

# Kofax Equitrac

Print Submission High Availability for DRE

Version: 6.1.0

Date: 2020-06-09

**KOFAX**

©2020 Kofax. All rights reserved.

Kofax is a trademark of Kofax, Inc., registered in the U.S. and/or other countries. All other trademarks are the property of their respective owners. No part of this publication may be reproduced, stored, or transmitted in any form without the prior written permission of Kofax.

# Table of Contents

Equitrac Print Submission High Availability for DRE.....	4
Overview.....	4
What is high availability?.....	4
Print submission high availability.....	4
High availability for DRE options.....	5
Hyper-V cluster.....	5
Create a highly available VM on a Hyper-V cluster.....	6
Network load balancer.....	6
Client and server network printing protocols.....	7
Client workstation traffic flow.....	8
DRE considerations.....	8
Layer 4 load balancing.....	8
Add the IP address variable.....	9
Add environment variable.....	10
Load balancing IPP print queues.....	10
Load balancing SMB print queues.....	11
Load balancing LPD print queues.....	12

# Equitrac Print Submission High Availability for DRE

## Overview

This guide details print submission high availability for Kofax Equitrac® and highlights the various options to setting up highly available printer servers via Equitrac DRE. High availability systems can be set up using Windows failover clustering, Windows Hyper-V clustering or Network Load Balancing (NLB).

This guide assumes that you have already created, configured and tested your Windows cluster or Network Load Balancer (NLB) appliance. For information on creating clusters, or setting up NLB, refer to your Microsoft documentation.

## What is high availability?

High Availability (HA) print environments are designed to provide full-time system availability. HA systems typically have redundant hardware and software that makes the system available in the event of a failure, and help distribute workload. To ensure that HA systems avoid having single points-of-failure, any hardware or software component that can fail has a redundant component of the same type.

Failures can be system faults, data errors and site outages. System faults can occur in the system itself (CPU, memory or power supply), the software (operating system, database or application), peripheral device error, power outage or the network disconnection. When failures occur, the processes performed by the failed component are moved (or failed over) to the redundant or backup component. This process resets system-wide resources, recovers partial or failed transactions, and restores the system to normal as quickly and as seamlessly as possible. A highly available system is almost transparent to the users.

## Print submission high availability

Print submission high availability provides server level protection to the print server and print spooler. The high availability solution in Equitrac ensures that there is always a print queue available for end users. When a user submits a print request to the print queue, the request is intercepted by DRE, and DRE forwards to request onto CAS for validation. The validated print request is sent back to DRE, and the job is sent to the printer.

In a highly available environment the print request workflow is uninterrupted by having a redundant or backup DRE available in the event of a failure, and to distribute workload.

## High availability for DRE options

High availability for DRE can be set up using failover clustering or Hyper-V clustering for Server 2012 R2, 2016 and 2019 or Network Load Balancing (NLB).

Clustering uses multiple servers, as though they were a single system. In clustering, one resource runs, and another redundant resource is available to step in if the first one becomes unavailable. The purpose of HA clustering is to reduce the consequences of hardware, software, or network outages.

Network Load balancing uses multiple independent servers that have the same setup, but do not work together as in a cluster setup. The network load balancer (NLB) forward requests to either one server or another, but one server does not use the other server's resources. Also, one resource does not share its state with other resources. In an NLB setup, all resources run at the same time, and a management layer distributes the workload across them.

## Hyper-V cluster

Windows Server 2012 R2, 2016 and 2019 use a high availability solution with virtual print servers within a clustered Windows Hyper-V environment. Hyper-V clustering is a combination of server virtualization and failover clustering.

Failover clustering with virtual machines (VM) provides a very highly available solution, providing fast and automated recovery, access to shared data and almost no noticeable downtime. It offers high availability for planned downtime, such as servicing, maintenance and upgrades, and for unplanned downtime, such as power outages, system failure or network disconnection.

By separating applications within each VM, the risk of a failed component impacting other services is reduced, and the number of components that need to be patched decreases, further reducing downtime from system reboots. Since it is easy to clone a VM, this simplifies deployment, minimizes configuration error and provides rapid scalability. Virtualization allows a VM to be moved to a different virtual host server with no downtime for the application running within the VM, so the host server can be patched or updated.

