VMware vCloud® Architecture Toolkit™
for Service Providers

Architecting a VMware vRealize® Log Insight™ Solution for VMware Cloud Providers™

Version 2.9

January 2018

Martin Hosken

VMware, Inc.
3401 Hillview Ave
Palo Alto, CA 94304
www.vmware.com

Contents

List of Figures

List of Tables

Overview

Within the data center, hardware and software systems are typically configured to forward log messages to an external and centralized system message logging (syslog) destination. To improve system administration and provide the VMware Cloud Provider™ community with the security and investigative capabilities it requires, VMware recommends configuring logging to external syslog servers from all hardware in the data center, including VMware ESXi™ hosts, storage, and network components. By facilitating the aggregate analysis of log messages on an external server, visibility is provided into events that affect multiple ESXi hosts and other data center components.

VMware provides several options for syslog target servers. A basic syslog server (VMware vSphere® Syslog Collector) is included as part of the VMware vCenter Server® package. A second option is to implement VMware vSphere Management Assistant for log consolidation. However, for a service provider that is looking for deeper insight into their global data center infrastructure, VMware vRealize® Log Insight™ provides a much more comprehensive and feature rich solution than either vSphere Syslog Collector or vSphere Management Assistant.

vRealize Log Insight gives administrators the ability to consolidate logs, monitor and troubleshoot vSphere and third-party infrastructure, and perform security auditing, compliance testing, log querying, aggregation, correlation, and retention. The vRealize Log Insight virtual appliance includes a syslog server, log consolidation tool, and log analysis tool that will work for any type of device that can send syslog data. vRealize Log Insight administrators can also create custom dashboards based on saved queries that can be exported, shared, and integrated with vCenter Server and VMware vRealize Operations Manager™ to provide a uniform approach to dashboard monitoring and operational management.

1.1 Audience

This document addresses key design and architectural considerations relating specifically to vRealize Log Insight as implemented by the VMware Cloud Provider Program based service provider community. While this document provides a design framework for the vCloud Network community, it does not provide a specific solution for either VMware Cloud Providers or enterprise business customers.

This component of the *VMware vCloud Architectural Toolkit™ for Service Providers* addresses the common questions that arise when designing a syslog solution with vRealize Log Insight. Its primary aim is to focus on the design aspects of syslog within a vSphere environment by providing sample reference architectures to further aid the design work of the VMware service provider partner community and to provide new ideas for syslog management strategies for provider based service offerings.

This document does not consider operational aspects, troubleshooting techniques, or the log analysis skills associated with syslog, because there are numerous resources available to assist with this. The goal of this document is to assist you in undertaking the syslog aspects of a vSphere design and to help you choose among the various options available for the architectural decisions you will be required to make.

1.2 Assumptions and Caveats

VMware, Cisco, and other third party hardware and software information provided in this document is based on the current performance estimates and feature capabilities of the versions of hardware and software described. These are subject to change by their respective vendors.

Understanding vSphere Logs

Before looking at syslog itself, it is important to address vSphere logging and understand how it operates outside of syslog.

2.1 ESXi Log Files

In vSphere 5.x, logging was significantly improved over earlier releases, making it far more straightforward to navigate and access logs, which in turn allows for improved troubleshooting and investigative analysis. In ESXi, all logs are now stored in the `/var/log` directory.

The physical location where logs are written depends on the device used during the ESXi installation. When the ESXi installation device is an SD card, USB key, or remote boot from an SAN environment, a local scratch partition is not created on the installation media automatically during the deployment. Despite its size, ESXi 6.x always sees this type of installation as remote, and as such, logs are stored in RAM disk (disk drive that is made up of a block of volatile memory) and lost when the host is rebooted.

The reason for this is that USB and SD devices are sensitive to high amounts of I/O, so the installer will not place the scratch partition on this type of device. The ESXi installer first scans for a local 4 GB VFAT partition. If it is unable to find one, it will then scan for a local VMFS volume to use to create a scratch directory. If no local VFAT partition or VMFS volume is found, the last resort is to put the scratch partition in the `/tmp/scratch` location on the local RAM disk.

After this type of installation, you will see a warning on the ESXi hosts in vCenter Server indicating that their log files are stored on non-persistent storage. (See *Syslog not configured messages on ESXi host console or in logs (1032460)* at http://kb.vmware.com/kb/1032460.

When this is the case, configure scratch space manually on the ESXi host using the VMware vSphere Web Client or CLI, or as part of a scripted installation procedure.

Because log messages that are stored on RAM disk are not retained after a reboot, troubleshooting information contained within the logs and core files will also be lost. If a persistent scratch location on the host is not configured properly, you might experience intermittent issues due to lack of space for temporary files, and the log files will not be updated. This can be problematic in low-memory hosts, but is not typically a critical issue for ESXi operation.

If the installation device is considered local during deployment, the ESXi host does not usually need to be manually configured with a scratch partition. The ESXi Installer creates a 4 GB FAT16 partition on the target device during the installation, if there is sufficient space to do so. If persistent scratch space is configured, most of these logs (see the following figure) are located on the scratch volume and the `/var/log/` directory contains symlinks (symbolic links) to the persistent storage location.

Figure 1. Host Log Files

```
/var/log # ls
Xorg.log              fdm.log               sdrsinjector.log      vmamqpd.log           vobd.log
auth.log              hostd-probe.log       shell.log             vmauthd.log           vprobe.log
boot.gz               hostd.log             smbios.bin            vmkdevmgr.log         vprobed.log
clomd.log             hostprofiletrace.log  storagerm.log         vmkernel.log          vpxa.log
configRP.log          ipmi                  swapobjd.log          vmkeventd.log         vsantraces
dhclient.log          lacp.log              sysboot.log           vmksummary.log        vsanvpd.log
esxcli.log            osfsd.log             syslog.log            vmkwarning.log
esxupdate.log         rhttpproxy.log        usb.log               vmware
/var/log #
```

Note A *symlink* is a special type of file that contains a reference to another file in the form of an absolute or relative path.

For more information, see *Creating a persistent scratch location for 4.x/5.x/6.0 (1033696),* available at http://kb.vmware.com/kb/1033696.

The local ESXi logs can be accessed and inspected in one of several ways directly on each host:

- VMware vSphere Client™ or VMware vSphere Web Client
- DCUI (View system logs option)
- Web browser (example: https://HostnameOrIPAddress/host)
- Power-CLI "`Get-Log`" cmdlets
- ESXi CLI (example: `cat /var/log/hostd.log`)

For more information about accessing ESXi logs, see *Location of ESXi 5.0 log files (2004201)* at http://kb.vmware.com/kb/2004201.

The following table provides a comprehensive list of ESXi 6.0 host logs, their persistent location, and a description. Many different log files are generated automatically by different ESXi components and services.

Table 1. Table Host Log Descriptions

Log File	Persistent Location	Description
/var/log/auth.log	/scratch/log/auth.log	ESXi Shell authentication information such as success and failures.
/var/log/dhclient .log	/scratch/log/dhclient.log	DHCP client log, including discovery, address lease requests and renewals.
/var/log/esxupdate .log	/scratch/log/esxupdate.log	ESXi patch and update logs (useful if you need to know why a patch failed).
/var/log/hostd.log	/scratch/log/hostd.log	Host management service logs includes virtual machine and ESXi host task and events, communication with the vSphere Client and VMware vCenter Server vpxa agent, and SDK connections.
/var/log/shell.log	/scratch/log/shell.log	ESXi shell usage logs that track commands that were run. ESXi shell usage logs include enable/disable, and every command entered.
/var/log/sysboot.log		VMkernel startup and module loading.
/var/log/boot.gz		A compressed file that contains boot log information. Can be read without unzipping, by using zcat.
/var/log/syslog.log	/scratch/log/syslog.log	Management service initialization, watchdogs, scheduled tasks and DCUI use.
/var/log/usb.log	/scratch/log/usb.log	USB device information such as discovery and pass-through to virtual machines.

Log File	Persistent Location	Description
/var/log/vobd.log	/scratch/log/vobd.log	VMkernel observation events. VOBD is a daemon that VMware and third-party applications use for monitoring and troubleshooting.
/var/log/vmkernel .log	/scratch/log/vmkernel.log	Core VMkernel logs, including device discovery, storage and networking device and driver information, and virtual machine startup information and driver events.
/var/log/vmkwarning .log	/scratch/log/vmkwarning. log	A summary of warnings and alert log messages excerpted from the VMkernel logs.
/var/log/vmksummary .log	/scratch/log/vmksummary. log	ESXi host startup/shutdown, and an hourly heartbeat with uptime, number of virtual machines running, and service resource consumption.
/var/log/vpxa.log	/scratch/log/vpxa.log	Present when a host is connected to a vCenter Server. vCenter Server VPXA agent logs, including communication with vCenter Server and the host management hostd agent.
/var/log/fdm.log	/scratch/log/fdm.log	Present when a host is connected to a vCenter Server. VMware vSphere High Availability logs, produced by the `fdm` service.
/var/log/lacp.log	/scratch/log/lacp.log	Link aggregation control protocol logs. ESXi 5.1 onwards only.
/var/log/hostd-probe.log	/scratch/log/hostd-probe.log	Host management service responsiveness checker.
/var/log/rhttpproxy .log	/scratch/log/rhttpproxy. log	HTTP connections proxied on behalf of other ESXi host web services.
/var/log/Xorg.log	/scratch/log/Xorg.log	Video acceleration.
/var/log/clomd.log	/scratch/log/clomd.log	CLOM daemon. Cluster level object manager (CLOM) logs. These logs are part of the VMware vSAN™ feature.
/var/log/esxcli.log		
/var/log/osfsd.log	/scratch/log/osfsd.log	OSFSD daemon. Object storage file system (OSFS) logs. These logs are part of the vSAN feature.

Log File	Persistent Location	Description
/var/log/ sdrsinjector.log	/scratch/log/sdrsinjector.log	VMware vSphere Storage DRS™ injector log. Profiles the capabilities of the datastores for vSphere Storage DRS.
/var/log/swapobjd .log	/scratch/log/swapobjd.log	SwapObj daemon logs.
/var/log/sysboot.log		Early VMkernel startup and module loading.
/var/log/vmamqpd.log		VMware AMQP daemon log.
/var/log/vmkeventd .log	/scratch/log/vmkeventd.log	Capture of VMkernel events.
/var/log/vprobe.log	/scratch/log/vprobe.log	Output of VMware VProbes™, which provides a facility for transparently instrumenting a powered-on guest operating system, its currently running processes, and VMware virtualization software.
/var/log/vsanvpd.log	/scratch/log/vsanvpd.log	vSAN logs.

The exact list of ESXi host system component logs can vary from the files listed in the table, depending on the version of hypervisor installed. This extensive logging on every single host in the infrastructure can become unmanageable for both operational teams and administrators. This is where the benefits of providing a centralized, consolidated, and searchable syslog solution can be appreciated.

In addition to the listed ESXi 6.x host logs, a number of additional component source logs are, by design, not forwarded to a remote syslog server, but instead exist only on local persistent storage or in RAM disk. The following is list of ESXi component log sources that are not part of the syslog mechanism:

- `configRP.log` – Resource pool changes
- `hostprofiletrace.log` – Host profiles information
- `smbios.bin` – Communication with smbios to provide hardware information
- `storagerm.log` – Storage resource management, including I/O throttling
- `vmauthd.log` – Authd information
- `vmkdevmgr.log` – Hardware device detection
- `vprobed.log` – Control and data logging

2.2 Logging Levels

It is possible to increase or decrease the levels of logging depending on the requirements. For example, you might choose, temporarily, to increase the logging level of a particular component from “Info” to “Trivia”, to obtain a more detailed view or greater insight into what is occurring within that component on the host in question. This technique can be used to help resolve a particular problem that is occurring within the environment. Alternatively, you might be required to reduce logging levels to reduce the amount of syslog traffic traversing a WAN link between the syslog source and target servers. This is described in greater depth in Section 3.4.1, Data Throttling.

Note Administrators must be cautious when increasing logging levels to either “trivia” or “verbose”. If you increase logging to one of these levels to help diagnose an issue, VMware highly recommends reverting to the default logging level immediately after any troubleshooting activity is complete.

2.3 Log Rotation

Rotation refers to the turnover of auto-numbered log files. For example, if you set the maximum number of log files to 10, after the 10th log file is written, the numbering sequence is reset to zero and the first of 10 log files is overwritten with the most current log entries.

The default setting for ESXi 5.x was to rotate 10 files with a size of 5120 KB each, whereas the default for ESXi 5.1.x, 5.5.x and 6.x is to rotate 10 files with a size of 10240 KB each. Log rotation relates only to the hosts’ locally-stored log files, and is designed to prevent the size of log files from getting overly large. Log rotation does not affect remote syslog server log message retention. Rotated logs are compressed at the persistent file location `/var/run/log/`.

For more information on log rotation and modifying logging levels, see *Increasing VMware vCenter Server and VMware ESX®/ESXi logging levels (1004795)* at http://kb.vmware.com/kb/1004795.

2.4 vCenter Server Log Files

In addition to ESXi host logs, vCenter Server maintains its own logs which, like host logs, can be used for auditing and troubleshooting, and operational and security analysis. The vCenter Server log file location varies by version and the underlying OS version. If you are employing a SUSE-based vCenter Server appliance, the logs are stored in the `/var/log/vmware/` folder.

If vCenter Server for Windows has been installed on a Windows 2008 SP2 or 2012 Server based operating system, the logs are located in `%ALLUSERSPROFILE%\VMWare\vappServer\logs`. If the vCenter Server service is running under a specific user name, the logs might be located in that user profile directory instead of `%ALLUSERSPROFILE%`.

