

IBM® Tivoli® Software

IBM and F5® Cloud Reference Guide

Overview, Capabilities and Use Cases for providing
end-to-end Infrastructure as a Service

Document version 1.0

Ron Aguirre
IBM Consulting IT Specialists, Tivoli Business Development



© **Copyright International Business Machines Corporation 2011.** US Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

CONTENTS

List of Figures.....	v
Revision History	vi
1 Purpose of this Document	1
2 Beyond Provisioning – Why a combined IBM and F5 solution bring real service management and agility	1
2.1 About this paper.....	1
3 The Main Pieces.....	2
3.1 Hardware used	2
3.2 Environment.....	2
3.3 IBM Service Delivery Manager Overview	4
3.4 IBM Service Delivery Manager uses the following software:	5
4 F5 BIG-IP.....	7
4.1 F5 BIG-IP Appliance Local Traffic Manager (LTM).....	7
BIG-IP Virtual Edition compatibility with VMware hypervisor products.....	7
VMware guest environment	7
Configure appliance	9
Setup appliance	10
5 Use Cases	12
5.1 Create Project with Virtual Servers.....	12
Sequence.....	12
5.2 Delete Project	13
Sequence.....	13
5.3 Add Virtual Server to Project	14
Sequence.....	14
5.4 Reserve Virtual IPs	15
Sequence.....	15
5.5 Release Virtual IPs	16
Sequence.....	16
5.6 Create Load Balancer Policy	17

	Sequence.....	17
5.7	Modify Load Balancer Policy	18
	Sequence.....	18
5.8	Delete Load Balancer Policy.....	19
	Sequence.....	19
5.9	Modify Firewall Policy	20
	Sequence.....	20
6	TSAM 7.2.2 Extensions	21
	Summary	22
	TSAM Roles	23
	Cloud Administrator (CA).....	23
	Cloud Manager (CM)	23
	Team Administrator (TA)	23
	Team User (TU).....	23
	Software Releases and Versions	24
	ISDM 7.2.1	24
	VMware.....	24
	F5 24	
	Bibliography	25
	References.....	26

LIST OF FIGURES

Figure 1: VMware environment	3
Figure 2: ISDM image contents	5
Figure 3: IBM Tivoli Virtualization Overview	6
Figure 4: F5 Initial config	9
Figure 5: F5 Final config	10

REVISION HISTORY

Date	Version	Revised By	Comments
August 31, 2011	0.1	RRA	initial draft
October 23, 2011	0.2	RRA	Reworked draft
December 21, 2011	1.0	RRA	Finalized

1 Purpose of this Document

IBM® and F5® have partnered to collaborate on a technical whitepaper that helps customers, partners and the marketplace appreciate the complimentary capabilities that exist today to support Infrastructure as a Service. While often the focus of partner solutions is a specific task or single use case, F5 and IBM bring to the market a broad set of use cases that solve many business issues, as well as the technical components necessary to realize the value.

This document is intended to be used by technical individuals to quickly stand up a running Cloud infrastructure using IBM and F5, and to demonstrate several scenarios that highlight the scope of what is possible with the combined vendors' products.

2 Beyond Provisioning – Why a combined IBM and F5 solution bring real service management and agility

The most common problem solved by Cloud is the ability to bring on quickly a service, with a consideration of multi-tenancy and billing to ensure a financial means to realize the business benefit. While this single use case still realizes great value to communication and enterprise organizations, business continuity, security, compliance and performance are that much more critical as cloud services is more financially tracked.

2.1 About this paper

- Section 2 describes basic information and the software used in this integration work.
- Section 3 describes the F5 BIG-IP® virtual appliance installation in the environment and also covers the initial setup and configuration of this appliance in IBM Service Delivery Manager (ISDM).
- Section 4 consists of the use cases that would be used to integrate the F5 BIG-IP virtual appliance in ISDM using the Simple Service Resource Manager portal.
- Section 5 is the Summary wrap-up.

3 The Main Pieces

This paper is focused on describing how the F5 BIG-IP virtual appliance is integrated, installed and managed by the IBM Cloud based on IBM Service Delivery Manager.

The Cloud environment used for this integration paper is based on IBM Service Delivery manager (ISDM). ISDM is the core of the IBM Cloud product Cloud Service Provider2 (CSP2). The integration work that IBM has worked with F5 on has been with this product. The CSP2 environment that is used for this is a minimal one for purposes of showing the integration. The ISDM product consists of four images running in a VMware ESX 4.1 server with VMware's vCenter 4.1 controlling the environment. Image storage is done with a local data store on the ESX server. All work was conducted on a single server running this setup. Running in the environment to show it can be proven that the F5 BIG-IP virtual load balancer can be integrated into any ISDM environment.