In Server 2012 R2, 2016 and 2019, the print spooler service is no longer a clustered resource and instead the entire VM with the DRE print server role installed can be migrated from one node in the Hyper-V cluster to the other. Hyper-V clustering provides no service level protection to the print server and print spooler.

A well-designed Hyper-V cluster transparently reroutes user sessions to an available node in the cluster. Highly available solutions allow a user to automatically reconnect to a database if the connection fails. Active transactions roll back, but the new database connection is made on a different node, which is identical to the original. A user notices no loss of connection as long as there is one instance left serving the application.

## Create a highly available VM on a Hyper-V cluster

This section details how to make an existing VM highly available on a Hyper-V cluster. It does not provide instructions on how to create a VM or how to create a Hyper-V failover cluster. Please refer to your Microsoft documentation for creating VMs and Hyper-V clusters.

Hyper-V and failover clustering must be installed on at least two physical servers. Once Hyper-V and failover clustering have been installed on both servers, ensure the VMs are stored on a Cluster Shared Volume.

To make an existing VM highly available on a Hyper-V cluster, do the following:

1. Open the **Failover Cluster Manager**.
2. Expand the cluster and select **Roles** from the menu, and then select **Configure Role** from the Actions pane. This command opens the High Availability Wizard.
3. Click **Next** on the High Availability Wizard Welcome screen to continue.
4. On the **Select Role** screen, select **Virtual Machine** and click **Next**.
5. On the **Select Virtual Machine** screen, select the VM you want to configure for high availability, and click **Next**.
6. On the **Confirmation** screen, verify the information is correct, and click **Next** to start the configuration.
7. A **Summary** screen opens when configuration is complete. Click **Finish** and the virtual machine is now listed in the Roles container in the Failover Cluster Manager.

## Network load balancer

With a load balancing solution, such as Windows Server Network Load Balancing, several VMs are configured identically and the load balancer distributes service requests across each of the VMs fairly evenly. This process reduces the risk of any single VM becoming overloaded. Using load balancing is an effective way to eliminate VM downtime because VMs can be individually rotated and serviced without taking the service offline. However, load balancing only works with identical VMs, which have no shared or centralized data.

Load balancing is an effective way of increasing the availability of critical applications. When server failure instances are detected, they are seamlessly replaced when the traffic is automatically redistributed to servers that are still running. Not only does load balancing lead to high availability it also facilitates incremental scalability. Network load balancing facilitates higher levels of fault tolerance within service applications.

The configuration for the Network Load Balancer (NLB) appliance varies depending on the vendor and type of appliance, and must be managed by the end customer's internal IT administrator. The generic requirements for each protocol configured on the NLB appliance are outlined.

The recommended load balancing method used is Layer 4 in a direct server return/N-Path/direct routing configuration. Layer 4 load balancing uses information defined at the networking transport layer as the basis for deciding how to distribute client requests across a group of servers.