Table 2. vCenter Server Log Files

vCenter Server	vCenter Server Appliance	Description
vmware-vpx\vpxd.log	vpxd/vpxd.log	The main vCenter Server logs. Includes all vSphere Client and web services, connections, and events.
vmware-vpx\vpxd-profiler.log	vpxd/vpxd-profiler.log	Profiled metrics for operations performed in vCenter Server. Used by the VPX Operational Dashboard (VOD) accessible at https://*VCHostnameOrIPAddress*/vod/index.html.
vmware-vpx\vpxd-alert.log	vpxd/vpxd-alert.log	Non-fatal information logged about the `vpxd` process.
perfcharts\stats.log	perfcharts/stats.log	Information about the historical performance data collection from the ESXi hosts.
eam\eam.log	eam/eam.log	Health reports for the ESX agent monitor extension and connectivity logs to vCenter Server.
invsvc	invsvc	VMware inventory service.
netdump	netdumper	VMware vSphere ESXi Dump Collector.
vapi	vapi	VMware vAPI endpoint. The vAPI endpoint provides a single point of access to vAPI services.
vmdird	vmdird	VMware directory service daemon.
vmsyslogcollector	syslog	vSphere syslog collector.
vmware-sps\sps.log	vmware-sps/sps.log	VMware vSphere profile-driven storage service.
vpostgres	vpostgres	VMware vFabric® Postgres database service.
vsphere-client	vsphere-client	VMware vSphere Web Client.
cim-diag.log and vws.log	vws	Common information model monitoring information, including communication between vCenter Server and a managed host's CIM interface.
workflow	workflow	VMware vCenter Server workflow manager
SSO	SSO	VMware vCenter Single Sign-On™.

vCenter Server	vCenter Server Appliance	Description
ls.log	–	Health reports for the licensing services extension and connectivity logs to vCenter Server.
sms.log	–	Health reports for the storage monitoring service extension, connectivity logs to vCenter Server, the vCenter Server database, and the xDB database for vCenter Server inventory service.
eam.log	eam.log	Health reports for the ESX agent monitor extension and connectivity logs to vCenter Server.
catalina.<date>.log and localhost.<date>.log	–	Connectivity information and status of the VMware web management services.
jointool.log	–	Health status of the VMware VCMSDS service and individual ADAM database objects, internal tasks and events, and replication logs between linked-mode vCenter Server instances.
vimtool.log	–	Dump of string used during the installation of vCenter Server with hashed information for DNS, username, and output for JDBC creation.
drmdump\	–	Actions proposed and taken by VMware vSphere Distributed Resource Scheduler™ (DRS), grouped by the DRS-enabled cluster managed by vCenter Server. These logs are compressed.

The specific logs generated by vCenter Server will vary to some extent depending on the version deployed. For more information, refer to *Location of log files for VMware products (1021806)* at http://kb.vmware.com/kb/1021806.

For additional information about modifying logging levels in vCenter Server, see *Increasing VMware vCenter Server and VMware ESX/ESXi logging levels (1004795)* at http://kb.vmware.com/kb/1004795 and *Enabling trivia logging in VMware vCenter (1001584)* at http://kb.vmware.com/kb/1001584.

Note When a log file reaches its maximum size, it is rotated and numbered similar to component-nnn.log files. The file might also be compressed.

2.5 Modifying Statistics Levels in vCenter Server

You can also modify the collection levels of host performance statistics on vCenter Server, specifying options that set the statistic collection intervals on hosts, and determine the frequency at which these statistic queries occur. In addition, option settings can also specify the length of time that statistical data is stored in the database and the type of statistical data that is collected.

Note The information provided here about performance statistic collection does not relate directly to the previous discussions about log files, but instead relates to vCenter Server statistics stored in a database. It is included here for completeness.

2.6 vRealize Log Insight and vCenter Server Integration

In addition to collecting syslog data from ESXi hosts, you can also configure vRealize Log Insight to capture tasks, events, and alarms data from vCenter Server. The collection of this information from vCenter Server 5.1 or later provides support for centralized auditing and deeper insight into overall host and site operations. vRealize Log Insight uses the VMware vSphere API to connect to vCenter Server systems and collect data over port 9000/TCP.

2.7 ESXi Syslog Service

As previously discussed, each ESXi host generates a large number of component logs. In an average day, with default logging settings, each ESXi host generates approximately 250 MB of data. Even with a relatively small number of hosts, querying and correlating this log data when troubleshooting a problem can quickly become very difficult. If you are tasked with maintaining hundreds or even thousands of ESXi hosts and other data center devices, managing logs locally and taking advantage of the information in the logs, that job can become nearly impossible. However, VMware provides a relatively straightforward solution to this problem, using a centralized syslog service.

Using a centralized syslog service, each ESXi host runs a local syslog daemon called `vmsyslogd`, which provides a standard mechanism for logging messages from VMkernel and other system components, and directing them to a centralized syslog target. By default, ESXi logs are stored on a local scratch volume, or in RAM disk, depending on the hosts installation device and configuration. To preserve the logs in a centralized location, the ESXi hosts and other devices must be configured to send their logs across the network, directly to a central syslog server. Alternatively, the logs must be directed to a syslog aggregation server, which in turn forwards the syslog messages to the central syslog location.

2.8 Syslog Messaging Overview

Syslog is a standard method by which computer devices can send event messages to a logging server, known as the syslog server. The syslog protocol is supported by a range of computer devices, but the focus of this document is the use of syslog to forward vSphere based logs to a centralized logging server for analysis, troubleshooting, and security auditing.

Unlike SNMP, syslog cannot be used to “poll” devices to gather information. Instead, syslog simply sends messages to a central location where special event handling can be triggered by receipt of specific log messages. It is possible to convert an SNMP trap into a syslog message by employing a service such as `snmptrapd`. Such a service can be installed on a system, which in turn, can forward converted SNMP messages to the remote syslog target. Such a solution can be implemented instead of, or in combination with, an SNMP monitoring system.

The structure of syslog messages is defined in RFC 5424, *The Syslog Protocol* at http://tools.ietf.org/html/rfc5424. (Third-party Web sites are not under the control of VMware, and the content available at those sites might change.) Every syslog message must contain five distinct fields as shown in the following figure. vRealize Log Insight can accept individual messages up to 100 KB in length.

Figure 2. Syslog Message Structure

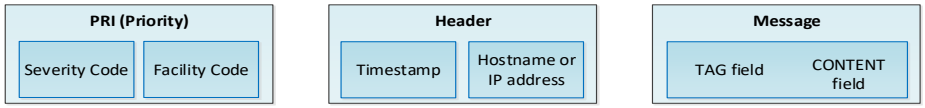

2.9 Syslog Message Severity

The log source (such as an ESXi host) that generates a syslog message will specify the severity of the message using a single-digit integer value between 0 and 7.

Table 3. Message Severity Codes

Severity Code	Severity Name	Description
0	Emergency	System is unusable
1	Alert	Action must be taken immediately
2	Critical	Critical conditions
3	Error	Error conditions
4	Warning	Warning conditions
5	Notice	Normal but significant condition
6	Informational (info)	Informational messages
7	Debug	Debug-level messages

2.10 Syslog Facility Codes

Syslog messages are broadly categorized based on the sources that generate them, such as an operating system, a process, or application, and are assigned a facility code represented by integer values ranging between 0 and 23. For example, Cisco devices use the local facility range values 16-23 (local0 – local7). By default, ESXi uses only the facility code local4. The facility configuration is located in the `config.xml` and `vpxa.cfg` files.

```
<syslog>
  <facility>local4</facility>
  <ident>Hostd</ident>
  <logHeaderFile>/var/run/vmware/hostdLogHeader.txt</logHeaderFile>
</syslog>
```

2.11 Syslog Timestamp

The syslog timestamp indicates the local time, in MMM DD HH:MM:SS format, when a message was generated.

2.12 Syslog Hostname Value

The hostname field consists of the host name of the device where a message originated, which is the name configured on the host itself, or the IP address.

2.13 Syslog Event Message Structure

The syslog message element includes additional information about the process that generated the message. The message element has a TAG field and a CONTENT field. The value in the TAG field is the name of the program or process that generated the message. The CONTENT field contains the details of the message.

Figure 3. Syslog Message Structure

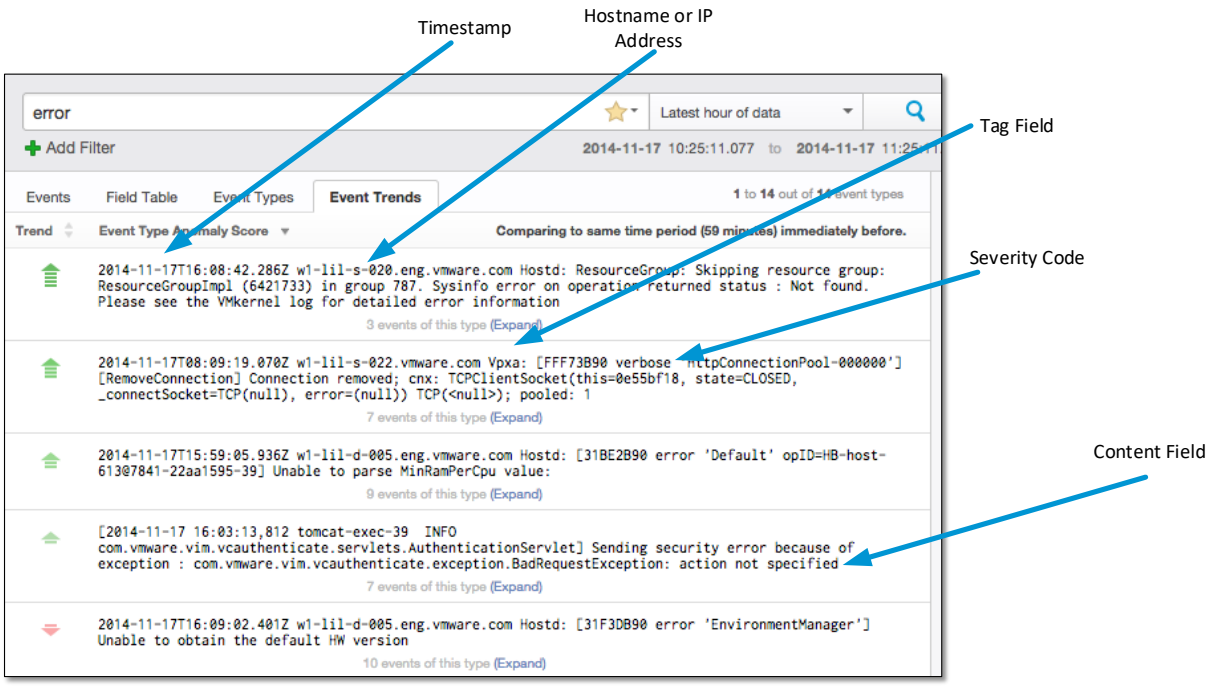

ESXi Host and Device Syslog Configuration

All ESXi hosts or other devices in the environment that are required to forward their log messages to the centralized syslog solution must first be configured to do so. Configuring syslog on ESXi 6.x is a two-step process:

1. Open the firewall to allow for the outgoing syslog traffic.
2. Configure the ESXi host with the IP address or FQDN of the remote syslog server where messages will be sent.

3.1 ESXi Host Firewall Configuration

The host firewall must pass traffic to the centralized syslog system, which is not allowed, by default, in ESXi 6.x.

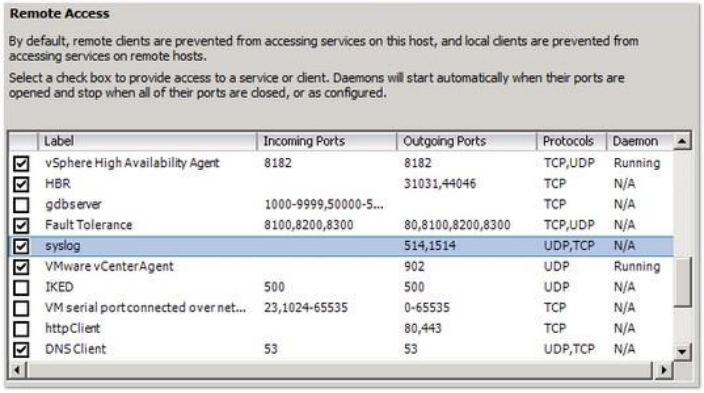

To configure an ESXi host to forward its logs to a centralized syslog server, the `Syslog.global.LogHost` value must be configured in the Advanced Settings panel. The value of this setting represents the remote host to which syslog messages will be forwarded and port on which the remote host will receive syslog messages. This can be configured with the hostname or IP address, followed by the port number, for example: `udp://sysloghost.domin.local:514`.

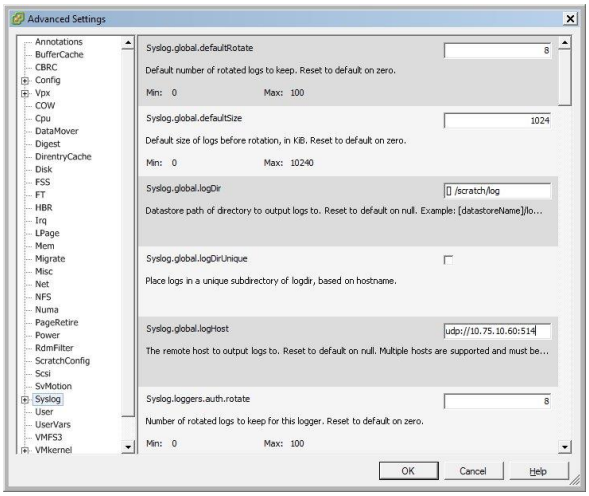

The transport protocol you choose, UDP by default, depends on a site's specific design requirements. TCP and SSL are also supported. The syslog server value can be configured with the hostname or IP address. The remote target system must have a syslog collector installed and be correctly configured to

receive the forwarded syslog messages before the hosts sending syslog messages are configured. Otherwise, configured settings will not take effect.

Note Considerations for the design of centralized syslog collection systems are detailed in Section 3.4, Remote Syslog Design Considerations,

Hosts can be configured in a number of ways to forward syslog messages to a centralized logging system such as vRealize Log Insight. vRealize Log Insight includes the Configure ESXi tool, which can be used to configure the entire ESXi environment for syslog message handling. Other alternatives include `esxcli`, VMware vSphere PowerCLI™, vCLI through the vSphere Management Assistant, vSphere Host Profiles, or the manual approach described previously.

Note For more information about configuring hosts for syslog, see *Configuring syslog on ESXi 5.x and 6.0 (2003322)*, at http://kb.vmware.com/kb/2003322.

If the target environment is licensed for vSphere Host Profiles (which is currently available only as an Enterprise Plus feature), the method of aligning all ESXi hosts to use an external syslog server is likely the preferred choice. Using host profiles allows you to standardize the host configuration throughout the vSphere clusters, and they can also be used for other aspects of host configuration and compliance monitoring for a site. vCenter Server reports on any element of configuration that drifts from its configured value.

Other techniques for applying the `Syslog.global.LogHost` value or modifying the ESXi firewall, such as using scripts or performing a manual alignment, can create risk of configuration drift or incorrect implementation. In addition, hosts provisioned using auto deploy typically do not have sufficient storage to save system logs locally and receive their entire configuration through the host profile and answer file. In those environments, they must be configured to use centralized syslog forwarding and configuration using host profiles as just described.

3.2 Syslog Transport Protocols

By default, syslog uses user datagram protocol (UDP) to transport messages from “client” to “server”. However, there might be specific design requirements that required you to consider other options for transporting messages across the LAN or WAN networks. The following port/protocol combinations are supported for the ingestion of syslog data by vRealize Log Insight.

Table 4. Syslog Transport Protocols

Transport Protocol	Port	Protocol
UDP	514	Syslog
TCP	514	Syslog
SSL over TCP	1514	Syslog-TLS (SSL)
SSL over TCP	6514	Syslog-TLS (SSL)

With UDP, syslog clients can send messages over an IP network without prior communication to set up special transmission channels or data paths. UDP has no handshaking process and, thus, provides no guarantee of delivery of the syslog messages to the syslog target. This protocol is suitable for syslog when error checking and correction are not necessary to meet site requirements. Time-sensitive packets are dropped, which avoids extra processing overhead at the network interface level. This method of handling is preferable to waiting to send delayed packets, which might not be an option in a real-time monitoring system.