The first part of this document shows the basics of the setup and configuring of the F5 BIG-IP vm in the ISDM environment. The ISDM install itself is a basic one with all four images installed and configured. In figure 2 this is described. For the saving of time the ISDM install is described at a high level. More detail is spent on the F5 information.

This paper explains that the F5 BIG-IP virtual appliance is installed as an image and can be supported in ISDM "Cloud" environment. The virtual appliance will not be deployed as an application image but will be managed as part of the environment. It is installed as virtual image in this paper but is available as hardware appliance that would be rack mounted. Two F5 BIG-IP appliances were installed. This document is to show integration into ISDM and not tuning of the F5 BIG-IP appliance.

3.1 Hardware used

- IBM x3650
- 2 x 2 Dual Core 3.2GHZ processors
- 48 GB Memory
- 876 GB DASD

3.2 Environment

VMware®

- VMware ESX 4.1
- vCenter 4.1, vCenter run on MS Windows Server 8 Enterprise 64bit virtual image

IBM Service Delivery Manager (ISDM)

- Tivoli System Automation Manager (TSAM) image / SUSE® Enterprise Server 10
- Tivoli Usage and Accounting Manager (TUAM) image / SUSE Enterprise Server 10
- IBM Tivoli Monitoring (ITM) image / SUSE Enterprise Server 10
- NFS image / SUSE Enterprise Server 10

F5

- BIG-IP® LTM® Virtual Edition version 10.2.1

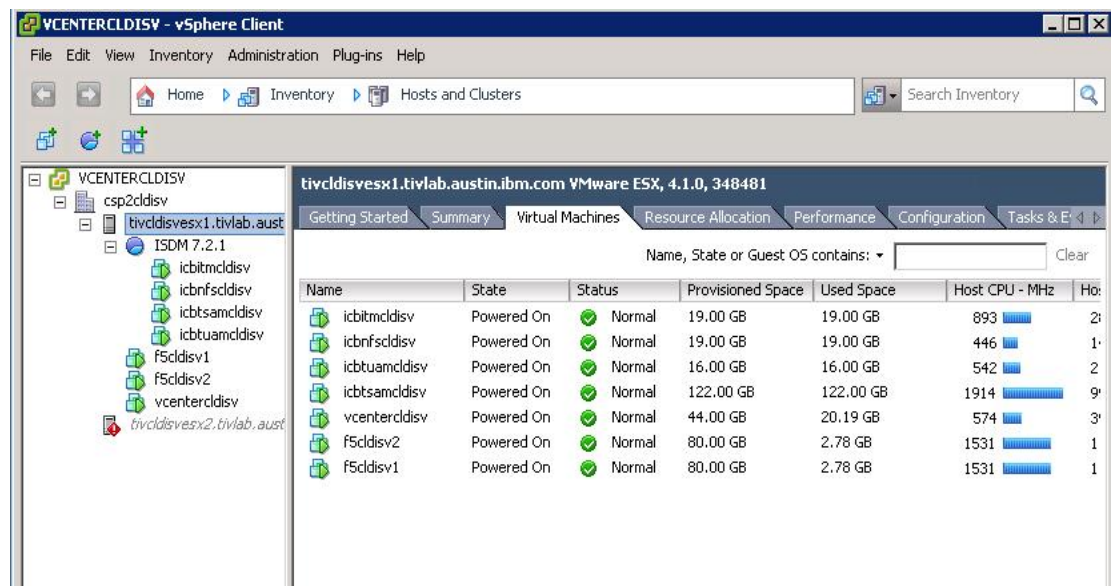


Figure 1: VMware environment

The ISDM information below was taken from the ISDM 7.2.1 Users Guide.

3.3 IBM Service Delivery Manager Overview

IBM Service Delivery Manager is prepackaged and self-contained software appliance that is implemented in a virtual data center environment.

It enables the data center to accelerate the creation of service platforms for a wide spectrum of workload types with a high degree of integration, flexibility, and resource optimization.

Use IBM Service Delivery Manager if you want to get started with a private cloud computing model. The product enables you to rapidly implement a complete software solution for service management automation in a virtual data center environment, which in turn can help your organization move towards a more dynamic infrastructure.

IBM Service Delivery Manager is a single solution that provides all of the necessary software components to implement cloud computing. Cloud computing is a services acquisition and delivery model for IT resources, which can help improve business performance and control the costs of delivering IT resources to an organization. As a cloud computing quick start, IBM Service Delivery Manager enables organizations to exploit the benefits of this delivery model in a defined portion of their data center or for a specific internal project. Potential benefits include:

- Reduction in operational and capital expenditures
- Enhanced productivity - the ability to innovate more with fewer resources
- Decreased time-to-market for business features that increase competitiveness
- Standardized / consolidated IT services to drive improved resource utilization
- Increased resiliency to market demands
- Improved quality of service for IT consumers
- IBM Service Delivery Manager provides preinstalled capabilities essential to a cloud model, including:
 - A self-service portal interface for reservation of computer, storage, and networking resources, including virtualized resources
 - Automated provisioning and de-provisioning of resources
 - Prepackaged automation templates and workflows for most common resource types, such as VMware virtual images and LPARs
 - Service management for cloud computing
 - Real time monitoring for elasticity
 - Backup and recovery

