ORACLE
Enterprise Manager

Oracle® Enterprise Manager Ops Center

Release 11gR1

Deployment Considerations

November 18th, 2010

Prepared by Mike Wookey, Nick Stephen

Oracle Corporation

Table of Contents

Product Design and Resource Utilization

Overview

Oracle Enterprise Manager Ops Center, hereinafter referred to as Ops Center, is designed to scale from small installations up through large installations. To enable this, key components of the product architecture must be deployed appropriately to meet the requirements. The following are the key components of the product with which you should be familiar:

1. **Browser user interface.** Ops Center provides a rich interface inside a browser. The user interface heavily leverages JavaScript and AJAX technology to provide a highly dynamic and asynchronously updating user view.

2. **Enterprise Controller.** The Enterprise Controller is is the core component of Ops Center, it provides a centralized place where management information is stored and from which operations are initiated. The browser user interface connects into a standard API that the Enterprise Controller provides.

3. **Proxy Controller.** The Proxy Controller is designed to be distributed throughout the networks. It provides proxy capabilities for operations that have to be logically located close to the targets due to network considerations, for example OS provisioning activities. It also provides a fan out for minimizing network load.

4. **Agent.** Agents run in two modes, the basic agent provides support for monitoring and updating of the host operating system. The virtualization management agent supports basic and virtualization control.

The following illustration shows the relationship between the browser console, Enterprise Controller, Proxy Controller, and multiple Agents:

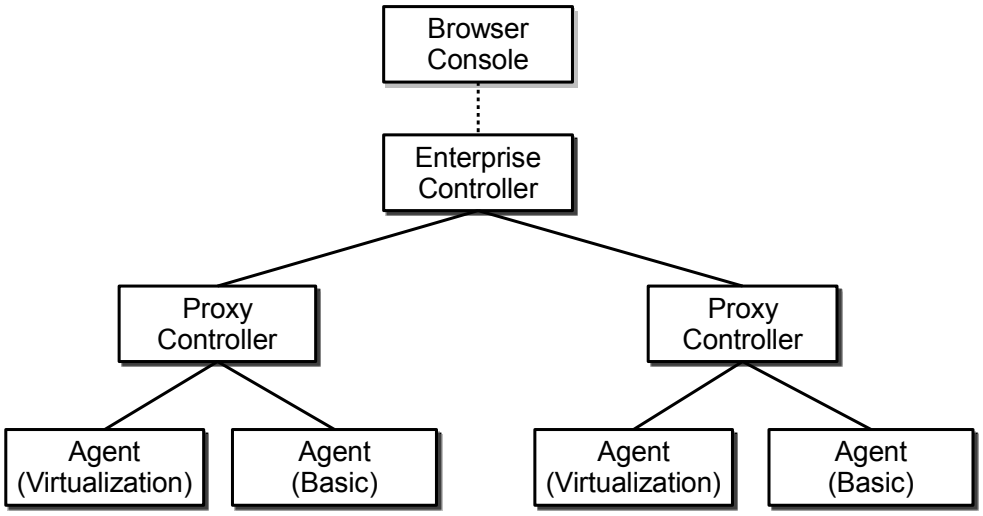

Product Design

To fully understand the best approach to a deployment, it is important to understand how the product uses resources and under what use conditions.

Ops Center provides such a wide range of potential use cases that it is difficult to provide an accurate and simple answer as to how far a particular product deployment can scale.

Ops Center is designed around two key design models:

1. A distributed state machine that maintains the information about all of the managed assets and what can be done to the assets under Ops Center's control. This state machine is incremented by discovering new assets, adding new services or metrics to the machine or creating new logical groups of assets. The state machine is decreased in size by deleting

assets. Each modification of a variable, for example a CPU utilization measure, causes a change in the state machine

2. An asynchronous job dispatch system. All actions that are requested through the Ops Center console are dispatched asynchronously through the system. The job system has a throttle that prevents the overloading of the system and manages queues of pending jobs.

Ops Center's State Machine

Ops Center's state machine is maintained in a combination of memory and database. As the managed environment of assets increase, so will the resource utilization (mostly in terms of memory required) of the system.