However, what if the solution design requires assurance that syslog messages are delivered to the syslog server for compliance or regulatory reasons? If this is the case, you will need to employ a transport protocol that uses reliable and ordered transmission methods and also provides error correction facilities on the data stream at the network interface level, which, for syslog, is the Transmission Control Protocol (TCP).

The network connection between the syslog source and the syslog target server is critical to confirming that events are delivered at the remote destination. A number of factors affecting the underlining network, such as latency, load, and the transport protocol used, can prevent log events from being delivered. With UDP, logs can be dropped by the network, without any easy way of determining why. TCP guarantees that the events are transported from the source to target server. TCP can be employed to overcome unreliable network issues, but it does not confirm that the syslog server itself will capture and keep the event. For example, if the vRealize Log Insight ingest pipeline is backed up, messages will be dropped, even if TCP is used. TCP is still preferred over UDP, however, when transporting events between the syslog source and target server, if traffic is traversing a WAN connection or another less reliable network.

What if a service provider requires syslog data to be encrypted between source and destination to meet solution requirements? After all, the information contained in syslog entries could be highly sensitive and easily be used by a hacker to map the network, uncover what hardware is being operated in the data center, and find vulnerabilities within the infrastructure. For these reasons, syslog also supports the ability to encrypt traffic with Transport Layer Security (TLS), allowing messages to be delivered to the target syslog server using TLS over TCP.

Note TLS is the more secure successor of SSL. When people talk about SSL encryption, they usually mean TLS encryption.

Typically, a solution architect will base design decisions regarding transport protocols on the following requirements and constraints:

- Use TCP when possible, especially between syslog aggregators and syslog servers or any traffic going over a WAN. Understand that TCP does not guarantee that events are not dropped. If the vRealize Log Insight ingest pipeline is backed up, messages will be dropped.
- ESXi supports all protocols. However, if other syslog client devices do not support TCP, then UDP is required. Use encryption when it is a specific service provider requirement to secure syslog data as it transverses the network, particularly over unsecured or WAN links.
- VMware NSX® Manager™ virtual appliance and the VMware NSX API™ centralize the provisioning of logical networking components and manage the connection of virtual machines and storage objects to the networking functions, which also includes syslog transport.

3.3 Configuration of Remote IP Address or FQDN

When configuring syslog to forward events to a remote syslog server, you have the option of specifying the target value as either an IP address or a Fully Qualified Domain Name (FQDN). If you are using FQDN, the process is dependent on Domain Name Service (DNS). The design decision about which method to use will require the architect to balance the requirement for flexibility in the data center against the potential risk of a DNS outage.

If DNS is experiencing problems or is otherwise unavailable, you might not be able to resolve the FQDN of the target syslog server, and events cannot be forwarded. Similarly, if the DNS servers are virtual machines residing on the hosts being configured, the ability to forward syslog messages also relies on the availability of the DNS virtual machines. This design decision will require the architect to balance the requirement for flexibility in the data center against the potential risk of a DNS outage.

Syslog agents, such as `vmsyslogd`, often include some level of DNS caching, making it possible for the syslog component to maintain a certain amount of name resolution even during a service disruption. However, if a dependency on key resources exist, such as the one described above for DNS, highlight it as a design risk. Also, if the source device for DNS is rebooted, or has its management agents restarted, any previously cached name resolution is lost and the syslog client will not be able to resolve the FQDN of the target server.

Despite these risks, the primary advantage of using the host name and DNS over an IP address is the flexibility it provides to change the syslog server application or the IP address associated with it, without the need to update all the syslog source configuration values on all devices that target that syslog server instance.

This flexibility can also be useful in a business continuity and disaster recovery (BCDR) design. The configuration of the remote syslog server location using the FQDN makes it possible to easily adjust DNS records to globally modify the syslog target. However, keep in mind that the `vmsyslogd` caching mechanism means that ESXi hosts will not update their local DNS cache automatically.

3.4 Remote Syslog Design Considerations

When VMware Cloud Providers have multiple data center locations, or their infrastructure is separated geographically, additional design considerations are often required. As highlighted previously, a typical vSphere environment will generate numerous log messages each day, and this amount will increase when issues or problems arise with systems. Being able to troubleshoot and find the root cause of a problem is critical, especially when the hardware or software in question is located in a different part of the country or in a different country from where the administrators and operational teams are working.

When designing a syslog solution that will span two or more physical locations, solution designers must understand the provider's requirements, the constraints (such as bandwidth across WAN links), and the security risks associated with lost syslog data. In addition, in a multi-site syslog architecture, solution designers may want to consider the following:

- The number of source devices sending events
- The expected amount of events generated by each source device
- Event retention, compliance, and regulatory requirements
- Role-base access control (RBAC)
- Security design considerations across firewalls and WAN links
- Operating requirement such as uptime, backup, disaster recovery, and performance

3.4.1 Data Throttling

vSphere 6.x includes fine-grained control over system logs, the location where logs are sent, and for each log, their default size, rotation policy, and logging levels. In a typical vRealize Log Insight deployment, the amount of local syslog data center traffic will not normally cause any issues and need not be throttled (although throttling might be desirable for other reasons such as consistency and efficiency). If you have remote sites with slower links, possibly in geographically dispersed locations, you might need to throttle log data at the source to reduce the amount of syslog traffic traversing the WAN connections.

The amount of information captured in the log files varies, depending on the level setting. When a log level is configured, only messages with that assigned log level and above are captured in the log files. For example, if the log level is set to Info, log messages will include Info, Warning, and Error level messages only. The following table shows log levels that are available in vSphere.

Table 5. Data Log Levels

Log Level Setting	Description
None	Disables logging.
Error	Logging limited to error messages.
Warning	Error messages plus warning messages are logged.
Info	Default setting on ESXi and vCenter Server systems. Errors, warnings, plus informational messages about normal operations are logged. Acceptable for production environments.
Verbose	Can facilitate troubleshooting and debugging. Not recommended for production environments.
Trivia	Extended verbose logging. Provides complete detail, including content of all SOAP messages between client and server. Use for debugging and to facilitate client application development only. Not recommended for production environments.

By default, host logging levels vary from service to service. For example, `vpxa` and the Host Agent are configured to log data at the Verbose level by default, while the `hostd` service has a default log-level setting of Info.

You can verify the logging level in different ways for different host services. For example, the Host Agent and `vpxa` current logging levels can be queried with the following vSphere PowerCLI commands:

```
Get-VMHost | Get-VMHostAdvancedConfiguration -Name "Config.HostAgent.log.level"
```

```
Get-VMHost | Get-VMHostAdvancedConfiguration -Name "Vpx.Vpxa.config.log.level"
```

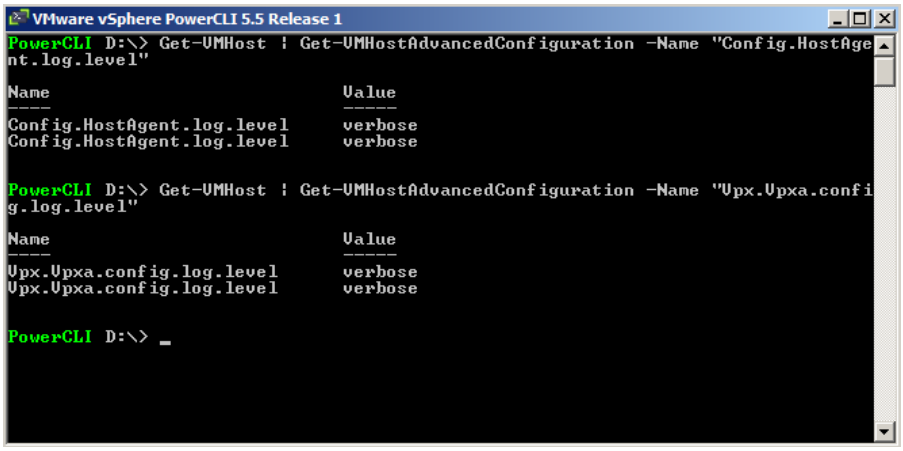

`hostd` logs are controlled by a setting in the config.xml file, located in the `/etc/vmware/hostd` subdirectory of an ESXi system.

```
<log>
  <directory>/var/log/vmware/</directory>
  <level>warning</level>
  <maxFileNum>8</maxFileNum>
  <maxFileSize>524288</maxFileSize>
  <name>hostd</name>
```

The logging level to choose for a design depends on the specific service provider requirements and whether vRealize Log Insight is intended to be used proactively or reactively.

If the syslog data is to be employed only reactively when a problem is detected, configuring the logging level to Warning might be sufficient. However, if you configure only the Error or Warning level, it might be too late to prevent a problem from occurring, and you might not be able to find the root cause of a problem without Info level log information.

If the provider's intention is to use the syslog data proactively, Info level logging is more appropriate as a way to gather lower level information, before a problem arises. It is also possible for messages to be logged at the wrong level, for example, error messages being logged as Info and informational logs being marked as Error. It is also possible that Error and Warning messages might be generated so infrequently that you do not know if logging is working properly.

In contrast, solution designers may also consider factors that might influence a solution to collect less log data. That is, requirements for a solution design might specify that only Error or Warning level logs are forwarded, for any of the following reasons:

- Too much storage space is required to keep the messages.
- It is too expensive to query the information (for example, with a product that charges fees per GB of log data).
- It is too much of a challenge to find the relevant logs, usually due to a lack of a scalable provider based logging solution.
- Logs are used as a monitoring tool only and the provider cares only about error and warning events.
- Unless an error or warning message is seen, logs are not usually or rarely analyzed.

Another aspect to consider for a remote syslog design is the bandwidth available between source and syslog target. When services are logged at the verbose level, a significant amount of data is typically directed to the syslog server. There could be as many as 5000 to 10,000 logs in each 5-minute period, depending on the size of the environment.

By switching the hostagent and vpx levels from verbose to warning, you might see a reduction of approximately 10 to 15 log messages for the same 5-minute time period. When a design is constrained by limited bandwidth between sites, this drop in log volume could provide a significant savings and have an even larger impact for traffic on the wide area network.

The following services can be easily modified to throttle remote logging on an ESXi 6.x host.

Table 6. Modifiable Component Logs

Component Service	Default Configuration	Target Configuration
Config.HostAgent.log.level	verbose	warning
vpx.vpxa.config.log.level	verbose	warning

You can use the Advanced Settings panel in the vSphere Web Client, as shown in the following figure, to modify individual log-level settings.

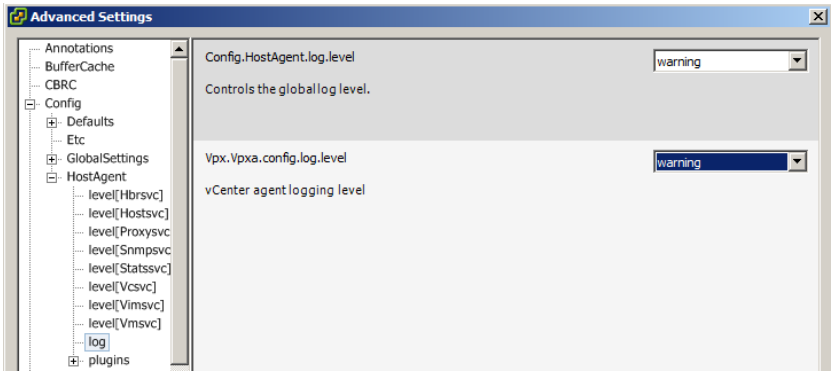

You can also use vSphere PowerCLI with the `Set-AdvancedSetting` cmdlet or host profiles to model log-level settings.

While it is possible to modify the logging levels of other components and services to further throttle syslog messages in remote architectures, changing the default configuration of the following files is not supported. Modify the following component levels only if directed by VMware to do so.

Table 7. Unsupported Log Level Changes

Component	Default Logging Level	Notes
hostd	Info	Modify to warning through config.xml
rhttpproxy	Verbose	Modify to warning through config.xml
fdm	Trivia	Modify to warning through fdm.cfg

For more information relating to modifying ESXi component logging levels, see *Increasing VMware vCenter Server and VMware ESX/ESXi logging levels (1004795)* at http://kb.vmware.com/kb/1004795.

For more information on *Enabling trivia logging in VMware vCenter Server (1001584) at* http://kb.vmware.com/kb/1001584.

3.4.2 Syslog Aggregators and vRealize Log Insight Forwarding

The ability to forward or aggregate syslog data is an important design feature when the architecture requires syslog data to traverse WAN interconnects across sites. This capability also allows you to process or store syslog events, as well as forward events, to another upstream syslog server. Syslog aggregators are typically used in multi data center environments or where message traffic needs to be controlled, for reasons typically relating to security.

vRealize Log Insight 2.5 introduced the ability for a vRealize Log Insight appliance to be used for message forwarding. Any vRealize Log Insight 2.5 or later instance can be employed to forward event messages, whether standalone or clustered. However, even when vRealize Log Insight forwards an event, the vRealize Log Insight instance still ingests and retains those events locally and even archives event messages if it has been configured to do so. In addition, a vRealize Log Insight instance that has been configured to forward events can still be used to run queries for locally stored data. This can be a significant design factor for storage and security when architecting a large scale, multi data center infrastructure for a service provider. There is no way to configure vRealize Log Insight to only forward events and not store them locally.

In environments where this might be an issue, and the provider does not want to store potentially sensitive event message data in specific locations, the two most commonly third-party syslog aggregators are syslog-NG and rsyslog. These daemons allow syslog sources to send them log messages and the aggregator's configuration determines how those messages are forwarded to the central vRealize Log Insight instance. Syslog-NG is available as either an open-source edition (OSE), or a premium edition (PE) at a cost per agent. The open-source edition is sufficient for most implementations. The following site provides a comparison of the two options: http://www.balabit.com/network-security/syslog-ng/comparing/detailed.

Both vRealize Log Insight forwarding and syslog aggregators provide some flexibility and control of how events are forwarded, the format they are forwarded in, the protocol used, and where they are forwarded to. They can also provide design advantages with their ability to limit the configuration changes required when reconfiguring syslog targets. For example, you might need to change only the configuration of the syslog aggregators, as opposed to every source device in the infrastructure, which can be useful from a BCDR design and planning perspective. Syslog aggregators also provide granular control and can be configured to send some or all of the messages that they receive and the messages they generate internally. The following figure shows a typical use case for employing vRealize Log Insight to forward syslog message events in the data center.