3.4 IBM Service Delivery Manager uses the following software:

- Tivoli Service Automation Manager
- IBM Tivoli Monitoring
- IBM Tivoli Usage and Accounting Manager

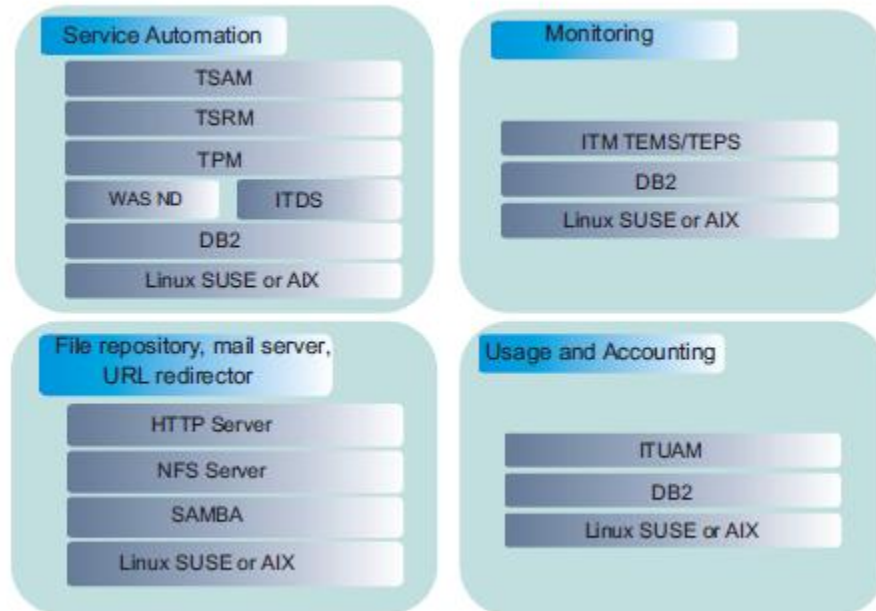


Figure 2: ISDM image contents

The next chart explains where ISDM fits in the overall scheme of IBM Tivoli Cloud Computing product offerings currently.



IBM Tivoli Virtualization Solution Overview

Cloud Product Comparison						
	Tivoli Service Automation Manager (TSAM)	IBM Service Delivery Manager (ISDM) x86	ISDM Power	CloudBurst Series x	CloudBurst Power	Cloud Service Provider Platform
Brief Description	<ul style="list-style-type: none"> A la carte software solution for maximum flexibility Provides a foundation / core building block for establishing a Cloud deployment strategy 	<ul style="list-style-type: none"> Pre-configured and optimized for virtual environments and cloud deployments Ships as a pre-configure service delivery software image. Can be installed on customers existing x86 server platform Includes all benefits of TSAM with additional ISDM features 	<ul style="list-style-type: none"> Pre-configured and optimized for virtual environments and cloud deployments Ships as a pre-configure service delivery <u>software image</u>. Can be installed on customers existing Power server platform Includes all benefits of TSAM with additional ISDM features 	<ul style="list-style-type: none"> Integrated <u>hardware, software and QuickStart services</u> optimized for clouds Pre-installed ISDM platform and preconfigured on IBM Series x Server. Optimized for virtual environments and cloud deployments Includes all benefits of TSAM and ISDM features 	<ul style="list-style-type: none"> Integrated hardware, software and QuickStart services optimized for clouds Pre-installed ISDM platform and preconfigured on IBM Power (AIX) Server. Optimized for virtual environments and cloud deployments Includes all benefits of TSAM and ISDM features 	<ul style="list-style-type: none"> Service Management Network Management Deep CI Management Security Management Storage Management
Products Included	<ul style="list-style-type: none"> TPM for Images 	<ul style="list-style-type: none"> Tivoli Usage & Accounting Manager (TUAM) IBM Tivoli Monitoring (ITM) TSAM <ul style="list-style-type: none"> TSRM TPM TPM for Images Tivoli System Automation (TSA) 	<ul style="list-style-type: none"> TUAM ITM TSAM <ul style="list-style-type: none"> TSRM TPM TPM for Images IBM Systems Director with VMControl 	<ul style="list-style-type: none"> ISDM <ul style="list-style-type: none"> TUAM ITM TSAM <ul style="list-style-type: none"> TSRM TPM TPM for Images TSA IBM Systems Director - ITM for Energy Management IBM System Storage SAN Volume Controller (SVC) Optional 	<ul style="list-style-type: none"> *ISDM <ul style="list-style-type: none"> TUAM ITM TSAM <ul style="list-style-type: none"> TSRM TPM TPM for Images *IBM Systems *SVC Included 	<ul style="list-style-type: none"> ISDM in Core offering Tivoli Monitoring System Automation Usage Metering Service Request Provisioning

Figure 3: IBM Tivoli Virtualization Overview

4 F5 BIG-IP

4.1 F5 BIG-IP Appliance Local Traffic Manager (LTM)

BIG-IP Local Traffic Manager Virtual Edition (VE) is a version of the BIG-IP system that runs as a virtual machine, packaged to run in a VMware® hypervisor environment. BIG-IP Local Traffic Manager VE includes all features of BIG-IP Local Traffic Manager, running on standard BIG-IP TMOS.

