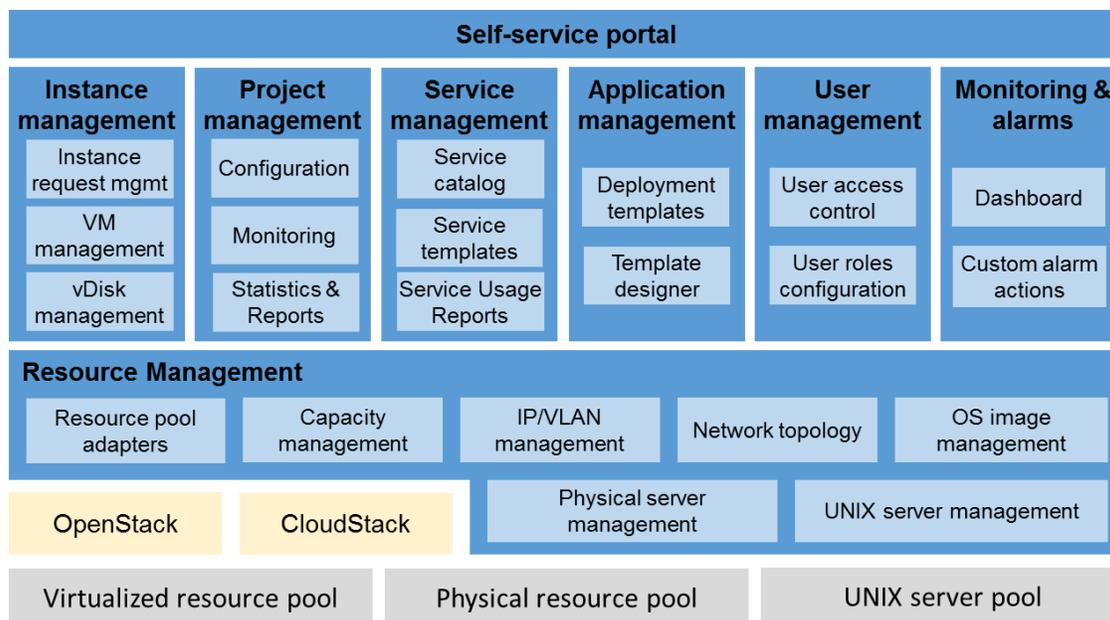


Figure 1, Teraproc Cloud Manager Functions

As shown in the figure, not only it handles the non-virtualized environment and UNIX server environment in addition to the virtualized x86 environment, it adds many functions and unique capabilities on top of the commonly used cloud computing infrastructure management solutions like OpenStack, Cloud Stack, etc.

Teraproc Cloud Manager unique capabilities

Teraproc Cloud Manager has the following unique capabilities

- Support All Workloads

Teraproc Cloud Manager supports all types of applications no matter whether the application runs on a physical server, a virtual server, or even a traditional UNIX server. This allows mixing traditional workloads with cloud workloads. Users can request a dynamic infrastructure on a physical server for high performance or on a virtual machine for flexibility. This allows the cloud solution to be applied to a broader range of application areas.

- Manage Multiple Resource Pools

Teraproc Cloud Manager can manage multiple heterogeneous resource pools including those managed by OpenStack, Cloud Stack, VMware, or bare metal physical servers, no matter these resource pools are connected via LANs or WANs.

- Application Service Templates

In addition to infrastructure services, Teraproc Cloud Manager allows a site to define and build application-specific templates, then publish them as services. This allows users to focus on applications rather than the infrastructure when they request services.

- Intuitive Management Console

Teraproc Cloud Manager is designed to allow quick and easy cloud deployment. Heavy training for system administrators is not required to operate the cloud environment powered by Teraproc Cloud Manager. This reduces the time to solution readiness. It also reduces the total cost of ownership (TCO) and the initial cost for small and medium IT shops to start benefiting from cloud computing. The management console is designed to provide all functionalities the Teraproc Cloud Manager has. It is also very intuitive and takes no time for administrators to get familiar with.