Figure 4. vRealize Log Insight Event Forwarder

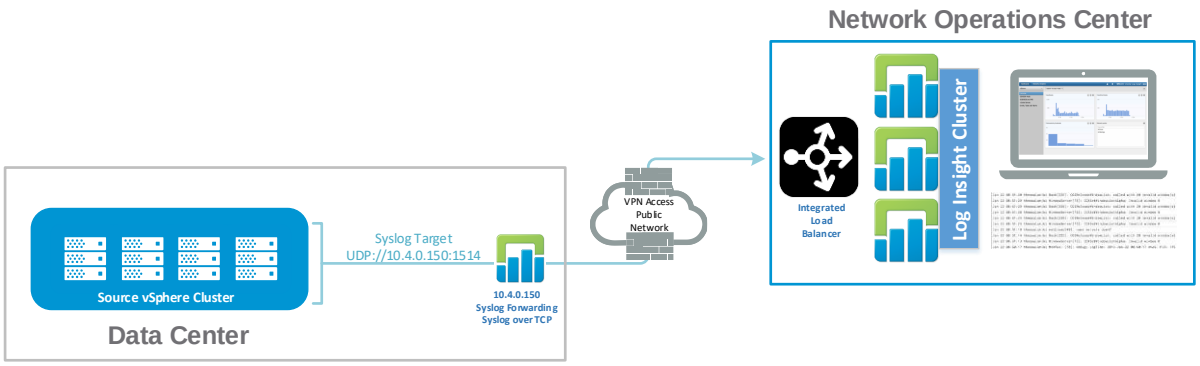

When implementing this type of forwarding solution, configure the syslog aggregators to maintain the original source IP or hostname of the device when forwarding the messages to the centralized syslog server. This is extremely important to provide successful auditing, querying, and analysis of logs.

Finally, syslog aggregators can also be used effectively in a scenario where the syslog target is located on a different network segment (VLAN) than the source host and a firewall sits between the two devices. This allows for significantly more granularity when it comes to configuring a firewall's access control list (ACL). For example, a network security team is likely to be far more comfortable configuring a firewall ruleset that allows traffic from a single source IP address to pass through to a single destination IP address, than allowing a whole network segment to pass traffic over the security appliance on port 514.

The following figure shows a use case where third-party syslog aggregators are employed.

Figure 5. Third-Party Syslog Aggregator

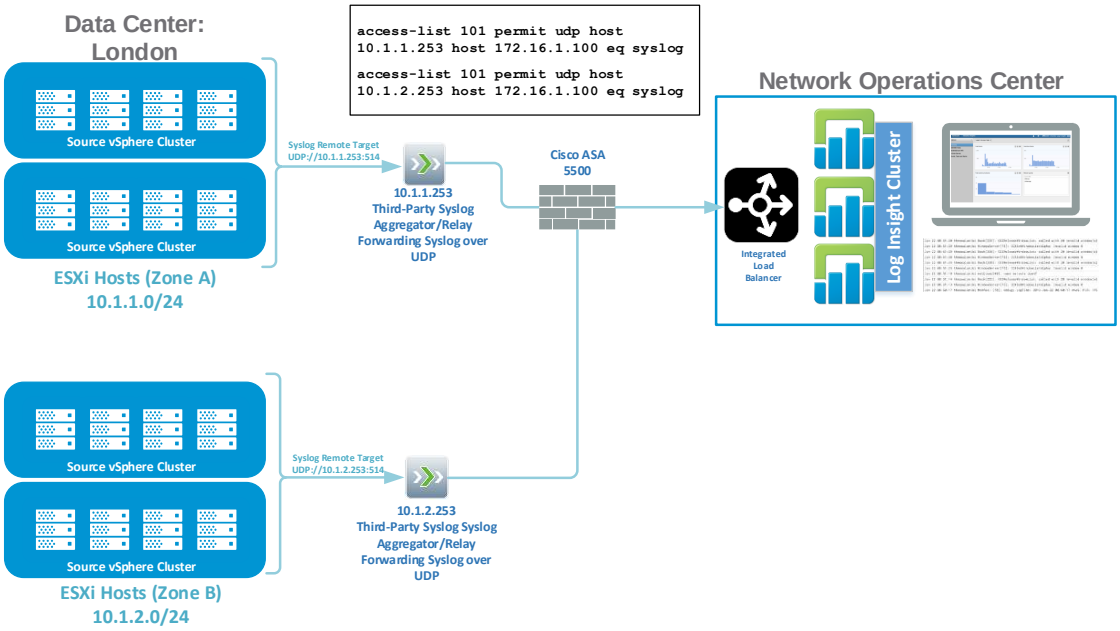

There is one final design consideration for the use of forwarders and syslog aggregators. In a smaller environment, sending syslog traffic directly to a vRealize Log Insight cluster would typically be fine. However, in a multi-site, multi-zone complex service provider environment, VMware recommends that designs employ forwarders in front of the core vRealize Log Insight instances at whatever points are appropriate in the architecture. For instance, forwarders might be deployed in every data center, at a minimum, including the local data center facility where the core vRealize Log Insight instance is located. Also, you might provide an additional level of granularity, depending on the solution design and network architecture, by including additional forwarder instances per network segment, per tenant, or per business unit.

Although syslog forwarders and aggregators increase architecture complexity with their initial setup and configuration (and the potential involvement of more third-party systems), they are often an architectural necessity to meet a service provider's complex mix of challenging requirements and environment circumstances such as network routing, firewalling, or providing a globally distributed architecture.

An additional complication to the solution design is that the VMware Cloud Provider might want such forwarders or aggregators to be highly available. In this case, vRealize Log Insight forwarders would have to be clustered for event message ingestion, providing a highly available configuration, requiring a minimum of three nodes per forwarder instance. When considering the virtual appliance sizing for dedicated forwarding instances of vRealize Log Insight, small deployments are typically sufficient to meet the needs of most environments, because the forwarders themselves are not going to be used as the primary mechanism to store ingested events or to perform operational queries of the environment.

vRealize Log Insight Design Factors

This section describes factors and considerations associated with vRealize Log Insight design.

4.1 NTP Design

A Network Time Protocol (NTP) server provides a precise time source to synchronize the system clocks of all the devices in the global data center design. NTP is transported over UDP and all NTP communications use Universal Time Coordinated (UTC) time. An NTP server receives its time from a reference time source, such as an atomic clock, and then distributes this time across the infrastructure.

Accurate time keeping and time synchronization is critical for a healthy vSphere infrastructure. All components, including ESXi hosts, vCenter Server, vCenter Single Sign-On, SAN and storage arrays, physical network infrastructure, and virtual machine guest operating systems must have consistent and accurate time sources.

vRealize Log Insight provides a consistent method of NTP configuration that synchronizes all the component clocks across the entire infrastructure and, as such, eases the task of reconciling, collating, debugging, and the tracing of log information. The following screen capture shows the vRealize Log Insight default NTP configuration.

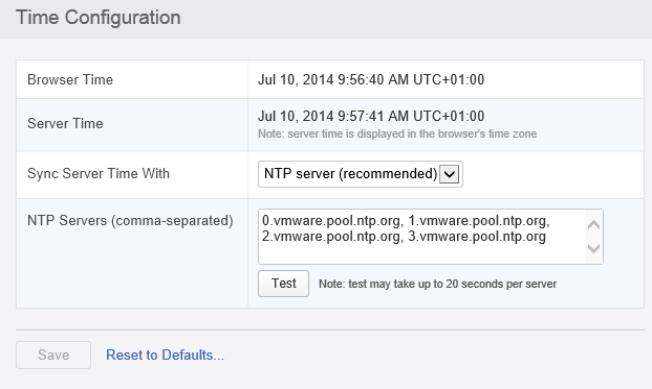

This configuration might require modification, depending on the design approach to NTP. It is important to maintain consistent NTP server configuration across the entire global infrastructure. When configuring the vRealize Log Insight virtual appliance with NTP sources, VMware recommends that you use a minimum of two reference clocks.

For more information on configuring NTP in vSphere environments, see *Configuring syslog logging for hostd and vpxa management agents on ESXi/ESX (1017658)* at http://kb.vmware.com/kb/1017658.

4.2 Clusters

vRealize Log Insight 2.0 introduced the ability to create multi-node application level clusters, providing support for high availability and, by scaling-out, increasing the log message ingestion rates to the scale required by large service provider customers.

When designing a vRealize Log Insight production environment, a minimum of three nodes is required to form an HA cluster. This supports highly available ingestion, the ability to easily scale-out, and the ability to perform maintenance on a live environment. Additionally, vRealize Log Insight clusters provide support for the following:

- A scale-out approach to designs using a centralized logging service.
- Multiple vRealize Log Insight instances working together and managed centrally.

- Log ingestion high availability that can be achieved either through the interpreted load balancer or by placing a third party load balancer in front of the cluster to handle the ingestion of syslog messages, evenly distributing the data across cluster nodes.
- Up to a 12-node cluster with one node as the master and up to eleven other nodes as workers.
- Both master and worker nodes capable of ingesting syslog messages.
- Querying of events from a single dashboard, even if the messages exist on separate nodes.

When designing a service provider management solution, VMware also recommends the use of the integrated load balancer, which allows for distributed ingestion of messages across the cluster nodes and provides high availability for the syslog service. All events are forwarded to the load balancer VIP address, as opposed to the individual vRealize Log Insight nodes. If a vRealize Log Insight node goes down, for example, as a result of host hardware failure, the load balancer will continue to direct events to other nodes within the same cluster. The failed appliance restarts on a different host through the standard vSphere HA mechanism, and there is no service downtime.

If possible, design the vRealize Log Insight cluster to distribute vRealize Log Insight nodes across multiple ESXi hosts in the management cluster to maximize service availability. You can apply a vSphere Distributed Resource Scheduler anti-affinity rule on the group of vRealize Log Insight virtual appliances to increase the availability of the syslog service provided by the cluster. You can also use this rule to specify the relationship between the virtual appliances in the group, so that they remain on separate hosts, thus minimizing the impact caused by a vSphere host restart in a standard vSphere HA environment.

This aligns with the standard VMware recommendation of separating servers and providing redundancy for other critical roles at the application layer to avoid service downtime if a vSphere host fails.

Figure 6. Management Cluster Environment

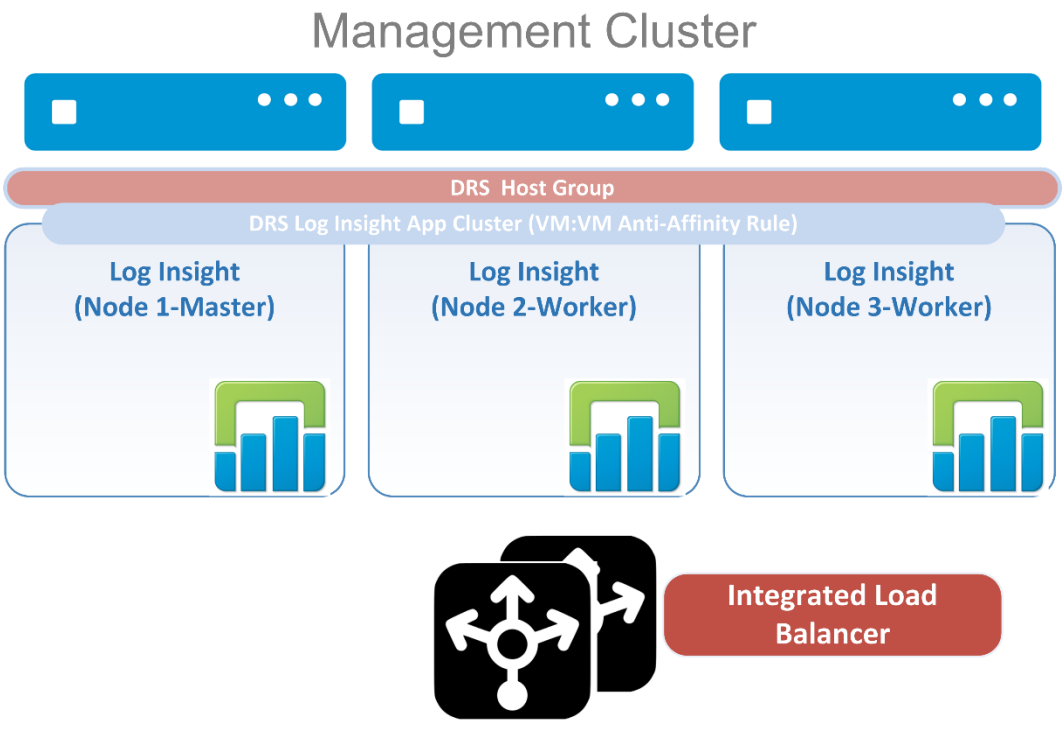

When evaluating the design factors for vRealize Log Insight clusters, keep in mind the following architectural design points:

- A minimum of three nodes is required in a cluster to provide high availability.

- All nodes must be in the same data center, because vRealize Log Insight does not support geo-clustering.
- All cluster nodes must be in the same Layer 2 network segment if an integrated load balancer is employed.
- Employ either the integrated load balancer or use an external load balancer to load balance traffic across nodes in the cluster.

4.3 Cluster Load Balancing

vRealize Log Insight requires a load balancing method when configured as a cluster. Prior to the release of vRealize Log Insight 2.5, the only option for load balancing syslog traffic across the clusters nodes was to use of an external load balancer. vRealize Log Insight 2.5 introduced an integrated load balancer feature that supported Layer 4 and Layer 7 load balancing, in addition to automated failover of the Virtual IP Address (VIP), thus providing good feature parity with external load balancers. The integrated load balancer also helps to simplify the vRealize Log Insight environment and reduces the configuration complexity associated with external load balancers, which in turn can help lower operational costs. The integrated load balancer also provides the following functionality:

- Native L4 integrated load balancer
- Automatic message rebalancing across nodes (L7 Functionality)
- Health check and automatic failover

Detailed configuration for external load balancing syslog message traffic across vRealize Log Insight nodes will vary from device to device and is beyond the scope of this document. The following is an overview of the principles and design considerations required from an architectural perspective:

- If UDP transport is used, a load balancer that supports UDP traffic is required. These are not as common as you might think, and could provide a reason for you to elect to send syslog messages using TCP.
- Your choice of load balancing algorithms depends on the choice of hardware or software appliances employed. To avoid issues caused by an unbalanced cluster node, use the *Least Connections* algorithm, if available.
- Verify that the throughput of the load balancer meets the needs of the planned environment. Some load balancing devices are licensed, based on the maximum MB per second throughput allowed, or on the maximum number of concurrent connections. Verify that the device choice can handle the number of anticipated connections and the required throughput, as well as allow for future growth.
- Avoid allowing “long-lived” TCP sessions in the load balancer configuration. The aim of this recommendation is to keep the load balancer from maintaining open sessions with the target server and to avoid unbalanced cluster nodes, if there is a prolonged outage between the syslog client and syslog server.
- Configure Source NAT (SNAT), if appropriate, in the load balancer configuration, to allow traffic to be forwarded correctly.

As with any design decision, whether to opt for the integrated load balancer, or an external load balancer, will come down to the specific customer use cases and requirements, and any other key design factors. When evaluating the load balancing options, consider the following key points:

- The integrated load balancer service runs on the existing vRealize Log Insight nodes and is simple to configure with only an additional IP address required. There is no additional cost or operational support required.
- The integrated load balancer supports all vRealize Log Insight ingestion protocols, including Layer 4 and 7 load balancing, with no additional configuration required.