Note: The BIG-IP Local Traffic Manager VE product license determines the maximum allowed throughput rate. To view this rate limit, you can display the BIG-IP Local Traffic Manager VE licensing page within the BIG-IP Configuration utility.

BIG-IP Virtual Edition compatibility with VMware hypervisor products

BIG-IP Virtual Edition (VE) is compatible with VMware ESX® 4.0 and 4.1, and VMware ESXi™ 4.0 and 4.1 hosts.

VMware guest environment

The virtual machine guest environment for BIG-IP Local Traffic Manager VE includes these characteristics:

- 2 virtual CPUs
- 2 GB RAM
- 3 virtual network adapters
- 1 - 40 GB LSI logic disk

Note: When you use the VMware vSphere client system to deploy BIG-IP Local Traffic Manager VE on the ESX or ESXi host system, it is important that you retain the guest environment characteristics as shown here. Modifying these characteristics can produce unexpected results. Also note that the guest environment does not support vmmemctl, the memory balloon driver.

The first steps in deploying BIG-IP Local Traffic Manager VE are to download the Zip file to your local system. You can then run the Deploy OVF Template wizard from within VMware vSphere Client. This wizard copies the file to the ESX/ESXi server and configures some network interface settings. Note that the Zip file contains a virtual disk image based on an Open Virtual Format (OVF) template. By following the steps in this procedure, you create an instance of the BIG-IP system that runs as a virtual machine on the host system.

Important: Do not modify the configuration of the VMware guest environment. This includes the settings for the CPU, RAM, and network adapters. Doing so can produce unexpected results.

- Two F5 BIG-IP virtual appliances were installed in the discussed environment.
- **Important:** the license of the F5 BIG-IP virtual appliance is tied to the virtual image identifier (UIID). If you delete the image and try to recreate it from the data store will fail.
 1. In a browser, open the F5 Downloads page, <https://1. downloads.f5.com/>
 2. Download the BIG-IP Local Traffic 2. Manager VE package.
 3. Extract the files from the Zip archive.
 4. Start VMware vSphere and log in.
 5. From the File menu, choose Deploy OVF Template. The Deploy OVF Template wizard starts.
 6. On the Source screen, click Deploy from file, and, using the Browse button, locate the OVA file. For example: `\MyDocuments\Work\Virtualization\BIG-IP-<BIG-IPnn. n.n.n-nn.ova>`
 7. Click Next. The OVF Template Details screen opens.
 8. Verify that the OVF template details are correct, and click Next. This displays the End User License Agreement.
 9. Read and accept the license agreement and click Next. The Name and Location screen opens.
 10. In the Name box, type a name for the BIG-IP virtual machine, such as: `F5_BIG-IP_LB1`.
 11. In the Inventory Location pane, select a data store name. Click Next.
 12. If the host system is controlled by VMware vCenter, the Host Cluster screen opens. Choose the desired host and click Next. Otherwise, proceed to the next step.
 13. Map the source network Management Network to the name of a destination management network in your inventory. An example of a destination management network is Management.
 14. Map the source network Internal Network to the name of a destination nonmanagement network in your inventory. An example of a destination internal network is Private Access.
 15. Map the source network External Network to the name of an external network in your inventory. An example of a destination external network is Public Access.
 16. Click Next. The Ready to Complete screen opens.
 17. Verify that all deployment settings are correct, and click Finish.

Configure appliance

Next step is to configure the virtual appliance.

1. Open the console of F5_BIG-IP_LB1 through Vsphere Client.
Default login = root/default
2. Run config command on the command line to run the command line-based Configuration Utility shell to configure an IP address, netmask, and default route for the management interface.