The diagram above shows a simple deployment of Ops Center. Each component of Ops Center manages part of the state machine. A basic agent will add the asset information for the operating system being managed and insert all of the relevant metrics, data, and operations that are supported as part of a discovery action. This part of the state machine will then be pushed up to the Proxy Controller, which manages the part of the state machine for all of the assets it is managing and ultimately up to the Enterprise Controller, which can see all of the state machine.

The state machine imparts the highest load on the managed systems during discovery actions. Every part of the system is involved, beginning with the initial location of the asset by way of network requests to the addition of the asset into the state machine, through to the propagation and automated correlation of the new asset in the Enterprise Controller.

Steady state operation of the state machine involves the changing of metrics on the assets themselves. For example, an operating system's asset will have metrics showing network I/O utilization, these are sampled periodically by the agent from the operating system and updated into the state machine on the agent. The agent caches to the state machine, per policy updates, and then sends the updates up to the Proxy Controller. The Proxy Controller follows a similar process up to the Enterprise Controller. The network load in steady state is very low because only changes to metrics are sent and connection management caching is in place to reduce the SSL handshake overhead between the tiers.

The browser user interface reflects a view of the state machine. Assets and their relationships as they exist in the state machine are displayed in the navigation panels, details of the assets are displayed in the center panels and actions associated with the assets are displayed in the Action panel. The Enterprise Controller manages the translation of data from the state machine into a form the browser can display. Viewing different assets or asset groups can place a processing load on the Enterprise Controller. Each

user performing these actions in parallel will increase the processing load on the Enterprise Controller. In addition, a small amount of memory is dedicated to each user session.

Ops Center's Job Management System

The job management system is used to perform actions that have been requested by the user. The job management system is completely asynchronous and active jobs are capped and queued to manage resource utilization by the entire Ops Center system. Depending on the particular action, job actions can run on the agent, the Proxy Controller, or the Enterprise Controller, or a combination of all three.

For example, an OS provisioning job is managed from the Proxy Controller that is managing the system that is the target of the action. The Proxy Controller will run the majority of the processes during such an action and will impart the majority of the CPU load that is used during such an action.

However, a software patch install will execute on the OS and be run by the agent on the hosting OS. Here the agent performs all of the requests required to fulfill the required action and will impart a minor CPU load on the host through job completion.

How Different Use Patterns Affect Resource Utilization

Use Patterns

Ops Center is designed to support an increasing number of administration and management actions. Each type of action and management analysis imparts a certain resource utilization on the underlying systems that support the Ops Center deployment. The management features that are used throughout a deployment will greatly affect the resource utilization and the demands of Ops Center on its host systems. It is therefore critical to understand the problems that are to be solved by Ops Center, as this will greatly affect the number of assets a deployment can manage.

The graph below shows the relationship between the number of services or use patterns that Ops Center supports in the blue, with the maximum number of assets a deployment can manage. The relationship between these two axis varies according to the particular services being used and the availability of resources from the underlying platform that is being used in the Ops Center deployment.

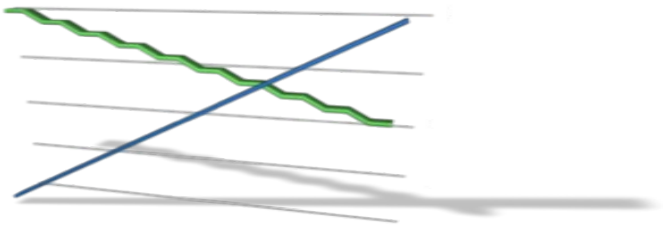

Which services are enabled is often related to the Ops Center deployment model. For example, deploying an agent to a Solaris 10 OS system will automatically enable virtualization management for Oracle Solaris Zones, which in turn increases the size of the state machine for that particular managed asset.

Use Pattern Categories

To simplify deployment, the following five use patterns are identified and compared:

- Hardware Management. Management of server and chassis hardware through Ops Center.
- Hardware Management and OS Provisioning. Management of server and chassis hardware and the ability to provision operating systems to the managed hardware.
- Hardware, OS and Update Management. Management of hardware, operating systems and deploying software updates to the operating systems.
- Hardware, OS and Virtualization Management. Management of hardware, operating systems, and virtualization features of Solaris 10.
- All Management features.

Hardware Management

Hardware management is performed from the Proxy Controller through the systems management Ethernet port on the servers or chassis. In this mode the majority of the utilization occurs at the proxy, which is responsible for performing initial discovery of the hardware and subsequent polling of the hardware for status and configuration changes.