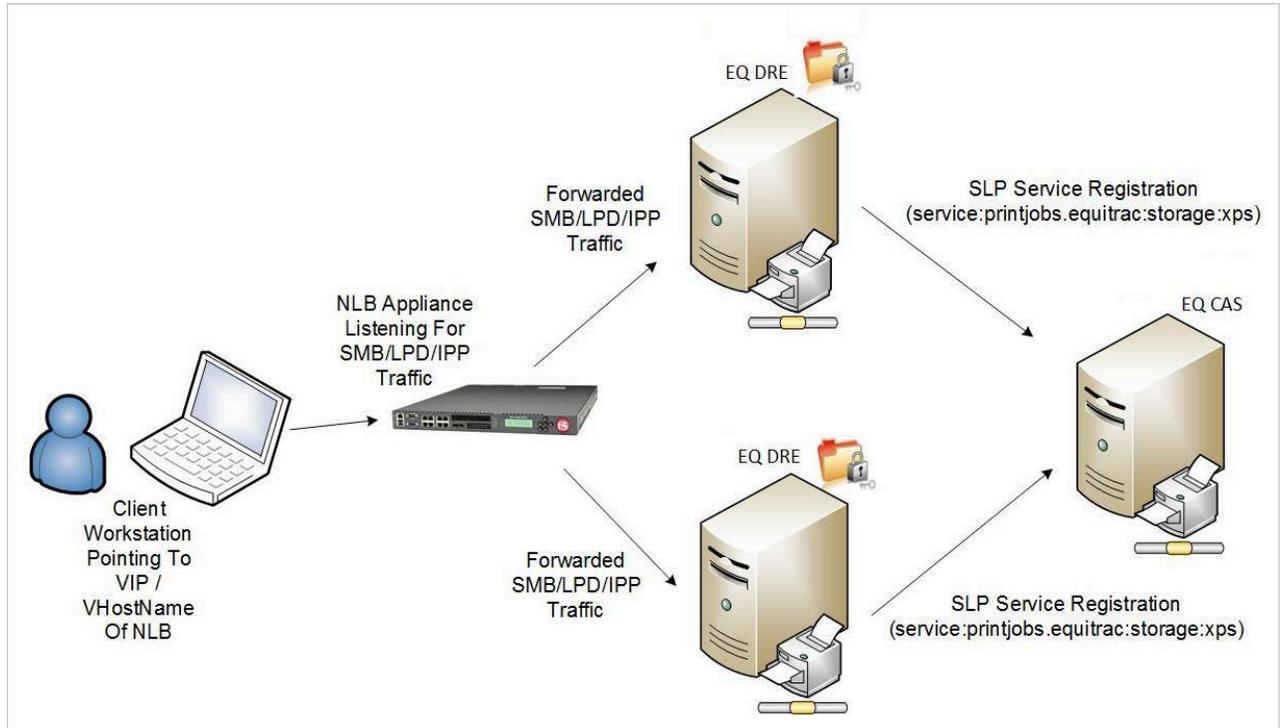
## Client and server network printing protocols

Microsoft print servers support different protocols to allow clients to submit print jobs over the network. The most common are IPP (Internet Printing Protocol), SMB (Server Message Block) and LPD (Line Printer Daemon). The pros and cons of each protocol when deployed behind an NLB appliance are described in the following table.

Protocol	Pros	Cons
IPP - (Internet Printing Protocol)	<ul style="list-style-type: none"> <li>Requires only one listener port configured on the NLB.</li> <li>Supports encryption using IPP over SSL.</li> <li>Strong integration between NLB technology and web based protocols.</li> </ul>	<ul style="list-style-type: none"> <li>Requires that the IPP client is installed/enabled on every workstation.</li> </ul>
SMB - (Server Message Block)	<ul style="list-style-type: none"> <li>Requires no client workstation modification.</li> <li>Can be deployed using OS standard tooling such as Group Policy Object (GPO).</li> </ul>	<ul style="list-style-type: none"> <li>Requires a large number of listener ports configured on the NLB. This is due to SMB generating multiple Remote Procedure Calls (RPC) from the client to the server side spooler on random destination ports.</li> </ul> <div style="background-color: #f0f0f0; padding: 5px;"> <p><b>Note</b> It is possible to force all RPC's to initiate using named pipes instead of TCP/IP by using the DisableRpcTcp flag. Setting this flag reduces the number of listener ports.</p> </div> <ul style="list-style-type: none"> <li>Requires the addition/modification of certain registry keys on each print server.</li> </ul>
LPD - (Line Printer Daemon)	<ul style="list-style-type: none"> <li>Requires only one listener port to be configured on the NLB.</li> <li>Works across non-trusted domains.</li> <li>Can be used to deliver print jobs from non-Windows clients.</li> </ul>	<ul style="list-style-type: none"> <li>Requires creation of a local print queue with an LPR port on every workstation.</li> </ul>

## Client workstation traffic flow

The following illustrates the basic traffic flow from a client workstation to multiple DREs through an NLB appliance.



## DRE considerations

The NLB appliance is the interface that all end users connect to, therefore the parameters for the print queue connection (the virtual Follow-You queue) must match across all Equitrac DRE servers. The queue name, driver and configuration must be identical on each DRE print server in the deployment. For redundancy, each DRE must have a physical device and output queue pointing to print devices in the deployment to ensure users can always release jobs to every device as well as submit them to the virtual Follow-You queue.