- Built-in Reports

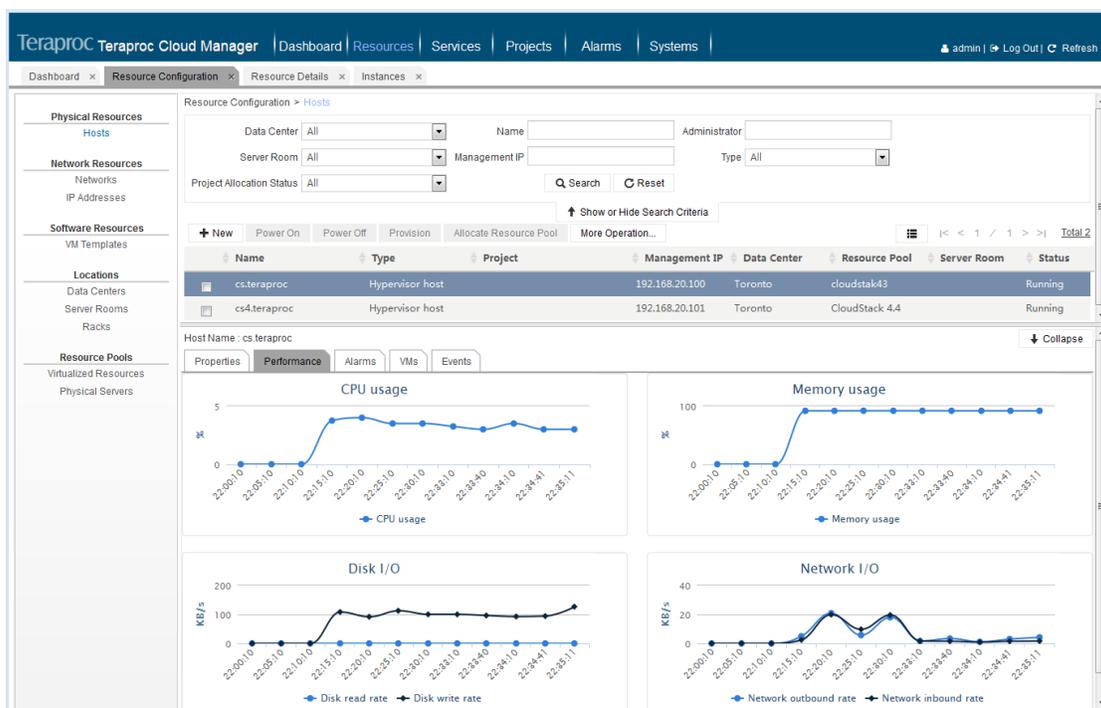
Built-in reports allow system administrators to analyze workload trends and project resource utilization, as well as to plan capacity expansions. Reports are categorized by resources, services, and projects.

Teraproc Cloud Manager Capabilities

Teraproc Cloud Manager has comprehensive management capabilities that allows multiple data centers to be managed through a single management console. The following paragraphs highlights its capabilities.

Resource Management

Resource management interacts with resource pools managed by cloud infrastructure management tool (OpenStack, CloudStack, Elaster, etc.), bare metal x86 servers, as well as UNIX servers. It is the foundation of Teraproc Cloud Manager.



Resource Configuration

Through the management console, administrator can configure:

- Location management

Teraproc Cloud Manager allows a site to input device location information for tracking purposes. The information helps system administrators to pin point where the device is physically located in a server room.

- Resource pool management

Resource pools in the system consist of virtualized resource pools and physical server pools. Hypervisors in virtualized resource pools are managed by CloudStack, OpenStack, or Elaster. Physical servers are placed in host resource pools. Teraproc Cloud Manager can connect and manage multiple heterogeneous resource pools simultaneously.

- Physical resource management

Physical servers, including bare metal hosts, and hypervisor hosts can be managed by Teraproc Cloud Manager. Hypervisor hosts managed by cloud infrastructure management software like OpenStack and CloudStack can be connected to Teraproc Cloud Manager automatically. Bare metal host information can be imported from Excel files. System administrators then are able to manage these physical resources by provisioning and powering cycle them.

Physical resource management capability helps data center management to oversee all cloud infrastructure devices via a “single-pane-of-glass” management console.