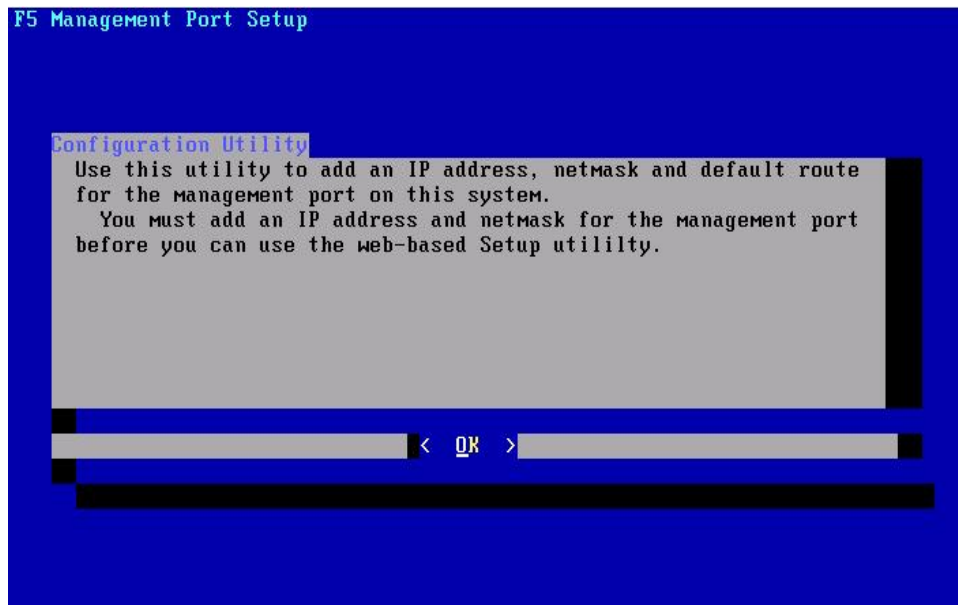


Figure 4: F5 Initial config




Figure 5: F5 Final config

Setup appliance

After configuration is complete you will need to setup the virtual appliance.

1. Access the UI console through a browser using **https://<management IP>**
2. At the logon prompt, type admin for the username and admin for the password.
3. Activate License Key. Suggest to use the Manual button.
4. Run setup utility on UI: Overview -> Welcome -> Run setup utility
5. Set High availability to Redundant Pair and ensure that Unit ID is set to 1 on F5_BIG-IP_LB1 and Unit ID is set 2 if configuring a F5_BIG-IP_LB2
6. Change root password and admin password for security purposes.
7. Ensure SSH access is enables and SSH IP Allow is set to All addresses.
8. Click Finished
9. Create new user with Advanced Shell permissions.
 - System Users User List
10. Create management VLAN
 - Network VLANs VLAN List
11. Use the High availability wizard to configure High availability between the two Load balancers
 - Templates and Wizards Device Wizards High Availability wizard

Unit ID = ID configured in setup utility

System Redundancy mode = Active/Standby

Redundancy State = Active on LB1, Standby on LB2

Network Failover = Network Failover

12. Next Create Failover VLANs

Failover Self IP for LB1 – x.x.x.x

Failover Self IP for LB2 – x.x.x.x

Network Failover

Peer Management IP IP of opposite LB on Failover VLAN

13. Add config for management network and Failover network

14. Create Floating IP on Failover VLAN as seen in screenshot.

Network Self IP

15. Add a routing Domain for management.

Network Route domains Create

16. Create management self IP

5 Use Cases

These use cases give examples of the integration of the F5 BIG-IP virtual appliance in the ISDM environment. These cases are examples of setting up and configuring the F5 BIG-IP appliance in the ISDM environment.

5.1 Create Project with Virtual Servers

Description

The user performs this use case to request a set of virtual servers with given characteristics in terms of computing power, disk space, operating system, and pre-installed software.

This is to provision network resources and optionally load balancer's virtual IPs (VIPs) to show the applications that will be installed on the virtual servers.

Sequence

1. The user selects the "Request a New Service Virtual Server Management Create Project with Virtual Servers" offering and follows the actual steps to
 - Specify reservation attributes
 - Add one or more virtual servers with a given configuration of CPU, memory, disk size, and additional software to install
2. The user specifies the expected number of virtual servers that will be created later, during service life-cycle
 - The System uses this information to select the suitable subnet type to reserve for the project: from a small to a higher number of IP addresses
3. [The user (optionally) selects the Load Balancer option, in which case she specifies how many VIPs should be reserved for load balancer (the System defaults this number to 1)]
 - The System uses this information to select the suitable subnet type to reserve for the project: from a small to a higher number of IP addresses
4. The user submits the request
5. CA/CSA approves the request when the requester is TA
6. The System performs the actual steps to provision the virtual servers. Only new steps are detailed below:
 - During reservation the System calls a user-exit to select the subnet and the VLAN-ID (in 1-1 association with the subnet) to be reserved for the project. This user-exit takes as input the expected size of the project and the number of VIPs to reserve for load balancer (if the option has been selected). This information is used to determine the optimal type of subnet for the project.