- While only a single VIP can be configured with the integrated solution, and all nodes must be on the same Layer 2 network segment, the native load balancer configuration is capable of supporting the maximum ingestion load of the vRealize Log Insight cluster.
- External load balancers, while capable of performing multiple load balancing tasks, are typically managed by a different operational team, which might introduce delays in configuration and operational support.

If you elect to use an external load balancer for the design, you can choose from a range of physical or virtual devices, however, VMware does not provide any specific recommendation. The following list includes some vendor technologies that can provide the required load balancing services:

- NSX Edge and VMware vCloud® Networking and Security Edge™ devices
- Cisco ACE/CSS (these devices are end-of-life)
- F5
- Kemp LoadMaster
- Riverbed
- Barracuda Load Balancer

If you are looking for a free open source solution for a proof-of-concept or lab environment, look at HAproxy/NGINX at http://www.haproxy.org/. If you require UDP load balancing, you can look at the Zen Load Balancer at http://www.zenloadbalancer.com/community/downloads/.

4.4 Ingestion Rates

When architecting a vRealize Log Insight solution to handle log file collection, the design must be able to handle the ingestion rate of all the configured source devices. vRealize Log Insight 3.0 can ingest more than twice the data per second than the vRealize Log Insight 2.5 release, with an ingestion rate of up to 15,000 events per second. With a maximum 12-node cluster, that comes to 2.7 TB of data per day.

Table 8. Ingestion Rates by vRealize Log Insight Cluster Deployment Size

Option	Log Ingest Rate	vCPUs	Memory	IOPS	Syslog Connections	Events per Sec
Extra Small	3 GB/day	2	4 GB	75	20	200
Small	15 GB/day	4	8 GB	500	100	1000
Medium	37.5 GB/day	8	16 GB	1000	250	2500
Large	112.5 GB/day	16	32 GB	1500	750	7500

Note Adding additional vCPU resources to vRealize Log Insight virtual appliances will improve query and ingestion performance.

With vRealize Log Insight 3.0 supporting up to 12 nodes in a cluster, the architecture can scale out to likely meet the needs of the most demanding VMware Cloud Providers that deploy VMware infrastructure across numerous geographically dispersed data centers.

Scaling out and providing high availability are the primary reasons to take a clustered approach to deployment. For example, it might be more appropriate to design a solution that employs three or four small nodes as opposed to a single medium or large node. This method provides the same or higher ingestion rate but has the added benefit of supporting high availability.

Designers should architect a solution to meet the current ingestion rate needs of the infrastructure, while also allowing for growth in the provider's requirements. In addition, always deploy the vRealize Log Insight virtual appliance with thick-provisioned, eager-zeroed disks. This allows for improved performance and operation of the virtual appliance.

4.5 Data Archiving

vRealize Log Insight supports the use of NFS mounts for data archiving, which might be necessary to include in a solution design if the service provider's design has requirements for long-term auditing, compliance, or data retention. The process of data archiving preserves old logs that would otherwise be removed from the vRealize Log Insight appliance due to storage limitations. Where the service consumer's business requirements demand it, plan to include vRealize Log Insight data archiving in the design.

Under normal operating conditions, vRealize Log Insight never runs out of disk space, because it checks the current storage status every minute, and if there is less than three percent of disk space remaining, it retires old "data buckets". If archiving is enabled, vRealize Log Insight archives the data buckets before retiring them.

As part of any vRealize Log Insight design proposal, include archiving calculations that outline the required NFS disk space needed each year to meet the provider's regulatory, compliance or data retention requirement.

vRealize Log Insight itself does not manage the NFS mount used, because this is maintained through the storage management tools typically provided by the service provider's storage vendor. To receive system notifications, you must configure vRealize Log Insight so it can send out an email alert when the NFS mount is getting low on available space. More importantly, if the NFS mount runs out of space or is not available for a period of time that is longer than the retention period of the vRealize Log Insight appliance, data ingestion will stop until the NFS mount has been restored and enough free space becomes available for archiving. Alternatively, you can disable archiving temporarily.

Log events that have been archived are no longer searchable. To search archived logs, you must first import the logs into a vRealize Log Insight instance. For more information on vRealize Log Insight data archiving, see the *VMware vRealize Log Insight Administration Guide* at http://pubs.vmware.com/log-insight-30/topic/com.vmware.ICbase/PDF/log-insight-30-administration-guide.pdf.

Another design factor to be considered is not to import archived data into the existing production vRealize Log Insight instance. This is due to the fact that data from the archive import will force the oldest ingested data to be deleted to make room, which might affect your inability to maintain the desired retention period for events. VMware recommends that you import archives into a dedicated vRealize Log Insight instance that is not ingesting message events and is dedicated to this function.

The following table lists sample log archival storage requirements for general guidance. Actual requirements will vary considerably from environment to environment depending on hardware, software, vSphere features used, and configuration settings.

Table 9. Sample Log Storage Requirements

Number of ESXi Hosts	Logging Level	Daily Storage Requirement	Monthly Archiving Requirement	1-Year Retention	3-Year Data Retention Policy	5-Year Data Retention Policy
1	Default Logging Levels	250 MB	7.8 GB	95 GB	280 GB	475 GB
x4 8-node clusters	Default Logging Levels	7.8 GB	242 GB	3 TB	9 TB	15 TB
1	Throttled: Warning and Error Logs Only	4 MB	124 MB	1.5 GB	4.5 GB	7.5 GB
x4 8-node clusters	Throttled: Warning and Error Logs Only	128 MB	4 GB	47 GB	140 GB	235 GB

After data archiving is configured, it is the operational team's responsibility to verify that the archive destination is cleaned up and maintained, and does not run out of space. vRealize Log Insight does not have any mechanism to manage or monitor the NFS destination, and will continue to attempt to archive data, even after the destination becomes full. Typically, a simple cleanup script can be leveraged to manage this.

Extending vRealize Log Insight Services

vRealize Log Insight can accept, process, and query logs from non VMware sources in addition to those from VMware infrastructure sources. vRealize Log Insight Content Packs are plug-ins with information about syslog messages from external sources such as Brocade switches, EMC or NetApp storage systems, Cisco UCS Systems, Windows Servers, or Windows applications such as Active Directory or Exchange.

Out of the box, vRealize Log Insight includes only the vSphere Content Pack. VMware or third parties can create Content Packs with information about specific events from other syslog sources that can be used by system administrators, operational teams, engineers, and CTOs. The content pack itself is made up of information that includes dashboards, fields, aggregations, alerts, and queries.

For example, if you have a design requirement to monitor syslog data from the provider's Cisco UCS Blade System, you can download the UCS Content Pack from the *VMware Solution Exchan*ge at https://solutionexchange.vmware.com/store/loginsight before importing the VLCP file into vRealize Log Insight. You can then configure UCS Manager to forward all syslog data collected on the fabric interconnects to vRealize Log Insight. The data becomes available for analysis and troubleshooting using the vRealize Log Insight user interface.

Note For the current list of available Content Packs, refer to the *VMware Solution Exchange* at the web address just listed.

To forward logs from Microsoft Windows Servers or Microsoft Windows-based applications such as Active Directory or Exchange, the vRealize Log Insight Windows agent must be installed on each source operating system, allowing messages from Windows event channels and log files to be forwarded to the vRealize Log Insight server. In vRealize Log Insight 3.0, there is limit of 60 Windows event log channels. (See *vRealize Log Insight 3.0 GA Configuration Limits* at http://pubs.vmware.com/log-insight-30/index.jsp?topic=%2Fcom.vmware.log-insight.administration.doc%2FGUID-0601A373-4B74-4B93-8C39-DA85F1D34FD4.html .)

Linux operating systems typically include a syslog agent. If not, you can usually install one. The most common syslog agents found on Linux operating systems are rsyslog and syslog-ng.

For more information about forwarding Microsoft Windows event logs to vRealize Log Insight, see the *VMware Realize Log Insight Administration Guide* available at https://www.vmware.com/support/pubs/log-insight-pubs.html.

vRealize Log Insight Security Design

This section provides an introduction to the security considerations addressing a vRealize Log Insight design for cloud service providers and cloud consumers.

6.1 Role-Based Access Control

Log file messages contain sensitive information about an infrastructure's design that could create a significant vulnerability to an organization's security if compromised. Securing this information and limiting its access to those required to perform specific and authorized job functions is a key part of any syslog system design. As with any service provider based IT platform, do not grant end users with access to the vRealize Log Insight infrastructure or details about its design.

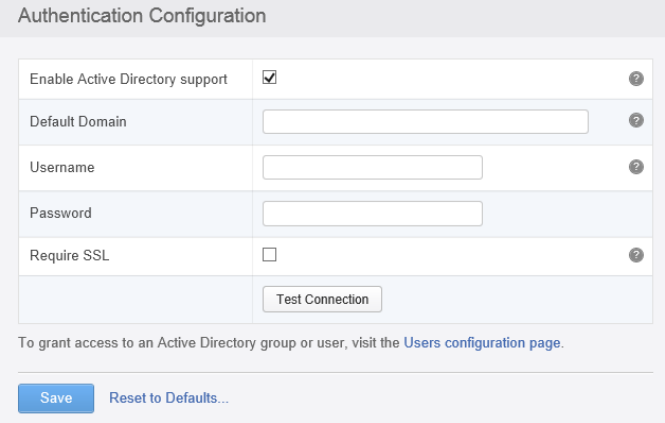

vRealize Log Insight supports the ability to authenticate users using Active Directory instead of the vRealize Log Insight built-in authentication method. This provide auditable role-based access control by way of group membership, and also eliminates the need for administrative users to remember additional user names and passwords.

Note For more information on integrating Active Directory authentication into vRealize Log Insight, see the *VMware vRealize Log Insight Administration Guide* at https://www.vmware.com/support/pubs/log-insight-pubs.html.

The following is a list of best practices to keep in mind when configuring role-based access control with vRealize Log Insight:

- Grant permissions using a privileged account and not a standard login account.
- Assign permissions to Active Directory groups, not individual user accounts. Avoid the use of vRealize Log Insight local accounts.
- Follow the "principle of least privilege". Grant permissions only when needed, and provide only the minimum permissions required to meet a group member's need.
- Create new Active Directory groups for vRealize Log Insight users and administrators. Avoid using built-in Windows groups or other existing groups.
- Use caution when granting root access to vRealize Log Insight.

When employing Active Directory in vRealize Log Insight you are required to provide credentials for a binding user. VMware recommends you use an Active Directory service account for this binding user mechanism. This will mitigate issues such as user credentials expiring or the user credentials becoming locked out, resulting in no Active Directory users being able to log into the vRealize Log Insight user interface.

6.2 Certificates

vRealize Log Insight supports the ability to upload a service provider's own certificate, obtained from a commercial or internal certificate authority, to replace the default VMware self-signed certificate. To replace the certificate, go to the SSL certificate section of the vRealize Log Insight user interface.

For more information on modifying certificates in vRealize Log Insight, see the *VMware vRealize Log Insight Administration Guide* available at https://www.vmware.com/support/pubs/log-insight-pubs.html.

6.3 Port Map

As part of a vRealize Log Insight design, include the specification of communications ports to be opened and the specific components of the architecture that will use them. Careful use of ports allows functioning across firewall boundaries, while also protecting the sensitive information gathered in the log messages from unsecure networks or unauthorized personnel. The best practice is to place the vRealize Log Insight appliances on a separated out-of-band management network segment protected by a firewall from the rest of the internal network. See Section 7, vRealize Log Insight Management Environment for more information on how to do that.

The following tables in this section list recommended communication port assignments. The exact port assignments might vary according to a specific service provider's vRealize Log Insight system design.

The first table lists ports that need to be open to vRealize Log Insight from sources that send syslog message data to vRealize Log Insight.

Table 10. vRealize Log Insight Ports – Source Message Data

Source	Destination	Port	Protocol	Service Description
System sending logs	Log Insight appliance	514/UDP, 514/TCP	Syslog	Syslog data
System sending logs	Log Insight appliance	1514/TCP, 6514/TCP	Syslog-TLS (SSL)	Syslog data over SSL
System sending logs	Log Insight appliance	9000/TCP	vRealize Log Insight Ingestion API	vRealize Log Insight Ingestion API
System sending logs	Log Insight appliance	9543/TCP	vRealize Log Insight Ingestion API (SSL)	vRealize Log Insight Ingestion API - TLS (SSL)

The following ports need to be open to vRealize Log Insight to allow access to the user interface for administrators and operational teams.

Table 11. vRealize Log Insight Ports – User Access

Source	Destination	Port	Protocol	Service Description
User workstation	Log Insight appliance	80/TCP	HTTP	HTTP: Web interface
User workstation	Log Insight appliance	443/TCP	HTTPS	HTTPS: Web interface
Admin workstation	Log Insight appliance	22/TCP	SSH	SSH: Secure Shell connectivity

The following ports are only used for internal cluster communication between vRealize Log Insight master and worker nodes and they are only required to be open if a solution design does not allow for direct Layer 2 communication between the vRealize Log Insight cluster nodes. Normally a solution design would allow direct layer 2 communication between nodes, however this might not be the case, for example, if the vRealize Log Insight cluster nodes exist on separate network segments with a firewall restricting traffic between them. The restriction on Layer 2 communication might also exist if the design employs a distributed software-based firewall solution that restricts traffic between the cluster nodes.

Table 12. vRealize Log Insight Ports – Internal Communication

Source	Destination	Port	Protocol	Service Description
Worker Log Insight appliance	Master Log Insight appliance	16520-16580/TCP	Thrift RPC	Log Insight cluster services
Worker Log Insight appliance	Master Log Insight appliance	59778/TCP	log4j server	Log Insight cluster services
Worker Log Insight appliance	Master Log Insight appliance	12543/TCP	database server	Log Insight cluster services

Other vRealize Log Insight communication ports might be required, depending on other features employed in the infrastructure design.