- Network resource management

A network resource refers to a virtual local area network (VLAN) and an IP address. This is used by hosts and VMs. Teraproc Cloud Manager allows system administrators to manage network resources and configure networks for hosts and VMs.

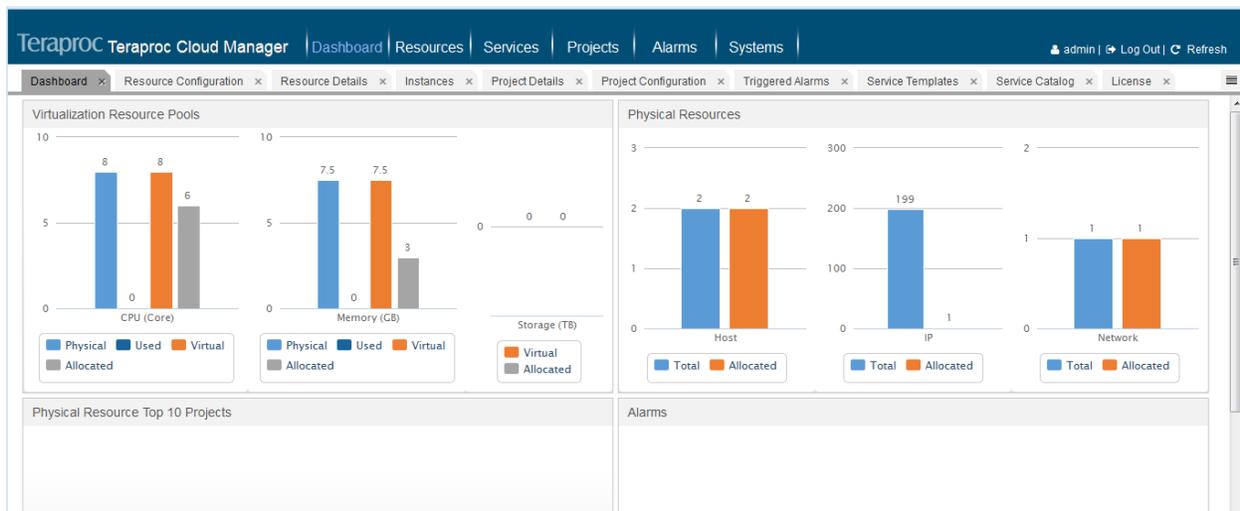
- Software resource management

The software resources described here refer to only VM templates. System administrators can manage these resources.

Resource Monitoring

A centralized monitoring system is included in Teraproc Cloud Manager. Monitored metrics are accessible via the management console. The topology of each data center is shown. Most devices in data centers are monitored including hosts, virtual machines.

The monitoring system collects information from individual resource pools then consolidates them into a single database. This allows a site to correlate information across multiple resource pools and data centers.



Resource Reports

In addition to the monitoring pages that show what's now, historical statistics could be reported. Resource reports display the usage and availability of physical resources, virtual resources, storage, and IP addresses at different time intervals. System administrators can query performance and usage for hypervisor hosts and non-hypervisor hosts.

Tailored for the needs of individual sites, custom reports could be developed by accessing data directly from the monitoring database.