- The default implementation of the user-exit expects the VLAN-ID and subnet stored into the DCM, and it simply reserves them. The Service Provider can modify the user-exit to interact with external OSS tools:
- During reservation the System calls a user-exit to configure the virtual firewall of the customer
- Configure the default rules for the project subnet blocking almost all the inbound traffic. This is an example of default rules:
 - Allow 80 (HTTP)
 - Allow 22 (SSH)
 - Allow 443? (RDP)
 - Deny *
 - During virtual server provisioning configure two vNICs: one for the management VLAN and one for the project VLAN

5.2 Delete Project

Description

The user initiates this to release all virtual servers and in general all resources (network and storage) provisioned during the project's life-cycle.

This is also triggered automatically by the System when the project expires. The System can be configured to require user approvals before deleting a project.

Sequence

1. The user selects the "Request a New Service Virtual Server Management Cancel Project" offering
2. The System shows all existing projects belonging to the customer with whom the user is working
3. The user selects the project to delete
4. The user selects options to a) keep saved images; b) backup virtual servers, and finally submits the operation
5. CA/CSA approves the request when the requestor is TA
6. The System performs the usual steps to remove saved images, backup virtual servers, depending on the options set by the user. Only new steps are detailed below
7. If the project was using the Load Balancer, then the System calls a user-exit to update the configuration of the Load Balancer: i.e. remove the virtual load balancer created for the project
8. The System calls a user-exit to update the configuration of the Firewall: i.e. remove the rules for the project's subnet

9. If the project was using the Load Balancer, the System calls a user-exit to release the VIP addresses reserved for the project
10. The System calls a user-exit to release the project's subnet and VLAN-ID

5.3 Add Virtual Server to Project

Description

The user performs this use case to add one or more (homogeneous) virtual servers to an existing project. This use case is updated from base TSAM in order to provision network resources.

Sequence

The user selects the "Request a New Service Virtual Server Management Modify Project Add "Virtual Servers" offering and follows the actual steps to

1. select the project she wants to modify
2. Add to it one or more virtual servers with a given configuration of CPU, memory, disk size, and additional software to install
3. The user submits the request
4. CA/CSA approves the request when the requestor is TA
5. The System performs the actual steps to provision the virtual servers. Only new steps are detailed below:
6. During virtual server provisioning configure two vNICs: one for the management VLAN and one for the project VLAN
7. The System notifies the user of the fulfillment of the request (e.g. e-mailing IP addresses and additional information)

5.4 Reserve Virtual IPs

Description

The user (a CSA, TA, or a CA on behalf of a customer) performs this use case to reserve one or more Virtual IP addresses to be used to advertise applications hosted by a cluster of virtual servers in a project. So, VIPs are reserved in the subnet of the selected project. Once reserved, they can be used to create load balancer policies (association between a VIP:port and a set of virtual server IP:port) on the load balancer device.

Sequence

1. The user selects the “Request a New Service Virtual Server Management Modify Project Load Balancer Reserve Virtual IPs” offering
2. The System shows existing projects belonging to the customer with whom the user is working.
3. Only projects that use the load balancer are selected by the System. The user cannot reserve VIPs on a project that has not been created to use the load balancer.
4. The user selects a project
5. The System displays to the user the following information
6. VIPs actually reserved
7. Existing Load Balancer policies: VIP:port {virtual server IP:port, ...}
8. Number of additional VIPs that can be reserved on the project’s subnet
9. The user verifies that there are rooms available and enters the number of additional VIPs to reserve
10. The System enforces a value up to the remaining number
11. The user submits the request
12. CA/CSA approves the request when the requestor is TA
13. The System calls a user-exit to reserve the number of VIPs on the project’s subnet
14. The System notifies the user of the fulfillment of the request (e.g. e-mailing IP addresses and additional information)

5.5 Release Virtual IPs

Description

The user (a CSA, TA, or a CA on behalf of a customer) performs this use case to remove the reservation of VIPs from the project's subnet.

Note that only VIPs that are not used in load balancer policies can be released.

The use case can be triggered by a custom automation (e.g. CronTask) for the VIPs that are no more referenced in any project.

Sequence

1. The user selects the "Request a New Service Virtual Server Management Modify Project Load Balancer Release Virtual IPs" offering
2. The System shows projects belonging to the customer with whom the user is working filtered by "using load balancer"
3. The user selects a project
4. The System displays to the user the following information
5. VIPs actually reserved
6. Existing Load Balancer policies: VIP:port {virtual server IP:port, ...}
7. Number of additional VIPs that can be reserved on the project's subnet
8. The System graphically enables "release option" only on VIPs that are not used in load balancer's policies
9. The user selects the "release option" for one or more VIPs
10. The user submits the request
11. CA/CSA approves the request when the requestor is TA
12. The System calls a user-exit to release the selected VIPs on the project's subnet
13. The System notifies the user of the outcome of the request

5.6 Create Load Balancer Policy

Description

The user (a CSA, TA, or a CA on behalf of a customer) performs this use case to advertise an application/service running on a cluster of virtual servers in a project (e.g. http traffic on port 80). This happens defining a load balancer policy and deploying on the load balancer device.