## Layer 4 load balancing

Layer 4 network load balancing (NLB) works by forwarding traffic to a specific server based upon the port or service the client wants to contact. With NLB, the load balancing appliance sits in between the client workstation and the server. The NLB appliance listens on its own port for requests from the client, and then makes decisions based on certain criteria which server to send them on to. The criteria could be traffic rate (i.e. if a particular server was busier than another), or it could simply be if the original server was no longer online or responding to requests.

This process is completely transparent to the clients requesting the service, they simply make the request and the NLB appliance routes the request to the appropriate server. The chosen server then replies directly to the client and offers the service.

## Add the IP address variable

Typical Layer 4 NLB deployments require that all servers placed behind a load balanced service have primary and secondary network interface cards (NIC) configured. The primary NIC provides the server with a dedicated, full-time connection to a network. The secondary NIC does not physically connect to the network.

When clients request a service via the NLB appliance, they contact an IP Address/Hostname that is configured on the NLB appliance specifically to listen for requests for that service. This is known as a Virtual IP or VIP. Since the NLB appliance forwards on these requests directly to the servers offering the service without altering the destination IP Address, the servers themselves must contain at least one NIC assigned with the same IP Address as the VIP. If they do not, then the request from the client is rejected as the servers assume that the request was not intended for them.

It is equally important that the secondary NIC added to each server does not actually connect to the production LAN. This ensures that when any client wishes to connect to the NLB appliance on its VIP, the servers with the secondary NIC also containing the VIP do not respond directly to the clients. This would initiate a direct connection between the client and the server and would avoid sending the traffic via the NLB appliance.

In order to avoid direct client to server connection, the majority of NLB appliance vendors advise to add the secondary NIC as a loopback adapter, as this is a virtual interface that does not physically connect to a network. Refer to your vendor's documentation for more information.

When a user prints via DRE, the DRE sends an SLP service registration message to CAS detailing how that print job can be accessed. That message includes:

- The hostname of the DRE
- The IP Address of the DRE
- The user ID that printed the job
- The port number the DRE service can be accessed on

When a secondary NIC containing the IP Address of the VIP is added to each DRE, the DRE sends both IP Addresses to CAS. When there are multiple registrations with the same IP Address, users do not see jobs that have been printed to different DREs. For example, if a user prints a job to DRE1, and then another job to DRE2, they only see the job on DRE1. This happens because both DREs have registered the same IP Address, and it is assumed that since DRE1 has been queried for jobs, DRE2 has also been queried.

In order to correct this behavior, and have DRE only send its production IP Address as part of the service registration message, a system environment variable must be added to each DRE containing the appropriate IP Address.

## Add environment variable

To add a system environment variable to each DRE, do the following:

1. Go to **Control Panel > All Control Panel Items > System** on the DRE, and select **Advanced system settings**.
2. On the System properties window, click **Environmental Variables**.
3. On the Environment Variables window, click **New** from the System variables section.
4. Create the Variable name **EQ\_IPADDRESSES** with a Variable value of the production IP Address of your Equitrac DRE server, and press **OK** and then **OK** again.
5. Repeat these steps above for every DRE in your deployment.

## Load balancing IPP print queues

To load balance IPP print queues, do the following:

1. Open the Print Management console and create identical Follow-You queues on every DRE print server. The Follow-You queues on all DREs must have the same name (including share name), driver (name and version), permissions and configuration.
2. Install the Internet Printing role on every DRE and verify that the print queues are visible via a web browser using the standard IPP URL – `http://printservername/printers` (e.g. `http://win2012/printers`).
3. Create a simple batch file with a command to stop the EQ DRE service, and save the batch file to a local directory on each DRE.
4. Go to **Administrative Tools > Services** and open the **Print Spooler** service and configure the **Recovery** options for the Print Spooler service to run the batch file upon failure.
5. Open Task Manager and end the Print Spooler process (`spoolsv.exe`), and confirm that once the Print Spooler service has stopped, the EQ DRE service has also stopped. Repeat this process on each EQ DRE.
6. On the NLB appliance, do the following: (consult your NLB appliance vendor for support):
  - Configure the **Virtual Service IP Address (VIP)** in the same subnet as the DRE server.
  - Set the **Ports** to **80**. A listener port for 443 TCP must also be created if using IPP over SSL
  - Set the **Protocol** to **TCP**.
  - Set the Load balancing **Forwarding Method** to **Direct Routing** (i.e. layer 4/direct routing/direct server return/N-Path).
  - Ensure the **Persistent** checkbox is selected. **Affinity/Persistence/Stickness** must be enabled.
  - Set the **Check Port** for server/service online to **2938**.
7. On the DNS server, create a hostname and corresponding “Host (A)” record to match the previously created VIP. This is the hostname/URL that end users connect to from their workstations to map to the Follow-You Queue (e.g. `http://PrintIPP/printers/EQFY/.printer`).
8. On a client workstation, enable the Internet Printing Client.
9. On the same client workstation, open a web browser and navigate to **`http://VIPName/printers`** where VIPName is the previously created hostname (e.g. `http://printipp/printers`). Verify that the print queues are visible.