Projects

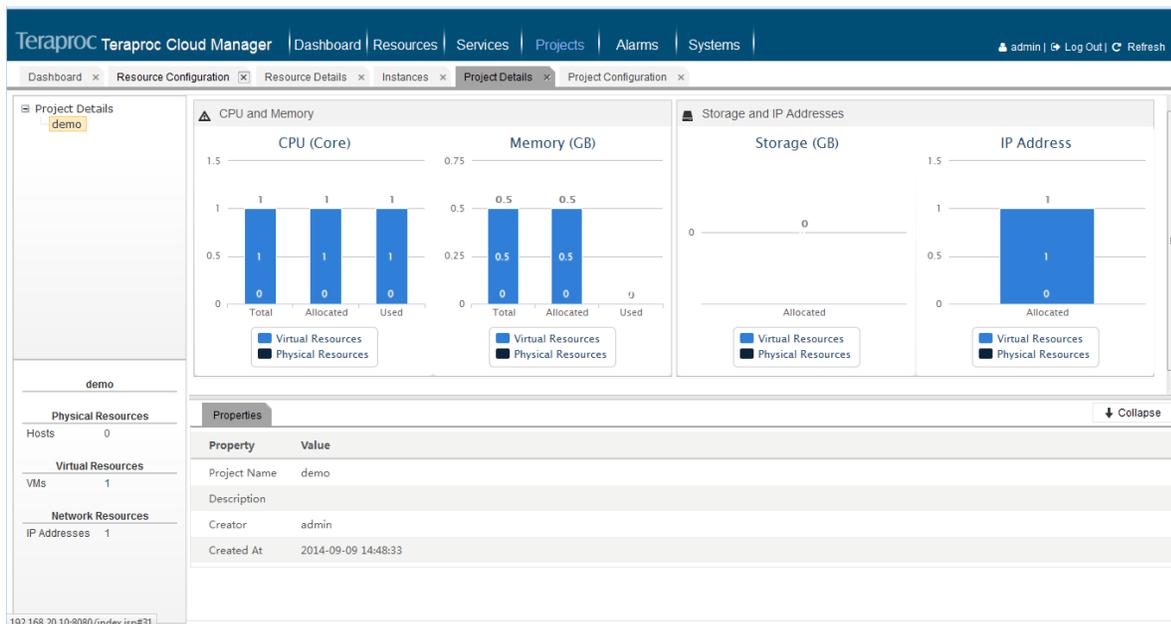
Teraproc Cloud Manager has a category called projects. Projects are used to represent groups of service instances that belong to a project, a business unit, or arbitrary groups. Such grouping allows a site to analyze and charge back cloud resource usage for specific business needs.

Project-based Operations

Teraproc Cloud Manager collects and displays capacity allocation and real usage based on projects.

On the monitoring page, projects that have the top allocation of CPUs, memory, storage, or IP addresses, are displayed in bar charts. Statistics about the allocation and usage of the resources associated with each project, including the CPU, memory, storage, and IP address resources, is also collected and displayed in bar charts. In addition, detailed information about each project is displayed.

Teraproc Cloud Manager tracks resource usage information for each project including the number of resources, such as hosts, VMs, and IP addresses. In terms of management actions, system administrators can provision, power on, or power off a host, or start, stop, create and assign, migrate, or delete a VM as well as allocate a project to a VM. Administrators can also cancel the project allocated to a VM, or modify the specifications of a VM.



Project Reports

Service resource usage reports can be generated by users or by services in a specified time period. The information helps administrators optimize resource capacity planning.

Instance & Service Management

Instances

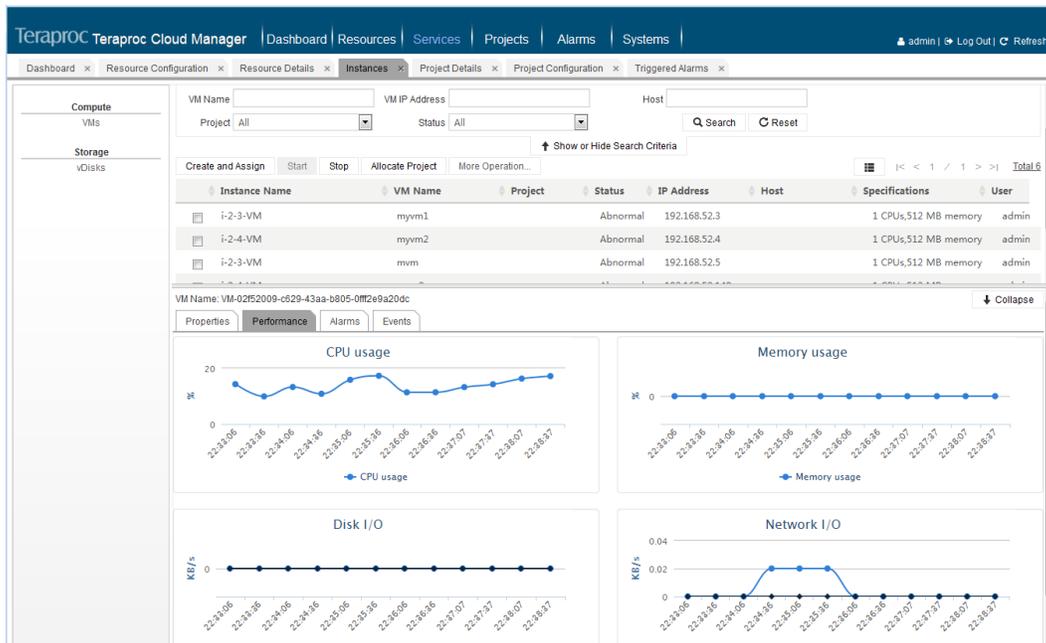
Teraproc Cloud Manager allows system administrators to track and manage service instances. Service instances are grouped according to instance types. They are compute resource instances (hosts and virtual machines), storage instances (vDisks), and application instances.