As outcome of this use case the load balancer device gets configured with the additional policy and then the advertised application/service can be turned on since it will be reached through the load balancer.

Sequence

1. The user selects the “Request a New Service Virtual Server Management Modify Project Load Balancer Create Policy” offering
2. The System shows projects belonging to the customer with whom the user is working filtered by “using load balancer”
3. The user selects a project
4. The System displays to the user the following information
5. VIPs actually reserved
6. Existing Load Balancer policies: VIP:port {virtual server IP:port, ...}
7. Number of additional VIPs that can be reserved on the project's subnet
8. The user selects a VIP and creates a policy: VIP:port
9. The System verifies the uniqueness of the VIP:port
10. The user then selects the virtual servers that must be balanced
11. The System creates a policy VIP:port {virtual server IP:port, ...} defaulting the virtual server's port to the port used for the VIP
12. The user can optionally change the port used for the virtual server
13. The user can specify other load balancing options (e.g. load balancing algorithm)
14. The user submits the request
15. CA/CSA approves the request when the requestor is TA
16. The System calls a user-exit to create the policy on the load balancer device
 - i.e. a TPM workflow in the F5 BIG-IP device. If this is the first policy, the workflow takes care of first creating the virtual load balancer on the physical device, and then add the policy
17. The System notifies the user of the outcome of the request

5.7 Modify Load Balancer Policy

Description

The user (a CSA, TA, or a CA on behalf of a customer) performs this use case to modify the cluster of virtual servers that are “load balanced”, or to modify the options of a load balancer policy.

Sequence

1. The user selects the “Request a New Service Virtual Server Management Modify Project Load Balancer Modify Policy” offering
2. The System shows projects belonging to the customer with whom the user is working filtered by “using load balancer”
3. The user selects a project
 - The System displays to the user the following information
 - VIPs actually reserved
4. Existing Load Balancer policies: VIP:port {virtual server IP:port, ...}
5. Number of additional VIPs that can be reserved on the project’s subnet
6. The user selects an existing policy (i.e. VIP:port)
7. The System displays the details of the policy
8. The user modifies the list of virtual servers that are load balanced: add / remove / change port
9. The user modifies some load balancing option (e.g. load balancing algorithm)
10. The user submits the request
11. CA/CSA approves the request when the requestor is TA
12. The System calls a user-exit to modify the policy on the load balancer device
 - i.e. a TPM workflow in the F5 BIG-IP tcdriver
13. The System notifies the user of the outcome of the request

5.8 Delete Load Balancer Policy

Description

The user (a CSA, TA, or a CA on behalf of a customer) performs this use case to delete a policy from the load balancer device.

This is also triggered if load balancer policies are active on the project being deleted.

Note that the VIP is not automatically released if no more used in other policies. See use case [] for comments about scheduled tasks to delete resources no more referenced.

Sequence

1. The user selects the “Request a New Service Virtual Server Management Modify Project Load Balancer Delete Policy” offering
2. The System shows projects belonging to the customer with whom the user is working filtered by “using load balancer”
3. The user selects a project
 - The System displays to the user the following information
 - VIPs actually reserved
4. Existing Load Balancer policies: VIP:port {virtual server IP:port, ...}
5. Number of additional VIPs that can be reserved on the project’s subnet
6. The user selects a policy (i.e. VIP:port) for deletion
7. The user submits the request
8. CA/CSA approves the request when the requestor is TA
9. The System calls a user-exit to delete the policy on the load balancer device i.e. a TPM workflow in the F5 BIG-IP tcdriver
10. The System notifies the user of the outcome of the request

5.9 Modify Firewall Policy

Description

The user (a CSA, TA, or a CA on behalf of a customer) performs this use case to customize the firewall policy for the project's subnet.

Sequence

1. The user selects the "Request a New Service Virtual Server Management Modify Project Firewall Modify Policy" offering
2. The System shows all the projects belonging to the customer with whom the user is working
3. The user selects a project
4. The System displays to the user the actual Firewall Policy for inbound trafficⁱ: a list of entries like:
 5. Allow <Adress|Range|ALLⁱⁱ>/Port
 6. Deny ALL/* (cannot be removed)
7. The user has the possibility to modify the ruleset: add / modify / delete rules, with the exception of Deny ALL/*, which cannot be touched.
8. The user also has the possibility to restore the default policy (a ui checkbox)
9. The user edits some changes
10. The System records any change made by the user to the ruleset
11. The user submits the request
12. CA/CSA approves the request when the requestor is TA
13. The System generates the changes to apply to the firewall in terms of: rules added; rules deleted; rules changed and then calls a user-exit to update the customer's virtual firewall
 - i.e. a TPM workflow in the Juniper tcdriver
14. The System notifies the user of the outcome of the request

6 TSAM 7.2.2 Extensions

The examples used in this paper use TSAM 7.2.1. In October 2011 TSAM v7.2.2 provided a new feature called Extensibility that allows any 3rd party to plug-in new code (following certain guidelines) that can extend the TSAM base functionality. These plug-ins are called TSAM Extensions. Each Extension is a self contained package that prereqs TSAM 7.2.2 and installs on top of it. An Extension does not modify or update TSAM modules but just installs additional modules to the TSAM base.