10. Click on the print queue link and select **Connect**, and then **Yes** when prompted to add the printer connection.
11. Print a job to the newly added IPP print queue and verify that it processed successfully. The job is listed in the Equitrac Print Queue Viewer on the first DRE.
12. On the DRE that was just printed through, stop the print spooler service, and verify that the EQDRE service also stops.
13. On the client workstation, verify that the load balanced print queue is still online.
14. Submit another print job and verify that it processed successfully. The job is listed in the Equitrac Print Queue Viewer on the second DRE.

**Note** In order to verify which DRE the job is sent to, you may need to access the NLB appliance's traffic logs.

15. Start the print spooler service on the first DRE, and verify that the EQDRE service restarts.
16. Log into a release device and verify both jobs are listed for release.

## Load balancing SMB print queues

To load balance SMB print queues, do the following:

1. Open the Print Management console and create identical Follow-You queues on every DRE print server. The Follow-You queues on all DREs must have the same name (including share name), driver (name and version), permissions and configuration.
2. Open the Registry Editor and add the following registry keys on each DRE:
  - a. Navigate to **HKEY\_LOCAL\_MACHINE\SYSTEM\CurrentControlSet\Services\LanmanServer\Parameters**, and create the following keys:
    - REG\_DWORD - **DisableStrictNameChecking** (set to 1)
    - REG\_MULTI\_SZ - **OptionalNames**. This is the name used to connect to the NLB appliance (e.g PrintLB). The name must resolve to the virtual service IP address (VIP).
    - REG\_SZ - **srvcomment**. This is a description of the print server (e.g WIN2012 In PrintLB)
  - b. Navigate to **HKEY\_LOCAL\_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa**, and create the following:
    - REG\_DWORD - **DisableLoopbackCheck** (set to 1)
  - c. Navigate to **HKEY\_LOCAL\_MACHINE\SYSTEM\CurrentControlSet\Control\Print**, and create the following:
    - REG\_DWORD - **DisableRpcTcp** (set to 1)
3. Create a simple batch file with a command to stop the EQ DRE service, and save the batch file to a local directory on each DRE.
4. Go to **Administrative Tools > Services** and open the **Print Spooler service** and configure the **Recovery** options for the Print Spooler service to run the batch file upon failure.
5. Open Task Manager and end the Print Spooler process (spoolsv.exe), and confirm that once the Print Spooler service has stopped, the EQ DRE service has also stopped. Repeat this process on each EQ DRE.

6. On the NLB appliance, do the following: (consult your NLB appliance vendor for support):
  - Configure the **Virtual Service IP Address (VIP)** in the same subnet as the DRE server.
  - Set the **Ports** to **445, 135-139**.
  - Set the **Protocol** to **TCP**.
  - Set the Load balancing **Forwarding Method** to **Direct Routing** (i.e. layer 4/direct routing/direct server return/N-Path).
  - Ensure the **Persistent** checkbox is selected. **Affinity/Persistence/Stickness** must be enabled.
  - Set the **Check Port** for server/service online to **2938**.
7. On the DNS server, create a hostname and corresponding “Host (A)” record to match the previously created VIP. This must match the value created for REG\_MULTI\_SZ - OptionalNames (i.e. PrintLB). This is the hostname that end users connect to from their workstations to map to the Follow-You Queue (e.g. \\PrintLB\FollowYouPrinting).
8. Add a printer on a client workstation using the newly created virtual name and the common print queue share name that resides on the DRE servers (e.g. \\PrintLB\EquitracFollowYou).
9. Print a job to the newly added SMB print queue and verify that it processed successfully. The job is listed in the Equitrac Print Queue Viewer on the first DRE.
10. On the DRE that was just printed through, stop the print spooler service, and verify that the EQDRE service also stops.
11. On the client workstation, verify that the load balanced print queue is still online.
12. Submit another print job and verify that it processed successfully. The job is listed in the Equitrac Print Queue Viewer on the second DRE.