Table 13. vRealize Log Insight Ports – Additional Communication Ports

Source	Destination	Port	Protocol	Service Description
User workstation	vRealize Log Insight appliance Tomcat service	9006-9007	TCP	Tomcat services
vAPI client applications	vRealize Log Insight appliance	9240	TCP	vRealize Log Insight vAPI service
vRealize Log Insight appliance	vRealize Log Insight appliance	111, 978	TCP, UDP	RPCbind service that converts RPC program numbers into universal addresses
vRealize Log Insight appliance	NTP server	123	UDP	NTPD: Provides NTP time synchronization
vRealize Log Insight appliance	Mail Server	25	TCP	SMTP mail service
vRealize Log Insight appliance	Mail Server	465	TCP	SMTPS: MTP mail service over SSL
vRealize Log Insight appliance	Mail Server	587	TCP	SMPT-MSA: mail submission agent
vRealize Log Insight appliance	DNS server	53	TCP, UDP	DNS
vRealize Log Insight appliance	AD server	389	TCP, UDP	Active Directory
vRealize Log Insight appliance	AD server	636	TCP	Active Directory over SSL
vRealize Log Insight appliance	AD server	3268	TCP	Active Directory global catalog
vRealize Log Insight appliance	AD server	3269	TCP	Active Directory global catalog SSL

vRealize Log Insight Management Environment

VMware strongly recommends separating management components into a logically or physically separate out-of-band management infrastructure, or management pod. The management pod can host the virtual machines that provide management services to the production workloads within the infrastructure. For example, services such as Active Directory, vCenter Server, vCenter Single Sign-On, syslog, storage array management tools, VMware Site Recovery Manager™, VMware vRealize Operations Manager™ or vCloud Director for Service Providers components can all reside in the management pod.

The primary benefits of separating the management components from the production resources they are managing include:

- Quicker troubleshooting and problem resolution, as management components are strictly contained within a relatively small and manageable cluster.
- Resource isolation between workloads running in the production environment and the actual systems used to manage the infrastructure.

Design the management infrastructure with the following host, storage and networking considerations:

- ESXi Local USB/SD boot. Do not use auto deploy or boot from SAN. (It is likely that the management tools and services required to boot will be running from this environment, creating a “circle of dependency”.)
- Highly available vSphere cluster configuration. Shared storage can be provided by Virtual SAN or a physical storage device to facilitate HA, VMware vMotion®, and DRS.
- High availability of virtual and physical network switching components.
- A minimum of N+1 resilience for all physical components.
- Simplified configuration of each component for quicker disaster recovery and improved RTOs.

The following figure shows the management cluster conceptual design.

Figure 7. Typical Management Environment Conceptual Design

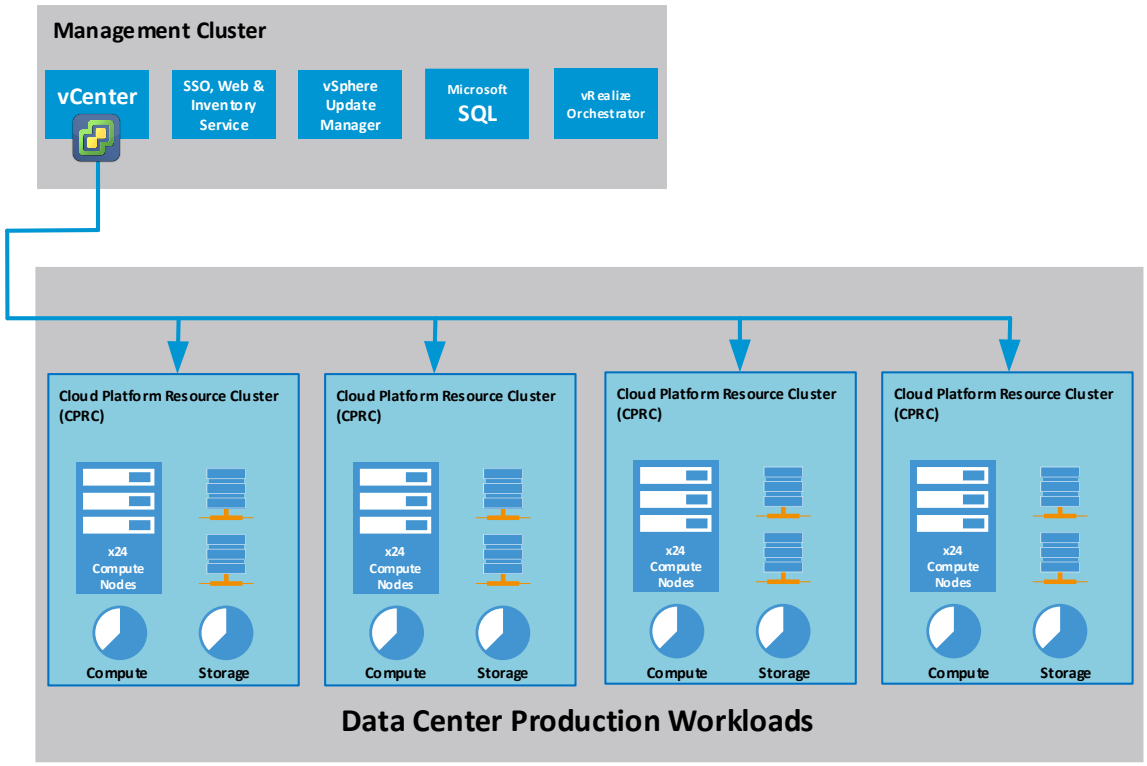

A primary goal of the management environment is simplicity. Designing a simple and static environment minimizes the risks of misconfiguration or human error. Furthermore, limiting the need for change and minimizing the technologies employed further reduces risk and speed up RTOs, should an outage or problem arise.

Running the management components on a large cluster or within a mixed production cluster complicates troubleshooting and makes it more difficult to track down management virtual machines in a recovery scenario. The optimum size of the management cluster will vary depending on the service provider's requirements for management components. VMware recommends a minimum of a three-node cluster to provide sufficient resources to manage the production infrastructure components while maintaining the recommended operational availability of N+1. You can scale out hosts further if the management cluster becomes resource constrained.

Physically separating the management cluster to different racks, hosts, and switches provides a clear demarcation point between the management and production workloads, and helps to prevent an outage in which one environment affects the other. It might also be appropriate to place a firewall between the two environments to enhance security. The ultimate aim of the management pod is that any type of incident or outage affecting the production system cannot impact the management tools, and any type of problem with the management components cannot impact the production workload systems.

7.1 Service Provider Management Design

The syslog system design forms only part of the service provider infrastructure management story. VMware also provides a full set of tools to manage and monitor any size or type of virtual, physical, heterogeneous, or hybrid end-to-end environment.

The following figure shows a VMware provider-based management network operations center (NOC). Although addressing the additional cloud management components in the figure is out of scope for this document, be aware of the relevance of these other management tools and how they integrate with and extend the functionality of a vRealize Log Insight deployment.

Figure 8. Sample Service Provider Management Design

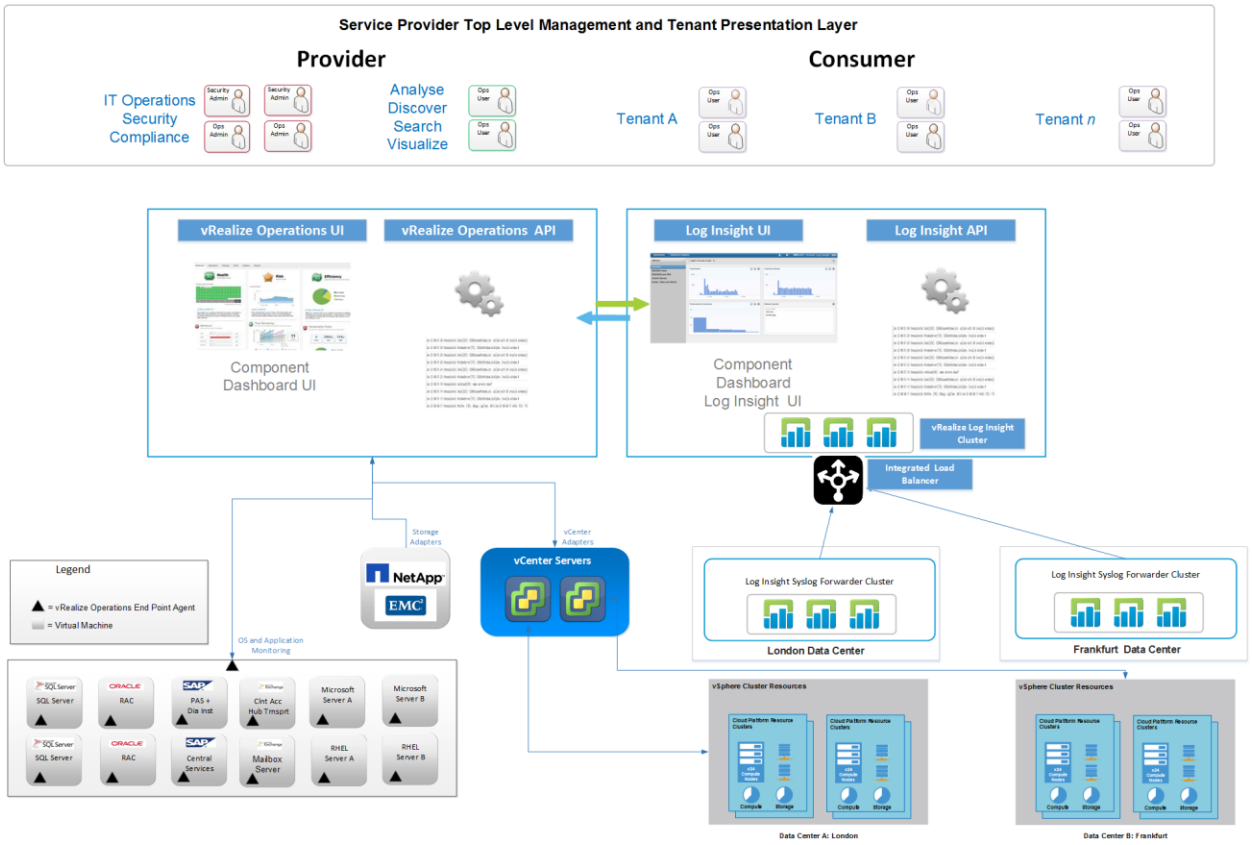

7.2 vCenter Server

vCenter Server and its supporting services are at the heart of the vSphere infrastructure. vCenter Server instances are used to provide a whole range of functionality including:

- Cloning of virtual machines
- Creating templates
- vSphere vMotion and VMware vSphere Storage vMotion
- Initial configuration of DRS and vSphere high availability clusters

vCenter Server also provides monitoring and alerting capabilities for hosts and virtual machines. System administrators can create and apply alarms to all managed objects in vCenter Server. These alarms include:

- Data center, cluster and host health, inventory, and performance
- Datastore health and capacity
- Virtual machine usage, performance, and health
- Virtual network usage and health

These events, tasks, alerts and alarms are collected in vRealize Log Insight as structured data with specific meaning attached to entries in individual fields of the data. In addition, vRealize Log Insight

ingests vCenter Server logs that contain unstructured data which can be queried, aggregated, correlated, and retained for auditing purposes as necessary.

7.3 vRealize Operations Manager

vRealize Operations Manager comes in multiple versions with different feature sets. The goal is to provide performance and capacity management for the vSphere infrastructure, making use of management packs and adaptors, storage, fabric, and third-party management tools such as Microsoft System Center Operations Manager (SCOM) or Oracle Enterprise Manager (OEM).

vRealize Log Insight and vRealize Operations Manager are integrated in the following ways:

- The launch of the context menu in vRealize Operations Manager can display actions related to vRealize Log Insight.
- vRealize Log Insight can send notification events to vRealize Operations Manager.

While is it possible to integrate multiple vRealize Log Insight instances to send alerts to a single vRealize Operations Manager, the Launch in Context feature only works in a one-to-one scenario.

7.4 vRealize Operations Manager 6.1 Agents

vRealize Operations Manager agents provide application monitoring and performance management for physical, virtual, and cloud environments. These agents, along with the syslog operating system agents, can be configured to send syslog data to Log Insight. This provides a comprehensive monitoring solution for both operating systems and applications, allowing for application owners and OS operational teams to query, analyze, and audit log data.

Cloud Services Syslog Management

VMware Cloud Providers offer a wide variety of different service offerings to their customers. This extensive list of potential offerings changes regularly in the fast moving cloud marketplace and the market will continually develop and evolve to meet the needs of a whole range of industries. For example, one service provider might specialize in providing services exclusively designed for the gaming industry or financial services, while another might target secure government customers for their service offerings. In addition, as well as targeting specific industries, VMware Cloud Providers are likely to offer a wide range of tier levels of cloud management services ranging from simply providing floor-space and power in a data center (perhaps in a dedicated cage), to a fully managed cloud service with a dedicated operational team for specific customers. The line of demarcation between provider and consumer will likely lay somewhere between these two examples.

That being the case, what if there is a design requirement for a specific service offering to provide access to host and other device syslog information directly to the service consumer for troubleshooting, security, or simply informational purposes? The specific design required to meet this architectural requirement would depend very much on the design factors for the explicit use case. For example, an architect or solution designer would need to consider:

- Does the customer have a dedicated hardware infrastructure?
- Does the customer have a dedicated vCenter Server?
- Does the architecture also include a requirement for vRealize Operations Manager?

Based on questions like these, it is possible to conceive of a number of different design scenarios. The remainder of this section details two possible scenarios for which a service provider delivers services.

Scenario A employs the built-in functionality provided by most devices, that is, ability to send syslog messages to two or more independent targets. In this example, ESXi hosts are configured to send syslog messages to two separate instances, consumer and provider, with the consumer instance possibly residing within the consumer's organizational virtual data center. Note that some devices, such as VMware ESXi 4.x, support only a single syslog target.

Figure 9. Multitenant Scenario A

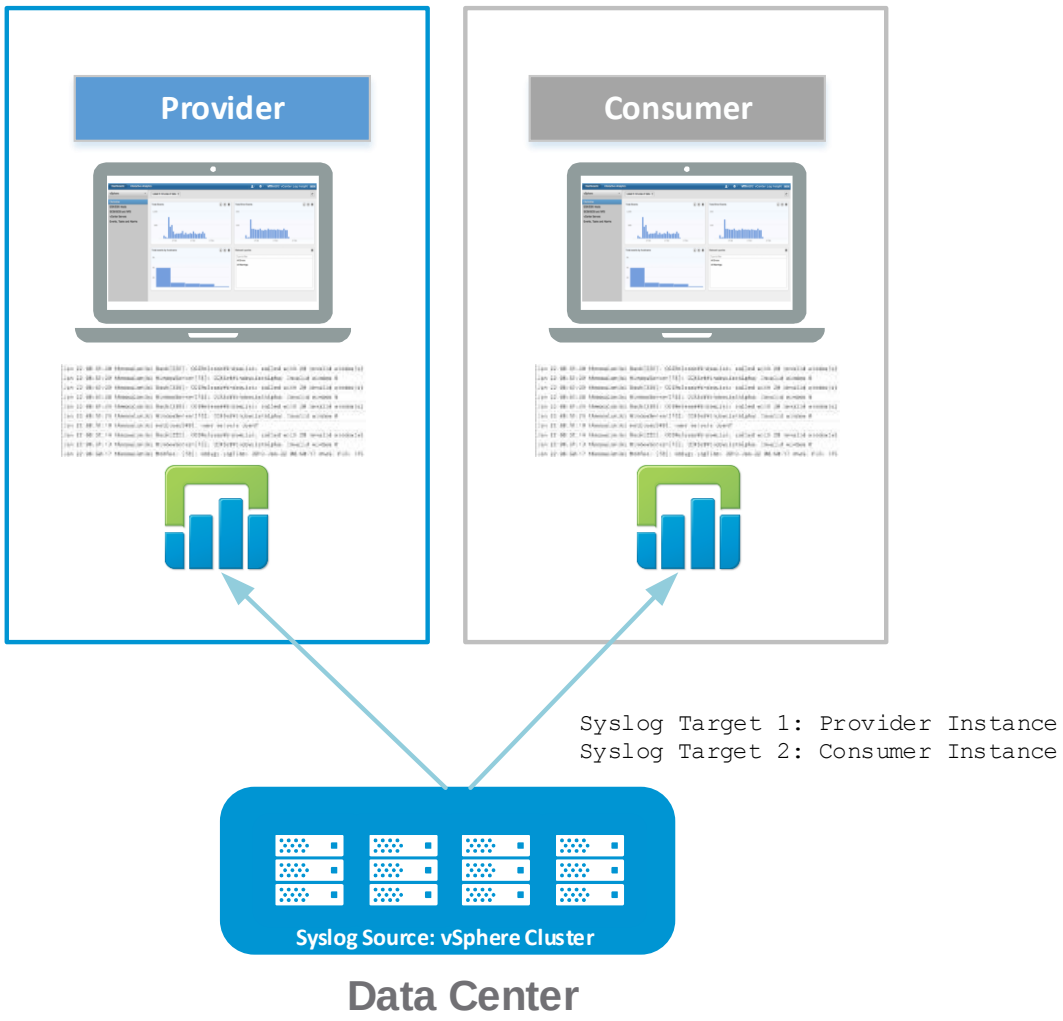

Scenario B vRealize Log Insight Forwards are being employed to forward syslog event messages to the secondary consumer instance. This has the advantage of allowing the provider to filter messages before they are forwarded.