- VM instance management

For VM instance management, Teraproc Cloud Manager supports start, stop, create, take over, restart, delete, or migrate a VM, allocate a project to a VM, cancel the project allocated to a VM, modify the specifications of a VM, and query detailed information, alarm and event information, or performance information about a VM.

- vDisk instance management

System administrators can create, attach, detach, or delete a vDisk, allocate a project to a virtual disk, and query detailed information about a vDisk.



User Request Management

After a user submits a request, system administrators can optionally modify the configuration of the request. System administrator can also approve, trace, or query a request. A request can be approved by a single administrator or hierarchically approved by multiple administrators. If a request is set to be hierarchically approved, the approval of each hierarchy can be set to be automatically or manually performed. The approval hierarchy and approval mode can be modified on the service parameter configuration page.

Service Templates

A service template encapsulates multiple system resources in an available service. Service templates in the system include VM service templates, vDisk service templates, and backup service templates. End users can query and apply for a service template only after the service template is associated with the service catalog. Basic service templates are pre-installed when the system is initially deployed. You are also allowed to create other service templates based on your service requirements. A self-defined template can be used only after it is approved and published by an administrator. An existing template can be modified, deleted, or unpublished.

Service Catalog

Service catalog is used to display service templates by service type. A common service catalog is pre-installed during initial system deployment. The system also allows you to create service catalog nodes to suit your needs. The service catalog modified by an administrator is displayed on the self-service portal (Teraproc Cloud Services) upon users' login only after it is published and approved.

Managing Traditional Workloads Together with Cloud Computing Workloads

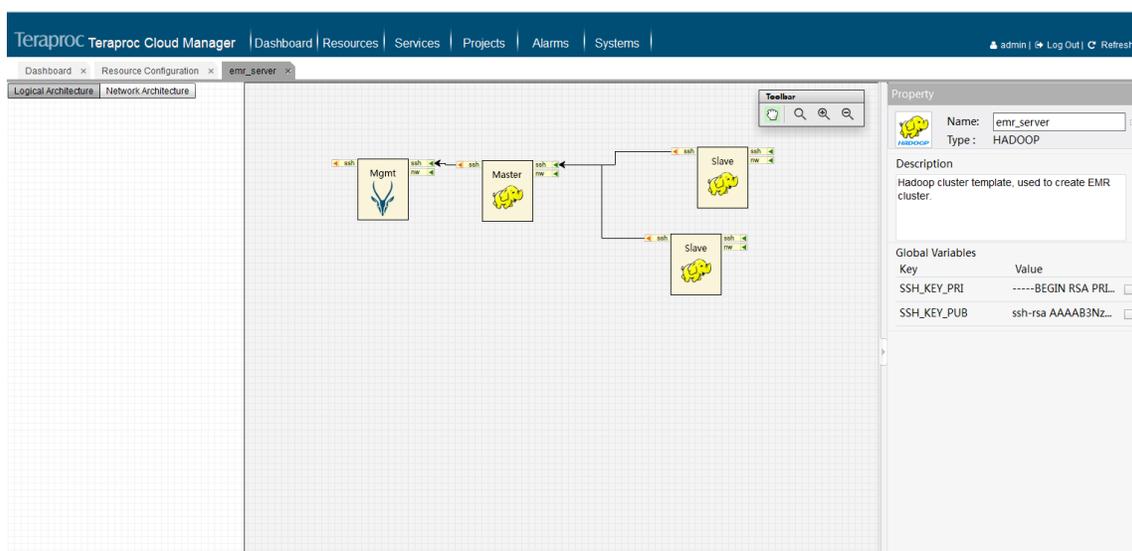
System administrators can edit the cloud service display page on the Teraproc Cloud Services by dragging and dropping icons to modify the content and layout of the page. After being edited, the page is automatically updated.