With Tivoli Service Automation Manager, you can complete several network configuration activities. Tivoli Service Automation Manager includes a set of extension nodes and a network data model that can be extended. You can access network artifacts such as network configuration templates, which contain network segments. These network configuration templates are configurable from the Tivoli Service Automation Manager user interface or they can be accessed programmatically through the API.

Tivoli Service Automation Manager includes two APIs: the Tivoli Provisioning Manager API to access data center model (DCM) objects and the Tivoli Service Automation Manager API to modify network configuration instances.

For more information:

<https://www-304.ibm.com/software/brandcatalog/ismlibrary/details?catalog.label=1TW10TS09>

Summary

To summarize what was demonstrated in this document, it has shown the F5 BIG-IP virtual appliance can be placed in an IBM Service Delivery manager (ISDM) "Cloud" environment and managed.

It demonstrated that the appliance can be placed and integrated into the ISDM environment. The Use cases show the steps that can be performed to accomplish this.

IBM Service Delivery manager is the base to IBM's Cloud strategy. It is the base of Cloudburst and Common Service Provider/2 products.

APPENDIX A

TSAM Roles

The self-service user interface identifies four user roles:

Cloud Administrator (CA)

Users in this role are the administrators of the cloud. Only cloud administrators can perform the following tasks:

- define new teams, user accounts and their associated roles
- modify their own information
- register and unregister software images
- allow resource allocations and changes
- check the status of projects and monitor the servers for all users
- approve or deny provisioning requests made by team administrators

Cloud Manager (CM)

Users in this role are the read-only administrators of the cloud. They can perform the following tasks:

- modify their own information (except for role and team membership)
- check the status of projects and monitor the virtual servers for any team

Team Administrator (TA)

Users in this role can perform the following tasks:

- create user accounts for other users
- modify their own information (except for role and team membership)
- modify information for team users
- place requests for provisioning servers and check the status of projects
- monitor the servers
- change server status or password
- log in and use the provisioned servers and applications

Team User (TU)

Users in this role can perform the following tasks:

- modify their own information (except for role and team membership)
- view projects available for their team
- check the status of the servers provisioned for their team
- log in and use the provisioned servers and applications

APPENDIX B

Software Releases and Versions

ISDM 7.2.1

- Tivoli Service Automation Manager (TSAM) 7.2.1
- IBM Tivoli Monitoring (ITM) 6.2
- Tivoli Usage and Accounting Manager (TUAM) 7.1

VMware

- vCenter 4.1 installed in MS Windows 2008 Server Enterprise x64-bit VM
- ESX 4.1

F5

- BIG-IP Local Traffic Manager Virtual Edition - 10.2.1.297.

Bibliography

- IBM Service Delivery Manager Version 7.2.1 (Revised February 2011) SC34-2614-01
- BIG-IP® Systems: Getting Started Guide version 10.1 MAN-0300-00
- Configuration Guide for BIG-IP® Local Traffic Manager™ MAN-0292-01
- BIG-IP® Local Traffic Manager™: Implementation version 10.2 MAN-0293-02
- Tivoli Service Automation Manager Version 7.2.2 Extensions Guide SC34-2658-00

REFERENCES

<http://www.f5.com>

<http://www.f5.com/products/big-ip/local-traffic-manager.html>

<http://www.ibm.com>

<https://www-304.ibm.com/software/brandcatalog/ismlibrary/details?catalog.label=1TW10TS01>

<https://www-304.ibm.com/software/brandcatalog/ismlibrary/details?catalog.label=1TW10TS09>

http://publib.boulder.ibm.com/infocenter/tivihelp/v10r1/index.jsp?topic=/com.ibm.tsamee.doc/c_supportedmgementenvironment.html&S_CMP=rnav



© Copyright IBM Corporation 2011
IBM United States of America
Produced in the United States of America
US Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.*

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PAPER "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes may be made periodically to the information herein; these changes may be incorporated in subsequent versions of the paper. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this paper at any time without notice.

Any references in this document to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
4205 South Miami Boulevard
Research Triangle Park, NC 27709 U.S.A.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at <http://www.ibm.com/legal/copytrade.shtml>.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

UNIX is a registered trademark of The Open Group in the United States and other countries.

F5, LTM and Big-IP are registered trademarks of F5 Networks, Inc.

Local Traffic Manager is a trademark of F5 Networks, Inc.

i

ii