Network sessions are initiated directly from the Proxy Controller to the hardware using specific server and chassis-type protocols that include, IPMI, SSH and SNMP.

Hardware management imparts relatively low resource impact on the Ops Center host OS. The largest impact is on the network traffic emanating from the Proxy Controller to the target servers and the Proxy Controllers should be scaled appropriately.

Hardware and OS Provisioning

Agents are not deployed for OS provisioning or pure hardware management, as no operating systems are being managed. OS provisioning is an action taken on managed servers and is executed primarily from the Ops Center Proxy Controller. The number of OS provisioning jobs that can occur in parallel is metered by the job management system, but OS provisioning does impart a greater load on the infrastructure.

Hardware, OS and Update Management

This mode applies to Solaris 8, Solaris 9, and Linux managed operating systems. Solaris 10 operating systems implicitly enable virtualization management. In this use pattern, agents are deployed to the managed operating systems and update / software provisioning jobs can be completed. This increases the size of the state machine, as operating systems and their configuration are modeled and made available to the update and monitoring capabilities of Ops Center.

Hardware, OS, and Virtualization Management

Solaris 10 Management adds in sophisticated virtualization management for Oracle Solaris Zones and VM Server for SPARC (formerly known as Sun Logical Domains.) Managing these virtualization features involves exposing significantly greater operating system configuration and metrics to the state machine. This will increase the memory utilization at the Enterprise and Proxy Controller tiers.

All Management Features

Enabling all management features of Ops Center will place the largest resource utilization burden on Ops Center's host infrastructure. It is recommended that a co-located Proxy Controller and Enterprise Controller are not deployed, unless the deployment is small.

Scaling Guidelines

As the previous sections have described, many factors can affect the resources used by Ops Center and ultimately its ability to scale to manage large deployments.

General Platform Considerations

Enterprise Controller

When selecting a platform for the Enterprise Controller consider the following factors:

- The Enterprise Controller has a number of parallel executing processes and a multi-core system will significantly improve performance
- Memory is consumed by the core Java virtual machine and its resource utilization will increase proportionally with the number of assets being managed. Should the Java virtual machine be forced into virtual memory, then Ops Center performance will reduce significantly.
- A database is present in Ops Center, as such high performing disk I/O is important in the implementation. Generally, disk striping or the use of a well configured high performing file system with such features built in, such as ZFS, is highly recommended.

Proxy Controller

The following factors should be considered when selecting a platform and environment for the Enterprise Controller:

- Proxy Controllers cannot be used in non-global zones. This is because the Proxy Controller includes the capabilities to perform OS provisioning, which are not supported by non-global zones.
- Network I/O. If the Proxy Controller is supporting Operating System provisioning then it is recommended that the Proxy Controller should be utilizing GigE network cards.

Reference Systems

Scalability figures are based on SPARC and Intel reference systems that are running Solaris 10 10/09 (update 8.)

For performance redundancy reasons, Hardware RAID 10 (0+1) is used throughout these configurations. The number of disks in the RAID set is a key factor in reducing I/O latency, which has a dramatic effect on overall Ops Center scalability. Memory access is affected by DIMM count and placement, therefore these configurations include the maximum number of DIMMs of lowest density available to achieve the best possible memory access times.

Any departure from the stated specifications might have a measurable effect on the overall performance and scalability of the reference configurations.

The scalability figures in System Sizing are quoted from the following reference systems:

- Enterprise Controller on a SPARC Platform
- Proxy Controller on a SPARC Platform
- Enterprise Controller on an Intel Platform
- Proxy Controller on an Intel Platform
    - 146 GB (whole Disk and its mirror)
    - Ops Center Data: /var/opt/sun/xvm directory

Enterprise Controller on a SPARC Platform

Hardware

- SPARC Enterprise T3-1 with 16-Core processor
- CPU Specification: 16-core 1.65-GHz UltraSPARC T3 processor
- Memory: 32 GB (8 x 4 GB DIMM)

Storage

OS root

- External FC/SAN Storage:
  - 4 x 300 GB Disks @ 15,000 RPM
  - Hardware RAID 10 (1+0)
  - 600 GB Volume, 600 GB LUN (whole Volume)