Application Template Management

Teraproc Cloud Manager allows administrators and users to create application deployment template. The template defines resources the application needs in each of its tiers, the network connections among all tiers, as well as the scalability of each tier.

Once defined, a user can easily launch an application instance at any time. Virtual servers will be provisioned and the application is deployed automatically.

User then can run the application and monitor its performance through the self-service portal.



User & System Management

Users

Users that have operation rights on Teraproc Cloud Manager or Teraproc Cloud Services are both created on the Users page on Teraproc Cloud Manager. However, users in Teraproc Cloud Manager and Teraproc Cloud Services are administrators and common users, respectively. It supports adding, modifying, deleting, or querying a user, as well as resetting the password of a user.

User Roles

A role defines the operation rights of an administrator. Therefore, when an administrator is added to the system, its role can be defined for access control. Role can be edited and managed.

System Parameters

This function is provided for querying and modifying system parameters. With this function, you can add system parameters during system deployment and commissioning, such as setting the request approval parameter and the request expiration notification parameter. In system operating, you can also modify system parameters as required.

Alarms & Events

- Alarm query and handling

All alarms generated in the system are listed. System administrator can ignore or handle an alarm, and query an alarm by multiple search criteria.

- Event query

All events occurred in the system are displayed by event type. These events include system events, VM events, request events, instance events, and host events.

- Alarm configuration

Both alarm triggering rules and alarm handling rules can be configured. An alarm triggering rule defines the indicator and threshold for triggering an alarm and the notification method of sending the alarm after it is generated. Alarms are triggered in the system only after alarm triggering rules are configured. An alarm handling rule determines the handling operation and priority. Alarm handling rules must be set before you handle alarms. In addition to adding alarm triggering or handling rules, you can also modify, delete, or query existing alarm generation and handling rules.

Self-Service Portal

The self-service portal provided by Teraproc Cloud Manager is called Teraproc Cloud Services. It enables users to apply for and use services in a way like online shopping and to monitor resources and operations. On this portal, users can also subscribe or unsubscribe resources to satisfy their services. In fact, the self-service portal can meet all requirements of all common users. With regard to service instance provisioning, service instance reclamation and resource release, or resource accounting, these processes are automatically performed by background programs in the system, requiring no human intervention.

Typical Scenarios for Teraproc Cloud Manager

Teraproc Cloud Manager has been used by many sites with mixed applications. The following are some commonly used scenarios for users to take advantages of Teraproc Cloud Manager:

Scenario 1: Infrastructure with applications running on physical servers

Sites like high performance computing centers often lease their computing capacities to internal and external users. Due to performance reason, physical servers are used by applications directly instead of virtual servers for applications to take performance advantages of high speed interconnects such as InfiniBand or high performance storage systems like Lustre and IBM GPFS.

In such environments, security is a key. One user's workload cannot be mixed with that of another user. Two users should not be able to access data or any other information each other. VLANs are often used to separate individual user's environment.

Scenario 2: Software development & testing environment

Software developers and QA engineers need different environments to develop and test their software. Such environments often involve different OSs, hardware architecture, as well as different virtualization hypervisors. A self-service system and automatically scheduled provisioning are required to increase development efficiency.

Scenario 3: Internal application hosting

Organizations with many applications hosted across multiple locations need to consolidate these applications with better managed infrastructure to improve reliability. The number of applications keeps changing with respect to the needs from different user groups.

Summary

When modernizing IT infrastructure, it is critical to deploy a solution that can manage both traditional workload and cloud workload. This way, the user can benefit from nimble and elastic infrastructure for their applications, while the return on investment cycle is short.

Teraproc Cloud Manager leverages expertise and practices from distributed computing built in the past two decades and brings more than what open source cloud computing management offers on the table. It is designed for organizations who are seeking safely migrating the existing infrastructure to the modern infrastructure.