**Note** In order to verify which DRE the job is sent to, you may need to access the NLB appliance’s traffic logs.

13. Start the print spooler service on the first DRE, and verify that the EQDRE service restarts.
14. Log into a release device and verify both jobs are listed for release.

## Load balancing LPD print queues

To load balance LPD print queues, do the following:

1. Open the Print Management console and create identical Follow-You queues on every DRE print server. The Follow-You queues on all DREs must have the same name (including share name), driver (name and version), permissions and configuration.
2. Install the LPD Service role on every DRE.
3. Verify that the LPD Service is running correctly and can receive print jobs via LPR. by pausing the Equitrac Follow-You Queue and then running the following command from another computer on the network that has the LPR port monitor installed. Modify the parameters to suit your environment.

```
LPR -S 1.2.3.4 -P PrinterName -o DocumentName.txt
```
4. Create a simple batch file with a command to stop the EQ DRE service, and save the batch file to a local directory on each DRE.
5. Go to **Administrative Tools > Services** and open the **Print Spooler** service and configure the **Recovery** options for the Print Spooler service to run the batch file upon failure.

6. Open Task Manager and end the Print Spooler process (spoolsv.exe), and confirm that once the Print Spooler service has stopped, the EQ DRE service has also stopped. Repeat this process on each EQ DRE.
7. On the NLB appliance, do the following: (consult your NLB appliance vendor for support):
  - Configure the **Virtual Service IP Address (VIP)** in the same subnet as the DRE server.
  - Set the **Ports** to **515**.
  - Set the **Protocol** to **TCP**.
  - Set the Load balancing **Forwarding Method** to **Direct Routing** (i.e. layer 4/direct routing/direct server return/N-Path).
  - Ensure the **Persistent** checkbox is selected. Affinity/Persistence/Stickiness must be enabled.
  - Set the **Check Port** for server/service online to **2938**.
8. If you wish to use hostnames when specifying the LPD server on each client, then on the DNS server, create a hostname and corresponding "Host (A)" record to match the previously created VIP. If you wish to use only IP Address, then you can skip this step.
9. On a client workstation, create a standard TCP/IP print queue, with an LPR port configured to point to the previously created hostname and VIP.

For mass creation and deployment of local LPR print queues, the following scripts can be run:

- Install the print driver (usually served from a shared folder on a file/print server):

```
cscript C:\Windows\System32\Printing_Admin_Scripts\en-US\prndrvr.vbs -a -m  
"Print Driver Name" -i \\server\folder\printdriver.inf
```

- Create the LPR port:

```
cscript C:\Windows\System32\Printing_Admin_Scripts\en-US\prnport.vbs -a -r  
PrintLPD -h 10.10.10.205 -o LPR -n 515 -q EQFY -2e
```

- Create the print queue, assign the print driver and attach the LPR port:

```
cscript C:\Windows\System32\Printing_Admin_Scripts\en-US\prnmngr.vbs -a -p  
Local_QueueName -m "Print Driver Name" -r "PrintLPD"
```

10. Print a job to the newly added LPR print queue and verify that it processed successfully. The job is listed in the Equitrac Print Queue Viewer on the first DRE.
11. On the DRE that was just printed through, stop the print spooler service, and verify that the EQDRE service also stops.
12. On the client workstation, verify that the load balanced print queue is still online.
13. Submit another print job and verify that it processed successfully. The job is listed in the Equitrac Print Queue Viewer on the second DRE.

**Note** In order to verify which DRE the job is sent to, you may need to access the NLB appliance's traffic logs.

14. Start the print spooler service on the first DRE, and verify that the EQDRE service restarts.
15. Log into a release device and verify both jobs are listed for release.