OS swap/dump

- Internal SCSI Storage:
  - 2 x 146 GB @ 10,000 RPM
  - ZFS mirrored for performance and redundancy
-
- External FC/SAN Storage:
  - 8 x 300 GB Disks @ 15,000 RPM
  - Hardware RAID 10 (1+0)
  - 1200 GB Volume, 1200 GB LUN (whole volume)

Proxy Controller on a SPARC Platform

Hardware

SPARC Enterprise T3-1 with 16-Core processor

CPU Specification: 16-core 1.65-GHz UltraSPARC T3 processor

Memory: 16 GB (4 x 4 GB DIMM)

-
-

Storage

OS root

- External FC/SAN Storage
  - 6 x 300 GB Disks @ 15,000 RPM
  - Hardware RAID 10 (1+0)
  - 900 GB Volume, 316 GB LUN (sub-Volume)

OS swap/dump

- Internal SCSI Storage:
  - 2 x 146 GB @ 10,000 RPM
  - ZFS mirrored for performance and redundancy
  - 146 GB (whole Disk)

Ops Center Data – /var/opt/sun/xvm directory

- External FC/SAN Storage:
  - 584 GB LUN (sub-Volume on same 900 GB Volume as root)

Enterprise Controller on an Intel Platform

Hardware

- Sun Fire X4270 with Quad-Core Nehalem processors
- CPU Specification: 2 x 2930 / 8MB
- Memory: 36 GB (18 x 2 GB DIMM)

Storage

OS root, swap/dump:

- Internal Storage
  - 8 x 146 GB @ 10,000 RPM
  - Hardware RAID 10 (1+0)
  - 584 GB Volume, 584 GB LUN (whole Volume)
  - root: 512 GB
  - swap: 72 GB

Ops Center Data – /var/opt/sun/xvm directory

- Internal Storage:
  - 8 x 300 GB Disks @ 10,000 RPM
  - Hardware RAID 10 (1+0)
  - 1200 GB Volume, 1200 GB LUN (whole Volume)

NOTE: This configuration requires two HW RAID Controllers, model SG-XPCIESAS-R-INT.

Proxy Controller on an Intel Platform

Hardware

- Sun Fire x4170 with Quad-Core Nehalem processors
- CPU Specification: 1 x 2930 / 8MB
- Memory: 8 GB (4 x 2 GB DIMM)

Storage

OS root, swap/dump

- Internal Storage
  - 4 x 146 GB Disks @ 10,000 RPM
  - Hardware RAID 10 (1+0)
  - 292 GB LUN
  - root: 256 GB
  - swap: 36 GB

Ops Center Data: /var/opt/sun/xvm directory

- Internal Storage
  - 4 x 300 GB Disks @ 10,000 RPM
  - Hardware RAID 10 (1+0)
  - 584 GB Volume, 584 GB LUN (whole Volume)

NOTE: This configuration requires one HW RAID Controllers, model SG-XPCIESAS-R-INT.

Disk Considerations

Sufficient storage disk space is critical in optimizing performance. The following are recommended disk configurations:

1. Place the /var/opt/sun/xvm directory on a different set of disk spindles from the operating system. Do not co-mingle with the OS.
2. Use at least 4 spindles in the hardware RAID cache configuration.
3. Place the OS on multiple spindles with hardware RAID cache . If the hardware does not allow this with internal drives, then boot from SAN.

The Proxy Controller software is considered relatively lightweight and, in all but the extreme scale cases, does not require a dedicated system.

The following configurations are applicable for SPARC platforms:

- The OS disk configuration is on a Fibre Channel hardware RAID.
- The Oracle Enterprise Manager Ops Center data is on a Fibre Channel hardware RAID.

The following configurations are applicable for Intel platforms:

- All storage is on an internal hardware RAID 10

System Sizing

It is generally recommended that the deployer use proxy fan out as it makes sense in the network to be deployed. For example, OS provisioning capabilities are often restricted to a network subnet.

Changing the resource configuration of the system will affect the numbers, as will more complex mixing of the services across the deployment.

System Matrix

The following table describes the number of systems that each Enterprise Controller and Proxy Controller can support based on the use pattern.

Use Pattern	Enterprise Controller	Proxy Controller
Hardware management	5000 systems	1000 systems
Hardware, OS, update and virtualization management	3000 systems	1000 systems