Figure 10. Multitenant Scenario B

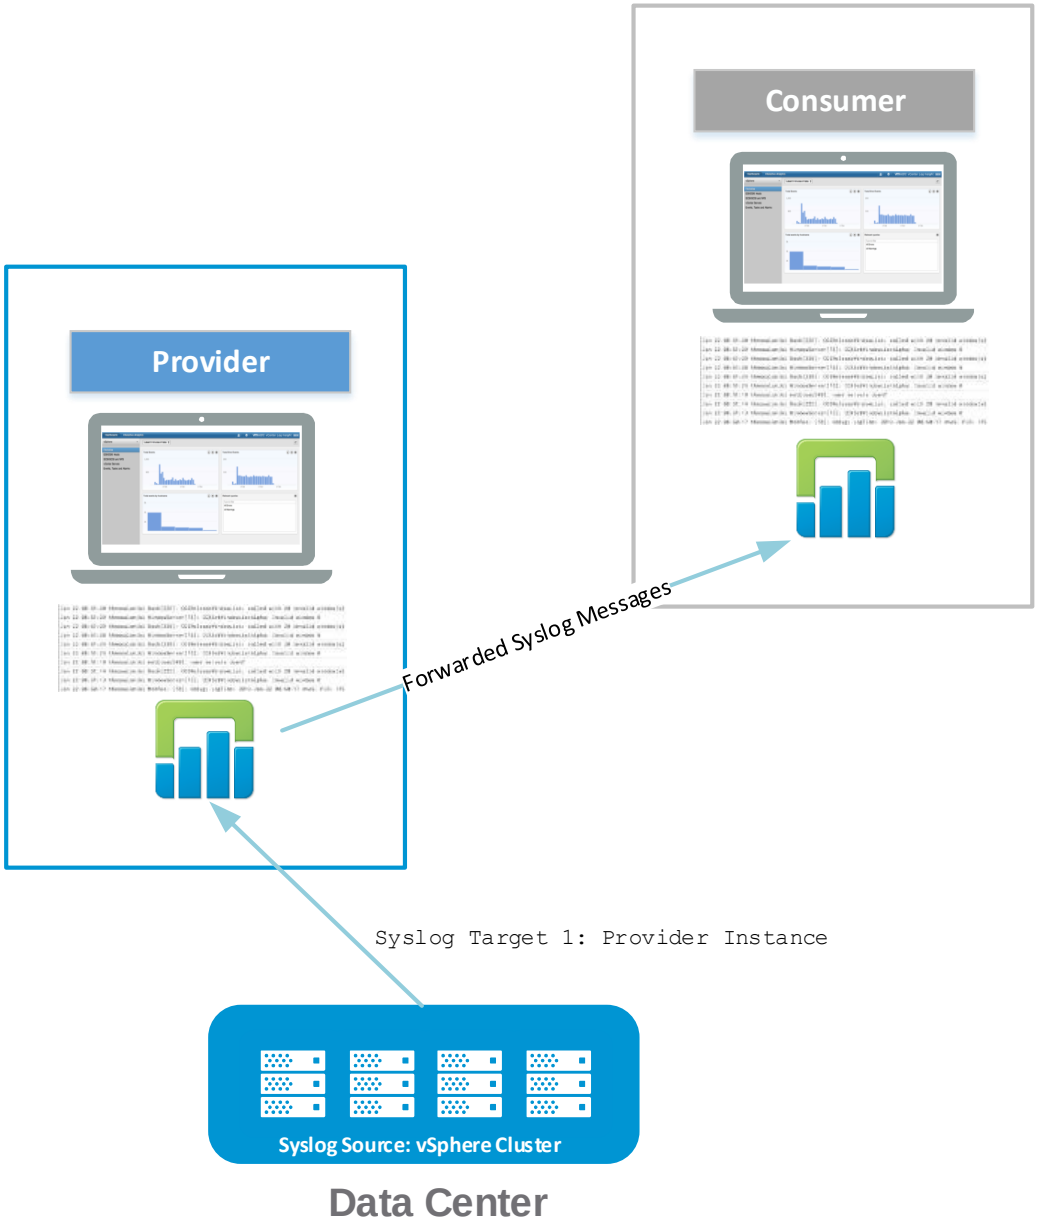

In addition to these two scenarios, vRealize Log Insight supports multi-tenant logging, in which you can control who sees what data. Through the built-in role-based access control (RBAC) mechanism, it is possible to control access to data sets, dashboards, interactive analytics, and administrative tasks. By configuring data sets, a static filter can be assigned to roles. Defining one or more static filters to restrict data helps provide the user interface restrictions that control what people can see and do in vRealize Log Insight.

vCloud Platform Component Logging

As with any cloud infrastructure, a VMware Cloud Provider Program platform is made up of a number of integrated components. As previously highlighted, vRealize Log Insight is able to analyze data from a wide variety of sources out-of-the box through its content packs. It provides a centralized means of monitoring an entire cloud stack made up from components not only from VMware, but also from a wide variety of other vendors.

This section addresses the syslog requirements for a few of the most common components used within the VMware Cloud Provider Program platform cloud stack vCloud Director for Service Providers and NSX for vSphere.

9.1 vCloud Director for Service Providers

9.1.1 vCloud Director Cell Logs

The log location of vCloud Director for Service Provider components is described in the following table. In addition to these application logs, redirect the operating system logs to the centralized syslog as well, through the vRealize Log Insight Linux agent.

Table 14. VMware vRealize Operations Manager

Log Type	Location	Collection Method

vCloud Director Info and debug logs	`%VCLOUD%/logs/*`	Periodic retrieval with
		`%VCLOUD%/logs/vcloudcontainer-debug.log`
		and
		`%VCLOUD%/logs/vcloud-container-info.log.`
		Most logs rotate with .1, .2, … extensions.
		While not for general use, these logs can be redirected to syslog by modification of the following properties file:
		`%VCLOUD%/etc/log4j.properties`
		http://kb.vmware.com/kb/2004564

vCloud Director syslog events	Sent through Syslog	Set syslog targets in
		`%VCLOUD%/etc/global.properties`
		and
		`%VCLOUD%/etc/responses.properties`

vCloud Director system logs	Standard Linux log locations: `/var/log/messages and /var/log/secure`	Syslog targets through `syslog.conf` or agent retrieval

To configure vCloud Director to forward log files to vRealize Log Insight

1. SSH into the vCloud Cells machine and open the following file to make setting changes:

   `/opt/vmware/vcloud-director/etc/log4j.properties`

2. Near the bottom of the file, append the following code block. For the 10.10.10.250 entry on the second line, substitute the actual IP or hostname of the vRealize Log Insight cluster VIP.

```
log4j.appender.vcloud.system.syslog=org.apache.log4j.net.SyslogAppender
log4j.appender.vcloud.system.syslog.syslogHost=10.10.10.250
#Logs go to port 514 unless you specify a port, as in the disable example below.
#log4j.appender.vcloud.system.syslog.syslogHost=remoteSyslogHost.example.com:5555
log4j.appender.vcloud.system.syslog.facility=LOCAL1
log4j.appender.vcloud.system.syslog.layout=com.vmware.vcloud.logging.CustomPatternL
ayout
log4j.appender.vcloud.system.syslog.layout.ConversionPattern=%d{ISO8601} | %-8.8p |
%-25.50t | %-30.50c{1} | %m | %x%n
log4j.appender.vcloud.system.syslog.threshold=INFO
```

3. In the code block, you can see that you can also control the level of logging being sent through to vRealize Log Insight by editing that last line and changing the threshold to **WARN** or **CRITICAL**.

4. Once that section has been added, go back to the top of the file and modify line 2 as shown here:

```
#Root logger
log4j.rootLogger=ERROR, vcloud.system.debug, vcloud.system.info
log4j.rootLogger=ERROR, vcloud.system.debug, vcloud.system.info,
vcloud.system.syslog
```

   These entries add the source configured in step 2.

5. Save the file and restart the `vmware-vcd` service.

In examining source logs, log entries related to a specific customer can be identified using their organization id, that is part of vCloud Director log entry contains: `currentContext.login.org.id=###`

The following code block provides an example of the type of information included in a typical log file.

```
Jul 05 11:53:56 Event [id=1e523973-0620-4dc3-8ffc-865c0d4c61c1,
timestamp=1341442436853, type=com/vmware/vcloud/event/task/complete, properties={

currentContext.org.name=ACME,

currentContext.user.id=acmeadmin(com.vmware.vcloud.entity.user:3ec7999f-d732-436c-
bfd4-4919228c5146),

entity.type=com.vmware.vcloud.entity.task,

currentContext.login.user.id=com.vmware.vcloud.entity.user:3ec7999f-d732-436c-bfd4-
4919228c5146,

currentContext.user.name=acmeadmin,

currentContext.login.member.id=acmeadmin(com.vmware.vcloud.entity.user:0c9ec4a6-
351d-40da-94a8-0209b73393cd), task.ownerType=com.vmware.vcloud.entity.vapp,

currentContext.user.clientIpAddress=,

task.name=VAPP_DEPLOY,

entity.name=VAPP_DEPLOY,

entity.id=VAPP_DEPLOY(com.vmware.vcloud.entity.task:f74b4e85-2608-4af6-b710-
30d62464a16f),

currentContext.cell.uuid=5252de5b-0a3c-4e75-822f-cbd9604b18f0,

currentContext.user.proxyAddress=,



currentContext.o...<13>...rg.id=ACME(com.vmware.vcloud.entity.org:3c8970fb-3b5b-
482a-8056-b229744ad6f8),

task.ownerId=659e986c-21d5-49ed-b718-8a8d38e99811,

currentContext.success=true,

currentContext.login.org.id=com.vmware.vcloud.entity.org:3c8970fb-3b5b-482a-8056-
b229744ad6f8,
```

9.2 NSX Manager Logs

Sitting alongside vCloud Director providing network and security services, VMware NSX for vSphere log components are described in following table.

Table 15. vCloud Networking and Security Manager Logs

Log Type	Collection Method
NSX Manager Events	Remote syslog
NSX Controller	The only supported method for configuring the syslog server on an NSX controller is through the VMware NSX API. See Configuring syslog server for VMware NSX for vSphere 6.x controllers (2092228) at http://kb.vmware.com/kb/2092228.
Tech Support Log	API
NSX Edge gateway logs	See the following section

9.2.1 NSX Edge Gateway Logs

When managed through vCloud Directory for service providers, the IP addresses of syslog servers for NSX Edge gateways are configured centrally through vCloud Director. (This is the same for all edges that might be NSX Edge gateways or edges on virtual application networks.) One of the external NSX Edge gateway interfaces is connected to the syslog network that routes to the syslog servers.

Table 16. vCloud Director for Service Providers Component Logging

	Provider Logs	Tenant Logs
vCloud Director cells	*	*
vCenter Server instances	*	
ESXi hosts	*	
vRealize Orchestrator	*	
NSX Manager	*	
NSX Edge	*	*

A tenant can also deploy their own vRealize Log Insight appliance to a dedicated Org in a virtual data center network. The subnet and IP address of the tenant's syslog server must match the secondary syslog address of NSX Edge gateways. This allows the tenant to see the NSX edge logs in real time, which is very useful for troubleshooting.

Figure 11. Edge Gateway Provider and Tenant Syslogs

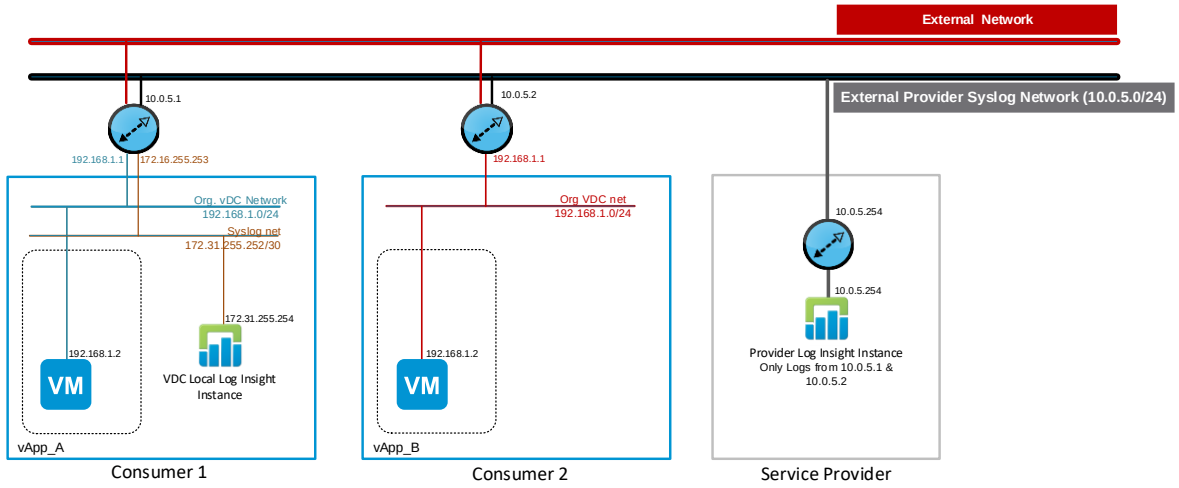

Table 17. NSX Edge Logs

Log Type	Collection Method
NSX Edge rule events. With the log box checked, the logs described here are interactions with firewall rules.	Remote or optional tenant syslog

The NSX Edge log has the following format:

```
<Date-Time> <NSXEdge-ID> <program/daemon-Name>[‘[‘PID’]’:][‘[‘Tenant/Organization-ID’]’] : [EdgeService/Action-Identifier] <Message>
```

The following code block provides an example log entry from NSX Edge.

```
Nov 26 12:29:42 vse-d61f13c3-b56e-42b7-9de3-f55e07f5a19a-0 firewall[]: [3c8970fb-3b5b-482a-8056-b229744ad6f8]: ACCEPT_1IN= OUT=vNic_0 SRC=10.0.2.71 DST=77.75.72.3 LEN=84 TOS=0x00 PREC=0x00 TTL=63 ID=0 DF PROTO=ICMP TYPE=8 CODE=0 ID=23042 SEQ=135 MARK=0x2
```

The vRealize Log Insight server adds the IP address of the syslog sender, which is the NSX Edge gateway network logging external interface.

The following table provides an overview of the format of NSX Edge Logs. For example, filtering on **Organization-ID** allows creation of dashboards specific to individual tenants.

Table 18. NSX Edge Log Format

Field Name	Description
Date-Time	Date and Time in format: Month Day HH:MM:SS, for example: Sep 5 10:50: 56
NSXEdge-ID	Unique Edge ID provided by NSX Manager (not vCloud Director).
Program/Daemon-Name	Name of the daemon or program that is logging this message (for example: DHCP, kernel, pluto, nginx)
PID	PID of the program logging this message. This is optional, especially for kernel and iptables related messages. PID is not logged.
Tenant/Organization-ID	This is the organization identifier.
EdgeService/Action-Identifier	This is optional. For some log messages (where daemon name does not specify the EdgeService uniquely), service identifier is prefixed to the log message (for example: DNAT, SNAT, Firewall-policyapplied-to-rule=ACCEPT|DROP).
Message	This is the actual log message.

Sample Syslog Design Scenarios

The design scenarios described in this section illustrate just a small number of possible deployment configurations and options. Use them for reference and guidance, keeping in mind that the design must meet the requirements of the specific VMware Cloud Provider or project.

10.1 Design Scenario A

This is a simple single-site deployment in which vRealize Log Insight is located on the same network segment as the source hosts being monitored. Each syslog source is configured to send events directly to a single vRealize Log Insight instance using UDP.

Figure 12. Design Scenario A

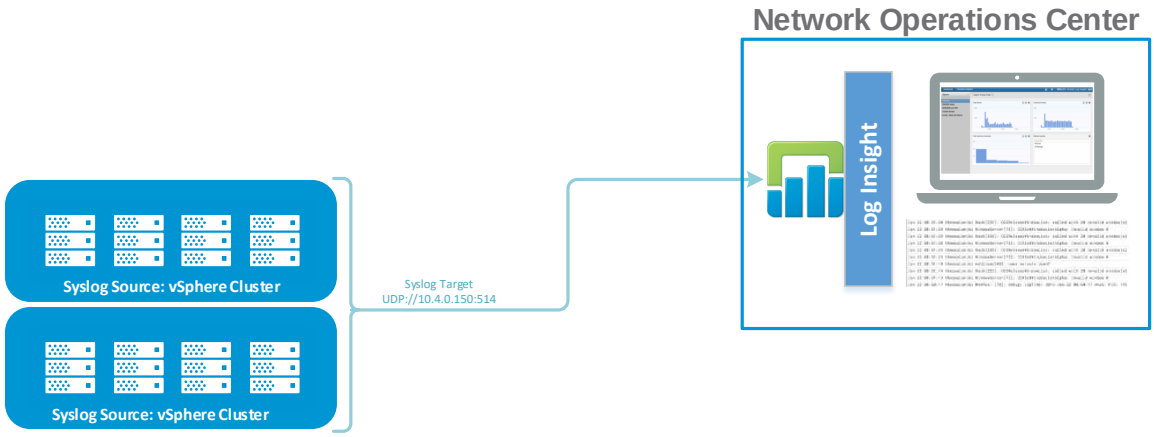

The following information provides specific comments about this scenario and accompanying solution design.

Design Quality	Architect Notes
Simple Configuration	Dependent only on the syslog agent, local network connectivity and vRealize Log Insight VA. There are no obvious drawbacks. However, this design is clearly limited in scope and would not meet the needs of most service provider level customers.
Syslog transport protocol: UDP	UDP is the most efficient protocol for local network syslog traffic most often used when there is no specific requirement to verify data packet delivery.
Single vRealize Log Insight instance	Sufficient for this design, as the number of source hosts will not go above the ingestion limit of 750. No vRealize Log Insight redundancy is provided other than standard vSphere HA protection, protecting against host failure. Another solution would be to provide a vRealize Log Insight cluster made up of two or more smaller instances. This would have the added advantage of providing an increase in syslog application availability.

10.2 Design Scenario B

This design scenario features a redundant NOC architecture. For legal compliance reasons, the provider has a requirement for log messages to be sent to and be stored at two different geographical locations. The data center is located at the same site at the primary NOC (London), where the out-of-band management infrastructure, including vRealize Log Insight and vRealize Operations Manager is located. The second NOC is located in Paris where an additional vRealize Log Insight cluster has been provisioned.

The data center houses a single syslog aggregator to forward logs to the remote secondary NOC in Paris. This requires that logs that transverse a public network be secured using SSL.

The service provider has 1350 source devices logging messages and requires two large vRealize Log Insight cluster nodes to meet the ingestion requirements. To allow for a single host failure without causing ingest pipeline congestion on the remaining instances, three nodes will need to be configured to meet the provider's requirement for syslog application high availability.

Some devices, such as VMware ESXi 4.x, support only a single syslog target. Depending on the provider's requirements, this might create a new challenge for the design, because syslog messages are dropped if the single syslog target becomes unavailable. One solution to this problem is to utilize a highly available load balancer pair configured with a VIP between two or more syslog servers. This way, if one syslog server is unavailable, the other syslog server can still ingest and process the events.

In this design, each of the two NOCs employs an external load balancer to distribute log messages evenly across the cluster nodes.

Figure 13. Design Scenario B

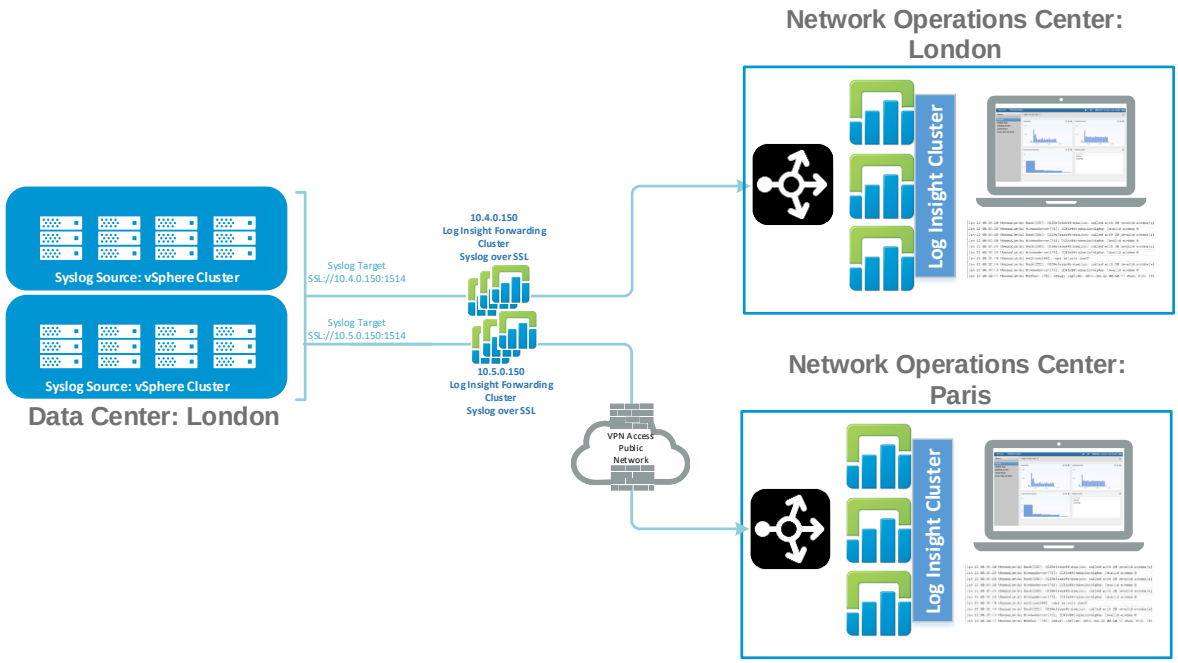

The following information provides some specific comments about this scenario and accompanying solution design.

Design Quality	Architect Notes
Redundant vRealize Log Insight appliances.	Provides a HA solution along with meeting the source ingest requirement.
Hosts are to be configured to forward syslog data to two target addresses.	`esxcli system syslog config set --loghost= 'tcp://10.4.0.150:514,tcp://10.5.0.150:514'`
Local syslog data is being transported over TCP to the local NOC.	TCP is the protocol of choice where the service provider requires assurance of message datagram delivery.
A syslog aggregator is used to convert TCP messages to SSL and then forward them across a secure WAN link to the remote NOC.	The use of a syslog aggregator in this design limits source traffic across the public network to a single address and provides a local target for client syslog configuration. In this design, traffic is also being repackaged from TCP to SSL/TCP to meet the providers' security requirements.

10.3 Design Scenario C

This is a highly secure design where the NOC is located in London alongside one of the data centers. The two other primary and secondary data centers are located at remote geographically dispersed sites. In this design, syslog data must be secured end-to-end between source and destination. Sensitive syslog messages are secured on the local network and MPLS network using SSL, and secured across the public Internet using a VPN tunnel connection.

The three data centers send log messages to a three-node vRealize Log Insight cluster that is employed as part of a full VMware service provider management infrastructure including vRealize Operations Manager components.

The VMware Cloud Provider has more than 1500 devices logging and requires a two-node vRealize Log Insight cluster to meet the ingestion requirements of the infrastructure. To allow for host failure without causing interruption to log message ingestion, a three-node vRealize Log Insight cluster is configured.

Each data center houses a single syslog aggregator to forward logs to the central NOC in London.

The external load balancing is provided by a NSX Edge HA pair device for syslog messages to be distributed evenly across the vRealize Log Insight nodes.

Figure 14. Design Scenario C

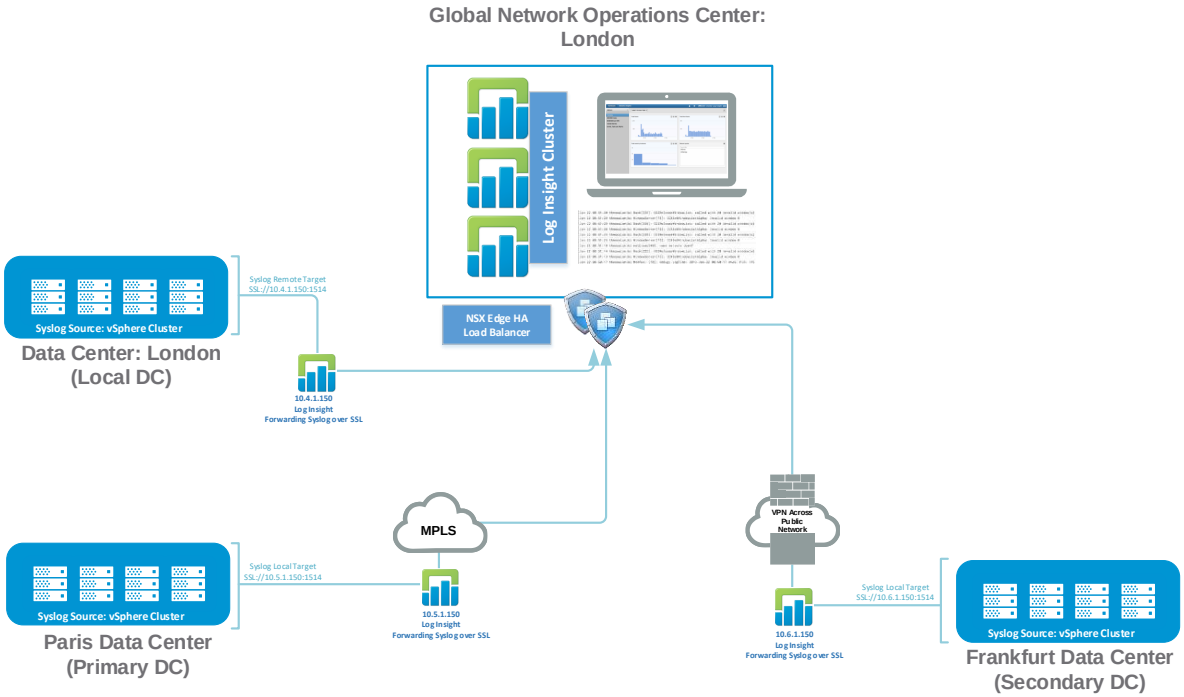

The following information provides some specific comments about this scenario and accompanying solution design.

Design Quality	Architect Notes
Redundant vRealize Log Insight appliances	Provides a HA solution along with meeting the source ingest requirement.
Secure Design	Meets a provider's requirement that all syslog data is secure while in transit across private and public networks.
Syslog aggregator or forwarder	The use of a syslog aggregator in this design limits source traffic across the public network to a single address and provides a local target for client syslog configuration.
Syslog transport protocol: SSL	SSL is employed to secure sensitive log data on internal networks and across a MPLS link. Connection to the secondary data center is through a VPN tunnel created, in this case, between two Cisco ASA 5500 firewall devices.
HA Load Balancer	HA Pair of NSX Edge devices configured to load balance traffic evenly across Log Insight nodes. Load balancing algorithm configured as least connections: "LEAST_CONN" on TCP 1514

Assumptions and Caveats

VMware, Microsoft and other third-party hardware and software information provided in this document is based on the current performance estimates and feature capabilities of the versions of code indicated. These are subject to change by their respective vendors.

Reference Documents

The following vRealize documents are available for additional information on vRealize Log Insight:

Document Title	Link or URL
VMware vRealize Log Insight Documentation	https://www.vmware.com/support/pubs/log-insight-pubs.html
VMware vRealize Log Insight 3.0 Release Notes	http://pubs.vmware.com/Release_Notes/en/LogInsight/3.0/log-insight-30-release-notes.html
VMware vRealize Log Insight Support	https://communities.vmware.com/community/vmtn/vcenter/vcenter-log-insight
VMware vRealize Log Insight Marketplace	https://solutionexchange.vmware.com/store/loginsight
The VMware vRealize Log Insight Community	http://loginsight.vmware.com/
Excellent external resources maintained by Steve Flanders	http://sflanders.net